

Hotelling Under Pressure*

Soren T. Anderson[†] Ryan Kellogg[‡] Stephen W. Salant[§]

June 26, 2014

Abstract

We show that oil production from existing wells in Texas does not respond to price incentives. Drilling activity and costs, however, do respond strongly to prices. To explain these facts, we reformulate Hotelling's (1931) classic model of exhaustible resource extraction as a drilling problem: firms choose when to drill, but production from existing wells is constrained by reservoir pressure, which decays as oil is extracted. The model implies a modified Hotelling rule for drilling revenues net of costs and explains why production is typically constrained. It also rationalizes regional production peaks and observed patterns of price expectations following demand shocks.

JEL classification numbers: Q3, Q4

Key words: crude oil prices; oil extraction; decline curve; oil drilling; rig rental rates; exhaustible resource

*For helpful comments and suggestions, we are grateful to Ying Fan, Cloè Garnache, Stephen Holland, Lutz Kilian, Mar Reguant, Dick Vail, Jinhua Zhao, and seminar and conference audiences at AERE, Case Western Reserve University, the Energy Institute at Haas Energy Camp, the Kansas City Federal Reserve, the MSU/UM Energy and Environmental Economics Summer Workshop, Michigan State University, NBER EEE/IO, the Occasional California Workshop in Environmental and Resource Economics, the University of Michigan, and the University of Montreal. Excellent research assistance was provided by Dana Beuschel and Sam Haltenhof.

[†]Michigan State University and NBER. Email: sta@msu.edu

[‡]University of Michigan and NBER. Email: kellogg@umich.edu

[§]University of Michigan and Resources for the Future. Email: ssalant@umich.edu

1 Introduction

Hotelling’s (1931) classic model of exhaustible resource extraction—featuring forward-looking resource owners that maximize wealth by trading off extraction today versus extraction in the future—holds great conceptual appeal. Ever since, economists have modeled the optimal extraction of an exhaustible resource as a cake-eating problem in which resource owners are able to reallocate extraction across different periods without constraint. A common application of Hotelling’s logic is to production from an oil reserve. In this paper, however, we argue that observed patterns of oil production and prices are not compatible with Hotelling (1931) nor any of its modifications in the literature. Instead, we show that to replicate structurally the dynamics of oil supply, economists should account for geological constraints on well-level oil production and recast the Hotelling model as a well-drilling investment problem—or as a keg-tapping problem, if one wishes to maintain an analogy to food and drink.

Using data from Texas over 1990–2007, we show that oil production from drilled wells declines asymptotically toward zero and is not affected by shocks to spot or expected future oil prices, even during 1998–1999 when oil spot prices were very low and oil futures markets implied that prices were expected to rise (temporarily) faster than the interest rate. This behavior is inconsistent with most extraction models in the literature, and we show empirically that it is not driven by common-pool problems, oil lease contract provisions, or other institutional factors. Instead, we argue that the production decline is rationalized by the cost structure of the oil industry and by the loss of underground reservoir pressure that results from cumulative extraction.

When a well is first drilled, the pressure in the underground oil reservoir is high. Production may therefore initially be rapid, since the maximum rate of fluid flow is roughly proportional to the pressure available to drive the oil through the reservoir, into the well, and then up to the surface. Over time, however, extraction reduces the reservoir pressure, so that the well’s maximum flow decays toward zero as the reserves are depleted. We hypothesize that the asymptotic decline in production in the Texas data occurs because extractors

never cut production below their declining capacity constraints, even though doing so would conserve pressure that is valuable in the future.

While production from drilled wells is insensitive to oil prices, our Texas data show that the drilling of new oil wells and the rental price of drilling rigs both respond strongly to oil price shocks. These results motivate us to recast Hotelling's (1931) canonical model as a drilling problem rather than a production problem. Extractors in our model choose when to drill their wells (or tap their kegs, per our analogy above), but the maximum flow from these wells is (like the libation from a keg) constrained due to pressure and decays asymptotically toward zero as more oil is extracted. In addition, the marginal cost of drilling in our model strictly increases with the rate of drilling, consistent with our data on rig rental prices and the notion that there is an upward-sloping supply curve for rig rentals.

We first characterize the extraction incentives implied by our model and explain why producing at the flow constraint can be optimal even when prices are expected to rise faster than the rate of interest and then plateau, as in 1998–1999. A well owner that attempts to arbitrage such prices by producing below his constraint cannot entirely recover the deferred production at the instant prices reach their zenith in present value terms. Instead, the pressure constraint forces him to recover this production gradually over the entire remaining life of the well, which will include periods when prices are lower in present value than today's price. Accordingly, our model implies that producers with price expectations matching those implied by futures markets never had an incentive to produce below their constraint in our sample, including 1998–1999 (the only exception being wells with very low production rates, for which fixed operating costs can generate an incentive to shut down). Our model also suggests that this outcome was not a historical fluke: mild sufficient conditions on primitives guarantee that production will be constrained along the entire equilibrium path.

We then characterize optimal drilling, showing that understanding investment incentives is central to understanding dynamics in oil markets. In the canonical Hotelling model, price net of marginal extraction cost rises at the rate of interest whenever production occurs

(“Hotelling’s Rule”). In our reformulated model, a modified Hotelling Rule holds: whenever drilling occurs, the discounted revenue stream that flows to the marginal well, net of the marginal drilling cost, rises at the rate of interest. In the limiting case with no resource scarcity, our model closely resembles a macroeconomic “Q-theory” model of investment, whose dynamics lead to a steady state in which the marginal discounted revenue stream from investing in a well equals the marginal cost of drilling a well.

We then show that the equilibrium dynamics implied by our model easily and naturally replicate a wide range of the crude oil extraction industry’s most salient qualitative features. First, our model implies that the flow constraint will typically bind in equilibrium so that production from drilled wells will be unresponsive to shocks, as we observe in our data. Thus, aggregate production will evolve gradually over time, following changes in the drilling rate, and will only respond to shocks with a significant lag. This result provides a foundation for a macro-empirical literature showing that aggregate oil production is price inelastic, at least in the short run (Griffin 1985; Hogan 1989; Jones 1990; Dahl and Yucel 1991; Ramcharan 2002; Güntner forthcoming).¹ This inelasticity has important implications for the macroeconomic effects of oil supply and demand shocks, since inelastic supply and demand lead to volatile oil prices (see Hamilton (2009) and Kilian (2009)). Second, within oil-producing regions, the model predicts the commonly observed phenomenon (Hamilton 2013) that production initially rises as drilling ramps up but then peaks and eventually declines as drilling slows and the flow from existing wells decays. Third, we show that positive global demand shocks lead to an immediate increase in oil prices, drilling activity, and rig rental prices, and that oil prices may subsequently be expected to fall if the increased rate of drilling causes production to increase. These results are reversed for negative demand shocks, which can—if large enough—lead to the expectation that oil prices will rise faster than the rate of interest following the initial drop in price. These predicted responses to demand shocks match our data on drilling activity, rig rental prices, and oil futures markets.

¹Rao (2010), however, finds evidence using well-level data that firms can shift production across wells in response to well-specific taxes.

We conclude by discussing how our model relates to the broader Hotelling literature and why prior work cannot explain the full array of real-world oil market phenomena that we document. In short, oil extractors face a keg-tapping problem, not a cake-eating problem. Thus, while we retain Hotelling’s (1931) conceptually appealing framework of forward-looking, wealth-maximizing agents, our core innovation is to impose the constraint that oil flow is limited by reservoir pressure (which declines with cumulative production), while allowing for convex costs of investment in new wells. Extraction decisions are therefore made well-by-well, not barrel-by-barrel. Our main contribution then is to show that a Hotelling-style model thus grounded in the actual cost structure of the oil industry *can* give predictions that are empirically valid. Our hope is that these results will renew interest in using Hotelling models to understand and predict the behavior of oil extractors and markets.

2 Empirical evidence from Texas

In this section we study how oil production and drilling in Texas respond to incentives generated by changes in current and expected future oil prices, as revealed in futures markets. We show that oil production exhibits nearly zero response to price shocks, whereas drilling activity—along with the cost of renting drilling rigs—responds strongly. We then discuss how these results derive from the fundamental technology of crude oil extraction.

2.1 Data sources

Our crude oil drilling and production data for 1990–2007 come from the Texas Railroad Commission (TRRC). The drilling data come from the TRRC’s “Drilling Permit Master” dataset, which provides the date, county, and lease name for every well drilled in Texas. A lease is land upon which an oil production company has obtained—from the (usually private) mineral rights owner—the right to drill for and produce oil and gas. Over 1990–2007, a total

of 157,271 new wells were drilled, along with 42,893 “re-entries” of existing wells.²

The production data come from the TRRC’s “Oil and Gas Annuals” dataset, which records monthly crude oil production at the lease-level.³ Individual wells are not flow-metered.⁴ Thus, we generally cannot observe well-level production, though for some analyses we will isolate the sample to leases that have a single flowing well.

Our analysis of the production data focuses on whether firms respond to oil price shocks by adjusting the flow rates of their existing wells, possibly all the way to zero, which is known as “shutting in” a well.⁵ Such adjustment would typically be accomplished by slowing down or speeding up the pumping unit.⁶ To distinguish these actions from investments in new production, such as drilling a new well, we discard leases in which any rig work took place.⁷ In the remaining data, there exist 16,148 leases for which production data are not missing for any month from 1990–2007 and production is non-zero for at least one month, yielding 3,487,968 lease-month observations. The typical oil lease in Texas has a fairly low rate of

²A re-entry occurs when a rig is used to deepen a well, drill a “sidetrack” off an existing well bore, or stimulate production by fracturing the oil reservoir. These interventions are all similar to drilling a new well in that they require a substantial up-front investment and provide access to new oil-bearing rock. A small share (< 10%) of the new wells were drilled to inject water or gas into the reservoir rather than extract oil. These injection investments can mitigate, but not eliminate, the rate of production decline that we document here. We abstract away from differences between new drilling, re-entries, and injection well drilling in our analysis.

³Due to false zeros for some leases in 1996 and December 2004–2007, we augmented these data by scraping information from the TRRC’s online production query tool, verifying that the two sources match for leases and months not affected by the data error.

⁴Direct production includes oil, gas, and often water. Separation of these products typically occurs at a single facility serving the entire lease, with the oil flowing from the separation facility into storage tanks. Oil is metered leaving the storage tank for delivery to a pipeline or tanker truck for sale. Although firms often assess well-level productivity by diverting each well’s flow into a small “test separator,” these data are not available from TRRC, nor would they give a particularly accurate measure of a well’s monthly flow.

⁵Throughout our analysis, we assume that oil price movements are exogenous to actions undertaken by Texas oil producers. This treatment is plausible given that Texas firms are a small share of the world oil market (1.3% in 2007) and evidence that oil price shocks during our sample were primarily driven by global demand shocks and (to a lesser extent) international rather than U.S. supply shocks (Kilian 2009). Moreover, the positive covariance between drilling activity and oil prices apparent in figure 4 strongly suggests that Texas drilling activity is responding to price shocks rather than vice-versa.

⁶The vast majority of the wells in the dataset are pumped and do not flow naturally. The average lease-month in the data has 2.02 pumped wells and 0.06 naturally flowing wells.

⁷Discarding these leases requires matching the drilling dataset to the production dataset. Since lease names are not consistent across the two datasets, we conservatively identify all county-firm pairs in which rig work took place and then discard *all* leases corresponding to such pairs (unlike leases, counties and firms are consistently identified with numeric codes in both datasets).

production, reflecting the fact that most oil fields in Texas are mature and have been heavily produced in the past. The average daily lease production in the data is 3.6 barrels of oil per day (bbl/d), with a standard deviation of 18.2 bbl/d. A total of 1,070,632 (31%) of the observed lease-months have zero production, while the maximum is 9,510 bbl/d.

Our oil price data come from the New York Mercantile Exchange (NYMEX) and measure prices for West Texas Intermediate (WTI) crude oil delivered in Cushing, Oklahoma—the most common benchmark for crude oil prices in North America—from 1990–2007. We use the Bureau of Labor Statistics’s All Urban, All Goods Consumer Price Index (CPI) to convert all prices to December 2007 dollars. We use the front-month (upcoming month) futures price as our measure of the spot price of crude oil and use prices for longer-term futures contracts to measure firms’ price expectations.⁸

Figure 1 shows the time series of crude oil spot prices (solid black line) as well as the futures curves as of December in each year (dashed colored lines).⁹ For example, the left-most dashed line shows prices in December 1990 for futures contracts with delivery dates from January 1991 through December 1992. As is clear from the figure, the futures market for crude oil is often backwarddated (meaning that the futures price is lower than the spot price) and was strongly backwarddated during the mid-2000s when the spot price was rapidly increasing. Kilian (2009) and Kilian and Hicks (2013) attribute the increase in spot prices

⁸The use of futures prices as measures of price expectations is not without controversy. Alquist and Kilian (2010), for example, shows that futures markets do not out-perform a simple no-change forecast in out-of-sample forecasting over 1991-2007. We use futures markets here for several reasons. First, NYMEX futures are liquidly traded at the horizons we consider here, and with many deep-pocketed, risk-neutral traders, the futures price should equal the expected future spot price. Second, a majority of oil producers in Texas claim to use futures prices in making their own price projections (Society of Petroleum Evaluation Engineers 1995). Third, Kellogg (2014) shows that firms’ drilling activity in Texas is more consistent with price expectations based on the futures market than with a no-change forecast. Finally, as we show in appendix A, we find that firms’ on-lease above-ground oil stockpiles increase when futures prices are high relative to spot prices, as one would expect if firms’ expectations aligned with the market.

⁹We have converted all of the price data so that the slopes of the futures curves in figure 1 and the expected rates of price change used in our analysis reflect real rather than nominal changes. To convert the futures curves from nominal to real, we adjust for both the trade date’s CPI and for expected annualized inflation of 2.50% between the trade date and delivery date (the average annual inflation rate from January 1990 to December 2007 is 2.50%, and inflation varies little over the sample). For example, we convert the nominal prices for futures contracts traded in December 1990 to real price expectations by multiplying by the December 2007 CPI, dividing by the December 1990 CPI, and then dividing each contract price by $1.025^{t/12}$, where t is the number of months between the trade date and the delivery date.

Figure 1: Crude oil spot prices and futures curves



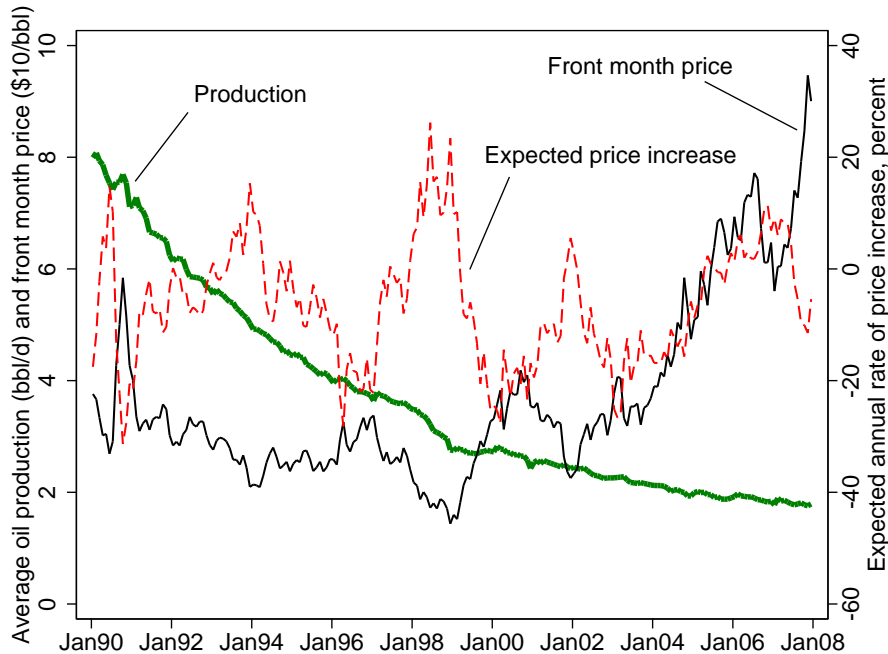
Note: This figure shows crude oil front month (“spot”) and futures prices, all in real \$2007. The solid black line is the spot price of oil. The dashed lines are the futures curves as of December in each year. See text for details.

during this period to a series of large, positive, and unanticipated shocks to the demand for oil, primarily from emerging Asian markets.

Figure 1 also reveals several periods of contango (meaning that the futures price is higher than the spot price) during the sample, particularly during 1998–1999 when the oil price was quite low. Kilian (2009) attributes the low oil prices during this period to a negative demand shock arising from the Asian financial crisis.

Finally, we obtained information on rental prices (“dayrates”) for drilling rigs from Rig-Data (1990–2013), a firm that collects information on U.S. onshore drilling activity and publishes rig rental rates in its Day Rate Report. As discussed in Kellogg (2011), the oil production companies that make drilling and production decisions do not drill their own wells but rather contract drilling out to independent service companies that own rigs. Paying to rent a rig and its crew is typically the largest line-item in the overall cost of a well. The

Figure 2: Crude oil prices and production from existing wells in Texas



Note: This figure presents crude oil front month (“spot”) prices and the expected percentage change in prices over one year, as well as daily average lease-level production from leases on which there was no rig activity (so that all production comes from pre-existing wells). All prices are real \$2007, and expected price changes are net of inflation. See text for details.

data provided by RigData are quarterly, covering Q4 1990 through Q4 2007, and are broken out by region and rig depth rating. We use dayrates for rigs with depth ratings between 6,000 and 9,999 feet (the average well depth in our drilling data is 7,425 feet) for the Gulf Coast / South Texas region. Observed dayrates range from \$6,315 to \$15,327 per day, with an average of \$8,008 (all real December 2007 dollars).

2.2 Production from existing wells does not respond to prices

Our main empirical results focus on production from leases on which there was no rig activity from 1990–2007, so that all production comes from pre-existing wells. Figure 2 presents daily average production (in bbl/d) for these leases in each month, along with monthly crude oil spot prices and the expected percentage change in spot prices over one year (i.e., the

percentage difference between the 12-month futures price and the spot price). Production is dominated by a long-run downward trend, with little response either to the spot price of oil or to expected future price changes. In appendix A, we present regression results confirming the lack of response to price incentives. In addition, we show that the pattern shown in figure 2 holds for subsamples of leases that have relatively high production volumes (in excess of 100 bbl/d) and for production from wells that are drilled during the sample period. Thus, our results are not specific only to the low-volume wells that are the norm in Texas.

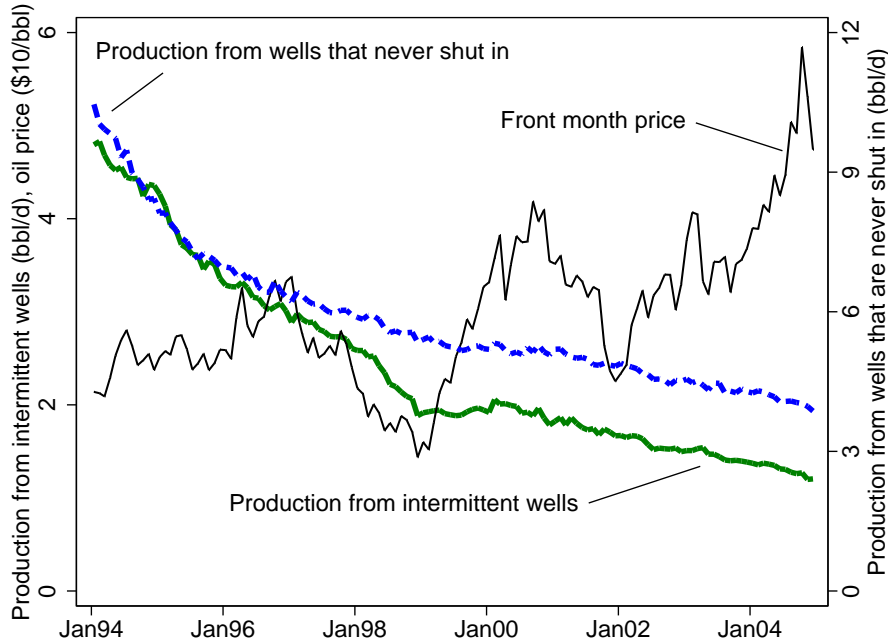
These results contrast sharply with predictions from standard Hotelling models, which would predict a complete shutdown of oil production during periods such as 1998–1999 when spot prices are lower than expected future prices in present value terms. Moreover, under a commonly used assumption of increasing marginal extraction costs, standard models will predict that production should increase following positive price shocks and decrease following negative price shocks. None of these predictions appears strongly in the time series of production from pre-existing wells.

Figure 2 does suggest that production may have deviated slightly from the long-run trend during the 1998–1999 period in which the spot price fell below \$20/bbl and the expected percentage price change over one year exceeded 10% (and sometimes 20%). In particular, it appears that production accelerated its decline rate in 1998 while prices were falling, leveled off in 1999 while prices were rising, and then resumed its usual decline in 2000.

To assess whether this deviation is real and what mechanism lay behind it, we study whether it arose from wells being shut in or from changes in production on the intensive margin. We first isolate the sample to leases that had no more than one flowing well over 1994–2004,¹⁰ so that observed lease-level production during this time can be interpreted as well-level production. We then split this sample into two groups: wells that are never shut in over 1994–2004 and “intermittent” wells that are shut in for at least one calendar month. Figure 3 plots the time series of production from these two samples. This figure makes clear

¹⁰We use a shorter sample window to increase the number of qualifying leases in the sample and to improve the visualization of the 1998–1999 period in figure 3.

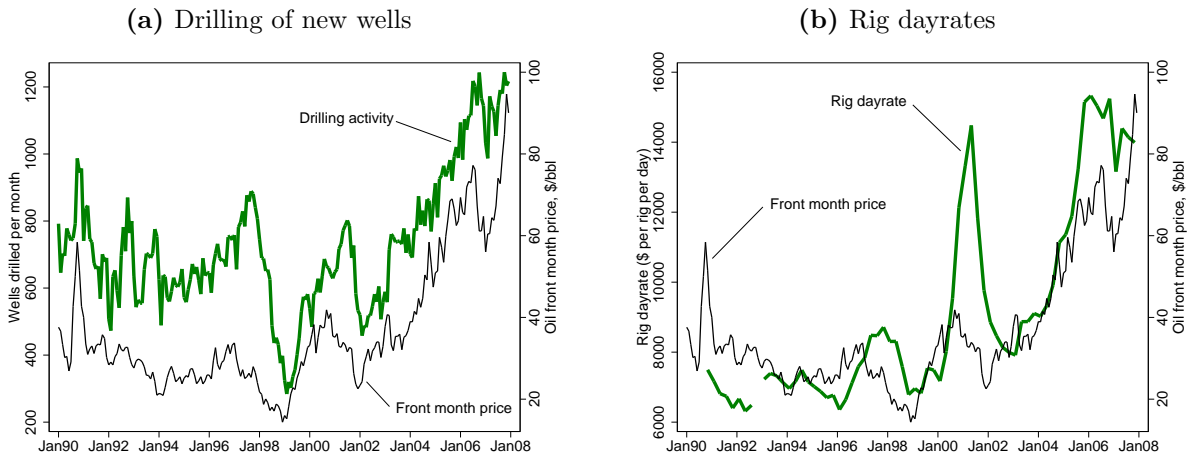
Figure 3: Intermittent wells versus wells never shut in



Note: This figure shows production from wells that shut in at least once during 1994–2004 versus production from wells that never shut in during 1994–2004. Production data come from leases that had no more than one productive well and never experienced a rig intervention over 1994–2004. Oil prices are real \$2007. See text for details.

that the 1998 deviation from trend was driven entirely by low-volume “marginal” wells that sometimes have zero production. For wells that always produce, there is no adjustment on the intensive margin, even though firms are able to adjust their pumping rate (see Rao (2010)). It appears that when prices fell in 1998, an unusually large number of marginal wells were shut in, temporarily accelerating the decline. Then, when prices recovered during 1999, many of these wells were returned to production, temporarily slowing the decline. Apart from these deviations, a significant response of production to price signals does not appear anywhere in the data. In particular, the most productive wells show no price response (see appendix A).

Figure 4: Texas rig activity versus crude oil spot prices



Note: Panel (a) shows the total number of new wells drilled across all leases in our dataset. Panel (b) shows dayrates for the Gulf Coast / South Texas region, for rigs with depth ratings between 6,000 and 9,999 feet. The dayrate data are quarterly rather than monthly. Data are available beginning in Q4 1990, and data for Q4 1992 are missing. See text for details.

2.3 Rig activity does respond to price incentives

These no-response results based on existing wells starkly contrast with new drilling activity in Texas. Figure 4(a) shows the total number of new wells drilled across all leases in our dataset, along with the spot price for crude oil. There is a pronounced positive correlation between oil prices and new drilling activity. Appendix A presents related regression results indicating that the elasticity of the monthly drilling rate with respect to the crude oil spot price is about 0.6 and statistically different from zero. We have also found that the use of rigs to re-enter old wells correlates with oil prices, though not as strongly as the drilling of new wells.

When oil production companies drill more wells in response to an increase in oil prices, more rigs (and crews) must be put into service to drill them. Figure 4(b) shows that these fluctuations in rig demand are reflected in a positive covariance between rig dayrates and oil prices. Regressions confirm that the elasticity of the rig rental rate with respect to oil prices

is large (0.79) and statistically significant. Thus, as the industry collectively wishes to drill more wells within a given time frame, the marginal cost of drilling those wells increases.

2.4 Industry cost structure explains these price responses

The analysis above documents two empirical facts about oil production and drilling in Texas from 1990–2007. First, production from drilled wells is almost completely unresponsive to changes in spot or expected future oil prices, with an exception being an increased rate of shut-ins during the 1998 oil price crash. Second, drilling of new wells responds strongly to oil price changes, and rig dayrates respond commensurately. Here, we argue that these empirical results reflect an industry cost structure with the following characteristics:¹¹

1. The rate of production from a well is physically constrained, and this constraint declines asymptotically toward zero as a function of cumulative production. This function is known in the engineering literature (Hyne 2001) as a well’s production decline curve.
2. The marginal cost of production below a given well’s capacity constraint, consisting of energy input to the pump (if there is one) and the cost of transporting oil off the lease, is very small relative to observed oil prices.
3. The fixed costs of operating a producing well are non-zero. There may also be costs for restarting a shut-in well, but they are not too large to be overcome.
4. Drilling rigs and crews are a relatively fixed resource, at least in the short run. Higher rental prices are required to attract more rigs into active use, leading to an upward-sloping supply curve of drilling rigs for rent.

The capacity constraint and low marginal production cost relate to the observation that production from existing wells steadily declines while not responding to oil price shocks.

¹¹For a particularly cogent discussion within the economics literature, see Thompson (2001).

Because oil production firms in Texas are price-takers,¹² production will be unresponsive to price shocks, as the data reflect, only if the oil price intersects marginal cost at a vertical, capacity-constrained section of the curve. While the marginal cost of production below the capacity constraint is not necessarily zero, it must be well below the range of oil prices observed in the data. There remains the question of why producers did not reduce production during 1998–1999 when oil prices were forecast to rise more quickly than the rate of interest. Section 3.2 discusses why the declining capacity constraint precluded firms from taking advantage of this seeming intertemporal arbitrage opportunity.

The existence of a capacity constraint for well-level production is consistent not only with the data presented above but also with standard petroleum geology and engineering.¹³ As noted recently in the economics literature by Mason and van't Veld (2013), the flow of fluid through reservoir rock to the well bore is governed by Darcy's Law (Darcy 1856), which stipulates that the rate of flow is proportional to the pressure differential between the reservoir and the well. In the simplest model of reservoir flow, the reservoir pressure is proportional to the volume of fluid in the reservoir. In this case, the maximum flow rate is proportional to the remaining reserves, consistent with the stylized fact reported in Mason and van't Veld (2013) that U.S. production has remained close to 10% of proven reserves since the industry's infancy, despite large changes in production over time. This proportionality yields an exponential production decline curve for drilled wells. More complex cases, which might involve the presence of gas, water, or fractures in the reservoir, may yield a more general hyperbolic decline. Regardless, the physical laws governing fluid flow place a limit on the rate at which oil can be extracted from a reservoir, and this limit declines with the volume of oil remaining.¹⁴

¹²The market for crude oil is global, and Texas as a whole (let alone a single firm) constitutes only 1.3% of world oil production (Texas and world oil production data for 2007 from the U.S. Energy Information Administration); thus, the exercise of market power by Texas oil producers is implausible.

¹³We give the geologic and engineering basis for well-level capacity constraints only a brief treatment here. For a fuller discussion of fluid flow and production decline curves, Hyne (2001) is an excellent source that does not require a geology or engineering background.

¹⁴Darcy's law governs oil flowing through the reservoir and into the bottom of the wellbore. Installing a pump on a well effectively eliminates the need for the oil to overcome gravity as it rises up the well, but

Per figure 3, some relatively low-volume wells were shut in during 1998. These shut-ins are consistent with the existence of fixed production costs, which intuitively arise from the need to monitor and maintain surface facilities such as pumps, flowlines, and separators so long as production is non-zero. When the oil price fell in 1998, production from these wells may no longer have been sufficient to cover their fixed costs, explaining the decision to shut in.¹⁵ When oil prices subsequently recovered, many (though not all) of these wells restarted, suggesting that start-up costs can be overcome.

In appendix A, we consider and rule out alternative explanations for the lack of response of oil production to oil prices. We show that the overall lack of price response cannot be explained by: (1) leasing agreements that require non-zero production (because multiple-well leases show the same results); (2) races-to-oil induced by open-access externalities within oil fields (because fields controlled by a single operator show the same results); (3) well-specific production quotas (because production quotas are not binding); or (4) producer myopia or price expectations that are not aligned with the futures market (because producers respond to high futures prices by stockpiling oil above ground).

While appropriate for most crude oil extraction, our model is not applicable to resources such as coal, metal ores, or oil sands that are mined rather than produced through wells. For these resources, fluid flow is not important, and marginal extraction costs are likely to be substantial and increasing with the production rate, so that modeling extraction requires a very different approach than that presented here.

it does not negate Darcy's Law, implying that one could never drain all of the oil in finite time, even if the pump pulled a vacuum on the bottom of the well. Enhanced oil recovery through the use of injection wells can slow the pressure decline in the reservoir, but again oil production must inevitably decline as the production wells produce more and more of the injected fluids rather than oil.

¹⁵These shut ins may also in part be explained by an incentive to postpone production (and the payment of fixed costs) to future periods when oil prices were expected to be higher. We discuss these arbitrage incentives in detail in section 3.2, noting their potential interaction with fixed costs in footnote 27.

3 Recasting Hotelling as a drilling problem

In this section, we develop a theory of optimal oil drilling and extraction that closely follows the industry cost structure described above. After setting up the problem, we derive and interpret conditions that necessarily hold at any optimum, focusing first on incentives to produce at the capacity constraint and then on incentives to drill new wells.

3.1 Planner's problem and necessary conditions

Because there are millions of operating oil wells in the world, we formulate our model as a decision problem in which there is a continuum of infinitesimally small wells to be drilled. We use continuous time to facilitate interpretation of the necessary conditions and the analysis of equilibrium dynamics. The planner's problem is given by:

$$\max_{F(t), a(t)} \int_{t=0}^{\infty} e^{-rt} [U(F(t)) - D(a(t))] dt \quad (1)$$

subject to

$$0 \leq F(t) \leq K(t) \quad (2)$$

$$a(t) \geq 0 \quad (3)$$

$$\dot{R}(t) = -a(t), \quad R_0 \text{ given} \quad (4)$$

$$\dot{K}(t) = a(t)X - \lambda F(t), \quad K_0 \text{ given}, \quad (5)$$

where $F(t)$ is the rate of oil flow at time t (a choice variable), $a(t)$ is the rate at which new wells are drilled (a choice variable), $K(t)$ is the capacity constraint on oil flow (a state variable), and $R(t)$ is the measure of wells that remain untapped (a state variable). The instantaneous utility derived from oil flow is given by $U(F(t))$, where $U(\cdot)$ is strictly increasing and weakly concave; we normalize $U(0) = 0$. The total instantaneous cost of drilling wells at rate $a(t)$ is given by $D(a(t))$, where $D(\cdot)$ is strictly increasing and weakly convex, and

$D(0) = 0$. We denote the derivative of the total drilling cost function as $d(a(t))$ and assume that $d(0) \geq 0$. Utility and drilling costs are discounted at rate r .

Consistent with our empirical results from Texas, we assume a trivially low (i.e., zero) marginal cost of extraction up to the constraint.¹⁶ We ignore any fixed costs for operating, shutting in, or restarting wells because such costs are only relevant for marginally productive wells or when oil prices are very low.¹⁷ These costs should therefore not substantially affect drilling incentives, as newly drilled wells will typically only become marginal many years after drilling.

Condition (4) describes how the stock of untapped wells $R(t)$ evolves over time. The planning period begins with a continuum of untapped wells of measure R_0 , and the stock thereafter declines one-for-one with the rate of drilling. Condition (5) describes how the oil flow capacity constraint $K(t)$ evolves over time. The planning period begins with a capacity constraint K_0 inherited from previously tapped wells. The maximum rate of oil flow from a tapped well depends on the pressure in the well and is proportional, with factor λ , to the oil that remains underground. Thus, oil flow $F(t)$ erodes capacity at rate $\lambda F(t)$. The planner can, however, rebuild capacity by drilling new wells. The rate of drilling $a(t)$ relaxes the capacity constraint at rate X , where we interpret X as the maximum flow from a newly drilled well (or to be more precise, a unit mass of newly drilled wells).¹⁸ If no new wells are being drilled at t ($a(t) = 0$) and production is set at the constraint ($F(t) = K(t)$), then oil flow decays exponentially toward zero at rate λ .

The total amount of oil in untapped wells is given by $R(t)X/\lambda$, so that the total amount

¹⁶As indicated above, marginal extraction costs are not literally zero. We ignore per-barrel extraction costs from existing wells because the lack of response to oil prices for such wells implies that marginal costs are low relative to oil prices.

¹⁷Accounting for these costs would complicate the analysis substantially. We would need to model, at each t , how the quantity of oil reserves remaining in tapped wells is distributed across the continuum of tapped wells, along with the shadow opportunity cost associated with extracting more oil from every point in this distribution.

¹⁸If the drilling cost function $D(a)$ is strictly convex, the planner would never find it optimal to set up a mass of wells instantaneously at $t = 0$ —or at any other time—and the stock of untapped wells and oil flow capacity constraint would both evolve continuously over time. When the drilling cost is linear, however, such “pulsing” behavior may be optimal, leading to discontinuous changes in these state variables.

of oil underground at the outset of the planning period is given by $Q_0 = (K_0 + R_0X)/\lambda$.¹⁹ Because the flow capacity constraint is proportional to the remaining reserves, the total underground stock of oil will never be exhausted in finite time.

We assume that there is no above-ground storage of oil to focus our analysis and discussion on the implications of our model for extraction and drilling dynamics. Extending the model to include above-ground storage with an iceberg storage cost is straightforward, and we do so in appendix F.²⁰

The solution to our planner’s problem can, via the First Welfare Theorem, also be interpreted as the competitive equilibrium that would arise in a decentralized problem with continua of infinitesimally small consumers and private well owners (and no common pool problems), each of whom discounts utility or profit flows at the rate r . In a market context, marginal utility $U'(F(t))$ is equivalent to the oil price, which we denote by $P(F(t))$, and the marginal drilling cost $d(a(t))$ is determined by the rental rate for drilling rigs.²¹ Consumers have an inverse demand function $P(F)$ (equal to $U'(F)$) for oil, and well owners maximize their profits by deciding when to drill their wells, taking as given the time paths of the oil price and the rig rental rate.²² We abstract away from modeling drilling rig owners and simply assume that they rent out their scarce equipment until the marginal cost of supplying additional rentals equals the rental rate.²³ In the discussion below, we will primarily use the

¹⁹A mathematically equivalent formulation of our problem would involve imposing resource scarcity directly on the recoverable oil stock remaining by replacing condition (4) with $\dot{Q}(t) = -F(t)$, where $Q(t) = (K(t) + R(t)X)/\lambda$ is the total amount of oil remaining underground at time t . We find that our current formulation leads to necessary conditions that are easier to interpret and manipulate.

²⁰Intuitively, the presence of costly above-ground storage places an upper bound on the rate at which the oil price may be expected to increase in equilibrium. Appendix F demonstrates that our main results—most notably the weak sufficient conditions necessary to guarantee capacity constrained production on the optimal path—hold when costly above-ground storage is available.

²¹There do exist non-rig costs associated with drilling; e.g., materials and engineering costs. Thus, $d(a(t))$ can be viewed as the sum of these costs, which are invariant to $a(t)$, with the drilling rig rental cost.

²²The price-taking assumption on both sides of the market is reasonable for onshore Texas, given the existence of thousands of oil producing firms. In other areas, such as the deepwater Gulf of Mexico, only very large “major” firms participate, and these firms may be able to exert monopsony power in the rig market even if they are oil price takers. Finally, large OPEC nations such as Saudi Arabia can potentially exert market power in the global oil market.

²³A richer model would allow for investment in durable drilling rigs or explicitly model rig heterogeneity; we save this extension for future work.

language of the planner's maximization problem, though we will find it convenient to use the competitive equilibrium language (and the notation $P(F)$ rather than $U'(F)$) when we discuss well owners' production incentives, taking the oil price path as given.

Following Léonard and Long (1992), the current-value Hamiltonian-Lagrangian of the planner's maximization problem is given by:

$$H = U(F(t)) - D(a(t)) + \theta(t)[a(t)X - \lambda F(t)] + \gamma(t)[-a(t)] + \phi(t)[K(t) - F(t)], \quad (6)$$

where $\theta(t)$ and $\gamma(t)$ are the co-state variables on the two state variables $K(t)$ and $R(t)$, and $\phi(t)$ is the shadow cost of the oil flow capacity constraint.

Necessary conditions are given by equations (7) through (14) and are interpreted in sections 3.2 and 3.3 below:

$$F(t) \geq 0, U'(F(t)) - \lambda\theta(t) - \phi(t) \leq 0, \text{ comp. slackness (c.s.)} \quad (7)$$

$$F(t) \leq K(t), \phi(t) \geq 0, \text{ c.s.} \quad (8)$$

$$a(t) \geq 0, \theta(t)X - d(a(t)) - \gamma(t) \leq 0, \text{ c.s.} \quad (9)$$

$$\dot{R}(t) = -a(t), R_0 \text{ given} \quad (10)$$

$$\dot{\gamma}(t) = r\gamma(t) \quad (11)$$

$$\dot{K}(t) = a(t)X - \lambda F(t), K_0 \text{ given} \quad (12)$$

$$\dot{\theta}(t) = -\phi(t) + r\theta(t) \quad (13)$$

$$K(t)\theta(t)e^{-rt} \rightarrow 0 \text{ and } R(t)\gamma(t)e^{-rt} \rightarrow 0 \text{ as } t \rightarrow \infty. \quad (14)$$

3.2 Implications of necessary conditions for production

We begin by focusing on condition (7), which characterizes production incentives. This condition involves the co-state variable $\theta(t)$, which denotes the marginal value of capacity at time t . This value is derived from the stream of future utility obtained by optimally

producing oil from the capacity. If the optimal program calls for the capacity constraint to be binding for all times $\tau \geq t$, then $\theta(t)$ is simply given by the value of the stream of future marginal utilities $U'(F(\tau))$ discounted at the rate $r + \lambda$. If, on the other hand, it is optimal to produce below the constraint either immediately or at some point the future, then this stream of discounted marginal utilities serves as a lower bound on $\theta(t)$. Thus, we have:²⁴

$$\theta(t) \geq \int_t^\infty U'(F(\tau))e^{-(r+\lambda)(\tau-t)}d\tau, \text{ holding with equality if } F(\tau) = K(\tau) \forall \tau \geq t. \quad (15)$$

Intuitively, owners of existing capacity can do no worse than produce from their capacity as rapidly as possible. If it is optimal for them to defer production, it must be that doing so enhances the value of their capacity.

This understanding of $\theta(t)$ facilitates the interpretation of condition (7). Increasing production at time t reduces the underground pressure and hence tightens the constraint on future oil flow at rate λ . Thus, the product $\lambda\theta(t)$ captures the opportunity cost of increased flow at t in terms of forgone future utility. When the constraint on oil flow is binding, condition (7) intuitively states that the marginal utility of increased oil flow strictly exceeds the marginal cost of the diminished capacity ($U'(F(t)) > \lambda\theta(t)$), so that there is no incentive to reduce production.

From the perspective of an individual extraction firm taking the expected future oil price path as given, it is intuitive that the capacity constraint will bind whenever prices are expected to rise strictly slower than the rate of interest r . Formally, it is clear from equation (15) that $P(F(t)) > \lambda\theta(t)$ in this case (substituting $P(F(t))$ for $U'(F(t))$ in (15)). But what if the oil price is expected to rise strictly faster than r ? If this rate of price increase is expected to persist forever, then equation (15) makes clear that we would have $P(F(t)) < \lambda\theta(t)$, so that the firm would reduce its production (in fact, since all firms would have this incentive, equilibrium prices would adjust so that $P(F(t)) = \lambda\theta(t)$).

²⁴Derivation: Use equation (7) to eliminate $\phi(t)$ in equation (13). The resulting linear first-order differential equation, in conjunction with the endpoint condition (14), can then be solved to obtain (15).

What if the oil price is expected to temporarily rise faster than r and then level off, as was the case during 1998–1999? In this case, the firm has an incentive to defer production to the point in time at which future prices are expected to be greatest in present value. However, the capacity constraint does not allow this arbitrage: any production that is deferred today cannot be completely recovered at the instant it is expected to be most valuable.²⁵ Instead, it must be recovered over the full remaining life of the well, including the time period when the oil price is expected to be lower than the current price in present value. Thus, the length of time for which the oil price is expected to rise faster than r must be fairly large in order for unconstrained oil production to be value-maximizing.

To verify that producers never had an incentive to set $F(t) < K(t)$ during our 1990–2007 sample, we use the futures price data and equation (15) to calculate $\theta(t)$ for each month of the sample. This calculation is carried out via a backward recursion procedure and is discussed in detail in appendix B.²⁶ Figure 5 plots our calculation of $\lambda\theta(t)$, the marginal value of deferred production, along with the spot price of oil from 1990–2007. At no time during the sample was withholding production for subsequent sale value-maximizing, including 1998–1999.

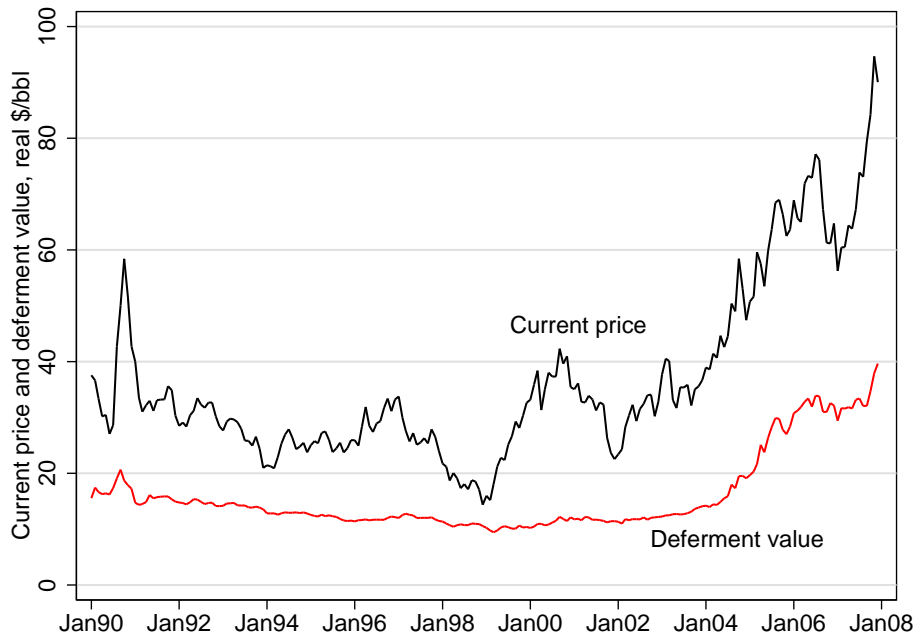
Thus, our model explains why producers did not respond on the intensive margin to in-sample expectations that the oil price would temporarily rise faster than r .²⁷ In general, however, our model does not rule out the possibility that, given appropriate initial conditions and functional forms for $U(F)$ and $d(a)$, unconstrained production may be optimal in the

²⁵Suppose that production is reduced below the constraint by an amount $\epsilon > 0$ for a time interval of length $\delta > 0$. Then, the total amount of oil production deferred equals $\epsilon\delta$, and the available production capacity after this time interval will be $\lambda\epsilon\delta$ greater than it otherwise would have been. This additional capacity is not infinite, so the entire deferred volume cannot be extracted immediately. The fastest way to extract the deferred production is to produce at the capacity constraint, in which case the rate of production declines exponentially at rate λ , and the deferred production is only completely recovered in the limit as $t \rightarrow \infty$.

²⁶The backward recursion procedure takes account of the possibility that firms may find producing below the constraint to be optimal at some future date. We use a value of 10% for both λ and r . Setting $\lambda = 0.1$ is consistent with our main empirical results and stylized facts reported in Thompson (2001) and Mason and van't Veld (2013). Setting $r = 0.1$ is consistent with a survey of oil producers during our sample period (see Society of Petroleum Evaluation Engineers (1995) and the discussion in Kellogg (2014)).

²⁷Non-zero fixed operating costs can rationalize shutting in production of wells with very low production rates when oil prices are low, as observed in 1998–1999. In this case, these wells' production revenue may not have been sufficient to cover their fixed costs. Alternatively, even if the fixed costs were covered, their presence may have made shutting in low-volume wells profitable by arbitraging the expected increase in future oil prices (given sufficiently low costs of re-starting production in the future).

Figure 5: Selling a barrel of oil vs. deferring its production



Note: This figure shows the crude oil front month (“spot”) price and the value of one barrel of deferred production in each month, all in real 2007 dollars. See text and appendix B for details.

solution to the planner’s problem. For instance, suppose that $K_0 > 0$, there are no new wells remaining to be drilled, and oil demand is constant elasticity, with an elasticity sufficiently small that if production declines exponentially at rate λ , marginal utility rises faster than r . In this case, it is optimal to produce below the constraint forever, following the standard Hotelling path.²⁸ In section 4.2, we discuss conditions on primitives—such as the value of K_0 and the shape of the oil demand curve—under which production is optimally constrained throughout the entire equilibrium path.

Note that whenever the constraint on oil flow is slack (assuming $F(t) > 0$), we have $\phi(t) = 0$, so that the marginal utility of increased oil flow exactly equals the marginal opportunity cost of increased extraction ($U'(F(t)) = \lambda\theta(t)$). Moreover, condition (13) then

²⁸By assumption, $U'(F) = aF^{-\eta}$, where $\eta\lambda > r$. If production is always unconstrained, $\dot{U}'/U' = r$, requiring $F(t) = F(0)e^{-rt/\eta}$. This production program is optimal if all reserves are extracted in the limit and $F(t) \leq K(t)$ for all t . Complete extraction in the limit requires that $K_0/\lambda = F_0\eta/r$, which implies that $F_0 < K_0$. This argument applies at any given starting time t_0 , so this production program is optimal.

implies that $\dot{\theta}(t) = r\theta(t)$, so that $\theta(t)$ rises at r , and then by condition (7) we have $U'(F(t))$ rising at r as well. Thus, whenever production is unconstrained, marginal utility rises at the discount rate as in the standard Hotelling model with zero extraction costs.

3.3 Implications of the necessary conditions for drilling: a modified Hotelling rule

Condition (9) characterizes drilling incentives. The $\theta(t)X$ term is the value of the X units of capacity created by drilling a new well, and the $d(a(t))$ term is the marginal cost of drilling a well at time t when the rate of drilling is $a(t)$. The term $\gamma(t)$ can be interpreted as the shadow value of the marginal undrilled well at t . Condition (11) implies that $\gamma(t) = \gamma_0 e^{rt}$, where $\gamma_0 \geq 0$ is a constant. Intuitively, the stock of undrilled wells acts as a store of value, so that the value of the marginal undrilled well must increase at the discount rate. Thus, when drilling occurs ($a(t) > 0$), conditions (9) and (11) together imply that the marginal return to drilling must rise at r :

$$\theta(t)X - d(a(t)) = \gamma_0 e^{rt}. \quad (16)$$

Condition (9) is analogous to the standard Hotelling rule, which states that, constrained to a fixed volume of oil, the planner should extract so that the net marginal value of extracting *barrels* rises at the rate of interest. Thus, every barrel extracted yields the same net payoff in present-value terms. Since this is a drilling problem, however, the Hotelling-like intuition applies to *wells*, not to barrels.

Equation (16) holds whenever drilling occurs. If the right-hand side is strictly larger than the left-hand side when $a(t) = 0$, it is optimal to refrain from drilling. A mundane way in which this situation may occur is if $d(0)$ is large relative to $U'(0)$. However, even if $d(0)$ is small relative to $U'(0)$, drilling must cease in finite time whenever marginal utility is bounded (as in the case of a substitutable backstop technology), for in that case we have

$e^{-rt}[\theta(t)X - d(0)] \rightarrow 0 < \gamma_0$ as $t \rightarrow \infty$.

Intuitively, if the marginal drilling cost is strictly positive, it seems that production should be capacity constrained whenever drilling is occurring: why spend resources to create capacity if that capacity is not being used? However, our model permits the oil flow constraint to be slack during drilling if the marginal drilling cost is increasing at r . To see this result, combine equation (16) with the result from section 3.2 that $\theta(t)$ must increase at r when production is unconstrained. In this scenario, the planner is indifferent between drilling immediately versus in the future, even though the new capacity is not fully utilized. Conversely, whenever $d(a(t))$ is rising at a rate strictly slower than r on the optimal path, production must be at capacity as the new wells come on line.²⁹

For the remainder of this section we focus on the empirically relevant case in which production is constrained during drilling. In this case, the necessary conditions can be manipulated to yield:³⁰

$$U'(F(t)) - \left[\frac{(r + \lambda)d(a(t))}{X} - \frac{d'(a(t))\dot{a}(t)}{X} \right] = \frac{\lambda\gamma_0}{X}e^{rt}. \quad (17)$$

Equation (17) holds whenever production is constrained and $a(t) > 0$, and it can be interpreted as the modified per-barrel Hotelling rule for our model. On the right-hand side of (17) we have the shadow value of wells ($\gamma_0 e^{rt}$) divided by the total amount of oil stored in a unit mass of untapped wells (X/λ), which we can interpret as the per-barrel shadow value of oil in untapped wells. On the left-hand side is the marginal utility of oil less a term in square brackets reflecting the cost of the additional production. The first term in square brackets is the amortized, per-barrel marginal cost of drilling a well at time t .³¹ The last term in square

²⁹In this case, condition (16) implies that $\dot{\theta}(t)/\theta(t) < r$, and (13) then implies that we must have $\phi(t) > 0$.

³⁰Derivation: If $F(t)$ is strictly positive, conditions (7) and (13) allow us to eliminate $\phi(t)$ from the system and obtain expression (i): $\dot{\theta}(t) = -U'(F(t)) + (r + \lambda)\theta(t)$. Differentiating condition (9) with respect to time and solving for $\dot{\theta}(t)$ yields expression (ii): $\dot{\theta}(t) = \frac{r\gamma(t) + d'(a(t))\dot{a}(t)}{X}$. Finally, we substitute for $\theta(t)$ and then $\dot{\theta}(t)$ in condition (16) using expressions (i) and then (ii), and assuming $a(t) > 0$, to yield equation (17).

³¹To clarify, consider the special case of constant marginal drilling costs: $d(a(t)) = \bar{d}$. If the planner drills one well (or rather, she marginally increases the rate of drilling) she obtains a marginal increase in oil flow h instants later of $Xe^{-\lambda h}$, assuming oil flow is set to the maximum. If each barrel of flow has imputed cost

brackets captures the opportunity cost of drilling now versus waiting, which arises due to the convexity in the drilling cost function. When $\dot{a}(t) < 0$, drilling activity and marginal drilling costs are falling over time, so that drilling immediately incurs an additional opportunity cost relative to delaying. One implication of this term is that, in contrast to standard Hotelling models in which production must decline over time (with the marginal extraction cost also declining if extraction costs are convex), oil flow in our model can *increase* over intervals during which the marginal cost of drilling is falling (i.e., $\dot{F}(t) > 0$ and $\dot{a}(t) < 0$ holding simultaneously is possible).

If drilling costs are affine rather than strictly convex, so that the last term in square brackets drops out of equation (17), then this equation becomes the Hotelling Rule for a standard barrel-by-barrel extraction model with a constant marginal extraction cost of $(r + \lambda)d/X$. In fact, appendix C shows that if $d'(a) = 0$, and if $a(t) > 0$ and $F(t) = K(t)$ throughout the entire optimal path, then our model and the standard Hotelling model yield identical production and marginal utility paths.

A standard Hotelling model can therefore yield the same predictions as our reformulated model, but the conditions necessary for this equivalence are not realistic. First, equivalence requires the absence of unanticipated shocks along the equilibrium path. Oil production will immediately respond to such shocks in a standard Hotelling model, but the binding capacity constraint precludes this response in our model (as we discuss in section 4.3). Second, the assumption that $d'(a) = 0$ is clearly at odds with the data on rig rental prices shown in figure 4. In addition, the requirement that $a(t) > 0$ throughout the entire optimal path requires that $U'(F)$ be unbounded—an assumption seemingly in conflict with the existence of alternative fuels and technologies.³² For the remainder of the paper, we therefore focus our attention on the dynamics of production, drilling, and prices assuming that $d'(a) > 0$

of c at that time, then at the time the well is drilled, the flow at h would have an imputed cost $cXe^{-(r+\lambda)h}$. Since such flows continue indefinitely, we want to find c such that $\bar{d} = \int_{h=0}^{\infty} cXe^{-(r+\lambda)h}dh$. Integrating, we obtain $\bar{d} = cX/(r + \lambda)$. Solving for c , we conclude: $c = \bar{d}(r + \lambda)/X$.

³²Appendix C discusses two additional necessary conditions for this special case— K_0 must be sufficiently small and λ must be sufficiently large—that are more likely to be satisfied empirically.

and that $U'(F)$ is bounded (with the former being more qualitatively important), obtaining implications that match the features of the data presented in section 2 and that differ from those of prior work.

4 Equilibrium production, price, and drilling dynamics

This section characterizes the dynamics implied by the solution to the optimal control problem presented in section 3. Throughout, we interpret our results in terms of a competitive equilibrium to relate our findings to data on oil prices and rig rental prices. We begin in section 4.1 by considering an individual oil field that is small relative to the global market, so that the path of oil prices is exogenous but the local rental market for drilling rigs clears at each instant. Section 4.2 then endogenizes the oil price path to consider the general equilibrium dynamics implied by our model.³³ Finally, section 4.3 considers the impact of unexpected demand shocks on equilibrium.

4.1 Exogenous oil prices

This section studies industry supply in isolation. This case is particularly relevant for interpreting drilling and extraction behavior in a small, local region such as Texas. We treat the time path of oil prices $P(t)$ as given exogenously, while the path of marginal drilling costs $d(a(t))$ is endogenously determined by the market-clearing condition for drilling rigs. This condition is given by equation (16) and requires that well owners be indifferent about when to drill for all times t such that $a(t) > 0$ (this intuition also applies in later subsections when we endogenize the oil price path).

We focus on the case in which the oil price is expected to equal a constant \bar{P} forever. In this case, equation (15) implies that $\theta(t)X = \frac{\bar{P}X}{r+\lambda}$. Since $\dot{\theta}(t)/\theta(t) = 0 < r$, conditions (8)

³³This section is “general” relative to the partial equilibrium analysis in section 4.1; however, like virtually all of the Hotelling literature, we do not model income effects generated by the rents earned by well owners (and in our case, by the owners of drilling rigs).

and (13) imply that $F(t) = K(t)$ for all $t \geq 0$: production is always capacity constrained. Moreover, drilling activity must cease in finite time.

Substituting for $\theta(t)$ in equation (16), the arbitrage condition is given by:

$$d(a(t)) = \frac{\bar{P}X}{r + \lambda} - \gamma_0 e^{rt}. \quad (18)$$

Equation (18) indicates that if drilling for oil is profitable ($\gamma_0 > 0$), then the marginal drilling cost must decline over time so that drilling remains equally attractive over all times with $a(t) > 0$. Because $d'(a) > 0$, the rate of drilling $a(t)$ must therefore decline over time.

Equilibrium also requires that drilling cease ($a(t) = 0$) when the initial supply R_0 of untapped wells is exhausted, and not beforehand.³⁴ This exhaustion condition determines the value of γ_0 , the present value of an undrilled well.³⁵ Equation (18) implies that the time of exhaustion T must satisfy $d(a(T)) = \bar{P}X/(r + \lambda) - \gamma_0 e^{rT}$, with $a(T) = 0$.

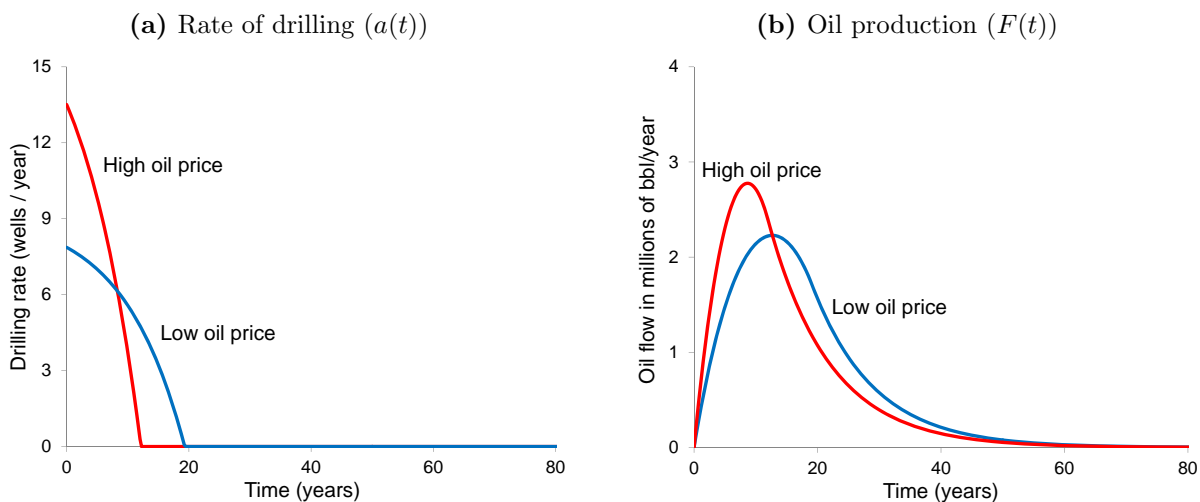
Figure 6 illustrates the equilibrium time paths of drilling and production given an initial capacity of zero ($K_0 = 0$) and an initial stock of undrilled wells $R_0 > 0$. Drilling activity ($a(t)$) follows the dynamics prescribed by equation (18) and ceases at T defined above. Drilling activity is intense at first but declines until it ceases at T , when all wells have been drilled. Oil production $F(t)$ follows the dynamics prescribed by equation (12): $F(t)$ begins at zero, initially increases as the stock of drilled wells increases, but then eventually declines asymptotically toward zero as the remaining volume of oil reserves (and reservoir pressure) declines. The resulting hump-shaped production profile is a well-known feature of production in oil fields across the world (Hamilton 2013). We show in section 4.2 below that a similar “peak oil” result also emerges in the case of endogenous oil prices.

Figure 6 also depicts how the equilibrium dynamics vary with the exogenous oil price \bar{P} . Because the total stock of wells is fixed, the oil price can affect the timing of drilling

³⁴The intuition for this statement is obvious, but we supply a formal proof for the general equilibrium case as part of appendix D.

³⁵To see this, observe that a very high value of γ_0 will result in the drilling of fewer than R_0 wells by the time T at which $a(T) = 0$. In contrast, a low value of γ_0 will cause the undrilled well stock R_0 to be exhausted at a time T for which $a(T) > 0$.

Figure 6: Optimal drilling in a local region



Note: This figure illustrates the equilibrium time paths of drilling rates (panel a) and oil production (panel b), with exogenously given constant oil price expectations of \$50/bbl (“low oil price”) or \$100/bbl (“high oil price”). The figure assumes an initial extraction capacity of $K_0 = 0$, an initial well stock of $R_0 = 100$, initial reserves per well of 0.5 million bbl, a decline rate of $\lambda = 0.1$, and a discount rate of $r = 0.1$. The marginal drilling cost curve is given by $d(a) = 1 + 5a$, with a in units of wells drilled per year and d in \$million/well.

and production but not total cumulative drilling and production. For a relatively high oil price, the initial rate of drilling (and the rig rental rate) must be relatively high, and γ_0 must therefore also be relatively high to satisfy the exhaustion constraint.³⁶ Thus, a higher oil price causes the stock of wells to be drilled relatively quickly, shifting oil production earlier in time (if the marginal drilling cost $d(a)$ is affine, as assumed in figure 6, it can be shown that peak production occurs sooner with a relatively high oil price). Moreover, because γ_0 is greater, a higher oil price increases the aggregate wealth derived from drilling the wells.³⁷

Note that this analysis comparing drilling paths under low and high price scenarios can be reinterpreted as an analysis of the response of drilling to an unanticipated price shock.

³⁶To see this, suppose γ_0 were unchanged (or reduced) from the low oil price case. In this case, equation (18) would require a uniformly higher rate of drilling over a longer period. This path would, therefore, call for more wells to be drilled than the available stock and cannot be an equilibrium. Increasing the value of γ_0 above its level in the low price case shortens the time interval over which wells are drilled so that the total number of wells drilled equals R_0 .

³⁷Each owner of an untapped well is indifferent over when he rents a drilling rig to drill his well. Hence, the wealth of the owners of the initially untapped wells is $\gamma_0 R_0$.

Suppose that the price of oil is expected to remain low forever but then, at $t = 0$, suddenly and unexpectedly increases to a higher level, where it is then expected to persist. Drilling and rig rental rates immediately jump up at $t = 0$, as the dynamics switch over to the drilling path corresponding to a higher oil price. Thus, the model replicates the covariance of oil prices, drilling, and rig rental rates that we observe in our Texas data.

4.2 Equilibrium dynamics with endogenous oil prices

We now close the model by endogenizing the path of oil prices and requiring that at every instant demand equals supply. We begin by characterizing the dynamics of oil production, drilling activity, oil prices, and marginal drilling costs using an intuitive phase diagram analysis. We then present two theorems that establish sufficient conditions under which the dynamics are well-behaved and production is always capacity constrained. After illustrating these equilibrium dynamics using a specific numerical example, we discuss the limiting case of the model in which the total number of wells to be drilled is infinite.

4.2.1 Phase diagram characterization of the equilibrium dynamics

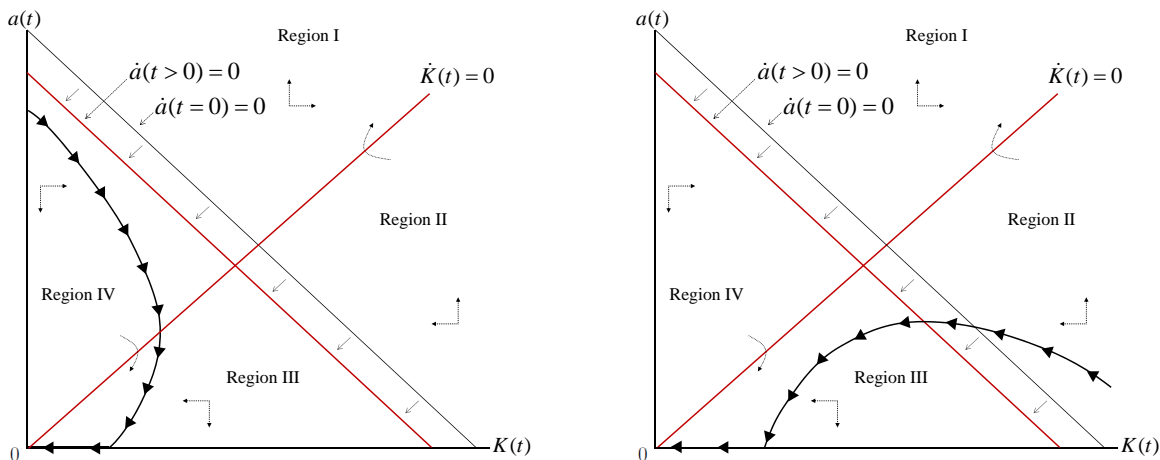
We construct our phase diagram in (K, a) space. The first step in this analysis is to identify the loci of points for which $\dot{a}(t) = 0$ or $\dot{K}(t) = 0$. We assume that inverse demand ($P(F)$) is strictly decreasing, that drilling supply ($d(a)$) is strictly increasing, that both functions are continuously differentiable, and that drilling occurs ($a(t) > 0$ for at least some t) in equilibrium (we discuss sufficient conditions for this last assumption in section 4.2.2 below). Given $a(t) > 0$ and $F(t) > 0$, we can deduce from conditions (12) and (17) equations that describe the $\dot{a}(t) = 0$ and $\dot{K}(t) = 0$ loci:³⁸

$$\dot{a}(t) = 0 : a(t) = d^{-1} \left(\frac{XP(F(t))}{r + \lambda} - \frac{\lambda\gamma_0}{r + \lambda} e^{rt} \right) \quad (19)$$

$$\dot{K}(t) = 0 : a(t) = \lambda F(t)/X. \quad (20)$$

³⁸We may use equation (17) here because production must be capacity constrained whenever $\dot{a}(t) = 0$.

Figure 7: Phase diagrams for endogenous price model



Note: This figure shows phase diagrams for the general equilibrium model. The left panel sketches a potential equilibrium drilling and extraction path starting from $K_0 = 0$, while the right panel sketches a path starting from a relatively large inherited capacity. See text for details.

Phase diagrams depicting these two loci, along with sketches of equilibrium paths, are presented in figure 7. The left panel depicts a path starting from $K_0 = 0$, while the right panel depicts a path starting from a relatively large K_0 .³⁹ Note that production must be capacity constrained along the $\dot{a}(t) = 0$ locus, since $a(t) > 0$ and $d(a)$ cannot be rising at the interest rate. Thus, we can replace $F(t)$ with $K(t)$ in equation (19), revealing that the $\dot{a}(t) = 0$ locus must be downward sloping in (K, a) space (since $P(F)$ is strictly decreasing).⁴⁰ This locus is not time-stationary. The presence of the $\gamma_0 e^{rt}$ term in equation (19) causes the $\dot{a}(t) = 0$ locus to shift downward over time, as shown in the figure. Below the $\dot{a}(t) = 0$ locus we have $\dot{a} < 0$, while above the $\dot{a}(t) = 0$ locus we have $\dot{a} > 0$.

The $\dot{K}(t) = 0$ locus is given by an upward-sloping ray from the origin, as shown, whenever capacity is constrained. Note that capacity must be constrained when the $\dot{K}(t) = 0$ locus is below the $\dot{a}(t) = 0$ locus, since drilling costs are decreasing in this region. Above the $\dot{a}(t) = 0$ locus, if $F(t) < K(t)$ then the $\dot{K}(t) = 0$ locus will curve to the right of the ray depicted in

³⁹Historically, we must of course have $K_0 = 0$. However, a sufficiently large negative demand shock could interrupt the original equilibrium path, causing the new path to begin in region II, region III, or even on the $a(t) = 0$ axis.

⁴⁰The $\dot{a}(t) = 0$ locus will be linear, as shown, under linear demand and marginal drilling costs.

figure 7. Below the $\dot{K}(t) = 0$ locus we have $\dot{K} < 0$, while above the $\dot{K}(t) = 0$ locus we have $\dot{K} > 0$.

The $\dot{a}(t) = 0$ and $\dot{K}(t) = 0$ loci, along with the axes, form boundaries that divide the phase diagram into four interior regions. Equilibrium dynamics are dictated by behavior at the boundaries and within these regions. Intuitively, neither region I nor its boundaries can be entered in equilibrium, since in this case drilling and capacity increase without bound, leading to a violation of the necessary conditions once all wells are drilled. Thus, starting from $K_0 \geq 0$ sufficiently small, as depicted in the left panel of figure 7, the equilibrium path begins in region IV (and on the $K = 0$ axis if $K_0 = 0$). The rate of drilling $a(t)$ is initially high but decreases over time as the extraction rate builds (and the oil price falls). Eventually, once the drilling rate becomes sufficiently small, region III is entered, and the extraction rate decreases while the oil price rises. Thus, there is a peak in oil production. Eventually, all wells are drilled (since $P(0)$ is finite), and the dynamics follow the $a = 0$ axis to the origin as capacity is gradually exhausted. If, on the other hand, K_0 is large, then the equilibrium dynamics will begin in region II or region III. If region II, the rate of drilling will initially increase over time as capacity falls, as shown in the right panel of figure 7. Eventually, however, the dynamics must enter region III, with drilling activity falling until all wells are drilled.

Note that this model nests the exogenous, constant oil price case from section 4.1. In this special case, the $\dot{a} = 0$ locus becomes a horizontal line that shifts downward over time. Region II can never be entered, so that we always have $\dot{a} < 0$, no matter the initial capacity K_0 . In fact, K_0 does not affect the time path of drilling at all when price is exogenous; all that matters is the initial stock of wells remaining to be drilled.

4.2.2 Formal rules governing the equilibrium dynamics

We formalize the restrictions (or “rules”) that equilibrium imposes on the dynamics of oil drilling, production, and prices in two theorems. These theorems establish sufficient con-

ditions under which equilibrium drilling and extraction paths are well-behaved and follow either the path sketched in the left panel of figure 7 with production constrained throughout, or in the right panel with production constrained beginning from some time before the drilling rate reaches its zenith. Here, we provide the key intuition for these results. See appendices D and E for formal proofs.

Theorem 1 only requires strict monotonicity and continuous differentiability of inverse oil demand and marginal drilling costs, along with an assumption that drilling costs are sufficiently small that drilling will actually occur. Formally, we assume that: (i) the marginal cost of drilling $d(a)$ is continuously differentiable, with $d'(a) > 0$ and $d(0) > 0$; (ii) the inverse demand for oil is given by $P(F) \geq 0$, with $P(0) > 0$, $P(F)$ continuous for $F > 0$ (and continuous at $F = 0$ if $P(0) < \infty$), $P(F) \rightarrow 0$ as $F \rightarrow \infty$, and $P(F)$ continuously differentiable, with $P'(F) < 0$, whenever $P(F) > 0$; (iii) inverse demand and drilling supply are such that $XP(0)/(\lambda + r) > d(0)$; and (iv) initial conditions are such that $\infty > R_0 > 0$ and $\infty > K_0 \geq 0$. Then the following rules hold:

Rule 1: $\gamma(t)$, $\dot{\gamma}(t)$, $\theta(t)$, $\dot{\theta}(t)$, $R(t)$, $\dot{R}(t)$, $K(t)$, $\dot{K}(t)$, $a(t)$, $F(t)$, $P(t)$, and $\phi(t)$ are continuous for all $t > 0$, as is $\dot{a}(t)$ when $a(t) > 0$. Continuity of these variables generally follows from the strict monotonicity and continuous differentiability of $P(F)$ and $d(a)$ and in some cases follows directly from the necessary conditions (e.g., for θ, γ and $\dot{\gamma}$).⁴¹

Rule 2: Capacity $K(t)$ must go to zero in the limit as $t \rightarrow \infty$. This rule follows from the fact that demand is strictly positive while the marginal extraction cost is zero. Thus, it is not optimal to let capacity be unused in the limit.

Rule 3: The drilling rate $a(t)$ and capacity $K(t)$ cannot both be weakly increasing (i.e., in the phase diagram, neither the interior of region I nor its boundaries defined by the $\dot{a} = 0$ and $\dot{K} = 0$ loci may be entered), nor can $a(t)$ and $F(t)$ both be weakly increasing. This

⁴¹In many phase portraits, the *costate* variable is on the vertical axis. Since the Pontryagin necessary conditions for optimality require that the costate variables be continuous functions of time, there can be no vertical jumps in an optimal trajectory depicted in the phase portrait. In our case, however, a *choice* variable $a(t)$ is on the vertical axis and we must *prove* that it is a continuous function of time to rule out such vertical jumps.

rule follows from the fact that region I cannot be escaped in the phase diagram without generating a discontinuity in drilling at the moment the last well is drilled.⁴²

Rule 4: All available wells will be drilled in the limit as $t \rightarrow \infty$. This rule follows immediately from rule 2 and the fact that strictly positive drilling is always strictly profitable as production approaches zero, since $P(0)X/(r + \lambda) > d(0)$.

Rule 5: At any time that drilling ceases, production must be constrained for a measurable period afterward. This rule follows from the intuition that it would be sub-optimal to drill a costly well and then not fully utilize that well's capacity immediately after drilling.

Rule 6: Given demand $P(F)$, marginal drilling costs $d(a)$, and initial reserves R_0 , there exists $\delta > 0$ such that if the initial capacity $K_0 < \delta$, then drilling will begin instantly and will initially be strictly decreasing over time. This rule follows from the fact that, with initial capacity sufficiently low, we must begin in region IV of the phase diagram.

Together, rules 1–6 imply that if $K_0 \geq 0$ is sufficiently small, then the drilling and extraction paths must start in region IV and then transition to region III. In both of these regions, production must be constrained. However, it is not clear whether production may enter region II from region III (since the $\dot{a}(t) = 0$ locus is shifting inward) and, if so, whether production is unconstrained while in region II. It is also not clear whether, once drilling ceases, production is constrained thereafter. Similarly, a path that begins in region II with $K_0 > 0$ sufficiently large must transition to region III, but it is unclear whether it could ever reenter region II for an interval.

Theorem 2 allows us to resolve these remaining questions with only a modest assumption on the shape of the demand curve. In addition to the assumptions of theorem 1, we assume that the inverse elasticity of demand given by $\eta(F) \equiv -F \frac{P'(F)}{P(F)}$ is weakly increasing in F (i.e., the inverse demand curve is weakly concave in logs). This property is satisfied by nearly all single-product demand curves used in applied work.⁴³ Given this additional assumption,

⁴²In appendix D, we show that region I cannot be entered even if reserves are infinite: entering region I must ultimately lead either to a violation of the transversality condition or to a discontinuity in $F(t)$, neither of which is possible in equilibrium.

⁴³In particular, it is easy to verify that this property is satisfied by any inverse demand curve of the form

the following rules hold:

Rule 7: If at any time production is constrained, then production must be constrained forever after. This rule follows intuitively from the fact that, if production is ever constrained in the optimal program, the incentive to defer production will only subsequently diminish. To see this, consider a special case in which production is constrained and no drilling occurs, such that $\dot{P}/P = \lambda\eta(K)$. In this case, as capacity declines over time, the demand assumption implies that the inverse demand elasticity will also decline, so that the rate of price increase declines and the incentive to defer production diminishes.

Rule 8: Given demand $P(F)$, marginal drilling costs $d(a)$, and initial reserves R_0 , there exists $\delta > 0$ such that if initial capacity $K_0 < \delta$, then production is constrained along the entire optimal path (so production is always constrained trivially if $K_0 = 0$). This rule follows immediately from rules 6 and 7.

Rule 9: Once the rate of drilling strictly declines it can never subsequently strictly increase, further implying that once drilling stops it cannot restart. This rule follows from the intuition that once the drilling and production path is in region III, the decline in the inverse demand elasticity provides no incentive to subsequently increase drilling, so that region II will not be entered.

Rule 10: Drilling will end in finite time if and only if there exists an $F > 0$ such that $\lambda\eta(F) < r$. This final rule follows from the demand assumption and the previous rules, which imply that capacity and production must approach zero in the limit. If price is constrained to eventually rise more slowly than the rate of interest as production declines (as must be true if $P(0)$ is finite), then all of the wells will be drilled in finite time. Otherwise, it is optimal to drag drilling out indefinitely.⁴⁴

Together, rules 1–10 imply that for $K_0 \geq 0$ sufficiently small, production must be con-

$P(F) = \alpha - \beta F^\delta$, with either $\alpha > 0, \beta > 0, \delta > 0$ or $\alpha \leq 0, \beta < 0, \delta < 0$. The first set of parameters encompasses a wide array of concave and convex demands with a finite $P(0)$, while the second set of parameters allows for $P(0)$ to be infinite (and if $\alpha = 0$, demand is constant elasticity).

⁴⁴Note that while we emphasize the empirically relevant case of $P(0) < \infty$, rule 10 and the other rules proven in the appendix actually permit the case of $P(0) = \infty$, which is sometimes considered in the literature.

strained throughout the entire equilibrium path depicted in the left panel of figure 7. In the right panel, with $K_0 > 0$ sufficiently large, it may be that production is initially unconstrained in region II. However, production must become constrained before the path enters region III, and from that point production will be constrained forever.

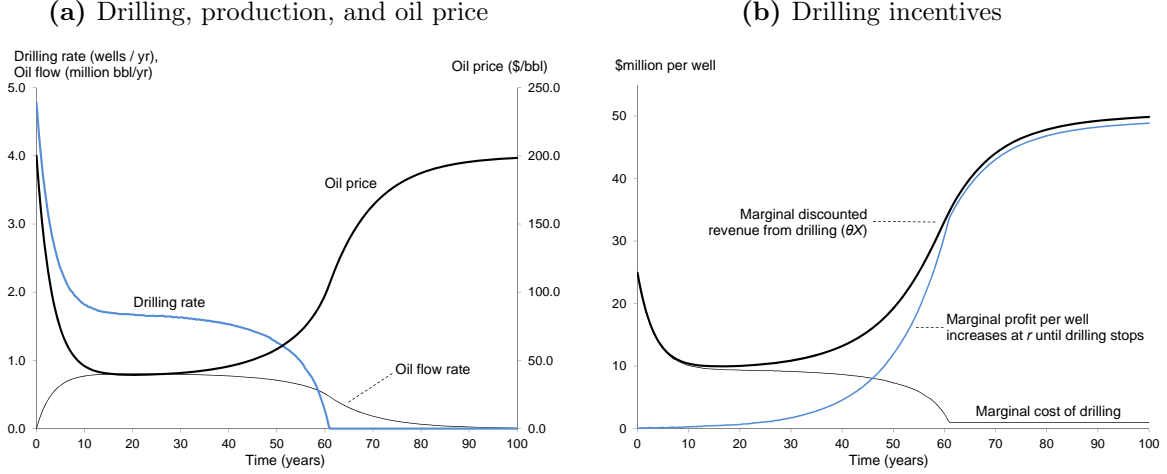
4.2.3 A specific example with linear demand and marginal drilling costs

Figure 8 depicts equilibrium dynamics for a specific example with $K_0 = 0$, linear demand, and linear marginal drilling costs, satisfying the conditions of theorem 2. This model does not permit an analytic solution, so we solve the model computationally using value function iteration; see appendix G for details. The model assumes $r = \lambda = 0.1$, consistent with survey evidence from Society of Petroleum Evaluation Engineers (1995) and with the decline rate estimated from our Texas dataset. The inverse demand curve is given by $P(F) = 200 - 200F$ (with F measured in million bbl/day), and marginal drilling costs are given by $d(a) = 1 + 5a$ (with d in million \$/well and a in wells/year). We simulate the optimal drilling and production of an initial well stock of $R_0 = 100$, with 0.5 million bbl of reserves per well.

Panel (a) illustrates that the rate of drilling is initially high but decreases rapidly. The rate of decrease slows to near-zero between years 10 and 40, but then accelerates so that all drilling is completed in year 61. Thus, the rate of oil production rapidly rises from zero at $t = 0$ to a plateau-like peak and then enters a decline as the drilling rate falls to zero. The path of oil prices is therefore U-shaped.

Panel (b) depicts the time path of drilling incentives. The marginal discounted revenue from drilling $\theta(t)X$ initially falls as the oil price falls but then rises, asymptotically approaching \$50 million per well (equal to the asymptotic oil price of \$200/bbl multiplied by $X = 0.05$ million bbl/d and divided by $r + \lambda$). Marginal drilling costs $d(a(t))$ fall with the rate of drilling so that the marginal profit per well (equal to $\theta(t)X - d(a(t))$) rises at the interest rate until drilling ceases, adhering to necessary condition (16).

Figure 8: Equilibrium paths with linear demand and marginal drilling costs



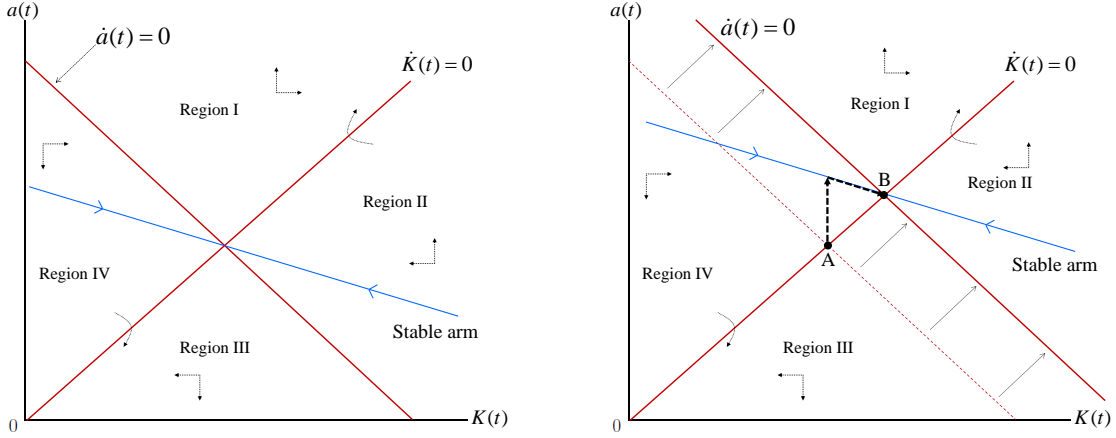
Note: This figure illustrates the equilibrium time paths of drilling, oil production, and prices for a model with a linear inverse demand curve given by $P(F) = 200 - 200F$ (F in million bbl/d) and marginal drilling costs given by $d(a) = 1 + 5a$ (with d in million \$/well and a in wells/year). The initial well stock is $R_0 = 100$, with 0.5 million barrels of reserves per well, and $r = \lambda = 0.1$. See text and appendix G for details.

4.2.4 Equilibrium dynamics with an unlimited number of wells

In the specific example depicted in figure 8, the period of nearly constant drilling and production between years 10 and 40 is driven by the fact that the initial reserves are large relative to the profitability of drilling. γ_0 is therefore very small so that, for small values of t , our modified Hotelling condition (16) can hold with nearly constant drilling and production rates. This outcome is closely related to the behavior of our model in the limit as $R_0 \rightarrow \infty$. In this limit, $\gamma_0 \rightarrow 0$ so that the $\dot{a}(t) = 0$ locus is stationary in our phase diagram. Thus, there is a steady state, with strictly positive and constant drilling and production rates, at the intersection of the $\dot{a}(t) = 0$ locus with the $\dot{K}(t) = 0$ locus.

The left panel of figure 9 presents a phase diagram depicting this steady state and the equilibrium dynamics for this limiting case. Since the production constraint binds when $\dot{a} = 0$, we can use equation (19) to solve for marginal drilling costs in the steady state, obtaining $d(a^*) = P(F^*)X/(r + \lambda)$, where the stars denote steady-state values. On the right we have the discounted stream of revenues generated by a well at constant price $P(F^*)$.

Figure 9: Phase diagrams for endogenous price model with infinite well stock



Note: This figure shows phase diagrams for the general equilibrium model in which there is an infinite supply of wells and therefore $\gamma(t) = 0$. The panel on the left shows the $\dot{a} = 0$ and $\dot{K} = 0$ loci, the steady state located at the intersection of these loci, the directions of motion within each region, and the stable arms leading to the steady state. The panel on the right illustrates the equilibrium path following a positive shift in oil demand $P(F)$. Such a shift causes the $\dot{a}(t) = 0$ locus to shift outward, moving the steady state from its original position at point A to point B, and causing the stable arm to shift outward. See text for details.

Thus, the marginal well earns just enough revenue to cover its drilling cost, and the market is in a long-run equilibrium.

Figure 9 also depicts the stable arm (saddle path) running northwest to southeast through the steady state. In equilibrium, from any initial K_0 , drilling and capacity will follow this stable arm to the steady state.⁴⁵

Finally, note that with an unlimited stock of wells our model bears a strong resemblance to a standard macroeconomic Q-theory model of optimal investment in an industry with convex adjustment costs, where K is capital and F is production. Both models feature a steady state in which the marginal cost of investment equals the marginal value of installed capital, which is derived from the capital's discounted marginal revenue stream. The key differences

⁴⁵A path starting above the stable arm will ultimately lead drilling and capacity into region I, from which there is no escape. A path starting below the stable arm will ultimately lead to a cessation of drilling and movement along the horizontal axis toward the origin. Given our assumption that $d(0) < \frac{XP(0)}{r+\lambda}$, any such path will eventually violate the condition that, in the absence of drilling, $d(0) \geq \theta X$ (this last point follows from equation (15), since on any such path $\theta \rightarrow P(0)/(r + \lambda)$).

are that: (1) production in the macroeconomic model typically features diminishing returns to capital, whereas in our model $F = K$ (there is no labor input); (2) in the macroeconomic model there is never a reason to produce less than what the available inputs permit; and (3) in our model capital depreciates with use rather than deterministically over time. This last difference effectively vanishes when production in our model is constrained.

4.3 Impacts of unanticipated demand shocks

In this section, we explore how persistent, unanticipated shocks to oil demand affect drilling, production, and prices in equilibrium. Our motivation derives from the large historical changes in oil prices shown in figure 1, which prior work has attributed primarily to demand shocks (Kilian 2009; Kilian and Hicks 2013). Our empirical analysis indicated that positive shocks are associated with increases in drilling activity and drilling costs and with backward-dated price expectations, while the reverse holds for negative demand shocks. In either case, production from existing wells is unaffected. Here, we demonstrate how our model naturally explains these phenomena.

Consider first the limiting case of an unbounded well stock, which is useful for building intuition. Suppose we are initially in a steady state, depicted by point A in the phase diagram in the right panel of figure 9. Then suppose there is a permanent, unanticipated increase in oil demand that shifts the $\dot{a}(t) = 0$ locus up and right, yielding a new steady state at point B . The optimal transition dynamics imply that we jump immediately to the new stable arm and then follow it gradually toward the new steady state, as depicted in the figure (thick black arrows). Oil prices, drilling, and rig rental rates therefore all increase immediately on impact. Oil flow then begins to increase as new wells come on line, causing prices to gradually fall until the new steady state is reached. These dynamics imply that following the positive shock, oil price expectations will be backwarddated.⁴⁶ Finally, because production must be constrained throughout this process, production from existing wells is

⁴⁶Note that following the demand shock, price expectations must be backwarddated, but the actual path of future spot prices need not be, since subsequent demand shocks may occur.

unaffected by the shock.

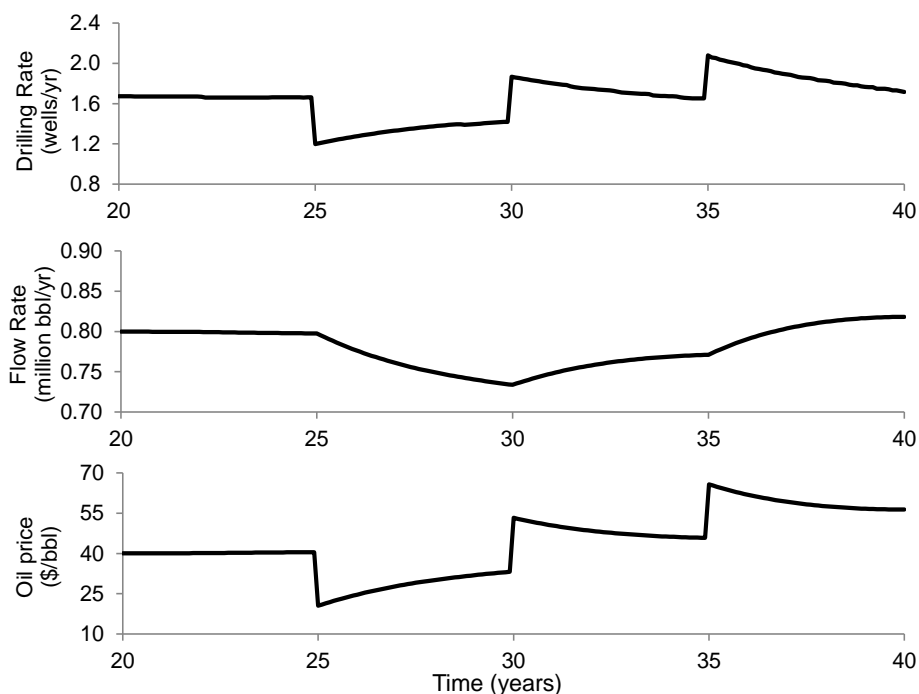
For a negative demand shock, the story is reversed. Oil prices, drilling, and rig rental rates all fall on impact and then rise along the new stable arm to the new steady state. Price expectations will be in contango following the shock. If the shock is sufficiently large, drilling will fall substantially, so that the capacity dynamics will be dominated by the exponential production decline from existing drilled wells. In this case, the expected price path can rise faster than the interest rate for a brief period following the shock (if oil demand is sufficiently inelastic).⁴⁷ This result helps explain the severe contango in oil futures markets in 1998–1999 when the Asian financial crisis led to a sharp fall in oil demand, per Kilian (2009).

Now suppose the stock of wells is finite. In this case, demand shocks will have a qualitatively similar effect on the equilibrium path, although the phase diagram does not permit a graphical analysis. To see this result, consider a capacity-constrained equilibrium path and an unexpected, permanent, vertical shift in $P(F)$ of magnitude Z . Suppose by contradiction that there is no change in the path of drilling. In this case, the entire path for $\theta(t)$ will increase by $Z/(r + \lambda)$, per equation (15). This shift implies that, under the original drilling path, $\theta(t)X - d(a(t))$ rises more slowly than the rate of interest, which is sub-optimal. Thus, the drilling rate must increase immediately on impact, just as in the no-scarcity case. If the shock is large enough, the jump in drilling will be sufficient to cause production to increase, yielding backwarddated price expectations (small shocks may not cause backwardation if production was falling prior to the shock). By the opposite mechanism, a negative demand shock will cause the rate of drilling to immediately decrease and may result in contango, with prices expected to rise at a rate potentially greater than r .

To illustrate these predicted responses, we expose the numerical model used to generate figure 8 to a series of three demand shocks: one negative shock followed by two positive

⁴⁷For an extremely large negative demand shock, production may actually fall below the capacity constraint upon impact, with the oil price and marginal cost of drilling both subsequently rising at the rate of interest. The fact that oil production from previously drilled wells does not respond to shocks during our 1990–2007 sample suggests that negative oil demand shocks during this period were not large enough to induce unconstrained production. The analysis underlying figure 5 confirms this result.

Figure 10: Impacts of unexpected demand shocks from the equilibrium model with linear demand, linear marginal drilling costs, and a finite well stock



Note: This figure uses the demand, drilling cost, and parameter specifications from figure 8. At time $t = 25$, demand is decreased from $P = 200 - 200F$ to $P = 180 - 200F$. At $t = 30$, demand is increased to $P = 200 - 200F$. Finally, at $t = 35$, demand is increased to $P = 220 - 200F$. See text for details.

shocks. See figure 10 for the results. The top panel shows that drilling “jumps” on impact for each shock. The middle panel illustrates that production does not jump following any shock but rather responds gradually to changes in the rate of drilling. The bottom panel shows that after the negative demand shock, the oil price jumps down on impact but then gradually rises, consistent with price expectations being in contango immediately following the shock. Moreover, in this example the oil price initially rises at a rate greater than r while production remains constrained, as observed in 1998–1999. Finally, following the positive demand shocks, the oil price initially jumps up but then declines as production increases, so that price expectations would be backwarddated immediately following each shock. This result relates to the increases in spot prices and the backwardation of oil futures markets

throughout the mid-2000s, as the oil market was repeatedly shocked by demand increases from emerging Asian markets, per Kilian (2009) and Kilian and Hicks (2013).

5 Our model’s relationship to the Hotelling literature

Within the Hotelling literature, our model is most closely related to models in which firms face convex investment costs to expand reserves and thereby reduce their marginal extraction costs (Pindyck 1978; Livernois and Uhler 1987; Holland 2008; Venables 2012), and to models in which production is directly constrained and firms face convex costs to expand production capacity (Gaudet 1983; Switzer and Salant 1986; Holland 2008).⁴⁸ Both types of models can generate initial periods of rising production and falling prices as reserves grow or as production capacity builds, followed by an inevitable decline (these models, like ours, also allow for periods of increasing production and backwarddated price expectations following positive demand shocks). In these models, however, the eventual decline in production arises due to an increasing Hotelling scarcity rent rather than a declining flow constraint. Thus, whenever production is declining it must be responsive to demand shocks—a prediction rejected by our data. Moreover, these models do not admit the possibility that forward-looking market participants may expect the oil price to rise faster than the rate of interest in equilibrium—something that we actually see in our futures data.

We show that when the capacity constraint declines with cumulative extraction due to the loss of underground pressure, the constraint will always bind under realistic conditions, even when production is falling. Thus, our model can explain why production from drilled wells declines steadily over time yet simultaneously does not respond to price shocks, and why futures markets sometimes forecast prices to rise faster than the rate of interest. In addition, our model links capacity expansion to drilling activity and the marginal cost of capacity expansion to the rental rates on drilling rigs, each of which can be observed empirically.

⁴⁸For a review of the theoretical Hotelling literature through the 1970s, see Devarajan and Fisher (1981). For more recent reviews, see Krautkraemer (1998), Gaudet (2007), and Gaudet and Salant (2014). For an accessible theoretical primer, see Salant (1995).

Other papers in the economics literature are, like ours, premised on the idea of an oil production constraint that decays with cumulative extraction. To the best of our knowledge, however, only two recent unpublished papers attempt to characterize equilibrium outcomes (Okullo, Reynès and Hofkes 2012; Mason and van’t Veld 2013).⁴⁹ The others all treat oil prices and drilling costs as exogenous, focusing on the valuation of existing individual reserves or on the optimal timing and development of new ones (Nystad 1987; Adelman 1990; Davis and Cairns 1998; Cairns and Davis 2001; Thompson 2001; Smith 2012; Cairns 2014).

Our paper differs in several crucial ways. First, we present empirical results to motivate every assumption in our model. Our key finding—that oil production from drilled wells is unresponsive to price shocks, consistent with a binding capacity constraint—has not, to the best of our knowledge, been documented in prior work. Second, motivated by our empirical results, we develop a new model that emphasizes the rate of drilling as the central choice variable rather than the rate of oil production. Third, we use our model to explain why the production constraint binds empirically. Fourth, we show that Hotelling’s logic still applies in our model, with the discounted marginal revenue stream from drilling minus the marginal cost of drilling (i.e., the rig rental price) rising at the interest rate. Finally, we show that the drilling, production, and price dynamics implied by our model can match those of the real-world oil extraction industry.

We also contribute to an extensive empirical literature that tests whether the canonical “Hotelling Rule” holds in practice, generally finding that it does not.⁵⁰ Rather than test the precise numerical implications of a model that we know must be unrealistic—particularly in the presence of shocks along the equilibrium path—we simply ask whether crude oil production from existing wells responds *at all* to prices in the way that a model adopting

⁴⁹Mason and van’t Veld (2013) only derive equilibrium outcomes in a simplified, two-period version of their model, while Okullo et al. (2012) focuses primarily on a graphical analysis of a stationary (and therefore incorrect) phase portrait.

⁵⁰Smith (1979), Slade (1982), and Berck and Roberts (1996) find limited evidence for an upward trend in exhaustible resource prices, but tests based on price alone are not correctly specified unless extraction costs are negligible. Structural econometric papers estimating *in-situ* values, including Miller and Upton (1985), Halvorsen and Smith (1984), Black and LaFrance (1998), and Thompson (2001), find mixed results. See Krautkraemer (1998) and Slade and Thille (2009) for recent reviews.

the cake-eating assumption would predict. Finding that it does not, we then show that our keg-tapping Hotelling model is capable of rationalizing this result as well as many other empirical regularities observed in oil markets.

6 Conclusion

Most Hotelling-style models take for granted that extractors are unconstrained in choosing the rate at which production flows to market. Our analysis of crude oil drilling and production in Texas shows this assumption to be inconsistent with the technology and cost structure of the oil industry. Oil is not extracted barrel-by-barrel. Instead, extractors drill wells, and the maximum flow from these wells is geologically constrained by underground reservoir pressure. Given this constraint, extractors always opt to produce at capacity under realistic conditions and will respond to oil prices only by varying their rate of drilling. Thus, crude oil extraction is best modeled as a dynamic drilling investment problem in which the flow of oil—though technically a control variable—is set equal to production capacity at all times and therefore behaves like a state variable in equilibrium.

We develop a new model of exhaustible resource extraction that accommodates these important features of the crude oil extraction industry. Our model—uniquely in the Hotelling literature—replicates several salient observations from oil markets: (1) production from pre-existing wells steadily declines over time and does not respond to oil price shocks; (2) drilling of new wells and drilling rig rental rates strongly co-vary with oil prices; (3) local oil-producing regions and fields exhibit production peaks; and (4) expected future oil prices can be backwardated after positive demand shocks and can rise faster than the interest rate (temporarily) following negative demand shocks. A main contribution of our paper is therefore to show that a Hotelling-style model *can* replicate qualitatively the oil price, extraction, and drilling dynamics that we observe in the real world, breathing new life into a theoretical literature that has, until now, largely failed to deliver empirically.

Our model could be extended in several logical ways. First, we currently model drilling costs with a fixed, upward-sloping supply curve. While the stock of drilling rigs and crews is fixed in the short run, with rigs allocated to their highest-valued use, this stock can change over time as new rigs are built and crews are trained, as old rigs are scrapped and workers retire, or as rigs and crews are moved from one region to another. These rig dynamics could be added to the model to effectively allow for a more elastic long-run drilling supply curve. Second, the model could incorporate uncertainty about future oil demand, so that firms rationally expect demand shocks to occur along the equilibrium path. In this case, each undrilled well would be characterized as a real option (see Kellogg (2014)). Third, since different locations typically have their own geological features, it would be natural to consider variation in drilling costs, production decline rates, and resource stocks across regions—or across individual wells within the same region. Heterogeneity in drilling costs may be particularly important to consider given our increasing reliance on deeper, more remote, and unconventional energy resources. We leave these extensions to future work.

References

- Adelman, M. A.**, “Mineral Depletion, with Special Reference to Petroleum,” *The Review of Economics and Statistics*, 1990, *72* (1), 1–10.
- Alquist, Ron and Lutz Kilian**, “What Do We Learn From the Price of Crude Oil Futures?,” *Journal of Applied Econometrics*, June/July 2010, *25* (4), 539–573.
- Benveniste, L. M. and J. A. Scheinkman**, “On the Differentiability of the Value Function in Dynamic Models of Economics,” *Econometrica*, 1979, *47* (3), 727–732.
- Berck, Peter and Michael Roberts**, “Natural Resource Prices: Will They Ever Turn Up?,” *Journal of Environmental Economics and Management*, 1996, *31*, 65–78.

- Black, Geoffrey and Jeffrey T. LaFrance**, “Is Hotelling’s Rule Relevant to Domestic Oil Production?,” *Journal of Environmental Economics and Management*, 1998, *36*, 149–169.
- Cairns, Robert**, “The Green Paradox of the Economics of Exhaustible Resources,” *Energy Policy*, February 2014, *65*, 78–85.
- Cairns, Robert D. and Graham A. Davis**, “Adelman’s Rule and the Petroleum Firm,” *The Energy Journal*, 2001, *22* (3), 31–54.
- Dahl, Carol and Mine Yucel**, “Testing Alternative Hypotheses of Oil Producer Behavior,” *The Energy Journal*, 1991, *12* (4), 117–138.
- Darcy, Henry**, *Les Fontaines Publiques de la Ville de Dijon* 1856.
- Davis, Graham A. and Robert D. Cairns**, “Simple Analytics of Valuing Producing Petroleum Reserves,” *The Energy Journal*, 1998, *19* (4), 133–142.
- Devarajan, Shantayanan and Anthony C. Fisher**, “Hotelling’s ‘Economics of Exhaustible Resources’: Fifty Years Later,” *Journal of Economic Literature*, 1981, *19* (1), 65–73.
- Elliott, Graham, Thomas J. Rothenberg, and James H. Stock**, “Efficient Tests for an Autoregressive Unit Root,” *Econometrica*, July 1996, *64* (4), 813–836.
- Gaudet, Gérard**, “Investissement Optimal et Coûts d’Ajustement dans la Théorie Économique de la Mine. (Optimal Investment and Adjustment Costs in the Economic Theory of the Mine. With English summary.)” *Canadian Journal of Economics*, February 1983, *16* (1), 39–51.
- , “Natural Resource Economics Under the Rule of Hotelling,” *Canadian Journal of Economics/Revue Canadienne d’Economie*, 2007, *40* (4), 1033–1059.

- **and Stephen Salant**, “The Hotelling Model with Multiple Demands,” in Robert Halvorsen and David Layton, eds., *Handbook on the Economics of Natural Resources*, Edward Elgar, 2014.
- Griffin, James M.**, “OPEC Behavior: A Test of Alternative Hypotheses,” *American Economic Review*, December 1985, *75* (5), 954–63.
- Güntner, Jochen H. F.**, “How Do Oil Producers Respond To Oil Demand Shocks?,” *Energy Economics*, forthcoming.
- Halvorsen, Robert and Tim R. Smith**, “On Measuring Natural Resource Scarcity,” *Journal of Political Economy*, 1984, *92* (5), 954–964.
- Hamilton, James D.**, “Understanding Crude Oil Prices,” *The Energy Journal*, 2009, *30* (2), 179–206.
- , “Oil Prices, Exhaustible Resources, and Economic Growth,” in Roger Fouquet, ed., *Handbook of Energy and Climate Change*, Cheltenham, United Kingdom: Edward Elgar Publishing, 2013.
- Hogan, William W.**, “A Dynamic Putty–Semi-putty Model of Aggregate Energy Demand,” *Energy Economics*, 1989, *11* (1), 53–69.
- Holland, Stephen P.**, “Modeling Peak Oil,” *The Energy Journal*, 2008, *29* (2), 61–79.
- Hotelling, Harold**, “The Economics of Exhaustible Resources,” *Journal of Political Economy*, 1931, *39* (2), 137–175.
- Hyne, Norman J.**, *Nontechnical Guide to Petroleum Geology, Exploration, Drilling, and Production*, 2nd edition ed., Tulsa, Oklahoma: PennWell Corporation, 2001.
- Jones, Clifton T.**, “OPEC Behaviour Under Falling Prices: Implications for Cartel Stability,” *The Energy Journal*, 1990, *11* (3), 117–130.

- Kellogg, Ryan**, “The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling,” *American Economic Review*, 2014, *104* (6), 1698–1734.
- Kilian, Lutz**, “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, 2009, *99* (3), 1053–1069.
- **and Bruce Hicks**, “Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003–2008?,” *Journal of Forecasting*, August 2013, *32*, 385–394.
- Krautkraemer, Jeffrey A.**, “Nonrenewable Resource Scarcity,” *Journal of Economic Literature*, 1998, *36* (4), 2065–2107.
- Léonard, Daniel and Ngo van Long**, *Optimal Control Theory and Static Optimization in Economics*, Cambridge University Press, November 1992.
- Livernois, John R. and Russell S. Uhler**, “Extraction Costs and the Economics of Nonrenewable Resources,” *Journal of Political Economy*, 1987, *95* (1), 195–203.
- Mason, Charles F. and Klaas van’t Veld**, “Hotelling Meets Darcy: A New Model of Oil Extraction,” June 2013. Paper presented at the 20th Annual Conference of the European Association of Environmental and Resource Economists (Toulouse, France).
- Miller, Merton H. and Charles W. Upton**, “A Test of the Hotelling Valuation Principle,” *Journal of Political Economy*, 1985, *93* (1), 1–25.
- Nystad, Arild N.**, “Rate Sensitivity and the Optimal Choice of Production Capacity of Petroleum Reservoirs,” *Energy Economics*, 1987, *9* (1), 37–45.
- Okullo, Samuel Jovan, Frédéric Reynès, and Marjan W. Hofkes**, “Modeling Peak Oil and the Geological Constraints on Oil Production,” March 2012. Technical report, Institute for Environmental Studies.
- Pindyck, Robert S.**, “The Optimal Exploration and Production of Nonrenewable Resources,” *Journal of Political Economy*, October 1978, *86* (5), 841–61.

- Ramcharran, Harri**, “Oil Production Responses to Price Changes: An Empirical Application of the Competitive Model to OPEC and Non-OPEC Countries,” *Energy economics*, 2002, *24* (2), 97–106.
- Rao, Nirupama S.**, “Taxation and the Extraction of Exhaustible Resources: Evidence from California Oil Production,” April 2010. MIT-CEEPR Working Paper 10-006.
- RigData**, “Day Rate Report,” 1990–2013. Available from <http://www.rigdata.com/news-and-analysis/day-rate-report.aspx> (accessed August 27, 2013).
- Rincón-Zapatero, Juan Pablo and Manuel S. Santos**, “Differentiability of the Value Function in Continuous-Time Economic Models,” *Journal of Mathematical Analysis and Applications*, 2012, *394*, 305–323.
- Salant, Stephen W.**, “The Economics of Natural Resource Extraction: A Primer for Development Economists,” *World Bank Research Observer*, February 1995, *10* (1), 93–111.
- Slade, Margaret E.**, “Trends in Natural-resource Commodity Prices: An Analysis in the Time Domain,” *Journal of Environmental Economics and Management*, 1982, *9*, 122–137.
- **and Henry Thille**, “Whither Hotelling: Tests of the Theory of Exhaustible Resources,” *Annual Review of Resource Economics*, August 2009, *1* (1), 239–259.
- Smith, James L.**, “Modeling the Impact of Taxes on Petroleum Exploration and Development,” 2012. IMF Working Paper No. 12/278.
- Smith, V. Kerry**, “Natural Resource Scarcity: A Statistical Analysis,” *Review of Economics and Statistics*, 1979, *61*, 423–427.
- Society of Petroleum Evaluation Engineers**, “Survey of Economic Parameters Used in Property Evaluations,” 1995.

Switzer, Sheldon and Stephen Salant, “Optimal Capacity Expansion by Foresighted Extractors and Producers,” in G. Gaudet and P. Lasserre, eds., *Problèmes de Ressources Naturelles et Théorie Économique*, University of Laval Press, Laval, Québec, 1986.

Thompson, Andrew C., “The Hotelling Principle, Backwardation of Futures Prices and the Values of Developed Petroleum Reserves—the Production Constraint Hypothesis,” *Resource and Energy Economics*, 2001, *23*, 133–156.

Venables, Anthony J., “Depletion and Development: Natural Resource Supply with Endogenous Field Opening,” September 2012. OxCarre Research Paper 62.

A Additional empirical results

This empirical appendix has three parts. First, we provide results from regressions that complement figures 2 and 4 in the main text. Second, we show that the primary features of figure 2—the deterministic production decline and the lack of response to price shocks—hold in subsamples of production from relatively high-volume leases and from wells drilled in-sample. Third, we present results that rule out alternative explanations for the lack of price response.

A.1 Regression analysis for main production and drilling results

Figure 2 indicates graphically that production from previously drilled wells does not respond to spot or future oil prices, while figure 4 indicates a strong response of drilling to these prices. Here, we demonstrate these results more formally via a regression analysis.

As a complement to figure 2, we seek to regress the log of oil production (bbl/day) from wells drilled before 1990 on both the logged front month price of oil and the expected annual rate of price increase. The oil production data are monthly, and we average the price data to the monthly level. We conduct the analysis in first differences because we cannot reject (using the GLS procedure of Elliott, Rothenberg and Stock (1996)) that both $\log(\text{production})$ and $\log(\text{front month price})$ are unit root processes.⁵¹

Our main specification is given by:

$$\begin{aligned} \Delta \log(\text{Production}_t) = & \alpha + \beta_0 \Delta \log(\text{Price}_t) + \beta_1 \Delta \log(\text{Price}_{t-1}) + \delta_0 \Delta \text{IncreaseRate}_t \\ & + \delta_1 \Delta \text{IncreaseRate}_{t-1} + \eta \cdot \text{Time} + \text{error}_t \end{aligned} \quad (21)$$

We include a lagged difference because doing so minimizes the AIC criterion. Removing this lag or adding additional lags does not qualitatively change the results. We include a time trend in our baseline specification to account for the possibility that the pressure-driven production decline curve is hyperbolic rather than exponential, so that $\Delta \log(\text{Production}_t)$ will not be constant over time even if the β and δ coefficients all equal zero (enriching the trend to a polynomial does not qualitatively affect the results). For inference, we use Newey-West with four lags (doing so only slightly increases the estimated standard errors; adding additional lags leaves the estimated errors essentially unchanged).

Column (1) of table 1 presents estimates from the specification in equation (21), while column (2) presents estimates from a specification that does not include the time trend. In both specifications, we find that oil price changes have neither an economically nor statistically significant effect on production from pre-existing wells, consistent with figure 2. In column (1), the sum of the coefficients on the current and lagged difference in $\log(\text{front month price})$ yields an insignificant elasticity of oil production with respect to front month price of -0.004 (with a standard error of 0.036). The sum of the coefficients on the current and lagged differences in the expected rate of price increase equals -0.0010, meaning that

⁵¹For $\log(\text{production})$, we obtain a test statistic of -0.644 with the optimal 12 lags, relative to a 10% critical value of -2.548. For $\log(\text{front month price})$, we obtain a test statistic of -1.337 with the optimal 11 lags, relative to a 10% critical value of -2.558. Results are similar for all other lags. These results include a time trend in the test, but results are similar when a trend is excluded.

Table 1: First-differenced regressions of Texas production and drilling on oil prices

VARIABLES	$\Delta \log(\text{Production})$		$\Delta \log(\text{Drilling})$	
	(1)	(2)	(3)	(4)
$\Delta \log(\text{Front-month price})$	0.047 (0.033)	0.058 (0.033)	0.322 (0.139)	0.315 (0.139)
Lagged $\Delta \log(\text{Front-month price})$	-0.052 (0.037)	-0.043 (0.037)	0.017 (0.159)	0.012 (0.157)
2nd lagged $\Delta \log(\text{Front-month price})$			0.264 (0.165)	0.259 (0.162)
Δ Expected rate of price increase	0.0001 (0.0004)	0.0002 (0.0004)	0.0036 (0.0023)	0.0035 (0.0023)
Lagged Δ Expected rate of price increase	-0.0010 (0.0005)	-0.0009 (0.0005)	-0.0009 (0.0024)	-0.0010 (0.0024)
2nd lagged Δ Expected rate of price increase			-0.0006 (0.0027)	-0.0007 (0.0026)
Time trend (in months)	0.0042 (0.0013)		-0.0029 (0.0073)	
Constant	-0.0070 (0.0008)	-0.0071 (0.0009)	-0.0005 (0.0045)	-0.0004 (0.0045)
N	214	214	213	213
R^2	0.096	0.074	0.107	0.081

Note: Data are monthly and the same as in figures 2 and 4. The expected rate of price increase is measured as a percentage. Standard errors in parentheses are Newey-West with 4 lags.

an increase in the expected rate of price increase of 10 percentage points (about one standard deviation) is associated with only a 1% decrease in production. A test against the null hypothesis that this sum equals zero yields a p-value of 0.102.

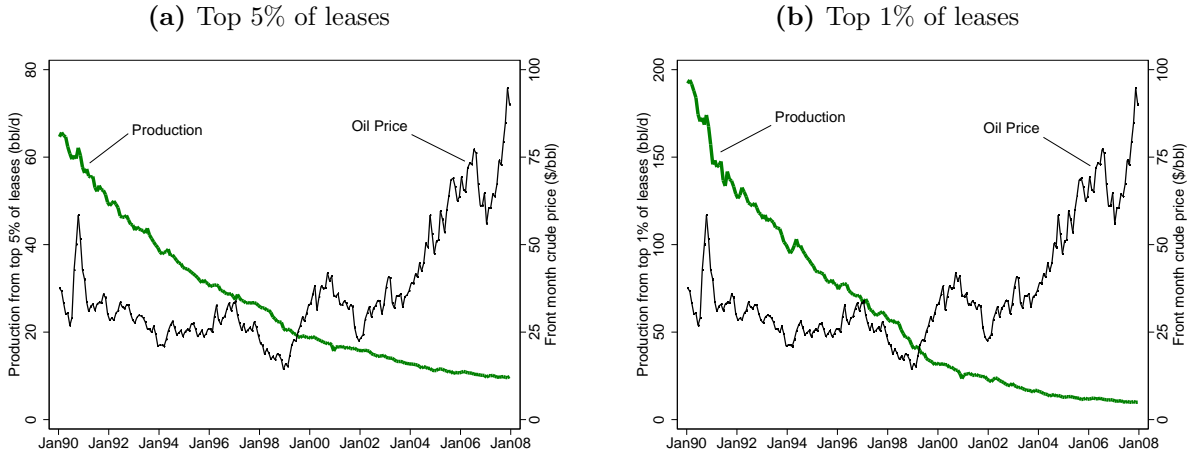
In columns (3) and (4) of table 1, we re-estimate equation (21) but use the log of wells drilled per month as the dependent variable. We also add an additional lag of the independent variables, minimizing the AIC criterion.⁵² Column (3) includes a time trend, while column (4) does not. In either case, the estimates follow what is clear from figure 4: drilling activity responds strongly to changes in oil prices. For column (3), summing the coefficients on the current and lagged front month price differences yields an elasticity of drilling with respect to the front month price of 0.604 (with a standard error of 0.191). The response of drilling to expected future changes in price, however, is estimated imprecisely (the sum of the three coefficients equals 0.002, with a standard error of 0.003).

A.2 Production decline in high-volume leases and wells drilled in-sample

Figure 2 presents average monthly production from all active oil leases in Texas for which there was no rig activity from 1990–2007. Average lease-level production in Texas is quite low, raising the question of whether our empirical results extend to higher-volume fields that

⁵²Intuitively, lags caused by engineering and permitting generate a lag in the response of drilling to price shocks.

Figure 11: Production from existing wells in high-volume Texas leases



Note: This figure presents crude oil front month prices and daily average oil production from leases on which there was no rig activity (so that all production comes from pre-existing wells). All prices are real \$2007. The left panel (a) includes leases that are in the top 5% of total production in the 1990–2007 sample, while the right panel (b) includes leases that are in the top 1%.

might be found elsewhere in the world.

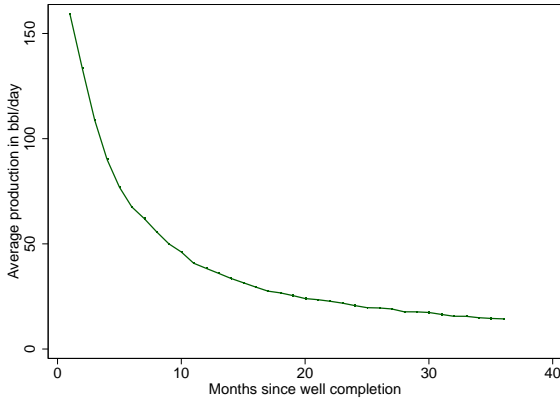
Figure 11 presents production data from subsamples of relatively high-volume leases in Texas. For each lease in the dataset, we obtain its total production by summing its production rate over the entire 1990–2007 sample. We then assign each lease to its appropriate percentile based on total production. The left panel of figure 11 includes leases within the top 5% of total production, and the right panel includes leases in the top 1%. The average top 1% lease produces nearly 200 bbl/day at the start of the sample, far larger than the overall average initial production of about 8 bbl/day shown in figure 2. Nonetheless, in both panels of figure 11, production declines deterministically and exhibits essentially no response to price signals. The production declines are steeper in figure 11 than in figure 2, suggesting either that high-volume leases have relatively high decline rates or that leases typically have decline curves that are hyperbolic rather than exponential.

We next examine production from wells drilled during the sample. To undertake this analysis, we first match drilling records, which come from TRRC drilling permits, to the TRRC lease-level production data. This match must be done based on lease names. Because naming conventions vary across the two datasets, we are able to match only 67.2% of drilled oil wells to a lease in the production data.

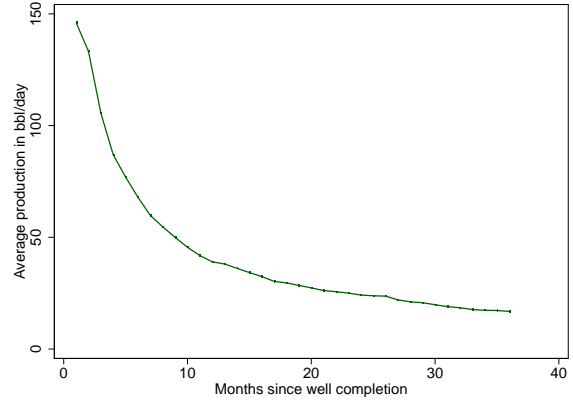
Because production can only be observed at the lease-level rather than at the well-level, we isolate the sample to drilled wells that are the only producing well on their lease for three years after the well was completed. The remaining sample consists of 4,105 wells drilled between 1990 and 2004 (inclusive). We break these 15 years into five 3-year periods. Each panel of figure 12 then plots, for wells drilled within a particular period, the average production for the first three years of the wells' lives. For each period of drilling, production

Figure 12: Production from wells drilled during the sample period

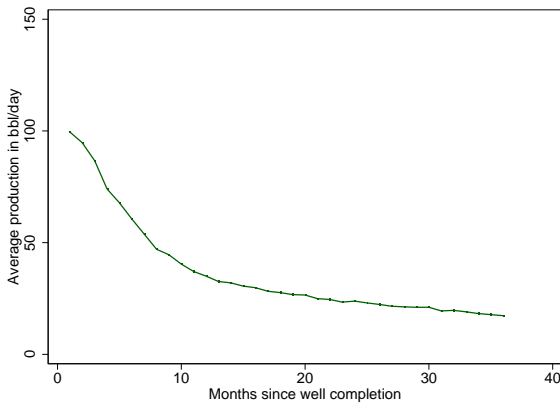
(a) Wells drilled from 1990–1992



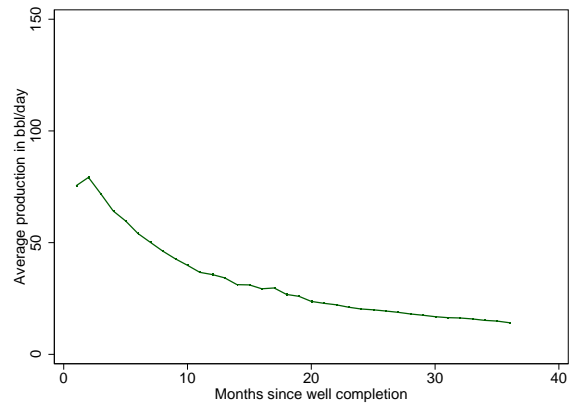
(b) Wells drilled from 1993–1995



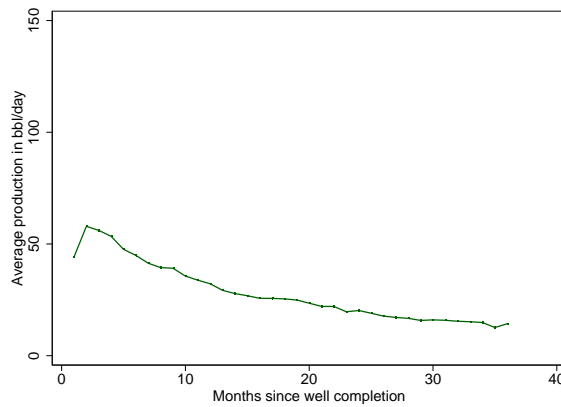
(c) Wells drilled from 1996–1998



(d) Wells drilled from 1999–2001



(e) Wells drilled from 2002–2004



Note: Each panel plots average monthly production for wells drilled during the indicated time interval. See text for details.

from drilled wells declines deterministically over time and again does not indicate any price responsiveness. Production from wells drilled in later periods is lower, on average, than production from wells drilled in early periods, suggesting that firms sensibly drill highly productive wells before drilling less productive wells. Incorporating well-level heterogeneity into the theoretical model is therefore likely to be a fruitful path for future research.

A.3 Ruling out alternative explanations for production’s lack of response to price incentives

There are several potential alternative explanations for the lack of price response among existing oil wells that we need to rule out. First, races to oil created by common-pool externalities in oil fields with multiple lease-holders would diminish the incentive that individual lease-holders have to defer production when prices are expected to rise, since much of the deferred production would be extracted by others instead. Figure 13(a) shows, however, that the long, downward trend in oil production manifests both for oil fields with multiple operators and for oil fields with just a single operator.

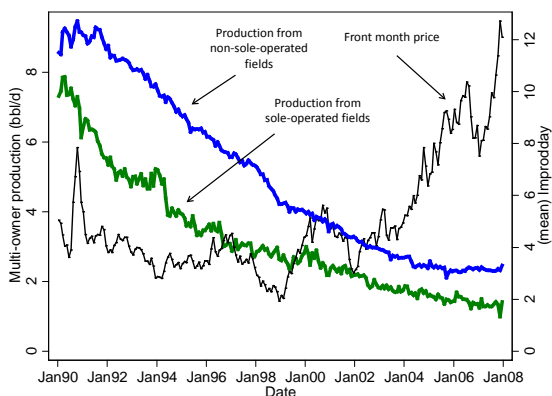
Second, a condition of many leases is that the lease holder produce oil; a firm that drops lease-level production to zero may therefore risk losing the lease. Figure 13(b) shows, however, that production from multi-well leases, on which producers may shut in at least some of their wells without risk of losing the lease, does not respond to price signals.

Third, oil production in Texas is subject to maximum allowable production quotas—or “allowables”—as determined by the Texas Railroad Commission. This system dates to the East Texas Oil Boom of the 1930s when a large share of world oil production came from Texas, and races to oil led to overproduction and collapsing world oil prices. Whether originally intended to end the race to oil, or simply to cartelize the Texas oil industry and boost prices, this system persists to this day, and every lease in our data has a monthly allowable, including on fields with just a single operator. One obvious concern is that these maximum production quotas are binding, leading to the lack of price response. Figure 13(c) shows, however, that the average production for leases in our main sample is well below the average allowable production. Thus, the allowables are not binding and therefore cannot explain the lack of price response that we observe in our data.

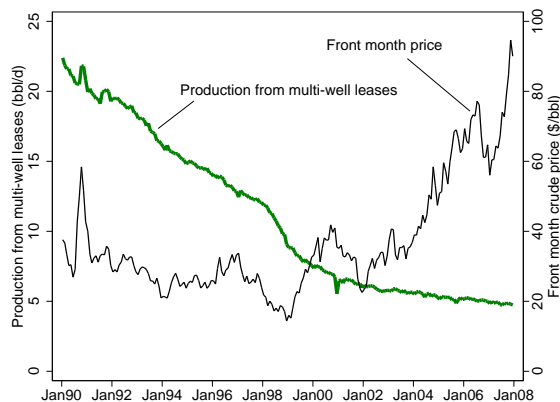
Fourth and finally, one possible concern is that the decision makers whose behavior we observe in our data could earn profits by delaying production during periods of extreme contango but do not do so because they dynamically optimize incorrectly or because their future price expectations are not aligned with the futures market. Figure 13(d) shows, however, that above-ground storage of crude oil on these leases increased notably during the 1998–1999 period of extreme contango. Lease-holders responded to this price incentive by accumulating inventories above-ground, deferring sales—not extraction—to take advantage of the expected increase in prices, and confirming that they appropriately respond to incentives generated by price expectations aligned with the futures market.

Figure 13: Graphical evidence ruling out alternative explanations

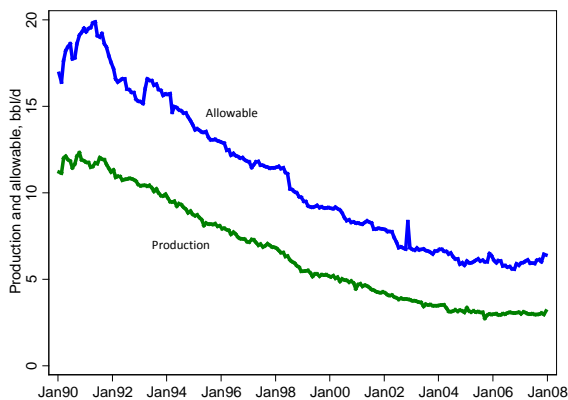
(a) Single vs. multiple-operator fields



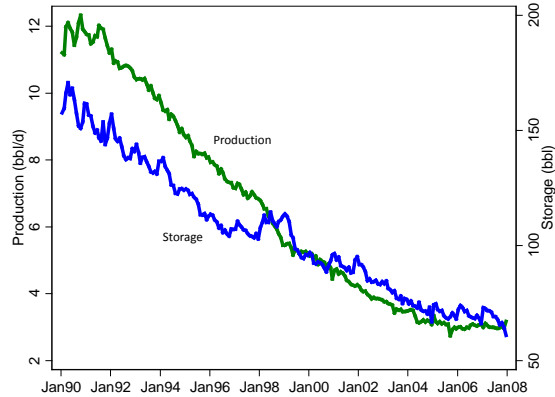
(b) Production on multi-well leases



(c) Actual versus allowable production



(d) Production and above-ground storage



Note: Panel (a) shows average daily production on fields with multiple operators as well as on fields with just a single operator, along with the spot price of crude oil. Panel (b) shows average production for leases with multiple wells in our main sample. Panel (c) shows average actual production as well as average allowable production for leases in our main sample. Panel (d) shows average oil production and average above-ground storage of crude oil on leases in our main sample. See text for details.

B Calculations showing that production below the constraint was never warranted

In this appendix we demonstrate that production below the constraint was never warranted for our sample of Texas oil producers in any of the 216 months between 1990–2007: withholding production in any given month t and selling it optimally over time was never anticipated to be as valuable as selling the production in month t .

To calculate the value of deferred production reported in figure 5, we assume an exponential production decline rate of 10% annually and a discount rate of 10% annually (see footnote 26). The corresponding monthly decline rate and monthly real rate of interest are then given, respectively, by $\lambda = 1 - (1 - 0.10)^{1/12}$ and $r = (1 + 0.10)^{1/12} - 1$. We assume that expected real prices at each date follow our futures data, which we describe in detail in the text. The longest futures contract is typically 60 months. We assume that prices more than 60 months in the future were expected at time t to plateau at the level of the 60-month futures price.⁵³ This assumption naturally follows from the long-run plateauing of the futures curves in the data, as is apparent in figure 1.

Given these assumptions, it is straightforward to calculate recursively the value of deferred production at each date t . Since prices were anticipated to plateau after 60 months, it would be optimal to produce at the constraint from that date onward, and so an additional unit of production capacity 60 months hence was anticipated at time t to be worth at $t + 60$:

$$\theta(t, 60) = \frac{P(t, 60)}{1 - \delta},$$

where $\delta = (1 - \lambda)/(1 + r)$ and $P(t, 60)$ is the price anticipated at t to prevail after 60 months. This is simply the value of inheriting one extra barrel of monthly production capacity in 60 months and producing at capacity forever while earning $P(t, 60 + j) = P(t, 60)$ per barrel for $j > 1$, all discounted back to time $t + 60$. Now suppose inductively that the shadow value on capacity $s + 1$ months in the future is anticipated at time t to be $\theta(t, s + 1)$. Then, the anticipated shadow value on capacity s months in the future is given by:

$$\theta(t, s) = \max \{ P(t, s) + (1 - \lambda)\theta(t, s + 1)/(1 + r), \theta(t, s + 1)/(1 + r) \},$$

where the well owner anticipates in month t choosing s months in the future the more lucrative of the two options: (1) producing at capacity and earning $P(t, s)$ that month and having fraction $1 - \lambda$ of initial capacity remaining in the following month, which is then valued at $\theta(t, s + 1)$; or (2) deferring production until the following month so that the full inherited capacity is saved until the following month, which is again valued at $\theta(t, s + 1)$. Thus, by backward recursion starting 60 months in the future, we can reconstruct the full sequence of shadow values, all the way back to $\theta(t, 1)$. Note then that it will be optimal to produce at the constraint at time t whenever $P(t, 0) + (1 - \lambda)\theta(t, 1)/(1 + r) > \theta(t, 1)/(1 + r)$, or equivalently $P(t, 0) - \lambda\theta(t, 1)/(1 + r) > 0$. This condition is the discrete-time analog of necessary condition (7) in the text. So the benefit from deferring one barrel of production

⁵³There are also periodic gaps in the futures data, particularly at longer time horizons. We linearly interpolate prices to fill these gaps.

at time t is $\lambda\theta(t, 1)/(1 + r)$ and the cost of deferring it is $P(t, 0)$ in month t .

Figure 5 in the text plots the marginal value of deferred production in each month based on these calculations, along with the spot price of oil. As shown in the figure, production below the constraint was never warranted, for the spot price of oil ($P(t, 0)$) exceeded the value of deferred production ($\lambda\theta(t, 1)/(1 + r)$) for every month t in our sample period—although the value of deferred production came within \$4 of the spot price during the 1998–1999 episode.

To test the sensitivity of our conclusions to parameter assumptions, we calculated the threshold values of r and λ that would lead the value of deferred production to just equal the spot price during the 1998–1999 episode. We find that if either the discount rate were as low as 4% annually or the production decline rate were as high as 21% annually, then it would have been optimal to defer production briefly during the 1998–1999 episode. Similarly, we calculated a threshold rate of increase in anticipated oil prices beyond the 60-month limit of our futures data. We find that if anticipated prices were to rise at a rate of 6% annually from the 60-month futures price, rather than plateau at that level, then it would again have been optimal to defer production briefly during the 1998–1999 episode. While this may seem like a small threshold rate of price increase, it is in fact equivalent to having the anticipated future price beyond the limits of our futures data be \$20 higher than what we observe for the 60-month futures price.⁵⁴

Finally, note that this analysis does not include any fixed operating costs. As discussed in the main text, fixed costs can generate an incentive to shut down the production of low flow rate wells when oil prices are low, as in 1998–1999.

⁵⁴Let X be the gap between the spot price and value of deferred production during the 1998–1999 episode. We are looking for the Z that solves $\lambda\delta^{60}Z/(1 - \delta) = X$, where again $\delta = (1 - \lambda)/(1 + r)$. On the right, X is simply the gap between spot price and deferral value under this calculation. On the left, we have the increase in the deferral value when anticipated prices beyond 60 months plateau Z dollars higher. Given our assumptions for the values of r and λ , the left side is $0.193Z$. Thus, the threshold Z is approximately $5X$. Since the gap between the spot price and deferral value was roughly $X = \$4$ at its narrowest, we conclude that the anticipated future price beyond 60 months would needed to have been \$20 higher.

C Standard Hotelling result

In this section, we briefly restate the standard Hotelling result before illustrating the conditions under which the path of Hotelling's planner and that of our planner coincide.

Assume that oil flow $F(t)$ at time t generates instantaneous utility flow of $U(F(t))$, with $U(0) = 0$, $U'(\cdot) > 0$, and $U''(\cdot) < 0$. Assume that an initial stock of Q_0 units of oil can be extracted at rate $F(t) \geq 0$, which is under the complete control of the oil extractor, and that the cost of this extraction is $cF(t)$, where $c \in [0, U'(0))$ is the constant marginal cost of extraction. Assume the social planner discounts utility and costs continuously at exogenous rate r and, if wealth maximizing agents are involved, they also discount profit flows at rate r . To maximize the discounted utility, the planner chooses $F(t)$ so that marginal utility less marginal cost of extraction grows exponentially at the rate of interest whenever oil flow is strictly positive:

$$F(t) \geq 0, U'(F(t)) - c - \gamma_0 e^{rt} \leq 0, \text{ with complementary slackness (c.s.).} \quad (22)$$

where γ_0 is a multiplier. Thus, quantity flows according to: $F(t) = U'^{-1}(\gamma_0 e^{rt} + c)$, where U'^{-1} is the inverse of the first derivative of $U(\cdot)$. In addition, the resource stock must be completely extracted either in finite time or asymptotically: $\int_0^\infty F(t) dt = Q_0$, which uniquely determines γ_0 and therefore the time path of extraction and marginal utilities.

If we assume that marginal utility is unbounded at zero ($U'(0) = \infty$), then marginal utility must rise forever and the resource stock will only be exhausted in the limit. If we instead assume that marginal utility is bounded at zero ($c < U'(0) < \infty$), then the resource stock will be exhausted in finite time at the precise instant that the rising marginal utility path reaches its upper bound. This is not only the planner's optimal extraction path, but it is also the aggregate extraction path that emerges in the competitive equilibrium of a decentralized market.

The main text discusses two reasons why a planner with the extraction technology described in section 3 would be unlikely to generate Hotelling's extraction path. First, oil flow in our model is constrained such that, even if the planner drilled every well immediately and produced at the maximum possible rate, oil would flow forever. Thus, price cannot rise at the rate of interest whenever oil is flowing, as Hotelling's path requires, unless marginal utility is also able to rise forever. Second, in our model, the incentive to drill in a given period depends, in part, on the cost of drilling in *other* periods, as captured by the $d'(a(t))\dot{a}(t)$ term in equation (17). One implication is that oil flow in our model can *increase* over intervals during which the marginal cost of drilling is falling, whereas the cost of extraction typically increases with production in standard models.

To overcome the first of these threats, we must assume that marginal utility is unbounded: $U'(0) = \infty$. To overcome the second, we must assume that the marginal cost of drilling is constant: $d(a) = \bar{d}$ for all $a \geq 0$. In this case, the imputed per-barrel cost of drilling is well-defined and is given by $c = \bar{d}(r + \lambda)/X$, which we assume is equivalent to the per-barrel extraction cost faced by Hotelling's planner. Given these two assumptions, condition (17) implies that the marginal utility of oil flow minus the per-barrel marginal cost of extraction rises at the rate of interest:

$$U'(F(t)) - c = \frac{\lambda\gamma_0}{X} e^{rt}. \quad (23)$$

Note, however, that we have implicitly assumed that the planner starts drilling wells at the outset of the planning period and never stops, producing at the constraint throughout, so that condition (17) always applies. For this to be the case, however, we need two other conditions to hold. First, the flow from drilled wells must decay sufficiently fast, so that the planner is able to achieve the Hotelling path via judicious control of the drilling rate while producing at her constraint. Second, the initial flow capacity K_0 cannot be too high, for otherwise the planner would delay drilling the first well—and may even produce below her constraint initially. If either of these conditions fails, then the planner is geologically constrained to an inferior path.

To illustrate the first of these two conditions, we assume for simplicity that drilling costs are zero ($c = 0$) and that utility from oil flow takes the constant elasticity form: $U(F) = \alpha F^\beta$ for $\alpha > 0, \beta \in (0, 1)$.⁵⁵ We also assume provisionally that $\lambda(1 - \beta) > r$. The total resource stock is given by $Q_0 = (K_0 + R_0X)/\lambda$. In this case, the optimal program is given by:

$$\theta(t) = \frac{\gamma_0 e^{rt}}{X}, \text{ where } \gamma_0 = \frac{\alpha\beta X}{\lambda} \left(\frac{rQ_0}{1-\beta} \right)^{\beta-1} \quad (24)$$

$$F(t) = \frac{rQ_0}{1-\beta} e^{-\frac{r}{1-\beta}t} \quad (25)$$

$$a(t) = \frac{\left(\lambda - \frac{r}{1-\beta} \right)}{X} F(t) \quad (26)$$

$$\gamma(t) = \gamma_0 e^{rt}. \quad (27)$$

Note that the planning period begins with a pulse of drilling such that flow capacity totaling $rQ_0/(1 - \beta) - K_0$ is immediately added to the inherited capacity. Equation (26) implies that $a(t) > 0$ for all $t \geq 0$, since $F(t) > 0$ and since we have provisionally assumed that $\lambda(1 - \beta) - r > 0$.⁵⁶ It is straightforward to verify that this program satisfies each of the necessary conditions (7)–(13) and is therefore optimal for our planner.⁵⁷ Since $a(t) > 0$, it achieves the same discounted utility as Hotelling's planner.

However, our planner cannot always accomplish this feat. Suppose that the rate of decay from drilled wells is too low, with $\lambda(1 - \beta) < r$. To achieve the result in (17) that $U'(F(t)) = \frac{\lambda\gamma_0}{X} e^{rt}$, the planner must set $\dot{F}(t) = \frac{rU'(F(t))}{U''(F(t))} = -\frac{rF(t)}{1-\beta}$. But then equation (12) implies that $a(t)X = \lambda F(t) - \frac{rF(t)}{1-\beta}$. Substituting and simplifying we conclude that $a(t)X = \frac{F(t)(\lambda(1-\beta)-r)}{1-\beta}$, which violates nonnegativity of $a(t)$. Intuitively, the planner is geologically constrained to a price path that rises more slowly than the rate of interest. Even if she drills all of the wells immediately and produces at the maximum possible rate, oil cannot be extracted quickly enough to satisfy Hotelling's rule.

So suppose instead that $\lambda(1 - \beta) > r$ and that the inherited flow capacity is too high,

⁵⁵Note that η , the inverse oil demand elasticity (which we treat as a positive number), is given by $1 - \beta$.

⁵⁶Since we have assumed that drilling and extraction are both costless in this example, any alternative drilling path such that the production path in (25) is feasible is also optimal. For any positive drilling cost, however, the planner would produce at the constraint and defer drilling until necessary.

⁵⁷Equations (24) and (27) imply that condition (9) holds with equality. Equations (24) and (25) imply that condition (7) holds. Equation (27) ensures that (11) holds. Finally, equations (25) and (26) ensure that (12) holds.

with $K_0 > rQ_0/(1 - \beta)$. In this case, the planner will choose not to drill any wells initially, and production will decline at rate λ until drilling commences (when $K(t) = rQ(t)/(1 - \beta)$). During this time, price will rise at a rate greater than r , and our planner will therefore not achieve the same utility as Hotelling's planner.⁵⁸

To summarize, when the extraction technology involves drilling wells rather than producing barrels, we should not expect the Hotelling path to be optimal, unless four conditions hold: (1) the marginal cost of drilling a well is constant, (2) marginal utility is unbounded at zero, (3) the decay in flow from drilled wells is sufficiently fast, and (4) the inherited rate of oil flow is not too high. While the latter two conditions seem reasonable (given that new wells are constantly being drilled in the real world) the former two conditions are not tenable. Our analysis of the Texas data shows clearly that marginal costs rise with the rate of drilling, while the viability of alternative fuels at *current* oil prices argues against an unbounded oil price.

⁵⁸Since the planner can produce below the production constraint, the standard Hotelling path can still be achieved. However, this possibility requires either that drilling is costless or that it is so expensive that no wells are ever drilled. Otherwise, it can be shown that, given $d'(a) = 0$, price would initially rise at the rate of interest (while production is below the constraint and drilling is zero), would then rise faster than the rate of interest (after the constraint starts to bind while drilling remains at zero), and finally would rise more slowly than the rate of interest (after drilling turns positive with the constraint continuing to bind).

D Proofs for rules (1)–(6) pertaining to the general case

This appendix provides the formal statement and proof of theorem 1 discussed in the main text:

Theorem 1. *Assume that the marginal cost of drilling $d(a)$ is continuously differentiable, with $d'(a) > 0$ and $d(0) > 0$. Assume that the inverse demand for oil is given by $U'(F) = P(F) \geq 0$, which has the following properties: $P(0) > 0$; $P(F)$ is continuous for $F > 0$ and continuous at $F = 0$ if $P(0) < \infty$; $P(F) \rightarrow 0$ as $F \rightarrow \infty$; and $P(F)$ is continuously differentiable, with $P'(F) < 0$, whenever $P(F) > 0$. Assume that the inverse demand curve and drilling supply curve satisfy the following: $XP(0)/(\lambda + r) > d(0)$. Lastly, assume $\infty > R_0 > 0$ and $\infty > K_0 \geq 0$. Then the following rules hold: (1) $\gamma(t)$, $\dot{\gamma}(t)$, $\theta(t)$, $a(t)$, $R(t)$, $\dot{R}(t)$, $K(t)$, $F(t)$, $P(t)$, $\phi(t)$, $\dot{K}(t)$, and $\dot{\theta}(t)$ are continuous for all $t > 0$, as is $\dot{a}(t)$ when $a(t) > 0$; (2) capacity must go to zero in the limit as $t \rightarrow \infty$; (3) the drilling rate and capacity cannot both be weakly increasing (in the phase diagram, neither the interior of region I nor its boundaries defined by the $\dot{a} = 0$ and $\dot{K} = 0$ loci may be entered), nor can the drilling and extraction rates both be weakly increasing; (4) all available wells will be drilled in the limit as $t \rightarrow \infty$; (5) at any time that drilling ceases, production must be constrained for a measurable period afterward; and (6) given demand $P(F)$, marginal drilling costs $d(a)$, and initial reserves R_0 , there exists $\delta > 0$ such that if initial capacity $K_0 < \delta$, then drilling will begin instantly and will initially be decreasing over time.*

The proof proceeds via the following series of lemmas.

Lemma 1. *If $K_0 = 0$ then drilling starts immediately.*

Proof. Suppose not. If it is optimal to drill zero wells and produce nothing initially, then the state of the system (R, K) remains unchanged. Since the problem is stationary, it is therefore never optimal to drill. This program is dominated, however, under our assumptions that $XP(0)/(r + \lambda) > d(0)$ and that $P(F)$ and $d(a)$ are continuous. There exists an $\epsilon > 0$ and a $\delta > 0$ such that drilling at rate ϵ for a time interval of length δ and then producing at capacity generates strictly positive net utility. \square

Lemma 2. *$K(t) > 0$ for all $t > 0$.*

Proof. Suppose $K(t) > 0$ for some t . Equation (12) specifies that, if $a(s) = 0$ for all $s > t$, then $K(s)$ declines at a rate that is no greater than that for exponential decay (since $F(s) \leq K(s)$). Such a decline curve will never reach zero in finite time. Moreover, any drilling activity will only further increase $K(s)$. Thus, once $K(t) > 0$ for some t , then $K(s) > 0$ for all $s > t$.

All that remains is to show that K becomes positive immediately. If $K_0 > 0$, the proof is complete. For the $K_0 = 0$ case, note from above that this case results in drilling taking place immediately, such that $a(0) > 0$. Further note with $K_0 = 0$ that $\dot{K}(0) = a(0)X > 0$. Thus, whether $a(0) = \infty$ (a pulse of drilling) or $a(0) < \infty$ and $a > 0$ for a measurable period starting at $t = 0$, it must be that $K(t) > 0$ for all $t > 0$. \square

Lemma 3. $\theta(t) > 0$ for all t .

Proof. Suppose by contradiction that $\theta(t) = 0$ at some t . By equation (13), and since we must have $\theta(t) \geq 0$ for all t , it must then be that $\phi(\tau) = \theta(\tau) = 0$ for all $\tau \geq t$. Then, for equation (7) to hold, we must have $P(F(\tau)) = 0$ for all $\tau \geq t$. However, this cannot happen: since the amount of available production capacity is finite, we must have $P(F) > 0$ at some future date. Once this happens, condition (7) will be violated.

Intuitively, capacity is always valuable because price is strictly positive for a sufficiently low consumption quantity and because the marginal cost of extraction is zero. \square

Lemma 4. It is never optimal to set F so high that $P(F) = 0$.

Proof. Proceed by contradiction. Whenever $P(F(t)) = 0$, we have $F(t) > 0$ such that equation (7) must hold with equality. This implies $\theta(t) = 0$, violating lemma 3 above. \square

Lemma 5. $\gamma(t)$, $\dot{\gamma}(t)$, $\theta(t)$, $a(t)$, $R(t)$, $\dot{R}(t)$, $K(t)$, $\dot{K}(t)$, $F(t)$, and $P(t)$ are continuous whenever $K > 0$.

Proof. Condition (11) immediately implies that both $\gamma(t)$ and $\dot{\gamma}(t)$ are continuous.

Note $\phi(t) \geq 0$ always, $\phi(t) = 0$ if $F(t) = 0$, and $\phi(t) < \infty$ if $F(t) > 0$ (since in this case equation (7) must hold with equality). Equation (13) then implies that $\dot{\theta}(t)$ exists for all t and therefore that $\theta(t)$ is continuous.

Given equation (9), the continuity of $\theta(t)$ and the assumption that $d'(a) > 0$ is continuous together imply that $a(t)$ is continuous.

Given equation (10), the continuity of $a(t)$ then immediately implies the continuity of $R(t)$ and $\dot{R}(t)$.

Given equation (12), the continuity of $a(t)$ and the boundedness of $F(t)$ together imply that $K(t)$ is continuous.

Now suppose by contradiction that $F(t)$ is not continuous. First consider what happens if $K(t) > F(t) > 0$ both before and after the discontinuity (“jump”), such that (7) holds with equality and $\phi(t) = 0$ both before and after the jump. This violates equation (7) because $P(F)$ is strictly monotonically decreasing by assumption and because $\theta(t)$ is continuous. Now consider what happens if $K(t) > F(t) > 0$ before the jump and $F = 0$ after the jump. This again violates equation (7) because the jump down in F causes P to jump up, while $\theta(t)$ is continuous and $\phi(t)$ is fixed at zero. Now consider what happens if $F = 0$ before the jump and $K(t) > F(t) > 0$ after the jump. This again violates equation (7) because the jump up in F causes P to jump down, while $\theta(t)$ is continuous and $\phi(t)$ is fixed at zero.

Now consider what happens when $F(t) = K(t)$ before the jump and $F(t) < K(t)$ after the jump. This again results in a contradiction. For equation (7) to hold (perhaps with strict inequality if $F = 0$ after the jump), the jump down in $F(t)$ must be matched with a jump up in $\phi(t)$, but $\phi(t) = 0$ after the jump. Finally, consider what happens if $F(t) < K(t)$ before the jump and $F(t) = K(t)$ after the jump. For equation (7) to hold with equality after the jump, the jump up in $F(t)$ must be matched with a jump down in $\phi(t)$, but $\phi(t) = 0$ before the jump. Since all possible discontinuities in $F(t)$ result in contradictions, $F(t)$ must be continuous.

The continuity of $F(t)$ immediately implies the continuity of $P(t)$, since $P(F)$ is continuous, and the continuity of $\dot{K}(t)$ (via equation (12)). \square

Lemma 6. $F(t) = 0$ is never optimal whenever $K(t) > 0$.

Proof. Suppose by contradiction that $F(t) = 0$. If we assume $P(0) = \infty$, then equation (7) is violated, leading to an immediate contradiction. So assume $P(0) < \infty$ instead. Since $F(t) = 0 < K(t)$, we must have $\phi = 0$ (equation 8), and therefore by equation (13) that $\dot{\theta}(t)/\theta(t) = r$. Note that we cannot then have that $F = 0$ forever, since in this case the transversality condition (equation 14) would be violated. So it must be that production eventually becomes strictly positive. Denote the time at which this happens by τ , and note that we must have equation (7) hold with equality at τ . Because $F(t)$ is continuous, we must have $F < K$ and therefore $\phi = 0$ immediately after τ , while P is decreasing and θ is increasing (since $\dot{\theta}(t)$ is continuous). In this situation, equation (7) cannot hold with equality immediately after τ , and we have a contradiction. \square

Lemma 7. $\phi(t)$ and $\dot{\theta}(t)$ are continuous whenever $K > 0$, as is $\dot{a}(t)$ when $a(t) > 0$.

Proof. The continuity of $F(t)$ and $\theta(t)$, along with the fact that equation (7) must always hold with equality (since $F(t) > 0$ per lemma 6) implies that $\phi(t)$ must also be continuous.

The results above, combined with equations (12) and (13) imply that $\dot{\theta}(t)$ is continuous.

Finally, the results above combined with equation (9) and the assumption that $d(a)$ is continuously differentiable imply that $\dot{a}(t)$ is continuous when $a(t) > 0$. \square

Recall that lemmas 5 and 7 establish the continuity of $\gamma(t)$, $\dot{\gamma}(t)$, $\theta(t)$, $R(t)$, $\dot{R}(t)$, $a(t)$, $K(t)$, $F(t)$, $P(t)$, $\phi(t)$, $\dot{K}(t)$, and $\dot{\theta}(t)$ whenever $K(t) > 0$, while lemma 2 establishes that $K(t) > 0$ for all $t > 0$. Thus, the seven preceding lemmas complete the proof of rule (1) stated in the theorem.

Lemma 8. Capacity $K(t)$ cannot be bounded away from zero as $t \rightarrow \infty$.

Proof. Suppose that, by contradiction, there is some $\hat{K} > 0$ and $\tau > 0$ such that $K(t) \geq \hat{K}$ for all $t > \tau$. Note that production cannot be bounded away from zero in this case, since this would cause $K(t)$ to eventually fall below \hat{K} (since the finite reserves would eventually be exhausted). Thus, production must eventually fall below \hat{K} , and once it does it will be unconstrained. Unconstrained production requires $\dot{F}(t) < 0$, so production will be unconstrained forever. This causes $\theta(t)$ to rise at the interest rate forever, violating the transversality condition (14) and generating a contradiction. \square

Lemma 9. Capacity $K(t)$ must go to zero in the limit as $t \rightarrow \infty$ (that is, for all $\epsilon > 0$, there exists $\tau > 0$ such that if $t > \tau$ then $K(t) < \epsilon$).

Proof. First note that if drilling permanently stops at some point in time, lemma 9 follows immediately from lemma 8, since in this case K can never increase. So suppose we are in a case in which drilling never stops permanently. Proceeding by contradiction, suppose that there exists $\hat{K} > 0$ such that for all $\tau > 0$, there exists $t > \tau$ with $K(t) \geq \hat{K}$.

To see that this leads to a contradiction, pick an $\epsilon \in (0, \hat{K})$. We know from lemma 8 that at some time, which we will denote as t_1 , we must have $K(t_1) < \epsilon$. Let t_2 denote the time at which capacity has risen to reach \hat{K} , so $K(t_2) = \hat{K}$. To achieve this result, wells must have been drilled with reserves equal to at least $(\hat{K} - \epsilon)/\lambda$. Now note that after t_2 , by lemma 8 there must be another future time, t_3 , for which $K(t_3) < \epsilon$, followed by some t_4 at which

$K(t_4) = \hat{K}$. Between t_3 and t_4 , we must also have drilled reserves of at least $(\hat{K} - \epsilon)/\lambda$. This process must repeat an infinite number of times, which yields a contradiction because the available reserves are finite.

Thus, it must be that $\lim_{t \rightarrow \infty} K(t) = 0$ (i.e., it is sub-optimal to leave valuable capacity unused). Lemmas 8 and 9 complete the proof of rule (2) stated in the theorem. \square

Lemma 10. *The drilling rate and capacity cannot both be strictly increasing ($\dot{a} > 0$ and $\dot{K} > 0$ is inadmissible).*

Proof. Suppose not. That is, suppose $\dot{a}(t) > 0$ and $\dot{K}(t) > 0$ for some time t . We will show that, no matter how drilling and production evolve from this point, a contradiction will result.

First, suppose production is constrained at t . In this case, equation (17) holds, and we have:

$$\dot{a}(t) = \frac{X}{d'(a(t))} \left[\frac{(r + \lambda)d(a(t))}{X} - P(F(t)) + \frac{\lambda\gamma_0}{X}e^{rt} \right], \quad (28)$$

whose sign depends on the term in brackets, which is increasing in a and F . Thus, if $\dot{a}(t) > 0$, $\dot{K}(t) > 0$, and production is constrained, we will continue to have $\dot{a} > 0$ for subsequent times so long as production is constrained. If production is constrained forever, then the rate of drilling will be shocked to zero at the time when reserves are exhausted, which contradicts the continuity of $a(t)$ (lemma 5).⁵⁹ Thus we cannot have $\dot{a}(t) > 0$ and $\dot{K}(t) > 0$ simultaneously with production at the constraint forever.

So suppose instead that production is unconstrained at t (similar logic applies if we allow production to fall below the constraint at some later time). In this case, we must have $\dot{P}/P = r$ and $\dot{\theta}/\theta = r$, implying via equation (9) that $\dot{d}(a)/d(a) = r$, in turn implying that $\dot{a} > 0$. Moreover, so long as production is unconstrained, we will continue to have $\dot{a} > 0$, and if production is unconstrained forever, we will again have a discontinuity in a once reserves are exhausted.⁶⁰ On the other hand, returning production to the constraint must result in a discontinuity in F , which contradicts lemma 5. Why? Unconstrained production implies that $\dot{F} < 0$. But note that since we had $\dot{K}(t) > 0$, we must continue in this case to have $\dot{K} > 0$ after time t . Thus, the only way to return production to the constraint is to induce a discontinuity in F .

Thus, if at any time $\dot{a}(t) > 0$ and $\dot{K}(t) > 0$, there must eventually be a discontinuity in either a or F , either of which is a contradiction. \square

Lemma 11. *The drilling rate and capacity cannot both be weakly increasing ($\dot{a} \geq 0$ and $\dot{K} \geq 0$ is inadmissible).*

Proof. Suppose not. Lemma 10 above handled the case in which both $\dot{a}(t) > 0$ and $\dot{K}(t) > 0$, so here we only need to consider cases involving equalities. If $\dot{K}(t) = 0$ and $\dot{a}(t) > 0$, then

⁵⁹Note that even with infinite reserves, this path will still yield a contradiction. Why? On this path, P is monotonically decreasing, which means that we must have $\theta(t) < P(t)/(r + \lambda)$. Equation 9 tells us that $\theta(t) = d(a(t))/X$, so that $(r + \lambda)d(a(t))/X = (r + \lambda)\theta(t) < P(t)$. However, per equation (28), this result contradicts the premise that $\dot{a}(t) > 0$.

⁶⁰Again note that even with infinite reserves, this path will still yield a contradiction. Why? K is forever increasing on this path, and θ is rising forever at the rate of interest. These contradict the transversality condition (14).

equation (12) implies that $K(t)$ must strictly increase immediately after t , leading to a situation in which both $\dot{a}(t) > 0$ and $\dot{K}(t) > 0$, which cannot happen per lemma 10. If $\dot{a}(t) = 0$ and $K(t) \geq 0$, then assuming $a > 0$ and $F > 0$, production must be constrained at t , and equation (28) implies that a must begin rising immediately after t , again leading to a situation in which both $\dot{a}(t) > 0$ and $\dot{K}(t) > 0$, which cannot happen per lemma 10.⁶¹ \square

Lemma 12. *The drilling and extraction rates cannot both be strictly increasing ($\dot{a} > 0$ and $\dot{F} > 0$ is inadmissible).*

Proof. Suppose not. That is, suppose $\dot{a}(t) > 0$ and $\dot{F}(t) > 0$ for some time t . Note that $\dot{F}(t) > 0$ implies that production is constrained at t . This situation cannot continue forever, as explained in lemma 10. If production falls below the constraint, there will again be a contradiction (a discontinuity in either a or F) for the same reason given in lemma 10. \square

Lemma 13. *The drilling and extraction rates cannot both be weakly increasing ($\dot{a} \geq 0$ and $\dot{F} \geq 0$ is inadmissible).*

Proof. Here again, we only need to consider cases involving equalities. If $\dot{F}(t) = 0$, then production must be constrained, such that $F(t) = K(t)$ and therefore $\dot{K}(t) = 0$ as well. This, combined with $\dot{a}(t) \geq 0$, gives us the same situation as in lemma 11. Similarly, if $\dot{a}(t) = 0$ and $\dot{F}(t) \geq 0$, then production again must be constrained, such that $\dot{K}(t) \geq 0$, and we have the same situation as in lemma 11.⁶² \square

The four preceding lemmas complete the proof of rule (3) in the theorem.

Lemma 14. *Once capacity is strictly decreasing over time, it must continue to strictly decrease.*

Proof. Suppose by contradiction that $K(t)$ transitions from being strictly decreasing to weakly increasing. Because $\dot{K}(t)$ is continuous, any such transition must involve a moment in which $\dot{K} = 0$. Let time \hat{t} denote this moment.

First, either production is constrained, in which case $\dot{F}(\hat{t}) = 0$, or production is unconstrained, in which case $\dot{F}(\hat{t}) < 0$. Thus, $\dot{F}(t) \leq 0$.

Second, we cannot have $\dot{a}(\hat{t}) \geq 0$, since this would contradict lemma 11. Thus, $\dot{a}(\hat{t}) < 0$.

Third, we cannot have $\dot{F}(\hat{t}) = 0$, since this must lead to a contradiction. Since $\dot{a}(\hat{t}) < 0$, then \dot{K} would be strictly negative after time \hat{t} and strictly positive before time \hat{t} , contradicting the premise that \dot{K} transitions from strictly negative to positive at \hat{t} . Thus, $\dot{F}(t) < 0$ with production unconstrained.

So we have narrowed the possibilities to $\dot{F}(t) < 0$ at \hat{t} with $\dot{a}(\hat{t}) < 0$. This, however, is an immediate contradiction, since we must have $\dot{a} > 0$ if production is unconstrained. \square

Lemma 15. *Once production is strictly decreasing over time, it must continue to strictly decrease.*

⁶¹With infinite reserves, it is permissible to have both $\dot{a}(t) = 0$ and $\dot{K}(t) = 0$, since this is the steady state. However, if either inequality is strict, the dynamics will lead to a situation in which both $\dot{a}(t) > 0$ and $\dot{K}(t) > 0$, which cannot happen per lemma 10.

⁶²With infinite reserves, it is permissible to have both $\dot{a}(t) = 0$ and $\dot{F}(t) = 0$, since this is the steady state. However, if either inequality is strict, the dynamics will lead to a situation in which both $\dot{a}(t) > 0$ and $\dot{F}(t) > 0$, which cannot happen per lemma 12.

Proof. Suppose not, and let \hat{t} denote the time at which production transitions to weakly increasing. First, production must be constrained at \hat{t} (since the price will be falling immediately afterward), so it must be that $\dot{K}(\hat{t}) \geq 0$ (since $K(t)$ is continuous). Thus, by lemma 14 above we cannot have $\dot{K}(t) < 0$ prior to \hat{t} . But then we also cannot have $\dot{F}(t) < 0$ just prior to \hat{t} without generating a discontinuity in F at \hat{t} , contradicting lemma 5. \square

Lemma 16. $\gamma(t) > 0$.

Proof. From lemma 9 above, we know that $K(t)$ approaches zero in the limit. Because reserves are finite, it must also be that $a(t)$ is not bounded away from zero in the limit. Thus, there must eventually be a time t at which $\frac{XP(K(t))}{\lambda+r} > d(a(t))$ will hold, with $K(\tau) < K(t)$ for all $\tau > t$. Recall that $\theta(t)$ denotes the shadow value of capacity (which is in turn equal to λ times the stock of remaining oil in the drilled wells). This value must be, at minimum, the discounted revenue stream that would be generated by producing the capacity at the constraint forever (per equation 15), since producing below the constraint is only optimal when doing so increases discounted revenues and therefore the value of capacity. As K goes to zero, the discounted revenue from a marginal unit of capacity at t must be at least $P(K(t))/(\lambda+r)$. Thus, for sufficiently large t , we must have $\theta(t)X > d(a(t))$, which violates equation (9) if $\gamma(t) = 0$. \square

Lemma 17. *Drilling will always occur at the optimum, and optimality requires complete exhaustion of reserves in the limit as $t \rightarrow \infty$.*

Proof. Suppose by contradiction that R is bounded away from zero in the limit. That is, there exists some $\hat{R} > 0$ such that $R(t) > \hat{R}$ for all $t > 0$. (Note that $R(t)$ must be weakly decreasing over time forever. Thus, it suffices to consider and rule out by contradiction the case in which $R(t)$ is bounded from below.) But since $\gamma_0 > 0$, this violates the transversality condition (14). \square

The preceding lemma completes the proof of rule (4) in the theorem.

Lemma 18. *At any time that drilling ceases, production must be constrained for a measurable period afterward.*

Proof. Let \hat{t} denote the time that drilling stops. To see that production must be constrained at this time, suppose by contradiction that it is not. In that case, we must have $\dot{\theta}(\hat{t}) = r\theta(\hat{t})$ (by equation 13), further implying by equation (9) and the continuity of $\dot{\theta}(t)$ that $\dot{d}(a(\hat{t})) = rd(a(\hat{t}))$. Note that $d(a)$ is strictly monotonically increasing and $d(0) > 0$, implying that the drilling rate must be rising at \hat{t} . But this is a contradiction, since $a(\hat{t}) = 0$ and the continuity of $a(t)$ implies that $a(t)$ must be decreasing at \hat{t} .

So production is constrained at \hat{t} , which is equivalent to $P(\hat{t}) > \lambda\theta(\hat{t})$. Because $P(t)$ and $\theta(t)$ are continuous, it must therefore be the case that production is constrained for a measurable time period after \hat{t} .

Intuitively, it is sub-optimal to stop drilling and immediately produce below the constraint because drilling is costly, and this combination of actions effectively wastes the capacity generated by the final amount of drilling. \square

The preceding lemma completes the proof of rule (5) in the theorem.

Lemma 19. *Starting from $K_0 = 0$, drilling must initially be strictly decreasing over time.*

Proof. We showed in lemmas 1 and 2 that drilling must begin immediately if $K_0 = 0$ and that capacity must therefore initially be strictly increasing. We also showed in lemma 11 that both drilling and capacity cannot simultaneously be weakly increasing; thus, drilling must initially be strictly decreasing. \square

Lemma 20. *Given demand $P(F)$, marginal drilling costs $d(a)$, and initial reserves R_0 , there exists $\delta > 0$ such that if initial capacity $K_0 < \delta$, then drilling must begin immediately and initially be strictly decreasing over time.*

Proof. Consider the path beginning from $R(0) = R_0$ and $K(0) = 0$. Per lemma 19, the drilling rate must initially be strictly positive and strictly decreasing over time. Consider drilling until time t_δ , at which point $K(t_\delta) = \delta > 0$, $R(t_\delta) = R_\delta < R_0$, and $\dot{a}(t_\delta) < 0$ (the continuity of $\dot{a}(t)$ guarantees that such a time exists). If the initial conditions of the problem were such that $K(0) = \delta$ and $R(0) = R_\delta$, then the optimal path from $t = 0$ would simply be the continuation of the original path beginning from $R(0) = R_0$ and $K(0) = 0$, so that we would have $a(0) > 0$ and $\dot{a}(0) < 0$. To complete the proof, we must show that we will still have $a(0) > 0$ and $\dot{a}(0) < 0$ for the path beginning at $K(0) = \delta$ and $R(0) = R_0$, for δ sufficiently small.

Note that $U(F)$ and $D(a)$ are continuously differentiable in their arguments with $U(F)$ strictly concave and $D(a)$ strictly convex. Meanwhile, the constraints all take the form of weak inequalities and are continuously differentiable in their arguments. Thus, the value function (equal to the present value Hamiltonian-Lagrangian evaluated at the optimal control) will be continuously differentiable in the state variables (K, R) in the interior of the state space (see Benveniste and Scheinkman (1979) and Rincón-Zapatero and Santos (2012)).

Let $\tilde{\theta}(t) = e^{-rt}\theta(t)$ and $\tilde{\gamma}(t) = e^{-rt}\gamma(t)$ denote the present values of the two co-state variables. By the necessary conditions for an optimal control, we therefore have that $\dot{\tilde{\theta}}(t) = -\partial H/\partial K = -\phi(t)$ and $\dot{\tilde{\gamma}}(t) = -\partial H/\partial R = 0$ are both continuous in the state variables. The continuity of $\dot{\tilde{\theta}}$ and $\dot{\tilde{\gamma}}$ then imply the continuity of $\tilde{\theta}$ and $\tilde{\gamma}$, which in turn imply the continuity of θ and γ . Thus, θ , γ , and ϕ are all continuous in the state variables. Equation (9), combined with the continuous differentiability and strict monotonicity of $d(a)$, then implies that $a(K, R)$ is continuous in (K, R) , as is $\dot{a}(K, R)$ whenever $a(K, R) > 0$.

It is possible to choose a sufficiently small t_δ , and therefore a sufficiently small δ , such that R_δ can be arbitrarily close to R_0 (since $R(t)$ is continuous). The continuity of $a(K, R)$ and $\dot{a}(K, R)$ then implies that the optimal path starting from initial conditions of $K(0) = \delta$ and $R(0) = R_0$ will have $a(0) > 0$ and $\dot{a}(0) < 0$. \square

Lemma 20 completes the proof of rule (6) in the theorem. With rules (1)-(6) proven, the proof of the theorem itself is also now complete.

E Proofs for rules (7)–(10) pertaining to the general case (including sufficient conditions for a binding flow constraint)

This appendix provides the formal statement and proof for theorem 2 discussed in the main text:

Theorem 2. *In addition to the assumptions of theorem 1, assume that the inverse elasticity of demand $\eta(F) \equiv -F \frac{P'(F)}{P(F)}$ is weakly increasing in F . Then the following rules hold: (7) if at any time \hat{t} production is constrained, then production must be constrained at all subsequent $t > \hat{t}$; (8) given demand $P(F)$, marginal drilling costs $d(a)$, and initial reserves R_0 , there exists $\delta > 0$ such that if initial capacity $K_0 < \delta$, then production is constrained along the entire optimal path (so production is always constrained trivially if $K_0 = 0$); (9) once the rate of drilling strictly declines it can never subsequently strictly increase, further implying that once drilling stops it cannot restart; and (10) drilling will end in finite time if and only if there exists an $F > 0$ such that $\lambda\eta(F) < r$.*

The proof proceeds via the following series of lemmas.

Lemma 21. *Suppose for some $F^* > 0$ that $\lambda\eta(F) < r$ whenever $F < F^*$. Further suppose that at some time t^* , $K(t^*) \leq F^*$ (such times must exist since $K(t)$ goes to zero in the limit). Then production must be constrained at t^* .*

Proof. First note that if $\dot{K}(t^*) \geq 0$, then $\dot{a}(t^*) < 0$ (per lemma 11), such that production must be constrained at t^* for (9) to hold. So consider the case in which $\dot{K}(t^*) < 0$. Suppose, by contradiction, that production is unconstrained at t^* . Then $P(t)$ must rise at rate r , so we have $P'(F)\dot{F}/P(F) = r$. Rearranging and substituting in the inverse demand elasticity, we have $\dot{F} = -rF/\eta(F) < -\lambda F$, where the inequality follows from the fact that $r/\eta(F) > \lambda$. Meanwhile, $\dot{K} \geq -\lambda F$ (holding with equality in the absence of drilling at t^*). Thus, we have F declining faster than K , so that production must be unconstrained forever (since $F(t)$ is continuous), declining at rate $r/\eta(F)$. Thus, total cumulative production from t^* onward equals $F(t^*) \int_{t^*}^{\infty} e^{-rt/\eta(F(t))} dt$. This value is strictly less than the remaining stock of oil that can be produced from the existing capacity at t^* , which is given by $K(t^*)/\lambda$. Thus, capacity does not go to zero in the limit, violating lemma 9. \square

Lemma 22. *Suppose that production is constrained at some time \hat{t} . Then production must also be constrained for all $t > \hat{t}$.*

Proof. We first consider the case in which $\dot{K}(\hat{t}) < 0$. Note that if demand is such that $\lambda\eta(F) < r$ for all $F < K(\hat{t})$, then by lemma 21 above production must be constrained at \hat{t} and, because K must continue to strictly decrease, for all subsequent t . Thus, we need only consider the sub-case in which there exists $F < K(\hat{t})$ such that $\lambda\eta(F) \geq r$.

Proceed by contradiction, and let $\tilde{t} > \hat{t}$ denote the time at which production becomes unconstrained. At and immediately after \tilde{t} we have that $\dot{F}(t) < \dot{K}(t)$ (that is, for production to fall below the constraint, the rate of production must decrease more quickly than capacity decreases). Note that throughout the period in which production is unconstrained, we must

have K sufficiently large that $\lambda\eta(K) \geq r$ (by lemma 21). Just before \tilde{t} , when production is constrained, there are two possibilities: either $\dot{P}/P < r$ or $\dot{P}/P \geq r$.

First consider the possibility that $\dot{P}/P < r$, which requires strictly positive drilling at \tilde{t} because when production is constrained and $a = 0$, $\dot{P}/P = \lambda\eta(K) \geq r$. Moreover, since we have assumed that production is unconstrained, we must have $\dot{a}(\tilde{t}) > 0$. The rate of drilling cannot strictly increase forever, and both $a(t)$ and $\dot{a}(t)$ are continuous while $a > 0$, so there must be some future time at which production becomes constrained again.

Since we must therefore eventually have $F = K$, since neither K nor F may ever increase, and since K must be sufficiently large that $\lambda\eta(K) \geq r$ throughout the period in which production is unconstrained, we must therefore have that F is sufficiently large that $\lambda\eta(F) \geq r$ throughout the period in which production is unconstrained (since K must eventually decrease to meet any F that is reached).

During the period in which production is unconstrained, we have that $\dot{F} = -rF/\eta(F)$ and $\dot{K} = aX - \lambda F$. Then, since $\dot{F}(\tilde{t}) < \dot{K}(\tilde{t})$, and since, throughout the unconstrained period, $a(t)$ is strictly increasing, $\lambda \geq r/\eta(F)$, and F is strictly decreasing, it must be that $\dot{F}(t) < \dot{K}(t)$ throughout the unconstrained period ($\dot{K}(t) - \dot{F}(t)$ is actually increasing throughout the period as F decreases). This result yields a contradiction, since in this case production can never again become constrained without inducing a discontinuity in $F(t)$. Thus, we cannot have production become unconstrained immediately after a period in which $\dot{P}/P < r$.

Now consider the second possibility, in which $\dot{P}/P \geq r$ just before \tilde{t} , and consider the object $P(t) - \lambda\theta(t)$. At \tilde{t} , since production is unconstrained, we must have $P(\tilde{t}) - \lambda\theta(\tilde{t}) = 0$ by equation (7). In addition, we have $\dot{\theta}(\tilde{t}) = r\theta(\tilde{t})$ and (from lemma 7) that $\dot{\theta}(t)$ is continuous. Further, for times t just before \tilde{t} , production is constrained, so $P(t) - \lambda\theta(t) > 0$.

The above facts generate a contradiction. To have $P(t) - \lambda\theta(t) > 0$ for t just prior to \tilde{t} , $P(\tilde{t}) - \lambda\theta(\tilde{t}) = 0$, $\dot{\theta}(\tilde{t}) = r\theta(\tilde{t})$, and $\dot{\theta}(t)$ continuous, it must be the case that $\dot{P}(t)/P(t) < r$. However, for times just before \tilde{t} , we have assumed that $P(t)$ must be rising at a rate weakly greater than r . Thus, we have a contradiction, and we therefore cannot have production become unconstrained immediately after a period in which $\dot{P}/P \geq r$.

We have now shown that, if production is constrained at \hat{t} and $\dot{K}(\hat{t}) < 0$, then production cannot subsequently become unconstrained, regardless of whether $\dot{P}/P < r$ or $\dot{P}/P \geq r$ at and after \hat{t} . Intuitively, the demand elasticity assumption in the theorem implies that the incentive to produce below the constraint increases with K . Thus, if production is already constrained at some time while capacity is decreasing, the incentive to produce below the constraint will only diminish as K decreases further.

We now consider the case in which $\dot{K}(\hat{t}) \geq 0$. This condition requires that $a(\hat{t}) > 0$ and $\dot{a}(\hat{t}) < 0$. Production can never be unconstrained so long as $a(t) > 0$ and $\dot{a}(t) \leq 0$. There are two ways this situation can end: either drilling may cease, or the rate of drilling may start to strictly increase. If drilling ceases, then production must be constrained at the time drilling ceases per lemma 18. At this time, we have $\dot{K} < 0$, so that production must always subsequently be constrained per the results above. If, on the other hand, $a(t)$ remains strictly positive and $\dot{a}(t)$ eventually becomes strictly positive, the continuity of $\dot{a}(t)$ implies the existence of a time t' at which $\dot{a}(t') = 0$. At this time, capacity must be constrained, with $\dot{K}(t') < 0$ (per lemma 11), so that the results above apply, and production must always subsequently be constrained. \square

The preceding lemma proves rule (7) in the theorem.

Lemma 23. *Given demand $P(F)$, marginal drilling costs $d(a)$, and initial reserves R_0 , there exists $\delta > 0$ such that if initial capacity $K_0 < \delta$, then production is constrained along the entire optimal path.*

Proof. With K_0 sufficiently small, we know from rule (6) that drilling will begin instantly and must initially be strictly decreasing over time, implying that production must initially be capacity-constrained. Rule (7) then implies that all subsequent production must be constrained as well. \square

The preceding lemma proves rule (8) in the theorem.

Lemma 24. *Over any interval in which $a(t) > 0$, once the rate of drilling strictly declines it can never subsequently strictly increase.*

Proof. We will again proceed by contradiction. We already showed in rule (3) that the rate of drilling and the rate of extraction (along with capacity) cannot weakly increase simultaneously. Thus, for $\dot{a}(t)$ to transition from being strictly negative to strictly positive, capacity and production must be strictly declining both at the transition itself and during any measurable interval in which $\dot{a}(t) \geq 0$. Since we also know that $a(t)$ cannot monotonically increase in the limit as $t \rightarrow \infty$, we must additionally have a subsequent transition in which $\dot{a}(t)$ changes from strictly positive to negative.

Let t_1 denote the time of the first transition, when $\dot{a}(t) < 0$ before t_1 and $\dot{a}(t) \geq 0$ after t_1 . Let $t_2 > t_1$ denote the time of the second transition, when $\dot{a}(t) > 0$ before t_2 and $\dot{a}(t) \leq 0$ after t_2 . At t_1 , production must be constrained, so by rule (7) it must be constrained through t_2 and beyond. We therefore also have that $\dot{F}(t)$ and $\dot{P}(t)$ are continuous, as is $\ddot{a}(t)$ (via differentiating equation (28)).

Since $\dot{a}(t_1) = \dot{a}(t_2) = 0$, equation (28) implies that the following two equations must hold:

$$\frac{(r + \lambda)d(a(t_1))}{X} + \frac{\lambda\gamma_0}{X}e^{rt_1} = P(t_1) \quad (29)$$

$$\frac{(r + \lambda)d(a(t_2))}{X} + \frac{\lambda\gamma_0}{X}e^{rt_2} = P(t_2), \quad (30)$$

where $a(t_2) > a(t_1)$ and $P(t_2) > P(t_1)$.

Moreover, since $\dot{a}(t) < 0$ before t_1 and $\dot{a}(t) \geq 0$ after t_1 , it must be that $\ddot{a}(t_1) \geq 0$. So if we differentiate equation (28) at t_1 and divide through by $P(t_1)$, we must have that:

$$\frac{r\lambda\gamma_0e^{rt_1}}{XP(t_1)} \geq \frac{\dot{P}(t_1)}{P(t_1)}. \quad (31)$$

On the other hand, by similar arguments, at t_2 we must have:

$$\frac{r\lambda\gamma_0e^{rt_2}}{XP(t_2)} \leq \frac{\dot{P}(t_2)}{P(t_2)}. \quad (32)$$

Given our functional form assumption for $P(F)$, it must be the case that $\dot{P}(t_2)/P(t_2) \leq \dot{P}(t_1)/P(t_1)$, implying that $e^{rt_1}/P(t_1) \geq e^{rt_2}/P(t_2)$. Why? Note that, when production is constrained and no drilling is taking place, the rate at which $P(t)$ increases is given by $\dot{P}(t)/P(t) = \lambda\eta(K(t))$. Since $K(t_2) < K(t_1)$ and $a(t_2) > a(t_1)$, it must be that $\dot{P}(t_2)/P(t_2) \leq \dot{P}(t_1)/P(t_1)$, since we have assumed that $\eta(F)$ is weakly increasing.

Thus, for equations (31) and (32) to hold, we must have that price rises, on average, weakly faster than r between t_1 and t_2 . However, this means that equations (29) and (30) cannot both hold, since $d(a(t))$ cannot rise faster than r and must be rising strictly slower than r in the neighborhoods of t_1 and t_2 (implying that $d(a(t))$ must be rising strictly slower than r , on average, between t_1 and t_2). Thus, we have a contradiction.

Intuitively, it would only make sense for the planner to increase the rate of drilling, after it had been decreasing, if a substantial increase in the rate of price increase is expected in the near future. The assumption that the inverse demand elasticity decreases as production decreases rules this possibility out. \square

Lemma 25. *Once drilling stops, it cannot restart. Thus, in general, once the rate of drilling strictly declines it can never subsequently strictly increase.*

Proof. This proof uses similar arguments as in lemma 24. We begin by showing that once drilling stops, it cannot restart. Proceed by contradiction. Let \hat{t} denote the time that drilling stops, and let \tilde{t} denote the time that drilling restarts.

At \hat{t} we must have $\dot{a}(\hat{t}) \leq 0$, and at \tilde{t} we must have that $\dot{a}(\tilde{t}) \geq 0$. Thus, by equation (28) the following two equations must hold:

$$\frac{(r + \lambda)d(0)}{X} + \frac{\lambda\gamma_0}{X}e^{r\hat{t}} \leq P(\hat{t}) \quad (33)$$

$$\frac{(r + \lambda)d(0)}{X} + \frac{\lambda\gamma_0}{X}e^{r\tilde{t}} \geq P(\tilde{t}) \quad (34)$$

By the continuity of $P(t)$, there must exist a $t_1 \in (\hat{t}, \tilde{t})$ such that:

$$\frac{(r + \lambda)d(0)}{X} + \frac{\lambda\gamma_0}{X}e^{rt_1} = P(t_1). \quad (35)$$

Since $a(t)$ cannot be monotonically increasing in the limit as $t \rightarrow \infty$, there must exist some time t_2 after \tilde{t} at which $\dot{a}(t)$ changes from strictly positive to negative. At t_2 , we must have $\dot{a}(t_2) = 0$ and therefore by equation 28 that

$$\frac{(r + \lambda)d(a(t_2))}{X} + \frac{\lambda\gamma_0}{X}e^{rt_2} = P(t_2). \quad (36)$$

Production must be constrained at both t_1 and t_2 . Thus, $P(t)$ is continuously differentiable at both t_1 and t_2 (since $K(t)$ and $P(F)$ are continuously differentiable), and following

the logic of lemma 24 we have:

$$\frac{r\lambda\gamma_0 e^{rt_1}}{XP(t_1)} \geq \frac{\dot{P}(t_1)}{P(t_1)} \quad (37)$$

$$\frac{r\lambda\gamma_0 e^{rt_2}}{XP(t_2)} \leq \frac{\dot{P}(t_2)}{P(t_2)}. \quad (38)$$

Because production is constrained at both t_1 and t_2 and because $K(t_1) > K(t_2)$, we have that $\dot{P}(t_2)/P(t_2) \leq \dot{P}(t_1)/P(t_1)$, which implies that $e^{rt_2}/P(t_2) \leq e^{rt_1}/P(t_1)$. Then, per the arguments given in lemma 24, equations (35), (36), (37), and (38) cannot simultaneously hold, and we have a contradiction. Thus, once drilling stops, it cannot subsequently restart. \square

The preceding lemma proves rule (9) in the theorem.

Lemma 26. *Suppose that, for some $F^* > 0$, if $F < F^*$ then $\lambda\eta(F) < r$. (Note that such an F^* must exist for any inverse demand function that satisfies the conditions of the theorem and has $P(0)$ finite. However, the converse is not true; that is, the existence of F^* does not imply that $P(0)$ is finite.) Then drilling must stop in finite time.*

Proof. Suppose, by contradiction, that drilling does not stop in finite time. We know that both $a(t)$ and $K(t)$ must approach zero in the limit, so there must be a point at which $K(t) < F^*$ and both $a(t)$ and $K(t)$ are decreasing forever (by rule 9 and lemma 14 respectively). Because $a(t)$ is decreasing, production must be constrained forever.

Note that, when production is constrained and no drilling is taking place, the rate at which $P(t)$ increases is given by $\dot{P}(t)/P(t) = \lambda\eta(t)$. Thus, when $K < F^*$, it must be the case that $P(t)$ is rising strictly more slowly than r (with or without drilling activity). Necessary conditions (7) and (13) then jointly imply that $\theta(t)$ must also be rising strictly more slowly than r . Then, because $\gamma(t)$ rises at r , we must have that for t sufficiently large, $\gamma(t) > \theta(t)X$. This is a contradiction, however, since we must have $\gamma(t) < \theta(t)X$ for drilling to occur, per equation (9).

Intuitively, if demand is such that price eventually must rise more slowly than the rate of interest while production declines, then all of the wells must be drilled in finite time. \square

Lemma 27. *Suppose that $\lambda\eta(F) \geq r$ for all $F > 0$. (Note that this condition implies that $P(0)$ is infinite, though again the converse is not true.) Then drilling will not stop in finite time.*

Proof. Suppose, by contradiction, that drilling stops in finite time. Denote the time drilling stops by \hat{t} . First, consider the case in which the inequality is strict; that is, $\lambda\eta(F) > r$ for all $F > 0$. At \hat{t} , it cannot be optimal to produce at the constraint. Why? Doing so causes $P(t)$ to rise at the rate $\lambda\eta(t) > r$. Moreover, at any future time, price must rise at a rate weakly greater than r (with a strict inequality any time production is constrained). Thus, it is profitable to deviate at \hat{t} by producing less at \hat{t} and producing more later. However, producing below the constraint immediately after drilling contradicts rule (5), so we have a contradiction.

Now consider the case in which there exists $F^* > 0$ such that $\lambda\eta(F) = r$ for all $F < F^*$. If drilling ceases at some capacity level $K(\hat{t}) > F^*$, then it will be optimal immediately to produce below the constraint, and we again have the contradiction above. What if drilling ceases at a capacity level $K(\hat{t}) \leq F^*$? In this case, producing at the constraint causes $P(t)$ to rise at exactly r . Rule (5) tells us that production must be at the constraint at \hat{t} and for a measurable period afterward. Note that production must in fact always be at the constraint in this case: the continuity of $F(t)$ implies that if production were to fall below the constraint, $P(t)$ would have to at least briefly rise faster than r , which cannot happen when production is unconstrained.

Thus, from \hat{t} onward, it must be that production is at the constraint and that $P(t)$ rises at r forever. For equations (7) and (13) both to hold forever, it must also be the case that $\theta(t)$ rises at r forever, and in particular that it rises at r at \hat{t} . But this leads to a contradiction: if both $\theta(t)$ and $\gamma(t)$ are rising at r at \hat{t} , then by equation (9) we have that $d(a(t))$ must also be rising at r at \hat{t} . However, the fact that $a(\hat{t}) = 0$ and the continuity of $a(t)$ imply that $a(t)$ must be decreasing at \hat{t} . Thus, if $\lambda\eta(F) \geq r$ for all $F > 0$, then drilling will not stop in finite time. \square

The two preceding lemmas together prove rule (10). With rules (7)-(10) proven, the proof of the theorem itself is also now complete.

F Extending the drilling model to allow for costly above-ground storage (including the proof that production is always constrained in such a model)

We have shown that our dynamics imply that it is possible, on the equilibrium path, for price temporarily to rise at a rate faster than the rate of interest while production is constrained. However, this result has not taken into account the possibility that oil may be stored above-ground. This appendix therefore asks how the availability of costly above-ground storage would affect our conclusions.⁶³

For tractability, we assume that storage costs take the form of an iceberg storage cost—that is, a continuous, proportional erosion at rate m of the quantity (and therefore the value) of stored oil. In this case, a stockpiler receiving capital gains but no convenience yield would be indifferent between buying or selling oil if the oil price is expected to rise at the constant percentage rate of $m + r$.

Below, we derive and briefly interpret the necessary conditions of this model. We then argue that allowing for costly storage does not affect our main conclusions from theorem 2: once production is capacity constrained, it must always subsequently be constrained, and if initial capacity is sufficiently low, production is always constrained along the equilibrium path. Intuitively, storage offers a substitute means by which consumption can be deferred to the future, so introducing storage into the model will not increase the incentive to produce below the available capacity constraint.

Note that this model still accommodates the possibility that the oil price can rise more quickly than r on the equilibrium path while production is constrained. In fact, during any measurable interval over which storage is occurring (and therefore the oil price is rising at the rate $m + r$), production *must* be constrained. Why? As we have seen, whenever production is unconstrained, the oil price must be rising at r , not $r + m$. Hence, even if aboveground storage is possible, the equilibrium may still involve intervals over which price rises in percentage terms at rate $r + m$, and production must be at capacity throughout each of these intervals. Such a steep rise in expected prices actually occurred in winter 1998–1999 (see figure 1) and was in fact accompanied by a surge in above-ground storage (see figure 13 in appendix A).

However, as figure 13 indicates, aboveground storage also occurred when expected capital gains were significantly smaller. This suggests that aboveground storage also provides a convenience yield that makes carrying inventory attractive to at least some stockpilers even when capital gains do not cover their interest and storage costs.⁶⁴ If we amend our account to include heterogeneous stockpilers, each with a convenience yield that is strictly concave in the amount stored, then aboveground storage would still put a ceiling on capital gains

⁶³It is clear that, if storage were costless and limitless, then the oil price could never be expected to rise more quickly than r . Real-world logistical costs and storage constraints argue against this hypothetical, as does the fact that we observe instances in futures market data when the price is expected to rise more quickly than r .

⁶⁴A convenience yield naturally arises in our setting from the fact that many leases are not connected to the pipeline network but instead sell their produced oil via tanker truck. Firms must store produced oil in on-site storage tanks between truck pick-ups.

of $r + m$ and would still be accompanied by production at the constraint whenever the percentage change in the oil price differs from r ; however, because of the convenience yields, storage could occur even when holding inventory results in net capital losses.

F.1 Necessary conditions and their implications

We now formally present the Hamiltonian-Lagrangean corresponding to the social planner's problem with costly storage and derive and interpret the implied necessary conditions. With storage, let the choice variable $C(t)$ denote oil consumption generating instantaneous utility flow of $U(C(t))$. Let $S(t) \geq 0$ denote a new state variable representing the quantity of oil stored at time t . The equation of motion for S with an iceberg storage cost of m is then $\dot{S}(t) = F(t) - C(t) - mS(t)$, where $F(t)$ is still the choice variable representing oil production, as above. The other variables and constraints—the drilling rate and drilling costs, production capacity and its equation of motion, and the stock of remaining wells and its equation of motion—are all the same as above.

The current-value Hamiltonian-Lagrangean of this maximization problem is then given by:

$$H = U(C(t)) - D(a(t)) + \theta(t)[a(t)X - \lambda F(t)] + \gamma(t)[-a(t)] \\ + \phi(t)[K(t) - F(t)] + \mu(t)[F(t) - C(t) - mS(t)], \quad (39)$$

where $\mu(t)$ is a new co-state variable on $S(t)$.

Necessary conditions are given by the following (re-numbering some conditions that match those from the original problem without storage):

$$C(t) \geq 0, \quad U'(C(t)) - \mu(t) \leq 0, \quad c.s. \quad (40)$$

$$F(t) \geq 0, \quad \mu(t) - \lambda\theta(t) - \phi(t) \leq 0, \quad c.s. \quad (41)$$

$$F(t) \leq K(t), \quad \phi(t) \geq 0, \quad c.s. \quad (42)$$

$$a(t) \geq 0, \quad \theta(t)X - d(a(t)) - \gamma(t) \leq 0, \quad c.s. \quad (43)$$

$$\dot{R}(t) = -a(t), \quad R_0 \text{ given} \quad (44)$$

$$\dot{\gamma}(t) = r\gamma(t) \quad (45)$$

$$\dot{K}(t) = a(t)X - \lambda F(t), \quad K_0 \text{ given} \quad (46)$$

$$\dot{\theta}(t) = -\phi(t) + r\theta(t) \quad (47)$$

$$\dot{S}(t) = F(t) - C(t) - mS(t), \quad S_0 = 0 \quad (48)$$

$$\dot{\mu}(t) \leq (r + m)\mu(t), \quad S(t) \geq 0, \quad c.s. \quad (49)$$

$$K(t)\theta(t)e^{-rt} \rightarrow 0, \quad R(t)\gamma(t)e^{-rt} \rightarrow 0, \quad \text{and} \quad S(t)\mu(t)e^{-rt} \rightarrow 0 \quad \text{as} \quad t \rightarrow \infty. \quad (50)$$

These necessary conditions differ in two main ways from those presented above. First, we now have necessary conditions that characterize production and consumption incentives separately. Condition (40) says that the price of oil will equal the shadow value on storage capacity whenever consumption is strictly positive and will be strictly less than this value when consumption is zero. Meanwhile, condition (41) is the same as above, but with the

price of oil replaced by the shadow value of storage $\mu(t)$. When consumption is strictly positive (and it always will be, as we show below), then $\mu(t)$ can be eliminated to yield the same condition characterizing production incentives that we had before. Second, the necessary conditions now include (48), which is the new equation of motion for storage, along with condition (49), which characterizes the evolution of its co-state variable. Since consumption is always strictly positive, this latter condition says that price must rise at rate $r + m$ during any interval over which storage occurs. Thus, during any such interval, the production constraint must bind.

F.2 Proof extending main theorems to the case of costly storage

We now prove that the assumptions of theorems 1 and 2 are sufficient to guarantee that once production is constrained, all subsequent production must also be constrained and that there exists a $\delta > 0$ (which depends on parameters) such that for any $K_0 < \delta$, production is constrained along the entire optimal path.

Formally, we present the following theorem:

Theorem 3. *Maintain the assumptions of theorems 1 and 2. In addition, suppose that above-ground storage technology is available, with an iceberg storage cost of $m > 0$ and an initial storage quantity of $S_0 = 0$. In this case, rules (1)-(2), (4)-(5), and (7)-(8) continue to hold.*

We prove the theorem by proceeding through nearly the same set of lemmas as were used to prove theorems 1 and 2. Most of these lemmas will go through with only minor or no changes. The proofs of lemmas 5 and 6 require more substantial modifications. Rules (3) and (6) no longer hold with storage, and the statements of lemmas 10 through 15 and lemmas 19 and 20 must therefore be modified.

We will also henceforth refer to $U'(C(t))$ as $P(C(t))$.

Lemma 1 goes through using identical arguments: in a system that starts with $K_0 = S_0 = 0$, if drilling does not start immediately, then drilling will never occur, violating our assumption that $P(0)X/(r + \lambda) > d(0)$.

Lemma 2 is as stated and proven without any changes.

Lemma 3 goes through with only a minor modification. Once we have $\phi(\tau) = 0$ and $\theta(\tau) = 0$ for all τ , it follows from equation (41) that $\mu(\tau) = 0$ for all τ . It must then also be the case, via equation (40), that $P(C(t)) = 0$ for all τ . As in the original lemma 3, this cannot happen, since the amount of available production capacity is finite, and at some point production and consumption must decline to a point at which $P(C) > 0$.

The statement of **lemma 4** now becomes: it is never optimal to set C so high that $P(C) = 0$. Proceed by contradiction. Whenever $P(C(t)) = 0$, we have $C(t) > 0$ such that equation (40) must hold with equality and $\mu(t) = 0$. Equation (49) then implies that μ is thereafter always zero. Then, for equation (41) to hold without having $\theta = 0$, it must be that $F = 0$ forever. Then we have a contradiction. Why? Note that $P(C(t))$ can never rise above zero, since this would violate condition (40) (since $\mu = 0$ forever). But if $P(C(t)) = 0$ forever, then $C(t)$ is bounded away from zero forever, and whatever is in storage will eventually run out. With nothing remaining in storage, and nothing more produced, $C(t)$ must go to zero, and therefore $P(C(t))$ must go above zero.

In **lemma 5**, the proofs that $\gamma(t)$, $\dot{\gamma}(t)$, $\theta(t)$, $a(t)$, $R(t)$, and $K(t)$ are continuous do not change. It is furthermore immediately true that $S(t)$ is continuous, since both $C(t)$ and $F(t)$ are bounded. Here, we prove that $C(t)$, and therefore $P(t)$, are continuous. We then prove that $C(t) = 0$ is never optimal when $K(t) > 0$, which immediately implies that $\mu(t)$ is always equal to $P(t)$ and is therefore also continuous. Finally, we prove that $F(t)$ is continuous.

First, note that $C(t)$ can never discontinuously “jump” down, since for equation (40) to hold, $\mu(t)$ would have to jump up. However, such a jump is forbidden by equation (49). So suppose by contradiction that $C(t)$ jumps up at t . Per equation (40), μ must then jump down at t . Such a jump is possible only if $S(t) = 0$. With $S(t) = 0$, the only way $C(t)$ can jump up is if $F(t)$ jumps up as well, since we must have $C(t) = F(t)$. Thus, F will be strictly greater than zero at or immediately after time t , meaning that equation (41) must hold with equality at this time. This is not possible, however, since μ jumps down at t , $\theta(t)$ is continuous, and ϕ cannot be decreasing at t (since $\phi = 0$ before the jump up in F). Thus, $C(t)$ cannot discontinuously jump either up or down, and must therefore be continuous.

The continuity of $C(t)$ and of $P(C)$ immediately imply that $P(t)$ is continuous.

Suppose by contradiction that $C(t) = 0$ while $K(t) > 0$. If we assume $P(0) = \infty$, then equation (40) is violated, leading to an immediate contradiction. So assume $P(0)$ is finite instead. First suppose that C continues to equal zero forever after time t . Note that we cannot have $F = 0$ forever, due to the transversality condition (50) for capacity, so it must be that $S > 0$ forever. But then this violates the transversality condition for S , as μ must be rising at the rate $r + m$, per equation (49).

So suppose that C increases away from zero at some time $\tau \geq t$. When this happens, P must decrease, and μ must also decrease (since it must be equal to P when $C > 0$). Note μ cannot be decreasing whenever $S > 0$, so it must be the case that there is no storage at time τ , and therefore that C and F are equal at and immediately after τ . Note $C(t)$ is continuous, and $K(t) > 0$, so it must be that, with C and F equal, production is unconstrained at and immediately after τ , with $\phi = 0$. In this case, equation (41) cannot hold with equality, since $\mu = P$ is falling and θ is rising at rate r . Thus, we have a contradiction, proving that we cannot ever have $C(t) = 0$ when $K(t) > 0$, and therefore whenever $t > 0$.

Since $C(t) > 0$ for all $t > 0$ (and at $t = 0$ if $K_0 > 0$), it immediately follows from equation (40) and the continuity of $P(t)$ that $\mu(t)$ is continuous.

To complete the parallel to lemma 5, we finally consider the continuity of $F(t)$. Suppose by contradiction that F jumps down at t . Because $K(t)$ is continuous and strictly greater than zero, we must have that production is unconstrained immediately after t , causing ϕ to equal zero. Given the continuity of $\mu(t)$ and $\theta(t)$, for equation (41) to hold it must be that production was also unconstrained (with $\phi = 0$) before the jump. Moreover, because $C(t)$ is continuous and strictly greater than zero, we must have that $F > C$ before t and/or $F < C$ after t . Either case requires $S > 0$ immediately after t . But then $\mu(t)$ will be increasing at the rate $r + m$ while $\theta(t)$ is increasing at only rate r , leading to a violation of condition (41).

What if F jumps up at t ? Again, for (41) to hold, production must be unconstrained both before and after the jump (the inequality cannot be strict, with $\phi > 0$, after the jump since we will have $F > 0$). The jump again implies that $S(t) > 0$, so that $\mu(t)$ will be increasing at the rate $r + m$ while $\theta(t)$ is increasing at only rate r , leading again to a violation of condition (41). Thus, a jump in $F(t)$ in either direction leads to a contradiction, implying the continuity of $F(t)$.

Lemma 6 holds, but with a modified proof. Suppose by contradiction that $F(t) = 0$ while $K(t) > 0$. Note that we cannot have $F = 0$ forever, since this would violate the transversality condition (50) for capacity. So F must ultimately increase away from zero. Let τ denote the time that F increases. With $C(t) > 0$, and both $C(t)$ and $F(t)$ continuous, it must be the case that for a measurable period after τ , production is unconstrained while $S > 0$. As was the case with the proof of the previous lemma, these circumstances will violate condition (41), and we have a contradiction. In general, production cannot be unconstrained while storage is strictly positive.

Lemma 7 goes through, with the addition that $\dot{S}(t)$ is continuous (since $F(t)$, $C(t)$, and $S(t)$ are all continuous).

Lemmas 8 and 9 are as stated and proven without any changes.

Lemmas 10 and 11 no longer hold. Instead, we show the following: It is possible to have $\dot{a}(t) \geq 0$ and $\dot{K}(t) \geq 0$ over some time interval $[t_1, t_2]$. However, this is admissible only if there exists an interval $[t_a, t_b]$ during which $S(t) > 0$ for all $t \in [t_a, t_b]$ with storage beginning at time $t_a < t_2$ (i.e., the interval of storage starts strictly before the *end* of the interval during which drilling and capacity are both weakly increasing) and if storage ends, it ends at some time $t_b > t_2$ at which $\dot{a}(t_b) < 0$ and $\dot{K}(t_b) < 0$. If on the other hand $S(t) > 0$ for all $t > t_a$ (i.e., $t_b = \infty$), then there must be some time t_3 such that $\dot{a}(t) < 0$ and $\dot{K}(t) < 0$ for all $t > t_3$. In addition, whether or not storage ends in finite time, production must be constrained over the entire interval $[\min(t_a, t_1), t_b]$.

First, we argue that production must be constrained throughout the entire interval $[t_1, t_2]$ during which $\dot{a}(t) \geq 0$ and $\dot{K}(t) \geq 0$. Suppose by contradiction that $F(\hat{t}) < K(\hat{t})$ for some $\hat{t} \in [t_1, t_2]$. Since production is unconstrained, we must have $S(\hat{t}) = 0$, which implies $F(\hat{t}) = C(\hat{t})$. With unconstrained production, price rises at the interest rate, so that $\dot{C}(\hat{t}) = \dot{F}(\hat{t}) < 0$. Meanwhile, with $a(\hat{t}) > 0$, marginal drilling costs $d(a(\hat{t}))$ must also be rising at the interest rate. Thus, everything in square brackets in equation (28) rises at the interest rate, such that $a(t)$ and $K(t)$ continue to increase, while $C(t) = F(t)$ continues to decrease, and production remains unconstrained. Note there is no way escape this situation without inducing a discontinuity in either $a(t)$ or $F(t)$: returning production to the capacity constraint induces a discontinuity in $F(t)$, while continuing on as is with production below the constraint and $a(t)$ rising induces a discontinuity in $a(t)$ the moment the last well is drilled. Since we have already shown that $a(t)$ and $F(t)$ must be continuous, we conclude that production must be constrained throughout the interval.

We next argue that the interval with $\dot{a}(t) \geq 0$ and $\dot{K}(t) \geq 0$ can only be escaped during an interval of continuous storage (i.e., $t_2 \in [t_a, t_b]$). Suppose first that no storage occurs at all on the interval $[t_1, t_2]$. Then the arguments in the original lemmas imply that the situation with $\dot{a}(t) \geq 0$ and $\dot{K}(t) \geq 0$ cannot be escaped without generating a discontinuity in either $a(t)$ or $F(t)$. So storage must occur on the interval. Now suppose that a period of storage occurs on the interval, but that this period of storage begins and ends before the interval concludes: $t_a < t_b \leq t_2$. Since $F(t) = K(t)$ during the period of storage, since $F(t)$ and $K(t)$ are both continuous, and since $S(t_b) = 0$, we must have that $C(t_b) = F(t_b) = K(t_b)$. Moreover, note from above that production must be constrained during the entire interval that $\dot{a}(t) \geq 0$ and $\dot{K}(t) \geq 0$. Thus, we must have $C(t) = F(t) = K(t)$ for all $t \in [t_b, t_2]$, during which $\dot{a}(t) \geq 0$ and $\dot{K}(t) \geq 0$. But then the arguments in the original lemmas again imply that this situation cannot be escaped without generating a discontinuity in either $a(t)$

or $F(t)$, and we again generate a contradiction. Thus, to escape an interval on which $\dot{a}(t) \geq 0$ and $\dot{K}(t) \geq 0$, we must have $t_2 \in [t_a, t_b]$.

Assuming that an interval of rising drilling and capacity can in fact be escaped during a period of storage (as we argue immediately below), it therefore immediately follows that production must be constrained over the entire interval $[\min(t_a, t_1), t_b]$.

To complete the revised lemma, it remains to show that: (1) it is in fact possible to escape the interval of $\dot{a}(t) \geq 0$ and $\dot{K}(t) \geq 0$ in the presence of storage, (2) if storage ends, it must end at some time $t_b > t_2$ at which $\dot{a}(t_b) < 0$ and $\dot{K}(t_b) < 0$; and (3) if $S(t) > 0$ for all $t > t_a$, then there must be some time t_3 such that $\dot{a}(t) < 0$ and $\dot{K}(t) < 0$ for all $t > t_3$. We proceed in order.

To see part (1), note that the time path cannot exit the condition in which $\dot{a} \geq 0$ and $\dot{K} \geq 0$ by passing through a point at which $\dot{a}(t) > 0$ and $\dot{K}(t) = 0$. At such a time, we have $\ddot{K}(t) = \dot{a}(t)X - \lambda\dot{K}(t) = \dot{a}(t)X > 0$, which would immediately lead back to a situation with $\dot{a} \geq 0$ and $\dot{K} \geq 0$. Escape must therefore involve a time \tilde{t} at which $\dot{a}(\tilde{t}) = 0$ while $\dot{K}(\tilde{t}) \geq 0$. This is not possible while $S(t) = 0$, for we have already shown that the production constraint must bind, and with $C(t) = F(t) = K(t)$, the rate of drilling and capacity would continue to grow. But an escape is feasible while $S(t) > 0$. At \tilde{t} , we have $\ddot{a}(t) < 0$ per equation (28), such that $\dot{a} < 0$ immediately after \tilde{t} , and moreover we will have $\dot{a}(t) < 0$ for as long as storage continues.⁶⁵ Thus, $\dot{a}(t)$ may transition from strictly positive to strictly negative while $\dot{K}(t) \geq 0$.

To prove part (2), that the drilling rate and capacity must be strictly decreasing when storage ends, recall that $C(t)$ must be strictly decreasing during storage and that, to avoid a discontinuity in $C(t)$ or $F(t)$, we must have that $K(t)$ approaches $C(t_b) < C(t_a)$ as storage ends. Thus, it must be that $K(t)$ is eventually strictly decreasing before storage ends. Finally, since once $\dot{a}(t) < 0$, we must continue to have $\dot{a}(t) < 0$ so long as $S(t) > 0$, we will also have that once $\dot{K}(t) < 0$, we will continue to have $\dot{K}(t) < 0$ so long as $S(t) > 0$.

To prove part (3), note that we cannot have $\dot{a}(t) \geq 0$ forever, since when reserves are exhausted such a path would cause a discontinuity in $a(t)$. Thus, we must eventually have $\dot{a}(t) < 0$, and we will then have $\dot{a}(t) < 0$ forever per the above arguments. Lemma 9 implies that we must eventually have $\dot{K}(t) < 0$ while $\dot{a}(t) < 0$, and once this situation occurs we must have $\dot{K}(t) < 0$ forever, completing the proof of the modified lemmas.

Lemmas 12 and 13 likewise no longer hold. Instead, the two modified lemmas from above apply, but with capacity $K(t)$ replaced by production $F(t)$. To see this, first note that $\dot{F}(t) \geq 0$ implies that production is constrained, for if production were unconstrained, then there would be no storage, and $C(t) = F(t)$ would be falling over time. With constrained production, we are in the situation of the previous lemma, with $\dot{a}(t) \geq 0$ and $\dot{K}(t) \geq 0$. As shown above, production will be constrained throughout the entire interval $[\min(t_a, t_1), t_b]$. Thus, the above results all carry over from $K(t)$ to $F(t)$.

⁶⁵Time differentiating equation (28) and assuming that $S(t) > 0$ yields: $\ddot{a}(t) = -[X\dot{a}(t)/d'(a(t))^2][(r + \lambda)d(a(t))/X - P(F(t)) + \lambda\gamma(t)/X] + [X/d'(a(t))][(r + \lambda)d'(a(t))\dot{a}(t)/X - (r + m)P(F(t)) + r\lambda\gamma(t)/X]$. Evaluating at $\dot{a}(\tilde{t}) = 0$ at time \tilde{t} then yields: $\ddot{a}(\tilde{t}) = [X/d'(a(\tilde{t}))][-(r + m)P(F(\tilde{t})) + r\lambda\gamma(\tilde{t})/X]$. This value must be strictly negative because, to have $\dot{a}(\tilde{t}) = 0$ while $a(\tilde{t}) > 0$, equation (28) implies that we must have $P(F(\tilde{t})) > \lambda\gamma(\tilde{t})/X$. Moreover, as long as storage continues, we still have $\dot{a}(t) < 0$ for $t > \tilde{t}$, since the $(r + \lambda)d(a(t))/X$ term in equation (28) is positive and falling over time, while the negative $P(F(\tilde{t}))$ term (growing at rate $r + m$) will continue to outweigh the positive $\gamma(\tilde{t})$ term (growing at rate r).

The statements of **lemmas 14 and 15** are modified to allow for the possibility that these variables may increase after strictly decreasing. However, at the moment \hat{t} in which either $\dot{K}(\hat{t}) = 0$ or $\dot{F}(\hat{t}) = 0$, it must be that $\dot{a}(\hat{t}) \geq 0$, for if $\dot{a}(\hat{t}) < 0$, then the logic of the original lemma 14 (starting at “Third, we cannot have ...”) implies a contradiction (note if $\dot{a}(\hat{t}) < 0$ that production must be constrained). Then, per the preceding modified lemmas, a period of storage would eventually ensue, with production constrained starting no later than the moment \hat{t} at which capacity or production reverse course and persisting at least through the end of the storage period. Should storage end, it would end with $\dot{K}(t) < 0$ and $\dot{a}(t) < 0$. Alternatively, should storage last indefinitely, then eventually $\dot{a}(t) < 0$ and $\dot{K}(t) < 0$ for all time, per the modified lemmas.

Lemma 16 goes through with a minor modification: the value of $\theta(t)$ must be, at minimum, the discounted revenue stream that would be generated by producing and consuming from the capacity at the constraint forever.

Lemma 17 is as stated and proven without any changes.

Lemma 18 is as stated and proven without any changes.

The statements of **Lemmas 19 and 20** are modified to allow for the possibility that drilling may initially be increasing, but that in this case production must be constrained immediately and throughout an interval that eventually must include a period of storage and that subsequently must conclude with drilling and capacity strictly decreasing, per the arguments in the modified lemmas above. In addition, we require the slight modification that the choice variables and co-state variables are continuous with respect to S in addition to K and R . Finally, it is important to note that, with storage, these lemmas rely on the assumption that $S_0 = 0$.

Lemma 21 goes through, with two minor modifications. First, note in the first sentence of the proof that if $\dot{K}(t^*) \geq 0$, it is now possible to have $\dot{a}(t^*) \geq 0$. However, production must then be constrained at t^* , per the arguments in the modified lemmas above. Second, for the remainder of the proof, note that there cannot be any storage in the proposed contradiction.

Lemma 22 also requires minor modification. The same three cases must be considered: (1) $\dot{K}(\hat{t}) < 0$ and $\lambda\eta(F) < r$ for all $F < K(\hat{t})$, (2) $\dot{K}(\hat{t}) < 0$ and $\lambda\eta(F) \geq r$ for some $F < K(\hat{t})$, and (3) $\dot{K}(\hat{t}) \geq 0$.

Case (1) is considered in the first paragraph of the proof. Note that if K were to subsequently increase, then we must have, per the modified lemmas above, that production is constrained at all times t for which $\dot{K}(t) \geq 0$. Moreover, this constrained production would persist through an interval of storage, during which we must eventually have $\dot{a}(t) < 0$ and $\dot{K}(t) < 0$. If the period of storage were to continue indefinitely, then we must have binding production forever. If the period of storage were to end with capacity below the critical threshold of lemma 21, then we would return to the conditions that define case (1). If the period of storage were to end with capacity above the critical threshold of lemma 21, then we move to case (2) below.

Case (2) is considered in paragraphs two (that begins “Proceed by contradiction”) through seven (that begins “The above facts generate a contradiction”). Note that there cannot be any storage in the proposed contradiction. Thus, production below the constraint is never optimal in this case.

Case (3) is considered in paragraph eight (that begins “We now consider the case in which $\dot{K} \geq 0$ ”). With storage, it is possible that $\dot{a}(t)$ becomes positive while $\dot{K}(t) \geq 0$. However,

per the modified lemmas above, this eventually leads to a scenario in which $\dot{K}(t) < 0$, with production constrained throughout. Thus, in the end we are back to either case (1) or case (2), which together imply binding production forever.

The proof of **lemma 23** requires minor modification. With K_0 sufficiently small, drilling must begin instantly but no longer need be initially strictly decreasing over time. However, the revised lemmas 19 and 20 imply that production must nonetheless be constrained immediately. Rule (7) then implies that all subsequent production must be constrained as well.

This completes the extension of rules (1)-(2), (4)-(5), and (7)-(8) to the case of costly storage.

G Details for the computational model

Section 4.2 introduces computationally solved drilling, extraction, and price paths for a specific case of the model in equilibrium. This appendix discusses the details of the computation procedure.

We obtain an approximate computational solution to the model by optimizing the social planner’s welfare function using value function iteration. The model has two state variables, R and K . Because the demand and drilling cost functions we specify satisfy the sufficient conditions of theorems 1 and 2, we have that $F = K$ so that there is only one choice variable, a .

A fairly dense state space is required to yield smooth-looking simulations of drilling, extraction, and prices. The state space for R consists of 600 steps, while that for F consists of 300 steps (where the maximum value for F is given by $X * R_{max}/4$). The action space for a consists of 600 steps (where the maximum value for a is given by $R_{max}/40$). Because numerical precision is particularly important when a , R , and F are small, we compress the discrete actions and states at the lower end of their support by making the step lengths equal in the square root of each state variable. For both solving and simulating the model, we use a time step of three months.

In the value iteration loop, for each iteration we calculate for each discrete state s the optimal discrete action a . Because each possible action is very unlikely to cause a discrete state to be hit exactly in the next period, we obtain the value for next period’s state via linear interpolation. For value function convergence, we use a tolerance of 10^{-7} (with the value function measured in units of \$million).

The solution to the model yields a policy function for the rate of drilling at all possible discrete states. To conduct the forward simulation, we linearly interpolate this policy function whenever the simulation requires the optimal drilling rate for a state that does not exactly match one of the discrete states.