Taxes and U.S. Oil Production: Evidence From California and the Windfall Profit Tax *

Nirupama S. Rao†

Abstract The widespread boom in U.S. oil production has prompted state debates on levying new taxes on oil. This paper uses new well-level production data and price variation from federal oil taxes and price controls to assess how taxes affect production. Empirical estimates suggest an after-tax price elasticity ranging between 0.295 (0.038) and 0.336 (0.042). Response along the extensive margin is minimal. There is no discernible evidence of spatial shifting of production to minimize tax liabilities. Taken together the results suggest that taxes reduced domestic production in the 1980s, and the response largely came from wells that continued to pump oil, albeit at a reduced rate.

Keywords: Taxes, Energy, Oil, Supply Elasticity.

JEL Classification Numbers: H23, H25, Q41, Q48.

---

*I thank Michael Greenstone, Jon Gruber, and especially Jim Poterba for valuable advice and encouragement. I also thank Maria Damon, Mike Golosov, Cynthia Kinnan, Matt Notowidigdo, John Parsons, Bob Pindyck, Mar Reguant-Rido, Steve Salant, Wolfram Schlenker, and seminar participants at MIT, NYU, Resources for the Future, University of Illinois, London School of Economics, Columbia University, University of Colorado Boulder, the Ford School of Public Policy, and the 2010 NTA annual meeting for helpful comments.

†nirupama.rao@nyu.edu, Robert F. Wagner School of Public Service, New York University. 295 Lafayette St New York, NY.

Introduction

The recent boom in U.S. oil production has been remarkable. Harnessing new drilling techniques and technologies to extract oil from rock formations formerly deemed too costly, the U.S. oil industry has effectively reversed two decades of decline in the last five years.[1] Last year the U.S. pumped 7.4 million barrels of oil daily—more than any point since 1989. More crude is now pumped in more states than ever before, raising questions of how states should tax this production.

The spread and scope of the boom has drawn the attention of policymakers at the federal level and in a wide array of states. Oklahoma doubled its tax rate on new wells in May 2014 while Governor John Kasich of Ohio has proposed raising the oil tax to 2.75 percent from 2.25 percent and California is debating imposing a severance tax on oil production for the first time. Advocates for higher taxes in these states and others point to North Dakota's 11.5 percent severance tax on oil production and accompanying billion-dollar surplus.[2] Interestingly, Alaska took the opposite tack in 2013, reversing course and repealing tax provisions that increased severance tax rates as oil prices rose. The recent policy reversal in Alaska was motivated by declining state production; the lower tax rates are aimed at encouraging production in Alaska.

In these states and others mulling altering their tax treatment of oil production, the impact of these taxes on production is a key consideration. The degree to which imposing higher taxes will alter producer behavior determines both the deadweight loss of such taxes and how effectively they will raise revenue. While drilling and casing a well may be an irreversible investment, production itself is often costly, meaning that taxes could potentially reduce production. The economic impact of oil taxes hinges critically on how producers respond to changes in after-tax price. If supply is inelastic, then these

[1]While domestic oil production fell every year between 1991 and 2008—a cumulative decline of a 33 percent—since 2008, daily production has increased by nearly 50 percent.

[2]Ohio's Governor Kasich argues "An increase—a modest increase—in the severance tax on Big Oil will allow us to reap some of the benefit of that oil, which they're pulling out of the ground and is a diminishing resource."

states could raise substantial revenue without triggering deadweight loss; if producers are sensitive to after-tax price, these taxes will be far less efficient. This paper examines the impact of an earlier federal excise tax to understand how taxes affect U.S. production.

Despite the importance of estimates of the elasticity of U.S. supply for assessing the impact of policy changes, consensus elasticity estimates have been lacking. Key prior studies are summarized in Table 1. Previous studies have relied exclusively on aggregate time-series variation and have mostly found very small and economically insignificant elasticities. Hogan (1989) and Ramcharran (2002) found statistically significant but economically minor supply elasticities of 0.09 (0.03) and 0.05 (0.02), respectively. Jones (1990) and Dahl and Yücel (1991) found insignificant elasticities of 0.07 (0.04) and -0.08 (0.06), and Griffin (1985) found a significant negative elasticity, -0.05 (0.02). Hogan (1989) also estimated a longer-run elasticity of 0.58 (0.18). More recently Anderson, Kellogg, and Salant (2014) used Texas lease-level data from 1990 to 2007 to examine the responsiveness of production to contemporaneous West Texas Intermediate (WTI) prices and future prices. They found that production was not sensitive to either spot or expected future prices.

Most policy studies regarding oil markets rely on a range of plausible elasticities due to the lack of consistent credible estimates. In fact, the 2006 Congressional Research Service (CRS) report on proposed windfall profit taxes stated, "few studies generate reliable estimates and in fact some studies estimate negative supply elasticities, which are not plausible."[3] Thus the CRS report, like other recent studies by federal agencies and the Organisation for Economic Co-operation and Development (OECD), employed a number of assumed elasticities—CRS used supply elasticities of 0.2, 0.5 and 0.8—rather than settling on a specific elasticity estimate.[4]

[3] Lazzari (2006)

[4] The OECD, in its 2004 Economic Outlook, based its projection of production by countries that are not members of the Organization of the Petroleum Exporting Countries on elasticities of 0.1, 0.3, and 0.5. The U.S. Department of Energy's Energy Information Agency does not explicitly state the elasticities it uses in its analyses, but its forecasts indicate that it used an elasticity of 0.2 over a ten-year window and virtually zero for one-year responses.

This paper estimates the responsiveness of domestic oil producers using a new rich dataset that reports monthly production for all onshore wells in the state of California—the third-ranking state in oil production—over a 31 year-year period beginning in 1977. The empirical strategy makes use of the variation in after-tax price induced by the end of price controls and the 1980 Windfall Profit Tax (WPT)[5], which levied substantial and varying excise taxes on U.S. wells—marginal tax rates ranged from 22.5 to 70 percent. The constructed dataset of 30,025,957 observations describing 140,672 wells includes wells that were already completed and wells completed during the period. In addition to monthly production, the data report monthly values, for each well, describing the quality of oil produced, the firm operating the well, the method of pumping, exact location, the field and pool it taps, and whether it is capable of producing or is shut-in. This level of detail is necessary for determining each well's correct regulatory and tax treatment, as prescribed by the *Code of Federal Regulations* for each year. Using this policy detail and monthly field-by-grade prices from Platt's *Oil Price Handbook and Oilmanac* for each year, the path of after-tax prices for each well is accurately traced over time, taking into account differential regulatory and tax treatment across wells.

This new data allows for two improvements over previous studies. First, well-level data allow for better controls for time-varying factors (such as changing price expectations) and well heterogeneity. Because the federal policies created substantial variation in after-tax price across wells and over time, the supply response is identified here using only within-well variation.[6] In fact, regulatory and tax policy generate enough across-well variation in after-tax price in each month-year that non-parametric controls for common unobserved time factors affecting well productivity can be employed. Second, because the data

[5]The 1980 Windfall Profit Tax was in fact an excise tax, not a tax based on profits. despite its name.

[6]As the analysis examines the within-well supply response, the exploration margin is not a part of the assessment. As new wells are completed they are added to the sample used to generate the empirical estimates, but since the analysis uses only within-well variation in after-tax price, the estimate does not measure the impact of new wells on aggregate production. For a theoretical investigation of the effect of taxes on exploration and development see Smith (2012).

track wells individually, the estimates can separate the extensive from the intensive response. Distinguishing between these margins is important; if the supply response is driven by the shutting-in of wells, the high cost of reversing shut-in makes this a potentially permanent loss of oil.

My estimates make clear that production from existing wells is price-responsive—much more so than previous estimates of the same era suggest. The main results show an after-tax price elasticity of oil production in California ranging between 0.295 (0.038) and 0.336 (0.042). Response along the extensive margin is minimal; a ten percent decrease in after-tax price would lead to at most a 1.11 percent increase in the shut-in rate. Like many currently proposed taxes, the WPT subjected different types of wells to different tax rates, creating an opportunity for producers to strategically shift production from tax-disadvantaged wells to tax-advantaged wells without changing or minimally changing total production. The empirical analysis does not find significant evidence of such strategic spatial shifting. Taken together the results suggest that the WPT did in fact reduce domestic production, and most of the response came from wells that continued to pump oil, albeit at a reduced rate.

While the estimates here are the first oil supply responses to be identified using plausibly exogenous variation, there are several factors to consider in applying these estimates to the current policy context. First the tax-based variation in after-tax price used here to identify the supply response dates back to the 1980s. Given technology changes and the effect of time on production, these estimates may not be entirely applicable today. The similarity of responses of wells of different ages does suggest that simple aging may not be a crucial factor. Second, the variation in after-tax price exploited here arises from a legislatively temporary excise tax—the WPT was slated to last only 11 years. Most oil excise taxes under consideration today would be legislated as permanent taxes. As temporary taxes incentivize retiming production while permanent taxes on exhaustible resources changes the opportunity cost of extraction identically for all future periods, it is reasonable to consider the estimates here an upper bound on the reaction of producers to permanent taxes. The empirical results can nonetheless help inform current policy considerations. Much like the WPT many state oil tax regimes attempt to tax different types of wells at differential rates, opening the potential for strategic spatial shifting. The estimates here suggest that this type of reallocation is not significant. Finally, though states may legislate taxes permanently, the experience of Alaska and current attempts in Oklahoma to reverse the recent tax increase suggest that oil taxes may be subject to considerable policy uncertainty, rendering even legislatively permanent tax change potentially temporary in the minds of producers. If current taxes are expected to be reversed with some probability, then the estimates reported here may be more applicable. Taken together the much higher elasticities estimated here relative to previous time-series estimates from the same era suggest that policymakers considering higher taxes on oil production should expect those taxes to slow production.

The paper proceeds as follows. The next section provides the relevant institutional knowledge regarding the decontrol of oil prices and the introduction of the WPT. Section 2 details the conceptual framework. Section 3 describes the new rich production and price data I assembled and details the estimation strategy. Section 4 presents the empirical results. Section 5 concludes and discusses the relevancy for oil tax policy.

1 Background: Decontrol and the 1980 Windfall Profit Tax

The estimation strategy makes use of price changes driven by price regulation, decontrol, and the imposition of federal excise taxes. These policies significantly altered producer prices and created considerable differences in producer price across wells. These policies are detailed below.

The decontrol of oil prices began in 1976 with marginally productive wells called stripper wells.[7] Rising prices and less stable foreign sources prompted concerns regarding U.S. oil independence and generated interest in increas-

[7] Vietor (1984) provides excellent detail on the timing and nature of these policy changes.

ing domestic oil production. The Carter administration began decontrolling non-stripper domestic crude in June 1979. Decontrol went forward with the understanding that the sudden increase in domestic producer prices would be taxed at the federal level.[8] The 1980 Windfall Profit Tax (WPT) was signed into law April 2, 1980, and virtually all non-Alaskan oil owned by a taxable private party was subject to the tax. The name is a misnomer. The WPT was not a profit tax. The WPT was an excise tax levied on revenues irrespective of costs. Purchasers withheld the tax from payments to producer and filed quarterly WPT tax returns with the Internal Revenue Service.

The timing of decontrol and WPT tax treatment varied by oil specific gravity, and by the age and productivity of the well from which oil was extracted. The WPT taxed oil that was typically more costly to extract and refine at a lower tax rate. American Petroleum Institute (API) gravity measures the specific gravity, or "heaviness" of oil, which determines how efficiently the crude can be refined into petroleum products. Heavier oil sells for lower prices.[9] Tax-favored oil included heavy oil that had an API gravity of 16 or less, and oil from stripper wells, which produce, on average, less than 10 barrels of oil per day for at least 12 months.

All taxable oil was divided into three tiers under the WPT; each tier corresponded to a different tax rate.[10] An operator's WPT tax liability was equal to the product of the WPT tax rate and the difference between the selling price and a tier-specific base price for each barrel of oil he sold. WPT payments were deductible from corporate taxable income, meaning that the after-tax price (ATP_{it}) received by the operator of well i at time t was:

$$\mathrm{ATP}_{it} = \begin{cases} \left(1-\tau_t^{Corp}\right)\left(P_{it}-\tau_{it}^{W}\left(P_{it}-B_{it}\right)\right) & \text{if } P_{it} > B_{it} \\ \left(1-\tau_t^{Corp}\right)P_{it} & \text{otherwise} \end{cases}$$

[8]According the Joint Committee on Taxation's General Explanation of the Crude Oil Windfall Profit Tax of 1980, "without such a tax, decontrol probably could not [have gone] forward." For more detail on the decontrol and levying of the WPT see Kalt (1981).

[9]API gravity is an inverse function of specific gravity:
$\text{API Gravity} = \frac{141.5}{\text{Specific Gravity}} - 131.5$

[10]Specific categories of oil, largely state-, Native American-, or charitable trust-owned oil, were exempt from the WPT. See Lazzari (2006) for further details.

where τ_t^{Corp} is the prevailing corporate tax rate, P_{it} is the real selling price, τ_{it}^W is the WPT rate, and B_{it} is the real base price for oil pumped from well i at time t. Note that real selling prices were common across all wells that produced oil of the same quality, while base prices and WPT tax rates potentially varied among wells producing the same grade of oil.

The WPT was legislated as a temporary tax. At its height, the WPT raised $44 billion in gross revenue (before corporate income tax deductibility), or roughly half the revenue raised by the corporate income tax. Statute required the tax to expire by 1991. In reality the tax became ineffective due to sharp decreases in oil prices in 1986; 1985 was the last year it raised any revenue. In fact, the WPT was repealed in 1988 to eliminate the administrative burden of a tax that did not raise revenue. The timing of decontrol and the simplified details of WPT treatment for each of the three tiers of oil follow and Figure 1 presents a timeline. [11]

Tier I Oil

Tier I oil was non-heavy oil extracted from a non-stripper well that produced oil in 1978. Tier I oil was subject to price controls through 1979. Price controls on Tier I oil were initially phased out gradually.[12] At the end of January 1981, the phase-out of price controls was abruptly ended and Tier I oil was fully decontrolled. The base price for Tier I oil was 21 cents less than the May 1979 price control price for the property. The tax rate on Tier I oil was 70 percent.

Tier II Oil

Tier II oil consisted of non-heavy oil from stripper wells that produced oil in 1978, and oil produced from a Naval Petroleum Reserve (NPR) field. A well is considered a stripper well if it has ever averaged less than 10 barrels of oil per day for 12 consecutive months after 1972. Oil produced from stripper wells was exempted from price controls in August 1976. An NPR field is one of four fields owned by the federal government to which access is leased to private

[11] For further detail see Joint Committee on Taxation (1981).

[12] Beginning in January of 1980, the selling price was a weighted average of the world market price and the price control price with the weight on the market price equal to 0.046 multiplied by the number of months since December 1979.

operators. The base price for Tier II oil was the December 1979 selling price of oil from the same property multiplied by 0.425, a conversion factor that achieved a statutorily set average base price of \$15.20. The tax rate on Tier II oil was 60 percent.

Tier III Oil

Tier III oil was composed of two types of oil, new oil from wells that did not produce oil in 1978 and heavy oil with an API gravity of 16 or less. New oil was fully decontrolled in June 1979. Price controls on heavy oil were lifted August 17, 1979. The base price for both new and heavy oil was the December 1979 selling price of oil from the same property multiplied by 0.462. Heavy and new oil were the most tax-favored types of oil; the tax rate on Tier III oil was 30 percent initially and was gradually reduced to 22.5 percent beginning in 1982.

The three tiers of oil, and even different categories of oil within Tier III, were treated very differently by government policies. Differences in the timing of decontrol and differential tax treatment provide the variation in after-tax price that generates the supply elasticities estimated here. These policies created cross-sectional variation in after-tax price allowing for flexible controls for underlying common time-varying factors, including future price expectations.

2 Conceptual Framework

2.1 Tax Incidence

The incidence of the WPT fell on U.S. oil producers. As they account for a small share of world production and operate in a market alongside a cartel, U.S. oil producers—including the California producers examined here—can reasonably be assumed to be price-takers.[13] Furthermore, refiners can freely purchase imported oil—which was exempt from the WPT. The availability

[13] Kilian (2009) asserts "the price of crude oil is determined in global markets." As in other empirical studies, such as Smith, Bradley, and Jarrell (1986), here domestic pre-tax prices are assumed to track world prices.

of tax-exempt imports fixed the refiner price at the world price; thus, U.S. producer prices were reduced by the full amount of the tax.[14]

2.2 Incentives to Shift Production

As first shown by Hotelling (1931), exhaustibility makes oil extraction a "pump today or pump tomorrow" decision.[15] Dasgupta and Heal (1979) among others have shown that a permanent, constant *ad valorem* tax does not create incentives to temporally shift production since the current and future opportunity cost of extraction are reduced identically, but can lead the well to ultimately produce less oil as production will now cease when the marginal cost exceeds the lower after-tax price. The WPT was a temporary *ad valorem* excise tax. It reduced the return to extraction in the near term, but left expected long-term returns for wells that outlive the tax unchanged, creating strong incentives to temporally shift production. Any currently considered taxes that are viewed as ultimately temporary will have similar incentives. If prices were constant, a producer who reacts to a temporary tax by delaying production loses the time-value of money but avoids paying the temporary tax. Price uncertainty, of course, complicates matters.[16] Different forms of price uncertainty have differing implications (see Weinstein and Zeckhauser (1975) and Lewis (1977)) but in general greater uncertainty, all else equal, enhances the option value of holding the reserves *in situ*.[17]

The temporary nature of the WPT meant that operators had much stronger

[14] Though transportation costs are small, roughly 5 percent of oil prices, domestic producers may have been able to pass a fraction of the tax, equal to the transport cost, on to purchasers. All oil produced in California is refined within the state, but refiner demand exceeds production so imports comprise the difference. Imports come largely from Canada and Mexico and average transport costs run roughly $1.30 per barrel according to Rodrigue, Comtois, and Slack (2013).

[15] Robert M. Solow's 1974 Richard T. Ely lecture (Solow 1974) provides an insightful (and humorous) summary of Hotelling's and subsequent work and their implications.

[16] Investigation of the statistical behavior of crude prices have found differing patterns with work by Pinkyck (1999) finding strong evidence of mean-reversion while more recent work by Hamilton (2009) concluded that crude prices follow a random-walk without drift.

[17] Kellogg (2014) shows that uncertainty also affects drilling decisions. In particular, when the expected volatility of futures prices increases, drilling activity decreases by a magnitude consistent with the real options model.

incentives to reduce production in reaction to contemporaneous taxes than they would under a permanent excise tax. Thus, the elasticities estimated here can be considered an upper-bound on the response of supply to permanently levied taxes.

The WPT, nonetheless, offers a unique opportunity to estimate the impact of taxes on oil production. Generally, all wells producing similar crude face a common price and are typically subject to a single tax rate. Thus regressions of production on after-tax price preclude flexible time controls and will yield coefficients biased by time-varying factors, such as the confounding effect of changing price expectations. The WPT taxed otherwise similar wells very differently, allowing for time fixed effects that net out the effect of evolving price expectations. As long as producers held common expectations for future price—which given the economic importance of the question and research budgets of oil producers seems reasonable—the cross-sectional tax variation of the WPT should afford the opportunity to isolate the impact of contemporaneous after-tax prices on production.

The spatial layout of oil extraction informs the empirics as well. Most operators do not pump their reserves with a single well. Instead a leaseholder typically taps his reserves using multiple wells. Wells on the same lease operated by the same firm can be subject to differential WPT treatment due to differences in well and oil characteristics. If production from well i on lease l at time t is denoted with q_{ilt}, then well-level regressions of production on after-tax price will yield the mean response of well production to contemporaneous price, $\frac{dq_{ilt}}{dp_{ilt}}$. The production of wells on the same lease, however, may be related; namely, a leaseholder could strategically reallocate production from high- to low-tax wells, leaving total lease production unaffected, or less affected, while minimizing tax liabilities. This type of shifting will lead to lease-level production responses that are smaller than suggested by well-by-well regressions, that is $\frac{dq_{ilt}}{dp_{ilt}}$ would exceed average lease production, $\frac{d\overline{q_{lt}}}{d\overline{p_{lt}}}$. If there is no spatial shifting, the supply elasticity estimated from well-level regressions of well production on after-tax price should be the equivalent to the response of average production of all the wells on a lease to average after-tax price, weighted by production.

Because wells simultaneously produce oil, gas and water, it is typically not possible to meter or measure production for each well. Instead, all of the wells on a lease are flowed to a separating facility, where production is metered at the lease-level. This aggregate lease production is then allocated back to the lease's wells based on periodic "well tests", in which each well is flowed into a small test separator to measure its production rate.[18] The lack of well-level metering raises the potential for leaseholders to allocate production across wells to minimize tax burdens without actually shifting real production to tax-advantaged wells.

The empirical investigation aims to estimate the impact of after-tax price on quantity, $\frac{dq_{ilt}}{dp_{ilt}}$, using using new well-level data and WPT tax variation, where the tax fully falls on producers. The empirics also assess the scope of spatial shifting since easy production shifting makes varying tax rates—often and currently used to encourage production from new or marginal wells—a potentially costly tax provision that may reduce revenue and encourage avoidance at best and noncompliance at worst.

3 Data and Empirical Design

The decontrol of oil prices and the introduction of federal excise taxes created substantial variation in after-tax price over time and across wells. These policies classified wells into different regulatory and tax tiers by the characteristics of the well and the oil it produced. Thus well-level data are necessary to account for and make use of this substantial variation. Wells within a field could be assigned very different after-tax producer prices depending on whether they produce the same kind of oil, share the same stripper status, or produced in 1978. To use this well-level variation, I assembled a new database of well-level production and after-tax producer prices that describes every onshore well in California starting in 1977, which encompasses the regulatory and tax periods.

[18]For more information regarding the oil production process please consult the thorough and accessible non-technical volume by Hyne (2001)

These data have not been used in previous studies.

3.1 Data

California remains the third highest oil producing state in the nation despite the recent surge in production elsewhere.[19] Onshore producers account for roughly 92.5 percent of state production, with offshore wells pumping the rest.[20] California crude is generally of lower quality than more prominent benchmark crudes such as WTI. California oil was more than 60 percent heavy or very heavy crude during the 1977-1985 period. Heavy oil is generally more expensive to extract and refine.[21]

The data used in this study cover all potentially active onshore oil wells in California, beginning in 1977. The main analysis regarding the impact of price regulation and excise taxes makes use of the more than 75,000 oil wells that were capable of producing at some point during the 1977 to 1985 period. The State of California Department Conservation Division of Oil, Gas and Geothermal Resources requires operators to report monthly production and characteristics for all completed wells that are potentially capable of production. Characteristics reported each month include the date of well completion, API gravity of the oil produced, the field and pool being tapped, operator name, and the status of the well. The data are particularly well suited for the analysis since they provide monthly information that allows more precision in the timing of price and tax changes relative to the annual or quarterly data used in other studies. More importantly, the data report the characteristics necessary to determine the timing of decontrol and WPT tax treatment for each well.

Some adjustments to the data were necessary. In months where oil produc-

[19] Interestingly Alaska has fallen from second to fourth due to the North Dakota oil boom

[20] Mauritzen (2014) examined Norwegian offshore fields and found no significant evidence of a concurrent reaction of field production to oil price, but found a lagged effect of roughly a 2% to 4% production increase for a 10 dollar per barrel real price increase.

[21] Heavy oil has an API gravity less than 20; very heavy oil has an API gravity less than 16. Higher API gravity oil is lighter and sells for a premium. During the 1977-1985 period, 11.6 percent of California crude was heavy while 49.8 percent was very heavy.

tion is zero either because the well is not yet complete or is shut-in, no API gravity data are reported; I assign these well-month observations the soonest future API gravity in the case of uncompleted wells and the most recent previous API gravity in the case of shut-in wells. Stripper well status is determined by examining production history within the data, so the share of wells qualifying for stripper status would rise mechanically at the end of 1977 if only production history determined stripper status. To address this data concern, I back-fill stripper status so that a well that is determined to be a striper well in January 1978 is classified as a stripper well in 1977 as well.

All oil does not trade at a single price; different grades, in terms of API gravity and sulfur content, trade at their own prices. The price data are from Platt's *Oil Price Handbook and Oilmanac*, which provides monthly field-by-field posted prices by API gravity for controlled and decontrolled oil. Fields for which price data are not available are assigned the average price for oil of the same API gravity for wells in California that month. Because the prices of different grades do not track the world price in parallel, using the more precise prices could potentially be important. Crude is globally traded and priced based on API gravity and location. Location provides information on the sulfur content of the oil since sulfur content is largely constant across the wells in a field.[22] Oil with low sulfur content, known as "sweet" crude, can be refined into light petroleum products such as gasoline or kerosene more cost effectively than high-sulfur, "sour" crude, which is typically processed into diesel or fuel oil.[23] For refining purposes, oil of the same API gravity and sulfur content is viewed as perfectly substitutable regardless of origin.

While various congressional acts created the systems of regulation, decontrol, and excise taxation that provide the identifying variation in producer

[22] Refiners with the lowest transportation costs, typically those with the closest refineries, will purchase from a given field. As individual purchase and production decisions are too small to move transport costs, the difference between price at the wellhead and price at the refiner is taken to be independent of the decisions of individual firms.

[23] When oil prices are referred to in the popular media, the price frequently quoted is that of WTI, or UK Brent, both of which are light and sweet. The OPEC basket, which is a weighted average of crudes produced by OPEC nations, is a third benchmark and is both heavier and sourer than WTI or Brent.

prices, the precise detailed rules of these legislative acts are found in the *Code of Federal Regulations* for each year. The details of price control assignment and WPT tax treatment are drawn from "Title 10: Energy" of the *Code of Federal Regulations* for each year, 1976-1980, and "Title 26: Internal Revenue" of the *Code of Federal Regulations* for each year, 1981-1985, which detailed the implementation of price control and WPT legislation.

Table 2 presents summary statistics for the full sample of 75,342 wells used to assess the impact of the regulatory and tax regimes of the late 1970s and 1980s. The average well produces 476.4 barrels of oil per month; conditioning on non-zero production raises the average roughly 40 percent. Approximately 28 percent of well-month observations report zero oil production either because the well is shut-in or because the well has not yet been completed. The production data are right skewed. The median well produces 113 barrels of oil per month, the 75th percentile well-month observation produces 428 barrels per month, and the 99th percentile observation produces 5,325 barrels per month. The within-well production variation, 868.5, is comparable to the overall standard deviation, 1,473.8. The average producer price during the period, \$18.3, is only 45 percent of the mean purchaser's price, with part of this difference attributable to the corporate income tax and part to the WPT. Producers for whom price controls were gradually phased out as they faced excise taxes under the WPT received the lowest—less than \$12.30—after-tax prices. Producers of lighter oil received the highest prices in the sample—exceeding \$32.00—at the end of 1979 and the beginning of 1980 prior to the introduction of the WPT. The within-well deviations in average after-tax price is 15 percent smaller than the overall variation in after-tax price, while the within-well and overall variation in pre-tax price is comparable. This discrepancy is driven by the differential regulatory and tax treatment of wells over the period. The average and median API gravities are 18.2 and 15.0, respectively, illustrating the heaviness of California oil. Finally, note that although there is considerable variability in API gravity in the sample (standard deviation of 6.8), each individual well has little variation in the API gravity of the oil it produces (standard deviation of 1.1).

3.2 Estimation Strategy

The way in which oil prices were decontrolled and oil production was taxed provide an unusual degree of variation in net-of-tax prices for often identical commodities across producers and over time. The decontrol of oil prices and the introduction of the WPT were policy changes implemented in tandem; oil prices were decontrolled by executive order while legislation enacting the excise tax was in committee in Congress. Figure 1 illustrates the timing of decontrol for different types of oil over the 1979 to 1981 period, starting with new oil and ending with old oil. These different categories of oil were also subject to different WPT tax rates and corresponding tax bases. Taken together these policy changes provide substantial deviations from the world market price. This policy-induced variation in after-tax price identifies the supply response estimated here.

The policy-driven incentives to shift production described above in Section 2 suggest a simple empirical framework. A natural regression model that would yield estimates of $\frac{dq_{it}}{dp_{it}}$ is a simple linear model of the form:

$$q_{it} = \beta \left(1 - \tau_t^{Corp}\right) \left(B_{it} + \left(1 - \tau_t^{W}\right)\left(P_{it} - B_{it}\right)\right) + X_{it}\gamma + \chi_t + u_i + \eta_{it} \quad (1)$$

where q_{it} is extraction per month, τ_t^{Corp} is the prevailing corporate tax rate, B_{it} is the real base price, τ_t^{W} is the WPT tax rate, P_{it} is the real selling price for oil pumped at well i at time t, X_{it} is a set of controls, and $u_i + \eta_{it}$ is the error term.[24]

Estimates of β from equation 1 capture the average well response to changes in after-tax price, including both the extensive and intensive margin. For wells near the end of their economic life, the post-tax profit from remaining reserves may not offset the losses they will incur during the tax period. Thus some

[24] The after-tax price here, $\left(1 - \tau_t^{Corp}\right)\left(B_{it} + \left(1 - \tau_t^{W}\right)\left(P_{it} - B_{it}\right)\right)$, captures the producer price under both price controls and the WPT. Producers bore the full incidence of both. Price controls and the WPT can both be described as taxes on a price basis, where the basis is the difference between the selling price of a barrel of oil and a statutory base price. In the case of price controls, the tax rate is 100 percent.

well operators may choose to exit by shutting-in their wells. In fact, there was notable concern regarding response along this extensive margin at the time the tax was introduced.[25] To assess the degree of extensive response, we want to estimate the impact of variation in after-tax price on the decision to shut-in a well. That is, a model of the form:

$$S_{it} = \delta\left(1-\tau_t^{Corp}\right)\left(B_{it}+\left(1-\tau_t^{W}\right)\left(P_{it}-B_{it}\right)\right)+X_{it}\gamma+\chi_t+u_i+\eta_{it} \quad (2)$$

where S_{it} is a dummy variable equal to one if the well is shut-in, and the regressors are as described above. The coefficient on after-tax price, δ, measures the percentage change in the probability of shut-in caused by a one-dollar after-tax price increase.

If the price ceilings and WPT tax rates were uncorrelated with the error term, the policy-based variation in after-tax price would yield an unbiased estimate of the tax response for both equations 1 and 2. But if after-tax price is correlated with an underlying well-specific component of the error term, u_i, then pooled ordinary least-squares estimation will yield biased estimates. The bias of the estimate will depend on the correlation between the omitted well-specific effect and the tax rate or price ceiling. Price ceilings and excise tax rates were not randomly assigned to wells by price controls and the WPT. Well characteristics (e.g. well age and stripper status) and oil characteristics (i.e. specific gravity), which can be factors in the cost of extraction, were used to determine regulatory and tax treatment. Regulatory and tax treatment varied along these dimensions, in part in an effort to favorably treat operators who would be most adversely impacted by the tax. Thus, pooled ordinary least squares (OLS) estimates of equation 1 or 2 would be inappropriate.

Because extraction costs vary across wells even within tier, controls for the factors that determine tax treatment may not be sufficient to fully address heterogeneity in extraction costs. Instead, to isolate variation in the after-tax

[25]For example, two months before the enactment of the tax, the *Wall Street Journal* ran a critical editorial about the proposed WPT titled "The Close-the-Wells Tax" Editorial (January 22, 1980)

price not related to underlying differences in extraction costs, the analysis uses only within-well variation. Because of the considerable across time variation in after-tax price generated by the decontrol of oil prices and the levying of the WPT, there remains sufficient variation for each well over time to identify the supply response.

3.2.1 Residual Variation in After-Tax Price

Figure 3 plots different price measures for two wells. The real posted price line reports the real purchase price of the oil. The upper plot describes a relatively tax-disadvantaged well, and the lower plot describes a relatively tax-favored well.

The upper plot tracks an initially non-stripper well that was decontrolled gradually beginning in January 1980, then fully decontrolled in January 1981. The gradual decontrol can be seen in the nearly linear upward slope of the Real Posted Price line starting in January 1980 and continuing until January 1981, when the price discontinuously jumps with full decontrol. This well was initially subject to a 70 percent WPT excise tax. The onset of the tax is the sudden downward jump in the After-Tax Price in March 1980. In October 1982, the well qualified as a stripper well and thus shifted to the slightly more tax-favored Tier II and became subject to a 60 percent excise tax rate; hence the uptick in the After-Tax Price. The decrease in posted price in January 1983 led to decreases in all price measures.

The estimation strategy removes well and time fixed effects. Purging the after-tax price measure of well fixed effects amounts to subtracting the well's average price over all periods from the price each period. Thus the Residual–Well FE line is the After-Tax Price line shifted downward by the well mean price. Further purging the post-well fixed effect residuals of time fixed effects amounts to then subtracting the average price each period over all wells. This two-way residual isolates relative within-well price variation, where relative means relative to all other wells in the sample that period. Thus, this well's two-way residual declines beginning in June 1979 as Tier III oil is fully decontrolled and market oil prices rise. The Residual–Well, Time FE line slopes upward between January 1980 and March of 1980 as the well began gradual decontrol, while already decontrolled wells faced less rapidly increasing prices. When the WPT is levied in March 1980, the two-way residual continues its upward trend because the increases in after-tax price due to continued decontrol more than offset the tax. Even after full decontrol in January 1981, the relative within-well after-tax price remains negative because this well faces the highest tax rate of all wells. The disadvantage narrows as posted prices in the Livermore field increased relatively faster than other fields. When the well is reclassified as a stripper well, there is a final uptick in the two-way residual as its WPT tax rate has fallen by 10 percentage points, which is short-lived as the Livermore price premium fades a few months later. From that point on, the two-way residual is near zero since declines in the posted price result in after-tax prices nearly equal to the average after-tax price for each well.

The lower plot tracks a relatively tax-favored well. The well did not produce oil in 1978 and is classified as a new well. The After-Tax Price line jumps upward in June 1979 when new oil was decontrolled and again several months later as posted prices reflected higher world prices. This Tier III well was initially subject to a 30 percent WPT tax rate, which was decreased by 2.5 percentage points each year starting in 1982 until the rate was 22.5 percent in 1984. Focusing on the two-way residual line, Residual–Well, Time FE, the fact that this well was tax-advantaged can be seen at several points. First, when this well was decontrolled in June 1979, the two-way residual is large and positive. The strong upward movement of posted prices beginning in 1980 is mitigated in the two-way residual since other wells were beginning decontrol and receiving higher after-tax prices during this time—the residuals do, however, remain above zero since this well was fully decontrolled. The residuals remains positive even after the introduction of the WPT because it was tax-favored.

Price variation generated by taxes is likely to be perceived as having a persistence that differs from that generated by movements in price. If producers perceive price changes as having greater persistence than tax-driven changes, then supply elasticities generated by price changes would overstate the supply response to temporary taxes. Thus within-well variation in after-tax price, which retains both price- and tax-driven changes in after-tax price may not be the appropriate price measure for the analysis. To isolate price differences due only to differential decontrol and tax treatment, the data are purged of time-series variation in price. The plot for each well tracks this process of isolating relative within-well variation in after-tax price.

The key exclusion restriction of an identification strategy that purges after-tax prices of well and time averages is that, outside a time-invariant fixed factor, wells respond identically over time to changes in relative after-tax price. In other words, there are no time-varying well-specific factors, besides after-tax price, affecting well production.

4 Results

4.1 Main Results

Table 3 presents OLS estimates of equation 1 using the full sample of California oil wells. The dependent variable is quantity of oil produced by well i in month t. All specifications include well-level fixed effects to absorb level differences across wells in operator responses to changes in after-tax price—namely, production cost heterogeneity. The sample includes all wells, whether or not they are shut-in. Month-by-year dummies absorb mean production and price variation each month. The tax-price elasticity is identified by within-well variation in after-tax price relative to the within-well variation of other wells. As wells age, their productivity may decline, so additional controls for the age of the well, measured from its date of completion, are also included.

Column 1 of Table 3 reports results from a model employing only time and well fixed effects. The estimated coefficient on the after-tax price term, β, implies that a one-dollar increase in the after-tax price leads the average well to produce 8.73 additional barrels of oil, a price elasticity of 0.335. Because well age is considered an important determinant of well productivity, column 2 adds a quadratic function of well age. The insignificant increase in the elasticity to 0.336 and the unchanged precision suggest that the well fixed effect provides sufficient controls and age does not matter very much within a well. Although over the course of a well's life there is little change in the API gravity of the oil extracted—the within-well standard deviation is only 1.1 degrees, less than 20 percent of the overall variation—changes in API gravity could lead to changes in lifting costs if the changes are concentrated and thus large for wells that do experience changing API gravity. Column 3 employs dummies and quadratic time trends for each decile of API gravity. The after-tax price coefficient is reduced by these added time-varying controls for oil quality, but the change, a reduction of the elasticity to 0.295, is statistically insignificant and economically minor.

Although the vast majority of wells in California are pumped in any given month, 33,198 wells produce oil based on their natural subsurface reservoir pressure for at least part of their lives. These flowing wells have low operating costs if they produce their natural flowing quantity, but it is very costly to adjust their production either upward or downward. Adjustment involves the installation of pumping equipment to either increase subsurface pressure to accelerate extraction or to exert downward pressure to reduce the flow rate. In other words, very high costs of extraction rate adjustment make the operators of flowing wells unlikely to adjust their production levels in response to temporary changes in after-tax price. Columns 4 and 5 of Table 3 present estimates of equation 1 separately for pumped and flowing wells, respectively.[26] Pumped wells—those for which production levels are more of a choice variable—are more price responsive than the average well. A ten percent increase in after-tax price leads to a 3.71 percent increase in production.[27] The higher point estimate however is not statistically significantly different from the specification of column 3 which includes both flowing and pumped wells. Flowing wells, on the other hand, do not show a statistically significant production response to changes in after-tax price.[28] The strong response of pumped wells and non-

[26]Because some wells may initially flow but then need to be pumped, the number of wells in the flowing and pumped regressions exceeds the total number of wells.

[27]All elasticities are evaluated at average price and quantity, separately for pumped and flowing wells.

[28]Using company-level aggregate data on reserves, Thompson (2001) found that many firms are operating at a corner solution given by capacity constraints. Here we only see this in the case of flowing wells.

response of flowing wells means that the tax responses estimated in columns 1 through 3 are driven by the types of wells that could in fact respond to changes in after-tax price.

4.2 Robustness

Table 4 examines the robustness of the baseline estimates. All specifications include well and time fixed effects as well as quadratic time trends by API gravity decile. To assess the role of outlier observations, column 1 drops wells that produce an excess of 100,000 barrels of oil per month. The elasticity estimate, 0.294 (0.975), is virtually identical to column 3 of Table 3.

The data cover all wells in the state of California, including wells located in the federally owned and privately leased NPR. The extracting firm in the NPR made productions decisions, but received less than the observable after-tax price for each barrel. Furthermore, as the firm only leased the reserves, it may not have taken the exhaustibility of the reserves into account in the same way that a reserve owner would. Thus, the production response of these NPR wells to changes in after-tax price may be smaller than the response for privately owned wells.[29] Column 2 presents estimates of a model identical to that of column 3 of Table 3, but drops the NPR wells from the sample. The point estimate is larger, 0.382 (0.027), which is consistent with the idea that the operator of the NPR wells was less price sensitive than other well operators. Though the estimated after-tax price response is larger in terms of the point estimate, the difference is statistically insignificant. The NPR wells, in other words, were not significantly biasing the overall estimate of column 1.[30]

Part of the variation in after-tax price comes from the Tier II tax rates

[29] The federal government opened the NPR to drilling in 1976. From 1976 until 1998 a private firm leased access to the field and extracted oil from the reserves. The oil was sold to private refiners at the after-tax price with the proceeds divided between the extracting firm and the federal government.

[30] The supply elasticity of the NPR wells, 0.168 (0.088) (not in table), is roughly 30 percent smaller than the non-NPR elasticity, but statistically indistinguishable from the overall or non-NPR elasticities.

applied to wells that qualify as stripper wells—a status that could be endogenously determined. That is, producing an average of 10 barrels per day for 12 consecutive months yields a classification as a stripper well, which has tax advantages. Since the incentive to reduce production to gain a lower tax rate yields a negative correlation between production and tax rates, we may be concerned that the estimates reported above are biased away from zero. Stripper status is the only mutable characteristic that affords tax advantages. To investigate the impact of this potential source of bias, column 3 of Table 4 drops all stripper well observations. The elasticity, 0.380 (0.040) is similar to the non-NPR estimate of column two and not statistically distinguishable from the main results. This is despite the fact that 2,180,132 observations are from stripper wells. Non-stripper wells do exhibit a stronger per-dollar response—a one dollar increase in after-tax price leads to a nearly 14 barrel production increase—but because these wells are by definition more productive than the marginally productive stripper wells, the elasticity is roughly the same. The results from column 3 suggest that endogenously determined stripper status does not significantly affect the estimates. There are two plausibly reasons why this potential source of bias is empirically innocuous. First, the tax advantage of stripper status is small—only 10 percentage points—relative to the tax benefit of new or heavy oil. Second, stripper status requires a substantial period of low production—producers may not be willing to curtail production to roughly 300 barrels a month for a year to gain the minor tax advantage.

Much of the identifying variation comes from wells that are tax advantaged (disadvantaged) because the oil they pump is new (old) or heavy (non-heavy). The full set of wells have a broad range of heaviness—API gravities range from 10.0 to 41.9—and age—well completions date back to 1901 in some cases. Columns four and five of Table 4 narrow the range of wells examined. Column four limits the sample to wells with API gravities between 13.0 and 19.0 degrees, inclusive; that is, three degrees above and below the API gravity, 16.0, which defines heavy oil. Limiting the sample to wells of more similar API gravities, yields a smaller elasticity, 0.152 (0.050), that is statistically different from the main specification. Nonetheless higher tax rates do lead to lower production. Column five limits the sample to wells capable of producing prior to 1980. These wells were drilled and completed prior to the WPT legislation and thus their development could not have be motivated by the specifics of the WPT definition of old vs. new oil. Again, in this sub-sample of more comparable wells, the estimated elasticity,0.223 (0.033), is smaller, but significant. Taken together the results of Table 4 show that the estimated elasticity may be as low as 0.152 (0.050) or as high as 0.382 (0.027), but that in all cases well production responds positively to increases in after-tax price where the identifying variation comes from within-well and after time fixed effects.

4.3 Well Closure Decisions

Oil taxes can motivate some producers to simply shut their wells if the costs of extraction exceed the after-tax revenue. If taxes motivate well operators to close their wells, then the short-run impact of the tax could translate into a permanent reduction in oil production as the reserves remaining in the shut wells are effectively lost.[31] As the WPT was a temporary tax, it is reasonable to think that fewer operators chose to close their wells than would under a permanent tax. The estimates of equation 2 reported in Table 5 can thus be considered lower bounds of the effect of a permanent tax on well closures.

Columns 1-4 of Table 5 report marginal effects and semi-elasticities from conditional logit models. For comparison purposes, columns 5 and 6 report results from fixed effect OLS models. All of the regression models include well and time fixed effects to partial-out cost heterogeneity at the well level and time-varying factors that affect production for all wells.

As the predicted values of conditional logit models must lie between one and zero, the conditional logit model excludes wells that experience no variation in shut-in status. Identification again comes from relative within-well changes in after-tax price and the exclusion restriction requires that no time-varying well-specific factors affect production and are correlated with after-tax

[31] Shut-in wells can be re-opened but rarely are because reopening is very costly and shut-in reduces the share of remaining reserves that is feasibly extractable. Only extraordinary price events typically trigger the re-opening of shut-in wells.

price. Approximately 15.6 percent of the full sample of well-month observations are shut-in during the 1977-1985 period; 27 percent of observations that describe wells that experience variation in shut-in status—the conditional logit sample—are at shut-in.[32] The estimated after-tax price coefficient reported in column 1 of Table 4 suggests that a 10 percent increase in the after-tax price only reduces the rate of shut-in by 0.96 percentage point. This small estimated response suggests that the WPT has a negligible impact on firms' shut-in decisions. This could be because the fixed costs of operating are small relative to profit from production or because few wells are near enough to the end of their economic life. Of the wells producing in 1977, 69 percent are still producing in 1987, 44 percent are still producing in 1997 and 34 percent are still producing in 2007.

Column 2 adds a quadratic term in well age to better adjust for any potential decline in productivity that occurs over the life of the well. The estimates are virtually identical, suggesting that a linear control for well age is sufficient. Adding API gravity decile fixed effects increases the semi-elasticity to -0.109, though the increase is statistically insignificant and economically minor. Column 4 excludes wells from the NPR field. Dropping wells from the NPR field increases the point estimate of price response along the extensive margin, suggesting again that firms that lease government reserves are less price responsive than other operators, though again the difference is statistically insignificant.[33]

The conditional logit model requires variation in the dependent variable for each well in the sample. To assess the impact of limiting the sample this way, I also report shut-in semi-elasticity estimates from fixed effect OLS models. For comparison, column 5 of Table 4 reports OLS estimates for the sample of wells

[32]The data report a well as shut-in if it is every shut during the month in question, even if it was only shut for a day for routine repairs. To distinguish between very short (days long) periods of shut-in and true shut-in where the operator has take the well offline, I define shut-in as a well in shut-in status for at least two months. Using single month shut-in as the dependent variable has a negligible impact on the magnitude of the estimates which remain statistically significant though less precise, as we'd expect given the random nature of very short-term shut-in.

[33]In fact the after-tax price semi-elasticity of shut-in among NPR wells is only -0.0002 (0.0002) and statistically insignificant.

with shut-in variation that is used to estimate the conditional logit model; column 6 reports OLS estimates from the full sample of wells. The estimate using the smaller sample is nearly three times as large as the estimate from the full sample and is similar to the conditional logit estimates. The estimates of columns 5 and 6 imply that, among operators that have meaningful discretion over the shut-in status of their wells, the effect of after-tax price on the shut-in decision is significantly larger. This suggests that the sample restrictions of the conditional logit model may be partly responsible for the higher semi-elasticity estimates of columns 1 through 4 relative to column 6. Though the conditional logit coefficients are twice as large as the full sample OLS coefficient, they remain small in magnitude. Taken together, these estimates suggest that taxes did not lead to economically important rates of shut-in.

4.4 Spatial Shifting

Tables 3 and 4 establish that wells facing higher tax rates produced less oil. These models estimate the mean response of well production to after-tax price. If producers strategically reallocate production from high- to low-tax wells on their lease, these estimated responses may overstate the overall response to variation in after-tax price. Specifically, if producers are strategically shifting production spatially, then responses at the lease level should be smaller than the responses suggested by the well-level regressions. Table 6 examines the degree of spatial shifting. Odd-numbered columns report estimates from well-level regressions while even-numbered columns report estimates from lease-level regressions.[34] Column 1 is the same specification as column 3 of Table 3 but omits wells with missing lease names. Dropping these wells does not affect the estimated elasticity, which is within rounding of the baseline in Table 3. Column 2 of Table 6 reports the lease-level model. The estimated elasticity is slightly smaller, 0.245 (0.075) versus 0.295 (0.035) and of lower precision though still significant at the 1%-level. The difference between the well-level

[34]For the lease-level regressions oil production data is summed for wells on the same lease that share the same operator. After-tax price is the production weighted average of after-tax price of wells in the lease with the same operator. Aggregating to purely to the lease-level rather than the operator-lease level does not substantively affect the results.

and lease-level elasticites, however, is not statistically or economically significant. The comparability of the estimates suggests that producers are not engaging in significant production reallocation and that such reallocation is not driving the estimates reported in Tables 3 and 4.

Columns 3 and 4 of Table 6 restrict the sample to leases that have wells that are classified into at least two different WPT tiers; this restricts the sample to only leases where strategic spatial shifting opportunities exist. The well-level estimates, column 3, suggests that this sample is somewhat more elastic that than the full sample of column 1. The lease-level estimate, column 4, is higher than the full sample estimate reported in column 2 as well. Comparing the point estimates of columns 3 and 4 suggests that wells in this restricted sample exhibit a slightly higher elasticity not fully mirrored by aggregate lease-level estimates. The difference, however, is not statistically significant. Columns 5 and 6 further require that the leases included in the sample feature both Tier I and Tier III wells—wells with the greatest rate disparity and thus the strongest incentives for spatial shifting. While the well-level and lease-level point estimates are somewhat higher, 0.441 (0.04) and 0.359 (0.199) respectively, their difference is again statistically insignificant.

In addition to different tax rates among wells on a lease, having more wells on a lease may facilitate strategic spatial production shifting. Columns 7 and 8 limit the sample to leases with at least three wells to ensure there is meaningful scope for spatial shifting. Once again the point estimate from the well-level regression, 0.307 (0.038), is larger than the lease-level estimate, 0.229 (0.115). Although both estimates are significant at at least the 5%-level their difference is too small to be statistically significant.

Taken together the regressions presented in Table 6 show a consistent pattern of lease-level point estimates that are smaller than the well-level estimates, but the difference is not large enough to be statistically significant. All the estimates show that higher taxes lead to less production, and the response is not primarily driven by strategic spatial shifting of production from tax-disadvantaged to tax-advantaged wells.

5 Conclusion and Policy Implications

This paper uses new detailed data on the quantity of oil produced by wells in California to estimate the effect of tax- and price control-induced variation in oil prices on production decisions. The unusual cross-sectional variation in after-tax price afforded by these government interventions allows for flexible controls for time-varying factors, like price expectations and underlying changes in technology, that affect oil production. The estimated coefficients imply an elasticity ranging between 0.295 (0.038) and 0.336 (0.042), meaning that a 10 percent excise tax would lead to a roughly 3 percent change in domestic oil production.

I find that while oil production from existing wells is responsive to the after-tax price, after-tax price has no appreciable impact on wells that flow in accordance with their natural subsurface pressure. Only pumped wells alter production in light of taxes. I also find no evidence of significant spatial shifting of production from tax disadvantaged wells to tax advantaged wells on the same lease, meaning that the estimated elasticity is largely a real reduction in production. Because these estimates imply that producers alter their behavior in response to tax changes, they suggest that the incidence of an oil excise tax cannot be modeled simply as a tax on the rents of oil producers. The higher elasticities estimated here make clear that producers react to taxes, though the elasticity is much less than unity. State taxes legislated today effectively raise the marginal cost of production, potentially reducing extraction and leading to deadweight loss. The elasticities estimated, however, are much below unity, suggesting the WPT and potentially more recent state taxes discourage production but not to a self-defeating extent.

The empirical findings bear on short-run production decisions, and it is important to remember several cautions about their broader interpretation. First, taxes are likely to delay or curtail exploration and development activities—the taxes delay or reduce profits, so firms will want to delay or curtail investments. This response margin is not captured by the analysis presented above. Second, California wells and the oil they produce have higher extraction costs than the average U.S. well. Because the oil is of such high specific gravity (low API gravity) it is costly to extract, or lift, to the surface. High marginal costs may make California producers more sensitive than most to taxation. Finally, the estimates presented here are identified by policies from the late 1970s and 1980s and are thus historic. Technological changes that have improved extraction efficiency may make these estimates less applicable to current policy proposals. Nonetheless, the elasticities estimated here, with caveats, can help inform current policy. The estimates here suggest that domestic oil production is not fixed but is in fact responsive to after-tax price. State and local governments considering oil taxes should view the revenues resulting from higher taxes as likely also entailing the cost of reduced production, though the reduction is unlikely to come from producers shutting in wells. Much like the WPT may state oil taxes attempt to tax different types of wells at differential rates. The limited spatial shifting seen under the WPT, despite considerable variation in after-tax price, suggests that these differential rates may not spur tax avoidance, and could help keep marginal wells in production.

Appendix (For Online Publication)

Further institutional details are provided here.

The United States is the third largest oil producer[35], behind only Saudi Arabia and Russia; California is the third largest oil producing state in the U.S.

Aggregate U.S. oil production comprised roughly 15 percent of total world production while price controls and windfall profit taxes were in place, a substantial but decidedly minority share. Domestic pre-tax prices are set by the global oil market. Unlike most other oil producing nations, oil extraction in the U.S. is a competitive market where large international oil firms operate alongside many smaller independent producers. Though the large international companies that operate in the U.S. also operate abroad, their market share was dramatically undercut by the establishment of the Organization of Petroleum Exporting Countries (OPEC) in 1960. By the mid-1970s, OPEC nations accounted for roughly half of world production and coordinated their production decisions in an effort to influence price. Though the evidence on OPEC's effectiveness as a cartel is mixed,[36] if any group of producers had the market share and coordination necessary to affect prices, it was and remains nationalized producers rather than the competitive fringe that operates in the U.S.[37]

During the price control era, a permit trading system allocated low-price domestic crude among refiners.[38] Refiners did not face shortages since imported oil was always available for purchase.

[35] The U.S. was the third largest producer in the 1970s and 1980s as well though U.S.S.R production totals were less accurately measured.

[36] Hamilton (2009) reviews recent production and quota discrepancies among OPEC nations and finds that OPEC members frequently cheat with respect to their quotas and there is little evidence of a clear enforcement mechanism. Also see Alhaji and Huettner (2000) for a review of 13 studies assessing the effectiveness of OPEC as a cartel.

[37] As the U.S., including California refiners, imports oil, within the range of transportation costs, domestic producers may have some pricing power. Given that transport costs comprise roughly 5 percent of oil prices, domestic producers have only a small scope of pricing power.

[38] Since only domestic crude was subject to price controls, refiners who procured domestic crude earned rents. The federal government created a system of tradable permits to allocate low-priced domestic crude among refiners to "fairly" distribute the potential windfall. Permits were allocated according to historic crude sourcing.

California is divided into six oil and gas districts. Each month between 1977 and 1985, total California production ranged between 2.37 million barrels in February 1978 and 3.20 million barrels in August 1985. Roughly 16.1 percent of wells are shut-in on average; there is some variation in shut-in rates, with the smallest share of shut-in wells, 14.5 percent, during October 1978 and the largest share, 17.5 percent, in December 1985. Each of the top five producing wells accounts for less than 0.5 percent of total production.

During the price control era oil from the same well was classified as lower and upper-tier oil, with upper-tier oil receiving a higher price. Lower-tier oil corresponded to what regulators believed was the "expected" level of production based on the property's production history. Until the well produced its lower-tier quota, all oil it produced would sell at the lower-tier price. If the operator exceeded his lower-tier quota, then all additional oil produced would sell at the higher upper-tier price. The determination of whether a barrel of oil subject to price controls was upper- or lower-tier is beyond the capacity of the data. This analysis assigns all price-controlled wells the upper-tier selling price, as it is the more likely price for marginal production from a California well.

References

Anderson, S. T., R. Kellogg, and S. W. Salant (2014): "Hotelling Under Pressure," NBER Working Paper No. 20280.

Dahl, C., and M. Yücel (1991): "Testing Alternative Hypotheses of Oil Producer Behavior," *The Energy Journal*, 12(4).

Dahl, C. A., and T. Duggan (1998): "Survey of Price Elasticities from Economic Exploration Models of US Oil and Gas Supply," *Journal of Energy Finance and Development*, 3(2).

Dasgupta, P., and G. M. Heal (1979): *Economic Theory and Exhaustible Resources*. Cambridge University Press, Cambridge, UK.

Editorial (January 22, 1980): "The Close-the-Wells Tax," *Wall Street Journal*, p. A20.

Griffin, J. M. (1985): "OPEC Behavior: A Test of Alternative Hypotheses," *American Economic Review*, 75(5).

Hamilton, J. D. (2009): "Understanding Crude Oil Prices," *Energy Journal*, 30(2), 179–206.

Hogan, W. W. (1989): "World Oil Price Projections: A Sensitivity Analysis," Energy and environmental policy center memo, John F. Kennedy School of Government. Harvard University.

Hotelling, H. (1931): "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39(2), 137–175.

Hyne, N. J. (2001): *Nontechnical Guide to Petroleum Geology, Exploration, Drilling, and Production*. PennWell Corporation, Tulsa, OK, 2nd edition edn.

Joint Committee on Taxation (1981): "General Explanation of the Crude Oil Windfall Profit Tax Act of 1980," Discussion Paper Discussion Paper H.R. 391, PL 96-223.

Jones, C. T. (1990): "OPEC Behaviour under Falling Prices: Implications for Cartel Stability," *The Energy Journal*, 11(3).

Kalt, J. P. (1981): *The Economics and Politics of Oil Price Regulation: Federal Policy in the Post-Embargo Era*. MIT Press, Cambridge, MA.

KELLOGG, R. (2014): "The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling," *American Economic Review*, 104(6), 1698–1734.

KILIAN, L. (2009): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99(3), 1053–1069.

LAZZARI, S. (2006): "The Crude Oil Windfall Profit Tax of the 1980s: Implications for Current Energy Policy," Report for Congress RL33305, Congressional Research Service.

LEWIS, T. R. (1977): "Attitudes towards risk and the optimal exploitation of an exhaustible resource," *Journal of Environmental Economics and Management*, 4(2), 111–119.

MAURITZEN, J. (2014): "The Effect of Oil Prices on Offshore Production: Evidence from the Norwegian Continental Shelf," NHH Dept. of Business and Management Science Discussion Paper No. 2014/7.

PINKYCK, R. S. (1999): "The Long-Run Evolution of Energy Prices," *The Energy Journal*, 20(2), 1–27.

RAMCHARRAN, H. (2002): "Oil Production Responses to Price Changes: An Empirical Application of the Competitive Model to OPEC and Non-OPEC Countries," *Energy Economics*, 24(2), 97–106.

RODRIGUE, J.-P., C. COMTOIS, AND B. SLACK (2013): *The Geography of Transport Systems*. Routledge.

SMITH, J. L. (2012): "Modeling the Impact of Taxes on Petroleum Exploration and Development," IMF Working Paper.

SMITH, R. T., M. BRADLEY, AND G. JARRELL (1986): "Studying Firm-Specific Effects of Regulation with Stock Market Data: An Application to Oil Price Regulation," *The RAND Journal of Economics*, 17(4), 467–489.

SOLOW, R. M. (1974): "The Economics of Resources or the Resources of Economics," *American Economic Review Papers and Proceedings*, 64(2), 1–14.

THOMPSON, A. C. (2001): "The Hotelling Principle, Backwardation of Futures Prices, and the Value of Developed Petroleum Reserves—the Production Constraint Hypothesis," *Resource and Energy Economics*, 23, 133–156.

Vietor, R. H. K. (1984): *Energy Policy in America Since 1945: A Study of Business-Government Relations*. Cambridge University Press, New York, NY.

Weinstein, M. C., and R. J. Zeckhauser (1975): "The Optimal Consumption of Depletable Natural Resources," *Quarterly Journal of Economics*, 89(3), 371–392.

Figure 1: Timeline of Price Decontrol and Enactment of 1980 Windfall Profit Tax

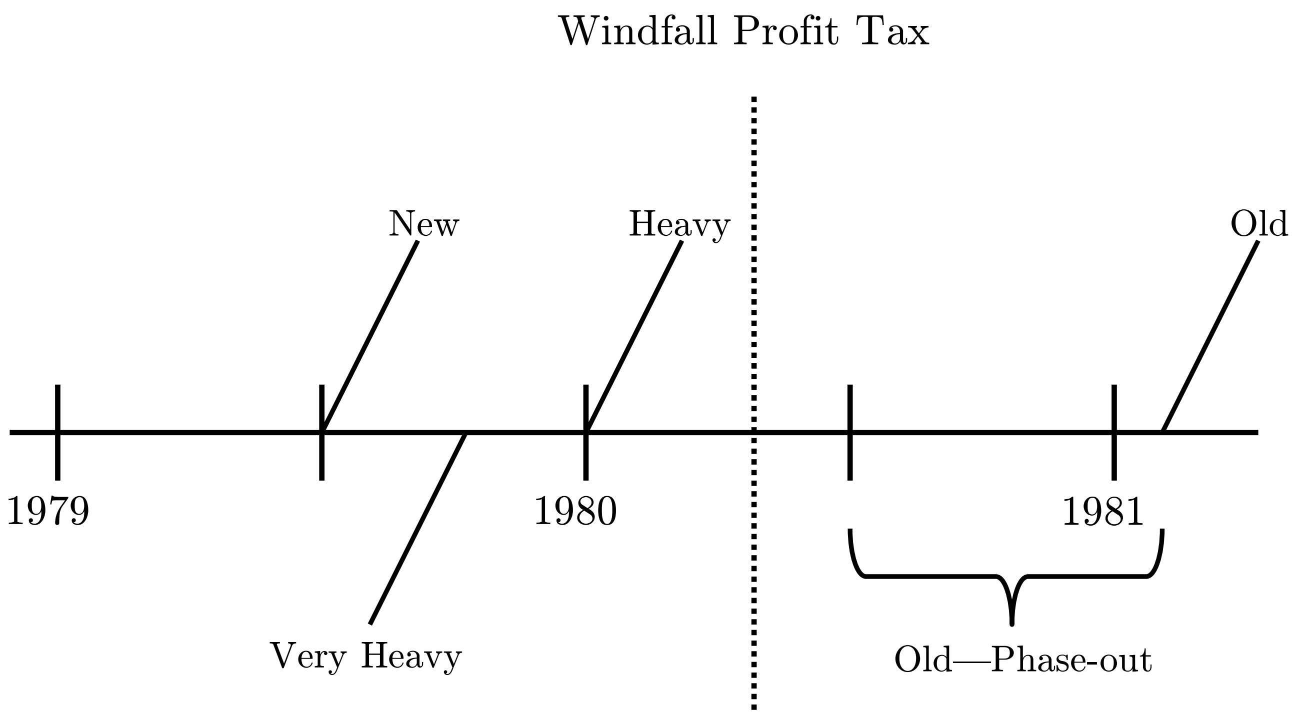

New oil (oil extracted from wells that did not produce oil in 1978) was decontrolled in June 1979.

Very heavy oil (oil with an API gravity of less than 16 degrees) was decontrolled in September 1979.

Heavy oil (oil with an API gravity of less than 20 but at least 16 degrees) was decontrolled in January 1980.

Old oil (oil extracted from wells that produced oil in 1978) was gradually decontrolled between January 1980 until January 28, 1981. During the phase-out period, old oil sold at a price that was equal to the weighted average of the world market price and the price control price ceiling, with the weight on the world market price growing by 0.046 each month. Old oil was fully decontrolled by President Reagan on January 28, 1981. February 1981 was the first full month in which old oil was decontrolled.

1980 Windfall Profit Tax was signed into law April 2, 1980 and went into effect immediately

Figure 2: Prices, Before and After Taxes and Fixed Effects, Two Wells

Well 120005: Livermore Field, Operator: Hershey Oil Corp.

Old oil, API gravity of 23; stripper starting Oct. 1982 (70% tax rate until Oct. 1982, then 60 percent)

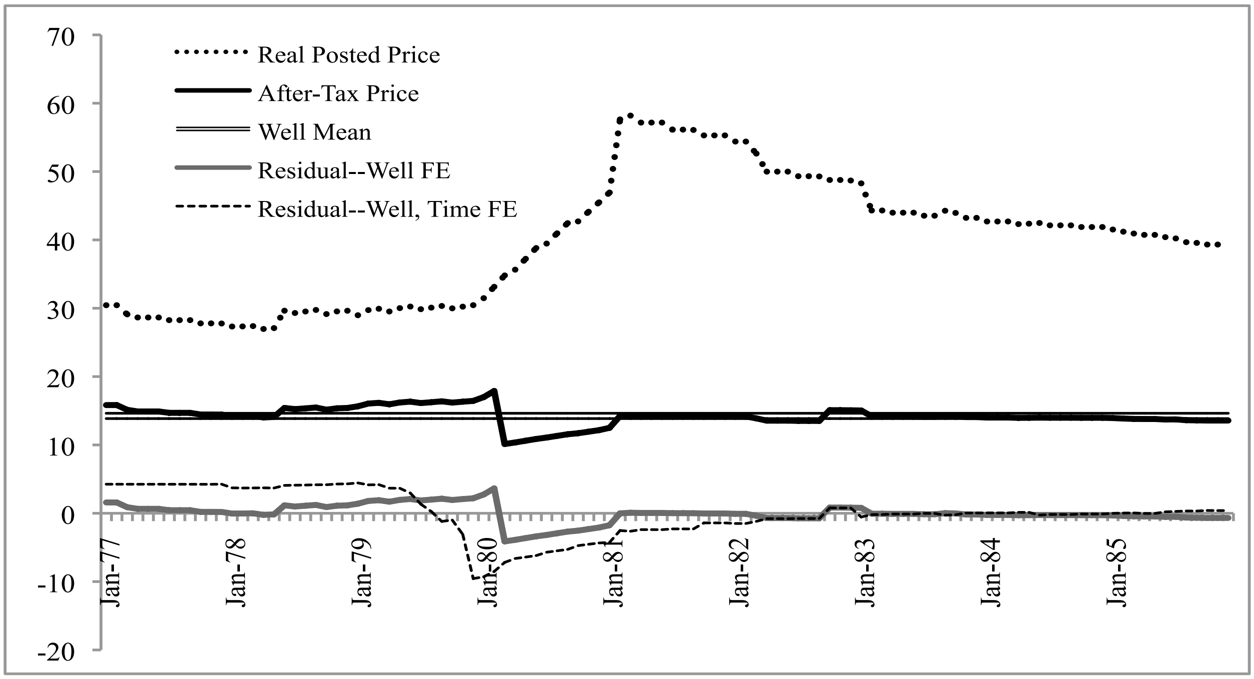

Well 1300071: Brentwood Field, Operator: Occidental Petroleum Corp.

New oil, API gravity of 40.7; never stripper (30% tax rate until 1982, then gradual decrease to 22.5%)

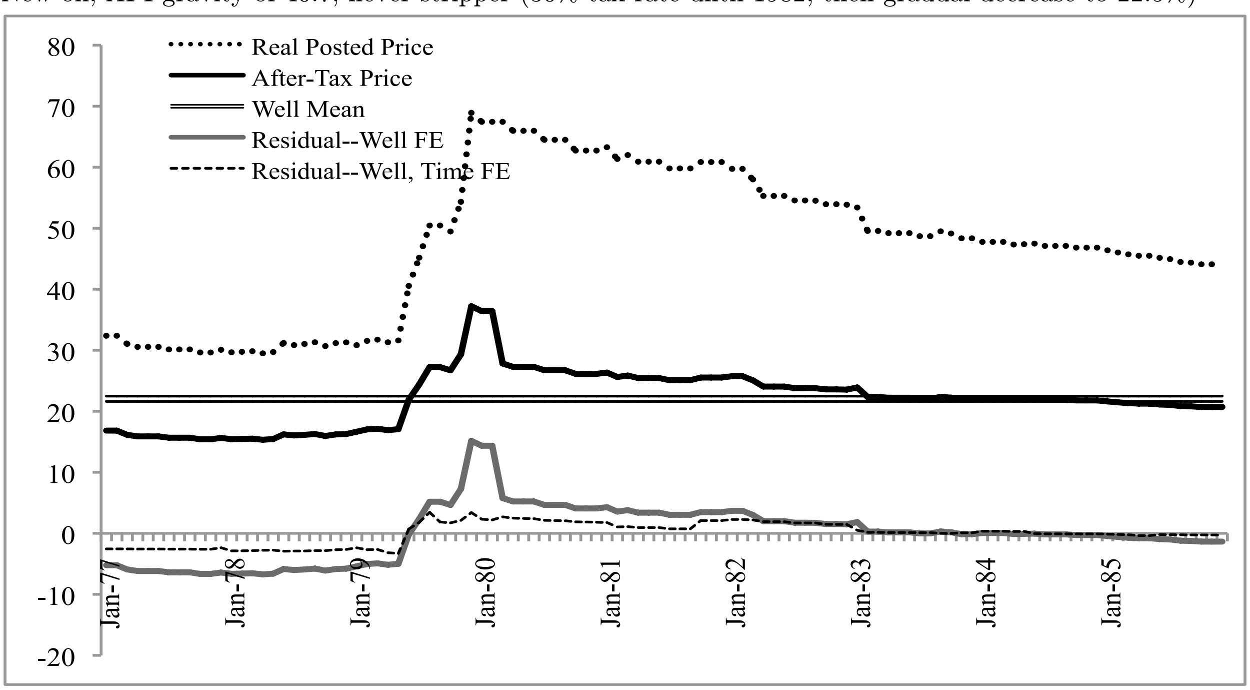

Table 1: U.S. Supply Elasticity Estimates from Previous Studies

Study	Sample Period	Data	Elasticity Estimate
Griffin (1985)	$1971Q_1$ - $1983Q_3$	Quarterly data on total U.S. production and average pre-tax posted price from $1971Q_1$ to $1976Q_2$, average pre-tax first purchase price from $1976Q_3$ to $1983Q_3$. No controls.	-0.05 (0.02)
Hogan (1989)	1966 - 1987	Annual data on total U.S. production and average pre-tax first purchase price.	0.09 (0.03)
Jones (1990)	$1983Q_3$ - $1988Q_4$	Quarterly data on total U.S. production and average pre-tax first purchase price. No controls.	0.07 (0.04)
Dahl and Yücel (1991)	$1971Q_1$ - $1987Q_4$	Quarterly data on total U.S. production and average pre-tax first purchase price. Added annual controls for production costs, number of wells drilled, U.S. income and world oil production.	-0.08 (0.06)
Ramcharran (2002)	1973 - 1997	Annual data on total U.S. production and average pre-tax first purchase price. Linear time trend included.	0.05 (0.02)

Note: Some of these studies estimated supply elasticities for total U.S. production as a part examinations of market structures among OPEC and non-OPEC countries; nonetheless they are the studies cited in supply elasticity surveys such as Dahl and Duggan (1998). All of these analyses rely on aggregate time-series data for the U.S. Standard errors are in parentheses.

Table 2: Summary Statistics

	Mean	Standard Deviation	
		Overall	Within-Well
Oil Production (barrels)	476.4	1,473.8	868.5
Oil Production if Producing	663.2	1,702.9	869.3
After-Tax Price ($)	18.3	4.1	3.5
WPT Tax Rate	.21	.24	.19
Purchase Price ($)	41.0	10.1	9.76
API Gravity (degrees)	18.2	6.8	1.1
Number of Wells	75,342		
Observations	6,517,139		

Note: The summary statistics above describe the well-month observations that comprise the sample for the main regression analysis. Not all 75,342 wells report 108 observations since new wells are drilled and old wells are abandoned during the sample period.

Table 3: Regressions of Quantity Produced on After-Tax Price

	(1)	(2)	(3)	(4)	(5)
After-Tax Price	8.730***	8.741***	7.673***	9.287***	-8.606
	(1.082)	(1.082)	(0.977)	(0.621)	(9.116)
Well Age		-1.228***	-1.335***	-3.164***	5.870
		(0.0809)	(0.441)	(0.690)	(6.326)
Well Age Squared		-0.000309	-0.000198	0.000125	7.50e-05
		(0.000229)	(0.000228)	(0.000288)	(0.000398)
After-Tax Price Elasticity	0.335***	0.336***	0.295***	0.371***	-0.262
	(0.042)	(0.042)	(0.038)	(0.025)	(0.278)
Observations	6,517,140	6,517,140	6,517,140	5,698,198	818,942
Number of Wells	75,342	75,342	75,342	72,797	33,198
Well FE	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y
API Gravity Time Trends			Y	Y	Y

Note: The table presents OLS regressions where the dependent variable is the quantity of oil produced by well i in month t. After-Tax Price is the posted price at which oil from well i sold during month t, net of corporate and Windfall Profit taxes. The coefficient on After-Tax Price, β in equation 1, reports the supply response of well operators to net price.

Column 1 is the baseline specification; it includes time and well fixed effects. Column 2 adds a quadratic function of well age. Column 3 includes separate quadratic time trends, slopes and coefficients, by API gravity decile. Column 4 restricts the sample to to only pumped wells. Column 5 restricts the sample to only flowing wells, which do not require mechanical lift to produce oil. The elasticities for all specifications are the product of the coefficient estimate and the ratio of after-tax price to average quantity for the estimation sample of producing wells.

All heteroskedasticity robust standard errors are clustered at the individual well level.

Robust standard errors in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 4: Regressions of Quantity Produced on After-Tax Price, Robustness

	(1)	(2)	(3)	(4)	(5)
After-Tax Price	7.629***	8.682***	13.85***	3.456***	6.480***
	(0.975)	(0.618)	(1.475)	(1.128)	(0.953)
Well Age	-1.618***	-1.652***	7.974***	-0.550*	-1.800***
	(0.348)	(0.329)	(2.871)	(0.299)	(0.400)
Well Age Squared	-0.000264	-0.000218	-0.000442	-0.000683**	0.000245
	(0.000218)	(0.000186)	(0.000415)	(0.000296)	(0.000321)
After-Tax Price Elasticity	0.294***	0.382***	0.380***	0.152***	0.223***
	(0.038)	(0.027)	(0.040)	(0.050)	(0.033)
Observations	6,517,137	6,350,819	4,170,687	3,079,546	5,030,912
Number of Wells	75,342	73,548	75,220	41,630	49,388
Well FE	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y
API Gravity Time Trends	Y	Y	Y	Y	Y

Note: The table presents OLS regressions where the dependent variable is the quantity of oil produced by well i in month t. After-Tax Price is the posted price at which oil from well i sold during month t, net of corporate and Windfall Profit taxes. The coefficient on After-Tax Price, β in equation 1, reports the supply response of well operators to net price.

All specifications include well and time fixed effects as well as quadratic time trends by API gravity decile. Column 1 drops wells that produce more than 100,000 barrels of oil per month. Column 2 drops observations from the federal Naval Petroleum Reserve. Column 3 drops stripper wells. Column 4 includes only observations with an API gravity between 13.0 and 19.0. Column 5 includes only wells producing before 1980.

All heteroskedasticity robust standard errors are clustered at the individual well level.

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 5: Shut-In Decisions and After-Tax Price

	(1) Cond. Logit	(2) Cond. Logit	(3) Cond. Logit	(4) Cond. Logit	(5) OLS	(6) OLS
After-Tax Price	-0.0052***	-0.0052***	-0.0059***	-0.0060***	-0.0043***	-0.0015***
	(0.0008)	0.0008)	(0.001)	(0.0009)	(0.0004)	(0.0002)
Well Age	0.0126***	0.0126***	0.0141***	0.0121***	0.0014***	0.0005***
	(0.0007)	(0.0007)	(0.0013)	(0.0007)	(0.0000)	(0.0000)
Well Age Squared		0.000				
		(0.0000)				
After-Tax Price Semi-Elasticity	-0.096***	-0.096***	-0.109***	-0.111***	-0.080***	-0.027***
	(0.015)	(0.015)	(0.018)	(0.017)	(0.0079)	(0.0034)
Observations	2,694,267	2,694,267	2,694,267	2,571,746	2,694,267	6,517,140
Number of Wells	29,297	29,297	29,297	27,989	29,297	75,342
Well FE	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y
API Gravity FE	N	N	Y	N	N	N

Note: The table presents conditional logit and OLS regressions where the binary dependent variable is one if well i is shut-in in month t and zero if it is not. After-Tax Price is the posted price at which oil from well i sold during month t, net of corporate and Windfall Profit taxes. The coefficient on After-Tax Price, δ of equation 2, describes the extensive response of operators to net price.

Column 1 includes a full set of month by year and well fixed effects and a linear control for well age. Column 2 adds a quadratic term in well age. Column 3 adds dummies for each API gravity decile. Column 4 excludes observations from the federally owned NPR. Column 5 estimates an OLS model with well and time fixed effects using the same sample of wells that experience variation in shut-in status. Column 6 estimates the fixed effect OLS model using the full sample of wells. The semi-elasticity calculations for all specifications is the product of the marginal effect estimate and average after-tax price.

All heteroskedasticity robust standard errors are clustered at the individual well level.

Robust standard errors in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 6: Spatial Shifting: Well and Field-Level Regressions of Quantity Produced on Price

	(1) well	(2) lease	(3) well	(4) lease	(5) well	(6) lease	(7) well	(8) lease
After-Tax Price	6.274***	5.066***	9.171***	7.325**	11.37***	12.41**	6.563***	5.225***
	(0.744)	(1.550)	(0.850)	(3.348)	(1.033)	(6.861)	(0.814)	(2.633)
Well Age	-0.770***	-3.229***	-2.383***	-3.204***	-2.145**	-4.470***	-0.692	-2.657***
	(0.553)	(0.940)	(0.891)	(0.767)	(1.049)	(1.942)	(0.577)	(0.520)
Well Age Squared	0.000193	0..0031***	0.000802***	0.00409***	0.000717	0.00855***	0.000126	0.00376***
	(0.000242)	(0.000660)	(0.000295)	(0.00104)	(0.000648)	(0.00371)	(0.000258)	(0.000834)
After-Tax Price Elasticity	0.295***	0.245***	0.405***	0.307**	0.441***	0.359**	0.307***	0.229***
	(0.035)	(0.075)	(0.038)	(0.140)	(0.040)	(0.199)	(0.038)	(0.115)
Observations	4,484,531	391,002	2,238,017	136,855	1,763,386	71,325	4,226,585	188,246
Number of Wells or Leases	51,153	4,804	23,812	1,373	18,766	706	50,497	2,267
Well FE	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y	Y	Y
API Gravity Time Trends	Y	Y	Y	Y	Y	Y	Y	Y

Note: The table presents OLS regressions where the dependent variable is either the quantity of oil produced by well i in month t or the average well production on lease l in month t. After-tax price is either the posted price at which oil from eel i sold during month t, net of corporate and Windfall Profit taxes, or the average of such prices for all wells on a lease l, weighted by oil production. The coefficient on the After-Tax Price variable is the coefficient of interest and describes the supply response of well operators to changes in net price.

All specifications include time fixed effects and either well or lease fixed effects as well as quadratic time trends by API gravity decile; wells from the National Petroleum Reserve are dropped. Odd columns are well-level regressions and even columns are lease level regressions. Columns 1 and 2 drop all wells with missing lease names. Columns 3 and 4 restrict the sample to leases that have wells that are classified into at least two different Windfall Profit Tax tiers. Columns 5 and 6 require that the leases included in the sample include both Tier 1 and Tier 3 wells—wells with the greatest tax rate disparity and thus the strongest incentives for spatial shifting. Columns 7 and 8 limits the sample to leases with at least three wells to ensure there is meaningful scope for spatial shifting.

All heteroskedasticity robust standard errors are clustered at the individual well or lease level.

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1