DISCUSSION PAPER

February 2012; revised December 2013 ■ RFF DP 12-12-REV

A Dynamic Model of Cleanup

Estimating Sunk Costs in Oil and Gas Production

Lucija Muehlenbachs

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

A Dynamic Model of Cleanup:
Estimating Sunk Costs in Oil and Gas Production

Lucija Muehlenbachs[1]

December 3, 2013

Abstract

The environmental remediation required to permanently decommission most industrial projects is an expensive, irreversible investment. Real options literature shows that temporary closure has value under uncertainty. However, even if there is no intention to restart operations, there is an incentive to label a closure as "temporary," to avoid having to remediate ongoing or future environmental externalities. I estimate a dynamic discrete choice model of closure under price and quantity uncertainty to evaluate the likelihood of reactivation. The model reveals that the option to temporarily close is being widely used to avoid environmental remediation of oil and gas wells in Canada.

Jel–Classification: Q30, Q47, Q58, C63

Keywords: environmental remediation, real options, structural estimation, oil and gas, reserve growth

[1]I am indebted to John Rust for his advice and encouragement throughout all stages of this project. I also thank Robert Cairns, Timothy Fitzgerald, Steven Gabriel, Adan Martinez-Cruz, Marc Nerlove, Lars Olson, Jim Quehl, Mar Reguant, Sirish Shah, Stefan Staubli, Chris Timmins, and Seth Wechsler for helpful suggestions and Kenneth Judd and Che-Lin Su for introducing me to the optimization software. A previous version of this paper was entitled "Testing for Avoidance of Environmental Obligations" and "Idle Oil Wells: Half Empty or Half Full?"

Resources for the Future, 1616 P Street, NW, Washington, DC 20036. Email: muehlenbachs@rff.org

1 Introduction

Once an industrial project, such as a landfill, nuclear power plant, mine, or oil field has reached the end of its life, the costs associated with permanently decommissioning operations tend to be very high. Literature on real options has shown that inertia is optimal in dynamic decisions involving sunk costs in uncertain environments [Dixit, 1989, 1992, Dixit and Pindyck, 1994]. If there is a chance that a project will be restarted in the future, there is value to temporarily closing it, and postponing the investment needed to decommission it. However, once a project is temporarily closed it could remain in a state of hysteresis; there is a sunk cost to reactivate so even if the forces behind the project's closure are reversed the project may still not be reactivated.

There are high sunk costs associated with decommissioning industrial projects because of the requirements to remediate existing environmental damages and implement measures to prevent ongoing or future damages. Therefore, by not decommissioning a project, a mothballed, or temporarily closed state carries environmental risks that might not be internalized by the owner of the project. A difficulty, however, arises because regulators not wanting to cut the life of a viable project short will allow projects to be temporarily closed. It therefore could be the case that the value of the option to reactivate is zero or negative, but an owner "temporarily" closes a project as a way to avoid paying for decommissioning. Using a temporary closure in lieu of a permanent closure is facilitated by regulators not having perfect information on the costs or expectations of the operators.

In this paper I find evidence to help disentangle whether a temporary closure is indeed being used to avoid environmental cleanup. To the best of my knowledge, this is the first attempt to understand the true motivations behind temporary closures. To do so, I first build a dynamic programming model of the choice to temporarily close, decommission, or reactivate under uncertainty in prices and quantities. I then estimate the structural parameters in the model using data on historical operating decisions and changes in prices and productivity. Assuming prices and productivity are believed to follow the same path as in the past, I can use the fully specified model to predict how likely a reactivation is when presented with ideal operating conditions. In doing so, I demonstrate that data on closure decisions can be used to structurally estimate a real options model, which can be used to test the likelihood that a temporary closure is in actuality permanent.

I apply this framework to the oil and gas industry. Currently, there are hundreds of thousands of "temporarily abandoned" oil and gas wells scattered across North America, including over 3,700 temporarily abandoned wells in the Gulf of Mexico.[1] Permanent decommissioning of wells is required,[2] however regulators, not wanting to impede production, make temporary closure an option, despite the potential environmental externalities of not decommissioning.[3] Postponing permanent decommissioning also increases the risk that a firm will declare bankruptcy before undertaking the expense of the environmental cleanup. The inventory of "orphaned wells," or wells without a party responsible for plugging, is quite startling: New York for example has 44,600, Pennsylvania has over 100,000, and Texas has roughly 10,000.[4] Furthermore, there is a current boom in drilling for oil shale and shale gas making it evermore important for policymakers to understand firms' incentives to environmentally remediate wells once they are exhausted. Therefore, I use data on the decisions made for 84,000 conventional oil and gas wells in Alberta, Canada to estimate the structural parameters of the dynamic programming-real options model of well operating decisions. This paper has two main contributions.

The first contribution is in testing the goodness-of-fit of a real options model to actual firm behavior. Real options models extend the Black and Scholes [1973] and Merton [1973] theory for financial options to that of irreversible real investments. Unlike the case of financial

[1] According to the Bureau of Ocean Energy Management's Borehole Dataset http://www.data.boem.gov/homepg/data_center/well/well.asp.

[2] This involves "plugging & abandonment," where equipment is removed, groundwater formations are sealed with cement and the surrounding land reclaimed.

[3] Not permanently decommissioning an oil or gas well increases the risk of contamination of the atmosphere, drinking water, vegetation, and soil; lost productivity of other wells in the same pool; erosion; forest fragmentation; and even explosions [Kubichek et al., 1997, Williams et al., 2000, Mitchell and Casman, 2011].

[4] According to the Interstate Oil and Gas Compact Commission's (IOGCC) Orphaned Wells State's Progress http://groundwork.iogcc.org/topics-index/orphaned-wells/state-progress and [Railroad Commission of Texas, 2006].

derivative models, empirical investigations testing the fit of real options models to data are rare. Gamba and Tesser [2009] note that this is due to two factors: the values of the state variables are often not observed and there is quantity uncertainty. In this paper, I circumvent these two issues through the use of a dataset on the reserve estimates of oil and gas pools over time. This is the first time these data have been used in the economics literature and they provide annual estimates for of the remaining reserves of over 42,000 oil and gas pools. These data provide me with information not used in previous studies: an estimate of a project's current productivity as well as a way to estimate uncertainty in its productivity in the future (i.e., technological advances in enhanced recovery methods that might increase recoverable reserves). This information is key in determining the likelihood that a well will be reactivated; without data on reserves I would not be able to distinguish whether a well is inactive because the decommissioning costs are high or because the remaining reserves are significant enough to warrant reactivation in the future.

Much of the literature on real options relies on examples from the natural resource industry and models of many different discrete decisions in the industry have been developed.[5] The empirical work on real options has relied on having data for the cost parameters rather than structurally estimating these parameters, thereby restricting investigations to small sample sizes and results to a comparison of stylized facts from the predictions of real options to the data (e.g., Moel and Tufano [2002], Paddock et al. [1988], Slade [2001], Harchaoui and Lasserre [2001], and Hurn and Wright [1994] examine irreversible investments in natural resource industries and all use data on fewer than 300 projects). Exceptions are from dynamic models of the decision to drill oil wells, Levitt [2009] focusing on the effect of learning and Kellogg [2010] on the effect of price volatility. However, these papers also do not measure quantity uncertainty.

The second contribution is in presenting a framework to understand the relative impor-

[5] Indeed, even the same three choices presented in this paper (to activate, inactivate, or decommission a project) have been modeled by Brennan and Schwartz [1985], Castillo-Ramirez [1999], Cortazar and Casassus [1998], Cortazar et al. [2001], Stensland and Tjostheim [1989], Dixit and Pindyck [1994], and Gamba and Tesser [2009]. However, these authors did not apply their models to real data and treated decommissioning costs as negligible.

tance of the value of investment flexibility and the perverse incentive to avoid cleanup costs. The extent to which permanent closures are being labeled as temporary has not been raised or investigated before. Throughout the real options literature, the permanent closure option is often downplayed. Decommissioning costs are treated as negligible or null [Brennan and Schwartz, 1985, Dixit and Pindyck, 1994], or the option of decommissioning is completely left out of the choice set [Moel and Tufano, 2002, Slade, 2001, Mason, 2001, Paddock et al., 1988]. By assuming away decommissioning costs, the previous literature has overlooked the case of firms continuing to maintain the option to reactivate a project, even when they have no intention, or there is no option value, to reactivate. When the costs from mothballing a project are small relative to the decommissioning costs, this behavior would be privately optimal, but when there are environmental externalities associated with mothballing, this behavior would not necessarily be socially optimal. If there is no potential or intention to reactivate a hazardous project, regulators have reason to implement policies to ensure that environmental obligations will be met.[6] The framework presented in this paper is important because the effectiveness of any policy relies on the underlying reasons for the temporary closures at hand.

The estimated model suggests that only with a drastic, arguably implausible, increase in prices and recovery rates will there be a significant increase in the number of reactivated oil and gas wells, implying that wells are typically left inactive not because of the option to reactivate, but rather to avoid costly environmental obligations. This is a function of the expected value of a reactivated well being less than the expected value of an inactive well, even when an operator is in an ideal state of nature. Considering that energy independence is frequently sought by policymakers, it is also important to consider the quantity of oil or gas that the reactivated wells might contribute to the energy supply. Under high oil and gas prices, the recoverable reserves increase (more so for gas than oil), but nonetheless, the

[6] One of the main reasons for a policy to induce prompt environmental cleanup is the risk that the firm will declare bankruptcy. The concern that oil and gas companies may walk away from their environmental obligations has been brought up by Boyd [2001], Parente et al. [2006], and Ferreira et al. [2003]. While these authors discuss bonding mechanisms, the model here can be used to quantify the effect of a bond on production as well as the choice to undertake cleanup.

number of reactivated wells remains minimal. Furthermore, the model predicts that the contribution to the oil and gas supply from these reactivated wells is only marginal. These findings have far reaching implications for the oil and gas industry. If decommissioning costs are not being internalized, the development of oil and gas reserves would be at a rate above what is socially optimal. The policy implication would be to create stronger mechanisms to internalize the costs of decommissioning. One such mechanism is to increase bonding requirements, which are arguably too low at present.[7] This paper demonstrates that in designing policies to decommission oil and gas wells unnecessary weight has been placed on not jeopardizing production.

2 Oil and Gas Well Background

This paper focuses on Alberta, the main oil and gas producing province of Canada. Extensive record keeping in Alberta has resulted in comprehensive data on the industry. Also, there is no limit to the length of time that a well can be left inactive in Alberta.[8] In the U.S., even when there is a time limit, permission for extended "temporary abandonment" is easily granted and the fine for leaving a well inactive without permission is usually small; for example, in Kansas the fine is only $100 [State Corporation Commission of Kansas, 2010]. Because of the externalities associated with not decommissioning a well, it is required that wells be decommissioned, but it is up to the producer to decide the time frame. The economic lifespan of a well is uncertain and by allowing temporary closure, the option to reactivate remains should prices or technology improve. Some wells have not produced any oil or gas in the last 60 years; nevertheless, this closure is still classified as temporary because the wells have not been permanently decommissioned. The cleanup costs associated with

[7]For example, a blanket bond of $150,000 will cover all wells drilled on federal land in the U.S. This bond amount was set in 1951 and does not reflect actual reclamation costs: between 1988 to 2009, BLM spent around $3.8 million to reclaim 295 orphaned wells [United States Government Accountability Office, 2010]. And in the case of Alberta, there are no bonding requirements.

[8]The regulator does have the authority to order that a wellsite be decommissioned; however, this is not a common occurrence and the order is often rescinded or amended. For example, in 2007 there were only 6 well abandonment orders and in 2006 there were 19 well abandonment orders, but as of June 2009 only 2 of these wells had been abandoned [Alberta Energy Regulator, 2013a].

decommissioning in Alberta range from $20,000 to several million dollars per well [Orphan Well Association, 2008], whereas the cost that a producer must pay to keep a well inactive is usually only the payment to the owner of the surface rights.[9] The sheer volume of wells that have been drilled (over 2.5 million in the United States[10] and in Alberta over 225,000 that will eventually need to be decommissioned) make examining the factors influencing the decision to environmentally remediate a worthwhile endeavor.

An operator might incur losses to maintain an inactive well when it is not currently profitable to produce oil or gas in the hopes that prices or technology improve. According to the data on reserves used in this paper, the percent of hydrocarbon in place that is recoverable (i.e., the recovery rate) ranges from .01% to 90% for oil and from 15% to 95% for gas. However, once a well is inactive, because of the sunk cost to reactivate or decommission, even if recovery rates improve or diminish, the well may remain inactive. This hysteresis is directly modeled in this paper, but there are other reasons for inactivity that are not explicitly modeled, but enter via an error term that compensates for unobservable states: (1) technical difficulties (for example, blockage in the wellbore, a leak caused by corrosion or erosion, an external fire, or a temperature change causing mechanical failure), (2) pipeline failure or pipeline capacity reached, (3) gas plant capacity reached, or (4) a mandated suspension for exceeding the maximum rate limit assigned to the well by the regulator.[11]

The development of enhanced recovery methods, including hydraulic fracturing and horizontal drilling, has brought wells back into production after many years of inactivity and

[9] The annual payment for a wellsite is based on "loss of use" and "adverse effects" only (not land value or entry fee which is paid one time in the first year). Compensation must be paid until the mineral rights owner has received a reclamation certificate. The annual payments range between $167 and $600 per acre for loss of land use and between $117 to $2,500 per acre for general disturbance [Alberta Agricultural and Rural Development, 2010].

[10] U.S. Energy Information Administration's U.S. Crude Oil, Natural Gas, and Dry Exploratory and Developmental Wells Drilled:
http://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET&s=E_ERTW0_XWC0_NUS_C&f=A

[11] Various wells must conform to maximum rate limitations set by the industry regulator. These limits are to ensure that the cumulative amount of oil or gas extracted is maximized. In this model I do not truncate by the maximum rate limit because only about 10% of the wells have limits placed on them, and for only a portion of those wells is the rate limit binding.

recoverable reserves have been seen to increase rather than decrease with time. Reserve growth was first examined by Arrington [1960] using his own company's reservoir data. And since then reserve growth has been studied using state or state subdivision estimates of initial established reserves from the American Petroleum Institute [Morehouse, 1997] or a small number of pools [Verma and Henry, 2004]. This is the first time that such a large dataset on reserves has been used to study reserve growth.

2.1 Environmental Impacts

The fact that technological advances can increase recoverable reserves over time increases the value of waiting to decommission a well. However, without proper decommissioning (and in some cases, even after proper decommissioning) a well poses a risk to vegetation, soil, surface water, and underground aquifers. Many wellbores extend thousands of meters underground, and it is often only a steel casing or cement that isolates the different formations. The casing might rust out or crack (especially when sand or salt water is lifted along with the hydrocarbons), and contaminants such as uranium, lead, salt, iron, selenium, sulfates, and radon [Kubichek et al., 1997] may enter into formations that contain fresh water. The likelihood of this occurring increases when injection from disposal or enhanced recovery builds pressure [Canter et al., 1987]. The most prevalent contaminant, methane, poses the risk of explosion if migration accumulates in adjacent buildings (for a list of explosions in Pennsylvania, some leading to fatalities, attributable to fugitive methane see Pennsylvania Department of Environmental Protection [2009]). Decommissioning a well does not guarantee that there will not be any leaks, but the risk is much reduced as compared to active or inactive wells. Unplugged wellbores would also prevent the future use of a reservoir for carbon capture and storage [Watson and Bachu, 2009]. Furthermore, by not reclaiming the land there is not only the opportunity cost of an alternative land use, but unreclaimed well pads and related infrastructure can contribute to habitat fragmentation [Schneider et al., 2010]. And finally, revegetating may also reduce any increased turbidity in downstream rivers and streams caused by wellsites [Olmstead et al., 2013].

In Alberta, it is required that all wells eventually be decommissioned; however, it is, in effect, left up to the firm to decide when to decommission. Decommissioning a well entails plugging and abandoning and reclamation. Plugging and abandoning refers to leaving the wellbore in a permanently safe and stable condition so that it can be left indefinitely without damaging the environment. It is required that all non-saline water formations are shutoff with cement [Alberta Energy Regulator, 2010b]. Reclamation includes removal of any structures, decontamination of land or water, and reconstruction of the land [Alberta Environment, 2000]. Because the risk of environmental contamination is lower after plugging and abandoning, a company's liability is lower, however it is not completely eliminated. Companies retain ongoing responsibility for wellbore integrity after plugging and abandoning a well [Province of Alberta, 2000].

The Alberta Energy Regulator estimates abandonment costs range from $9,067 to $84,659 and land reclamation costs from $13,200 to $33,700 [Alberta Energy Regulator, 2013b] however abandonment and reclamation costs can dramatically surpass these figures. For example, the Orphan Well Association spent over $2 million to re-enter and repair one orphan well [Orphan Well Association, 2008] and has spent on average $23,000 per site in reclamation costs [Orphan Well Association, 2013]. While plugging and abandoning a well might prevent litigation or remediation costs associated with fluid or gas leakage [National Petroleum Council, 2011], it is difficult to quantify the cost of the ongoing environmental damages from unplugged, unreclaimed wells. The externalities associated with these wells depends on the well location (for example, if the well intersects the range of the woodland caribou [Schneider et al., 2010] or is near any houses[12]), whether there is any groundwater contamination, or fugitive emissions of methane (a potent greenhouse gas) to the atmosphere. However, methane leakage is poorly quantified [Alvarez et al., 2012] as are the existence values of threatened species. Litigation for groundwater contamination[13] provide insights into the cost of groundwater contamination in the worst cases, however, it is difficult to value the externalities from a given inactive well. Therefore this paper proceeds under the assumption that the regulation requiring decommissioning is in place because the averted environmental impacts outweigh the costs of decommissioning.

[12] "Leaky Calmar well forces demolition of homes," CBC News, December 6, 2010.

[13] "Calgary judge hears $33M lawsuit over natural gas drilling," CBC News, January 18, 2013.

3 Data

The data collected on the oil and gas industry in Alberta are unrivaled in their comprehensiveness and accessibility. Here, five datasets of the Albertan oil and gas industry are used. The first dataset is a panel of production from the universe of oil and gas wells in Alberta. Obtained through IHS Incorporated, which distributes the records collected by the Alberta Energy Regulator[14] this dataset contains monthly oil and gas production information dating back to 1924, with complete records starting after 1961. There is information on a well's location (latitude and longitude as well as the name of the field and pool it is on[15]), depth, license date, spud date (the day the drill hit the ground), and on-production date, plus the names of the current and original operators (unfortunately there is no information on whether a well switched hands between these operators).

The second dataset is a panel of official reserve estimates of all nonconfidential pools in Alberta from both the Alberta Energy Regulator and the National Energy Board of Canada.[16] The dataset spans 2000 to 2007 and contains 67,142 oil and gas pools, although not observed in every year. The year that the estimate was last reviewed is listed, and therefore the data are extended to years prior to 2000 if the last review date of the pool was before 2000. This dataset contains (1) initial oil or gas in place; (2) recovery factor, which is the fraction of the oil or gas in place that can be extracted "under current technology and present and anticipated economic conditions" [Alberta Energy Regulator, 2008]; (3) initial established reserves, which is equal to the initial oil or gas in place multiplied by the recovery factor; and (4) remaining established reserves, which is the initial established reserves minus the cumulative production and surface loss. Each pool contains information on characteristics of the pools and hydrocarbons in those pools, such as porosity, initial pressure, area, density, temperature, and water saturation among others.

The third dataset is a list of all wells that were permanently decommissioned (plugged

[14] Formerly called the Energy Resources Conservation Board (ERCB).

[15] An oil field is the geographical area that a well is drilled. A field can have multiple pools, but each pool is a distinct reservoir that is confined within impermeable rock or water.

[16] All pools eventually lose their confidential status (usually after one year), and so this dataset contains nearly all pools in Alberta.

and abandoned and reclaimed). To decommission a well entails that the well has met abandonment standards set by the Alberta Energy Regulator [Alberta Energy Regulator, 2010b] and reclamation standards set by Alberta Environment [Alberta Environment, 1995], and received a reclamation certificate from Alberta Environment or Alberta Sustainable Resource Development or was exempted from certification. The dataset contains both wells that were abandoned, along with the date of abandonment, and the wells that received a reclamation certificate or were reclamation exempt.

The fourth dataset consists of GIS shape files that designate areas that, according to the Petroleum Services Association of Canada (PSAC), have similar costs in production and drilling (the areas are further described in the Appendix Figure 3 and Table 6). The PSAC boundaries and well locations were entered into ArcView GIS to assign a PSAC area to each well.

The final dataset is the average wellhead price of crude oil and natural gas in Alberta, obtained from the Canadian Association of Petroleum Producers' Statistical Handbook [CAPP]. The wellhead price is inflated to 2007 dollars using Statistics Canada's quarterly machinery and equipment price index for mining, quarries, and oil wells.

A panel is created where each well is classified as active, inactive, or decommissioned for each year from when it was drilled until 2007. A well is classified as active if it produced any volume of oil or gas within that year; classified as inactive if it did not produce oil or gas in 12 months or more; and classified as decommissioned if it appeared in the dataset of decommissioned wells.

The full dataset of the universe of wells in Alberta is pared down to a subsample that is used for the estimation. Excluding coalbed methane, heavy oil, injection, and water wells, there are 350,457 wells in the production dataset. The decision to decommission, stop production, or reactivate a well depends on the remaining oil and gas reserves, and so the full sample is restricted to only those wells that have a reserve estimate. Of the 350,457 wells, 105,207 are in a pool that is listed in the reserves dataset. The result of this restriction is that the analysis corresponds to wells that are, or once were, deemed producible, i.e. not "dry holes." Wells that are drilled but do not tap into an oil or gas pool are more likely to be decommissioned without being completed, and they will also not show up in the subsample. More than 45% of the wells that are decommissioned in Alberta are decommissioned immediately after being drilled. The results from an estimation using the subsample, cannot be generalized to all wells in the full sample, but could be generalized to wells in the full sample that at one time produced.[17] Whether to complete a well for production is a separate decision from whether to produce from an already completed well. And indeed, it is more challenging to determine the future of wells that have, or once had, a potential for production as opposed to those that definitely cannot produce.

The subsample is further reduced by deleting wells that traverse both oil and gas pools. Doing so does not significantly reduce the size of the subsample (from 105,207 to 94,009); however, it does significantly reduce the computational complexity because modeling the choice to produce oil or gas is avoided without losing much insight into the choice of operating state.

The majority of the wells have small reserves and only a few have large reserves, some being extremely large–for example, the largest gas reserve is 1,500 times larger than the mean gas reserve (Table 1). The pools with large reserves have more than one well–as many as 4,151 wells in a gas pool and 711 in an oil pool.

[17]The similarity between age at decommissioning for wells in full sample that produced and all wells in the subsample are shown in Appendix Figure 7.

Table 1: Summary Statistics

Variable	No. of Obs.	Mean	Std. Dev.	Min	Max	Unit
$\overline{Q}_{\text{gas}}$	118187	15.23	62.95	0	8800	E^6m^3
$\overline{Q}_{\text{oil}}$	54523	27.09	237.07	0	43871	$E^3 Barrels$
Wellhead Price$_{\text{gas}}$	37	111.90	64.76	25.89	293.91	2007C\$ $/E^3m^3$
Wellhead Price$_{\text{oil}}$	37	30.16	12.12	12.88	64.45	2007C\$ $/Barrel$
Age$_{\text{gas}}$	61876	19.86	15.76	1	104	Years
Age$_{\text{oil}}$	31430	16.58	12.15	1	94	Years
Q_{gas}	118187	32.97	297.79	0	51271	E^6m^3
Q_{oil}	54523	135.64	1026.83	0	104866	E^3Barrels
No. of Wells in Pool$_{\text{gas}}$	118187	3.59	47.02	1	4117	Wells
No. of Wells in Pool$_{\text{oil}}$	54523	4.44	17.91	1	699	Wells
q_{gas}	322907	1.68	8.27	0.0001	568.39	E^6m^3
q_{oil}	155773	7.27	14.51	.001	822.95	$E^3 Barrels$
Depth	93239	1197.53	690.36	90.9	6552	m
Porosity$_{\text{gas}}$	22452	.20	.08	0.01	0.4	Fraction
Porosity$_{\text{oil}}$	25894	.16	.07	0.01	0.36	Fraction
Density$_{\text{gas}}$	22452	.64	.08	0.54	2.03	kg/m^3
Density$_{\text{oil}}$	25894	868.64	48.01	708	999	kg/m^3
Initial Pressure$_{\text{gas}}$	22452	9038.08	7564.57	130	99625	kPa
Initial Pressure$_{\text{oil}}$	25894	12568.80	5688.37	1442	61097	kPa
Temperature$_{\text{oil}}$	25894	50.14	20.34	9	350	C°
Water Saturation$_{\text{oil}}$	25894	.31	.12	0.06	0.82	Fraction
Wells per Firm	1196	281.93	2015.26	1	44095	Wells
Pool Discovery Year$_{\text{gas}}$	22452	1989.67	13.90	1904	2007	Year
Pool Discovery Year$_{\text{oil}}$	25894	1988.31	12.40	1910	2006	Year
Area of Pool$_{\text{gas}}$	22452	854.43	9908.34	1	598512	Acres
Area of Pool$_{\text{oil}}$	25894	183.94	554.36	1	17890	Acres
Duration Inactive$_{\text{oil}}$	9556	8.39	8.33	0	73	Years
Duration Inactive$_{\text{gas}}$	12298	9.58	10.23	0	78	Years
Duration Active$_{\text{oil}}$	14472	10.11	9.11	0	46	Years
Duration Active$_{\text{gas}}$	34047	11.35	11.85	0	46	Years

Notes: Statistics for wells in the subsample. Data on remaining reserves (Q) are listed for pools, 1993 - 2007. Extraction (q) is listed for wells, 1993 - 2007. The pool-specific variables– depth, porosity, density, initial pressure, temperature, water saturation, and discovery year are time invariant in the data. Data on the age of the wells, and duration active and inactive are a snapshot of 2007. Price data are the wellhead price from 1971 to 2007. $E^3 = 1000$.

The production dataset contains firm-reported volumes to which the accuracy is difficult to attest. The Alberta Energy Regulator identifies cases when there is *any* difference in the reported production of oil from a production company and a pipeline company. When the difference is 5% to 20% of reported gas volumes, the penalty is only a warning message. Further misreporting results in a fee of \$100 if a well does not report in a given month, and upon persistent noncompliance the firm might be subject to increased audits or inspections, or partial or full suspension [Alberta Energy Regulator, 2010a]. Nonetheless, to the best of my knowledge there is no other dataset of this size or comprehensiveness of any natural resource industry. And with these data the composition of active, inactive, and decommissioned wells can be replicated to match reality closely.

Table 2: Distribution of operating choice for inactive wells by age

Age (in years)	Number of Observations		Proportion Reactivated		Proportion Stay Inactive		Proportion Decommissioned	
	(Oil)	(Gas)	(Oil)	(Gas)	(Oil)	(Gas)	(Oil)	(Gas)
$1 \leq age < 10$	18963	21091	0.110	0.194	0.874	0.790	0.016	0.015
$10 \leq age < 20$	19892	14965	0.056	0.066	0.922	0.915	0.022	0.019
$20 \leq age < 30$	7997	11234	0.057	0.075	0.914	0.897	0.029	0.028
$30 \leq age < 40$	2340	4791	0.046	0.054	0.928	0.913	0.026	0.033
$40 \leq age < 50$	2176	3135	0.030	0.042	0.945	0.929	0.025	0.030
$50 \leq age < 60$	704	1461	0.024	0.027	0.953	0.955	0.023	0.018
$60 \leq age < 70$	131	503	0	0.010	1	0.990	0	0
$age \geq 70$	50	157	0	0	1	0.987	0	0.013

Notes: Data from 2000-2007 subsample.

Table 2 shows the proportion of inactive oil and gas wells that have been reactivated, left inactive, or decommissioned by different age intervals. The table illustrates that the hysteresis of inactivity increases as wells age. The proportion of inactive wells that are reactivated decreases with the age of the well and the proportion of inactive wells that are decommissioned increases then decreases with age.

4 Model

In order to capture the value of leaving a well inactive, I construct a real options model that includes the following features: the operating state is dynamic and can be changed now, or at some later date; there are unrecoverable sunk costs to changing operating states; and future prices and recovery are uncertain. The producer's decision to extract, 1; temporarily stop extraction, 2; or permanently decommission and remediate environmental damages, 3, is modeled as an infinite time Markov Decision Process [Rust, 1994]. It is assumed that the producer is rational and follows a decision rule, $d_t = \delta_t(s_t, \epsilon_t)_{t=0}^{\infty}$, that maximizes the expected discounted sum of profits, $V(s, \epsilon) = \max_\delta \mathrm{E}_\delta \left[\sum_{t=0}^{\infty} \beta^t \pi(s_t, d_t, \epsilon_t | s_0 = s, \epsilon_0 = \epsilon)\right]$, where V is the value function for the well when choosing the optimal choice, δ, and depends on observed state variables s, and an unobserved random "payoff shock," ϵ, different for each choice. The instantaneous profit, $\pi(\cdot)$, is discounted by discount factor β, $0 \leq \beta \leq 1$. The observed state variables, s, include the age of the well, A, the wellhead price of the hydrocarbon, P, the per-well remaining reserves, $\overline{Q}$, and the current operating state, o. The current operating state (o = active, 1; inactive, 2; or decommissioned, 3) is endogenous to the decision, and the remaining reserves per well are endogenous (when the operator extracts oil or gas) and exogenous (upon technology change or if another well is in the same pool).

The model assumes that the producer maximizes lifetime profits only through the extensive decision for the operating state, but not through the intensive decision for how much to extract.[18] If the producer decides to extract, the per-period quantity recovered is a random draw from a distribution that depends on the remaining per-well reserves, the age of the well, and parameters, α, estimated in a separate estimation outside the dynamic programming model (i.e., the expected production is $\mathrm{E}q = \int_0^Q q f_q(y|\overline{Q}, A, \alpha)dy$). The profit (equation (1)) if the producer decides to produce is equal to the expected quantity recovered, $\mathrm{E}q$, times the price of the hydrocarbon, P, less the per unit extracted royalty rate (or severance tax), R, and a per-unit lifting cost to extract, C.[19] The per-unit lifting costs also depend on age, per well-reserves, and parameters, θ, to be estimated, $C = C_g(\overline{Q}, A, \theta)$.[20] The royalty rate in Alberta adjusts according to price and quantity produced, $R = R(P, q)$. This profit is then reduced by the corporate income tax, τ, assumed flat for all wells. If the current state of the well is inactive, there is a switching cost to activate, $SC_{(2\to1)}$.

[18] The implication of not simultaneously modeling the decision of whether to extract and the decision of how much to extract is that there will be a disconnect between the lifetime profit-maximized quantity and the expected quantity I estimate outside of the model, which in turn could bias the predicted duration in a given operating state. However, the assumption rests on the assertion that extraction is mainly driven by reserve size and geologic factors which producers do not have full control over (further discussed in Section 5.1.1).

[19] Note that I do not include a fixed cost of extraction. When I include a fixed extraction cost, the estimated parameters converge to the lower and upper bounds specified for the parameters. Presumably this is because fixed costs are difficult to identify (as further discussed in Section 5.2).

[20] Chermak and Patrick [1995] and Foss et al. [2002] show how the lifting cost of natural gas depends on quantity extracted and remaining reserves. Chermak and Patrick [1995] use data from 29 gas wells in Wyoming and Texas from 1988 to 1990, and Foss et al. [2002] use data from 22 gas wells in Alberta for roughly three years. They both find that operating costs increase with quantity extracted and decrease with remaining reserves. It is expected that extraction costs rise as reserves are depleted; however, Livernois and Uhler [1987] explain that the discovery of new reserves can increase the reserves by more than what is extracted, but these new reserves are more costly to extract. This is how Livernois and Uhler [1987] explain a positive relationship between extraction costs and reserves using aggregate data from the Albertan oil industry. However, upon disaggregation, they find the typical results of extraction costs increasing with reserve depletion and quantity extracted.

If the producer instead chooses that the well be inactive, the producer pays a annual inactivity cost, M, and if the current state of the well is active, a switching cost $SC_{(1\to 2)}$.

To decommission a well is to enter an absorbing state for which the producer pays a one-time switching cost, $SC_{(1,2\to 3)}$, assumed to be the same for active and inactive wells.[21] Leaving the well in its current state entails no switching costs, $SC_{(1\to 1)} = 0$, $SC_{(2\to 2)} = 0$, and $SC_{(3\to 3)} = 0$. The expected profit from a single period is:

$$\pi(s,d,\epsilon) = \begin{cases} ((1-R)P - C)\,\mathrm{E}q - \tau \max\{((1-R)P - C)\,\mathrm{E}q, 0\} - SC_{(o\to 1)} + \epsilon_1 & \text{if d=1} \\ -M - SC_{(o\to 2)} + \epsilon_2 & \text{if d=2} \\ -SC_{(o\to 3)} + \epsilon_3 & \text{if d=3} \end{cases} \tag{1}$$

The expected present discounted value of the well can be expressed as the unique solution to the Bellman equation[22]:

$$V(s,\epsilon) = \max_d [\pi(s,d,\epsilon) + \beta \int_{s'} \int_{\epsilon'} V(s',\epsilon') h(s',\epsilon'|s,\epsilon,d) \mathrm{d}\epsilon' \mathrm{d}s']$$

The state transition probability density function, $h(s',\epsilon'|s,\epsilon,d)$, is assumed to be a Markov process. The Conditional Independence assumption, as per Rust [1987, 1988], is adopted, allowing for the factorization, $h(s',\epsilon'|s,\epsilon,d) = f(s'|s,d)\rho(\epsilon'|s')$. Specifically in the case of this model, the state transition probability density function can be written:

$$h(\overline{Q}', P', A', \epsilon'|\overline{Q}, P, A, \epsilon, d) = f_q(\overline{Q} - \overline{Q}'|\overline{Q}, A, d) f_{\overline{Q}}(\overline{Q}'|\overline{Q}, P) f_P(P'|P) f_A(A'|A) \rho(\epsilon'|\overline{Q}', P')$$

Price is assumed to follow the exogenous process $f_P(P'|P, \varsigma)$, characterized by parameters ς. Recoverable reserves decrease from extraction but also increase or decrease from new discoveries, revisions, or technological change and follow $f_{\overline{Q}}(\overline{Q}'|\overline{Q}, P, \phi)$. The quantity extracted, q, is modeled as a random draw from the density, $f_q(q|\overline{Q}, A, \alpha)$. The payoff shocks follow the transition probability $\rho(\epsilon'|s')$, but I assume that $\rho(\epsilon'|s')$ is independent of s so the payoff shocks are assumed to be independent and identically distributed (i.i.d.) across

[21] It is not necessary to make decommissioning an absorbing state, but in the sample there are only 261 observations of a switch from decommissioned to active, whereas there were 22,308 observed deactivations, 15,369 reactivations of inactive wells, 1,917 active wells decommissioned, and 3,664 inactive wells decommissioned

[22] Following Blackwell's theorem (outlined in Rust [1994], Theorem 2.3).

choices, state variables, and time. They capture unobserved, current period events that make a choice more or less expensive for the operator. For example, the shocks would capture a blowout that forces the operator to shut in an active well or capture an unexpected surplus of available workover rigs that make it cheaper to reactivate or decommission a well. While it is possible to allow serial dependence in the process $\{\epsilon\}$ (see Norets [2009]), it is computationally burdensome and as a first approximation, the main dependencies are captured by the serial dependence in the observed state variables.

This paper abstracts away from modeling the decision to decommission as a function of the strategic interaction between agents competing for the same oil. The common pool resource problem would arise if more than one company competes for a migratory hydrocarbon [Libecap and Wiggins, 1984, 1985]. However, each oil and gas pool in the sample is a distinct reservoir that is confined within impermeable rock or water and so only within pools would we expect to see behavior influenced by the common-pool. For the majority of the pools in this paper only one firm has access to the pool and the number of wells in a pool is small (Table 1). The average number of wells in a gas pool is 3.5 (and 4.4 for oil pools), but the majority of the time there is only 1 well per pool (the median and mode are 1 for both oil and gas). This may be driven by the fact that 81% of the mineral rights in Alberta are owned by the Crown, and then leased to companies, whereas in the U.S., mineral rights are mainly determined by surface landownership resulting in more fragmentation of mineral rights.[23] Nonetheless, observations of wells that are in pools that have no other wells are modeled and estimated separately from wells that are in pools with other wells. Wells in single-well pools and multi-well pools are modeled differently through their transition probability of recoverable reserves from extraction, f_q. For wells on single-well pools, only when the decision is to extract, $d = 1$, is the transition probability of reserves dependent on the probability of how much can be extracted, $f_q(\overline{Q} - \overline{Q}' | \overline{Q}, A, \alpha, d)$, while for wells on multi-well pools, reserves transition according to this probability whether the operator extracts or not,

[23] In Alberta, most leases are for one quarter section (160 acres). Only one oil well can be drilled on a quarter section, and only one gas well on one section (a company must obtain the mineral leases for all four quarters of the section).

$f_q(\overline{Q} - \overline{Q}'|\overline{Q}, A, \alpha)$. Exogenous to whether the well is active or not, recoverable reserves also follow another process, $f_{\overline{Q}}(\overline{Q}'|\overline{Q}, P, \phi_g)$, that accounts for the probability of change from improved technology, discoveries, reassessment, and additions. The per-well reserves, $\overline{Q}$, also decrease whenever another well is drilled in the pool. It is assumed that the number of new wells drilled is an exogenously determined random shock.[24] The probability of a decrease in per-well reserves by another well being drilled is incorporated into the exogenous change dictated by the transition probability density $f_{\overline{Q}}$. Although there may be strategic interactions between different firms extracting from the same pool, modeling these in the probability of reserve decreases of multi-well pools is beyond the scope of this paper. A justification for this simplification is that a large fraction of the multi-well pools are operated by a single firm[25], alleviating much of the concern of strategic extraction driven by the common-pool. However, it could still be the case that firms are making joint decisions for all the wells in the same pool, yet the model is treating these as separate decisions. Given that not all wells in the same pool have the same operating state, this is not a very large concern, however if joint decisions were being made it would be because of increasing returns to scale in which case I would be under-estimating the per-well costs.

Assuming that ϵ is drawn from the type I extreme value distribution, the Bellman equation becomes [Rust, 1988]:

$$V_\theta(s, \epsilon) = \max_d [v_\theta(s, d) + b\epsilon(d)]$$

where θ are the parameters to be estimated for each well group (including the cost parameters in the profit equation and the parameters in the transition probability density functions) and v_θ is the fixed point of $v_\theta = \Gamma(v_\theta)$, where Γ_θ is a contraction mapping:

$$\Gamma_\theta(v)(s, d) = \pi(s, d, \theta) + \beta \int_{s'} b \log \sum_{d'=1}^{3} \left[\exp\left\{\frac{v_\theta(s', d')}{b}\right\}\right] f(s'|s, d)\mathrm{d}s' \tag{2}$$

[24] For a model of where to drill for oil and gas, see Levitt [2009].

[25] On average there are 1.8 (1.9) firms per oil (gas) pool but 54.1% (49.9%) of the oil (gas) pools only have one firm and pools that have more than one firm are larger (2.6 (4.2) times for oil (gas)).

with location parameter of the extreme value distribution of ϵ normalized to zero and scale parameter, b, to be estimated.[26]

The assumption of the extreme value distribution[27] allows for a closed form solution of the choice probabilities–that of the multinomial logit:

$$p(d|s,\theta) = \frac{\exp \frac{v_\theta(s,d)}{b}}{\sum_{d'} \exp \frac{v_\theta(s,d')}{b}} \tag{3}$$

5 Estimation

The estimation consists of three stages. First, I estimate the parameters of the producer's subjective belief for how the state variables progress over time as a standard parametric estimation. Second, the parameter estimates from the first stage are taken as given, and the remaining parameters in the Bellman equation, the costs in the profit function, are estimated via the Nested Fixed Point Algorithm [Rust, 1987]. Nested within the algorithm to maximize the likelihood function of the choice probabilities (equation (3)), there is an inner algorithm to compute the fixed point, v_θ, of equation (2). The outer loop of the algorithm, the maximization of the likelihood, was submitted to the solver KNITRO [Byrd et al., 2006]. The inner loop, which solves the fixed point of equation (2), consists of successive approximations followed by Newton-Kantorovich iterations. The third step is to obtain consistent standard errors from the full likelihood function. The parameter values from the first stage contain a measurement error, but they are treated as the true parameters in the second stage, and so the standard errors for the second stage parameters are inconsistent. To obtain consistent standard errors, the consistent parameter values from the first and second stages are used as starting points for one Gauss-Newton step of the full likelihood function [Rust, 1994].

Well-level heterogeneity is accounted for by estimating the dynamic programming model separately for different well-types, g. All wells of the same type are treated as homogeneous,

[26] A location parameter of zero means that ϵ has a mean zero; a scale parameter of one would mean ϵ has a variance of $\pi^2/6$. As b approaches zero, $V_\theta(s,\epsilon)$ converges to the ordinary Bellman equation.

[27] Dagsvik [1995] showed that the generalized extreme value class is dense; choice probabilities from any distribution can be approximated arbitrarily closely by choice probabilities from the generalized extreme value class.

and wells of the same type that also have the same reserve size and the same age are assumed identical. The well types are determined by (1) whether the well is an oil or gas well, (2) whether the well is in a single-well pool or a multi-well pool, (3) the royalty regime applicable, (4) PSAC area, and within these groups, (5) clusters based on time invariant characteristics (depth, initial pressure, density, water saturation, and temperature). The group is divided into clusters only if the likelihood ratio test confirms that clustering improves the fit over not clustering. This results in 88 different types of wells. The royalty regime depends on when the pool was discovered: there is an "old" category for oil from pools discovered before 1974, "new" for oil from pools discovered between 1974 and 1992, and "third tier" for oil from pools discovered after 1992. For gas wells, "old" refers to gas from pools discovered before 1974 and "new" to gas from pools discovered after 1974. Within each type, the royalty depends on price and the quantity extracted. The royalty regime remained the same from 1993 to 2009 [Province of Alberta, 2008], coinciding with the study period.

5.1 First Stage Estimates

The producer's beliefs about future prices and recoverable reserves are estimated in the first stage. These beliefs are unobservable and subjective, but here I assume that the producer's beliefs are recoverable from objective probability measures estimated from the data. I estimate the parameters, $\theta_{1^{st}}$, that maximize the first stage partial likelihood function:

$$L_1(\theta_{1^{st}}) = \prod_{i=1}^{N_i} \prod_{t=1}^{T_i} f(s^i_{t+1} | s^i_t, \theta_{1^{st}}) \tag{4}$$

These parameters and their transition probability densities are described in the next three subsections.

5.1.1 Transition in Remaining Reserves from Extraction, f_q

Reserve changes due to extraction are such that when the well is active, the quantity extracted is modeled as a random draw from a distribution that depends on the per-well remaining reserves, $\overline{Q}$ and age of the well. It implies that the producer only chooses whether to extract and does not have control over the quantity extracted. Although producers do have control over extraction rates, ultimately extraction is also driven by geological constraints that the producer does not have control over. In a regression of the annual quantity extracted on factors that the producer does not have control over, I find that the remaining reserves, porosity, temperature, depth, density, water saturation, and initial pressure are all statistically significant in determining extraction. Furthermore, these exogenous factors explain more of the variation in extraction quantities than the price of oil or gas.[28]

Therefore, the quantity extracted is modeled as a random draw from a distribution estimated using an equation that describes the production by well w in year t:

$$\log q_{wt} = \alpha_0 + \alpha_1 \log \overline{Q}_{wt} + \sigma \varepsilon_{wt} \tag{5}$$

assuming an independent and identically distributed $N(0,1)$ error, ε. The regression is estimated separately for each well type-age group combination.[29] Extraction from a well is truncated to fall in the interval $[q^L, q^U]$ where the lower bound, q^L, is 10^{-8} (not zero because of the subsequent logarithm), and the upper bound, q^U, is equal to the well's per-well remaining reserves, $\overline{Q}$, multiplied by a factor, κ_m, which depends on whether the well is in a single-well pool, $m = 0$, or a multi-well pool, $m = 1$. In the dataset there are a few observations where the amount produced in a year is greater than the per-well remaining reserves even for wells that are on their own pools (6% of the production data would be classified as such). Evidently the reserve size is sometimes an underestimate. Therefore, the factor κ_m is equal to the 99th percentile of the observed fractions $q_w/\overline{Q}_w$ (different for

[28] The adjusted R^2 of a regression of the quantity extracted on a constant and price is much smaller than the adjusted R^2 of a regression that instead of price includes the remaining reserves (specifically, price results in 0.0068 for oil wells and 0.0002 for gas wells whereas remaining reserves result in 0.1275 for oil and 0.3635 for gas). Per-well remaining reserves also depend on price, however the adjusted R^2 from a regression that includes the truly exogenous characteristics of porosity, temperature, depth, density, water saturation, and initial pressure is also larger than the adjusted R^2 when only including price (i.e., 0.0836 for oil and 0.2405 for gas).

[29] If there are less than 30 observations of production within a given type's age group, then observations from the age group without clustering was used.

single-well pools and multi-well pools).[30] The weighted average (across well-types) of the coefficients in equation (5) are displayed in Table 3. As expected, the older the well, the less production is expected from the well, for both oil and gas wells.

5.1.2 Exogenous Transition in Reserves, $f_{\overline{Q}}$

Apart from production there are exogenous changes in reserves from, for example, more wells being drilled, reassessments, improved technology, or new discoveries. Changes in reserves from production are already accounted for by including f_q, so to estimate exogenous changes, $f_{\overline{Q}}$, the estimates of initial established reserves (IER) are used. A pool's IER is an estimate of the initial oil or gas in place multiplied by the recovery factor, and does not include what has been extracted.

The current price of oil or gas might have different effects on the transition probability of remaining reserves. For example, under high prices, one would expect there to be more research and development into extraction technology, which would in turn increase the recoverable reserves. Once developed these technologies would remain available, even under low gas prices, and therefore, one would not expect there to be symmetrical decreases in reserves under low prices, but rather only smaller increases. On the other hand, higher prices would also result in the drilling of more wells, which would in turn reduce the per-well remaining reserves. I therefore estimate the distribution of the size of reserve increases separately from the distribution of the size of reserve decreases. As well as the likelihood of the size of a change, the transition probability, $f_{\overline{Q}}$, also includes the probability to increase, decrease or remain the same. Depending on the type of pool (i.e., PSAC area, single- or multi-well, year discovered), 63% to 84% of the pool-year observations have no change in per-well initial established reserves.

The distribution of the natural logarithm of changes that did occur are depicted in Figure 7 in the Appendix, and can be approximated by two exponential distributions spliced together. When there is an increase in reserves, the size of the increase, $\ln(\overline{Q}_{t+1}^{IER}/\overline{Q}_t^{IER})$, can

[30] The 99th percentile is used because there are a few outliers where q_w dramatically exceeds $\overline{Q}_w$. (That is, the 99th percentile of $q_w/\overline{Q}_w$ for gas wells on single-well pools, k_0, is 2.9 compared to a maximum of 66 and for gas wells on multi-well pools, k_1, is 25.2 compared to a maximum of 807.)

be approximated by an exponential distribution, as when there is a decrease in reserves, and the size of the decrease, $-\ln(\overline{Q}_{t+1}^{IER}/\overline{Q}_t^{IER})$, can be approximated by a different exponential distribution. That is, when there is an increase (or decrease), $\Delta = \left|\ln(\overline{Q}_{t+1}^{IER}/\overline{Q}_t^{IER})\right|$ follows a distribution with density function:

$$f_{\overline{Q}}(\Delta|\lambda) = \lambda \exp(-\lambda\Delta)$$

The current price is incorporated into the probability via λ, where:

$$\lambda = \begin{cases} (\phi_{0U} + \phi_{1U}P)^{-1} & \text{when increase,} \\ (\phi_{0D} + \phi_{1D}/P)^{-1} & \text{when decrease.} \end{cases}$$

Therefore, using any observations of a reserve increase ($\overline{Q}_{it}^{IER} > \overline{Q}_{it-1}^{IER}$), I maximize the likelihood of the size of a increase as:

$$L(\phi_{\mathbf{U}}) = \prod_i \prod_t \frac{1}{(\phi_{U0} + \phi_{U1}P_{t-1})} \exp\left(-\frac{\left|\ln(\overline{Q}_{it}^{IER}/\overline{Q}_{it-1}^{IER})\right|}{(\phi_{U0} + \phi_{U1}P_{t-1})}\right)$$

And likewise, using any observations of a reserve decreases ($\overline{Q}_{it}^{IER} < \overline{Q}_{it-1}^{IER}$), I maximize the likelihood of the size of a decrease as:

$$L(\phi_{\mathbf{D}}) = \prod_i \prod_t \frac{1}{(\phi_{D0} + \phi_{D1}/P_{t-1})} \exp\left(-\frac{\left|\ln(\overline{Q}_{it}^{IER}/\overline{Q}_{it-1}^{IER})\right|}{(\phi_{D0} + \phi_{D1}/P_{t-1})}\right)$$

5.1.3 Transition in Price

Analysis of the price of oil is a well-researched area, although there is little consensus for the best- fitting model. Models differ by allowing for mean reversion, non-stationary unit roots, underlying market fundamentals, unexpected jumps, or time-varying volatility, for example. To limit the number of states included in the model, I assume price follows an exogenous first order Markov process, which is an assumption for the formation of expectations for future prices: the producer bases their expectation on current prices. At the same time, producers have experienced periods of both low and high prices in the past. Therefore, I include a switching process between a high price regime and a low price regime in the Markov process. The regimes are determined depending on whether the price is above or below the average price observed from 1971 to 2007. Because we know which regime we are in simply by knowing whether the price is above or below the average price, there is no need to include "price regime" as a state variable in the model. That is, the final matrix of transition probability weights for price depends only on the current price. The transition probability of switching from regime H to L is simply the number of times that regime H was followed by regime L divided by the number of times the process was in regime H:

$$\hat{p}_{HL} = \frac{\sum_{t=1}^{T} I\{r_t = L, r_{t-1} = H\}}{\sum_{t=1}^{T} I\{r_{t-1} = H\}}$$

While the opposite holds true for switching from L to H. For each regime, the parameters from a regression with deviations from the mean logarithm of price, $\wp_{t,r} = \log P_t - \mu_r$, are estimated:

$$\wp_{t,r} = \vartheta_r \wp_{t-1,r} + \varsigma_r \varepsilon_t \tag{6}$$

where ε is independent and identically distributed $N(0,1)$. The process is truncated so that price does not fall below $\underline{P} = 1E - 6$ (and not zero because of the subsequent logarithm). The transition probability of price in regime r is:

$$F_P(P_t|P_{t-1}, r) = \frac{\Phi\left(\left(\wp_{t,r} - \vartheta_r \wp_{t-1,r}\right)/\varsigma_r\right) - \Phi\left(\left(\underline{\wp}_r - \vartheta_r \wp_{t-1,r}\right)/\varsigma_r\right)}{1 - \Phi\left(\left(\underline{\wp}_r - \vartheta_r \wp_{t-1,r}\right)/\varsigma_r\right)}$$

where Φ is the standard normal cumulative distribution function and $\underline{\wp}_r = \ln(\underline{P}) - \mu_r$.

The transition probability matrix is derived from a mixture of the distributions under high and low price regimes, including the transition probabilities between the two regimes:

$$F_P(P_t|P) = \begin{cases} p_{HH} F_P(P_t|P_{t-1}, r = H) + p_{HL} F_P(P_t|P_{t-1}, r = L) & \text{if } P > \overline{P}, \\ p_{LH} F_P(P_t|P_{t-1}, r = H) + p_{LL} F_P(P_t|P_{t-1}, r = L) & \text{if } P \leq \overline{P}. \end{cases}$$

.

5.1.4 Transition in Age

To save computing time the age variable is discretized into the intervals $A = 1, 5, 15, 30$ for $1 \leq age < 5$, $5 \leq age < 15$, $15 \leq age < 30$, $age \geq 30$, under the assumption that wells within these age intervals are similar. Oil and gas production over the life of a well typically follows an exponential decline, with the steepest decline within the first several years. Therefore, I make smaller intervals for younger wells, when the difference in age matters more.[31] These intervals also divide the sample into roughly four even groups (of 19%, 31%, 29% and 18% of the observations in each age group, respectively). The transition probability of entering the next interval is $1/n_{years}$ where n_{years} is the number of years in the current interval.

5.2 Second Stage Estimation

For each different well type, g, a different set of structural parameters, $\theta_{2^{nd}} = (C,\ M,\ SC_{(2\to1)},\ SC_{(1,2\to3)})$, is estimated. The likelihood of observing the decisions d that were made for each well ($w = 1...W_g$) in the well group is maximized:

$$L(\theta_{2^{nd}}) = \prod_{t=1}^{T} \prod_{w=1}^{W_g} p(d_t^w | P_t, Q_t^w, A_t^w, \theta_{2^{nd}}, \hat{\theta}_{1^{st}})$$

where p is the multinomial logit probability, given in equation (3), that the choice for well w at time t is decision d. For each iteration of the likelihood there is a nested subroutine to find the fixed point to the Bellman equation (2). The model is in discrete time and the producer chooses the operating mode on a yearly basis. In reality, this decision is in continuous time; however, a well is classified as an inactive well by the Alberta Energy Regulator if it has not reported any volumetric activity (production, injection, or disposal) within the last 12 months. Therefore, the data are assigned as follows: for a well in 2000, the current operating state, o, is the operating state in 1999, where the decision, d, is the operating state in 2000, given the average wellhead price of oil (or gas for gas wells) in 2000, the reserve size in 2000, and the age of the well in 2000.

The royalty rate is calculated using formulas specified by the Alberta Department of

[31] Also, when estimating production, q, from a well, the shorter intervals for younger wells helps capture this nonlinearity in production over the life of the well.

Energy [Alberta Energy, 2006]. The rates range from 5% to 35% depending on the price of oil (or gas), when the reserve was discovered, and the volume of oil (or gas) produced. As this model is based on the expected production, and not the actual production, the royalty rate is the expected royalty rate.[32] The Alberta corporate income tax rate is 10% of taxable income while the federal corporate income tax rate is 22.12%. The combined federal and provincial tax rate on corporate income, τ, is set at 32.12% [Alberta Department of Energy, 2007].

Estimating the discount factor, β, along with the cost parameters is difficult. For example, both a high reactivation cost and a low discount factor will prolong reactivation. Therefore, for each well group I estimate the cost parameters under seven candidate discount factors, ranging from .7 to .99. By examining the sum of the log-likelihoods across well groups for each candidate discount factor, a discount factor of .90 results in the highest total log-likelihood (illustrated in Appendix A.3).

The estimation requires specifying the functional form of the profit equation. I estimated the model under many different specifications that seemed reasonable and were flexible to incorporate features such as costs increasing as age increases and remaining reserves decrease. A parsimonious specification that led to timely convergence and high likelihood values is a specification where the lifting cost depends on the reserve size and age, $C = \theta_1 + \theta_2/\overline{Q}^{\theta_3}$; the fixed inactivity cost is a constant, $M = \theta_4$; the reactivation cost depends on age, $SC_{2\to1} = \theta_5(1 + \theta_6)^A$; and deactivation and decommissioning costs are constant, $SC_{1\to2} = \theta_7$ and $SC_{1,2\to3} = \theta_8$. By construction of the multinomial logit (equation (3)), identifying all fixed costs of the model is not possible. To identify the absolute costs external information on actual well sale prices would be needed. The cost parameters are interpreted in million dollars, but also in relation to a normalized reactivation cost of zero for active wells (i.e., $SC_{(1\to1)} = 0$). Furthermore, the scale of the profit equation is already normalized because the coefficient on price is normalized (to one) which also means that I can estimate the scale parameter of the type I extreme value error term, ϵ.

[32] Formulas can be found in Alberta Energy [2006].

5.3 Third Stage Estimation: Maximizing the Full Likelihood

The parameter estimates from the partial likelihood estimation are used as starting values in the maximization of the full likelihood function. I allow for one iteration of the maximization routine of the full likelihood function to determine a consistent estimate of the asymptotic covariance matrix for the estimates, which is used to determine consistent standard errors. The average estimates and standard errors of the 88 different well-types, weighted by the number of observations, are displayed in Table 3. The coefficients were derived using price data scaled by one million dollars. The coefficient on price in the probability distribution for an exogenous increase in reserves, $\phi_{1,U}$, appears to be very different for oil and gas wells; however, because the prices are in millions of dollars, inverted, and taken to the negative power of Euler's number, the difference in the probability of an increase is not as large as these coefficients suggest. Nonetheless, gas reserve growth is more responsive to price changes than oil reserve growth, which is most likely because gas reserves are more difficult to estimate than oil reserves [Vanorsdale, 1987].[33]

It is important to note that the estimates in Table 3 do not represent the costs for switches that were actually made, but are the costs for a hypothetical well to switch to an arbitrary operating state [Kennan and Walker, 2011]. Operators only choose to switch when the costs net of the payoff shocks are favorable, which would be less than the estimates in Table 3.

I find that on average the costs to decommission a well are higher than the costs to reactivate a well, and so one might expect that wells are more likely to be reactivated than decommissioned. However, the lifting costs increase as the remaining reserves decrease, and

[33] Gas reserves are not only made up of "free" gas, but also the more difficult to measure "adsorbed" gas attached to the rock surface. There are other less probable reasons for differences in reserve-growth elasticities. For example, by construction of the model, newly drilled wells reduce per-well recoverable reserves, which combined with four times stricter well-spacing limits for gas wells (see footnote (23)) would reduce the probability of a reserve decrease for gas compared to oil. However, gas pools are on average four times larger than oil pools, so this is not likely to be the case. It is also not likely that the difference in reserve-growth elasticities is driven by a difference in the age of pools. Gas pools are slightly younger (84% of the gas pools were discovered after 1974 compared to 82% for oil) however not significantly.

therefore, reactivating the well for a lower range of $\overline{Q}$ would result in paying more to extract than the price of gas or oil.

Firms could be postponing decommissioning because they intend to declare bankruptcy in the future, however, the number of orphan wells is relatively small in Alberta implying that bankruptcy is not a very important option.[34] In other empirical contexts in which bankruptcy is an important option, the model would have to take the full portfolio of wells that a firm owns into consideration (e.g., a firm would not declare bankruptcy if they have many producing wells). In this case, to the extent that bankruptcy is an option, I am omitting an option that is available to active and inactive wells (i.e., the option to declare bankruptcy, at some cost, and avoid decommissioning costs) and by omitting this option, the estimated costs of an inactive or active well are biased downwards and the decommissioning costs biased upwards as compared to a model that incorporates this option. However, more important than interpreting the cost estimates is correctly simulating counterfactual scenarios, in which case the biased cost parameters to some extent account for the omitted bankruptcy option.

[34]For example, in 2012, a year noted for having a "large increase in the number of new orphan wells," 50 new wells were added to the inventory of 14 orphan wells [Orphan Well Association, 2013].

Table 3: Weighted Average Parameter Estimates from the Full Likelihood

Parameters	Oil		Gas	
	Estimate	Std.Err.	Estimate	Std.Err.
Reserves Transition				
$\alpha_{0,1}$	0.312	(0.596)	3.333	(2.324)
$\alpha_{0,5}$	0.3	(0.581)	2.522	(1.797)
$\alpha_{0,15}$	0.07	(0.430)	2.454	(1.036)
$\alpha_{0,30}$	-0.103	(0.411)	2.288	(1.369)
$\alpha_{1,1}$	0.43	(0.163)	0.438	(0.263)
$\alpha_{1,5}$	0.274	(0.147)	0.47	(0.194)
$\alpha_{1,15}$	0.317	(0.128)	0.492	(0.115)
$\alpha_{1,30}$	0.404	(0.161)	0.51	(0.179)
σ_1	1.288	(0.178)	1.501	(0.100)
σ_5	1.267	(0.078)	1.52	(0.132)
σ_{15}	1.273	(0.077)	1.504	(0.231)
σ_{30}	1.171	(0.242)	1.344	(0.302)
$\phi_{0,U}$	0.373	(0.026)	0.268	(0.071)
$\phi_{1,U}$	2.00E-07	(3.81E-01)	571.55	(371.032)
$\phi_{0,D}$	0.69	(0.082)	0.432	(0.040)
$\phi_{1,D}$	1.00E-08	(3.13E-03)	6.56E-05	(7.56E-06)
Price Transition				
ϑ_L	0.427	(0.972)	0.707	(0.747)
ς_L	0.155	(0.477)	0.17	(0.357)
ϑ_H	0.603	(0.498)	0.594	(0.660)
ς_H	0.145	(0.206)	0.213	(0.378)
Lifting Cost (C)				
θ_1	0.032	(0.066)	1.64E-04	(7.76E-05)
θ_2	0.046	(0.202)	0.147	(0.286)
θ_3	1.588	(1.229)	3.775	(1.066)
Inactivity Cost (M)				
θ_4	0.26	(0.259)	0.386	(0.415)
Cost to Reactivate ($SC_{(2\to1)}$)				
θ_5	4.86	(1.131)	5.957	(2.629)
θ_6	0.087	(0.095)	0.024	(0.019)
Cost to Temporarily Deactivate ($SC_{(1\to2)}$)				
θ_7	1.774	(2.267)	2.594	(1.311)
Cost to Decommission ($SC_{(1,2\to3)}$)				
θ_8	7.923	(3.673)	9.703	(1.333)
Scale Parameter				
b	1.134	(0.376)	1.346	(0.300)
Not Estimated in Likelihood				
p_{HL}	0.272		0.062	
p_{LH}	0.160		0.100	
$\overline{P}$	0.031		1.143e-4	
μ_H	-3.203		-8.784	
μ_L	-3.772		-9.726	
β	.90		.90	

Notes: These are the weighted averages of the estimates across 88 well groups. Using specification $C = \theta_1 + \theta_2/\overline{Q}^{\theta_3}$; $M = \theta_4$; $SC_{2\to1} = \theta_5(1+\theta_6)^A$; $SC_{1\to2} = \theta_7$ $SC_{1,2\to3} = \theta_8$. The standard errors are derived from the White [1982] misspecification consistent information matrix.

5.4 Goodness-of-Fit Tests

To test the dynamic programming model's ability to fit the data, the choice probabilities from the estimated dynamic programming model $p(d|s,\hat{\theta})$ are compared to the observed (non-parametric) estimates of the conditional choice probability function $\hat{p}(d|s)$. The non-parametric estimate $\hat{p}$ is the sample histogram of choices made in the subsample of wells with state s. Following Rust and Phelan [1997] and Rothwell and Rust [1997], by sample enumeration, if S is a collection of s cells, the nonparametric estimate of the choice probability is computed as:

$$\begin{aligned} \hat{p}(d|S) &= \int_{s\in S} \hat{p}(d|s)\hat{F}(\mathrm{d}s|S) \\ &= \frac{1}{N_S}\sum_{i=1}^{N} I\{d_i = d, s_i \in S\} \end{aligned}$$

where $\hat{F}(\mathrm{d}s|S)$ is the nonparametric estimate of the conditional probability distribution of s given S, equal to the number of observations in cell ds divided by the total number of observations in all cells that comprise S. This is compared to the estimates of the choice probability from the dynamic programming model:

$$\begin{aligned} p(d|S,\hat{\theta}) &= \int_{s\in S} p(d|s,\hat{\theta})\hat{F}(\mathrm{d}s|S) \\ &= \frac{1}{G}\sum_{g=1}^{G}\frac{1}{N_g}\sum_{i=1}^{N_g} p(d|s,\hat{\theta_g})I\{s_{ig}\in S\} \end{aligned}$$

where $p(d|s,\hat{\theta_g})$ is the probability given by equation (3) and $\hat{\theta_g}$ are the estimates of the structural parameters for group g.

Table 4 shows the observed choice probabilities alongside the expected choice probabilities from the dynamic programming model for oil and gas wells. The three panels in Table 4 show the cases that S is a collection of all possible s cells, that S is a collection of wells that are active, and that S is a collection of wells that are inactive. The dynamic programming model does a very good job at predicting the overall observed choice probabilities. In predicting the operating state all wells, the chi-square test cannot reject the dynamic programming model

Table 4: Actual versus Predicted Choice Probabilities

	Oil		Gas	
Current State: Active or Inactive	Observed	Expected	Observed	Expected
Pr(Active)	0.6207	0.6205	0.6878	0.6822
Pr(Inactivate)	0.3683	0.3684	0.3026	0.3042
Pr(Decommission)	0.0110	0.0111	0.0096	0.0136
No. Obs.		150,078		186,274
χ^2		0.15		227.37
Marg.Sig.		0.93		0
Active				
Pr(Active)	0.9253	0.9256	0.9421	0.9330
Pr(Inactivate)	0.0691	0.0690	0.0534	0.0572
Pr(Decommission)	0.0056	0.0054	0.0046	0.0097
No. Obs.		96,880		129,322
χ^2		0.90		402.53
Marg.Sig.		0.64		0
Inactive				
Pr(Activate)	0.0660	0.0652	0.1105	0.1128
Pr(Inactive)	0.9132	0.9133	0.8685	0.8650
Pr(Decommission)	0.0209	0.0216	0.0210	0.0222
No. Obs.		53,198		56,952
χ^2		1.67		7.25
Marg.Sig.		0.44		0.03

Notes: The chi-square test statistic was calculated as $\chi^2 = N\sum_{d=1}^{3}\{(\text{Obs}_{Pr(d)} - \text{Exp}_{Pr(d)})^2/\text{Exp}_{Pr(d)}\}$, where N is the number of observations.

at the 93% significance level for oil wells and at the 70% significance level for gas wells. In the case of gas, as well as when the sample is separated by current operating state, the chi-squared test rejects the dynamic programming model, however the observed and expected probabilities do not differ by more than .01.

5.5 Comparing Actual and Simulated Data over Time

Using the state of the industry in 2000 as a starting point, I simulate the progression of wells, quantity extracted, and remaining reserves over seven years (Figure 1). The data used in the estimation of the parameters are unbalanced for some wells because observations were missing for some pools in some years. The simulations only include wells that are observed in every year from 2000 to 2007. Each well's path is simulated individually by a series of pseudorandom draws from its type-specific probability density of the state transitions and its subsequent type-specific probability density of operating-state choice. Each well path is simulated 30 times to obtain a 90% confidence interval around the average simulation.

The dynamic programming model is able to match the data closely for the first year of the simulation for both oil and gas wells, but over time slightly overpredicts the number of inactive oil wells and overpredicts the number of inactive gas wells. The purpose of these simulations is to show that the dynamic programming model can predict decisions over time (not just a one time snapshot). Divergence between the predicted and actual data could be driven by the parameters being estimated on the full sample, but this comparison uses only wells with a balanced panel of observations. Also, it is worth noting that we see production and reserves drop off over time, but this is not a prediction for Albertan production or reserves; this model only predicts outcomes for those wells already drilled (and does not include the drilling of new wells).

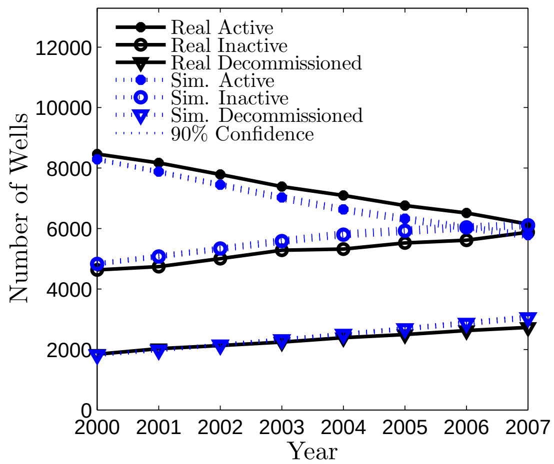

(a) Composition of Oil Wells

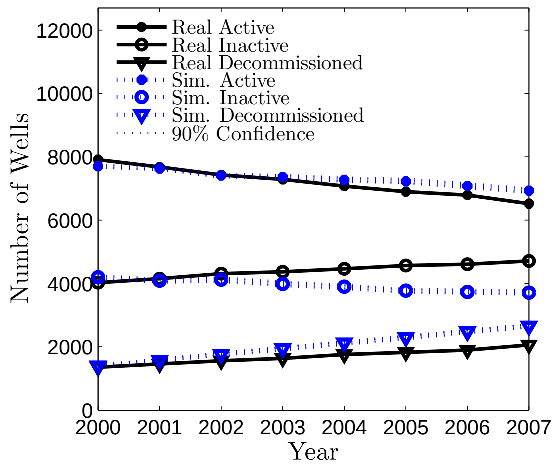

(b) Composition of Gas Wells

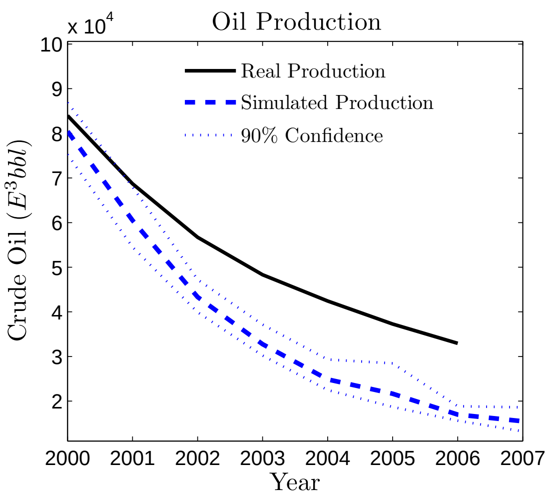

(c) Production from Oil Wells

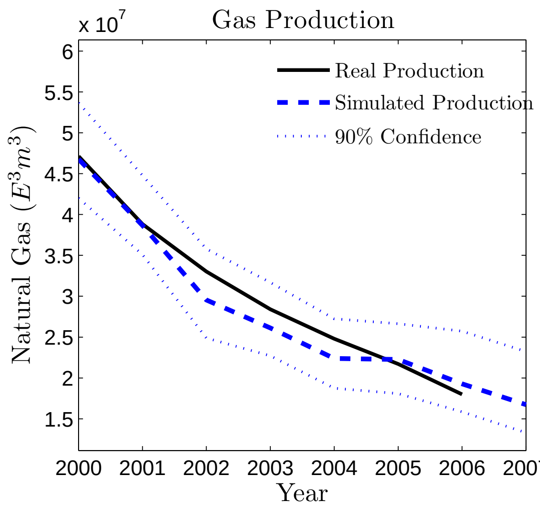

(d) Production from Gas Wells

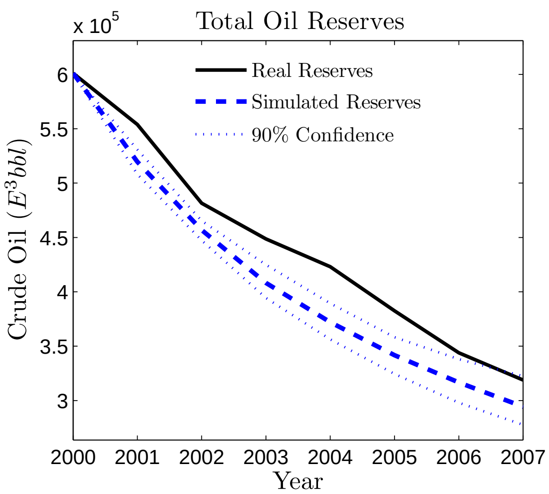

(e) Reserves for Oil Wells

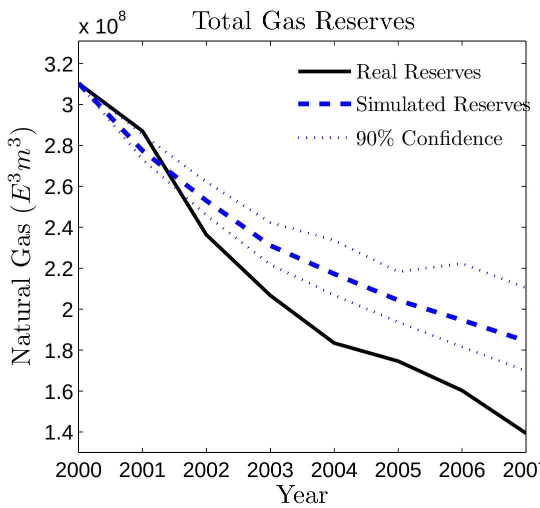

(f) Reserves for Gas Wells

Figure 1: Comparison of Actual Data and Simulated Data

5.6 Policy Application of the Model

If the option to mothball a well did not exist, then there would be a state of nature, L, of low enough prices, remaining reserves, and expectations for their increase that the value of operating the well, $V_L(1)$, would be less than the value of decommissioning, $V_L(3)$, and the well would be decommissioned. When the option to mothball exists, then such a well (with $V_L(1) < V_L(3)$) would not necessarily be decommissioned because leaving it inactive could be more valuable (less costly) than decommissioning, $V_L(1) < V_L(3) < V_L(2)$. However the values of course depend on the current state of nature, and it could be the case that when prices or recovery rates become high enough, H, and these same wells would be reactivated in a high state, $V_H(3) < V_H(2) < V_H(1)$. However, there is a problem when there are wells for which even in a high state it is preferable to decommission them, but even more preferable to leave them inactive, $V_{L,H}(1) < V_{L,H}(3) < V_{L,H}(2)$.

It could be the case that if the externalities associated with mothballing are internalized, then the option to decommission would become preferred and wells that would not be reactivated would be decommissioned, $V_{L,H}(1) < V_{L,H}(2 \text{ \& internalized costs}) < V_{L,H}(3)$. There are various policies that a regulator could try that would persuade the inactive wells, $V_H(1) < V_H(3) < V_H(2)$, to be decommissioned, $V_H(1) < V_H(2) < V_H(3)$, such as by increasing inactivity costs or subsidizing decommissioning costs. But without knowing the social cost of an inactive well, policies might be too bold and result in too many wells being prematurely decommissioned. Specifically, the worry of regulators is that these policies might encourage wells that would otherwise be producing, $V_H(3) < V_H(2) < V_H(1)$, to be decommissioned, $V_H(2) < V_H(1) < V_H(3)$. This might occur because the value of an active well depends on the option to mothball, so when the option to mothball becomes more expensive, the value of an active well also decreases. Moreover, wells that would have otherwise been temporarily deactivated in a low state, might be decommissioned and then not able to reactivate in a high state. Premature decommissioning is an undesirable outcome especially when the regulator places a large value on domestic production.

To evaluate a policy based on its ability to encourage decommissioning without jeopardizing future production really depends on the well in question. Importantly, the value of a well in each of its operating states depends on the current state of nature as well as the payoff shock, ϵ. Because there is a wide distribution in the age of wells and the remaining reserves, one policy will affect each value differently. Therefore to evaluate a policy, it is important to look at the current distribution of wells in question and to look at how the policy affects the sum total of all wells. To do so, I simulate the choices for wells under different counterfactuals to determine the total effect on wells in the sample. If a significant proportion of the wells in the sample fall into the category of $V_{L,H}(1) < V_{L,H}(3) < V_{L,H}(2)$, then under a high state of nature we would not see many reactivations. We can also see how changes to parameters in the model affect the operating choices, current production, and cumulative production overall. An important caveat is that these simulations are made under the assumption that the estimated transition probabilities reflect the actual transition probabilities of the operator. It could be the case however that an operator is more optimistic about reserves or prices and puts higher weight on future high states relative to those estimated in this paper. Therefore the hypothesis that operator behavior can be explained by their waiting for better conditions is tested under the assumption that operators forecast the future using the same transition probabilities that are estimated in the paper. The hypothesis that operator behavior can be explained by optimism about the future is not tested.

5.6.1 Twelve Year Forecasts of Ideal Scenarios

The model is used to simulate the industry under different scenarios that operators claim to be waiting for: high prices, improved recovery rates and reduced reactivation costs. These ideal scenarios are compared to a baseline scenario where prices, recovery factors and the state of the industry are simulated to progress using only the estimated parameters.

In the first scenario, each well-type faces the costs estimated for the type, but they now receive a constant "high price" of \$197.72/bbl for oil and \$462.44/e^3m^3 for gas produced.[35] At the end of the 12 year forecast, averaged over 50 simulations, the high price for oil is 3.2 times the average forecasted price of the baseline. As illustrated in Figure 2 and listed in the first two rows of Table 5, the high price case results in 19% more oil wells that are

[35]This is equal to the U.S. Energy Information Administration's Annual Energy Outlook of 2009 "high price" case in 2030 for oil, and 1.5 times the "high price" for gas.

active than the baseline prediction. The high price for gas wells is on average 2.0 times the average forecasted price of the baseline after 12 years, but only leads to 6% more wells that are active. The annual oil production by the twelfth year is 21% higher, and there is 21% more cumulative oil production over the prior 12 years. In the case of oil, the growth in reserves does not compensate for the increased production, so that after 12 years there are 24% fewer oil reserves than in the baseline case. For gas reserves, the high price results in more reserve growth showing that the expected returns from investments in exploration or enhanced recovery are greater for gas than oil. At the end of the 12 years, there are 120% more gas reserves and 78% more production than in the baseline case. It is fascinating that in spite of the active wells being more productive, the increased reserves and higher prices are not sufficient to induce many inactive gas wells to be reactivated. This is particularly striking in the case of gas wells where, with 78% more productive wells and 120% increased remaining reserves, there are only 6% fewer inactive wells than in the baseline case.

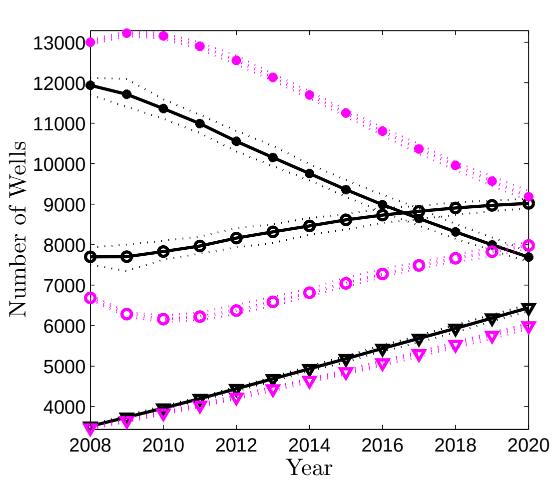

(a) Composition of Oil Wells

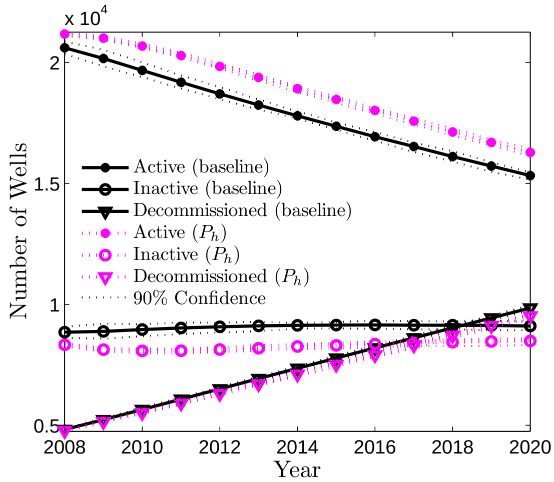

(b) Composition of Gas Wells

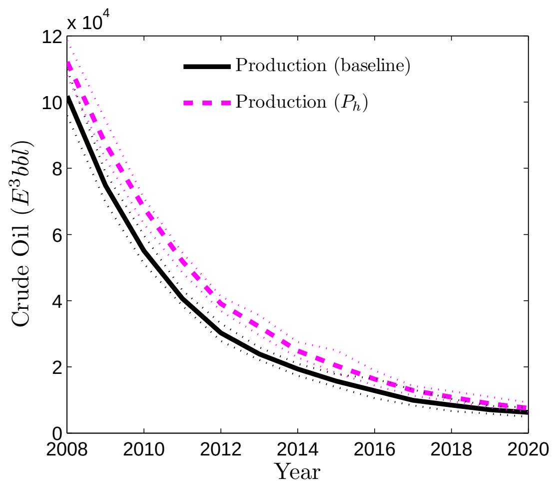

(c) Production from Oil Wells

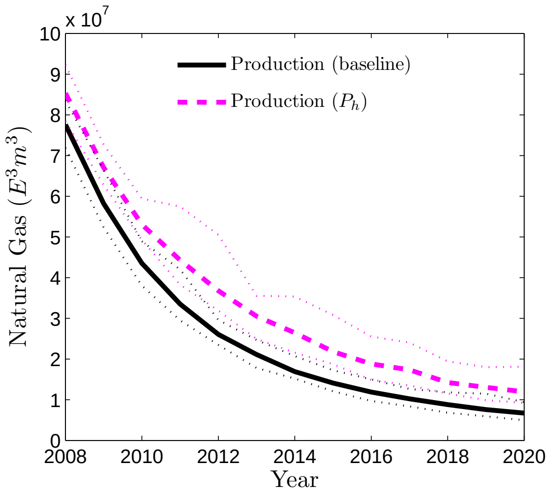

(d) Production from Gas Wells

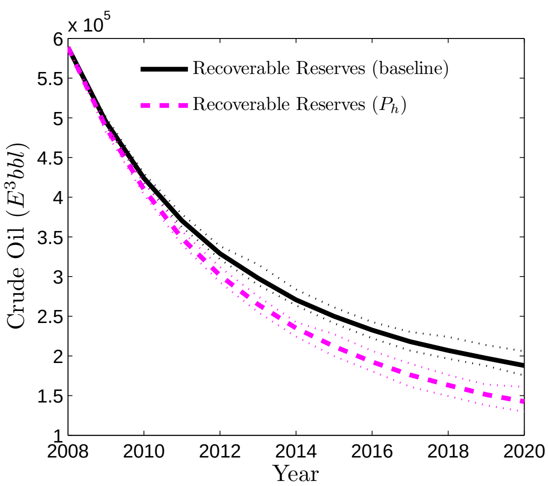

(e) Reserves for Oil Wells

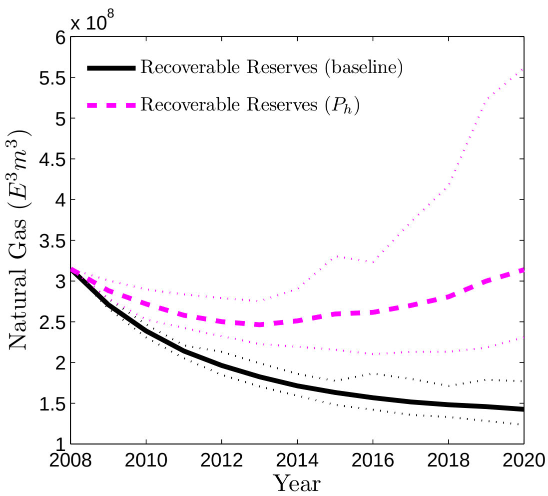

(f) Reserves for Gas Wells

Figure 2: Forecast under Baseline and High Price Scenarios

Table 5: Twelve Year Forecast of Counterfactual Scenarios (% $\triangle$ from Baseline)

Counterfactual Scenario	Well Type	No. Active	No. Inactive	No. Decomm.	Remaining Reserves	Annual Prod.	Cumulative Production
High prices (2.1-3.3×baseline P)	Oil	19%	-12%	-7%	-24%	21%	21%
	Gas	6%	-7%	-3%	120%	78%	31%
High recovery factors (100% recovery)	Oil	17%	-10%	-7%	514%	411%	164%
	Gas	7%	-6%	-6%	418%	387%	275%
Low reactivation ($.75\times SC_{(2\to 1)g}$)	Oil	19%	-7%	-12%	-13%	20%	9%
	Gas	9%	-4%	-10%	-11%	14%	7%
Low decommissioning ($.75\times SC_{(1,2\to 3)g}$)	Oil	-17%	-20%	48%	14%	-17%	-5%
	Gas	-18%	-20%	46%	10%	6%	-2%
Low annual inactivity ($1.25\times M$)	Oil	6%	-9%	5%	-4%	-5%	2%
	Gas	5%	-13%	3%	3%	11%	2%

Notes: Values represent the percent difference between the counterfactual scenarios in the first column and the baseline scenario. Values are averages over 50 simulations. The columns of values represent the difference in: the number of active wells in year 12; inactive wells; decommissioned wells; remaining reserves in year 12; annual production in year 12; and cumulative production over 12 years.

In a second scenario (Table 5, rows 3 and 4) a hypothetical technology change allows for all of the oil or gas-in-place to be recovered. To date, according to the data, recovery rates range from 15% to 95% with an average of 67% for gas and from .01% to 90% with an average of 12% for oil. In the hypothetical scenario, recovery rates are simulated to be 100% of the oil and gas that is in place. In the case of gas, increasing the recovery rate to 100% increases the remaining recoverable reserves by 418%, increases the annual production from producing wells by 387% (275% cumulative). But this significant increase does not induce the reactivation of many inactive gas wells (there are only 7% more active wells at the end of the period than in the baseline case). In the case of oil, the increased recovery rate increases remaining reserves by 514%, increases annual production by 411% (164% cumulative), but the number of active wells only increases by 17%. A higher recovery rate alone has less of an effect on increasing the number of producing wells than a higher price of oil or gas.

However, technology might not only improve recovery factors, but might also decrease the cost to reactivate a well. Therefore, I simulate the industry when the reactivation costs for all well groups are reduced by 25% (Table 5, rows 5 and 6). In the case of gas, after 12 years of lower reactivation costs, there are 9% more active wells, 4% fewer inactive wells and 10% fewer decommissioned wells. Interestingly, the additional production from these wells is only marginal: there is only 14% more production in the last year of the simulation (7% more cumulative production). The reactivated oil wells are slightly more productive than gas wells: after 12 years there are 19% more active wells, and 20% more production (only 9% more cumulative production than in the baseline case). Corresponding with the increased cumulative production, there are less remaining reserves at the end of the 12 year period (a 11-13% reduction). And in the case of gas wells, the production by the end of the 12 years is less than in the baseline case. However due to lack of data on reactivation costs, it is more difficult to assess the probability of a reduction in reactivation costs than to assess the probability of an increase in prices or recovery rates.[36]

It is also interesting to look how responsive the model is to changes in the cost of decommissioning the well. I simulate the 12 year forecast with decommissioning costs being 25% cheaper than in the baseline case. I find that the number of decommissioned wells is very elastic to decommissioning costs, as a 25% reduction in cost results in 46-48% more decommissioned wells. Decreasing decommissioning costs might be unappealing to a regulator because not only the number of inactive wells decrease, but also the number of active wells. This scenario results in 20% fewer active wells, and 2-5% less production.

Instead, increasing the costs of leaving a well inactive could increase the number of decommissioned wells without decreasing the number of active wells. To examine how responsive the operating choice is to the cost of leaving a well inactive, I simulate a 12 year forecast under a scenario where inactivity is 25% more expensive per year. As long as the externalities associated with leaving a well inactive are accounted for, ad infinitum, then leaving a well inactive could be socially optimal. This simulation can be likened to a tax on inactive wells. Such a policy would be more appealing to a regulator in favor of maximizing production because under this scenario the number of decommissioned wells increases (by 3-5%) as do the number of active wells (by 5-6%). As expected with the reactivation of wells, this scenario results in an increase in the cumulative oil and gas produced over the 12 year period. However, the increase in cumulative production is less than the increase in reactivated wells

[36]Cost data are very hard to come by. For example, PSAC publishes average drilling and completion costs for 46 "typical" wells for $750 per year, however these data would not shed light on the reactivation costs which could vary widely depending on the well.

(cumulative production only increases by 2%), indicating that the average reactivated well is considerably less productive.

6 Conclusion

The decision that oil and gas producers make for the operating state of their wells can be categorized as a classic example of an irreversible investment under uncertainty. Restarting production or finally decommissioning a well is an expensive endeavor and is made with uncertainty in future recovery technology and prices. I show that this decision can be modeled by a real options formulation, and the structural parameters of the model can be estimated using data on operating decisions from oil and gas wells in Alberta. Indeed, the operating decisions made for 84,000 wells in Alberta can be replicated by modeling well operators as dynamic optimizers. Within-sample goodness of fit tests show that the model is able to closely predict actual operating choices.

The example of whether to activate, temporarily deactivate, or permanently decommission is used to demonstrate real options in textbooks and classrooms, however, the case of firms using temporary deactivation as a way to avoid paying for permanent decommissioning has been ignored. The motivation of this paper was to determine the rational for leaving oil and gas wells inactive; either they could be a blessing, if they eventually are reactivated and contribute to our energy supply, or they could be a curse, if they are never reactivated, cause environmental degradation, and must undergo costly decommissioning. With the estimated structural parameters I can predict how the operating choices might change under different conditions. I find that increased oil and gas prices and recovery rates might increase per well annual production, but will not substantially increase the number of active wells. For example, doubling the gas price results in a 120% increase in recoverable reserves, but only a 6% increase in the number of inactive wells. On the other hand if it became cheaper for wells to be reactivated, we would have more active wells, but these these wells would likely not be very productive, and total production would not increase as much. The cost of decommissioning plays an important part in determining the number of decommissioned wells, however a policy to decrease the cost of decommissioning a well would result in not only fewer inactive wells, but also fewer active wells, making such a policy less appealing to a regulator in favor of production. On the contrary, a policy of increasing the cost of mothballing a well would increase the number of active wells as well as decrease the number of inactive wells.

If optimistic conditions are not enough to induce well reactivation, this implies that wells are left inactive not because of the option to reactivate, but rather the sunk cost of decommissioning is too high to warrant undertaking. Should there be externalities from idling the wells (such as continued contamination of groundwater) that are not accounted for in the decision, then this behavior may not be socially optimal. This paper demonstrates that for the majority of inactive wells, temporary closure is, in effect, permanent closure. Advances in hydraulic fracturing and horizontal drilling have prompted a new surge in drilling for oil and gas which has resulted in much debate and controversy, particularly in regards to the immediate environmental impacts. Less attention has been given to concerns for the final cleanup and reclamation of these wells—unwarrantedly so given that in many of the areas that wells are being drilled, the financial bonds meant to ensure cleanup were set over half a century ago. The occurrence of firms not internalizing their decommissioning costs is likely generalizable to other industries that have high clean up costs and are also allowed extended temporary closures.

References

Alberta Agricultural and Rural Development. Compensation for surface leases. Technical report, 2010. URL http://www1.agric.gov.ab.ca/$department/deptdocs.nsf/all/ofa11701.

Alberta Department of Energy. Royalty information briefing 6: Corporate income tax, Alberta royalty review. Technical report, 2007. URL http://www.assembly.ab.ca/lao/library/egovdocs/2007/aleo/159246.pdf.

Alberta Energy. Oil and gas fiscal regimes: Western Canadian provinces and territories. Technical report, Alberta Department of Energy, 2006. URL http://www.energy.gov.ab.ca/Tenure/pdfs/FISREG.pdf.

Alberta Energy Regulator. Alberta's energy reserves 2007 and supply/demand outlook 2008-2017. Technical report, 2008. URL http://www.aer.ca/documents/sts/ST98/st98-2008.pdf.

Alberta Energy Regulator. Compliance assurance, directive 19. Technical report, Alberta Energy Regulator, 2010a. URL `http://www.aer.ca/documents/directives/Directive019.pdf`.

Alberta Energy Regulator. Well abandonment, directive 20. Technical report, Alberta Energy Regulator, 2010b. URL `http://www.aer.ca/documents/directives/Directive020.pdf`.

Alberta Energy Regulator. Abandonment and closure orders. Technical report, Alberta Energy Resources Conservation Board, June 2013a. URL `http://www.aer.ca/documents/orders/ibo/AbandonmentOrders.pdf`.

Alberta Energy Regulator. Licensee liability rating (LLR) program: Updated industry parameters and industry costs, directive 011. Technical report, Alberta Energy Regulator, 2013b. URL `http://www.aer.ca/documents/directives/Directive011_May2013.pdf`.

Alberta Environment. Reclamation criteria for wellsites and associated facilities. Technical report, Alberta Environment, 1995. URL `http://environment.gov.ab.ca/info/library/6855.pdf`.

Alberta Environment. *Environmental Protection and Enhancement Act*. 2000.

Ramón A Alvarez, Stephen W Pacala, James J Winebrake, William L Chameides, and Steven P Hamburg. Greater focus needed on methane leakage from natural gas infrastructure. *Proceedings of the National Academy of Sciences*, 109(17):6435–6440, 2012.

J.R. Arrington. Predicting the size of crude reserves is key to evaluating exploration programs. *The Oil and Gas Journal*, 58(9):130–134, 1960.

Fischer Black and Myron S. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–54, 1973.

James W. Boyd. Financial responsibility for environmental obligations: Are bonding and assurance rules fulfilling their promise? Discussion Papers dp-01-42, Resources for the Future, 2001.

Michael J. Brennan and Eduardo S. Schwartz. Evaluating natural resource investments. *Journal of Business*, 58(2):135–157, 1985.

Richard H. Byrd, Jorge Nocedal, and Richard A. Waltz. KNITRO: An integrated package for nonlinear optimization. In G. di Pillo and M. Roma, editors, *Large Scale Nonlinear Optimization*, volume 4, pages 35–59. Springer Verlag, 2006.

Larry W. Canter, Robert C. Knox, and Deborah M. Fairchild. *Ground water quality protection.* Lewis Publishers, Inc., Chelsea, MI, 1987.

CAPP. Statistical handbook for Canada's upstream petroleum industry. Technical report, Canadian Association of Petroleum Producers.

Augusto Castillo-Ramirez. An application of natural resource evaluation using a simulation-dynamic programming approach. *Journal of Computational Finance*, 3(2):91–107, 1999.

Janie M. Chermak and Rorbert H. Patrick. A well-based cost function and the economics of exhaustible resources: The case of natural gas. *Journal of Environmental Economics and Management*, 28:174–189, 1995.

Gonzalo Cortazar and Jaime Casassus. Optimal timing of a mine expansion: Implementing a real options model. *The Quarterly Review of Economics and Finance*, 38(3, Part 2):755–769, 1998.

Gonzalo Cortazar, Eduardo S. Schwartz, and Jaime Casassus. Optimal exploration investments under price and geological-technical uncertainty: a real options model. *R&D Management*, 31(2):181–189, 2001.

John K. Dagsvik. How large is the class of generalized extreme value random utility models? *Journal of Mathematical Psychology*, 39(1):90–98, 1995.

Avinash Dixit. Entry and exit decisions under uncertainty. *Journal of Political Economy*, pages 620–638, 1989.

Avinash Dixit. Investment and hysteresis. *The Journal of Economic Perspectives*, 6(1):107–132, 1992.

Avinash K. Dixit and Robert S. Pindyck. *Investment Under Uncertainty.* Princeton University Press, Princeton, NJ, 1994.

Doneivan F. Ferreira, Saul B. Suslick, and Paula C.S.S. Moura. Analysis of environmental bonding system for oil and gas projects. *Natural Resources Research*, 2003.

Matthew Foss, Daniel V. Gordon, and Alan MacFadyen. The cost of lifting natural gas in Alberta: A well level study. *Discussion Paper, University of Calgary*, 2002.

Andrea Gamba and Matteo Tesser. Structural estimation of real options models. *Journal of Economic Dynamics and Control*, 33(4):798–816, 2009.

Tarek M. Harchaoui and Pierre Lasserre. Testing the option value theory of irreversible investment. *International Economic Review*, 42(1):141–166, 2001.

Stan Hurn and Robert E. Wright. Geology or economics? Testing models of irreversible investment using North Sea oil data. *Economic Journal*, 104(423):363–371, 1994.

Ryan Kellogg. The effect of uncertainty on investment: evidence from Texas oil drilling. Technical report, National Bureau of Economic Research, 2010.

John Kennan and James R. Walker. The effect of expected income on individual migration decisions. *Econometrica*, 79(1):211–251, 2011.

Robert F. Kubichek, Jerry J. Cupal, W. Iverson, S. Choi, and M. Morris. *Identifying Ground-Water Threats from Improperly Abandoned Boreholes*. 1997.

Clinton Levitt. Learning through oil and gas exploration, University of Iowa manuscript. 2009.

Gary D. Libecap and Steven N. Wiggins. Contractual responses to the common pool: Prorationing of crude oil production. *The American Economic Review*, 74(1):87–98, 1984.

Gary D. Libecap and Steven N. Wiggins. The influence of private contractual failure on regulation: the case of oil field unitization. *The Journal of Political Economy*, 93(4):690–714, 1985.

John R. Livernois and Russell S. Uhler. Extraction costs and the economics of nonrenewable resources. *The Journal of Political Economy*, 95(1):195–203, 1987.

Charles F. Mason. Nonrenewable resources with switching costs. *Journal of Environmental Economics and Management*, 42:65–81, 2001.

Robert C. Merton. Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4(1):141–183, 1973.

Austin L. Mitchell and Elisabeth A. Casman. Economic incentives an regulatory framework for shale gas well site reclamation in pennsylvania. *Environmental Science and Technology*, 45:9506–9514, 2011.

Alberto Moel and Peter Tufano. When are real options exercised? An empirical study of mine closings. *The Review of Financial Studies*, 15(1):35–64, 2002.

David F. Morehouse. The intricate puzzle of oil and gas “reserves growth”. *Energy Information Administration, Natural Gas Monthly*, pages vii–xx, 1997.

National Petroleum Council. Plugging and abandonment of oil and gas wells, paper 2-25, working document of the npc north american resource development study. Technical report, 2011.

Andriy Norets. Inference in dynamic discrete choice models with serially orrelated unobserved state variables. *Econometrica*, 77(5):1665–1682, 2009.

Sheila Olmstead, Lucija Muehlenbachs, Jhih-Shyang Shih, Ziyan Chu, and Alan Krupnick. Shale gas development impacts on surface water quality in Pennsylvania. *Proceedings of the National Academy of Sciences*, March 2013.

Orphan Well Association. Orphan well association 2007/08 annual report, Alberta oil and gas orphan abandonment and reclamation association, 2008.

Orphan Well Association. Orphan well association 2012/13 annual report, Alberta oil and gas orphan abandonment and reclamation association, 2013.

James L. Paddock, Daniel R. Siegel, and James L. Smith. Option valuation of claims on real assets: The case of offshore petroleum leases. *Quarterly Journal of Economics*, 103(3):479–508, 1988.

Virginia Parente, Doneivan Ferreira, Edmilson Moutinho dos Santos, and Estanislau Luczynski. Offshore decommissioning issues: Deductibility and transferability. *Energy Policy*, 34:1992–2001, 2006.

Pennsylvania Department of Environmental Protection. Stray natural gas migration associated with oil and gas wells (draft report). Technical report, 2009.

Province of Alberta. *Oil and Gas Conservation Act, O-6 RSA*. 2000.

Province of Alberta. Mines and Minerals Act: Petroleum Royalty Regulation, Alberta regulation 248/90. Technical report, 2008.

Railroad Commission of Texas. Railroad commission plugs record number of abandoned wells in fiscal year 2006, 2006. URL `http://www.rrc.state.tx.us/news-releases/2006/100606.html`.

Geoffrey Rothwell and John Rust. On the optimal lifetime of nuclear power plants. *Journal of Business and Economic Statistics*, 15(2):195–208, 1997.

John Rust. Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, 55(5):999–1033, 1987.

John Rust. Maximum likelihood estimation of discrete control processes. *SIAM Journal on Control and Optimization*, 26(5):1006–1024, 1988.

John Rust. Structural estimation of Markov decision processes. In R.F. Engle and D. McFadden, editors, *Handbook of Econometrics*, volume 4, chapter 51, pages 3081–3143. 1994.

John Rust and Christopher Phelan. How social security and Medicare affect retirement behavior in a world of incomplete markets. *Econometrica*, 65(4):781–831, 1997.

Richard R. Schneider, Grant Hauer, W.L. (Vic) Adamowicz, and Stan Boutin. Triage for conserving populations of threatened species: the case of woodland caribou in alberta. *Biological Conservation*, 143(7):1603–1611, 2010.

Margaret E. Slade. Valuing managerial flexibility: An application of real-option theory to mining investments. *Journal of Environmental Economics and Management*, 41:193–233, 2001.

State Corporation Commission of Kansas. General rules and regulations for the conservation of crude oil and natural gas, 2010. URL `http://www.kcc.state.ks.us/conservation/`.

Gunnar Stensland and Dag Tjostheim. Optimal investments using empirical dynamic programming with application to natural resources. *Journal of Business*, 62(1):99–120, 1989.

United States Government Accountability Office. Oil and gas bonds: Bonding requirements and BLM expenditures to reclaim orphaned wells. Technical report, January 2010. URL `http://www.gao.gov/assets/310/300218.pdf`.

C.R. Vanorsdale. Evaluation of initial reservoir data to estimate devonian shale gas reserves. In *SPE Annual Technical Conference and Exhibition*, 1987.

Mahendra K. Verma and Mitchell E. Henry. Historical and potential reserve growth in oil and gas pools in Saskatchewan. Miscellaneous Report 2004-4.1, Saskatchewan Geological Survey, 2004.

Theresa Watson and Stefan Bachu. Evaluation of the potential for gas and co2 leakage along wellbores. *SPE Drilling & Completion*, 24(1):115–126, 2009.

Halbert White. Maximum likelihood estimation of misspecified models. *Econometrica*, 50(1):1–25, 1982.

Michael L. Williams, Charles R. Matthews, and Tony Garza. Well plugging primer. Technical report, Railroad Commission of Texas, Well Plugging Section, Oil and Gas Division, 2000. URL `http://www.rrc.state.tx.us/forms/publications/plugprimer1.pdf`.

A Appendix

A.1 PSAC Areas

Table 6: Characteristics of PSAC Areas

PSAC Area	Surface	Hydrocarbon	Characteristics
1	Rocky Mountains	Deep gas	Strict environmental regulations
2	Ranching, farming and forest	Oil and gas	Easily accessed
3	Agricultural prairie grassland	Gas and medium/heavy oil	Easily accessed
4	Prairie and woodland	Gas and heavy oil	Easily accessed
5	Agricultural	Oil and gas	Most densely populated area
6	Prairie and woodland	Shallow gas	Only winter drilling
7	Agricultural and logging	Oil and gas	Often no road access and winter drilling

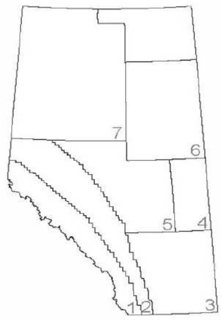

Figure 3: PSAC Areas

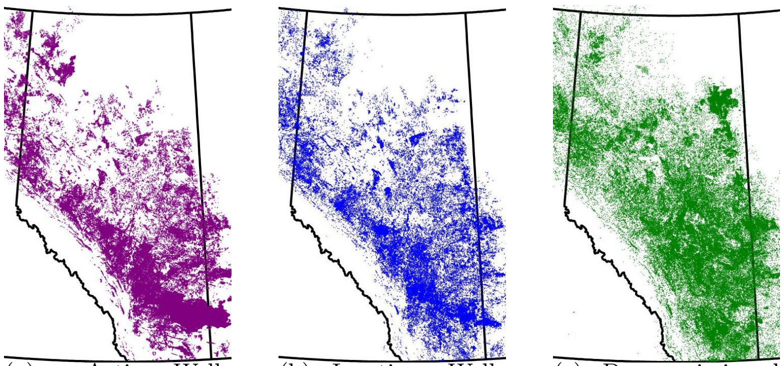

(a) Active Wells (150,997)

(b) Inactive Wells (95,575)

(c) Decommissioned Wells (103,885)

Figure 4: The Location of Oil and Gas wells in Alberta in 2007 by Operating State

A.2 Subsample as Compared to Full Sample

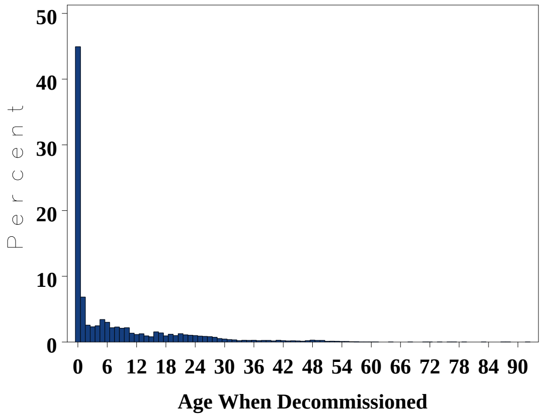

(a) Full Sample

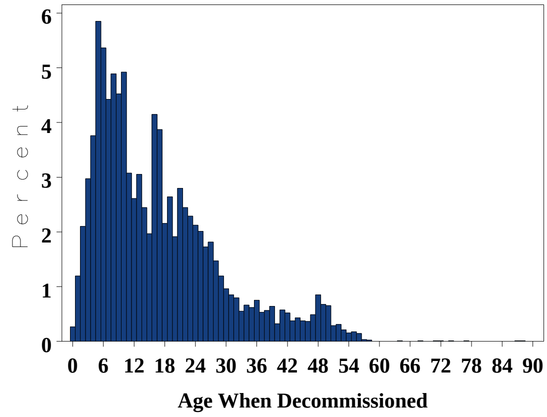

(b) Wells in Full Sample that Produced

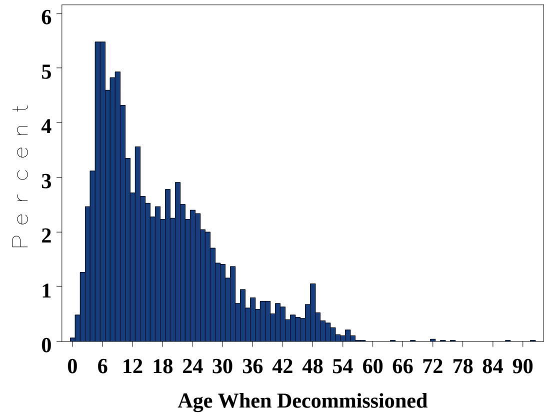

(c) Subsample

Figure 5: Histogram of Age when Decommissioned

A.3 Choice of Discount Factor

The optimal discount factor varies by the different well groups and no one discount factor would result in the highest likelihood for all groups separately. However, by summing the log likelihoods across the well groups by different fixed discount factors, a discount factor between .90 and .95 has the highest log likelihood. Therefore, the estimation presented in this paper uses a constant discount factor of .90 across all of the different well types. The summed log likelihoods by discount factor for all gas well groups are displayed in Figure 6. Also depicted are the average estimated costs across the well groups by discount factor. The reactivation costs are for inactive wells, 5 years of age and the lifting costs are for wells with per well reserves of 3 million m^3).

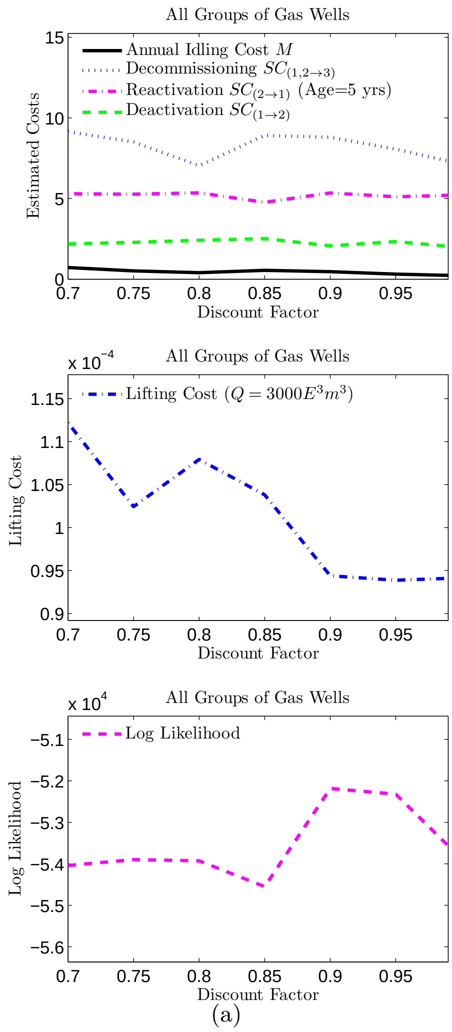

Figure 6: Effect of Using Different Discount Factors on Results from All Gas Wells

A.4 Modeling the Probability of Transition of Recoverable Reserves

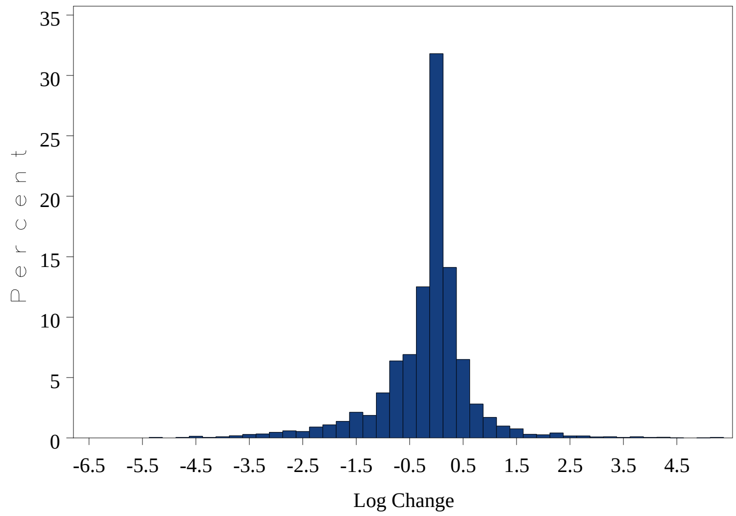

(a) $\log\left(\frac{\overline{Q}^{IER'}}{\overline{Q}^{IER}}\right)$ for Oil Pools

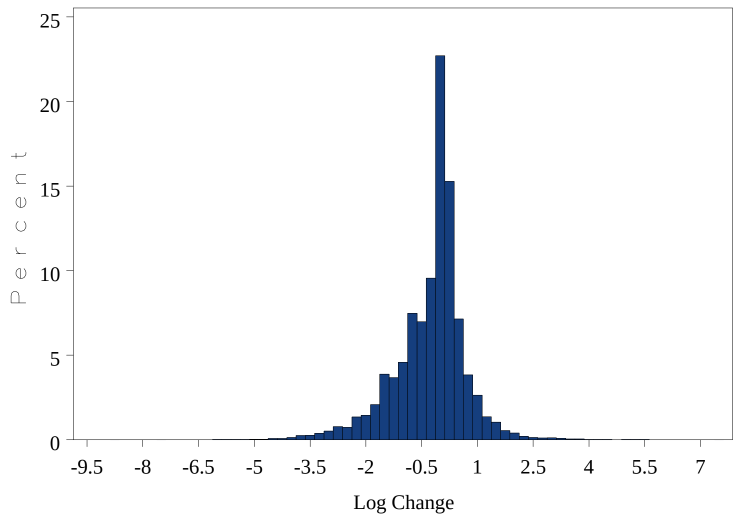

(b) $\log\left(\frac{\overline{Q}^{IER'}}{\overline{Q}^{IER}}\right)$ for Gas Pools

Figure 7: Histograms of the Natural Logarithm of Annual Change in Initial Established Reserves (not including occurrences of no change)