

DISSERTATION

AN OPTION VALUE ANALYSIS OF HYDRAULIC FRACTURING

Submitted by

Joshua H. Hess

Department of Economics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Fall 2017

Doctoral Committee:

Advisor: Terrence Iverson

Harvey Cutler
Stephan Weiler
Dale Manning

Copyright by Joshua H. Hess 2017

All Rights Reserved

ABSTRACT

AN OPTION VALUE ANALYSIS OF HYDRAULIC FRACTURING

Many uncertain public policy decisions with sunk costs can be optimally timed leading policymakers to delay implementing a policy despite positive expected net present value. One salient example of this is hydraulic fracturing (fracking), a recently developed oil and gas extraction technology, that has increased fossil fuel reserves in the US. However, many municipalities have seen fit to ban its use despite seemingly positive expected net benefits. We hypothesize that an option value framework that values the ability to delay and learn about an uncertain project may explain fracking bans in practice where the neoclassical net present value rule does not. We test this by developing a stochastic dynamic learning model parameterized with a computable general equilibrium (CGE) model that calculates the value of learning about uncertainty over damages and uncertainty over benefits. Applying the model to a representative Colorado municipality, we quantify the quasi-option values (QOV), which create an additional incentive to ban fracking temporarily in order to learn. To our knowledge, this is the first attempt to quantify an economy-wide QOV associated with a local environmental policy decision.

In Chapter 1 we argue that a numerical, option value approach is the appropriate way to examine uncertain public policy issues involving sunk costs. This method allows for an optimal timing of the public project rather than the 'now or never' approach of the ubiquitous net present value rule. We present local fracking policy as an excellent application for an option value approach as has positive expected net benefits but has been subject to local bans seemingly despite the net present value rule. We also defend our use of a CGE model to estimate the local economic benefits of fracking.

Chapter 2 presents the option value model associated with *epistmelogical uncertainty* over environmental damages. Also, this chapter presents damage values parameterized to the City of Fort Collins for application in this and the subsequent chapter. With this in hand, we solve the model and demonstrate the results.

Chapter 3 has a similar structure to Chapter 2. First, it discusses the literature on stochastic oil movements, then it presents the option value model associated with *stochastic uncertainty* over local benefits. Then, assuming the same parameterized expected damage as in Chapter 2, we solve the model and display the results.

ACKNOWLEDGEMENTS

Thanks Mom and Dad! I love you. I hope this makes up for my teenage years. Also, thanks Dr. Chris Blake. You are a good friend and colleague but the worst officemate ever.

Charlize and Pierson, although you don't know it, you've sacrificed much for this; and, although I'm very proud of it, I am still more proud of you. Tiffany, I could not have done this without you. No one could ask for a better friend and partner. You are patient and supportive beyond belief. I love you madly!

TABLE OF CONTENTS

Abstract	ii
Acknowledgements	iv
Chapter 1. Option Values in Public Policy	1
1.1. Introduction	1
1.2. Option Values	3
1.3. Application of Methodology	7
1.4. Predicting Economic Impacts of Fracking	18
1.5. Conclusion	25
Chapter 1 References	27
Chapter 2. The Value of Learning about Hydraulic Fracturing Damages	35
2.1. Introduction	35
2.2. Model	36
2.3. Results	45
2.4. Discussion and Conclusion	55
Chapter 2 References	59
Chapter 3. The Value of Learning about Hydraulic Fracturing Benefits	60
3.1. Introduction	60
3.2. Oil Price Dynamics	61
3.3. Model	64
3.4. Results	72
3.5. Conclusion	79

Chapter 3 References.....	82
---------------------------	----

CHAPTER 1

OPTION VALUES IN PUBLIC POLICY

1.1. INTRODUCTION

Uncertain public projects are typically evaluated by comparing the present value of the expected costs and benefits. If the expected benefits exceed the expected costs – i.e. the expected net present value is positive – then it is optimal to do the project. This process treats the decision as one that must be made now or never. It places no value on the ability, if it exists, to revisit the decision at a later date, possibly with a better understanding of the risks. When a project is uncertain, has upfront sunk costs, and can be delayed it is optimally made by comparing the expected net present value of doing it now to the expected value of forgoing the project today and revisiting it at a later date – i.e. the option value.

To illustrate the intuition of this option value, briefly consider a two-period example. Suppose doing a project that involves sunk costs will give profit π in period one. There are two possible (present) values for period two profit: ϕ or ψ , and there is no cost to forgoing the project. Further suppose the policymaker initially thinks that the probability of ϕ is p_1 and the probability of ψ is $1 - p_1$. However, she also expects information to arrive before the decision is revisited. So, she expects the period two probability, p_2 , to be different from the period one probability. We can now write the expected net present value of doing the project in period one, and the period one value of delaying the project and revisiting it with the updated probabilities. The expected net present value is $NPV = \pi + [p_1\phi + (1 - p_1)\psi]$ and the value of delaying the project is $V_d \equiv 0 + [p_2 \max\{\phi, 0\} + (1 - p_2) \max\{\psi, 0\}] \geq 0$. In both the net present value and the value of delay expressions, the last two terms are bracketed to

emphasize that they are expectations. This highlights the intuition that the value of delay comes from the expectation of maximizing with updated information where as the net present value assumes the policymaker cannot revisit the decision. As to policy, the net present value rule would have the project go forward if $\max\{NPV, 0\} > 0$. However, since this project can be delayed and has sunk costs, it is optimally done if $\max\{NPV, V_d\} = NPV$, i.e. if the net present value exceeds the value of delay.

Generalizing the process to a public project has large obstacles not present in this simple example. In reality, there are likely many possible states of profit, not just two. Moreover, they are likely to be time-dependent responding to ever-changing economic conditions. Thus, incorporating option values into a public policy decision requires calculating the potential costs and benefits for many states of uncertainty and for many policy paths. As there is no generally agreed upon functional form for a the time-dependent impacts of a policy that affects many sectors, we turn to numerical methods.

The next section of this chapter discusses the literature on option values, their development and applications, and the appropriate numerical methodology for valuing a public project. The following section introduces a salient policy issue to which an option value model is applicable and provides valuable insight: hydraulic fracturing. The subsequent section discusses numerical models of the economic impacts of hydraulic fracturing and the final section concludes.

1.2. OPTION VALUES

Two approaches to option values developed independently but are related. The quasi-option value approach matured in the environmental and resource literatures and has been widely applied, perhaps most notably by the climate change literature for its ability to value learning about uncertain parameters. The contingent claims approach was brought to the economics literature from the finance literature and shown to be equivalent to a dynamic programming approach. This has come to be known as the real options approach. These two approaches differ in their inception and uses but are fundamentally related.

1.2.1. DEVELOPMENT. While pointing out that some private goods have public good characteristics, Weisbrod (1964) informally initiated the quasi-option value. The literature spent the next 25 years pinpointing its source and defining its features (see Hanemann (1989) for a summary). Importantly, Arrow and Fisher (1974) and Henry (1974) develop its use in the environmental and resource economics literature and Hanemann (1989) formalizes it into what is now commonly called the *Arrow-Fisher-Hanemann-Henry Quasi-Option Value* (QOV).

In contrast to the development of the quasi-option approach, the real options approach borrowed a well-defined concept (the call option) from the finance literature to explain observed phenomena. Pindyck (1991) suggested that firms may not be using the net present value rule in decision making, which might explain why “econometric models have had limited success in explaining and predicting changes in investment spending.” Dixit and Pindyck (1994) broadcast the approach to a wide audience, including the mainstream environmental

and resource literature, demonstrating how to include Wiener process price movements in a dynamic programming framework.

The Dixit and Pindyck option value (DPOV) and the QOV are not necessarily equivalent, as pointed out by Mensink and Requate (2005), and Traeger (2014) provides a general relationship between the QOV and DPOV. The QOV is the value of learning conditional on the ability to postpone a decision and the DPOV is the net value of postponing a decision conditional on the ability to learn about the uncertain element. The difference between the QOV and DPOV is due to their origins. Resource decisions, like instituting a carbon tax, can be delayed but investment decisions, like purchasing a competitors company may not arise again. Moreover, epistemological uncertainty can resolve (e.g. honing in on the ‘true value’ of climate sensitivity), which makes learning valuable, but stochastic uncertainty (e.g. what will be the price of widgets in the future?) may be constant (like in the case of a Markov process) and learning – observing the uncertain element prior to making the decision – is likely less valuable.

1.2.2. QOV APPLICATIONS. This QOV models are routinely applied to climate policy. Chichilnisky and Heal (1993) argue that the failure of global warming models to account for irreversibility has led to an understated need for immediate action, though Ulph and Ulph (1997) show that Epstein’s (1980) irreversibility conditions are not met for even a simple, two-period model of global warming. Kolstad (1996a) jointly examines capital investment in abatement (sunk-cost irreversibility) and environmental damages from the stock of carbon (emissions irreversibility) and argues that either sufficiently fast learning or sufficiently slow carbon decay makes either decision irreversible. Kolstad (1996b) concludes that capital investment irreversibility increases initial optimal emissions levels (lowers abatement). Fisher

and Narain (2003) support this result, finding that the negative effect of capital irreversibility on optimal abatement outweighs the opposing impact of irreversible environmental damages. Lemoine and Traeger (2014) model irreversible changes in climate sensitivity as crossing an unknown threshold, but learning is strictly ex post in their model. They find that the existence of ‘tipping points’ raises the optimal first-period carbon tax.

1.2.3. CALCULATING OPTION VALUES. Realistic calculations of the DPOV and the QOV require valuing the project over all the states of the stochastic variable. This is readily done at the firm level when costs are known but price (Pindyck 1991) or demand (Dixit 1991) are uncertain. DPOV approaches typically value the project by assuming a functional form for the stochastic element (e.g. price follows geometric Brownian motion) and using Itô calculus to analytically solve the resulting differential equation. On the other hand, applied QOV approaches typically use numerical methods (e.g. recursive dynamic programming versions of DICE) to calculate optimal policy under a limited number of states (e.g. two climate “tipping points”). With no generally agreed upon functional form that describes the long run effect of changes in one part of a complex economy on its total, we turn to numerical methods.

There are several possible numerical methods for estimating policy impacts. Obviously, econometric methods should be preferred whenever the data is available. This may not be the case, though, if the project has not yet been done in an economically similar location. Of the remaining methods input-output (IO) and computable general equilibrium (CGE) modeling have become the most popular.

IO models fix input coefficients rendering them static models. This results in two problems. First, the IO construction process can take some time which could result in a considerable difference between the analysis period and the base year. Second, the model is unable to handle relative price changes over time, technological advances, or changing returns to scale. These issues are best addressed by a dynamic framework (West 1995).

The assumption of linearity in production is prevalent in IO models. This implies a strictly proportional relationship between input coefficients and output (West 1995). Moreover, household income is an average propensity, employment is determined by average productivity, and consumption depends on average expenditures. This fails to capture any nonlinearities which are well-handled by CGEs (West 1995).

IO models are Keynesian in nature, assuming fixed prices and perfectly elastic supply. Inputs are perfectly elastically supplied in production because IO models have no supply-side constraints (Partridge and Rickman 1998). This eliminates the role of price disconnecting the value added of primary factors and final demand. Because of this, IO models almost always predict that economic impact will be proportionate to the exogenous change.

On the other hand, CGE models are Walrasian allowing for imperfectly elastic labor supply and prices to flex so that markets clear. They typically use nested constant elasticity of supply functional forms for value added in production and intermediate goods. This allows for the specification of differing elasticities of substitution between good types and region of production. Then, in CGE models the value added factor usage responds to factor costs and imports of intermediate goods respond to price (Partridge and Rickman 1998). Consequently, the various elasticities of supply and demand determine the economic response to a shock,

which is not inevitably proportionate (Partridge and Rickman 1998). Because of this, CGE models are commonly used to empirically analyze the welfare impacts of policies whose effects may be transferred through multiple markets (Wing 2004).

1.3. APPLICATION OF METHODOLOGY

We contend that hydraulic fracturing is a good application for an option value approach to public policy. It is a relatively new technology for extracting oil and gas from shale deposits and there is wide public (Kohut et al. 2012) and scientific (Shonkoff et al. 2014; Jackson et al. 2014) uncertainty as to its costs. On the other hand, it has greatly increased fossil fuel reserves across the United States and beyond often times bringing substantial economic benefits (Hausman and Kellogg 2015). Since 2010, though, there have been more than 50 bans related to fracking enacted or adopted in the US and abroad (e.g., Longmont, CO; New York State; Germany) (see Figures 1.1 and 1.2). This is despite, in some cases, seemingly positive expected net local benefits (Wobbekind and Lewandowski 2014; Feyrer et al. 2017). Some of these bans, like the ones in Delaware River Basin and Quebec, were explicitly temporary citing the need for further study.

Like Pindyck (1991) we hypothesize that an option value framework may explain the seeming failure of the net present value rule to explain observed phenomena. First, we provide a brief history of the technology, the opposition to its uses, and the adoption of bans and moratoria¹. Then, we survey the economic literature the benefits and damages.

¹A list of fracking bans in practice can be found at Food and Water Watch

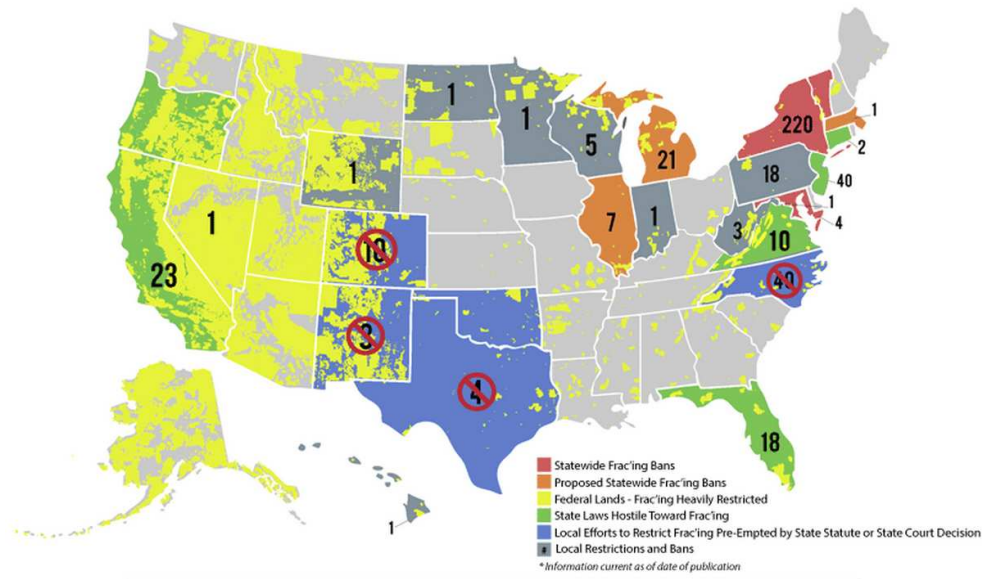


FIGURE 1.1: US Bans, Restrictions, and Hostile Resolutions
Source: Energy In Depth; July, 2015

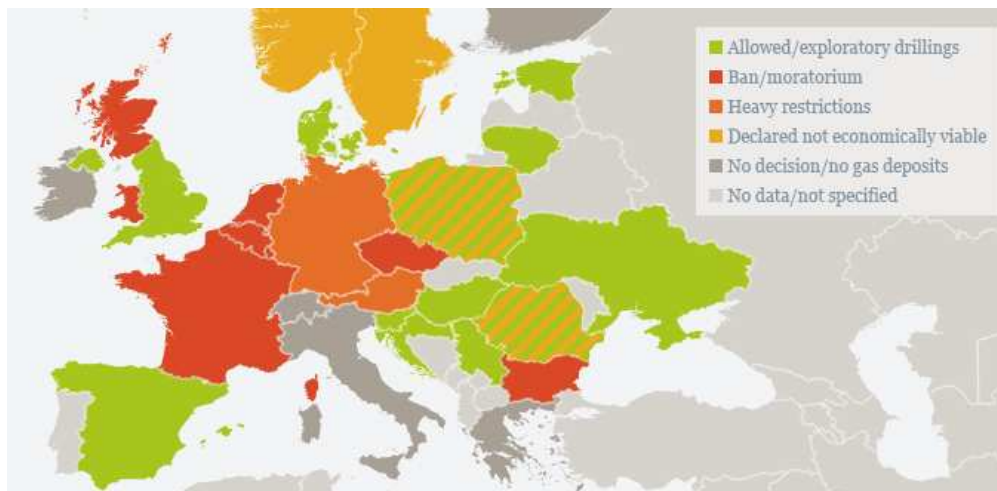


FIGURE 1.2: Fracking in Europe
Source: Deutsche Welle; July, 2015

1.3.1. THE HISTORY OF FRACKING, OPPOSITION, AND BANS. During 1862 Battle of Fredericksburg, Union officer Colonel Edward A. L. Roberts witnessed Confederate artillery shells plunging into a canal that obstructed the bloody battlefield. This experience planted the seed for an idea that he would later call “superincumbent fluid tamping”. In 1866, he was

awarded a U.S. Patent (No. 59,936) for the Roberts Torpedo (AOGHS 2015). The torpedo was lowered into a borehole filled with water and detonated to more efficiently fracture the surrounding strata. The technique was quickly and widely adopted. Some wells reported flow increases 1,200% within a week of being ‘torpedoed’ (Manfreda 2015). Despite safety concerns, the nitroglycerin torpedo remained in use for oil and gas development until the 20th century.

In 1947, Stanolind Oil and Gas began to study the amount of pressurized treatment used in their wells. This led to experimentation in Grant county, Kansas which saw a thousand gallons of gelled gasoline injected some 2,400 feet down into limestone followed shortly by a gel breaker Manfreda (2015). Halliburton later conducted two more successful experiments and fracking became commercialized

In the mid-1970s, amidst a petroleum shortage, the US Department of Energy launched the Eastern Gas Shales Project whose goal was developing techniques to extract ‘unconventional’ natural gas reserves. Shale development is particularly challenging due to the relatively low permeability of the rock. This project saw the development of many techniques in use today: horizontal drilling, multi-stage fracturing, and ‘slick’ water fracturing which dramatically increased the pressures delivered to rock formations.

Fracking, as it is practiced today, began in the Barnett shale in the mid 1990s. Nick Steinsberger, a newly promoted executive at Mitchell Energy, had just figured out that his area – roughly 200 wells in central Texas – was to be shut-down (Smith 2016). In the face of necessity, he decided to make the field more cost effective by watering down the explosive gel. Despite the predictions (and objections) from the gel’s manufacturers, it worked. The

conventional wisdom at the time was that excessive over-watering would make the formation swell and the wells wouldn't produce (Smith 2016). The conventional wisdom was shown to be wrong and the modern 'slickwater' fracture job was born. Mitchell Energy combined fracking and sideways drilling throughout the Barnett shale demonstrating economic feasibility at a level not seen before. The merger of Mitchell's company with Devon Energy in 2002 ignited the modern shale gas boom.

In Tuscaloosa County, Alabama, a family living near the River Gas coalbed methane (CBM) development enlisted the help of the Legal Environmental Assistance Foundation (LEAF) to fight nearby fracking operations believing it to be the cause of their contaminated well. A 1990 task force, with participation by state, federal, and industry entities had acknowledged that contamination was possible and recommended guidelines for underwater injection control (UIC). However, Alabama's Oil and Gas Board did not meaningfully incorporate these guidelines into its fracking rules (although, as Drilling Contractor (2000) points out, there was widespread compliance among development agencies according to the industry). Because of this, LEAF petitioned the EPA to remove Alabamans 'primacy' in the oversight of underwater injection control. The EPA denied the petition claiming there was no evidence that hydraulic fracturing endangered underground water and should not be regulated by its UIC program. LEAF then filed with the 11th Circuit Court of Appeals. In 1997, the court ruled in favor of LEAF deciding underwater injection did apply to all hydraulic fracturing, not just CBM fracking, the issue of the initial complaint. The EPA began a study on CBM in 1999 and released the results in 2004. They concluded that CBM operations posed no risk to underwater drinking supplies. This led to the 2005 Energy Policy Act, which amended the 1974 Safe Drinking Water Act to exclude fracking injection fluids other than diesel fuels and

also exempted extraction companies from disclosing the chemicals involved in their fracking operations. This Act has since been criticized as being unduly influenced by oil and gas concerns through the Energy Task Force led by then vice president and former Halliburton CEO Dick Cheney.

Josh Fox's 2010 documentary *Gasland*, ignited widespread public opposition to fracking. It received a Primetime Emmy as well as awards from the Sundance Film Festival and the Sarasota Film Festival and was nominated for both an Academy Award and a Writer's Guild Award. The film follows Fox as he speaks with citizens of Colorado, Utah, Wyoming, and Texas about their chronic health problems alleged to come from air, ground water, and surface water contamination. The film was criticized as factually incorrect and the Energy in Depth organization produced an associated film, *Truthland*, as a rebuttal. Journalist Phelim McAleer directed *FrackNation*, which attempts a more balanced portrayal of fracking. It is more factually oriented than *Gasland* and attempts to depict both the economic benefits to impoverished rural Americans as well as the issue of contamination.

The *Gasland* ignition was literal. Its depiction of a homeowner setting tap water on fire, indicating high levels of methane alleged to have come from a fracking operation, has over 170,000 youtube views and a related video, *Light Your Water on Fire from Gas Drilling, Fracking* attributed to the Gas Drilling Awareness Coalition has over 1.6 million views.

The anti-fracking organization Food & Water Watch, maintains an up to date list of fracking bans and resolutions (Grant 2017). Opposition to fracking has occurred at the city/municipality, county, state, and country levels; however, one of the earliest bans came under the authority of the Delaware River Basin Commission (DRBC). The Delaware basin contains about

13,539 square miles, about one third of which sits atop the Marcellus shale (the second shale play in the US after the Barnett) in Pennsylvania, New Jersey, and New York. Over 15 million people rely on the Delaware for drinking water including the populous New York City. In 2010 the five commissioners unanimously voted to delay any drilling decisions until new regulations could be implemented (Collier 2009).

In November 2010, Pittsburg became the first US city to ban fracking. Buffalo, NY quickly followed with a ban in February 2011, but it was largely symbolic as there was no interest in drilling. Dryden, NY enacted a ban in August 2011 and quickly became a center of contention although several courts have found in favor of the small town. Emboldened by Dryden's success, several New York towns enacted or adopted bans: Syracuse, Albany, Woodstock, Rochester, Wawarsing, Kirkland, and Canandaigua culminating with a statewide ban issued by Governor Cuomo in December of 2014. The New York ban was not the first statewide ban, only the most significant. Vermont, which had no oil and gas development, had banned fracking in March of 2012 and Connecticut enacted a three-year ban on storage and handling of fracking waste in August 2014. Other notable US city bans include Los Angeles, Beverly Hills, Philadelphia, and Denton, TX. The Denton ban is particularly interesting as Denton, located on top of the Barnett shale, is considered the birthplace of fracking. However, on May 18th, 2015, in the wake of an Ohio ruling maintaining state superiority, Governor Greg Abbott signed a bill that made local bans illegal in Texas. Denton overturned their ban the following month.

A similar battle occurred on the Colorado Front Range, which sits atop the Niobrara shale formation. The town of Longmont changed its charter to prohibit fracking in November 2012 and was sued by the Colorado Oil and Gas Association – a large industry group –

alongside the principal development interest, TOP Operating, and the Colorado Oil and Gas Conservation Commission, the state agency in charge of regulation. Seemingly unfazed by this litigation, the cities of Fort Collins², Lafayette, Boulder³, and Broomfield passed similar measures in November 2013. The case is situated between *Voss*, in which the court found that cities (Greeley in this case) could not ban fracking due to “statewide interest in the efficient development and production of oil and gas resources.”⁴, and *Bowen/Edwards* where the court found counties (La Plata) could add regulations (permits) if they did not conflict with the state’s interest. Compounding the legal difficulties of the issue, both rulings were handed down the same day eliminating reliance on precedence. On May 2, 2016, however, the Colorado Supreme Court ruled in favor of the State’s rights, overturning the Longmont ban and the Fort Collins moratorium, though the Court maintained local rights to unspecific regulation⁵.

Mora County, NM became the first US County to have a ban in May 2013, but a federal judge overturned it in January 2015. Hawaii County enacted a ban in October of 2013 and was followed by three California counties: Santa Cruz County in May 2014 and then Mendocino and San Benito Counties in November elections that year. Internationally, France became the first country with a moratorium in June 2011⁶. Since then, Bulgaria, Luxembourg, Germany, Scotland, and Wales have banned fracking. In addition, due to concerns about

²There were only eight wells in Fort Collins

³With no oil and gas interests, this ban was largely symbolic.

⁴Justice Joseph Quinn said in the opinion.

⁵NY Times

⁶It was upheld in 2013.

local environmental impacts, the Cantabria Region of Spain, Nova Scotia, Quebec, and Five City Breaks municipality of Argentina have implemented fracking bans.

1.3.2. THE LITERATURE ON FRACKING. For the most part, the existing economics literature on fracking focuses on either costs or benefits, though there are a few notable exceptions: Jackson et al. (2014) discuss the environmental benefits and costs, Walsh et al. (2015) uses a spatial econometric approach to examine the political, socioeconomic, and social factors influencing local bans in New York State, and Loomis and Haeefele (2017) perform a benefit cost analysis at the national level. In contrast, we build on the prior literatures on fracking benefits and costs, while embedding local fracking policy within an explicit decision-theoretic framework. To our knowledge, this analysis is the first to quantify the value of learning about the full costs of fracking in practice, and the first to carefully assess both the costs and benefits underlying observed policies.

1.3.2.1. *Environmental Damages.* Despite a consensus that risks exist, there remain significant gaps in the literature connecting risk pathways to economic impacts (Jackson et al. 2014; Shonkoff et al. 2014). Krupnick and Gordon (2015) prescribe important pathways through which routine fracking activities may impact humans and the environment (see Figure 1.3). Their survey of 215 experts from industry, academia, government, and NGOs finds a high degree of consensus concerning the most important pathways that arise during normal fracking activities. Hedonic property valuation studies have attempted to quantify the cost of these local impacts (Boxall et al. 2005; Gopalakrishnan and Klaiber 2014; James and James 2014). Typically, these studies find negative impacts on housing values from nearby oil and gas development, especially if the home depends on groundwater (Muehlenbachs et al. 2015). Bennett and Loomis (2015) estimate that each well drilled within a half-mile of

a house in Weld County, Colorado decreases the home value by \$1,805 in urban areas. Many human health studies have focused on benzene pollution outside the fracking context (Chen et al. 2000; Aguilera et al. 2009; Slama et al. 2009; Zahran et al. 2012) and have highlighted significant health damages, including low birth weight. Hill (2013) shows that fracking in Pennsylvania increased the occurrence of low birth weight (by 25%) and term birth weight (by 18%). She estimates a lower bound of the public cost to be \$4.1 million due to low infant birth weights caused by benzene air pollution from fracking in Pennsylvania in 2010. Crime is another potential burden of fracking. James and Smith (2017) find a positive effect on various property and violent crimes around active fracking plays inducing an estimated \sim \$2 million dollar cost at the county level.

Boudet et al. (2014) argue that fracking bans accompany a high degree of public uncertainty about local damages. A 2012 Pew Center poll found that 26% of Americans had heard a lot about fracking, 37% had heard a little, and 37% had heard nothing. Regardless of the level of information, it is perceptions about risk (informed or otherwise) that drive local policy. Schenk et al. (2014) find that the public is most concerned about water quality and seismic activity, despite a general scientific consensus that best practices manage these particular risks well.

1.3.2.2. *Economic Benefits.* While it is obvious that exploiting valuable fossil fuel reserves will confer gross economic value, it is less clear how much of the created value will accrue to the local economy where fracking occurs. Also, it is unclear if short run employment gains will result in long run welfare gains.

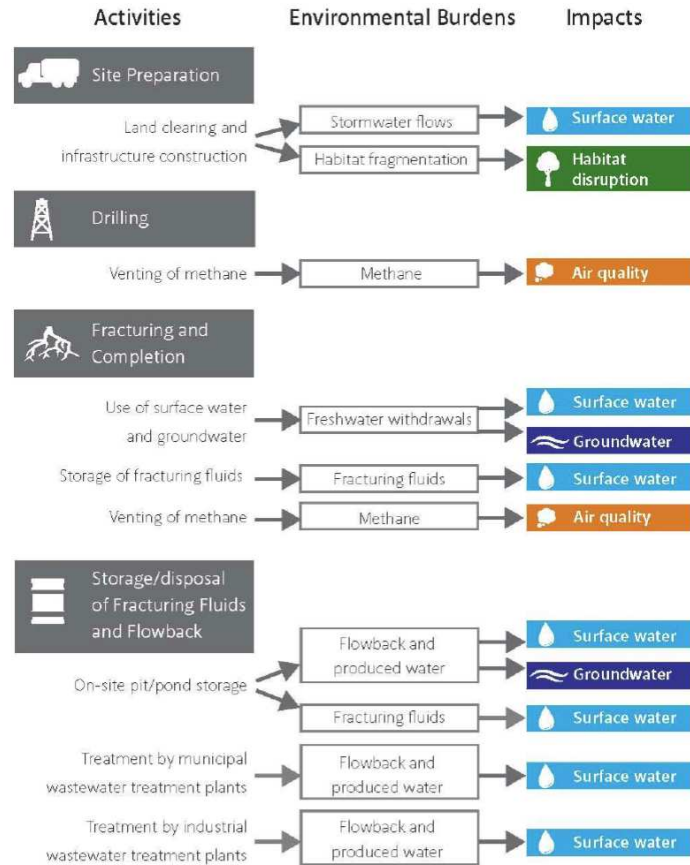


FIGURE 1.3: Routine Risk Pathways

Source: Krupnick and Gordon (2015)

Hausman and Kellogg (2015) use an econometric approach and find economy-wide total surplus gains of a third of a percent in the US. Regionally focused work has estimated significant employment gains – *ex post* – to Pennsylvania (Considine et al. 2010), Colorado (Wobbekind and Lewandowski 2014), and Arkansas (CBER 2008). However, some of these studies use input-output models whose results are sensitive to assumptions (i.e. household spending and savings behavior, labor supply elasticity, mineral rights ownership) and the benefits are likely to be significantly overstated (for a review of several studies see Kinnaman 2011). More recent *ex post* analyses find much smaller employment impacts than predictions (Weber 2012). Maniloff and Mastromonaco (2014) find that *ex post* job growth fell ‘well short’

of predictions. Furthermore, natural resource development could reduce economic growth through a “resource curse” (Corden and Neary 1982; Sachs and Warner 1995). Allcott and Keniston’s (2014) national analysis and Weber’s (2014) analysis of the Arkansas, Louisiana, Oklahoma, and Texas find no evidence of Dutch Disease, contrary to work by Jacobsen and Parker (2016). Maniloff and Mastromonaco (2014) reconcile this by looking at so-called ‘boom counties’. They find that counties with tight labor markets and little prior industry presence saw growth in tradable sector wages, a precursor to Dutch Disease. Using a comprehensive data set of oil and gas production alongside LBS and IRS data spanning seven years, Feyrer et al. (2017) find that each million dollars of new oil and gas production is associated with \$80,000 increase in wage income and 0.85 new jobs within the county. Further, roughly 40 percent of the income increases are occurring in tangentially related industries like construction, hospitality, and local government. Additionally, they find \$132,000 in royalty payments and increased business income. These impacts are relatively persistent over a two year time period and spill over into the rest of the economy.

This leaves the question, what are the general equilibrium impacts of beyond this time frame? As discussed in above, CGE models are an excellent tool for a data-driven analysis of aggregate welfare for policies that affect many sectors of a Walrasian economy (West 1995; Partridge and Rickman 1998; Weber 2014). In the next section, we reassert the necessity of using a CGE model, rather than an IO model, to estimate the impacts of fracking on a jurisdiction.

1.4. PREDICTING ECONOMIC IMPACTS OF FRACKING

In order to use an option value framework to analyze fracking policy, we must calculate the economic benefits of the possible policies, allowing fracking and maintaining, a ban across decision periods, and the results must be empirically-grounded to explain observed phenomena. One way to do this would be to assume that local economic benefits are a constant fraction of private benefits measurable by the value of unconventional reserves only economically accessible by fracking. Another way would be a different numerical approach.

We assert that CGE models are the best way to do this and argue against the two alternatives. This section first reviews previous attempts at predicting the economic impacts using IO models and discusses their shortcomings. Then, it presents our CGE model and demonstrates how our CGE model aligns with current econometric findings where assuming local benefits are a fraction of private benefits does not. Finally, we demonstrate that a CGE is capable of calculating a 'multiplier' effect in line with current econometric findings, and that sensitivity of policy results to its specification.

1.4.1. SUPERIORITY OF CGE MODELS FOR FRACKING POLICY. There have been several attempts at numerically predicting the economic impacts of fracking using IO models (CBER 2008; Considine et al. 2009; Considine et al. 2010; Scott and Huang 2009; Weinstein and Clower 2009), though they were generally either conducted or solicited by industry. Kinnaman's (2011) review of these studies includes the critiques mentioned in Section 1.2.3 summarized in Figure 1.1.

Instead we use a CGE model (discussed in the subsequent section) to estimate the impact of fracking that includes, among 17 other sectors, oil and gas production. Although there may be uncertainty as to the size of local benefits, we assume they are deterministic and the portion that accrue locally follows from assumptions about nonlocal resource ownership. CGE models are commonly used to empirically analyze the welfare impacts of policies whose effects may be transferred through multiple markets (Wing 2004) in Walrasian economies (Partridge and Rickman 1998), like allowing fracking in a municipality. Recent econometric analyses of direct effects indicate that there are employment spillovers in retail, construction, and transportation (Maniloff and Mastromonaco 2017), and Feyrer et al. (2017) estimates that 40 percent of local income increases occur in industries unrelated to fracking like construction, hospitality, and local government. These income shocks can lead many consumers to increase consumption (Brown et al. 2017). Basic economic theory implies consumption shocks can affect an entire economy and can have a multiplier effect. Our simulated general equilibria agree (Figure 1.5). We find that a million dollar increase in the economically recoverable reserves increase the present value of household consumption by \$1.84 million. Although a functional form might result in an analytically tractable model, slight misspecifications of this multiplier would dramatically alter the calculated optimal policy and misstate the importance of learning (Table 1.2). Rather than this, we directly combine the CGE and the dynamic program. CGE models have already been used to understand the impacts of environmental regulation (Bovenberg and Goulder 1996; Weyant 1999; Goulder 2002) and begun to include environmental amenities as a sector (Carbone and Smith 2013). However, to our knowledge, there has yet to be a municipal-level CGE analyzing the impact of fracking.

TABLE 1.1: A Comparison of IO Studies and Assumptions

Source: Kinnaman (2011)

Shale play	Estimated impact	In the year	To the economy of	Assumptions
Marcellus	\$4.2B in output 48,000 jobs	2009	Pennsylvania	100% royalties spent immediately "The locations of all these suppliers and income recipients were determined using the company profile databases Reference U.S.A. and Manta, which also provided the economic sector for each purchase" (95% of direct spending in state)
Marcellus	\$8.04B in revenues 88,588 jobs	2010	Pennsylvania	100% royalties spent immediately "The locations of all these suppliers and income recipients were determined using the company profile databases Reference U.S.A. and Manta, which also provided the economic sector for each purchase" (95% of direct spending in state)
Barnett	\$11B in revenues 111,131 jobs	2008	Dallas/Ft. Worth Area	"The amounts were fully adjusted to reflect those funds that are paid outside the region (and state) and are further reduced to account for out-of-area spending, savings, and taxes."
Hayensville	\$2.4B in revenues 32,742 jobs	2008	Louisiana	All direct spending in state Assumes households spend 5% of lease and royalty payments in 2008.
Fayetteville	\$2.6B in revenues 9533 jobs	2007	Arkansas	Survey asks firms to report state of residence of employers, but not whether spending occurs in state or out of state.
Marcellus	\$760M in revenues 810 jobs	2000 wells over 10 year period	Broome County, NY	Assumptions regarding percentage of drill spending in local economy not stated
Marcellus	\$2.06B in revenues 2200 jobs	Gas production per year	Broome County, NY	Assumes 15% of royalty earnings remain in local economy

1.4.2. THE CGE MODEL. The model is an adaptation of Cutler and Davies (2007) who built a model for Fort Collins, CO. Here, we parameterize the model to represent an oil-and-gas-producing Colorado municipality with 50,000 residents whose policymaker is considering the removal of a fracking ban. These changes to the Fort Collins model were made in order to isolate the determinants of a ban, looking across economic parameters that vary across regions where bans have been implemented. The land, labor, and capital employment in each of seventeen production sectors is parameterized using census and county assessor's data from Fort Collins, CO. Data to calculate input-output coefficients for intermediate inputs come from IMPLAN. Fort Collins is large relative to the Colorado average, so the economy is scaled down to 50,000 people⁷, holding constant: production technologies, labor supply per household, and per capita demand.

All production sectors of the CGE model include intermediate inputs, land, capital, and labor. Output and factor prices are endogenous, with perfectly mobile labor in five household

⁷The average size of Colorado cities above 10,000 people is 55,000, excluding the capital city of Denver (United States Census Bureau / American FactFinder. "Annual Estimates of the Resident Population: April 1, 2010 to July 1, 2014" and United States Census Bureau. "B01001 Sex by Age. 2010 - 2014 American Community Survey. U.S. Census Bureau's American Community Survey Office).

groups. Land and capital are quasi-fixed but respond over time to differences in rental rates. Importantly, this implies that returns to land and capital are sector-specific in any time period. Local differences between demand and supply are met by imports (or exports when production exceeds demand). The CGE model also contains local, state, and federal government sectors.

In addition to the standard factors of production, output in the oil and gas sector depends on a natural capital factor that captures the remaining, economically-accessible natural resource stock in the ground. The size of this factor depends on whether or not a fracking ban is in place. Since reserves depend on allowable technology, a ban means the oil and gas sector has access to a smaller stock than if fracking were allowed. Thus, there is some (conventional) production even with a ban. Fitzgerald and Rucker (2016) estimate average annual royalty rates for oil (13.3%-13.8%) and gas (10.5% - 12.7%). Based on this, we assume that 12.5% of oil and gas production value is paid to the owners of those rights⁸.

The simulations computed are inspired by the 2013 Fort Collins moratoria, which was enacted to be five years. We calibrate the CGE model using annual data but in the QOV model, we use 5-year periods by summing the output of the CGE model over 5 years. In addition, we assume that policymakers revisit the ban/frack decision every (five year) period for five model periods (25 years), after which oil and gas companies lose interest in developing the municipality's unconventional reserves and the option to allow fracking vanishes. Regardless of when fracking is allowed, benefits and damages accrue for ten periods. In the base scenario, the policymaker is risk averse with a constant relative risk aversion coefficient of 2, though we also consider the implications of risk neutrality. The base oil and gas prices are assumed

⁸See Appendix for a complete description of the data.

to be \$40 per barrel and \$2.50 per thousand cubic feet respectively. Finally, the annual discount rate is set to 5 percent.

The simulations are constructed in the following way. To generate the no-fracking baseline, we simulate normal growth where total factor productivity and export demand are assumed to increase by one percent separately in the first period. The model then moves to a new steady state after a 50-year time horizon. The second simulation assumes that fracking occurs in the first period, along with normal growth, and stimulates a threefold increase in oil and gas development. This is based on EIA estimates of oil reserves in Colorado, which climbed from 386 million barrels in 2010 to 1200 million barrels as of 2014. This observed increase in oil reserves occurred almost exclusively because of the introduction of fracking, and natural gas reserves experienced a similar increase (Colorado Oil & Gas Conservation Commission). When the unconventional reserves are exhausted, the excess extraction capital immediately exits the local economy. In the simulation, this occurs in the period after next (i.e. ten years later in year 11). Separate simulations are computed for policy scenarios when fracking is allowed in years 6, 11, 16 and 21 (periods 2, 3, 4, and 5) to reflect the decision to allow fracking in each of these periods. This generates levels of household consumption associated with all policy scenarios used in the QOV model.

The recursive CGE model is solved with an annual time step and a 50-year time horizon. Consumption across years is aggregated to obtain 5-year consumption values, used as an input into the dynamic programming model. Figure 1.4 illustrates the consumption paths with spikes occurring at the time of fracking. After the initial shock, fracking stops as reserves are depleted and consumption falls, reaching a new steady state. To test for sensitivity to the price of oil and gas, the CGE model is used to evaluate the impact of fracking bans over




FIGURE 1.4: Consumption Paths From CGE Model

The x-axis is the policy period, which are five-year increments. The y-axis is consumption in billions of dollars. All simulations include normal growth shocks in year 1.

a range of prices. Oil and gas prices are assumed to move together, increasing and decreasing from base values in equal proportions. Given our base specification, the annuitized value of the consumption benefits of fracking in the initial period is equal to $\sim \$113.6$ million per (5-year) period.

1.4.3. SENSITIVITY OF RESULTS. We use the CGE to compute the present value of the change in household consumption from land, labor, capital, and mineral rights payments if fracking is allowed in the first period for different rental values of the unconventional reserves. The results are presented in Figure 1.5.

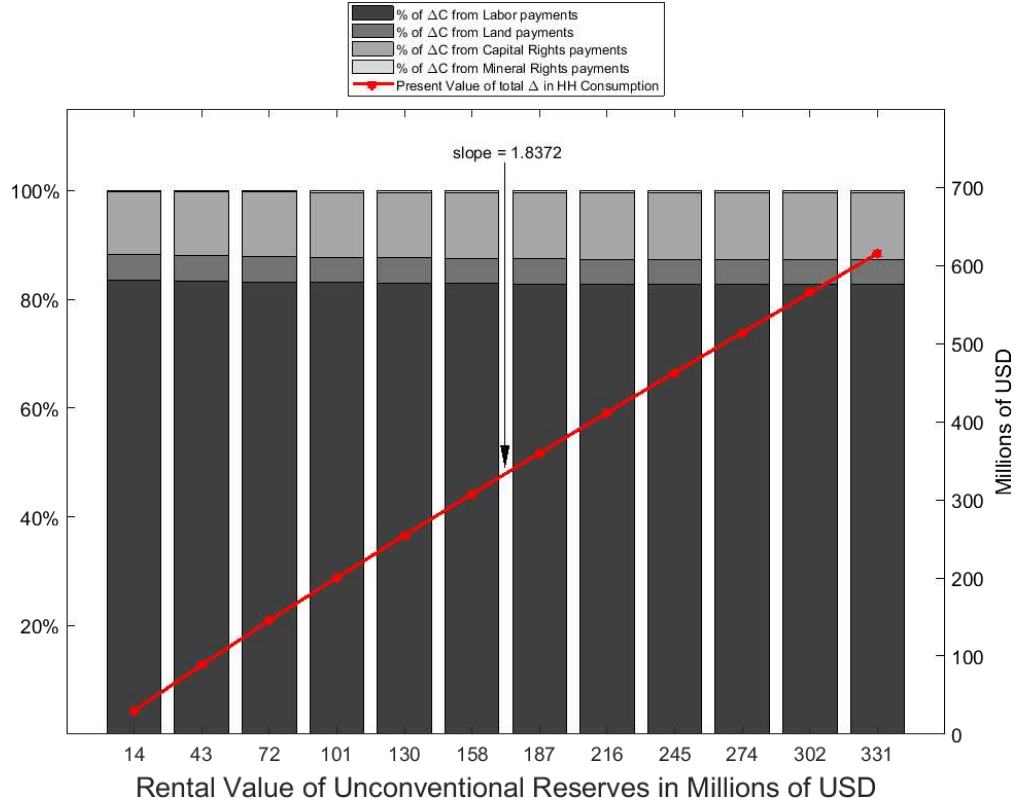


FIGURE 1.5: The Simulated General Equilibrium Impacts of Fracking on the Local Economy

The x-axis is the value of the unconventional reserves, which require fracking for economic development, measured in millions of USD. The left hand y-axis is the percentage of the present value of the total change in consumption benefits attributed to Labor, Land, Capital, and Royalties. Capital and Royalty payments account for less than 15%. The right hand y-axis is the present value of the total change in household consumption.

The bar graphs shows that roughly 95% of the change in consumption is attributable to labor and land payments. Assuming that local benefits are a constant fraction of private benefits – the returns to capital and mineral rights – would dramatically underestimate the benefits. We also compute the present value of the total change in household consumption and plot it with the right side axis. It is roughly linear (OLS $R^2 = .995$) with a slope of 1.8372. This can be interpreted to mean that fracking an unconventional reserve whose rental value is one million dollars will generate \$1.84 million of local consumption.

TABLE 1.2: The Simulated General Equilibrium Impacts of
Local Fracking

The no uncertainty mean where policy switches $\hat{\mu}$ and maximum QOV are reported for different multipliers all in units of millions of USD. The estimated multiplier is highlighted.

Multiplier	$\hat{\mu}$	Maximum QOV
1.5	103.95	293.70
1.7	117.61	267.20
1.8372	126.98	251.29
1.9	131.27	237.95
2.1	144.93	178.51

Policy and QOV calculations are sensitive to misspecifications of this multiplier. We calculate the present value of the total change in household consumption for different multipliers and use it to construct an annuity and solve the finite horizon dynamic program for the no uncertainty policy switchpoint ($\hat{\mu}$) and maximal QOV. We do this under risk-neutrality for two reasons. First, it makes the QOV interpretable as dollar amounts (in fact, millions of dollars). Second, a CRRA decision maker would view an annuity differently than a lumpier, but equivalent in terms of present value, stream preferring whichever brings larger payments quicker. Using alternative multipliers leads to policy and maximum QOV estimates that have wide ranges in comparison to the highlighted estimate. The policy range is 32% of the highlighted base value and the maximum QOV range is of 46% of the base value.

1.5. CONCLUSION

The ability to delay an uncertain public project and revisit it is an important part of optimal decision making. Doing this requires valuing the project across time and policy permutations. Computable general equilibrium models are the appropriate tool for doing this for projects that affect many sectors in a Walrasian economy.

The recent multitude of bans on fracking, which seem to be at odds with the expectation of net economic gains, is an excellent place to apply a numerical option value approach. We hypothesize, like Pindyck (1991), that an this approach may explain these fracking bans where the net present value rule does not. To test this empirically, we have developed a data-driven methodology that links a computable general equilibrium model with a dynamic program to calculate a realistic quasi-option value. The use of the CGE is necessary since results are heavily dependent on parametric specifications and other numerical models are not suitable for Walrasian economies. The next two chapters describe two related implementations of this methodology. Chapter 2 treats benefits as deterministic and damages as uncertain. In this model, learning occurs via Bayesian updating as the policymaker observes a noisy signal during a ban. In Chapter 3, damages are deterministic, but benefits are stochastic depending on price of oil and gas, which follows a geometric Brownian motion. This work contributes a methodology that uses an empirically-oriented, option value framework to asses local policy. This methodology can be widely applied to policies that have uncertainty over costs or benefits, sunk costs, and can be delayed. It also contributes to the discussion on fracking policy by presenting it as an application for public option value analysis and demonstrating a methodology capable of understanding bans in practice. Finally, we contribute a CGE model to the literature on ex ante estimations of the economic impacts of fracking.

CHAPTER 1 REFERENCES

- I. Aguilera, M. Guxens, R. Garcia-Esteban, T. Corbella, M. J. Nieuwenhuijsen, C. M. Foradada, and J. Sunyer. Association between gis-based exposure to urban air pollution during pregnancy and birth weight in the inma sabadell cohort. *Environmental health perspectives*, 117(8):1322–1327, 2009.
- H. Allcott and D. Keniston. Dutch disease or agglomeration? the local economic effects of natural resource booms in modern america. Technical report, National Bureau of Economic Research, 2014.
- K. J. Arrow and A. C. Fisher. Environmental preservation, uncertainty, and irreversibility. *The Quarterly Journal of Economics*, pages 312–319, 1974.
- A. Bennett and J. Loomis. Are housing prices pulled down or pushed up by fracked oil and gas wells? a hedonic price analysis of housing values in weld county, colorado. *Society & Natural Resources*, pages 1–19, 2015.
- H. Boudet, C. Clarke, D. Bugden, E. Maibach, C. Roser-Renouf, and A. Leiserowitz. “Fracking” controversy and communication: Using national survey data to understand public perceptions of hydraulic fracturing. *Energy Policy*, 65:57–67, 2014.
- A. L. Bovenberg and L. H. Goulder. Optimal environmental taxation in the presence of other taxes: general-equilibrium analyses. *The American Economic Review*, 86(4):985–1000, 1996.
- P. C. Boxall, W. H. Chan, and M. L. McMillan. The impact of oil and natural gas facilities on rural residential property values: a spatial hedonic analysis. *Resource and energy economics*, 27(3):248–269, 2005.

- J. P. Brown, T. Fitzgerald, and J. G. Weber. Asset ownership, windfalls, and income: evidence from oil and gas royalties. 2017.
- J. C. Carbone and V. K. Smith. Valuing nature in a general equilibrium. *Journal of Environmental Economics and Management*, 66(1):72–89, 2013.
- Center for Business and Economic Research. Projecting the economic impact of the fayetteville shale play for 2008-2012. Technical report, The University of Arkansas, 2008. URL <http://www.aogc.state.ar.us/sales/FayettevilleShaleEconomicImpactStudy2008.pdf>.
- D. Chen, S.-I. Cho, C. Chen, X. Wang, A. I. Damokosh, L. Ryan, T. J. Smith, D. C. Christiani, and X. Xu. Exposure to benzene, occupational stress, and reduced birth weight. *Occupational and environmental medicine*, 57(10):661–667, 2000.
- G. Chichilnisky and G. Heal. Global environmental risks. *The Journal of Economic Perspectives*, pages 65–86, 1993.
- C. R. Collier. Determination of the executive director concerning natural gas extraction activities in shale formations within the drainage area of special protection waters, 2009. URL <http://www.state.nj.us/drbc/library/documents/EDD5-19-09.pdf>.
- T. Considine, R. Watson, R. Entler, and J. Sparks. An emerging giant: Prospects and economic impacts of developing the marcellus shale natural gas play. *The Pennsylvania State University*, 2009.
- T. J. Considine, R. Watson, and S. Blumsack. *The economic impacts of the Pennsylvania Marcellus shale natural gas play: an update*. The Pennsylvania State University Department of Energy and Mineral Engineering, 2010.
- W. M. Corden and P. J. Neary. Booming sector and de-industrialisation in a small open economy. *The economic journal*, 92(368):825–848, 1982.

- H. Cutler and S. Davies. The impact of specific-sector changes in employment on economic growth, labor market performance and migration. *Journal of Regional Science*, 47(5): 935–963, 2007.
- A. Dixit. Irreversible investment with price ceilings. *Journal of Political Economy*, 99(3): 541–557, 1991.
- A. K. Dixit and R. S. Pindyck. *Investment under uncertainty*. Princeton UP, Princeton, 1994.
- Drilling Contractor(2000). Alabama lawsuit poses threat to hydraulic fracturing across us, 2000.
- L. G. Epstein. Decision making and the temporal resolution of uncertainty. *International economic review*, pages 269–283, 1980.
- J. Feyrer, E. T. Mansur, and B. Sacerdote. Geographic dispersion of economic shocks: Evidence from the fracking revolution. *American Economic Review*, 107(4):1313–34, 2017.
- A. C. Fisher and U. Narain. Global warming, endogenous risk, and irreversibility. *Environmental and Resource Economics*, 25(4):395–416, 2003.
- T. Fitzgerald and R. R. Rucker. Us private oil and natural gas royalties: estimates and policy relevance. *OPEC Energy Review*, 40(1):3–25, 2016.
- S. Gopalakrishnan and H. A. Klaiber. Is the shale energy boom a bust for nearby residents? evidence from housing values in pennsylvania. *American Journal of Agricultural Economics*, 96(1):43–66, 2014.
- L. H. Goulder. *Environmental policy making in economies with prior tax distortions*. Edward Elgar Publishing, 2002.
- M. Grant. Local resolutions against fracking, 2017.

- W. M. Hanemann. Information and the concept of option value. *Journal of Environmental Economics and Management*, 16(1):23–37, 1989.
- C. Hausman and R. Kellogg. Welfare and distributional implications of shale gas. no. w21115. Technical report, 2015.
- C. Henry. Option values in the economics of irreplaceable assets. *The Review of Economic Studies*, pages 89–104, 1974.
- E. L. Hill. Shale gas development and infant health: evidence from pennsylvania. Working Paper. Available at, 2013.
- R. B. Jackson, A. Vengosh, J. W. Carey, R. J. Davies, T. H. Darrah, F. O’Sullivan, , and G. Pétron. The environmental costs and benefits of fracking. *Annual Review of Environment and Resources*, 39:327–362, 2014.
- G. D. Jacobsen and D. P. Parker. The economic aftermath of resource booms: evidence from boomtowns in the american west. *The Economic Journal*, 126(593):1092–1128, 2016.
- A. James and J. James. *A canary near a gas well: Gas booms and housing market busts in Colorado*. Work. Pap, 2014.
- A. James and B. Smith. There will be blood: crime rates in shale-rich us counties. *Journal of Environmental Economics and Management*, 84:125–152, 2017.
- T. C. Kinnaman. The economic impact of shale gas extraction: A review of existing studies. *Ecological Economics*, 70(7):1243–1249, 2011.
- A. Kohut, C. Doherty, M. Dimock, and S. Keeter. As gas prices pinch, support for oil and gas production grows: Those aware of fracking favor its use. *Washington, DC: Pew Research Center*. Retrieved from Link: [http://www. people-press. org/files/legacy-pdf/3-19-12% 20Energy% 20release. pdf](http://www.people-press.org/files/legacy-pdf/3-19-12%20Energy%20release.pdf), 2012.

- C. D. Kolstad. Fundamental irreversibilities in stock externalities. *Journal of Public Economics*, 60(2):221–233, 1996a.
- C. D. Kolstad. Learning and stock effects in environmental regulation: the case of greenhouse gas emissions. *Journal of environmental economics and management*, 31(1):1–18, 1996b.
- A. J. Krupnick and H. G. Gordon. What experts say about the environmental risks of shale gas development. *Agricultural and Resource Economics Review*, 44(2):106–119, 2015.
- D. Lemoine and C. Traeger. Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy*, 6(1):137–166, 2014.
- J. Loomis and M. Haefele. Quantifying market and non-market benefits and costs of hydraulic fracturing in the united states: A summary of the literature. *Ecological Economics*, 138:160–167, 2017.
- J. Manfreda. The real history of fracking, 2015. URL <http://oilprice.com/Energy/Crude-Oil/The-Real-History-Of-Fracking.html>.
- P. Maniloff and R. Mastromonaco. *The local economic impacts of unconventional shale development*. Work. Pap, 2014.
- P. Maniloff and R. Mastromonaco. The local employment impacts of fracking: A national study. *Resource and Energy Economics*, 49:62–85, 2017.
- P. Mensink and T. Requate. The dixit–pindyck and the arrow–fisher–hanemann–henry option values are not equivalent: a note on fisher (2000). *Resource and Energy Economics*, 27(1):83–88, 2005.
- L. Muehlenbachs, E. Spiller, and C. Timmins. The housing market impacts of shale gas development. *American Economic Review*, 105(12):3633–59, 2015.
- M. D. Partridge and D. S. Rickman. Regional computable general equilibrium modeling: a survey and critical appraisal. *International Regional Science Review*, 21(3):205–248, 1998.

- R. S. Pindyck. Irreversibility, uncertainty, and investment. *Journal of Economic Literature*, pages 1110–1148, 1991.
- J. D. Sachs and A. M. Warner. Natural resource abundance and economic growth. Technical report, National Bureau of Economic Research, 1995.
- O. Schenk, M. H. Lee, N. H. Paydar, J. A. Rupp, and J. D. Graham. Unconventional gas development in the us states: Exploring the variation. *European Journal of Risk Regulation*, 5(4):436–458, 2014.
- L. C. Scott and J. Huang. The economic impact of the haynesville shale on the louisiana economy in 2008. *Prepared for the Louisiana Department of Natural Resources*, 2009.
- S. B. C. Shonkoff, J. Hays, and M. L. Finkel. Environmental public health dimensions of shale and tight gas development. *Environmental Health Perspectives (Online)*, 122:8, 2014.
- R. Slama, O. Thiebaugeorges, V. Goua, L. Aussel, P. Sacco, A. Bohet, A. Forhan, B. Ducot, I. Annesi-Maesano, J. Heinrich, et al. Maternal personal exposure to airborne benzene and intrauterine growth. *Environmental health perspectives*, 117(8):1313, 2009.
- S. V. Smith. Oil # 3: How fracking changed the world [radio broadcast etpisode], 2016. URL <http://www.npr.org/templates/transcript/transcript.php?storyId=490375230>.
- A. O. . G. H. Society. Shooters – a fracking history, 2015. URL <http://aoghs.org/technology/hydraulic-fracturing/>.
- C. P. Traeger. On option values in environmental and resource economics. *Resource and Energy Economics*, 2014.
- A. Ulph and D. Ulph. Global warming irreversibility and learning. *The Economic Journal*, 107(442):636–650, 1997.

- P. J. Walsh, S. Bird, and M. D. Heintzelman. Understanding local regulation of fracking: a spatial econometric approach. *Agricultural and Resource Economics Review*, 44(2):138–163, 2015.
- J. G. Weber. The effects of a natural gas boom on employment and income in colorado, texas, and wyoming. *Energy Economics*, 34(5):1580–1588, 2012.
- J. G. Weber. A decade of natural gas development: The makings of a resource curse? *Resource and Energy Economics*, 37:168–183, 2014.
- B. L. Weinstein and T. L. Clower. Potential economic and fiscal impacts from natural gas production in broome county, new york. *report prepared for Broome County, NY, available at: <http://www.gobroomecounty.com/countyexec/broome-county-releases-natural-gas-economic-impact-study> (accessed 14th July, 2010)*, 2009.
- B. A. Weisbrod. Collective-consumption services of individual-consumption goods. *The Quarterly Journal of Economics*, 78(3):471–477, 1964.
- G. R. West. Comparison of input–output, input–output+ econometric and computable general equilibrium impact models at the regional level. *Economic Systems Research*, 7(2):209–227, 1995.
- J. P. Weyant. *The cost of the Kyoto Protocol: a multi-model evaluation*. The Energy Journal, 1999.
- I. S. Wing. *Computable general equilibrium models and their use in economy-wide policy analysis*. Technical Note, Joint Program on the Science and Policy of Global Change, MIT, 2004.
- R. Wobbekind and B. Lewandowski. *Hydraulic Fracturing Ban: The Economic Impact of a Statewide Fracking Ban in Colorado*. The Business Research Division, Leeds School of Business, The University of Colorado Boulder, 2014.

S. Zahran, S. Weiler, H. W. Mielke, and A. A. Pena. Maternal benzene exposure and low birth weight risk in the united states: A natural experiment in gasoline reformulation. *Environmental research*, 112:139–146, 2012.

CHAPTER 2

THE VALUE OF LEARNING ABOUT HYDRAULIC FRACTURING DAMAGES

2.1. INTRODUCTION

While hydraulic fracturing has the potential to bring substantial economic benefits (Section 1.3.2.2), there is tremendous uncertainty as to their potential to incur environmental and human costs (Section 1.3.2.1). In this chapter, we hypothesize that the ability to learn about these costs during a moratorium may explain why so many jurisdictions have temporarily banned fracking (Section 1.3.1), despite, in some cases, seemingly positive expected net local benefits. To empirically assess the influence of the QOV associated with uncertain cost on fracking bans, we develop a data-driven, finite horizon, dynamic program that incorporates Bayesian learning. We calibrate plausible beliefs about environmental damages and demonstrate the rationality of a ban in the face of positive expected net benefits. Specifically, we find that a policymaker anticipating $\sim \$113$ million in benefits per period but expecting $\$100$ million in damages per period – with a standard error of $\$60$ million – will optimally implement a ban if she expects to revisit the decision with a standard error of $\$57.8$ million.

The next two sections of the chapter present the model (Section 2.2) and results (Section 2.3). Section 2.4 discusses the results and concludes.

2.2. MODEL

To examine the role of uncertainty and learning in local fracking policy, we develop a dynamic information model that incorporates data-driven, economy-wide fracking benefits and uncertain damages that can become better understood over time. Economic benefits are calculated with a CGE model, a standard tool for empirical welfare analysis of policies that impact many sectors of an economy Wing (2004).

2.2.1. DYNAMIC LEARNING FRAMEWORK. The policymaker faces a decision in discrete time periods . In each period, a local policymaker chooses to ban or allow fracking. Specifically, she chooses $\chi_t \in \{0, 1\}$ where “0” denotes **BAN** and “1” denotes **FRACK**. The policymaker in t observes a $(1 \times t)$ vector of the full history of past decisions: $H_{t-1} = (\chi_0, \chi_1, \dots, \chi_{t-1})$. For example, a ban followed by two periods without a ban would be represented as $H_3 = (0, 1, 1)$. H_t evolves according to $H_t = H_{t-1} \hat{\cup} \chi_t$ where $\hat{\cup}$ indicates vector concatenation. The model requires carrying the history of decisions since the current period consumption depends on it.

Locality-wide consumption is represented as a series $\{C_{t+j}\}_{j=0}^{T-t}$ where C_{t+j} is the deterministic local consumption in period t . C_{t+j} depends on the prior fracking history – in particular, *if* and *when* fracking began. Baseline consumption is defined as the case in which fracking is always banned. If fracking is allowed, there is a surge in economy-wide consumption stemming from royalties on extracted resources. Economy-wide consumption in time t is a function of the history, H_{t-1} , and the present choice, χ_t : $C(H_{t-1}, \chi_t)$. Since the transition equation for history is, $H_t = H_{t-1} \hat{\cup} \chi_t$, it is written as $C(H_t)$.

Let η be the environmental damages of fracking, expressed in dollars. If fracking has not occurred, η is unknown to the policymaker and is modeled as a stochastic variable with a normal distribution. If, and when, fracking occurs the true value is revealed to the policymaker and is denoted η^* . Going forward, this distinction is important for interpreting equations and results. The parameters of the normal distribution are not known with certainty but decision-makers have a belief in each time period about the values of the mean and variance of the distribution. Learning from the noisy signal causes beliefs about to evolve so that in time t , $\eta \sim \mathcal{N}(\mu_t, \sigma_t^2)$. We use a constant relative risk-aversion (CRRA) utility function over net consumption, $C(H_t) - \eta$ and denote the coefficient of relative risk aversion as ρ .

2.2.1.1. Current Net Benefit Flow. As discussed above, consumption depends on if and when fracking begins, which is summed up in the history vector, H_t . Therefore, the consumption function maps elements from the set of feasible (i.e. monotonic) history vectors, denoted \mathcal{H} , into locality-wide consumption values: $C(H_t) \rightarrow C_t \forall H_t \in \mathcal{H}$. If fracking is allowed ($\chi_t = 1$), then the expected current flow of net benefits is the expected utility of the higher consumption level less damages, expressed as: $\mathbb{E}_\eta \left[\frac{(C(H_t) - \eta)^{1-\rho}}{1-\rho} \mid \mu_t, \sigma_t^2 \right]$, where \mathbb{E}_η is the expectation over damages, η . If fracking is banned ($\chi_t = 0$), the current flow of net benefits is the utility of baseline consumption: $\frac{(C(H_t) - 0)^{1-\rho}}{1-\rho}$. For succinctness, we write the utility function as $U(C(H_t), \eta\chi_t) = \frac{(C(H_t) - \eta\chi_t)^{1-\rho}}{1-\rho}$ where the choice variable controls whether or not there are damages. This implies that the expected current flow is $\mathbb{E}_\eta \left[U(C(H_t), \eta\chi_t) \mid \mu_t, \sigma_t^2 \right]$. Note the expectation over damages is trivial if the fracking ban is maintained – and there is no damages – or if fracking occurred in a previous period revealing the true damages and collapsing uncertainty: $(\mu_t, \sigma_t^2) = (\eta^*, 0)$.

2.2.1.2. *Information.* Once allowed, fracking results in a constant flow of health and environmental damages. Damages persist for ten 5-year periods including the period when fracking begins⁹ requiring the computation of 14 periods of consumption benefits. Learning brings better information about the distribution of η . In the initial period, the policymaker has prior beliefs about the mean, μ_0 , and variance σ_0^2 , of $\eta \sim \mathcal{N}(\mu_0, \sigma_0^2)$.

The flow of information is modeled as an observed, time-dependent, noisy signal (s) on the true (but unknown) value of damages, η^* . From the perspective of the policymaker, η^* has not been realized so the signal is expressed as¹⁰

$$(1) \quad s_t = \eta + \epsilon_t$$

where ϵ_t is a normally distributed i.i.d. random variable with mean $\mu_\epsilon = 0$ and variance σ_ϵ^2 . As the sum of two normally distributed i.i.d. random variables, s_t is normally distributed $s_t \sim \mathcal{N}(\mu_t, \sigma_t^2 + \sigma_\epsilon^2)$. Therefore, posterior beliefs are

$$(2) \quad \mu_{t+1} = \frac{\sigma_\epsilon^2 \mu_t + \sigma_t^2 s}{\sigma_\epsilon^2 + \sigma_t^2}, \quad \sigma_{t+1}^2 = \frac{\sigma_\epsilon^2 \sigma_t^2}{\sigma_\epsilon^2 + \sigma_t^2}$$

⁹We could alternatively model fracking damages as a one-time event. The crucial assumption is that conditional on fracking, future damages are exogenous from the perspective of the current decision-maker.

¹⁰Note that the signal is produced by a draw around the true damages so that the process generating the signal is $s_t = \eta^* + \epsilon_t$.

Note that $\lim_{t \rightarrow \infty} (\mu_t, \sigma_t^2) \rightarrow \infty$ provided $\sigma_\epsilon < \infty$. The rate of learning depends on the precision of future information, i.e. the variance of the signal noise. If σ_ϵ^2 is large, the signal is relatively uninformative and learning is slow. As σ_ϵ^2 shrinks, the signal becomes more informative and uncertainty is resolved faster.

2.2.1.3. Bellman Equation. The decision is posed as a recursive problem with three state variables. The first is the history of past decisions: $H_{t-1} = (\chi_0, \chi_1, \dots, \chi_{t-1})$. The other states characterize beliefs about the normally distributed damages: the mean (μ_t) and variance (σ_t^2). Irreversibility is modeled by a restricted choice set $\chi_t \in \{\chi_{t-1}, 1\}$, and we assume fracking has not yet occurred, $H_0 = \vec{0}$. Since the value function depends only on the state variables, and not explicitly on time, it can be written as a Bellman equation in time t as follows:

$$\begin{aligned}
(3) \quad V_t(H_{t-1}, \mu_t, \sigma_t^2) = & \max_{\chi_t \in \{\chi_{t-1}, 1\}} \left\{ \mathbb{E}_\eta \left[U(C(H_t), \eta \chi_t) \mid \mu_t, \sigma_t^2 \right] \right. \\
& + (1 - \chi_t) \beta \mathbb{E}_s \left[V_{t+1} \left(\underbrace{(H_{t-1}, 0)}_{H_t}, \underbrace{\frac{\sigma_\epsilon^2 \mu_t + \sigma_t^2 s}{\sigma_\epsilon^2 + \sigma_t^2}}_{\mu_{t+1}}, \underbrace{\frac{\sigma_\epsilon^2 \sigma_t^2}{\sigma_\epsilon^2 \sigma_t^2}}_{\sigma_{t+1}^2} \right) \mid \mu_t, \sigma_t^2 \right] \\
& \left. + \chi_t \beta \mathbb{E}_\eta \left[V_{t+1} \left(\underbrace{(H_{t-1}, 1)}_{H_t}, \underbrace{\eta}_{\mu_{t+1}}, \underbrace{0}_{\sigma_{t+1}^2} \right) \mid \mu_t, \sigma_t^2 \right] \right\}
\end{aligned}$$

The first term on the right-hand side of Equation 3 is the expected current flow of utility, conditional on beliefs, μ_t and σ_t^2 . The second and third terms describe the continuation value if fracking is banned or allowed, respectively. Although we are mainly interested in situations for which the option to ban remains (i.e. fracking has not occurred), the equation also depicts

the value in time t if fracking has already occurred. In this case, the choice set is $\chi_t \in \{1, 1\} \Rightarrow \chi_t = 1$. Moreover, the true value of damages, η^* , is realized in the period in which fracking occurred implying $\mu_t = \eta^*$ and $\sigma_t^2 = 0$. Then, if fracking has already occurred, the value function as of time t collapses to $V_t(H_{t-1}, \eta^*, 0) = \frac{(C_t(H_t) - \eta^*)^{1-\rho}}{1-\rho} + \beta V_{t+1}(H_t, \eta^*, 0)$

2.2.2. OPTION VALUES. Traeger (2014)¹¹ suggests a convenient way to summarize the determinants of optimal policy. For the current setting, this so-called *Quasi-Option Value Rule* can be summarized as follows:

$$(4) \quad \begin{aligned} &\mathbf{FRACK} \text{ if } NPV_t > QOV_t + SOV_t = V_t^{soph} \\ &\mathbf{BAN} \text{ if } NPV_t \leq QOV_t + SOV_t = V_t^{soph} \end{aligned}$$

V_t^{soph} is the *full value of sophistication*, QOV_t is the *quasi-option value*, SOV_t is the *simple option value*, and NPV_t is the present value of the expected net gain from fracking. All values are expressed in utility units and functional arguments are suppressed. V_t^{soph} captures the presumption that a fully sophisticated decision-maker would value both the ability to simply delay a project (SOV_t) and the ability to learn about the project (QOV_t) conditional on and during the delay.

Traeger (2014) shows that can be constructed from three present values: *learning*, *postponement*, and *now or never*. These can be defined in the context of our model as follows:

- $V_t^l(\cdot | \chi_t = 0)$: the present value of a **ban** by a policymaker who anticipates *learning*;

¹¹Building on Arrow and Fisher (1974), Henry (1974), and Hanneman (1989).

- $V_t^p(\cdot|\chi_t = 0)$: the present value of a **ban** by a policymaker who anticipates the ability to revisit the decision – to *postpone* it – but does not anticipate the ability to learn;
- $V_t^n(\cdot|\chi_t = 0)$: the present value of a **ban** to a policymaker who does not anticipate the decision will be revisited — a *now or never* perspective.

The respective value functions become

$$\begin{aligned}
 V_t^l(H_{t-1}, \mu_t, \sigma_t^2 | \chi_t = 0) &= U\left(C(H_t), 0\right) + \beta \mathbb{E}_s \left[V_{t+1}\left(H_t, \mu_{t+1}, \sigma_{t+1}^2\right) \mid \mu_t, \sigma_t^2 \right] \\
 (5) \quad V_t^p(H_{t-1}, \mu_t, \sigma_t^2 | \chi_t = 0) &= U\left(C(H_t), 0\right) + \beta \mathbb{E}_\eta \left[V_{t+1}\left(H_t, \mu, \sigma^2\right) \mid \mu, \sigma^2 \right] \\
 V_t^n(H_{t-1}, \mu_t, \sigma_t^2 | \chi_t = 0) &= U\left(C(H_t), 0\right) + \beta V_{t+1}\left(H_t, 0, 0\right)
 \end{aligned}$$

The value of **fracking** is, $\mathbb{E}_\eta \left[U\left(C(H_t), \eta\right) + \beta V_{t+1}\left(H_t, \eta, 0\right) \mid \mu_t, \sigma_t^2 \right]$ and is the same in each case: $V_t^l(\cdot|1) = V_t^p(\cdot|1) = V_t^n(\cdot|1)$. The first and second equations in Equation 5 differ in the stochastic variable over which the expected continuation values are calculated. The first equation takes the expectation of s , the signal, and anticipates updated beliefs about the damage distribution. In the second equation, no signal is anticipated, so beliefs do not change over time and there is only uncertainty over the damage parameter, η . Consequently, the state variables in this case are not time-dependent and the expectation is taken with respect to η rather than s . The right hand side of the equations need no superscripts since the arguments of the expectation and value functions indicate whether the decision maker: anticipates learning and the expectation is conditional on updated beliefs, does not anticipate learning and the expectation is conditional on static beliefs, or does not anticipate revisiting the decision.

Following Traeger (2014) we calculate $NPV_t = V_t^n(\cdot|1) - V_t^n(\cdot|0)$ and $V_t^{soph} = V_t^l(\cdot|0) - V_t^n(\cdot|0)$ and decompose the full value of sophistication into the option values:

$$(6) \quad \underbrace{V_t^l(\cdot|0) - V_t^n(\cdot|0)}_{\text{full value of sophistication}} = \underbrace{V_t^l(\cdot|0) - V_t^p(\cdot|0)}_{QOV_t} - \underbrace{V_t^p(\cdot|0) - V_t^n(\cdot|0)}_{SOV_t}$$

Using the Arrow-Fisher-Henry-Hanneman Quasi-option Value Rule (Equation 4) we can express our current period value function (Equation 3) as $V_t(\cdot) = \max \{NPV_t, QOV_t + SOV_t\}$ and see the impact of learning, captured by the QOV, on welfare. Since QOV_t is non-negative (Traeger 2014) and increasing with more precise information, the ability to learn weakly increases the value function in Equation 3.

2.2.3. PARAMETERIZATION OF DAMAGES. In order to solve the model, we also parameterize the initial beliefs about the distribution of the monetary value of fracking damages, η . Given additive separability, this can occur independently from the parameterization of economic benefits described above. Recall that η is normally distributed with initial beliefs about the mean and standard error equal to and μ_0 , σ_0 respectively. The magnitude of these damages must be weighed against the consumption benefits of fracking, calculated in Section 1.4.2.

To parameterize current beliefs about this distribution, we define a plausible range within which the monetary value of damages is likely to fall. First, we assume a 5% chance of negative damages (i.e. benefits) from fracking. This could occur if, for example, fracking allows natural gas to displace coal in local energy production, leading to cleaner air. At

the other extreme, we assume that damages could exceed a high-damage scenario with a probability of 5%.

To define the high-damage scenario, we calculate the monetary cost of purchasing water rights to permanently replace the current surface water supply. We assume the municipality would obtain shares in the Colorado-Big Thompson (C-BT) System, where water rights have sold for \$50,000 per acre-foot. To calculate the number of shares that the municipality would need to purchase, we use the Colorado City of Fort Collins's water use as an example. If Fort Collins Utilities had to purchase C-BT shares to cover 100% of its population of 150,000, it would need to purchase 51,805 acre-feet¹². Assuming the representative municipality of 50,000 people would use the same initial mix of C-BT and non-C-BT water sources and consumption per capita as Fort Collins, it would need to purchase 17,268 acre-feet from C-BT (valued at \$50,000 per acre-foot) for a total one-time cost of \$863 million. Amortized over 10 (5-year) periods with a discount rate of 5%, this high-damage estimate becomes approximately \$200 million per period.

The 5% upper (\$200 million) and lower (\$0) tails of fracking damages allow us to calculate the mean and standard error of the distribution. Figure 2.1 displays the distribution of fracking damages given the specified two tails. From Figure 2.1, we can see that the mean of the per-period fracking damages is ~\$100 million per period. The standard deviation of this distribution is \$60 million. Therefore, our initial beliefs are $\mu_0 = \$100$ million and $\sigma_0 = \$60$ million.

¹²Fort Collins Utilities (FTCU) delivers water to 130,200 people out of the approximately residents 152,061 (2013 estimate). Approximately 19% of FTCU's average raw water supply comes from the C-BT. The rest comes from the Poudre River, assumed polluted beyond use. If FTCU were to obtain the remaining 81% while scaling supplies to the entire Fort Collins population, this would require the purchase of 51,805 acre-feet of water from the C-BT.

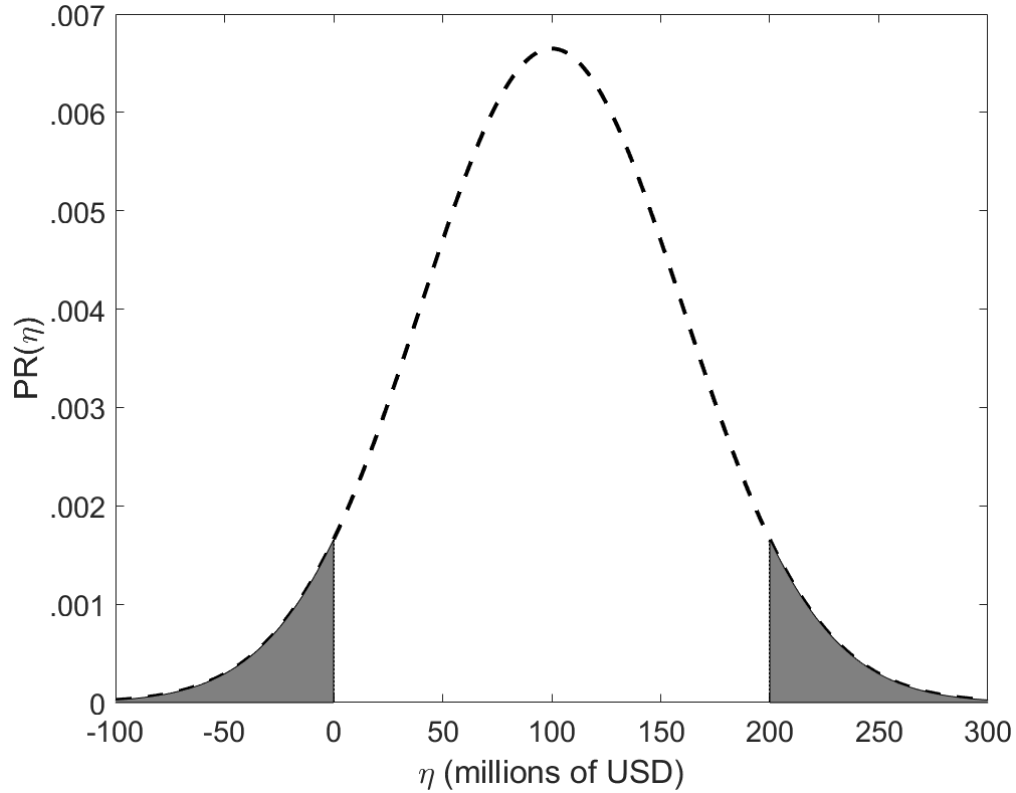


FIGURE 2.1: Distribution of Environmental Damages Given Calibrated Initial Beliefs

The x-axis is the value of the damage parameter, in millions of dollars. The y-axis is the probability. The 5% tails are shaded.

A policymaker who holds these calibrated initial beliefs thinks allowing fracking will incur a cost that is drawn from this distribution. If, on the other hand, fracking is banned, the policymaker anticipates receiving a signal that will provide information about the distribution of fracking damages.

Recall that the annuitized present value of the consumption benefits of fracking today is approximately \$113.6 million per period. This suggests that, in expectation, the net benefits of fracking are positive. Under a naive net-present value rule, a risk neutral policymaker would allow fracking because the expected benefits exceed the expected costs. Nevertheless, a policymaker that anticipates learning about the distribution of damages will wait before

allowing fracking if the QOV is sufficiently valuable. To find the conditions under which a policymaker would continue to ban fracking we solve the parameterized dynamic option value model using the rule described in Equation 4.

2.3. RESULTS

We first solve the model for a range of initial beliefs about damages and a range of assumptions about the rate of learning and examine if the ability to learn influences the optimal policy decisions. Then, we compare the role of learning to other factors such as the value of oil and gas reserves. Next, we quantify a monetary value of the QOV and use Monte Carlo simulations to show that faster learning leads to better decision-making over time. We finish with tests for the robustness of the results.

2.3.1. OPTIMAL POLICY. To explore the impact of learning on the optimal policy decision, we solve the model for a range of values for the standard error of the signal noise, presented in Equation 1. Specifically, we are interested in estimating Equation 4 which compares the NPV_t (the present value of the expected net increase in utility units from fracking) with estimates of QOV_t (*quasi-option value*) and (*simple option value*). If the NPV is greater (less) than the sum of QOV and SOV, then it is optimal to FRACK (BAN).

Using the results from the CGE model discussed above and the estimates for QOV and SOV, Figure 3 presents the optimal initial-period policy for fast ($\sigma_\epsilon = \$1$), medium ($\sigma_\epsilon = \200 million) ,and slow ($\sigma_\epsilon = \$500$ million) rates of learning as a function of initial beliefs about environmental damages. We also display the no-learning case. These curves mark the initial belief combinations where the policymaker is indifferent between allowing fracking and

maintaining the ban. Intuitively, as learning becomes faster, there are fewer combinations of initial beliefs (less area under the curve) for which fracking is optimal. The parameterized initial beliefs are labeled and it becomes clear that the optimal policy in the calibrated model is sensitive to the rate of learning. When learning is slow, the optimal policy is to allow fracking since new information is relatively uninformative and unlikely to change next period's policy. In this case, the expected gains from fracking dominate the value of learning in influencing the optimal policy. Under the parameterized beliefs, it is optimal to maintain a ban, receive precise information about fracking damages, and revisit the decision next period if learning is fast. The ban occurs despite positive expected net benefits from fracking, so a simplistic benefit-cost framework that does not anticipate learning — a *now or never approach* — would allow fracking. When $\sigma_\epsilon = \$211$ million, the risk averse policymaker is indifferent between fracking and banning given the parameterized initial beliefs. This implies that the policymaker should implement a ban if she anticipates at least a 7.5% reduction in σ_t^2 over the first 5 years (1 model period).

Notice that all three policy boundary curves converge to the same mean, $\mu_0 = \$111.8$ million as $\sigma_0 \rightarrow 0$. This occurs because information is only valuable under uncertainty. We denote this mean belief where policy switches under no uncertainty as $\hat{\mu}$ so that $\hat{\mu} = \$111.8$ million. Beliefs that $\mu_0 < \hat{\mu}$ with low initial uncertainty (southwest corner) will result in a **FRACK** policy. Similarly, beliefs that $\mu_0 > \hat{\mu}$ and the initial uncertainty is high (northeast corner) will result in a **BAN** policy. Increasing uncertainty while holding $\mu_0 = \hat{\mu}$ makes a ban more likely, due to the ability to learn about the true distribution. Even with no learning, risk aversion means that higher uncertainty can push towards a **BAN**.

Two other interesting features are visible. First, the annualized consumption benefits are \$113.6 million, which is larger than $1\hat{\mu} = \$111.8$ million. This difference arises because the intertemporal elasticity of substitution is constant and equal to $\frac{1}{\rho}$. As fracking today gives a larger increase in period 2 consumption than in period 1 consumption (See Figure 1.4), there is a greater than one percent change in the marginal rate of substitution. Consequently, the percentage change in the ratio of consumption $\frac{c_1}{c_2}$ is larger than one. Therefore, the risk-averse policymaker needs higher first period consumption benefits than the risk-neutral policymaker thus banning fracking at lower damages. Second, there is little effect on the optimal decision as the uncertainty increases enormously. This is because the baseline consumption of the community is approximately \$1.159 billion. The marginal utility at this point is $7.4445\text{e-}19$, which is quite flat and roughly risk-neutral.

2.3.2. LEARNING AND THE VALUE OF ENERGY RESOURCES. Figure 2.2 reveals that learning can play a pivotal role in the policy decision holding other economic factors constant. Now, we compare the impact of improved learning to changes in the resource value. We use the CGE model to compute the benefits of fracking for four oil and gas prices. Then, we fit a benefit function, $b(t, f, p)$, that maps current time period (t), when fracking began (f), and price of oil (p) into a dollar value of consumption benefits, (C): $b : (t, f, p) \rightarrow C$. With this, we populate a consumption benefit matrix for a range of oil prices from \$30 to \$90 per barrel in \$5 intervals and solve for both the value of learning and the value of fracking using Equation 5. The results are presented in Figure 2.3. The speed of learning is expressed as the number of signals required to reduce the standard error by half of its initial level. The figure highlights the policymaker's willingness to trade off faster learning for decreased economic

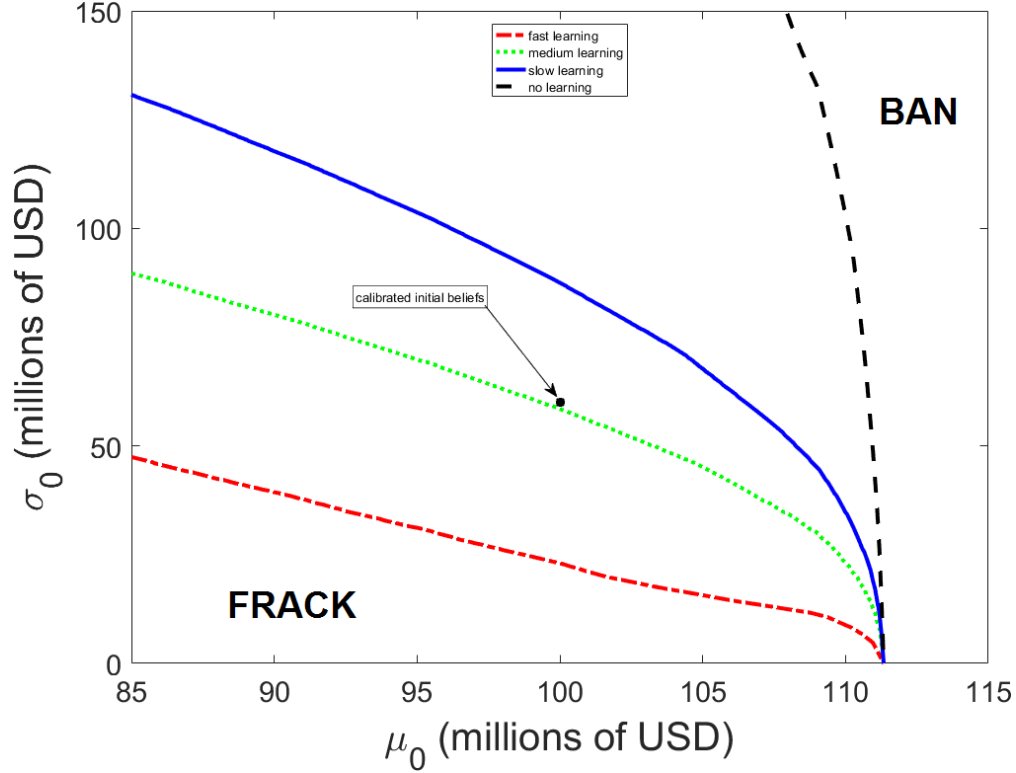


FIGURE 2.2: First Period Optimal Policy

The x-axis is the initial mean damage belief. The y-axis is the initial standard error belief. Each curve demarks where policy switches. Fast learning is $\sigma_\epsilon = \$1$, medium learning is $\sigma_\epsilon = \$200$ million, and slow learning is $\sigma_\epsilon = \$500$ million; and the initial beliefs used in the calibration exercise are $(\mu_0, \sigma_0) = (100, 60)$ million dollars. The slope of the line is the willingness to pay for lower uncertainty with decreased net benefits.

benefits through a lower reserve value. We display the policy boundary curve with calibrated beliefs as well as for higher and lower initial standard errors for the damages distribution.

Figure 2.3 shows that improvements in the rate of learning can tilt the decision towards a temporary ban, but that the price of oil – and the value of reserves – has a substantial influence on the decision. Improving the rate of learning only affects the first period decision in a small range of prices. This is true even in the high-initial variance case (dotted line Figure 2.3). On the other hand, a price change from \$45 to \$55 per barrel likely changes optimal policy in this context. In our parameterized scenario (solid black line with dot markers in Figure 2.3), a five period reduction in the number of periods required for the

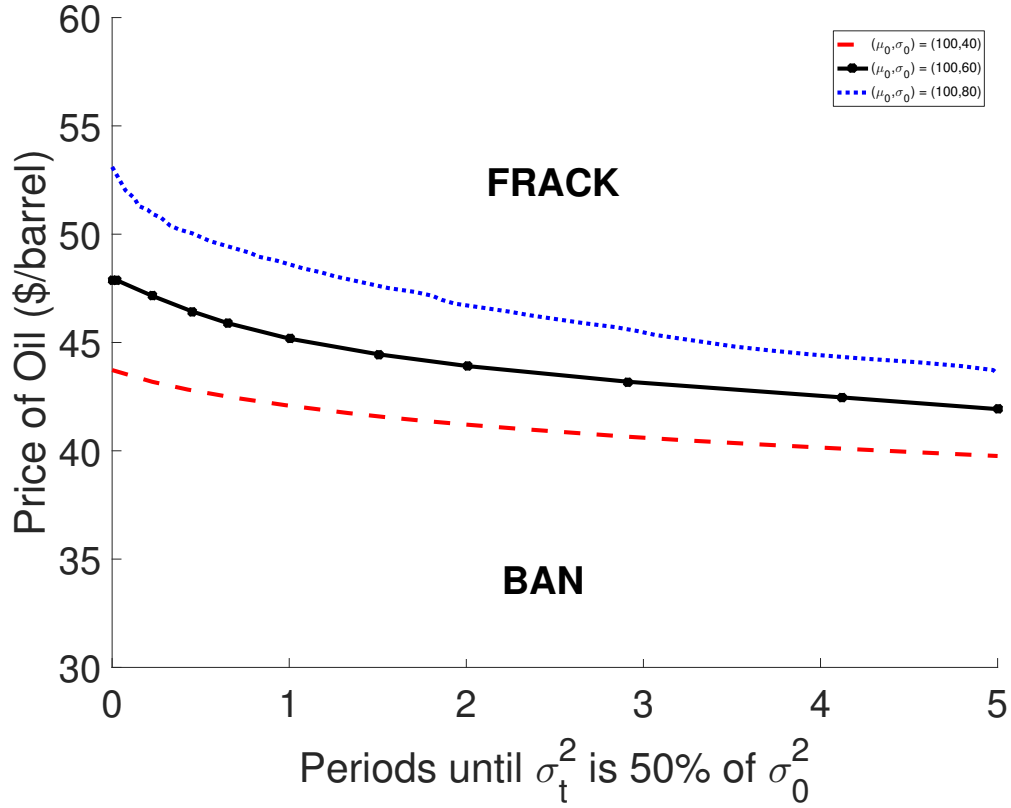


FIGURE 2.3: First Period Optimal Policies for 3 Different Initial Beliefs

The x-axis represents the speed of learning as the number of signals for $\sigma_t = .5\sigma_0$ where a time period is 5 years. The y-axis is the price of oil in dollars per barrel. The 3 curves are a mean-preserving spread. The slopes are the willingness to avoid a slower rate of learning in terms of reduced net benefits.

variance belief to reduce by half is equivalent to an increase in oil price from \$41.92 to \$47.88 per barrel – around a \$6 change. For comparison, between August 2015 and August 2016, oil prices ranged from ~\$30 per barrel to nearly \$55. This suggests that changes in price expectations consistent with existing oil price volatility could have a much larger impact on decision-making than improvements in the speed of learning. More generally, this suggests that the decision to frack or ban hinges on the local value of reserves, which, in turn, is influenced by the size of the reserve, the price of oil and gas, and local mineral rights

ownership. Our calibrated example shows that the ability to learn can play a pivotal role, provided the value of reserves falls within a relatively narrow range¹³.

2.3.3. QUANTIFYING THE OPTION VALUE. Despite the relative importance of the value of reserves, the calibrated locality-wide QOV remains large, even in comparison to the consumption benefits of fracking. In order to calculate a monetary value of the QOV, we solve the model under risk neutrality. This enables us to express all value functions and option values in monetary units. The more common approach is to solve the model under risk-aversion and transform welfare into dollars by dividing by the marginal utility. However, this approach is typically applied to scalar values and would distort the behavior of the value functions. To be specific, CRRA utility gives negative values and, of course, marginal utility is positive. This technique would yield negative dollar values. Obviously, we could get positive dollar values by -1 multiplication, but this would flip the curves and the focus of this paper is the behavior of the QOV under different rates of learning in relation to policy. Moreover, the monetary value of the QOV increases with risk aversion so these results represent a lower bound.

Figure 2.4 shows how the numerical value of the first-period QOV depends on the standard error of the signal (σ_ϵ). Recall that the QOV is the difference between V^{soph} and V^p (Equation 6). All three reflect present values denominated in initial-period monetary units. The QOV is largest (\$101.5 million - \$40.4 million = \$61.1 million) when σ_ϵ is smallest (learning is fast) and decreases, as learning slows. With risk neutrality, policy switches in our calibrated setting when $\sigma_\epsilon = \$177$ million (labeled in Figure 2.4). At this learning rate, the initial

¹³This also suggests that the option value associated with learning about the price of oil may be quite large compared to learning about damages. This should be explored in future work.

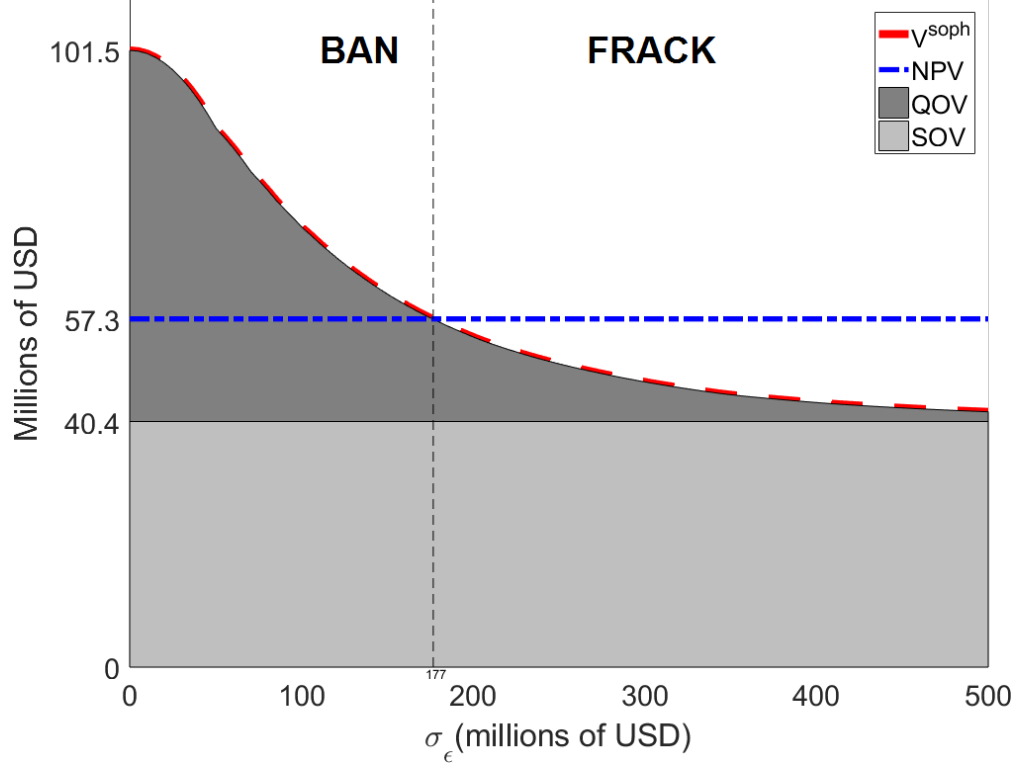


FIGURE 2.4: Value Functions and QOV under Risk Neutrality

The x-axis is the standard error of the noisy signal. The y-axis is the monetary value in millions of dollars. As learning becomes slower, the QOV decreases pulling the full value of sophistication, which approaches the SOV asymptotically. The maximum value of the QOV occurs where $\sigma_\epsilon \rightarrow 0$ and is 61.1 million dollars.

variance drops 10% after the first signal. Recall that the cutoff under risk aversion ($\rho = 2$) in Figure 2.2 was $\sigma_\epsilon = \$211$ million, showing that a risk neutral policymaker requires faster learning (all else equal) to justify a fracking ban.

These QOV represents the value of information acquired during a temporary ban and its numerical value indicates that it is an economically important social value. When learning is fast, the QOV represents 12.76% of the \$478.99 million in gross consumption benefits that fracking brings (in present value terms). To our knowledge, this is the first attempt to quantify a numerical, locality-wide QOV associated with an environmental policy decision.

2.3.4. POLICY SIMULATION. Thus far, our results indicate that faster learning increases the QOV, incentivizing a moratorium on fracking, in the initial period. However, this does not show that policymakers interested in supporting economic development activities should prefer slow learning. To illustrate this, we simulate policy decisions that occur after a first-period ban over a range of η^* that contains both instances where benefits exceed damages (fracking is optimal *ex post*) and instances where benefits are less than damages (banning is optimal *ex post*). This allows us to demonstrate that, as is intuitive, faster learning leads to better decision-making over time, both increasing fracking instances when it is beneficial and decreasing instances when it is not.

We use the base model ($\rho = 2$) and let the range of η^* go from \$60 million to \$180 million. This spans the calibrated $\hat{\mu}$ so that when η^* is greater (less) than $\hat{\mu}$ banning (fracking) is optimal *ex post*. The signals (described in Equation 1.2) depend on η^* even when initial beliefs are held constant. We consider the fast and medium learning rates defined in section 2.3.1 and presented in Figure 2.2. This means that the first period decision is always BAN ($H_1 = 0$) For each learning rate, σ_ϵ , and then for each η^* , we draw 1000, four-element, signal sequences, $\{s_t\}_{t=1}^{t=4}$, from the distribution $(\eta^*, \sigma_0^2 + \sigma_\epsilon^2)$. The calibrated beliefs (100, 60), in millions of dollars, are the initial damage beliefs (μ_0, σ_0^2) , which evolve over time according to Equations 2, depending on the random signal sequence. We evaluate and compare $V^{soph}(H_{t-1}, \mu_t, \sigma_t^2)$ and $NPV_t(H_{t-1}, \mu_t, \sigma_t^2)$ for $t = 2, 3, 4, 5$ to find the optimal policy in accordance with Equation 4.

The results of the simulations are presented in Figure 2.5. Results are displayed as the probability of making the correct (or incorrect) decision by the end of the 5-period decision horizon. In the left panel, damages are less than the benefits so fracking is beneficial. Under

fast learning, fracking occurs 94% of the time by period 2 while with medium learning, it takes all 5 periods before at least 79% of simulations result in beneficial fracking. Note that 97% of the fast learning simulations frack by period 5, the terminal decision period. Recall that the policy switches at $\hat{\mu} = \$111.8$ million but the annualized consumption benefits are \$113.6 million. The reason that 3% of the fast-learning simulations do not frack when it is beneficial is risk aversion. That is, regardless of the level of certainty, a risk averse policymaker with mean beliefs $\mu_t \in (\$111.8, \$113.6)$ will ban fracking under uncertainty even though it would bring an increase in net welfare *ex post*. The right panel presents the result of simulations in which damages exceed the benefits, so a fracking ban optimizes *ex post* welfare. In this case, under fast learning an initial ban results in the optimal decision in every instance (i.e., the probability of fracking is always zero). On the other hand, under medium learning, there is a 45% chance fracking will eventually be allowed, even though it would reduce *ex post* welfare.

Figure 2.5 illustrates that, despite incentivizing a moratorium in the first period, faster learning results in more fracking when it is beneficial and less when it is not. Therefore, a policymaker whose sole interest is economic development might optimally enact a moratorium.

2.3.5. SENSITIVITY ANALYSIS. In Figure 7, we explore the sensitivity of our conclusions to assumptions about initial damage distributions and risk aversion. First, we test robustness of the results to changes in initial beliefs. The results in Figure 4 present the policy boundary curves under a mean-preserving spread of the initial damage distribution. Here, we hold fixed the standard error and let the mean change. The left panel of Figure 7 shows that the main conclusions are robust to these changes. Conditional on the initial beliefs, increasing the




FIGURE 2.5: Simulation Results: Probability of Fracking In or Before each Period

The x-axis is the time period in which fracking begins. Each time period represents 5 years. The y-axis is the percentage of simulations in which fracking occurs in time, t . The left panel is the simulations for which η^* is less than the present value of consumption benefits. The right panel is the simulations for which η^* exceeds the present value of consumption benefits. In comparison to medium learning ($\sigma_\epsilon = \$200$ million), faster learning ($\sigma_\epsilon = \$1$) results in more fracking and sooner when it's beneficial and no fracking when it is not.

rate of learning has a relatively small impact on the policy decision. Initial beliefs about the mean do have a notable impact on the (still narrow) price range within which increasing the rate of learning plays a pivotal role in determining optimal policy.

In the right panel of Figure 7, we assess the impact of changing the coefficient of relative risk aversion. The panel replicates the base curve in Figure 4 using a range of values for ρ . As expected, increasing ρ raises the oil prices for which a ban is optimal. It also increases the importance of the rate of learning in influencing the policy decision, indicated by the steeper slope of the high risk aversion boundary curve in Figure 7. Despite this, even at high levels of risk aversion, the range of prices where learning is influential remains narrow.




FIGURE 2.6: Model Sensitivity

The x-axis is imputed number of signals until initial beliefs about standard error is reduced by half, a transformation of signal variance. The y-axis is the price of oil in dollars per barrel. The overall results hold under a variety of initial beliefs and risk aversion specifications.

2.4. DISCUSSION AND CONCLUSION

Uncertainty about fracking damages and the ability to learn create a QOV that can impact the economic rationale for imposing a temporary ban on fracking activities. In our calibrated setting, we show that a moratorium can be justified if beliefs about environmental damage variance are expected to drop at least 7.5% before the decision is revisited (or 10% for a risk neutral policymaker). Faster learning also leads to better decision-making over time. Though learning can influence optimal policy, we find that its role is relatively unimportant when compared to plausible (indeed historical) fluctuations in the price of oil.

Although we emphasize a model of fracking policy, the developed methodology expands the class of problems that can be quantitatively approached with an option value framework. Juxtaposing a detailed CGE model with a dynamic learning framework makes it possible to quantify the impact of uncertainty and learning within an empirically grounded general equilibrium setting. The approach could be useful in other policy contexts, including public infrastructure investment or public safety measures.

In addition to the policy-relevant observations above, several other policy implications can be drawn from the analysis. First, uncertainty may push local policymakers to temporarily ban fracking until better information about associated damages becomes available. Consider the 2005 Energy Policy Act, which amended the 1974 Safe Drinking Water Act to exclude fracking injection fluids (other than diesel fuels) from the EPA's oversight, while exempting extraction companies from disclosing the chemicals involved in fracking operations. This change could create public uncertainty about the safety of drinking water near fracking fluid storage sites and make dangers of fracking which makes adopting a ban more attractive *ceteris paribus*.

Next, the rate of learning influences local policy decisions in a context of uncertain fracking damages. A high rate of learning makes a first period ban more appealing but makes fracking, if beneficial, more likely in subsequent periods. The value function when the decision remains (Equation 3) is weakly increasing in the rate of learning, implying that faster learning cannot decrease welfare. Consequently, the public has an interest in reducing the noisiness around fracking information through, for example, research and improved industry transparency. Although the potential for learning could push a community to implement a temporary ban, it also creates the incentive to remove the ban if this is in their interest. Many policy options

exist to support the opportunity for learning. These include funding for scientific research on impacts, information provision that enables homeowners to better negotiate with oil and gas companies (see Timmins and Vissing 2014), encouraging municipalities to fund their own studies , and providing assistance with local impact studies.

Our quantification of the QOV highlights an intriguing dimension of local fracking policy. The information-revealing signal about fracking damages is a public good. The ability to ban fracking and learn from others' experiences in similar, perhaps nearby, regions implies a free-rider problem where local jurisdictions obtain the benefits of information without contributing to its production. The full value of sophistication represents the local jurisdiction's willingness to pay for the ability to ban fracking and learn but there are currently no institutions that allow for its capitalization.

While useful, the model presented here has some important limitations. First, we assume that fracking policy is a binary (yes/no) decision. Feasibly, policymakers could choose both when to frack and at what intensity. When decisions are adaptable over time, the ability to learn tends to increase the level of development in early periods (Karp and Zhang 2006). Allowing a small amount of fracking in certain areas of a given jurisdiction could result in very precise information about the true value of damages. Then, policymakers could adjust the amount of fracking to ensure optimality. This is similar to the result in Karp and Zhang (2006) that the ability to learn about climate sensitivity can increase early emissions levels. Despite this, binary policies such as local bans are common in practice and likely reflect political or legal constraints that prevent policymakers from employing more delicate instruments. Indeed, many bans have arisen through the blunt instrument of local referenda.

Second, we ignore the stochastic nature of energy prices leaving it for future work. In reality, policymakers also learn about the value of the reserves they control. If prices have an upward drift, for example, this would create a further incentive to wait before fracking is allowed and oil and gas reserves are exploited. Future work should consider the interaction between stochastic energy prices and uncertain environmental damages.

Another limitation is that consumption benefits do not capture distributional effects. It could be that the economic benefits accrue to a small fraction of the local population. Routine burdens, including noise and light pollution or increased traffic, tend to affect those most closely located to fracking operations (Gopalakrishnan and Klaiber 2014), but as Hill (2013) points out these are often socio-economically disadvantaged groups that may not receive the benefits from fracking. A mismatch between those that benefit and those that incur the costs from fracking is not considered here but future work should investigate how this could affect the local political economy of fracking policy decisions.

CHAPTER 2 REFERENCES

- S. Gopalakrishnan and H. A. Klaiber. Is the shale energy boom a bust for nearby residents? evidence from housing values in pennsylvania. *American Journal of Agricultural Economics*, 96(1):43–66, 2014.
- E. L. Hill. Shale gas development and infant health: evidence from pennsylvania. Working Paper. Available at, 2013.
- L. Karp and J. Zhang. Regulation with anticipated learning about environmental damages. *Journal of Environmental Economics and Management*, 51(3):259–279, 2006.
- C. Timmins and A. Vissing. Shale gas leases: Is bargaining efficient and what are the implications for homeowners if it is not? Working Paper, 2014.
- C. P. Traeger. On option values in environmental and resource economics. *Resource and Energy Economics*, 2014.
- I. S. Wing. *Computable general equilibrium models and their use in economy-wide policy analysis*. Technical Note, Joint Program on the Science and Policy of Global Change, MIT, 2004.

CHAPTER 3

THE VALUE OF LEARNING ABOUT HYDRAULIC FRACTURING BENEFITS

3.1. INTRODUCTION

Previously, we assumed that the stream of local consumption benefits was deterministic. In reality, the royalties driving these benefits are a percentage of revenues (Brown et al. 2016) and dependent on oil and gas prices, which follow random processes. Though the damages of fracking constitute environmental externalities, royalties accrue privately to the owners of mineral rights. The prospect of a future ban, technological improvements that make oil and gas an obsolete energy source, or the discovery of more economically accessible reservoirs elsewhere may drive the owners of mineral rights to view the development decision with a *now or never* framework. Because of this, they will not internalize the positive externality of timing development to coincide with a higher oil and gas prices, and consequently higher royalties. This constitutes a market failure and justifies a moratorium until prices rise.

Therefore, this chapter introduces uncertainty to economic benefits by allowing oil and gas prices to follow a geometric Brownian motion process. This creates an option value associated with banning fracking until prices are sufficiently high. Using a model parallel to the one in Chapter 2, we calculate an empirically-based, social QOV which can be compared to the QOV calculated in Chapter 2. We find that the QOV associated with the epistemological

uncertainty surrounding environmental damages, calculate in Chapter 2, and the QOV associated with stochastic uncertainty calculated in this chapter are remarkably similar, differing only by 6.1%.

Finally, basic economic principles indicate that extraction firms will ‘cap’ fracked wells if prices fall below average variable cost. This can severely impact local economic benefits but would not affect damages. This arguably more realistic model, which allows for stoppage of economic benefits, increases the QOV by 29.2%.

The next section (Section 3.2) reviews the literature on oil price dynamics. Section 3.3 presents the dynamic model, Section 3.4 discusses the results, and Section 3.5 concludes.

3.2. OIL PRICE DYNAMICS

The lion’s share of local economic benefits stem from royalty payments (Brown et al. 2017) which, in turn, depend on the price of oil. The appropriate way to model oil price fluctuations (see Figure 3.1) has been discussed at length. This discussion focuses on models that use pricing theory and continuous time processes.

3.2.1. CONTINUOUS TIME PROCESSES. Early real options approaches utilized geometric Brownian motion (GBM) in stochastic optimal control problems with applications to: development, management, and abandonment of exhaustible natural resources (Brennan and Schwartz 1985); firm shut-down conditions with stochastic price fluctuations (McDonald and Siegel 1985); and valuing claims on offshore petroleum leases (Paddock et al. 1988). GBM is a continuous-time, but nowhere differentiable, stochastic process that satisfies $dP = \alpha P dt + \sigma P dZ$. It is a specific form of Brownian motion, observed by Robert Brown in

1827 (Brown 1828) and developed mathematically by Einstein (1905). It has since become ubiquitous in stock pricing (Black and Scholes 1976) because under GBM: (1) the expected returns are independent of the stock prices, as is realistic (Hull 2006); (2) prices are strictly positive; (3) movements are discontinuous ‘jumps’; (4) calculations are relatively easy and solutions are often analytic and tractable. A possible drawback, though, is that the instantaneous standard deviation, σ , is constant, which embeds high volatility in predicted prices and implies (arguably unreasonably) high uncertainty.

Other works (Laughton and Jacoby 1995; Schwartz 1997), though, argue that a mean reverting process (MRP) is more appropriate. An Ornstein-Uhlenbeck process, developed to explain the motion of massive Brownian particles (Uhlenbeck and Ornstein 1930), is a common specification of an MRP: $dP = \eta(\bar{P} - P)dt + \sigma dZ$. In MRPs, there is a long-run equilibrium level (\bar{P}), or perhaps a historical trend, to which the stochastic element revert with speed η . This specification allows for negative prices and Dixit and Pindyck (1994) modify it into what is commonly called a Geometric Ornstein-Uhlenbeck (GOU) process $dP = P [\eta(\bar{P} - P)dt + \sigma dZ]$. GOU allows price to oscillate around \bar{P} which, unlike GBM, leads to bounded expectations and lower uncertainty. This model’s usefulness stems from observations that supply and demand need time to adjust to price shocks. Also, from an econometric standpoint, adopting a GOU rather than GBM process means estimating one less parameter.

GBM and GOU are extremes in terms of uncertainty and prediction. GBM has unbounded expectation and very high uncertainty whereas GOU has a finite expectation (\bar{P}) and very low uncertainty. Other, less extreme, processes have been proposed for the stylized facts of the Oil market. Gibson and Schwartz’s (1990) two-factor model, inspired by the Oil




FIGURE 3.1: Spot Price of West Texas Intermediate Crude in Cushing, OK

Source: U.S. Energy Information Administration

Derivatives Market, has price following GBM ($dp = \mu P dt + \sigma_1 P dZ_1$) and its convenience yield¹⁴ following a MRP ($d\delta = k(\alpha - \delta)dt + \sigma_2 dZ_2$)¹⁵. Pilipovic's (2007) models long run uncertainty by embedding a stochastic, long run equilibrium level (E where $dE = \mu_E E dt + \sigma_E E dZ_E$) into a GBM process for price ($dP = \eta(E - P)dt + \sigma_p dZ_P$). Dias and Rocha (1999) explain price trends toward the mean with unexpected changes by developing a one-factor mean reversion with jumps model. Here, price evolves as an MRP but can 'jump' abruptly with given probability ($\frac{dP}{P} = \eta(\bar{P} - P)dt + \sigma_P dZ_P + (\theta - 1)dq$).¹⁶ Hilliard and Reis (1998) extend Schwartz's (1997) three-factor model by including jumps in the asset price and Miller and Zhang (1996) model the influence of global factors (i.e. the Gulf War) with a GBM model that includes positive Poisson jumps in peacetime and negative Poisson jumps in wartime.

¹⁴The futures price, F_0 , of a commodity is given by $F_0 = S_0 e^{(c-y)T}$ where S_0 is the spot price, c is cost of carry (the interest rate less the income earned plus the storage costs), y is the **convenience yield**, and T is time until the delivery date.

¹⁵where Z_1 and Z_2 are correlated Wiener process.

¹⁶where $dq = 1$ with probability λdt and $dq = 0$ with probability $1 - \lambda dt$.

3.2.2. **ECONOMETRIC APPROACHES.** Postali and Picchetti (2006) econometrically evaluate the accuracy of these processes for describing oil price movements. They conclude that GBM is a good approximation as the average half-life of oil (between four to eight years) is long enough. But, Kaffel and Abid (2009) present a methodology for determining the best continuous-time stochastic process and conclude that GBM with jumps is the best for oil prices and GOU is appropriate for the convenience yield, the implied return on holding a commodity rather than a derivative product that arises through a no-arbitrage condition. Hahn et al. (2014) examine crude oil prices from 1990 to 2013. As they are unable to statistically determine if historic prices are mean-reverting or not, they support the rationale for a two-factor model. However, they also find that including data since 2005 (where previous work – i.e. Askari and Krichene 2008; Meade 2010 – undermining continuous-time models did not) supports a simple, one-factor GBM process. They conclude that long-term forecasting may be well done with one-factor GBM. Moreover, they estimate the drift and volatility parameters for the full data set to be $\mu = -0.0582$ and $\sigma = 0.2563$.

3.3. MODEL

This section develops the framework for incorporating oil and gas price volatility into the dynamic learning framework. We first discuss the incorporation of yearly price fluctuations into the five-year decision framework. Then, we review the mathematics of geometric Brownian motion (discussed in Section 3.2). Next, we demonstrate our strategy for imputing consumption benefits as a function of time-dependent oil prices. Finally, we develop the Bellman equation and present the associated option values.

3.3.1. **POLICY CHOICES AND HISTORIES.** Parallel to the problem studied in Chapter 2, the policymaker faces irreversible decisions (**Frack** or **Ban**) every five years. We again use the results generated by the CGE that was described in Section 1.4.2. However, the CGE model generates yearly consumption flows and much of the work on oil and gas prices movements estimates yearly drift (α) and volatility (σ) parameters. So, the model is constructed so that t can be interpreted as a year.

The model again requires carrying the history of decisions, which is a $(1 \times t)$ vector, since the current period consumption depends on it. As the decision to allow fracking will incur sunk costs, it is assumed to be irreversible. This permits only monotonic history vectors like $(0, 0, 1)$ but not $(0, 1, 0)$. We denote the set of monotonic vectors as \mathcal{H} so that $H_t \in \mathcal{H} \forall t = 1, 2, \dots, T$.

History will again be a state variable and again evolve by vector concatenation. However, defining this evolution is subtly different than before as we wish to define the model in yearly terms. A choice made in a decision period cannot be revisited until the next decision period and can be thought of as being repeated throughout the intervening five years. Supposing year t is a decision year, we make the definition $\chi_t * \vec{1}_{(1 \times 5)} \equiv \vec{\chi}_t$, where $\vec{1}$ represents a vector of ones. Then, history evolves as $H_t = H_{t-1} \cap \vec{\chi}_t$ provided t is a decision period. Also, we assume that $H_0 = 0$ so that the option to ban fracking is initially in place.

3.3.2. **PRICES.** Recall that the consumption benefits in Section 1.4.2 depend on the value of local unconventional reserves which is set by an average of oil and gas prices, henceforth denoted as P . The literature on oil prices, discussed in Section 3.2, suggests that geometric Brownian motion (GBM) is an appropriate way to model the stochasticity of oil prices (Hahn

et al. 2014). Below, we make use of Itô calculus to solve the stochastic differential equation (SDE) associated with GBM for the one-year probability density function. With this, we are able to numerically estimate a discrete time Markov chain according to Tauchen (1986). More information on SDEs and Itô calculus can be found in Oksendal (2000).

P is stochastic and follows geometric Brownian motion (GBM). The equation of motion is $dP_t = \mu P_t dt + \sigma P_t dz_t$ where z is an increment of the Wiener process meaning: 1) any change in z , Δz , that occurs during Δt satisfies $\Delta z = \epsilon \sqrt{\Delta t}$ where $\epsilon_t \sim N(0, 1)$, and 2) $\mathbb{E}[\epsilon_t \epsilon_s] = 0$ for $t \neq s$ (i.e. ϵ_t are serially uncorrelated). We can rewrite the stochastic differential equation as $\frac{dP_t}{P_t} = \mu dt + \sigma dz_t$ and then as an Itô integral: $\int_0^t \frac{dP_t}{P_t} = \int_0^t \mu dt + \int_0^t \sigma dz_t = \mu t + \sigma z_t$ assuming $z_0 = 0$. Since P_t is an Itô process, we use Itô's formula $d(\ln P_t) = \frac{dP_t}{P_t} - \frac{\sigma^2}{2} dt$. The Itô integral is then $\int_0^t d(\ln P_t) = \mu t + \sigma z_t - \frac{\sigma^2}{2} t = (\mu - \frac{\sigma^2}{2})t + \sigma z_t$. Therefore, $\ln P_t - \ln P_0 = \ln(\frac{P_t}{P_0}) = (\mu - \frac{\sigma^2}{2})t + \sigma z_t$ and $(\frac{P_t}{P_0}) = e^{(\mu - \frac{\sigma^2}{2})t + \sigma z_t}$. Finally, we have $P_t = P_0 e^{(\mu - \frac{\sigma^2}{2})t + \sigma z_t}$. This means that for any t , P_t is log-normally distributed with expected value $\mathbb{E}[P_t] = P_0 e^{\mu t}$ and variance $\mathbb{V}[P_t] = P_0^2 e^{2\mu t} (e^{\sigma^2 t} - 1)$ (Geometric Brownian Motion 2017). The parameters μ and σ are commonly referred to as the drift and volatility, respectively. Also, the probability distribution of P_t for any given initial price is given by

$$(7) \quad f(P_t; P_0, \mu, \sigma, t) = \frac{1}{P_0 \sigma \sqrt{2\pi t}} e^{-\frac{\ln P_t - \ln P_0 - (\mu - \frac{\sigma^2}{2})t}{\sigma \sqrt{2\pi t}}^2}$$

To solve the model numerically, we discretize prices into an m -element Markov chain $P \in P^1, P^2, \dots, P^m$, and the pdf described by equation 7 into a right-stochastic ($m \times m$) matrix \mathbf{T} according to Tauchen (1986).

3.3.3. CONSUMPTION. As in Section 2.2, locality-wide consumption is computed in the general equilibrium model as a yearly time series, $\{C_{t+j}\}_{j=0}^{T-t}$ where C_{t+j} is the deterministic local consumption in year $t + j$. For each possible policy path, we use the CGE model to compute local consumption as a function of the **initial** rental value of oil and gas reserves, denoted \hat{P} . That is, the CGE model generates consumption values, $C_{t+j}(\hat{P})$, that depends on the prior policy history – in particular, *if* and *when* fracking began – as well as the **initial** price \hat{P} . If fracking is allowed, there is a surge in economy-wide consumption stemming from royalties on extracted resources. If fracking is not allowed, the economy experiences normal growth. Therefore, economy-wide consumption in time t is a function of average oil and gas price, \hat{P} , and the history vector in time t and is written as $C(H_t, \hat{P})$. Regardless of the policy, it is economically sensible that the value of consumption at any time t depends on the sequences of consumption, and hence the path of price, prior to time t . Therefore, we approximate consumption in time t by assuming that the growth rate depends only on the previous period's price rather than the full history of price movements. For example, if the policymaker were to expect a price drop from \bar{P} to \underline{P} immediately after period one, then she would expect her period two growth rate to be the same as if the price did not drop (\bar{g}), but applied to the second period consumption value associated with \underline{P} (See Figure 3.2).

In general and for any policy path, the consumption in period t_{k+1} for price movement $P_k \rightarrow P_{k+1}$ is approximated as

$$(8) \quad C(t_{k+1}, P_{k+1}) = C(t_k, P_{k+1}) + \left[\frac{C(t_k, P_1) - C(t_{k+1}, P_k)}{C(t_k, P_k)} \right] C(t_k, P_{k+1}) \forall H_t \in \mathcal{H}$$

The effect of this approximation is to increase consumption in t_{k+1} if prices were previously higher ($P_{k+1} < P_k$) and decrease consumption if prices were previously lower ($P_{k+1} > P_k$).

Recalling that time and policy path are summed by the history in time t , H_t , consumption is a function the history in time t , H_t , and the price variable, P : $C(H_t, P)$.




FIGURE 3.2: How Consumption is Approximated

The x-axis is time and the y-axis is the dollar value of consumption. The figure shows consumption as a function of time for two different prices \bar{P} and \underline{P} . The policymaker expects that prices will move immediately after period one with some probability. Then, consumption in period three depends on the growth rate under the higher price, \bar{g} , and the consumption value under the lower price in period two. The effect is that the consumption is higher in period three than it would have been on the low price trajectory.

We discretize prices into an m -element Markov chain $P \in P^1, P^2, \dots, P^m$, and the pdf described by equation 7 into a right-stochastic $(m \times m)$ matrix \mathbf{T} according to Tauchen (1986) where element t_{ij} of matrix \mathbf{T} is the probability of that price will transition from state i to state j . Stochastic matrices have two properties that make them a powerful tool in stochastic dynamic programming. First, for any integer k and right-stochastic matrix \mathbf{T} , \mathbf{T}^k is also right-stochastic. Second, the probability of the random element transition from state i to state j in k -steps is the ij -th element of \mathbf{T}^k .

Consumption is approximated using Equation 8 and we generate an $m \times m$ matrix, \mathbf{C} , where element $[c_{ij}]$ is the period $t+1$ consumption for j^{th} price $p_{t+1}^j \in P^1, P^2, \dots, P^m$ conditional on the i^{th} t -period price. The construction of \mathbf{T} and \mathbf{C} mean that only elements on the diagonal of the product are true expectations. To understand this, consider a simple 2-price situation. Then $\mathbf{T} = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}$ where p_{12} is the probability that prices will move from state 1 to state 2 in one year. As well, $U(\mathbf{C}) = \begin{bmatrix} u_{11} & u_{21} \\ u_{12} & u_{22} \end{bmatrix}$ where u_{12} is the utility of the consumption that is generated when the price move *from state 2 to state 1*. Matrix multiplication gives $\mathbf{T} \cdot U(\mathbf{C}) = \begin{bmatrix} p_{11}u_{11}+p_{12}u_{21} & p_{11}u_{21}+p_{12}u_{22} \\ p_{21}u_{11}+p_{22}u_{21} & p_{21}u_{21}+p_{22}u_{22} \end{bmatrix}$. Obviously, only the diagonal elements are expectations. So for the general case, the approximate expected utility is computed as $\mathbb{E}_{P_{t+j}}[U(\mathbf{C})] = \text{diag}(\mathbf{T}^j \cdot U(\mathbf{C}))$ for $j = 0, 1, 2, 3, 4$.

3.3.4. CURRENT NET BENEFIT FLOW. Typically, the current flow occurs between the current decision and the next. However, decisions occur every five years whereas the CGE model generates yearly consumption flows and much of the work on oil prices estimates yearly drift (μ) and volatility (σ) parameters. So, working with years rather than periods, the current flow is the expected present value of five years of consumption benefits. To calculate this, we use a constant relative risk-aversion (CRRA) utility function over net consumption, $C(H_t, P) - \bar{\eta}$ with deterministic damages, annual discount factor β , and transition matrix \mathbf{T} as described above.

If fracking is allowed ($\chi_t = 1$), then the expected current flow of net benefits is the expected utility of the higher consumption level less deterministic damages, expressed as $\sum_{j=0}^4 \beta^j \mathbf{T}^j \frac{(C(H_t, P) - \bar{\eta})^{1-\rho}}{1-\rho}$. If the ban remains in place ($\chi_t = 0$), the current flow is the expected utility of the baseline consumption, $\sum_{j=0}^4 \beta^j \mathbf{T}^j \frac{(C(H_t, P) - 0)^{1-\rho}}{1-\rho}$.

Using the current choice χ_t to control whether or not there are damages, the utility function is succinctly written as $U(C(H_t, P), \chi_t \bar{\eta}) = \sum_{j=0}^4 \beta^j \mathbf{T}^j \frac{(C(H_t, P) - \chi_t \bar{\eta})^{1-\rho}}{1-\rho}$ the current flow is $\sum_{j=0}^4 \beta^j \mathbf{T}^j U(C(H_{t+j}, P), \chi_t \bar{\eta})$.

3.3.5. THE BELLMAN EQUATION. The problem is recursively posed over two state variables. The first is the prior history of decisions, H_{t-1} , as in Section 2.2. The second is the average price of oil and gas, P , which follows GBM as described above in Section 3.3.2. Irreversibility is, again, modeled by the restricted choice set, $\chi_t \in \{\chi_{t-1}, 1\}$ assuming a ban is initially in place ($H_0 = \vec{1}$). Since the value function depends only on state variable, and not on time, it is stationary and can be written as a Bellman equation in time t as

$$(9) \quad V_t(H_{t-1}, P) = \max_{\chi_t \in \{\chi_{t-1}, 1\}} \left[\sum_{j=0}^4 \beta^j \mathbf{T}^j U(C(H_{t+j}, P), \chi_t \bar{\eta}) + \beta^5 \mathbf{T}^5 V_{t+5}(H_{t+5}, P) \right]$$

The first term on the right hand side captures the present value of the expected utility stream generated by the decision today χ_t , which is concatenated into H . If the ban remains in place, $\chi_t = 0$, then the second term on the right hand side captures the expected value of being able to revisit the decision in year $t + 5$ after observing the average price of oil and gas P_{t+5} .

3.3.6. OPTION VALUES. Equation 4 gives Traeger's (2014)¹⁷ so-called *Quasi-Option Value Rule*. We now follow Section 2.2.2 and construct present three present values: *learning*, *postponement*, and *now or never*, which we will subsequently use to calculate the *QOV*. There

¹⁷Building on Arrow and Fisher (1974); Henry (1974); Hanemann (1989)

are subtle differences between the present values of this chapter and the present values presented in Section 2.2.2. These differences arise because the model in Chapter 2 concerns *epistemological uncertainty* whereas the model in here addresses *stochastic uncertainty*. In the context of this chapter, the present values are:

- $V_t^l(\cdot|\chi_t = 0)$: the present value of a **ban** by a policymaker who anticipates observing P prior to the subsequent decision;
- $V_t^p(\cdot|\chi_t = 0)$: the present value of a **ban** by a policymaker who anticipates the ability to revisit the decision but does not anticipate observing P first;
- $V_t^n(\cdot|\chi_t = 0)$: the present value of a **ban** to a policymaker who does not anticipate the decision will be revisited.

The respective value functions become

$$\begin{aligned}
 V_t^l(\cdot|0) &= \sum_{j=0}^4 \beta^j \mathbf{T}^j U(C(H_{t+j}, P), 0) + \beta^5 \mathbf{T}^5 V_{t+5}(H_{t+5}, P) \\
 (10) \quad V_t^p(\cdot|0) &= \sum_{j=0}^4 \beta^j \mathbf{T}^j U(C(H_{t+j}, P), 0) + \beta^5 \max \left\{ \begin{aligned} &\mathbf{T}^5 \sum_{j=0}^{T-t} \beta^j \mathbf{T}^j U(C(H_{t+j}, P), \bar{\eta}), \\ &\mathbf{T}^5 V_{t+5}^p(H_{t+5}, P) \end{aligned} \right\} \\
 V_t^n(\cdot|0) &= \sum_{j=0}^{T-t} \beta^j \mathbf{T}^j U(C(H_{t+j}, P), 0)
 \end{aligned}$$

The value of **fracking** is: $V_t(\cdot|\chi_t = 1) = \sum_{j=0}^{T-t} \beta^j \mathbf{T}^j U(C(H_{t+j}, P), \bar{\eta})$ and is the same in each case: $V_t^l(\cdot|1) = V_t^p(\cdot|1) = V_t^n(\cdot|1)$. Notably, the second equation, the value of postponement, is non-stationary since the expectation is inside the max operator. To address this, we solve for V_t^p from the perspective of all the time periods prior to t . To understand this, consider the final period decision $\max\{TVF, TVB\}$ where TVF and TVB are the terminal values of fracking and banning, respectively. In the penultimate decision period, a policymaker who does not anticipate learning views the value of postponement as $V_t^l(\cdot|\chi_t = 0) = \sum_{j=0}^4 \beta^j \mathbf{T}^j U(C(H_{t+j}, P), 0) + \beta^5 \max\{\mathbf{T}^5 TVF, \mathbf{T}^5 TVB\}$. So, we calculate the value of fifth period postponement from the perspective of period 4 and make the definition: $\max\{\mathbf{T}^5 TVF, \mathbf{T}^5 TVB\} \equiv V_{5,4}^p$. Similarly, $V_{5,3}^p \equiv \max\{\mathbf{T}^{10} TVF, \mathbf{T}^{10} TVB\}$ and $V_{5,2}^p \equiv \max\{\mathbf{T}^{15} TVF, \mathbf{T}^{15} TVB\}$.

As in Section 2.2.2, we calculate $NPV_t = V_t^n(\cdot|1) - V_t^n(\cdot|0)$ and $V_t^{soph}(\cdot|0) = V_t^l(\cdot|0) - V_t^n(\cdot|0)$ and decompose the full value of sophistication into the option values with Equation 6 and calculate the QOV.

3.4. RESULTS

We solve the model over a range of drift and volatility parameters with an initial price of oil of $P = 45 \text{ \$}/\text{bbl}$. Next, we use the drift and volatility specifications for GBM estimated by Hahn et al. (2014) and the one year value of our damage calibration from Section 2.2.3, $\bar{\eta} \cong \$21$ million, to solve for the present values described by Equation 10, the price point at which policy switches, and the maximal QOV. Then, with same drift, and volatility specifications, we solve the model over a range of $\bar{\eta}$ s for the price at which policy switches. This allows us to estimate a willingness to trade higher economic benefits (via higher P) for higher damages.

As basic economics indicates that firms will temporarily shut-down operations if the sale price falls below the average variable cost, we repeat these exercises for the case in which firms stop extraction (after fracking the well) in response to a low price of oil. This behavior is supported, anecdotally, from experience in the Bakken play in North Dakota and Devonian Shale play in Oklahoma. In early 2016, several development companies – including the giants Chesapeake Energy, Continental Resources (commonly called the “King of the Bakken” by insiders) and Whiting Petroleum – ceased operations (and hence royalty payments) when oil prices fell below 30 \$/bbl. During this stoppage, North Dakota and Oklahoma both projected shortfalls of a billion USD¹⁸.

3.4.1. QOV AND POLICY. We first solve the model over a range of drift, $\mu \in [1, 1]$, and volatility, $\sigma \in (0, 1]$, parameters for $P = 45$ \$/bbl. We solve for the first period present values of Equation 10 and the optimal policy using Equation 4 and present the results in Figure 3.3. Recall that drift is the expectation of the rate of increase (east side) or decrease (west side) and that high volatility (north side) means high risk. Therefore, a policymaker who is not risk-loving will wish to ban fracking when there is high risk and the expectation is that prices will decrease (northwest) and allow fracking when there is little risk and the expectation is that prices will increase (southeast). Also, the shaded area is where fracking is optimal and the QOV contour lines are in millions of USD. The maximal QOV occurs in the area where banning is optimal, as in the Chapter 2 model, and where the drift is $\mu = -0.0047$ and volatility is $\sigma = 0.3444$. For future purposes, we have also indicated Hahn et al.’s (2014) parametric estimates of the drift and volatility parameters for the GBM of oil prices.

¹⁸DESMOG

Figure 3.3 has several interesting implications. First, we can see the policymaker's willingness to trade higher expected returns for more risk in the slope of the outer line demarking where the optimal policy switches. We should also note that the policymaker finds it optimal to frack even for negative drift values, provided there is low volatility. In fact, these expectations incentivize fracking immediately in order to garner the most benefits. This brings us to another interesting feature – the 'bubble' in the 'Frack Zone'. Here, the low-risk expectation that prices will rise incentivizes delaying fracking in order to maximize the value of the reservoir and, consequently, the economic benefits (see Figures 1.4 and 1.5). However, further increasing the drift parameter while holding volatility constant increases the expected benefits enough to offset the patience of policymaker. Next, comparing Figure 3.3 to Figure 2.3 reveals the impact of including price movements. Figure 2.3 indicates that policymaker would allow fracking if $P = 50 \text{ \$/bbl}$ and damages are expected to be \$100 million for even the fastest rate of learning. However, Figure 3.3 shows that if $P = 50 \text{ \$/bbl}$, damages are expected to be \$100 million, and the policymaker is only *slightly* more optimistic than Hahn et al. (2014) (i.e. a drift of $\mu = -.050$ instead of $\mu = -.058$), she will choose to allow fracking. Finally, perhaps the most powerful outcome of our methodology is the ability to quantify the QOV. The maximal value here is \$78.3 million, 27% higher than the \$61.1 million from the fastest rate of learning in Figure 2.4.

3.4.2. PRESENT VALUES AND POLICY. To present policy results in the state space, we select Hahn et al.'s (2014) parameterization and solve the model over a range of prices and display the results in Figure 10. We can immediately see the impact of using the *Quasi-Option Value Rule* (Equation 4) rather than the *Net Present Value Rule*. The NPV rule

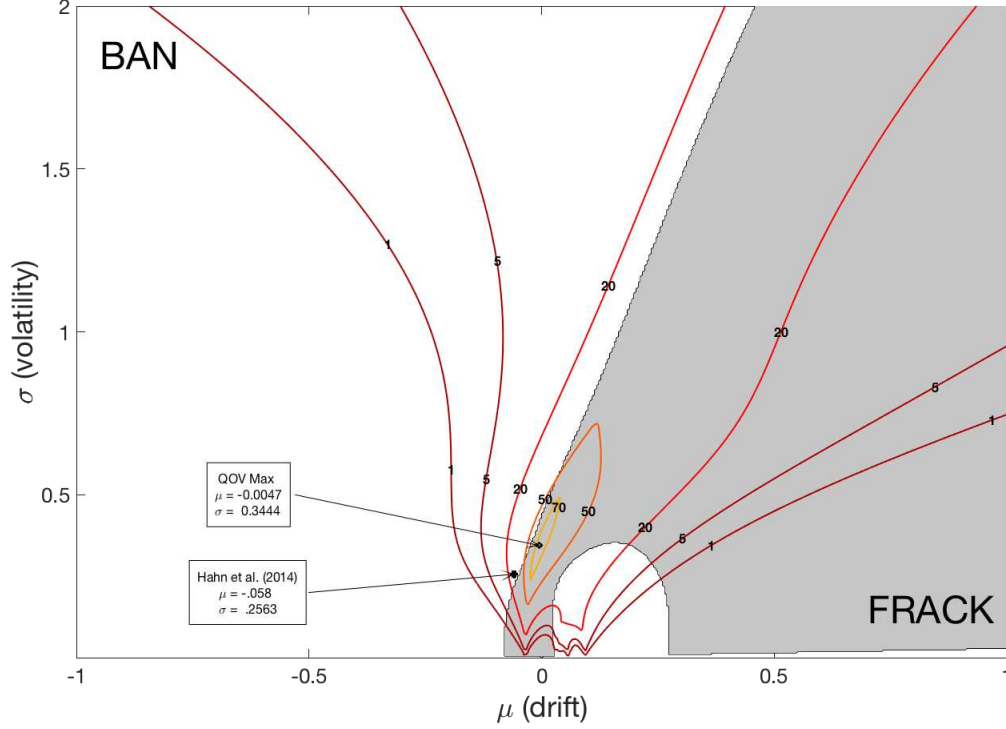


FIGURE 3.3: QOV contour lines and Policy for $P = 50 \$/bbl$

The x-axis is the drift parameter, the expected rate of change of P , and the y-axis is the volatility of the diffusion process. Price is set to be $50 \$/bbl$. The QOV contour lines are in millions of USD and the shaded area is where the optimal policy is to allow fracking. We also indicate Hahn et al.'s (2014) estimated GBM drift and volatility parameters for oil price movements.

would suggest fracking whenever $P \geq 10 \$/bbl$ since the expected present value of fracking, $V_0(\cdot|1)$, is positive in this range. On the other hand, the QOV rule suggests banning fracking unless $P \geq \tilde{P}$ since the present value of learning, $V_0^l(\cdot|0)$, exceeds $V_0(\cdot|1)$ when $P \in [10, 50.77) \$/bbl$. The value function (Equation 9) can also be written using option values as $V_t(\cdot) = \max \{NPV_t, QOV_t + SOV_t\}$ and is displayed as the upper envelope of $V_0(\cdot|1)$ and $V_0^l(\cdot|0)$. The figure also indicates the maximal QOV of ~ 64 million USD, 4.7% higher than the ~ 61.1 million USD (see Figure 2.4). This figure highlights that uncertainty over economic benefits could sway a policymaker to ban fracking despite expected positive net economic benefits. However, in contrast to the results of Figure 2.4, the QOV is highest when the policy is to allow fracking.

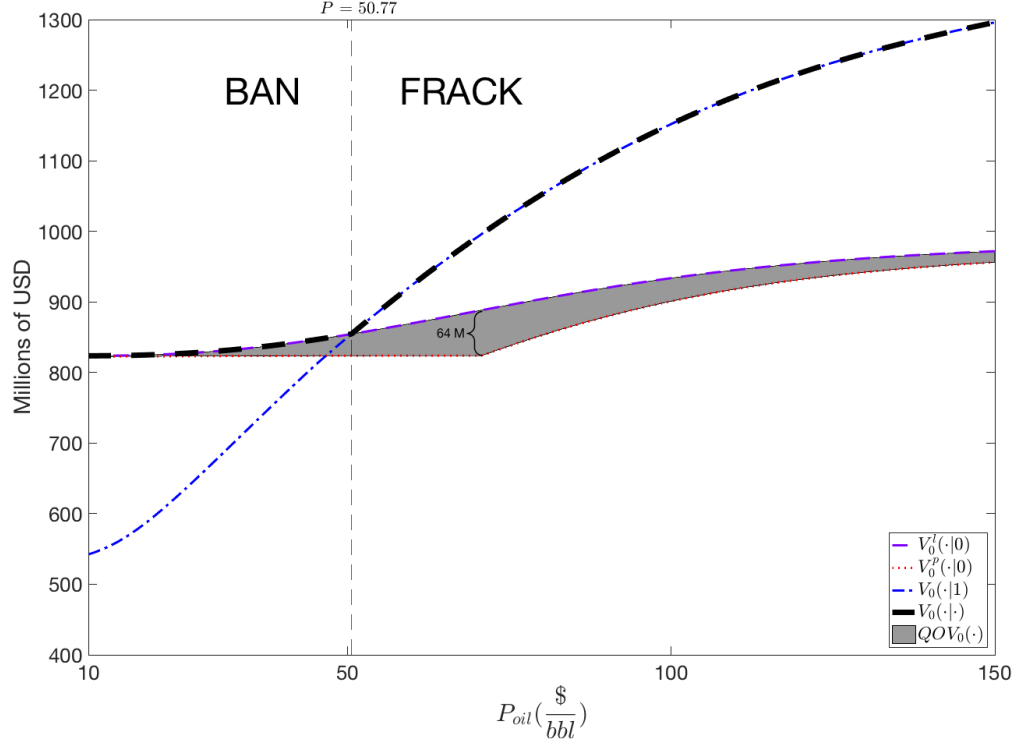


FIGURE 3.4: Present Values, Value Function, and QOV

The x-axis is the price of oil in USD per barrel and the y-axis is millions of USD. The QOV is the shaded area between V^{soph} and V^p and reaches a maximal value of $\sim \$64\text{M}$ USD. The policy is to ban fracking for oil prices below $50.77 \$/bbl$ and allow fracking otherwise.

3.4.3. LEARNING AND THE COSTS OF ENVIRONMENTAL DAMAGES. To understand the trade off between economic benefits and damages, we solve the model for a range of damages $\bar{\eta} \in (0, 50]$ million USD per year and find the price at which first period policy switches. We display the results in Figure 3.5. When damages are relatively low and the price is high (northwest) the policymaker will wish to allow fracking. Conversely, when damages are high and prices are low (southeast) banning fracking is optimal. The figure displays the policy switch prices as a scatter plot. To understand the trade off, we first fit a linear model (displayed as a dashed line) to the data. This model suggests that the policymaker is willing to accept $\$1$ million more in damages per year if the initial price of oil is $2.88 \$/bbl$ higher; or perhaps more intuitively, an increase in initial price of $1 \$/bbl$ is roughly equivalent to $\$347,200$. However, the linear model is disconcerting in that it implies some amount of

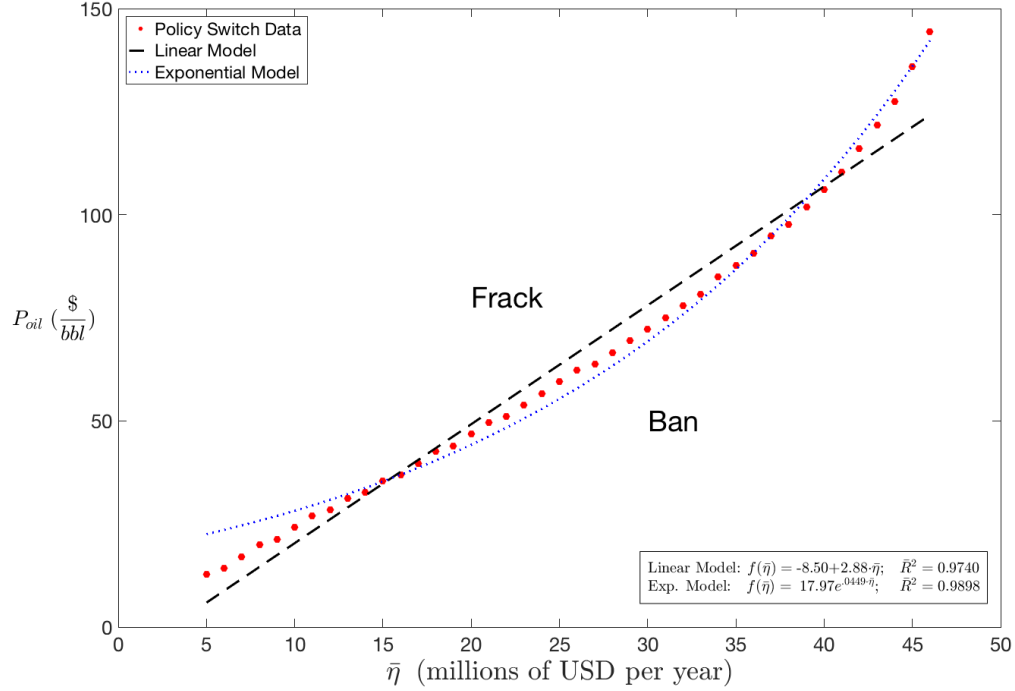


FIGURE 3.5: Policy Switch Curve

The x-axis is the cost of damages in millions of USD and the y-axis is the initial price of oil in $\$/bbl$. The data is fit with two models: linear and exponential.

positive damages is acceptable if the price of oil is $0\$/bbl$. Moreover, a visual inspection of the data reveals convexity in the tail. So, we also fit an exponential model (displayed as a dotted curve). Visually, this seems a much better match to the convexity of the tail and has higher adjusted R-squared. The coefficient on $\bar{\eta}$ is .0449 meaning that the policymaker is willing to accept \$1 million more in damages per year if the initial price is $\sim 45\%$ higher or that a 1% increase in the price of oil is roughly equivalent to \$22.2 million more in damages per year.

3.4.4. WELL SHUT-DOWN. We next solve the model with a well shut-down condition. Basic economics indicates that firms will temporarily shut-down operations if the sale price falls below the average variable cost. As the operational costs of a well vary by company, fracking play, and individual location, we assume an average shut-down price of $40\$/bbl$.




FIGURE 3.6: Policy Switch Curve with Well Shut-down

The x-axis is the price of oil in USD per barrel and the y-axis is millions of USD. The QOV is the shaded area between V^{soph} and V^p and reaches a maximal value of $\sim \$83$ M USD. The policy is to ban fracking for oil prices below $67.87 \$/bbl$ and allow fracking otherwise.

When the price of oil falls below this, the economy returns to its normal growth path (see Figure 1.4). If the price rises, fracking begins again. The results are presented in Figure 3.6. Now, the policymaker wishes to ban fracking for any price at or below $\tilde{P} = 67.87 \$/bbl$, a 33.7% increase with respect to the solution without the well shut-down condition. For reference, the price of oil has not been that high since November 2014. Furthermore, the maximal QOV has increased to $\sim \$83$ million, 35.8% higher than in Figure 2.4.

As before, we also solve the model over a range of damage parameters in order to estimate a willingness to trade economic benefits for damages. We present the results as a scatter plot in Figure 3.7 and compare to Figure 3.7. The linear model coefficient of 2.88 indicates that policymaker is willing to accept \$1 million more dollars of damages in exchange for P

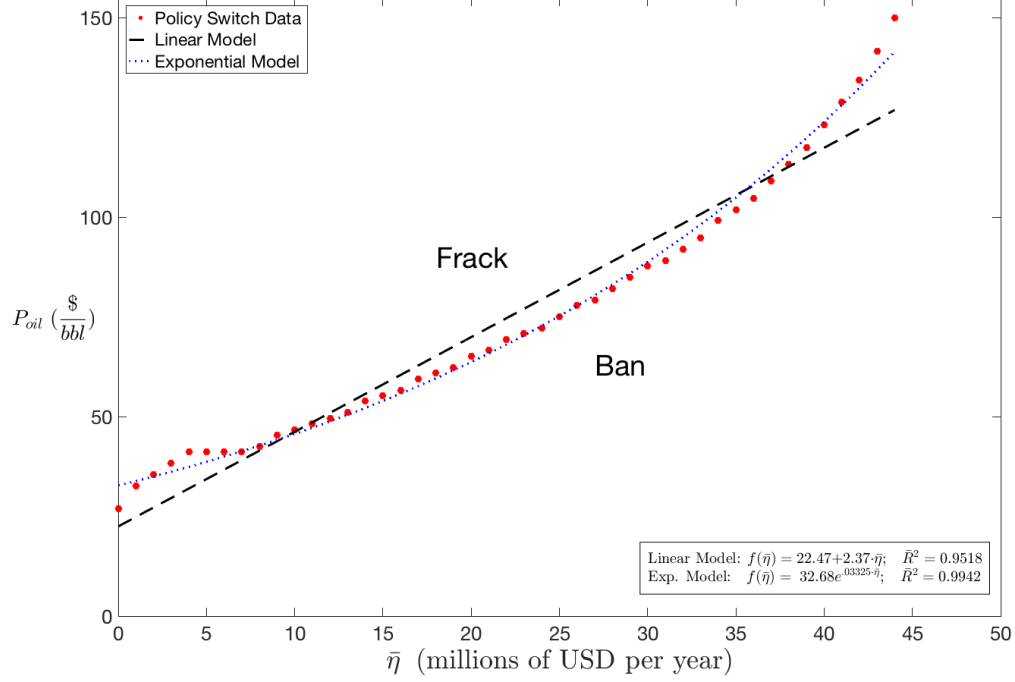


FIGURE 3.7: Policy Switch Curve

The x-axis is the cost of damages in millions of USD and the y-axis is the initial price of oil in $\$/bbl$. The data is fit with two models: linear and exponential.

to be $2.37 \$/bbl$ initially higher. This means that an initial increase in P of $1 \$/bbl$ is worth \$421,940 more in costs. In comparison to the solution without the shut-down condition, this is \$74,740 higher, or 21.5%. The linear model presents similar issues here as well, so we also fit an exponential model. The fitted coefficient, .03325 means a $\sim 3.3\%$ increase in the initial P is worth and extra \$ 1 million in damages or that 1% increase in the initial P is worth \$30.07 million, 35% higher than without the shut-down condition.

3.5. CONCLUSION

The prospect of a future ban, technological improvements that make oil and gas an obsolete energy source, or the discovery of more economically accessible reservoirs elsewhere may drive the owners of mineral rights to view the development decision with a *now or never*

framework. Our calculations show that, under Hahn et al.'s (2014) specification, the optimal policy using the NPV decision rule is to allow development when $P \in (10, 150) \text{ \$/bbl}$. However, we have also shown that the QOV rule indicates that fracking should be delayed if $P \leq 50.77 \text{ \$/bbl}$. Therefore, mineral rights owners using a NPV rule would not internalize the positive externality of optimally timing development as captured by the QOV, valued at \$83 million at its maximum. This constitutes a market failure and justifies government intervention in the form of a temporary ban on fracking.

Using feasible calibrations of fracking costs and realistic price moments for oil and gas, the results of this chapter seem to support the Chapter 2 assertion that the value of local reserves is likely more influential in decision making than the ability to learn about environmental damages. The QOV related to *stochastic uncertainty* over oil prices is larger than the QOV related to *epistemological uncertainty* over prices and, intuitively, more so when the policymaker accounts for the developers willingness and ability to cap an active well ceasing the royalty payments that drive the local economic boom.

We should also note that a risk-neutral policymaker concerned with *epistemological uncertainty* over environmental damages would allow fracking at $50 \text{ \$/bbl}$, whatever the rate of learning for the preferred specification. That is, the results of Figure 2.4 were generated assuming $P = 40 \text{ \$/bbl}$. Using $P = 50 \text{ \$/bbl}$ would find that the NPV exceeds V^1 for any precision of future information ($V_0(\cdot|1) > V_0^l(\cdot|0) \forall \sigma_\epsilon^2 > 0$)¹⁹. Here, though, the policymaker concerned with *stochastic uncertainty* over local economic benefits would ban fracking under the preferred specification at this price only allowing it if $P > 50.77 \text{ \$/bbl}$. Moreover, if this policymaker anticipates that local firms will cap fracked wells if P falls below $40 \text{ \$/bbl}$ then

¹⁹That the policy outcome is not dependent on the rate of learning is why those results were not presented.

she will only allow fracking if $P > 67.87 \$/bbl$. This conclusion is so surprising that it bears reiterating concisely: a policymaker concerned with environmental uncertainty is more likely to allow fracking than a policymaker concerned with uncertainty over local benefits.

This work also sets the stage for further inquiry. Future research should incorporate both types of option values into a single framework. It is possible to assume a functional form for benefits derived from the 'multiplier' effect as shown in Table 1.2. Further assuming GBM for P and Bayesian updating for a normally distributed η , a closed-form solution may be feasible. This model would allow for a direct, apples-to-apples, comparison of the epistemological and stochastic QOVs in order to conclude which is more influential.

CHAPTER 3 REFERENCES

- K. J. Arrow and A. C. Fisher. Environmental preservation, uncertainty, and irreversibility. *The Quarterly Journal of Economics*, pages 312–319, 1974.
- H. Askari and N. Krichene. Oil price dynamics (2002–2006). *Energy Economics*, 30(5): 2134–2153, 2008.
- F. Black and M. Scholes. Taxes and the pricing of options. *The Journal of Finance*, 31(2): 319–332, 1976.
- M. J. Brennan and E. S. Schwartz. Evaluating natural resource investments. *Journal of business*, pages 135–157, 1985.
- J. P. Brown, T. Fitzgerald, and J. G. Weber. Capturing rents from natural resource abundance: Private royalties from us onshore oil & gas production. *Resource and Energy Economics*, 46:23–38, 2016.
- J. P. Brown, T. Fitzgerald, and J. G. Weber. Asset ownership, windfalls, and income: evidence from oil and gas royalties. 2017.
- R. Brown. *A Brief Account of Microscopical Observations Made... on the Particles Contained in the Pollen of Plants, and on the General Existence of Active Molecules in Organic and Inorganic Bodies*. 1828.
- M. A. G. Dias and K. M. C. Rocha. Petroleum concessions with extendible options using mean reversion with jumps to model oil prices. In *3rd Real Options Conference*, pages 1–27, 1999.
- A. K. Dixit and R. S. Pindyck. *Investment under uncertainty*. Princeton UP, Princeton, 1994.

- A. Einstein. Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen. *Annalen der physik*, 322(8):549–560, 1905.
- Geometric Brownian Motion. Geometric brownian motion — Wikipedia, the free encyclopedia, 2017. URL https://en.wikipedia.org/wiki/Geometric_Brownian_motion. [Online; accessed 29-April-2016].
- R. Gibson and E. S. Schwartz. Stochastic convenience yield and the pricing of oil contingent claims. *The Journal of Finance*, 45(3):959–976, 1990.
- W. J. Hahn, J. A. DiLellio, and J. S. Dyer. What do market-calibrated stochastic processes indicate about the long-term price of crude oil? *Energy Economics*, 44:212–221, 2014.
- W. M. Hanemann. Information and the concept of option value. *Journal of Environmental Economics and Management*, 16(1):23–37, 1989.
- C. Henry. Option values in the economics of irreplaceable assets. *The Review of Economic Studies*, pages 89–104, 1974.
- J. E. Hilliard and J. Reis. Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions in the spot. *Journal of financial and quantitative analysis*, 33(1):61–86, 1998.
- J. C. Hull. *Options, futures, and other derivatives*. Pearson Education India, 2006.
- B. Kaffel and F. Abid. A methodology for the choice of the best fitting continuous-time stochastic models of crude oil price. *The Quarterly Review of Economics and Finance*, 49(3):971–1000, 2009.
- D. G. Laughton and H. D. Jacoby. The effects of reversion on commodity projects of different length. *Real options in capital investments: Models, strategies, and applications*. Ed. by L. Trigeorgis. Westport: Praeger Publisher, pages 185–205, 1995.

- R. L. McDonald and D. R. Siegel. Investment and the valuation of firms when there is an option to shut down. *International economic review*, pages 331–349, 1985.
- N. Meade. Oil prices—brownian motion or mean reversion? a study using a one year ahead density forecast criterion. *Energy Economics*, 32(6):1485–1498, 2010.
- M. Miller and L. Zhang. Oil price hikes and development triggers in peace and war. *The Economic Journal*, pages 445–457, 1996.
- B. Oksendal. Stochastic differential equations, an introduction with applications, 2000.
- J. L. Paddock, D. R. Siegel, and J. L. Smith. Option valuation of claims on real assets: The case of offshore petroleum leases. *The Quarterly Journal of Economics*, 103(3):479–508, 1988.
- D. Pilipovic. *Energy risk: Valuing and managing energy derivatives*. McGraw Hill Professional, 2007.
- F. A. Postali and P. Picchetti. Geometric brownian motion and structural breaks in oil prices: a quantitative analysis. *Energy Economics*, 28(4):506–522, 2006.
- E. S. Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3):923–973, 1997.
- G. Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters*, 20(2):177–181, 1986.
- C. P. Traeger. On option values in environmental and resource economics. *Resource and Energy Economics*, 2014.
- G. E. Uhlenbeck and L. S. Ornstein. On the theory of the brownian motion. *Physical review*, 36(5):823, 1930.