

A Run on Oil: Climate Policy, Stranded Assets, and Asset Prices

Michael Barnett *

University of Chicago

[Please click here for latest version](#)

June 27, 2019

Abstract

I study the dynamic implications of climate policy risk and uncertainty on macroeconomic outcomes and asset prices. I find that incorporating climate policy that restricts oil use and has an unknown arrival time in an otherwise standard climate-economic model with oil extraction generates a run on oil; meaning oil firms dynamically accelerate extraction as climate change increases and oil reserves decrease due to the risk of future climate policy actions stranding oil reserves. Furthermore, the risk and uncertainty of the climate policy, and the run on oil it causes, leads to a downward shift and dynamic decrease in the oil spot price and value of oil firms compared to the setting without climate policy risk. Empirical evidence based on cross-sectional and time series analysis suggests that the effects of climate policy risk on observable market outcomes such as oil production, stock returns, and oil prices are consistent with the model predictions about the effects of climate policy risk and uncertainty.

Keywords: Climate Policy, Oil Prices, Production-Based Asset Pricing, Stranded Assets

*Michael Barnett. Email: mdbarnett@uchicago.edu. Address: University of Chicago, 1126 E. 59th Street - Saieh Hall, Chicago, IL 60637. I am very grateful for the guidance and support of my advisors Lars Peter Hansen and Pietro Veronesi, and my committee members Michael Greenstone and Bryan Kelly. I want to thank Ufuk Akcigit, Buz Brock, Stefano Giglio, Amir Jina, Ryan Kellogg, Paymon Khorrami, Ralph Koijen, Moritz Lenel, Stefan Nagel, Alan Sanstad, Rob Townsend, Willem Van Vliet, Amir Yaron, and participants from the Capital Theory working, the Economic Dynamics Working Group, the Booth Finance Brownbag, and the Booth Finance Seminar for their comments and suggestions. I'd also like to think Pietro Veronesi and Yoshio Nozawa for providing code for the CME data. I am grateful for the financial support of the National Science Foundation, the Fama-Miller Center, the Energy Policy Institute at the University of Chicago (EPIC), the Stevanovich Center for Financial Mathematics, and the University of Chicago.

1 Introduction

Climate change has become one of the most significant and complex challenges currently facing our planet. Because climate change has the potential to significantly impact the well-being of households and the production decisions of firms, the possible macroeconomic and financial implications in the near and distant future could be substantial. However, the potential physical consequences of climate change are only part of the story. Governments and policymakers around the world are making various decisions with respect to climate change, and these policy actions carry risk and uncertainty as to how and when they will play out and the effect they will have. Thus, the risk and uncertainty of climate policy, together with the possible physical risks of climate change, should be a central component to any analysis of the consequences of climate change on the global economy and financial markets.

With these concerns in mind, my paper examines the dynamic implications of climate policy risk and uncertainty for economic and financial outcomes, focusing particularly on the oil sector. To study this issue, I use a general equilibrium, production-based asset pricing model that incorporates climate change as well as uncertain climate policy, or climate policy that carries significant risk and uncertainty because it restricts the use of oil and has an unknown arrival time. In the model, oil firms make decisions about oil extraction and exploration based on the demand for oil, remaining oil reserves, and climate change impacts. Uncertain climate policy, the key and novel component of my analysis, is modeled as a stochastic jump in the energy input share of oil for final output production. The arrival rate of this policy shock increases as climate change increases. This type of policy is meant to capture in a reduced-form manner the types of policies currently being proposed, as well as those used historically, that impose mandates or restrictions on the use of oil and fossil fuels, incentivize green-technology innovation, and carry significant uncertainty as to how and when exactly the policy will be implemented. The risk and uncertainty associated with climate policy generates what I call a run on oil, that is even as temperature rises and oil reserves diminish oil firms increasingly run up production of oil to avoid having reserves become stranded. This risk of stranded assets and the run on oil lead to dynamic reductions in the price of oil and the value of oil firms. Ignoring the impact of climate policy risk would substantially inflate oil prices and oil firms values because they do not incorporate the likelihood of oil reserves becoming stranded and the run on oil production that the stranded assets risk induces.

To provide further intuition about the model mechanism, I investigate a number of key counterfactual scenarios and model extensions. In the first counterfactual, where there is no climate component and therefore no risk of a climate policy shock, oil production gradually decreases as the remaining level of oil reserves decreases, leading to a gradual decline oil in firm prices and a gradual increase in the spot price of oil. These outcomes are consistent with a more standard Hotelling-model result. In the second counterfactual, where there are climate impacts from oil emissions and a constant arrival rate that is independent of climate change for the climate policy shock, the risk of climate policy leads only to a shift up in oil production and shift down in oil firm values, compared to the no uncertain

climate policy setting. Again, the dynamic outcomes in this constant policy risk are consistent with the more standard Hotelling model-type dynamics.

The model extensions I explore focus on the impact of two alternative scenarios. The first scenario assumes there is no oil exploration, providing insight about the significance of exploration for the run on oil. The second scenario assumes an alternative climate policy shock that only affects the oil sector input demand share and not the green sector input demand. This oil-sector only alternative policy demonstrates the role that the increasing benefit of green energy as an input has on the model mechanism demonstrated in the baseline model setting. In these cases I show that the same intuition and forces exist, though the levels of the quantity and price impacts may differ due to changes in the social costs of the climate policy shock because of the non-renewability of oil reserves or the lack of a green sector innovation with the policy shock.

To better understand the welfare implications of the realization of a climate policy shock, I provide a policy regime welfare comparison exercise to determine when the climate policy shift improves social welfare. I do this for the baseline policy setting, as well as the alternative model specifications. Such a welfare comparison demonstrates when a social planner would accept a policy regime change, if they were allowed to make such a choice in the model. These welfare comparisons provide intuition for the types of policies that are more likely to be acceptable to voters and elected officials representing voters interests and therefore implementable in practice, for what states of the world these policies are more acceptable, as well as why countries may respond differently to uncertain climate policy.

I also provide empirical evidence which suggests observable outcomes are consistent with the model mechanism. I first do this using an event-study analysis of climate policy events that shift the likelihood of future climate policy actions taking place. For events that imply a downward shift in the likelihood of future climate policy occurring, such as the 2016 US presidential election or the US Supreme Court decision to put a stay on the Clean Power Plan, the model would predict these events should increase the value of firms with high climate policy risk exposure, such as oil firms, and also increase the price of oil. The opposite should hold for events that increase the likelihood of future climate policy actions, such as the announcement of the Clean Power Plan or the UN's Paris Climate Accord. I estimate the effect of shifts in the likelihood of future climate policy due to climate policy events by regressing sectors' cumulative abnormal returns after the event on their exposure to climate policy risk, proxied for by exposure to oil price shocks as motivated by the model prediction. I find sectors with the highest climate policy risk exposure experienced the largest increases in cumulative abnormal returns for events that decreased the likelihood of future climate policy action and the largest decreases in cumulative abnormal returns for events that increased the likelihood of future climate policy actions, consistent with the model predictions.

Finally, I construct a "climate policy" event index from realized climate policy, energy sector, and climate-related events to estimate the dynamic impact of changes in climate policy shocks. In estimated reduced-form regressions, I find that increases in the likelihood of major climate policy

measured by my index lead to increased global and regional oil production. I also find that positive climate policy shocks lead to increasingly negative returns for the US oil sector and the spot price of oil. Finally, I estimate a structural VAR for the global oil market that includes the climate policy index, and calculate impulse response functions for a shock to climate policy. The results suggest that increases in the likelihood of significant climate change policy leads to long-term and permanent increases in crude oil production and a statistically significant decreases in the oil spot price, consistent with the dynamic predictions of my model. For each index-based empirical test, the statistical and economic significance are greater during the more recent, policy-focused time period (1996-2017) than for the entire available time sample (1973-2017), further validating the temperature dependence of outcomes implied by the model and the dynamic effect of uncertain climate policy the model predicts.

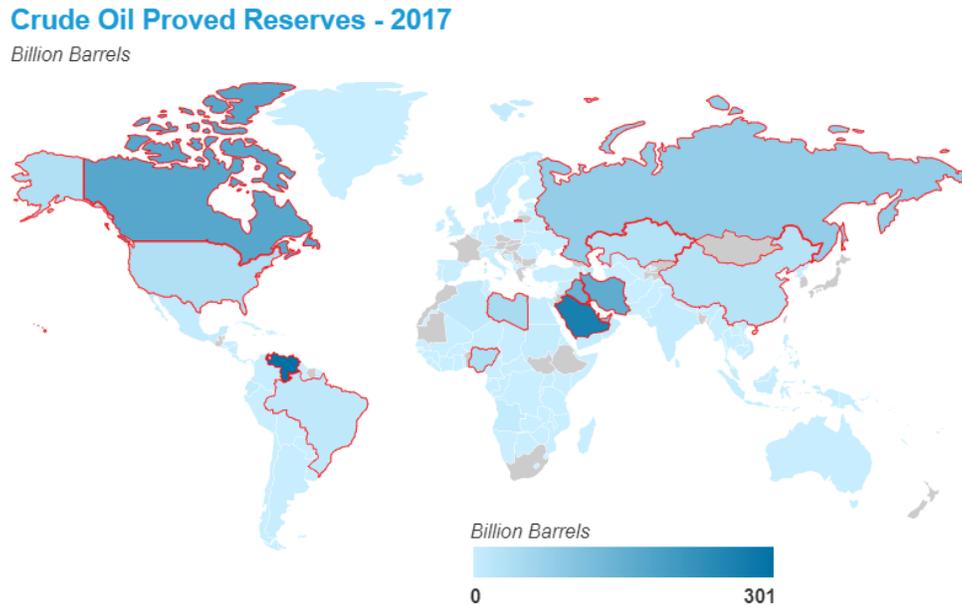
2 Example of Uncertain Climate Policy and Responses

To highlight the importance and potential impact of uncertain climate policy, consider the recent global climate policy agreement established in 2015 known as the Paris Climate Accord. This agreement came about as a result of the UN Framework Convention on Climate Change in order to limit change in the global mean temperature (GMT) from the pre-industrial level to no more than 2° C. Though many have seen this “temperature ceiling” agreement as a significant step toward limiting climate change, substantial uncertainty exists about when and if the necessary policy actions to achieve this target will occur. This uncertainty is due in large part to the considerable difficulty in coordinating such global policy actions, highlighted in this case by the fact that countries’ actions to limit climate change are self-determined and self-reported and that there is no centralized enforcement mechanisms in place to hold members accountable for failing to achieve proposed targets.

The policy responses to this agreement are particularly enlightening in regards to the potential impacts of uncertain climate policy. Consider first the policy responses of Norway, a climate-conscious country with sizeable oil reserves. First, Norway proposed moving forward their carbon neutrality goal from 2050 to 2030 as a response to the Paris Accord. Expediting this goal will require reductions to the emissions the country produces, which it plans to achieve through policies such as only selling electric vehicles by 2025, and purchasing carbon emission licenses to offset remaining emissions. However, at nearly the same time, Norway also proposed policy to increase oil drilling and development of their considerable reserves (New York Times, June 17, 2017, "Both Climate Leader and Oil Giant? A Norwegian Paradox"). Although these policies appear to be antithetical to each other, my model demonstrates that the risk of oil reserves becoming stranded from future climate policy could justify this policy response.

The US, who holds even more significant oil reserves and is now one of the largest producers of oil in the world due the fracking boom, has responded even more directly. The US was initially a significant force in helping establish the Paris Climate Accord, establishing in conjunction with the

Figure 1: Global Oil Resources

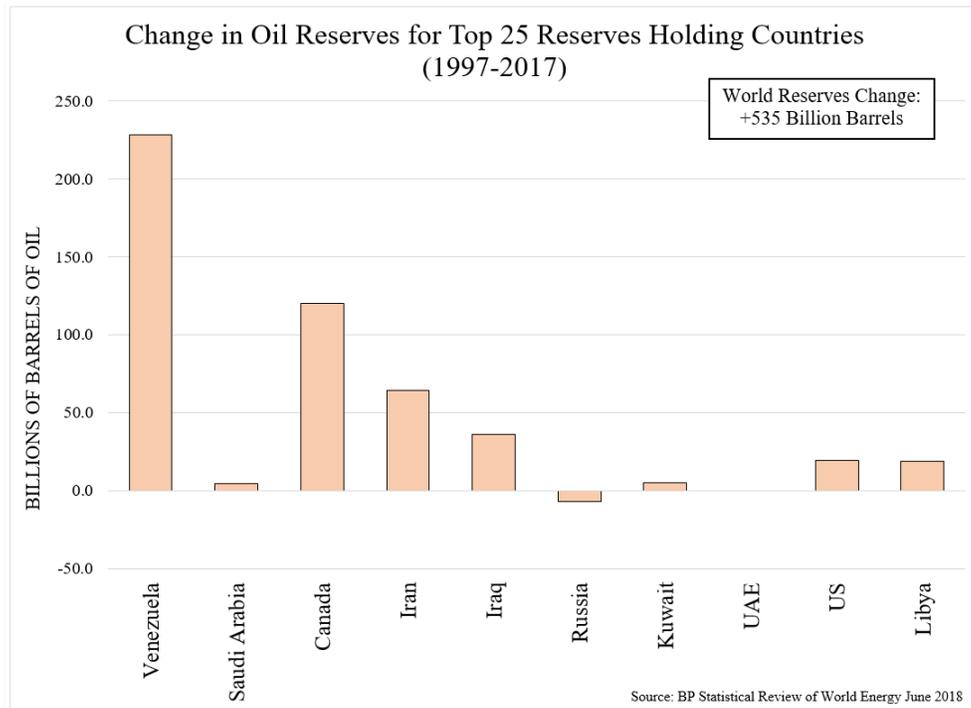


Source: US Energy Information Administration

Paris Agreement the Clean Power Plan, a policy that promoted reductions in greenhouse gas emissions and set renewable portfolio standards requiring increases in the fraction of energy and electricity produced from low-emission and renewable resources while phasing out high-emissions sources like coal and oil. However, after the Paris Agreement the US elected President Donald Trump, who has worked to fulfill campaign promises such as repealing the Clean Power Plan, pulling out of the Paris Agreement, rebuilding and supporting the coal sector, and repealing policies with strict greenhouse gas emissions regulations. Although other motivations exist for why the US elected Donald Trump as president, my model demonstrates why uncertain climate policy and the risk of oil becoming a stranded asset could have played an important role in this response.

Next, consider the recent proposal of Saudi Aramco, the state-owned oil producing firm of Saudi Arabia, to partially privatize through an initial public offering. The Saudi Kingdom lists a desire to diversify itself away from oil as the motivation for this decision (Financial Times, December 14, 2016, "The privatisation of Saudi Aramco"). While Saudi Aramco has proposed an over \$ 2 trillion valuation for the firm, many analysts and investors believe this assessment could be appreciably too high due to the potential risk of the countries oil reserves becoming a stranded asset (Financial Times, August 13, 2017, "Saudi Aramco's value at risk from climate change policies"). The fact that a country that has thrived on oil production from its massive oil reserves is considering an IPO, and the debate surrounding how to value this firm if it does go public, again indicate that even the largest oil producers may be taking climate policy risk seriously and that uncertain climate policy could be

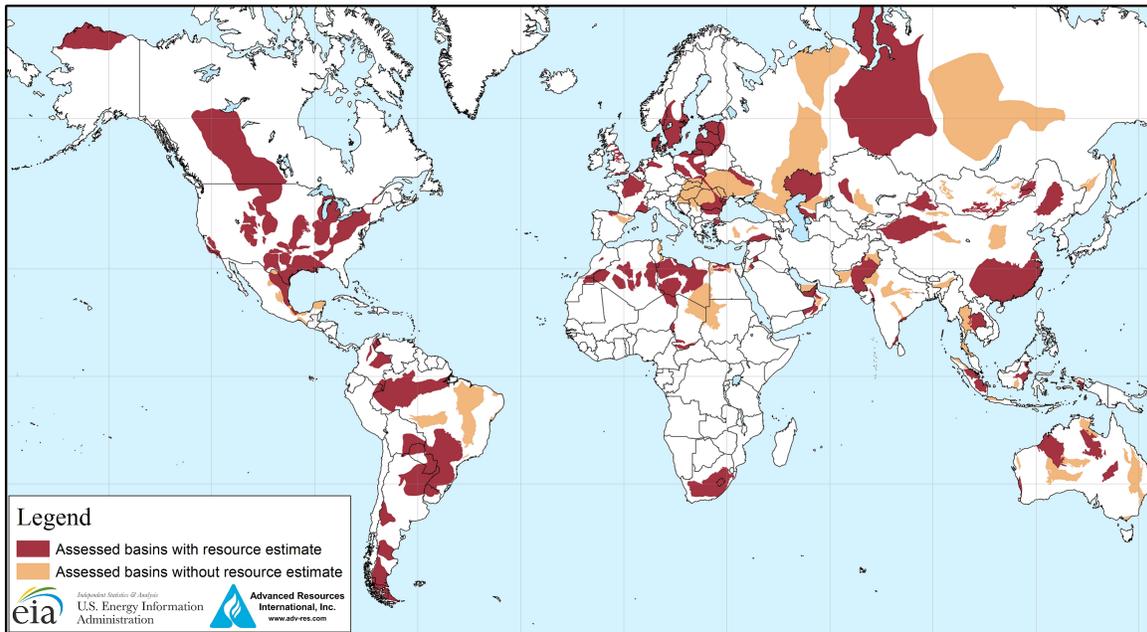
Figure 2: Change in Oil Reserves



playing a key role in firm decisions and how oil firms are valued in the future.

Each case above shows how the increasing likelihood of global climate change policy actions that could lead to stranded oil reserves may already be influencing decisions of oil producers and the value of oil firms. To further highlight the magnitude of stranded assets risk and why it may be so important in these situations, it is important to understand how much oil reserves there actually are. Figure 1 shows a heat map of global oil reserves held by each country. Not only does this map show that there are still over a trillion barrels of proven oil reserves that exist, it also shows that there are many countries holding these reserves. Furthermore, figure 2 shows changes in oil reserves over the last 20 years for the top oil reserve holding countries. As can be seen in this figure, many of the major reserve holders have had significant oil discoveries during this time, demonstrating that in addition to the significant recoverable oil reserves already known, there are still substantial discoveries of oil reserves being made. Finally, figure 3 shows a map of shale oil reserves by country across the globe. Because shale has only recently become economically feasible to drill, this figure shows again that there are significant new oil reserves across many countries which are just beginning to be tapped. Thus, future climate policy that would ban or severely restrict oil production would generate a significant stranded assets risk for countries all over the world. Understanding the optimal decisions for firms and countries facing such climate policy risks, as well as the impact they will have on oil prices and stock prices for energy firms, which are some of the largest firms in the world based on market capitalization, is of vital importance for understanding the full economic and financial impacts of climate change.

Figure 3: Global Shale Resources



3 Related Literature

This paper contributes to a number of important lines of literature in economics and finance. The first is research studying the interaction between economics and climate change. Nordhaus (2014), Golosov et al. (2014), Acemoglu et al. (2016), Stern (2007), Pindyck and Wang (2013), Hambel et al. (2015), and Cai et al. (2015) theoretically explore this link by examining the social cost of carbon and optimal carbon taxation, directed technological change, and other key climate-economic elements. Kelly and Kolstad (1999) and Crost and Traeger (2011) focus the impact of risk, or unknown outcomes with known probabilities, in climate economics models. Lemoine and Traeger (2012), Anderson et al. (2016), Brock and Hansen (2017), and Barnett et al. (2018) incorporate developments in decision theory in economics, such as those developed by Anderson et al. (2003), Klibanoff et al. (2005), Hansen et al. (2006), Maccheroni et al. (2006), Hansen and Sargent (2011), and Hansen and Miao (2018), to account for the impact of ambiguity and model misspecification about climate change and climate models. Deschenes and Greenstone (2007), Dell et al. (2012), and Hsiang et al. (2017) empirically estimate climate damages in different economic sectors and regions and the impact of climate change on economic growth. I add to this literature by examining the link between the risks from uncertain climate change policy and the oil sector, with an emphasis on the dynamic implications for quantities and asset prices.

Two particularly relevant areas in the climate-economics literature that my paper builds on are stranded assets risk and the Green Paradox. Stranded assets are assets that become worthless due to technological or policy changes. McGlade and Ekins (2015) study the potential magnitude of

stranded assets for fossil fuels based on a 2° C temperature ceiling using least-cost analysis. The Grantham Research Institute at the London School of Economics is also exploring the stranded assets issue and its implications for a potential “carbon bubble,” or possible overvaluation of oil firms from not accounting for stranded assets risk. The Green Paradox, a theory proposed by Sinn (2007) and recently extended by Kotlikoff et al. (2016), suggests the possibility that climate policy intended to mitigate climate change on the demand side may cause firms to alter the timing of their fossil fuel production in a possibly harmful way. By studying economic outcomes and asset prices in a setting where uncertainty exists about policy that has yet to be implemented, and allowing for the potential that fossil fuels become stranded and unused in a fully dynamic, stochastically uncertain environment, my paper synthesizes and extends the scope of these two areas in significantly important ways.

This paper also connects to important areas in the asset pricing literature. First is the growing literature on production-based asset pricing. Seminal papers in this area of research include Brock (1982), Cochrane (1991), and Jermann (1998), as well as Gomes et al. (2003), Gomes et al. (2009), Papanikolaou (2011), and Kogan and Papanikolaou (2014). These address key questions such as the equity- and value-premium puzzles using models with both quantity and pricing implications. Another area is the work on the interaction between government and asset prices. Belo et al. (2013), Kelly et al. (2016), Pastor and Veronesi (2012), and Santa-Clara and Valkanov (2003) are key examples in this area that focus on the impact of election and policy uncertainty in the US on asset prices that face differential exposures to these risks. Furthermore, Sialm (2006) and Koijen et al. (2016) are two examples in this literature that are closely related to this paper. They explore the impacts of uncertain policy on asset prices in the form of taxes and healthcare. Pástor and Veronesi (2009) provides an important example of the asset pricing impacts that arise from a shift in the production function, in their case due to learning about and adopting a new technology, related to the type of policy event I explore here. This paper also contributes to the nascent literature exploring the link between climate change and asset prices. Examples here include Bansal et al. (2016), Dietz et al. (2017), Hong et al. (2016), and Barnett (2017). These papers explore the impact of climate change and long-run risk on the social cost of carbon and asset prices, the elasticity of climate damages, the reaction of stock prices in the food sector to climate change, and the cross-sectional and time series implications of climate change and climate model uncertainty on economic and asset pricing outcomes, respectively. The current paper adds to these areas by building climate change into a production-based asset pricing model that incorporates important characteristics of the oil sector in order to explore the impacts of climate policy outside the socially optimal tax framework, and the uncertainty associated with this type of policy, on the energy sector.

Finally, this paper contributes to the extensive literature on resource extraction and oil prices. Hotelling (1931) and Dasgupta and Heal (1974) are the seminal works on natural resource extraction. Important contributions have been made by Hamilton (2005), Hamilton (2008), Kilian (2008), Kilian and Park (2009), Kilian (2009), and Baumeister and Hamilton (2017) in exploring the link between

oil prices and economics and financial shocks. Carlson et al. (2007), Casassus et al. (2009), Kogan et al. (2009), David (2015), Ready (2015), and Bornstein et al. (2017) each focus on varying stylized facts of oil prices and identify important model mechanisms required to match those outcomes. Salant and Henderson (1978) explore the impact of uncertain, exogenous government policy on commodities such as auctions on gold prices. My model allows for various components from the models in this literature, while also incorporating the impact of uncertain climate policy, linked directly to the state of the climate, to study the impact on oil production and prices of oil and oil firms.

4 The Model

Having framed the question of interest and highlighted the contribution of this paper to the literature, I can now lay out the model I will use to study the implications of uncertain climate policy. The model consists of three components: households, production, and the climate and climate policy component. The following section outlines the details of each of these components.

4.1 Households

Households have recursive preferences of the Duffie-Epstein-Zin type, given by

$$h(C, V) = \rho(1 - \xi)V(\log C - \frac{1}{1 - \xi} \log((1 - \xi)V))$$

where ρ is the subjective discount rate, ξ is risk aversion, V is the value function or continuation value, and C_t is consumption. These preferences allow for the separation of risk aversion and the elasticity of intertemporal substitution (EIS), meaning the concern agents have for consumption over time is not inversely related to how they view risk across states of nature. Furthermore, the recursive nature of the preferences means that agents' concerns about the resolution of future uncertainty are incorporated into the decision-making process. Because of these features, these preferences have shown to be extremely useful in helping match asset pricing outcomes observed in the data. Thus, even while imposing the assumption of unit EIS that allows for tractability when solving the model, the assumed preferences allow for a more realistic analysis of the outcomes of interest. Given this preference structure, the household maximizes discounted lifetime utility subject to their budget constraint:

$$V = \max_{C_t} E\left[\int_0^\infty \rho(1 - \xi)V(\log C_t - \frac{1}{1 - \xi} \log((1 - \xi)V))dt\right]$$

subject to

$$C_t \leq \Pi_t + w_t + \mathcal{T}_t$$

where Π_t is profits from the firms which the households own, w_t is wages from labor, and \mathcal{T}_t are any taxes that are rebated to the households. The household inelastically supplies a unit of labor.

4.2 Production

4.2.1 Final Output

The final output firm produces the consumption good using a Cobb-Douglas technology with capital, labor, and energy as the inputs:

$$Y_t = A_C K_t^\gamma L_{C,t}^\alpha E_t^{1-\gamma-\alpha}$$

where A_C is total factor productivity (TFP), $L_{C,t}$ is final output labor, K_t is capital, E_t is energy, and γ and α are the factor input shares of capital and labor. Energy is a Cobb-Douglas aggregate of oil and green energy:

$$E_t = O_t^{\nu_t} G_t^{1-\nu_t}$$

where O_t is oil, G_t is green energy, and ν_t is the energy input share of oil. The state of climate policy governs the value of ν_t and is determined by a Poisson jump process. For simplicity there are only two possible values of ν_t , and so a climate policy shock permanently shifts ν_t from ν in the pre-policy state to 0 in the post-policy state. When the value of ν_t is high, it represents loose climate policy and a high demand for oil to be used in production of the final good, and when the value of ν_t goes to 0, it represents strict climate policy where final output production can only be done with green energy. The jump process is characterized by a climate-dependent arrival rate. Modeling policy in this way is a reduced-form representation of policy mandates that limit fossil fuel use and incentivize green-technology innovation that also captures the uncertainty that accompanies implementation of global climate policies. The full characterization and motivation for the climate policy set-up I use in the model are in section 4.3.

The final output sector is perfectly competitive, and so firms in this sector maximize discounted, expected lifetime profits by optimally choosing investment, labor, and energy inputs subject to state variable evolution, market clearing, and taking prices as given:

$$\begin{aligned} V_C &= \max_{O,G,L,I} E \int \pi_t (\tilde{Y}_t - w_t L_C - P_{I,t} I_t - P_{O,t} O_t - P_{G,t} G_t) ds \\ &\text{subject to} \\ dK_t &= K_t (\ln B + \delta_1 \ln I_t - \delta_2 \ln K_t) dt + \sigma_K K_t dB_K \\ w_t, P_{O,t}, P_{G,t} &: \text{wages, oil price, green price, taken as given} \end{aligned}$$

Note that \tilde{Y}_t is used in the firm problem above, which is final output after accounting for climate change damages. I explain this important climate impact in detail in section 4.3. π_t is the stochastic discount factor (SDF) that provides the necessary discounting across time and states of nature in order to derive firm values, which I derive and elaborate on in section 6.3. The evolution of the

capital stock is subject to a specific case of the adjustment costs used by Jermann (1998) and others, highlighted in recent work by Anderson and Brock (2017). This case of adjustment costs is empirically indistinguishable from other common forms used in the literature for observable outcomes in the data and allows for tractability when solving the model. The unit supply of labor is divided between use in the final good production and use in the green input production.

4.2.2 Oil Input

The oil firm produces using the linear technology

$$O_t = N_t - i_{R,t}R_t$$

where O_t is the oil used for final output production, N_t is oil extracted, R_t is oil reserves, and $i_{R,t}R_t = I_{R,t}$ is investment in oil reserves exploration. Oil firms maximize discounted expected lifetime profits by choosing extraction and exploration subject to evolution of state variables and market clearing:

$$V_O = \max_{N, I_R} E \int \pi_t(P_{O,t}O_t)ds$$

subject to

$$dR_t = (-N_t + \Gamma R_t i_{R,t}^\theta)dt + \sigma_R R_t dB_R$$

$$dT_t = \varphi N_t dt + \sigma_T dB_T$$

Again π_t is the SDF used for discounting firm profits. T_t is atmospheric temperature, discussed in detail in section 4.3. The evolution of reserves is determined by investment in exploration, exploration adjustment costs, and the extraction of oil. Adjustment costs are such that there are diminishing returns to exploration, and there are no explicit costs of extraction. However, because oil firms take into account the shadow value of holding reserves, this implicit cost limits the amount of extraction done at any given time. I assume a competitive oil sector so that oil firms take prices as given.

There are two critical features of the role of oil in the model. First, oil has a negative environmental impact, as seen in the equation for dT and which I elaborate on in section 4.3. I explore solutions in which this climate impact externality is internalized in a socially optimal way. Second, the demand for the oil input for final output production is initially higher given its higher energy input demand share than green energy. Thus, a trade-off between climate impacts and productivity is considered when making optimal choices about the use of the oil energy input. The novel feature of my model is exploring how a shift in the energy input demand share due to uncertain climate policy influences financial and economic outcomes.

4.2.3 Green Input

The green firm produces using a decreasing returns to scale (DRS) technology:

$$G_t = A_G L_{G,t}^\omega$$

where A_G is the green sector TFP, L_G is green labor, and ω is the DRS parameter for the labor input. The unit supply of labor is divided between use in the green energy production and final output production, as mentioned previously. This setting is chosen to maintain simplicity as the main focus of this paper is on the oil sector. Adding a “green” capital stock or making the the green sector TFP stochastic could be done as straightforward extensions, while having little to impact on the key results of the paper. I provide details of such extensions for a particular case of the model in the appendix. Green firms operate in a perfectly competitive sector, and so maximize discounted expected lifetime profits by optimally choosing labor subject to market clearing and taking prices as given:

$$V_G = \max_{L_G} E \int \pi_t (P_{G,t} G_t - w_t L_{G,t}) ds$$

subject to

$$w_t, P_{G,t} : \text{wages, green price, taken as given}$$

π_t is the SDF used for discounting profits as before. The green energy input has two key features that I briefly highlight here. First, the green input has no negative environmental impact, since green energy does not generate emissions. Second, the demand for green energy is initially lower than for oil because it has a lower energy input share before the realization of the climate policy shock.

4.3 Climate and Climate Policy

Atmospheric temperature in excess of pre-industrial levels, or simply temperature, evolves as

$$dT_t = \varphi N_t dt + \sigma_T dB_T$$

where φ is the carbon-climate response (CCR) to emissions from oil. This climate process is a stochastic version of the the relationship estimated and studied by Matthews et al. (2009), Matthews et al. (2012), and MacDougall and Friedlingstein (2015), who show that an affine relationship connecting carbon emissions to changes in atmospheric temperature closely approximates complex climate dynamics. Though the “Matthews approximation” model is typically seen as being best suited for longer-term time scales, I use it in place of more complex climate dynamics for a few reasons. First, it allows for greater tractability. Second, it accounts for the essentially permanent impact of emissions on the atmosphere, an important climate model feature given that the estimated rate of decay for atmospheric carbon is on the order of hundreds or possibly even thousands of years. Third, the

longer-run nature of the approximation fits the long-term climate change impacts I am interested in, as opposed to short-term weather fluctuations.

An important element of the climate-economic model is the damage function. The relationship assumed in my model, which is common across many climate-economic models, is that the damage function $D(T_t)$ multiplicatively scales final good output. Furthermore, the damage function has the properties $D(T_t) \in [0, 1] \forall T_t$, $D(0) = 1$, $D(\infty) = 0$, and $\frac{dD}{dT} < 0$. The functional form for the damage function and final output available for consumption are given by

$$D(T_t) = \exp(-\eta T_t) \quad \text{and} \quad \tilde{Y}_t = D(T_t)Y_t$$

The central and novel feature of the analysis in this paper is the uncertain climate policy. Change in policy is modeled by a permanent jump in the energy input share of oil, ν_t , which occurs according to a Poisson jump process. This defining component of the model produces the key, driving mechanism for the results in my model. By explicitly modeling the climate policy shock as I do here, I am able to study the role of uncertain policy and the risk of stranded assets on oil production, exploration, oil prices, and firm prices in a way not done previously.

The arrival rate of the shock to ν_t , or climate policy shock, is given by

$$\lambda(T_t) = \psi(1 - \exp(-\varpi T_t^p))$$

A critical element for the arrival rate is that it is dependent on the endogenously evolving level of climate change due to emissions generated by oil use. The interpretation for this climate-linked arrival rate is that the probability of significant climate policy being enacted increases as climate change becomes more pronounced. Also, the functional form is very similar to the damage functions used in climate-economics models. Thus, the choice of this functional form captures the fact that the realization of climate policy is likely strongly linked with observed climate damages.

4.3.1 Interpretation and Motivating Policy Examples

As uncertain climate policy is central to this paper, I elaborate on the interpretation and motivation for the policy structure assumed. Apart from the example of the Paris Climate Accord previously highlighted, numerous examples of climate policies in the US can be used to motivate the type of policy set-up that should be used. The Energy Policy Conservation Act (EPCA) was one of the first fuel economy goals passed in the US. The policy led to the development of catalytic converters and unleaded gas in order to reach the required vehicle emissions levels specified by the policy. The Clean Air Act (CAA), another early policy act that has been amended and updated in more recent times, gives air pollution and vehicle emissions standards while providing technical and financial assistance to state and local governments in order to enforce and achieve these standards. The Diesel Emissions Reduction Act (DERA) set increased diesel engine emissions standards with regards to greenhouse

gases, leading to innovations in diesel engine technology spearheaded by Cummins. The Energy Independence and Security Act (EISA) and Corporate Average Fuel Economy (CAFE) standards have helped lead to the development of hybrid and electric vehicles such as the Toyota Prius, Nissan Leaf, and Tesla vehicles. Such policy examples are particularly relevant because they typically set target goals for future deadlines, which leads to some uncertainty about such policies being achieved, and focus on emissions from crude oil use in motor vehicles, which makes up over 70% of crude oil consumption in the US (“Use of Oil,” EIA Independent Statistics and Analysis, September 19, 2017).

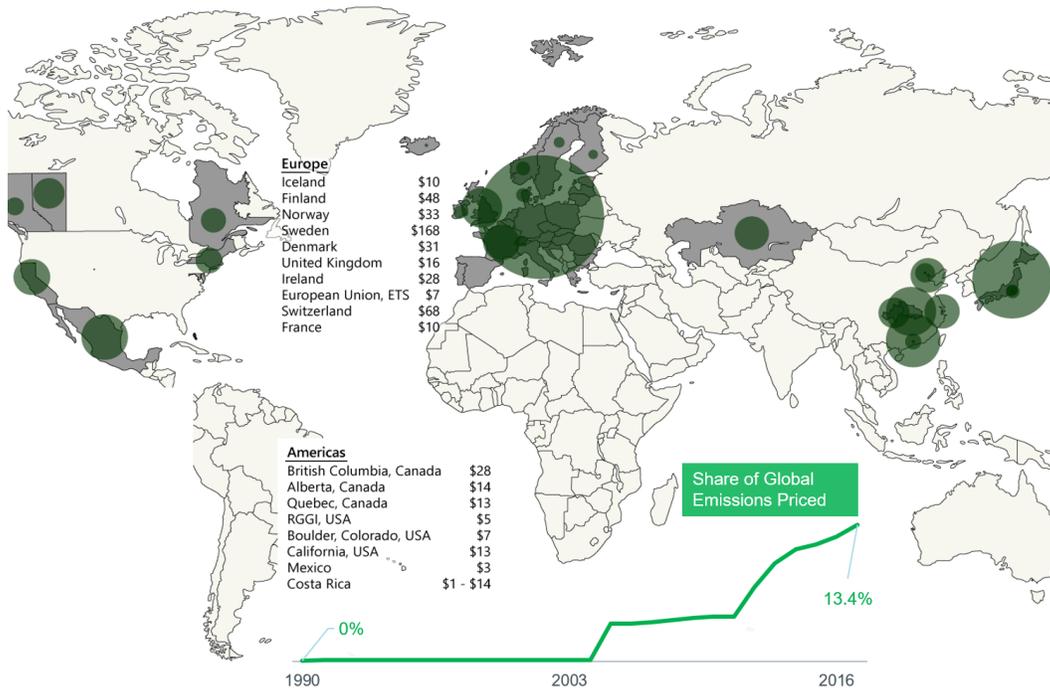
Beyond vehicle emissions, Renewable Portfolio Standards (RPS) are another example of the type of policy being used with regards to climate change. RPS policies, such as the Clean Power Plan (CPP) established by President Obama in conjunction with the Paris Climate Accord, which I previously mentioned, requires that a certain fraction of electricity be produced from renewable sources to increase green production/productivity by a proposed future deadline. Though RPS policies are meant to be mostly market-based, they also include multipliers to help direct revenue, investments, and jobs towards renewable sectors to help drive the necessary innovations in the green sector to make the target goals feasible.

Further evidence can be found in the annual 10-K filings for US oil producers. Each year US firms are required to include Section 1.A - “Risk Factors” in their 10-K’s filed with the SEC, where they are requested to list the “most significant factors” that affect the future profitability of the firm. Further details on the “Risk Factors” section of firms’ 10-K filings can be found in Koijen et al. (2016). Examining these filings for the 10 largest oil firms in terms of reserves held (Anadarko, Chevron, ConocoPhillips, EOG Energy, ExxonMobil, Halliburton, Marathon, Occidental, Phillips 66, and Valero) makes clear that climate policy risks are becoming increasingly more relevant for oil producers. Between 2004 and 2010, each of these firms began including sections about climate change policy and regulation. These sections include key words and phrases such as climate change, climate change policy, climate regulation, carbon-constrained economy, mandate, greenhouse gases, carbon emissions, increasing competition, reduced demand, alternative energy/fuels, renewable energy/fuels, and Paris Agreement. The key risks associated with climate policy that firms list include uncertainty about its impact, timing, and form, as well as potential mandates and shifts in demand away from oil and towards alternative clean energy sources. Consider the following excerpts from the 2018 10-K filings of ExxonMobil and Chevron, respectively:

“...the ultimate impact of GHG emissions-related agreements, legislation and measures on the company's financial performance is highly uncertain... because the company is unable to predict with certainty... the outcome of political decision-making processes...”

...even with respect to existing regulatory compliance obligations... it [is] difficult to predict with certainty the ultimate impact...” – ExxonMobil

Figure 4: Carbon Prices by Country and Global Average Carbon Price



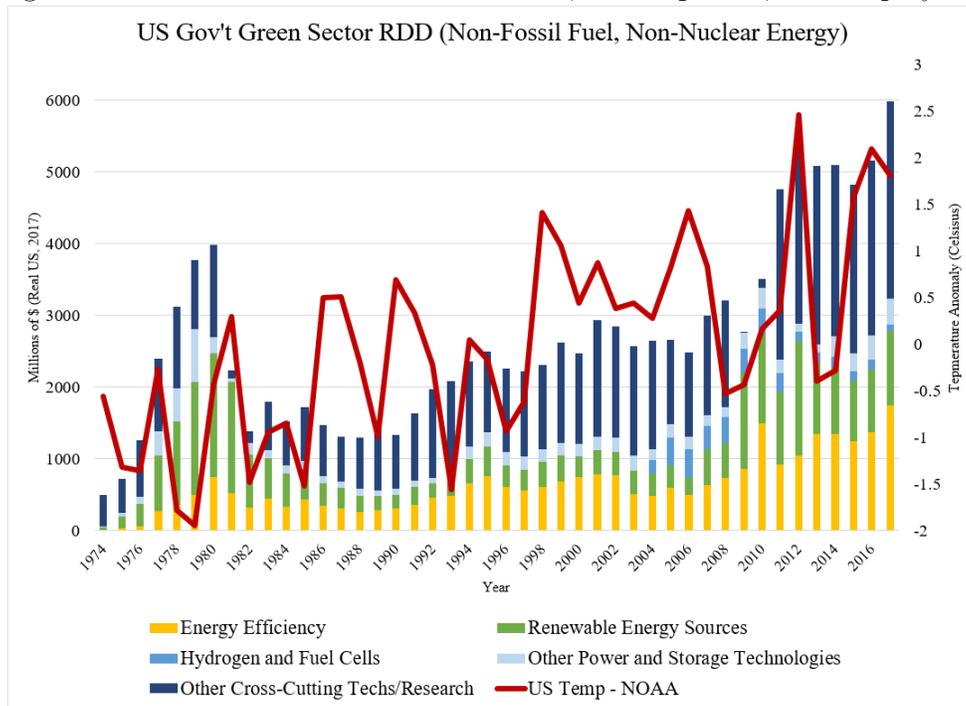
Source: Energy Policy Institute at the University of Chicago (EPIC)

“These requirements could... reduce demand for hydrocarbons, as well as shift hydrocarbon demand toward relatively lower-carbon sources...

...governments are providing tax advantages and other subsidies to support alternative energy sources or are mandating the use of specific fuels or technologies...” – Chevron

Finally, the assumption that the likelihood of climate policy is related to increasing climate change is another important assumption I make. Figures 4 and 5 provide empirical support for this relationship. Figure 4 shows a map of carbon prices for various countries, as well as the global average trend in price. The increasing likelihood of significant climate policy as climate change increases in my model is in line with the increasing positive trend in the the average global carbon price shown, which positively relates to observed temperature increases. We also see from this map that carbon prices exist in relatively few places in the world. Figure 5 shows the time series of US Government Research, Development, and Deployment for different green sectors and technologies. The time series for RDD has a correlation of 0.40 with US temperature anomaly (the red line) and 0.72 with global temperature anomaly (not show), consistent with the assumption that there is an increasing likelihood of significant climate policy and shift to greener production as climate change increases.

Figure 5: US Government Green Research, Development, and Deployment



Source: International Energy Agency and NOAA

These policy examples and their relation to the major uses of oil, the connection of policy to increases in temperature, and the types of policy concerns firms are listing as major risk factors demonstrate that concerns about climate policies that include a restriction on the use of oil as a production input, a shift in the energy demand share of green energy in the final output production function, and a positive correlation with temperature are consistent with historical and current policy actions and are particularly relevant to consider. Furthermore, by exploring the extension of an oil-sector only policy scenario, I can extend the analysis to explore the case where policy is in line with simply shutting off the oil sector input without the accompanying technological change in the green sector. This provides further insight about the policies mentioned above in that there are likely to be cases where governments seek to impose oil production mandates even if there is no significant green innovation. I will be able to show that this alternative setting the same mechanisms and dynamic impacts as in the baseline policy setting still exist, though the levels of the outcomes may differ. Furthermore, I'll address the differences each policy outcome has in terms of welfare implications.

5 Equilibrium Solutions

With the set-up of the model components now given, I put forward the equilibrium concept used throughout the paper and derive the equilibrium solutions for each case of interest. The type of equi-

librium considered for each scenario is that of a recursive Markov equilibrium, where optimal decisions are dependent only on the current value of the state variables. Thus, the equilibrium definition is given by optimal choices of quantities $\{C_t, I_t, O_t, G_t, L_{C,t}, L_{G,t}, N_t, I_{R,t}\}$ and prices $\{P_{I,t}, P_{O,t}, P_{G,t}, w_t\}$, which are functions of the state variables T_t, R_t, K_t , such that

1. Households maximize lifetime utility
2. The household budget constraint holds
3. Firms maximize discounted, expected lifetime profits
4. Market clearing in the goods and labor markets holds

I begin by deriving equilibrium solutions for the special case in which there is no climate component and then when there is a climate component but the policy arrival rate is constant and climate independent as counterfactuals for comparison. I then extend the model to the uncertain policy case for a perfectly competitive oil sector and derive the equilibrium solution for this scenario. Finally, I derive solutions for alternative policy scenarios and model extensions. In each case, I solve for the socially optimal choices, up to the constraint that the planner does not directly control the climate policy shock. To derive the prices that support these socially optimal outcomes, I will consider the decentralized economy which requires a decentralizing tax to generate the prices that support these socially optimal quantity outcomes. I discuss the pricing results in section 6.

5.1 Counterfactual Comparisons

5.1.1 No Climate Interaction

To start, consider the setting where we ignore the interrelationship between economic and climate processes as well as any possible uncertain climate policy in the model. Given this simplification, we develop intuition for how the underlying model compares with standard resource extraction models. This no climate case is the first step in establishing the counterfactual for comparison with the uncertain climate policy setting. The first proposition, which characterizes this setting, is as follows:

Proposition 1. *For the no climate, no policy shock setting, the agent's value function is given by:*

$$V(K_t, R_t) = c_0 K_t^{c_1} R_t^{c_2}$$

The optimal choices of investment, labor, extraction, and exploration are given by

$$\begin{aligned} I_t &= C_1 Y_t & L_C &= \bar{L} \\ N_t &= D_1 R_t & I_{R,t} &= E_1 R_t \end{aligned}$$

The value function constants c_0, c_1, c_2 and the FOC constants C_1, D_1, E_1, \bar{L} are functions of the model parameters only (given in the appendix).

I note a few important results pertaining to the optimal choices for this solution. First, the choice of labor is constant and investment in capital is proportional to final good output. The respective constants for these choices, given in the appendix, are functions of input shares, capital evolution parameters, and the value function constant for capital respectively. This structure of constant labor choice and capital investment proportional to total output will continue to hold for each version of the model considered because of the assumed form of adjustment costs, final output production technology, and the utility function.

Second, exploration and extraction of oil are both proportional to the level of reserves. Note that the proportionality of exploration and extraction to reserves is consistent with the standard Hotelling model-type outcome common to many resource extraction problems (see Hotelling (1931)). As reserves decrease, the marginal value of reserves increases and leads to lower returns to production today compared to future production, and so extraction is limited today as a result. Thus, in this version of the model without uncertain climate policy, we see decreasing extraction as the level of oil reserves decreases. Similar intuition explains the trade-off of the marginal benefit and marginal cost of exploration. Introducing the climate component and uncertain climate policy will break this proportional extraction and exploration relationship and generate a non-monotonic relationship between the marginal value of reserves and the level of reserves. This change in the marginal value of oil reserves, connected to the level of climate change, is the key mechanism of the model that can trigger a run on oil production and a significant decrease in the value of oil firms.

5.1.2 Constant Policy Arrival Rate

Now I introduce the climate component of the model, but assume the arrival rate of the policy is constant, i.e., $\lambda(T_t) = \bar{\lambda}$. The model under these assumptions further completes the counterfactual for comparison with the uncertain climate policy setting by demonstrating how a non-climate-dependent policy arrival rate continues to generate the the Hotelling model-type results that underlie the basic model. As a result, this setting demonstrates the critical feedback effect associated with the climate dependent policy arrival rate that generates the dynamics implications of the model. The following proposition characterizing this setting is as follows:

Proposition 2. *With uncertain climate policy where the arrival rate is given by $\lambda(T_t) = \bar{\lambda}$ and $\nu_t = \nu$ before the policy shock (pre) and $\nu_t = 0$ after the policy shock (post), the value functions for the two*

policy regimes are given by:

$$V_{pre}(K_t, R_t, T_t) = K_t^{c_1} \exp(c_3 T_t) f(R_t) \quad V_{post}(K_t, T_t) = \tilde{c}_0 K_t^{c_1} \exp(c_3 T_t)$$

where investment and labor decisions are given by

$$\begin{aligned} I_{pre,t} &= C_1 \tilde{Y}_{pre,t} & L_{pre,C} &= \bar{L}_{pre} \\ I_{post,t} &= C_1 \tilde{Y}_{post,t} & L_{post,C} &= \bar{L}_{post} \end{aligned}$$

Exploration and extraction are given by

$$i_{R,t} = \left(\frac{f_R \Gamma \theta}{f_R - \varphi c_3 f} \right)^{1/(1-\theta)} \quad N_t = \frac{f \vartheta}{f_R - \varphi c_3 f} + i_{R,t} R_t$$

$f(R_t)$ is the solution to the simplified HJB equation characterizing the planner's problem (given in the appendix). The value function constants c_0, c_1, c_3 and the FOC constants $\vartheta, C_1, \bar{L}_{pre}, \bar{L}_{post}$ are functions of the model parameters only (also given in the appendix).

The main contributors determining oil extraction are the marginal value of atmospheric temperature ($c_3 f(R_t)$), the marginal value of oil reserves (f_R), and the model primitives (risk aversion, production, and investment parameters). Note that while allowing for climate change in the model introduces a temperature-related adjustment to the optimal choice of oil production and exploration, the temperature state variable itself cancels out of these expressions. Therefore, the climate impact in this setting is a shift down due to the scaled, reserves-related component of the value function, $c_3 f(R_t)$. As a result, the dynamics of the optimal production and exploration decisions are qualitatively similar to those found in the no climate, Hotelling-type case. As oil reserves diminish, the marginal value of reserves still increases. The temperature-related adjustment $c_3 f(R_t)$ grows in magnitude as the value function becomes increasingly negative with lower reserves, and thus the climate adjustment simply shifts down the level of production and exploration in order to put off climate change until the future when climate costs are more heavily discounted.

When $\bar{\lambda} = 0$, the temperature-related adjustment to the choice of oil extraction is the only impact of incorporating climate change in the model. When $\bar{\lambda} > 0$, there is an additional impact that acts like an increase to the subjective discount rate, similar to an overlapping generations models with a constant arrival rate of death. As a result, agents in the $\bar{\lambda} > 0$ setting value current consumption more, leading to a shift up in oil production. However, this constant discount rate adjustment does not alter the qualitative dynamics of the model, which are tied to an increasing marginal value of oil, and thus decreasing oil production as reserves diminish. As a result, the main intuition of the standard resource extraction model still prevails in this setting.

5.2 The Impact of Uncertain Climate Policy

With the solutions for the model settings with no climate component and with the climate component but a constant arrival rate in mind, I now derive the equilibrium solution to the model that incorporates the full climate specification of the model, which includes uncertain climate policy where the policy arrival rate $\lambda(T_t)$ is temperature-dependent and therefore changes with climate change. The solution for this setting is as follows:

Proposition 3. *With uncertain climate policy where $\nu_t = \nu$ before the policy shock and $\nu_t = 0$ after the policy shock, and where the arrival rate of policy is given by the temperature dependent function $\lambda(T_t)$, the value functions for the two policy regimes are given by:*

$$V_{pre}(K_t, R_t, T_t) = K_t^{c_1} v(R_t, T_t) \quad V_{post}(K_t, T_t) = \hat{c}_0 K_t^{c_1} \exp(c_3 T_t)$$

where investment and labor decisions are given by

$$\begin{aligned} I_{pre,t} &= C_1 \tilde{Y}_{pre,t} & L_{pre,C} &= \bar{L}_{pre} \\ I_{post,t} &= C_1 \tilde{Y}_{post,t} & L_{post,C} &= \bar{L}_{post} \end{aligned}$$

Exploration and extraction are given by

$$i_{R,t} = \left(\frac{v_R \Gamma \theta}{v_R - \varphi v_T} \right)^{1/(1-\theta)} \quad N_t = \frac{v \vartheta}{v_R - \varphi v_T} + i_{R,t} R_t$$

Note $v(R_t, T_t)$ is the solution to the simplified HJB equation characterizing the planner's problem (given in the appendix). The value function constants \hat{c}_0, c_1, c_3 and the FOC constants $\vartheta, C_1, \bar{L}_{pre}, \bar{L}_{post}$ are functions of the model parameters only (also given in the appendix).

The main contributors determining the oil extraction and exploration are the marginal value of atmospheric temperature (v_T), the marginal value of oil reserves (v_R), and the model primitives (risk aversion, production, and investment parameters). These are similar contributors to the previous model settings. A key difference with uncertain climate policy is the role temperature now plays in those contributions. Directly, temperature impacts climate damages and the likelihood of climate policy occurring, but now in a way that the temperature-related adjustment to extraction and the subjective discount rate adjustment depend on the state of climate change. Indirectly, temperature has greater influence on the marginal value of reserves because the value function is no longer separable.

First, consider the direct effects of uncertain climate policy. The influence on the subjective discount rate from uncertain climate policy is now state dependent. One way to see this discount rate adjustment is by looking at the HJB equation. Normally, in simple CRRA utility settings for example, the multiplier scaling the value function in the HJB equation is just the subjective discount factor, which is ρ in this case. However, if we were to gather the terms that scale the value function

in this uncertain climate policy setting, there is an additional term for the policy arrival rate, $\lambda(T_t)$. Therefore, as T_t increases, $\lambda(T_t)$ also increases, and so the uncertain climate policy increases the degree to which agents discount future outcomes. Intuitively, firms and households are more concerned about reserves being stranded as temperatures increase due to the increased likelihood of climate policy being implemented. Second, the marginal cost of climate change, which directly determines the choice of oil extraction changes as well. Since increasing temperatures lead to an increased discount rate and increased expectations of the policy arriving, agents worry less about additional climate change they may cause because the climate policy arrival will stop any additional emissions. Therefore, even as oil reserves are decreasing and climate change is increasing, the increasingly likely arrival of policy that prohibits the use of oil to stop additional anthropogenic climate impacts reduces the marginal cost of climate change, a key determinant in the optimal choice of oil extraction.

The indirect effect, due to the non-separable effect of temperature on the marginal value of reserves, also means decreasing oil reserves do not necessarily imply increasing marginal value of reserves. Lower reserves are now associated with a decreasing cost of oil reserves being stranded and a potentially increasing likelihood of climate policy occurring due to climate change. As a result, the marginal value of reserves may now significantly decrease as the likelihood of oil reserves becoming stranded increases due to increased climate change. This effect on the marginal value of reserves amplifies the potential for a run on oil production, as oil firms value current profits more and more and ramp up oil production as a result to avoid holding reserves that become worthless after policy is enacted.

Thus the key impact of the climate-linked policy risk, through both the direct and indirect effects, is that it creates a dynamic, climate-related feedback mechanism in the model. The risk of a policy shock that restricts the use of oil leads to an increased level of oil extraction. This is true whether λ is constant or temperature dependent. However, as the arrival rate is state dependent, increased oil extraction caused by the stranded asset risk leads to increased climate change. This further exacerbates the stranded assets risk, and thus provides motivation for oil firms to further increase their oil extraction. The link between climate change and climate policy that strands oil reserves generates a feedback loop that can lead to a dynamic increase in oil production, not simply a level shift up. This can occur even as oil reserves are decreasing and temperature is increasing.

The final effect to consider is the potential desire to avoid climate policy. This is particularly prevalent when oil reserves are high and temperature is low, and thus the arrival rate of policy is low. When oil reserves are high, the cost of stranding oil is high. When temperature is low, and so the likelihood of a policy arrival is low, there is an incentive to delay oil production in order to try and delay the arrival of climate policy. However, as temperature increases and the likelihood of policy occurring gets larger, a run becomes more likely. This is due to the fact that the incentive to run on oil to run down reserves and minimize the cost of stranding oil reserves now exceeds the benefit of trying to delay policy by delaying production. Thus, for low temperatures values and high levels of reserves the level of oil production may actually be pushed down, amplifying the magnitude of the

dynamic run up in oil production before the arrival of climate policy.

5.3 Alternative Model Settings

To enrich the theoretical analysis, I explore alternative settings which provide further insight into the impact of uncertain climate policy with respect to the financial and economic outcomes of the model. The alternative settings also help provide intuition for how this policy setting compares to other possible outcomes of proposed policies that governments and policy makers are seeking to implement to try to stave off climate change impacts in the future. The focus here centers in particular on the alternative of an oil-sector only policy shock and the case when there is no oil exploration. I leave the derivations, details, and numerical results of these cases for the appendix. I also provide additional characterizations without numerical results in the appendix as well.

5.3.1 Oil-Sector Only Policy Impact

The first extension explores the setting where the arrival of policy only impacts the energy input demand share of oil. This oil-sector only case provides a comparison to show how much the shift up in the green energy input demand share when the policy shock occurs influences the run on oil mechanism in the model. For this case, I use a final output production function of the following form:

$$Y_t = A_C L^\alpha K^\gamma O^{\nu(1-\alpha-\gamma)} G^{\beta(1-\alpha-\gamma)}$$

Policy shocks are still given by a permanent shift to ν_t , determined by a Poisson jump process, however the policy does not alter β . As this setting alters the input share of oil as before without altering the input share of the other energy sources, it provides a different cost trade-off as compared to the baseline oil restriction and green innovation policy case. With these changes, the following proposition provides the solution for this case:

Proposition 4. *With uncertain climate policy that only impacts the oil sector input demand share, and $\nu_t = \nu$ before the policy shock and $\nu_t = 0$ after the policy shock, the value functions for the two policy regimes are given by:*

$$V_{pre}(K_t, R_t, T_t) = K_t^{c_1} v(R_t, T_t) \quad V_{post}(K_t, T_t) = \bar{c}_0 K_t^{c_1} \exp(c_3 T_t)$$

where investment and labor decisions are given by

$$\begin{aligned} I_{pre,t} &= C_1 \tilde{Y}_{pre,t} & L_{pre,C} &= \tilde{L} \\ I_{post,t} &= C_1 \tilde{Y}_{post,t} & L_{post,C} &= \tilde{L} \end{aligned}$$

Exploration and extraction are given by

$$i_{R,t} = \left(\frac{v_R \Gamma \theta}{v_R - \varphi v_T} \right)^{1/(1-\theta)} \quad N_t = \frac{v \vartheta}{v_R - \varphi v_T} + i_{R,t} R_t$$

Note $v(R_t, T_t)$ is the solution to the simplified HJB equation characterizing agent's problem (given in the appendix). The value function constants \bar{c}_0, c_1, c_3 and the FOC constants $\vartheta, C_1, \tilde{L}$ are functions of the model parameters only (also given in the appendix).

The results here are similar to the original specification. The key difference in the oil-sector only policy shock case is that there is no shift in labor supply from the final output sector to the green sector after the policy shock occurs. Therefore, $L_C = \frac{\alpha}{\alpha + \omega \beta (1 - \gamma - \alpha)}$ before and after the policy shock in this case, whereas in the baseline policy case the choice of L_C goes from $L_{C,pre} = \frac{\alpha}{\alpha + \omega (1 - \nu) (1 - \gamma - \alpha)}$ before the policy shock to $L_{C,post} = \frac{\alpha}{\alpha + \omega (1 - \gamma - \alpha)}$ after the policy shock. Because the oil-sector only policy setting has no shift in the energy input demand share of the green sector, there is no incentive to shift labor to the green sector after the policy. This fixed choice of labor and no impact on the green sector demonstrates the additional cost that arises in this setting from the policy because there is no innovation in the green sector to offset the mandate to stop oil use. Yet, even with this difference, the same mechanism that alters the marginal value of reserves and marginal cost of temperature change are still present. Policy leads firms to value future profits from oil production less because of the concern about stranded assets. Therefore, the motivation to increase production as temperature increases, driving down spot prices and the value of oil firms, still exists.

5.3.2 The Impact of No Exploration

Another important alternative setting of the model to think about is the case where oil exploration is not allowed. The no exploration case helps us to understand whether or not the motivation to run on oil is driven by the fact that oil is essentially a renewable resource in the model. The setting without exploration is equivalent to the case where the exploration adjustment cost parameter Γ is set to 0. Under this assumption, the choice of exploration is given by $i_R = 0$ and the optimal choice of extraction is given by $N_t = \frac{v \vartheta}{v_R - \varphi v_T}$. However, everything else from the baseline uncertain climate policy model setting will remain exactly the same. Thus, the main difference in the outcomes for this setting is that the choice of extraction no longer has the additional boost from exploration. However, the same fear of stranded assets and expectation of climate policy limiting future climate change play a role here, and so the mechanism for a run on oil is still in place. However, as the lack of exploration is likely to increase the marginal value of holding oil reserves, as reserves are not renewable in any way, and because we lose the additional bump in extraction from the exploration component, it is reasonable to expect that in this case the run on oil and pricing impacts will likely be muted.

5.4 Welfare Implications of Policy Shocks

A particularly interesting and valuable tool that this general equilibrium model offers is a way of determining when, if at all, a climate policy shock is welfare improving. The value function for a given regime characterizes the welfare of the economy, and thus comparing the pre- and post-policy value functions that are solved for previously allows me to do this carry out this welfare comparison in a straightforward manner. In particular, to determine if a climate policy shock is welfare improving, I can simply compare the value functions for the two different policy regimes for any combination of the state variables in the model:

$$V_{pre}(R_t, T_t, K_t) \stackrel{?}{\leq} V_{post}(R_t, T_t, K_t)$$

As a result, not only does this model allow us to determine the financial and economic implications of uncertain climate policy, but I can also quantify when policy changes are actually welfare improving and how that varies across different model settings. In the numerical results, I provide a characterization of the optimal policy regions for each type of policy to show how the welfare implications change for the various types of policies.

6 Asset Prices

Having derived the solutions to the macroeconomic side of the model, I can now derive the asset pricing outcomes. As mentioned previously, the focus of this paper is on the implications of uncertain climate policy for the oil sector, however I will also derive and discuss results about the final output firm and green energy for completeness. Furthermore, because the oil firms are only valued before the climate policy shock in my main specification, I focus on asset prices in the pre-policy state.

6.1 Decentralization

As I stated before, I focus on the solution to the planner's problem. In order to derive the input prices and asset prices from this setting I must characterize the decentralized counterpart to the planner's problem where an optimal tax incentivizes the internalization of the climate externality. The externality arises from the fact that individual consumers and firms do not account for their individual contribution to climate change and climate damages that result from the use of oil in production. I briefly characterize the necessary components of the decentralized setting to show the asset pricing results, including specifying the decentralization mechanism that generates prices that correspond to the planner's solution, leaving the full derivation for the appendix. The following proposition provides the optimal oil tax:

Proposition 5. *A decentralized market with an oil production tax, $\tau_{optimal}$, lump-sum rebated back to households gives the socially optimal outcomes, and the prices that support market clearing equilibrium. This tax is given by*

$$\tau_{optimal} = \frac{-\varphi v_T}{v_R - \varphi v_T}$$

where the oil firm production problem is given by

$$\max_{N_t, i_{R,t}} \int_0^\infty \frac{\pi_t}{\pi_0} P_{O,t} \{(1 - \tau_{optimal})N_t - i_{R,t}R_t\} dt$$

where only the revenues from oil extraction are taxed since that is the only portion that contributes to carbon emissions, and therefore climate change.

The tax policy is a simple expression in terms of the value function for the planner's problem. What matters here is the marginal cost of emissions ($-\varphi v_T$) and the marginal value of oil reserves (v_R). From the numerical results we will see that $v_R \geq 0$ and $v_T \leq 0$, meaning there is a benefit to holding reserves and a cost to increasing temperature. In a standard setting without uncertain policy, the marginal cost of climate change would increase with temperature, the marginal benefit of reserves would decrease with temperature, and the marginal benefit of reserves would decrease with reserves reflecting increasing concerns for climate damages and increasing scarcity of a valued commodity such as oil. However, as the uncertain climate policy alters these marginal benefits, motivating a run on oil as discussed before, the optimal tax will reflect these changes as well.

6.2 Spot Prices

The spot prices are calculated directly from the first order conditions for the final output firm's profit maximization problem, applying the planner's optimal choices of O_t , G_t , I_t , and L_t . The energy spot prices are given by $P_{O,t}$ for oil and $P_{G,t}$ for green energy. Equity represents a claim to the stream of future dividends, which is revenues minus cost. As given firm wants to maximize shareholder value, or profit, taking the SDF π_t as given, the representative final goods output firm solves:

$$\begin{aligned} \max_{L_{C,t}, I_{K,t}, O_t, G_t} E \int_0^\infty \pi_t (\tilde{Y}_t - w_t L_{C,t} - P_{I,t} I_t - P_{O,t} O_t - P_{G,t} G_t) dt \\ \text{subject to } dK_t = K_t (\ln B + \delta_1 \ln I_t - \delta_2 \ln K_t) \end{aligned}$$

From this firm problem, spot prices can be derived from the first order conditions for $L_{C,t}$, O_t , and G_t , and are given as follows:

Proposition 6. *Wages, the spot price for oil, and the price for green energy are given by*

$$\begin{aligned} w_t &= \alpha \tilde{Y}_t L_{C,t}^{-1} \\ P_{O,t} &= \nu(1 - \alpha - \gamma) \tilde{Y}_t O_t^{-1} \\ P_{G,t} &= (1 - \nu)(1 - \gamma - \alpha) \tilde{Y}_t G_t^{-1} \end{aligned}$$

which come from the first order conditions of the final output firm.

This representation helps demonstrate the inverse relationship between oil production and the spot price of oil. As a result, we expect a run on oil production will lead to a drop in oil spot prices because of the significant supply of oil in the market.

6.3 SDF, Prices, and Returns

Following the derivation of Duffie and Skiadas (1994), the stochastic discount factor (SDF) for preferences of the Duffie-Espstein-Zin type is given by $\pi_t = \exp(\int_0^t h_V ds) h_C$. The SDF is essential to deriving asset prices because it incorporates the information necessary to properly discount firm profits over time and across states of nature. For this reason, the risk-free rate and the compensations required for holding certain risks, or the prices of risk, are derived from the SDF's drift and volatility, respectively. Specifically, an application of Ito's lemma to π_t provides us with the evolution of the SDF, $\frac{d\pi_t}{\pi_t}$, and the aforementioned prices in the following proposition:

Proposition 7. *The evolution of the stochastic discount is given by*

$$\frac{d\pi_t}{\pi_t} = -r_{f,t} dt - \sigma_{\pi,K} dB_K - \sigma_{\pi,R} dB_R - \sigma_{\pi,T} dB_T - \Theta_\pi dJ_t$$

where $r_{f,t}$ is the risk-free rate, $\sigma_{\pi,K}, \sigma_{\pi,R}, \sigma_{\pi,T}$ are the compensations for the diffusive risks of capital, oil reserves, and temperature, respectively, and Θ_π is the compensation for the jump risk of uncertain climate policy. Note that J_t is the Poisson process for the jump transition of ν_t . Expressions for these compensations are given by

$$\begin{aligned} \sigma_{\pi,K} &= (\gamma - c_3) \sigma_K \\ \sigma_{\pi,R} &= \left\{ \nu(1 - \alpha - \gamma) \frac{O_R}{O} - \frac{v_R}{v} \right\} \sigma_{RR} \\ \sigma_{\pi,T} &= \left\{ \nu(1 - \alpha - \gamma) \frac{O_T}{O} - \frac{v_T}{v} - \eta \right\} \sigma_T \\ \Theta_\pi &= \left\{ 1 - \frac{V_{post} \tilde{Y}_{post}^{-1}}{V_{pre} \tilde{Y}_{pre}^{-1}} \right\} \end{aligned}$$

I leave the expression for the risk-free rate for the appendix since it is fairly cumbersome. The expressions for the risk prices are useful for intuition, even though most are not in closed form

and therefore numerical solutions are needed to fully characterize the outcomes. First, each diffusive component follows the standard asset pricing result of scaling a risk aversion component by a volatility component. The expression for the capital risk price is in closed form and is constant. For reserves and temperature risk, we see the value function, production choices, and derivatives of the value function and production choices matter for the risk aversion component, which is our first clue about how uncertain climate policy will impact asset prices. As mentioned previously, uncertain climate policy alters the production choices and value function, as well as the marginal values or derivatives of these outcomes of interest. Thus, the same mechanisms driving the run in production of oil, the non-linear, non-monotonic behavior of the marginal values of reserves and temperatures due to fear of stranded assets, determine the risk compensations as well. Furthermore, the run on oil itself affects the risk prices since the production of oil and its derivatives influence these expressions as well. We also see that the uncertain climate policy contributes directly to risk prices through the compensation required for the jump risk of changes in the energy share coming from policy. The size of this jump risk depends of how significant the welfare change and production change we be between the pre- and post- policy economies.

Now, given the SDF and corresponding risk-free rate and prices of risk, we can derive the firm prices. This derivation requires using the solutions for input prices and quantities derived from the macroeconomic side of the model to compute the profits or dividends for each firm. I assume the firms are 100% equity-financed firms and so profits and dividends correspond one-to-one. Once we have the firm prices, we can also derive the risk premia for the model. The following proposition provides the firm prices and risk premia for the oil firm, green energy firm, and final output firm in the baseline uncertain policy setting:

Proposition 8. *The prices for the final output firm (S^C), green energy firm (S^G), and oil firm (S^O) in the pre-policy state can be derived from the value function by application of the envelope theorem, and the resulting firm values are given by:*

$$S_t^C = a_C \tilde{Y}_t, \quad S_t^G = a_G \tilde{Y}_t, \quad S_t^O = a_O \frac{v_R R}{v} \tilde{Y}_t$$

where a_C, a_G, a_O are constants which are given in the appendix.

Risk premia for firms $X = C, G, O$ in the pre-policy state, $RP^X = -cov(\frac{d\pi_t}{\pi_t}, \frac{dS^X}{S^X})$, are given by

$$RP^X = \gamma(\gamma - c_1)\sigma_K^2 + \sum_{\chi=R,T} \left(\left(\frac{\partial}{\partial \chi} S^X \right) / S^X \right) \sigma_\chi(\chi) \sigma_{\pi,\chi} + \lambda(T_t) \Theta_\pi(S_{post}^X / S_{pre}^X - 1)$$

where the expressions for the functions used here are given in the appendix.

Full details of the derivations and expressions used here can be found in the appendix. For intuition, note that for prices, returns, and risk premia, the impacts of K are independent of the

impact of the remaining state variables T and R . Therefore, the capital impact simply scales the firm prices and provides a constant additive contribution to the risk premia, but does not interact with the uncertain climate policy. Though the asset pricing impacts associated with oil reserves and temperature can only be determined numerically, we know the impact of climate policy will matter because of the expressions for the risk prices and the firm values.

Consider first the final output sector and the green energy sector. The prices of the final output firm and green energy firm are both proportional to final output scaled by climate damages, \tilde{Y}_t . Thus, we can expect two forces to play a role here. First, over time we expect climate change to increase and thus the damages to increase, bringing down the value of the final output firm and green energy firm. However, due to the risk of uncertain climate policy we expect there to be a run on oil production. This run on oil leads to an increase in the oil used in final output production and thus an increase in final output production itself. As a result, this force should increase the final output firm value and green energy firm value. The numerical results will help us determine which of these forces dominates.

The price of the oil firm also includes the damage-scaled final output, and so forces related to the impact of climate damages and the impact of the run on oil production for the damage-scaled final output that impact the final output firm and green energy firm prices still matter here. However, the price of the oil firm is also scaled by the marginal value of reserves v_R . Thus, the oil firm has an additional force impacting its price. We know from the macroeconomic outcomes that the risk of stranded assets from the uncertain climate policy cause the marginal value of oil reserves to decrease over time as reserves diminish and climate change increases. Therefore, we expect that the price of the oil firm will be lower than without uncertain climate policy, due to the reduced value of holding oil reserves in this setting, and that the price will decrease dynamically as well, due to the increasing likelihood of policy occurring that is also driving the run on oil.

Characterizing the impact of uncertain climate policy on asset prices, as I have done here, is an important contribution of my analysis. Given the relatively small realizations of climate change we have experienced so far, the fact that asset prices are forward looking in nature and incorporate expectations about future uncertainty is critical for understanding the impacts of the model mechanism I study in this model and providing testable model predictions. While the measurable impact on macroeconomic outcomes due to uncertain climate policy may still be quite small so far in the macroeconomic data, or even negligible, the asset pricing outcomes should provide greater power to identify the slow-moving, long-term risk and concerns about climate change and uncertain climate policy due to the additional forward-looking information they incorporate.

7 Numerical Solutions

Given the theoretical results above, I now discuss the numerical results of the model. I first discuss briefly the model parameters and numerical method used to solve the model, and then delve into the

solutions based on the parameters and solution method given.

7.1 Model Parameters

The parameters I use for the solutions are given in table 1. Although the theoretical model is designed to qualitatively demonstrate the novel uncertain climate policy mechanism, I choose parameter values in order to provide reasonable values for the economic and financial outcomes of interest. The discount rate and risk aversion parameters are chosen to be relatively conservative, consistent with other values used in the production-based asset pricing literature and climate-economics literature such as Papanikolaou (2011) and Golosov et al. (2014). The choices for initial TFP in each sector and the capital and labor input shares are monthly counterparts to the values used in Golosov et al. (2014). The choices for the capital adjustment cost parameters are consistent with Anderson and Brock (2017). The capital and oil reserves volatility are chosen to be monthly counterparts within the range of values used in Carlson et al. (2007), Casassus et al. (2009), and Kogan et al. (2009) and consistent with the 2018 BP Statistical Review of World Energy data on oil reserves. The values for exploration adjustment costs are chosen to be conservative, such that there are diminishing returns to exploration and a decreasing oil reserves over the time series simulations. For the energy input demand shares, I choose values to reflect high current oil demand in production.

The parameters relating to the climate part of the model are chosen as follows. The temperature volatility is a monthly counterpart to the value estimated by Hambel et al. (2015). The damage function parameter is chosen to match with Golosov et al. (2014). The climate sensitivity parameter comes from the estimate provided by Matthews et al. (2009) and Matthews et al. (2012).

Lastly, the parameters for the uncertain policy arrival rate have no clear counterparts that I am aware of in the literature. I choose values that allow for demonstration of the model mechanism in a setting where a large policy shock is likely to occur within 30-40 years from the beginning of the model simulations. Work calibrating these parameters so that model-generated quantity and asset pricing outcomes are consistent with values observed in the data, although very interesting, is left for future work. Furthermore, the appendix contains details on necessary parameter restrictions that must hold for the model and for convergence of the numerical results.

7.2 Numerical Method

I now briefly discuss the numerical method used to solve the model for each of the different specified frameworks mentioned previously. I use the Markov chain approximation method proposed by Kushner and Dupuis (2001) to solve the partial differential equations (PDEs) that characterize the model. As the name implies, this method uses a Markov chain approximation to discretize the continuous-time problem. Under fairly simple-to-verify conditions for the Markov chain approximation, convergence of the approximated solution to the true solution is guaranteed. The solution method, in many ways,

Table 1: Model Parameters

Discount Rate	ρ	0.005
Risk Aversion	ξ	1.5
Final Output Capital & Labor Shares	γ, α	{0.6, 0.3}
Green DRS Parameter	ω	0.9
Oil Energy Input Share	ν	{0.9, 0}
Final Output & Green TFP Values	A_C, A_G	{240.76, 10.93}
Capital Adjustment Costs	B, δ_1, δ_2	{1.13, 0.03, 0.03}
Capital Volatility	σ_K	0.1
Reserves Volatility	σ_R	0.02
Temperature Volatility	σ_T	0.03
Climate Sensitivity	φ	0.0024
Policy Arrival Rate Parameters	ψ, ϖ, p	{0.5, 0.01, 4}
Climate Damages Parameter	η	0.01
Exploration Adjustment Costs	Γ, θ	{0.05, 0.5}

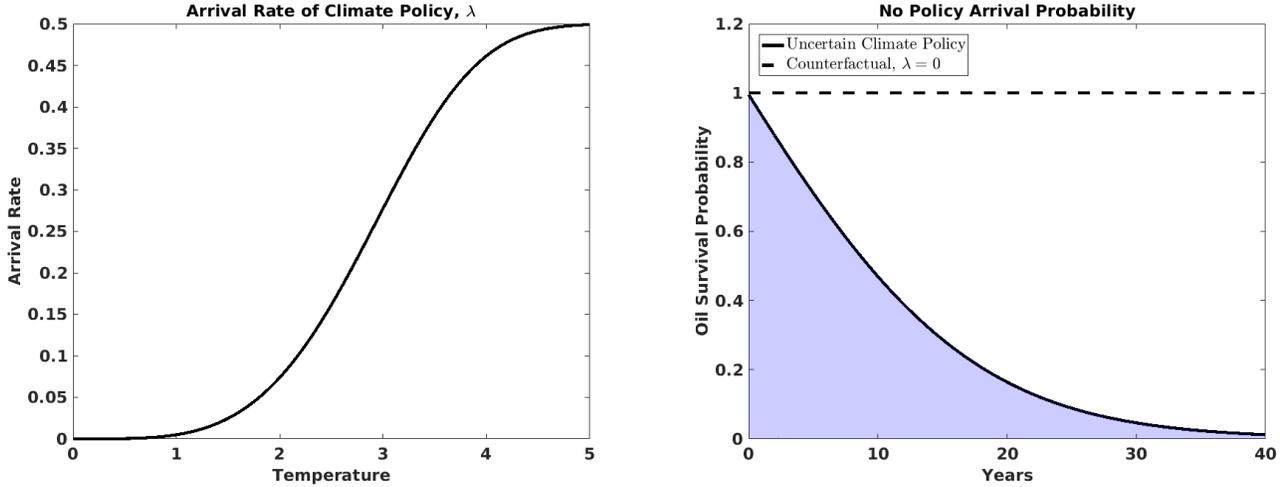
is similar to discrete-time value function iteration and provides a fairly intuitive and robust method for solving the PDEs. More details can be found in the appendix.

7.3 Simulated Time Series Comparisons

For the numerical results, I plot the simulated time series using the model solutions for the key outcomes of interest. I focus on the two central cases: the counterfactual comparison setting without uncertain climate policy and the uncertain climate policy setting central to my model. The results are generated by averaging over 10,000 simulations, each for 600 months, starting from the same initial conditions for each of the state variables. I choose the initial conditions for reserves and temperature to be $R_0 = 450$, $T_0 = 1$ and show the results for 40 years. T_0 is approximately the current global mean temperature anomaly from pre-industrial levels, and R_0 fits within the estimates for global oil reserves available according to values cited by Golosov et al. (2014), the 2018 BP Statistical Review of World Energy data, and estimates of existing reserves by Rystad Energy. I provide a brief discussion below of the impact of exploration and leave the plots for these results to the appendix. The counterfactual, no uncertain climate policy case (labeled as "Counterfactual, $\lambda = 0$ ") is given by the dashed line and the main model outcome with uncertain climate policy (labeled as "Uncertain Policy") is given by the solid line.

Figure 6 provides two plots for the policy arrival rate function. The first plot, "Arrival Rate Function of Climate Policy, λ ," shows the arrival rate of climate-dependent policy $\lambda(T_t)$ as a function of temperature. This function captures the main force in my model, and therefore understanding this function will help provide intuition for the solution to the model. We can see the function is strictly increasing in T , and the shape goes from a relatively gradual increase to a fairly steep and dramatic increase as temperature starts to get quite high. Note also that the arrival rate is only a

Figure 6: Uncertain Policy Arrival Rate Function & Cumulative Probability



function of T , and so there are no changes to policy arrival due to changes in the level of reserves. The second plot, “No Policy Arrival Probability”, begins the time series results for the model simulations by showing the cumulative probability of no policy shock occurring. Even though the change in the arrival rate realized in the simulations is relatively modest, the cumulative probability of the policy shock not occurring approaches zero after about 40 years. For the no uncertain policy setting, which by definition has $\lambda(T_t) = 0$, this cumulative probability of no policy occurring is one. I include this “oil survival probability” as the blue shaded region for each of the time-series plots that follows in order to provide an idea of how likely the time-series realizations are, given that I suppress the realization of policy shocks from occurring to provide a full time-series representation.

The next set of plots, Figures 7 and 8, show oil extraction, oil exploration, atmospheric temperature change, oil reserves, the spot price of oil, the oil firm price, the final output firm price, and the green energy firm price. Note that firm prices are normalized so that the counterfactual-setting firm value is zero at the initial period for convenience in interpretation. Beginning with extraction, we see that for the uncertain policy setting a run-up in extraction occurs over the time series, gradually flattening out around year 40. The run-up comes as a result of the increasing likelihood of the arrival of climate policy, causing the marginal value of reserves and the marginal cost of temperature to decrease even as reserves are decreasing and temperature is increasing. The tapering off is a result of the increasing significance of the effect that as reserves start to diminish, and as policy has not yet been implemented, the firm decides to slow production to save some oil for potential future profits. For the counterfactual setting, not only is oil extraction gradually decreasing as in the standard resource extraction model, but the level of oil extraction is lower as well. The increased oil production over time, combined with the level increase in oil production, is the run on oil mentioned previously for the uncertain policy setting. This choice to run up oil production is optimal for the planner in the model, even given

the potential climate and policy implications, because the planner wants to maximize gains from oil production now before oil reserves become a stranded asset.

The exploration in level terms is substantially lower in the baseline case than for the counterfactual setting. This level difference in exploration reflects the increased discounting that results from the uncertain climate policy, whereas in the counterfactual without uncertain climate policy the oil firm behaves as if it will produce forever and thus exploration is extremely valuable in that setting. However, over the region where there is a run on oil production there is an increase exploration as well, likely reflecting the desire to maximize the benefit of oil reserves before they become stranded. Thus, exploration appears to augment the dynamic effects of uncertain climate policy.

The consequences of the run on oil can be seen in the level of reserves and the change in temperature observed in the model-simulation time series. As expected, the uncertain policy setting where extraction is higher leads to lower reserves and substantially higher temperature than the counterfactual setting. The temperature difference after 40 years is more than a quarter of a degree Celsius. Given the concern in the scientific and policy-making communities about reaching the 2° C temperature threshold, and as we are already at 1° C, this difference is meaningful. However, one has to take into account the probability of achieving the full temperature change shown is quite low. Note that the no uncertain policy case is significantly different from the uncertain policy setting. The level of reserves remains quite high throughout the simulated time series, and the magnitude of climate change is relatively small because of the increasing value of holding reserves and increasing cost of climate change. These increasing marginal values are due to the fact that oil reserves become increasingly scarce, climate change becomes increasingly costly, and so oil use is taxed increasingly more and there is no expectation that oil will be replaced. Thus, climate change impacts are pushed to the future, whereas the run on oil brings forward some of that impact.

Next consider the pricing outcomes in the simulated-model results. I focus here on the firm values and spot price of oil in particular, but the results for the market prices of risk can be found in the appendix. The financial outcomes reflect what we see in terms of the run on oil in the macroeconomic side of the model and the concerns about the risk of stranded assets. First, in the uncertain policy setting where extraction is quite high, the spot price of oil is quite low, reflecting the substantially increased supply of oil in the market. Furthermore, the run on oil is seen in the spot price as the price has a downward sloping path over the time series. For the no uncertain policy setting, the spot price gradually and monotonically increases as extraction gradually decreases.

The evolution of the final output firm price reflects the change in spot prices and expectations of future policy change. Since the cost of the oil input diminishes, the value of the final output firm increases, coinciding in shape and timing with the run on oil production. Again, the dynamics of the no uncertain policy setting markedly differ. For this counterfactual setting, there is a gradual decrease that occurs in the final output firm price over the time series, due in large part to increasing oil prices and reduced oil output. Therefore, in the uncertain policy setting there is a portion of the economy

Figure 7: Uncertain Policy Arrival Comparison - Quantities

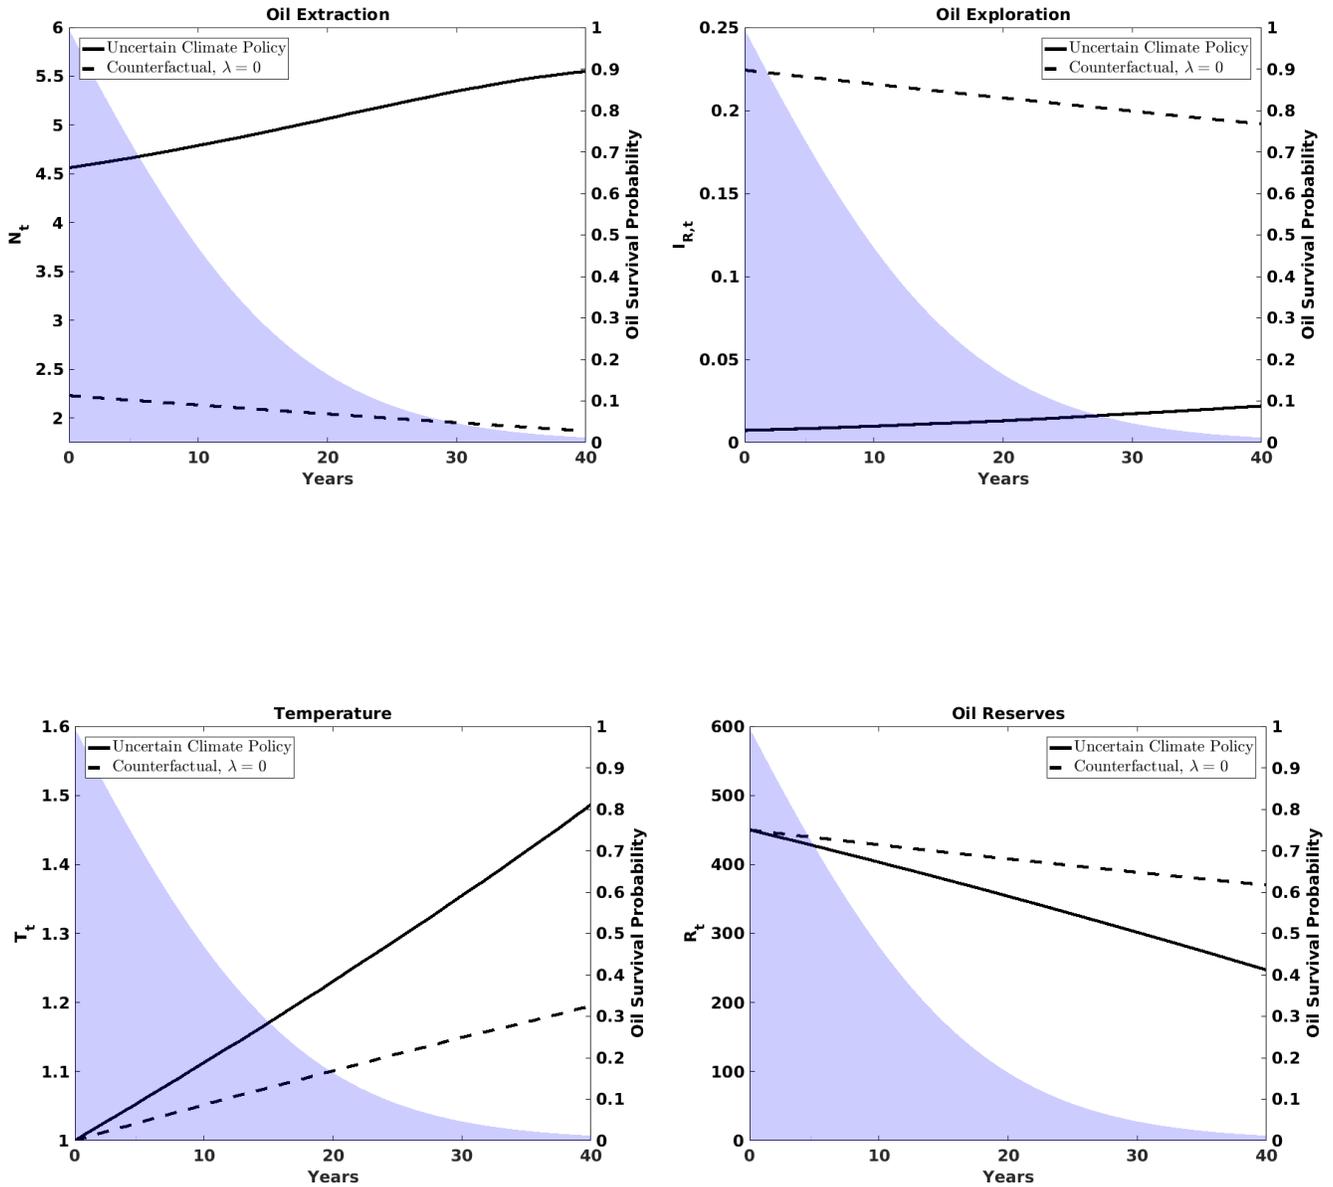
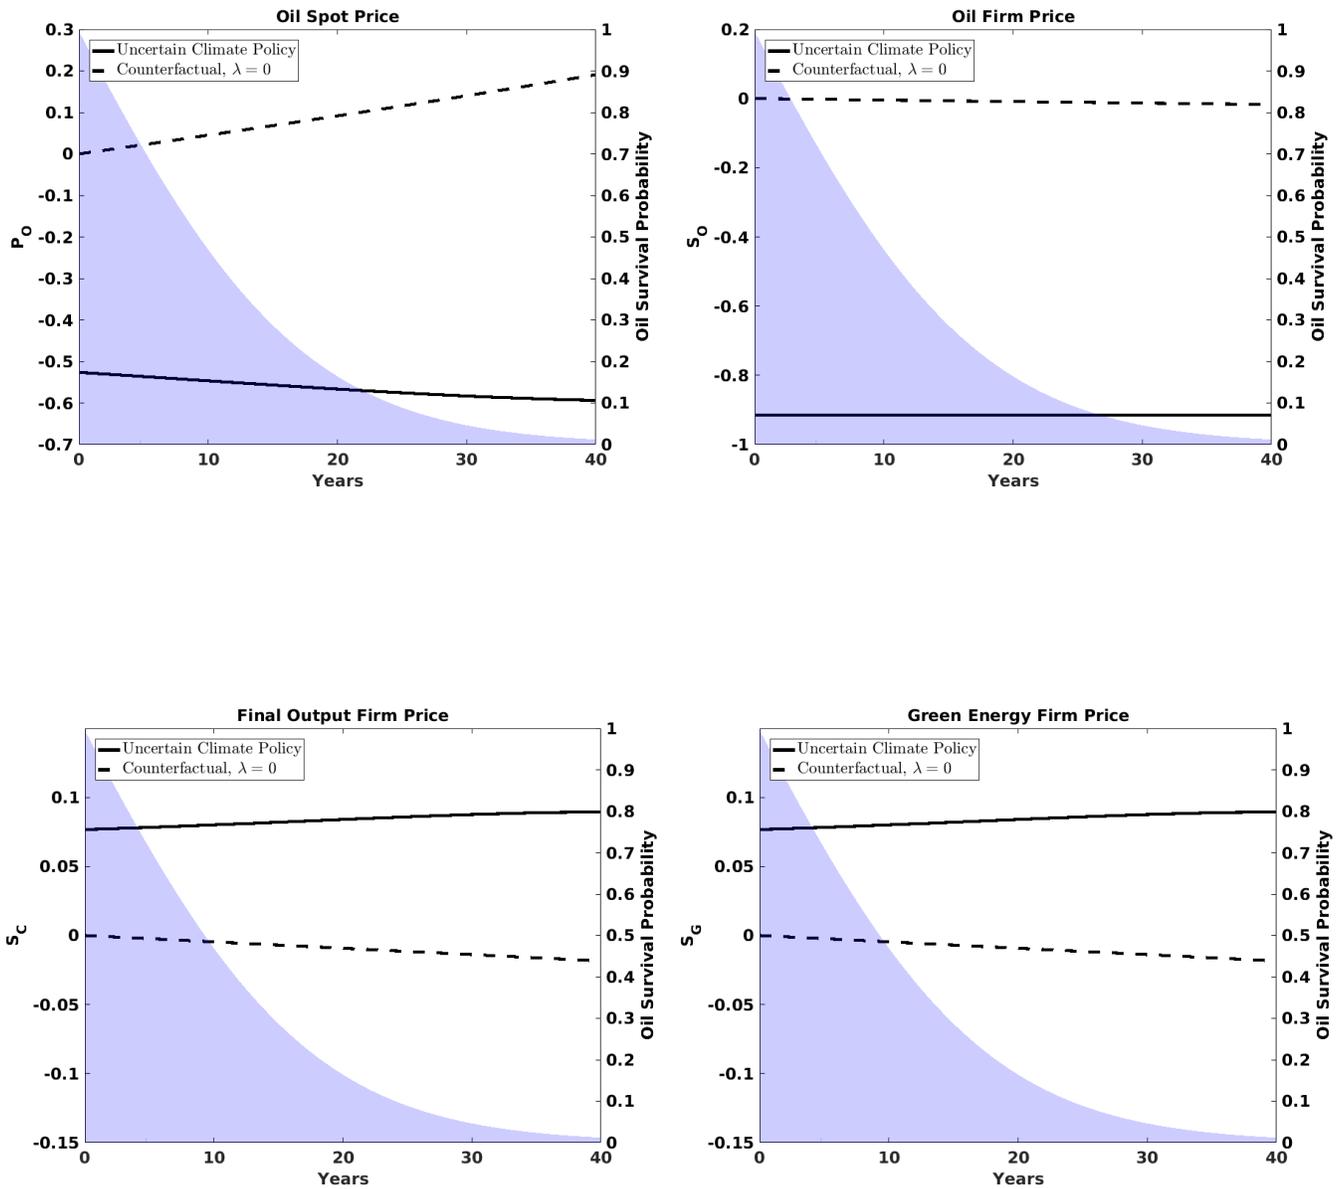


Figure 8: Uncertain Policy Arrival Comparison - Prices



that appears to benefit from the risk of uncertain climate policy, whereas in the counterfactual case without uncertain climate policy there is no such benefit for firm prices.

The behavior of the green firm is similar to the final output firm. Again, for the green firm price the impact of the price change is due to expectations about future policy change and changes in the wage from the run on oil. Because the choice of labor for the green input sector does not change until after the policy shock occurs, these changes are not due to changes in green production. The change in the green firm price without changes in green production demonstrate the value of using asset prices to understand the full impact of uncertain climate policy because they incorporate forward-looking policy expectations that alter firm prices even if output for a given sector does not change.

The impact of uncertain policy on oil firm prices is particularly significant. For the uncertain policy setting, there are two key impacts of uncertain climate policy. There is a significant shift down in the level of the oil price when compared to the counterfactual case. This shift down in price reflects in part the fact that oil reserves will likely become stranded, and thus a portion of the oil reserves that the oil firm has are actually worthless. The level difference also reflects the fact that oil reserves that will be used are less valuable because the oil firms are producing at a higher rate and so the spot price of oil is markedly lower. The second impact of uncertain policy is the dynamic decrease in the oil firm price. This dynamic decrease is seen in the slight downward slope in the path of the oil firm price, which is a result of, and thus coincides with, the run on oil and drop in spot prices. Again, the counterfactual setting with no uncertain policy differs substantially. Because there is no risk of stranded assets, the value of the oil firm is substantially higher. Moreover, because spot prices and reserves remain higher in this setting the value of the oil firm is further boosted.

7.3.1 Understanding the “Carbon Bubble”

The significant difference in the price of the oil firm when not accounting for uncertain climate policy is a key outcome that has an important interpretation. Work done by the Grantham Institute at LSE and others has highlighted the potential existence of a “carbon bubble,” meaning that oil and other fossil fuel firms are potentially overvalued due to the fact that implementation of a temperature ceiling policy such as the Paris Climate Accord will likely require that a substantial amount of fossil fuel reserves must remain in the ground and are therefore worthless. Oil firm prices determined based on expectations that there is no uncertain climate policy are therefore substantially higher because of the expectation that all the oil reserves these firms hold can eventually be used.

An important contribution of my paper is that it provides a dynamic, general equilibrium model setting that can be used to characterize this potential “carbon bubble.” Because of this dynamic, production-based asset pricing modeling framework, I am able to identify a novel, additional factor exacerbating this potential “carbon bubble” beyond the static effect that reducing oil reserves to account for the fact that a fraction of oil reserves may not be usable due to climate policy would generate. By incorporating the dynamic effects of uncertain climate policy and stranded assets on

oil firm decisions and valuations, my model shows that the counterfactual setting without uncertain climate policy not only misses the impact on prices resulting from the fact that significant amounts of oil reserves may in fact be worthless, but the mispricing in the counterfactual setting is amplified by also not accounting for the impact of a run on oil production induced by the stranded assets risk. This additional mispricing effect exists because the counterfactual, no uncertain policy setting, as well as other models of resource extraction, implies slowly diminishing oil extraction over time, leading to a gradual increase in oil prices. The decrease in production and increase in the spot price of oil in these alternative model settings helps keep the oil firm price higher. The run on oil induced by uncertain climate policy in the main setting of my model drives up oil extraction, leading to a dynamic decrease in the price of oil which negatively impacts the value of oil firms and leads to an even larger difference in the oil firm price when compared with the counterfactual, no uncertain climate policy case. This impact on oil firm values provides an additional component not previously captured in previous work studying a potential “carbon bubble,” and highlights the importance of incorporating the dynamic, non-linear implications my model captures.

7.4 Model Extensions Comparisons

The other cases I consider are the case with no oil reserves exploration and the case with oil-sector only policy effects, which I briefly discuss here and provide plots for these results provided in the appendix (Figures 13, 14, 15, and 16). There is little difference in the dynamics of the no exploration case when compared to the “Uncertain Policy” case with exploration. The level of oil extraction, and therefore temperature and reserves is nearly identical. The biggest difference in these two settings is the value of the oil firm. The value of the oil firm is higher in the no exploration case as extraction is nearly identical but there is the obvious difference in exploration and therefore costs. As such, the oil firm is actually more valuable for the given set-up in the no exploration setting as a result. Furthermore, because the cost of stranded assets is diminished for the no exploration case as there are no potential discoveries that can also be stranded, this also helps to keep the value of the oil firm higher when there is no exploration, given the similar choice of oil extraction for the two settings.

For the oil-sector only climate policy case, the outcomes are also quite similar to the baseline model case, with minor increases in extraction, exploration, temperature change, and oil firm value and minor decreases in the oil price and reserves. The difference in this model setting, which comes from the policy shock no longer including an increase in the input demand share of green energy, appears to have little effect on outcomes before the policy shift. Though the change in policy means there are somewhat larger costs of the policy switch, there is no significant change in outcomes before that switch happens. This is likely due in part to the small role green energy plays in this economy. However, the similarity of the results implies that the same fear of stranded assets drives up oil production and influences the asset prices observed before the realization of the policy change.

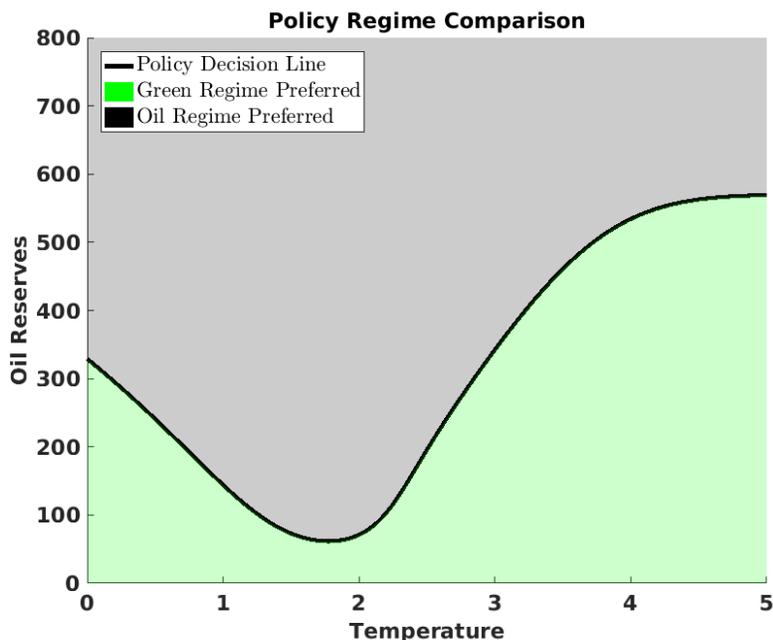
7.5 Policy Welfare Comparison

Next, I examine whether the climate policy shocks can be welfare improving compared to remaining in the pre-policy, oil-producing regime. These comparisons are given in Figure 9 and in the appendix for the model extensions (Figures 17 and 18). In the figures, the state space where the pre-policy, oil-producing regime is better in terms of welfare is shaded in black and is labeled the “Oil Regime” region. The state space where the post-policy, no-oil regime is preferred in terms of welfare is shaded in green and is labeled the “Green Regime” region. For Figure 9, which is the regime comparison for the baseline climate-dependent arrival rate setting, we see that for higher values of temperature the threshold for oil reserves where the green regime is preferred is higher. The exception to this is at the lowest temperature values, due at least in part to initial increases in damages being fairly costly and initial increases in the arrival rate of policy being quite small because of the functional forms chosen. The frontier dividing these preferred regions is similar to the shape of the arrival rate function. However, we can see that even for high temperature levels the planner would still prefer to be in the oil-producing regime when there are still significant enough reserves remaining. The reason for this preference is that the planner knows there is a significant likelihood that the policy regime will switch, given temperature is so high, and so being able to use more of the oil reserves for production brings current production benefits and reduces the stranded assets cost, while expectations are that continued climate change will be limited because of policy implementation.

The oil-sector only policy case and the no exploration case welfare comparisons look quite similar. Plots for these results are contained in the appendix. As before, lower temperature and higher reserves mean that the oil regime is preferred in terms of welfare for each setting. The state space where the oil regime is preferred is slightly larger for the oil-sector only case because the policy shock has a somewhat higher cost as there is no increase in the green energy input share when the policy occurs. While the impact is relatively small here, this welfare comparison demonstrates that policies that support increases in green technological innovation with the mandate to reduced emissions through reduced oil production are preferred to policies that simply increase the cost of using oil in the social welfare sense. As a result, policies based on increasing the benefit of using green energy are also likely to translate into policies that are preferred in terms of public approval and voter support because the expected cost of such policies are lower than for policies without accompanying green innovation.

In the no exploration case, the state space where the oil regime is preferred is also slightly larger due to the fact that without exploration, the cost of stranding oil reserves is lower because there is no renewability of reserves. Again, while the difference is quite small, this welfare comparison demonstrates that expectations about the level of potentially recoverable reserves, and not just existing recoverable reserves, influences the welfare impact of the policy shock as well. In countries that have high expectations of oil discoveries, a mandate to stop using oil is even more costly because of the stranded assets risk for currently recoverable reserves and expected reserves from future oil discoveries.

Figure 9: Climate Dependent Arrival Rate Comparison



8 Empirical Analysis

The solution to the model I use for my analysis of uncertain climate policy provides a number of important predictions that I now examine empirically. The first prediction is that uncertain climate policy generates a run on oil. The model also predicts that uncertain climate policy depresses the spot price of oil because of the increased oil production. Empirically, this corresponds to an observed increase in the likelihood of future climate policy leading to an increase in current and future oil production and a decrease in current and future oil spot prices. Also, the model predicts that the value of oil firms decreases due to the risk of stranded assets leading to expectations that not all oil reserves held by firms can be used and because the run on oil reduces oil prices. Finally, the model predicts that the value of the final output firm and the green energy firm will increase due to decreased oil prices and policy expectations. These final two predictions can be re-stated as empirically observed increases in the likelihood of future climate policy occurring should lead to oil firms experiencing negative returns and non-oil firms experiencing positive returns.

I examine these predictions using the following empirical exercises. The first exercise examines the impact of events that shift the likelihood of future climate policy action occurring. Estimating cross-sectional regressions for returns of US sector portfolios on a proxy for climate policy exposure shows sectors with greater climate policy exposure experienced larger increases in returns from climate policy events that decreased the likelihood of future climate policy action and larger decreases in returns from climate policy events that increased the likelihood of future climate policy action. Thus, the

post-event outcomes provide evidence that the shift in expectations for future climate policy did impact asset prices as my model predicts.

I then extend this event-type analysis of the model predictions by exploring the impact of the time series of “climate policy” events. I construct an index of the time series of “climate policy” events by aggregating lists of key climate- and energy-related events from non-partisan, informational websites. With this time-series index of relevant events, I first test the model predictions using reduced-form regressions of the impact of “climate policy” shocks on oil production in different regions, returns of US oil sector firms, and returns of the oil spot price. The second approach focuses on the dynamic impact of climate policy shocks by estimating a vector autoregression that incorporates the climate policy events index into a standard global oil market model to examine the impact of climate shocks on current and future oil production and oil prices. Results from each of these exercises again appear to be consistent with the model predictions.

8.1 Data Sources

The data I use for oil production and oil prices comes from the US Energy Information Administration (EIA). I use global atmospheric temperature as the variable for climate change, which is available from NOAA and NASA. Data on returns for the oil sector and the market come from Ken French’s website and the Compustat/CRSP merged database available from WRDS. Macroeconomic variables such as GDP growth, deflators, and other indicators come from FRED, the BEA, Lutz Kilian’s website, and James Hamilton’s website. Additional oil spot price data come from FRED and the CME’s End of Day database, made available to me through the University of Chicago Booth School of Business Fama-Miller Center. Finally, I construct a proxy for changes in the probability of climate policy occurring in my model by compiling a time series of significant climate, climate policy, and energy events (major fossil fuel and alternative energy events, IPCC meetings, US presidential election results, and lists of major climate policies and US energy policies) from non-partisan government, academic, and non-profit informational websites (ProCon.org, IPCC website, and Wikipedia.org). Table 19 shows the list of events since 1997, though the full list extends back to match the full range of dates available from the EIA (1973). As indicated in the table, events can either be positive or negative policy shocks in terms of an increased or decreased likelihood, respectively, of a shift to the production function. The variable contains values of 0 for no event, 1 for a positive event, or -1 for a negative event at the daily level that are then aggregated to a monthly count for the empirical analysis. Section 8.3 provides further details on the index.

8.2 Climate Policy Event Study Analysis

The first empirical exercise I do is to estimate the impact of events that in expectation would shift the likelihood of future climate policy action occurring on stock returns for different sectors in the

US economy based on their estimated exposure to climate policy. I focus on events for which signal changes in expectations for future policy outcomes rather than actual policy implementation as this more directly tests the mechanism in my model which relates to the likelihood of future climate policy action. Within the context of my model, these events can be compared to comparative statics in the model where $\lambda(T_t)$ is shifted up or down. Moreover, the more unexpected the shock to the value of $\lambda(T_t)$, the more cleanly we can identify the impact as a comparative static shock rather than a more prolonged response to an expected outcome. Such shocks have clear implications in the model: a shock that increases (decreases) the likelihood of future climate policy action should lead to increased (decreased) oil production, negative (positive) realized returns for the oil sector, and negative (positive) realized returns for the spot price of oil due to increased (decreased) stranded assets risk.

To formalize this exercise using an event-study analysis, I estimate the impact of the unexpected shift in climate policy risk expectations on cumulative abnormal returns after the election by exploiting the cross-sectional variation in climate policy risk exposure across different sectors. To estimate this cross-sectional regression, I use daily returns for the 49 sector portfolios provided on Ken French's website. I derive abnormal returns as unexplained differences with respect to the market portfolio, or the CAPM model. I estimate the following regression for each sector i using daily returns for the year leading up to the event date to estimate abnormal returns:

$$R_{i,t} = \alpha_i + \beta_{i,Mkt}(R_{Mkt,t} - R_{f,t}) + \epsilon_{i,t}$$

The residual for this regression $\epsilon_{i,t}$ is then the abnormal return. I aggregate these residuals for each sector in order to get the cumulative abnormal returns:

$$CAR_{i,t} = \sum_0^t \epsilon_{i,t} = \sum_0^t (R_{i,t} - \alpha_i + \beta_{i,Mkt}(R_{Mkt,t} - R_{f,t}))$$

Next, I need to derive a measure of climate policy risk exposure. My model predicts that changes in climate-policy expectations influence oil prices and oil production, as well as firm values. Therefore, the model predictions suggest a sector's exposure to climate policy risk can be proxied for by the sector's exposure to oil price innovations or oil production innovations. Given that oil prices are available at a daily frequency, are in direct units of comparison, and are closely linked to oil production, I use exposure to oil price returns as the proxy for exposure to climate policy risk. I estimate this exposure as the beta for oil price returns from the following regression for each sector i over the full available time series of oil prices:

$$R_{i,t} = \alpha_i + \beta_{i,Mkt}(R_{Mkt,t} - R_{f,t}) + \beta_{i,OilPrice}R_{OilPrice,t} + \epsilon_{i,t}$$

To estimate the impact on returns of climate policy risk exposure after the event, I run the

following cross-sectional regression:

$$CAR_{i,event} = \delta_0 + \delta_1 \frac{\beta_{i,OilPrice}}{\sigma(\beta_{i,OilPrice})} + e_i$$

Because I normalize the climate policy risk beta by the cross-sectional standard deviation of the beta estimates ($\sigma(\beta_{i,OilPrice})$), the coefficient δ_1 can be interpreted as the percent change in cumulative abnormal returns due to a one-standard deviation increase in the climate policy risk beta resulting from the change in climate policy expectations from the event outcome.

Table 2 provides estimates for a number of recent climate policy-related events. The estimates are the cumulative abnormal return response one day and 4 weeks after the policy events for value- and equal-weighted sector portfolios with t-stats for the heteroskedastic-robust standard errors and the z-stats for the bootstrap errors for the two-stage estimation to account for the inclusion of a generated regressor. Events in the table are those recent events where the impact on returns was statistically significant. Other events tested (such as other US presidential elections since 1996 and the Kyoto Protocol) provided null results.

In each case with significant estimates, the results are consistent with the model predictions. For events that increased the likelihood of future climate policy (the publication of the Clean Power Plan and date of the Paris Climate Agreement) there is a negative CAR response for sectors with higher climate policy exposure. For the events that decreased the likelihood of future climate policy action (the Trump Presidential elections, the announcement date of the US plan to withdraw from the Paris Climate Agreement, and the US Supreme Court ruling to put a hold on implementing the Clean Power Plan) there is a positive CAR response for sectors with higher climate policy exposure.

In the appendix, I provide scatter plots of the one-day and four-week cumulative abnormal return responses for the value- and equal-weighted sector portfolios, sorted by climate policy risk exposure beta, and the estimated climate policy risk exposure impact slope coefficient δ_1 (with the t-statistic and z-statistic for the estimate) for two events that I want to highlight briefly. The first event was the 2016 US presidential election. This event was a surprise shift down in the likelihood of future climate policy action given President Trump's campaign statements about supporting the coal and oil sectors, withdrawing from the Paris Climate Agreement, removing emissions regulations policies such as the Clean Power Plan, and doubting the impact of human behavior on the climate. After the election, for the value-weighted (equal-weighted) portfolios a one-standard deviation increase in climate policy risk beta would have resulted in a 1.24% (1.11%) increase in cumulative abnormal returns after one day and a 2.06% (2.81%) increase in cumulative abnormal return after four weeks. For the value-weighted portfolio estimates, both estimates are statistically significant using heteroskedasticity-robust standard errors (t-statistic), while only the one-day cumulative abnormal return response is statistically significant for the bootstrapped standard errors (z-statistic). For the equal-weighted portfolio es-

Table 2: Event Study Analysis of Significant Climate Policy-Related Events

	1-Day, VW	4-Weeks, VW	1-Day, EW	4-Weeks, EW
Clean Power Plan	-1.08	-0.77	-1.01	-2.08
T-Stat: Robust SE	-4.62	-1.14	-2.64	-1.80
Z-Stat: Bootstrap SE	-4.56	-0.98	-2.63	-1.51
Paris Climate Accord	-0.68	-0.39	-0.79	-1.03
T-Stat: Robust SE	-2.49	-0.35	-5.30	-2.48
Z-Stat: Bootstrap SE	-2.49	-0.32	-4.50	-1.55
USSC Hold on CPP	0.55	4.49	0.42	6.82
T-Stat: Robust SE	2.38	2.17	0.48	8.84
Z-Stat: Bootstrap SE	1.89	2.12	0.46	6.96
Trump 2016 Election	1.24	2.06	1.11	2.81
T-Stat: Robust SE	3.05	1.82	1.96	1.99
Z-Stat: Bootstrap SE	2.47	1.57	1.93	1.90
US Paris Withdrawal	0.71	0.30	0.71	-0.07
T-Stat: Robust SE	4.25	0.43	5.52	-0.22
Z-Stat: Bootstrap SE	3.91	0.41	4.42	-0.16

This table shows the relationship between the cumulative abnormal returns of sectors after a given climate policy related event and their standardized exposure to climate policy risk. The events are major recent events which had significant responses for cumulative abnormal. The regression specification is given by $R_{i,t} = \alpha_i + \beta_{i,Mkt}(R_{Mkt,t} - R_{f,t}) + \beta_{i,OilPrice}R_{OilPrice,t} + \varepsilon_{i,t}$. Cumulative abnormal returns are normalized to zero at the election date, and are estimated with respect to the value-weighted excess market return. Estimates are for value- and equal-weighted sector portfolio cumulative abnormal returns one day and 4 weeks after the election. I provide the t-state for the coefficient for heteroskedasticity-robust standard errors and the z-state for bootstrap standard errors of the two-stage estimation procedure to account for the use of a generated regressor. See text for full definition of variables.

timates, both estimates are statistically significant using heteroskedasticity-robust standard errors (t-statistic) and bootstrapped standard errors (z-statistic).

The second event was the 2016 US Supreme Court decision to put a stay on the Clean Power Plan. This event was also a surprise decrease in the likelihood of future climate policy action given the lower courts had yet to rule on the constitutionality of the policy. After the court decision, for the value-weighted (equal-weighted) portfolios a one-standard deviation increase in climate policy risk beta would have resulted in a 0.552% (0.423%) increase in cumulative abnormal returns after one day and a 4.489% (6.823%) increase in cumulative abnormal return after four weeks. For the value-weighted portfolio estimates, both estimates are statistically significant using heteroskedasticity-robust standard errors (t-statistic) and the bootstrapped standard errors (z-statistic). For the equal-weighted portfolio estimates, only the 4-week estimated impacts are statistically significant, for both the

heteroskedasticity-robust standard errors (t-statistic) and bootstrapped standard errors (z-statistic).

I highlight these two examples as they help link the observed outcomes to the model. First, these two events are arguably two of the more unanticipated outcomes listed, meaning the estimates are more likely to capture the full impact of the event whereas the other events being more anticipated mean that the estimates are likely lower bounds on the estimated event impacts. This corresponds with the fact that these events had the largest estimated impacts on cumulative returns. Second, these events also saw dynamic effects that played out for up to four weeks that were statistically significant. While the use of the climate policy risk exposure measure based on oil prices and production is one way that helps link these outcomes to my specific model, the dynamic responses, corresponding to the dynamic responses related to production and pricing impacts my model predicts, provide further formal evidence consistent with my model mechanism in the direction of the impact and dynamics.

8.3 Climate Policy Events Index

I now extend the empirical analysis to the time series of climate policy related shocks. To identify the impact of climate policy shocks to oil production and oil sector returns, I first estimate reduced-form regressions focused on the link between changes in the probability of climate policy occurring, as measured by climate- and climate-policy-related events, and oil production decisions, oil sector returns for US firms, and oil spot price returns. The climate policy shocks measure is labeled as *ClimPol*, the index variable tracking different climate related events discussed previously. For example, events in *ClimPol* include the establishment of the Paris Accord in 2015, as well as the election of Donald Trump as the President of the United States in 2016. The Paris Accord is considered a positive shock to the arrival rate of a significant climate policy action and so a positive one in the index and the election of Trump is considered a negative shock and so a negative one in the index. I identify these events at a daily level, and then aggregate them up to monthly values for my analysis.

The goal of this exercise is to identify whether events related to changes in the likelihood of future climate policy action lead to changes in production and prices consistent with the model predictions. My model predicts that a positive shock to the arrival rate should cause an increase in oil production and negative oil firm returns and oil price returns, whereas a negative shock to the arrival rate should lead to a decrease in oil production and increase in oil firm returns and oil price returns.

The reduced-form regression approach provides a estimate of how climate policy driven demand shocks influence economic and financial outcomes. To determine whether the empirical outcomes from this simple analysis are consistent with the model, I focus on the signs and statistical significance of the estimates, and compare those with the qualitative results of the model. Furthermore, I will estimate each regression on the full time sample (1973-2017) and on a shorter, more recent subsample (1996-2017). The recent sub-sample I refer to as the policy-relevant sample. I choose 1996 as the starting year for this sample as it is near the time when major climate policy begins to take place, such as the Kyoto Protocol, which was an early global climate agreement similar to the recent Paris

Climate Accord. The model would predict that impacts estimated in the policy relevant subsample should be higher as temperature has increased and the likelihood of climate policy occurring is higher.

8.3.1 Oil Production

Table 3: Climate Policy Impact on Oil Production (1973-2017)

	OPEC	US	Non-OPEC	World
ClimPol	-0.022	0.026	0.099	0.069
S.E.	(0.034)	(0.063)	(0.012)	(0.072)
# Obs.	536	536	536	536
R^2	0.995	0.979	0.991	0.988

Table 4: Climate Policy Impact on Oil Production (1996-2017)

	OPEC	US	Non-OPEC	World
ClimPol	0.035	0.041	0.141	0.176
S.E.	(0.055)	(0.015)	(0.047)	(0.074)
# Obs.	260	260	260	260
R^2	0.980	0.972	0.988	0.987

These tables show the impact of climate policy events as measured by the *ClimPol* index on oil production for the Non-OPEC, OPEC, US, and World regions. The top table are estimates using the full time sample of data (1973-2017), and the bottom table are estimates using the policy-relevant time subsample (1996-2017). The regression specification is given by $Y_t = \alpha + \beta_Y Y_{t-1} + \phi_{ClimPol} ClimPol_t + \epsilon_t$. I omit the constant and lag variable coefficients from the table. See text for full definition of variables.

I begin by focusing on the impact of climate policy on oil production. To estimate the effect of climate policy shocks, I estimate the following regression:

$$Y_t = \alpha + \beta_Y Y_{t-1} + \phi_{ClimPol} ClimPol_t + \epsilon_t$$

where Y_t is crude oil production, Y_{t-1} is the one-period lag of crude oil production, and *ClimPol* is the index for the climate events mentioned above. I exploit the time series and panel dimensions of the data by estimating this regression using information from the changes in the dependent variable across time for different regions of interest. In Tables 3 and 4 show the results for oil production

across four different regions using the full time sample (1973-2017), and the results for oil production across four different regions using the more recent, policy-focused time sample (1996-2017).

Controlling for the lag value, meant to capture relevant market conditions and current market effects or trends, oil production increases for the US, non-OPEC countries, and globally for an event that increases the likelihood of climate related policy. The effect is larger and more statistically significant in the recent policy-focused time period estimates. Climate policy events have essentially no impact on the production of the OPEC-countries region, though the estimate goes from negative to positive when comparing the full time sample to the policy-relevant time sample. These results are in line with two predictions of the model. First, they are consistent with the result of a run on oil occurring for an increased likelihood in climate policy occurring, at least for all regions that are not exclusively the OPEC region. Second, the increased magnitude for the most recent time period is consistent with the prediction that an increased likelihood of climate policy, tied to higher temperatures, should generate larger impacts on production.

8.3.2 Oil Sector and Oil Price Returns

Next, I test the model implication for the impact of climate policy on oil sector returns and oil price returns. To do this, I estimate whether shocks to climate policy predict negative changes to oil sector returns by estimating the following regression:

$$r_{i,t+1,t+h} = a_i + b_i X_t + c_i ClimPol_t + \varepsilon_{i,t}$$

where $r_{i,t+1,t+h}$ are 1-, 6-, 12-, 18-, and 24-month ahead cumulative returns, i is for cumulative excess returns for the oil sector portfolio or cumulative returns for the WTI oil price, $ClimPol_t$ is the climate policy dummy, and X_t includes non-contemporaneous controls for the market portfolio, economic productivity, spot price and sector returns, oil production innovations, and log OECD industrial production innovations. Tables 5 through 8 show the results across the five different cumulative return scenarios using the full time sample (1973-2017), and the results across the five different cumulative return scenarios using the more recent, policy-focused time sample (1996-2017).

After including controls to capture relevant market and macroeconomic conditions and trends, I find that shocks to climate policy have a negative impact on oil sector and oil price returns, as can be seen across the different horizons and time samples used for estimation. The coefficients are negative or insignificant for all horizons of cumulative returns. The effect is more negative and is statistically significant for the estimates based on the more recent time period sample at the 18- and 24-month horizons. The magnitude of the impact and predictability are also increasing with the horizon of the cumulative returns, as the coefficients become more negative and the R^2 's become larger. These results are in line with three more predictions of the model. First, the results are consistent with

Table 5: Climate Policy Impact on Oil Sector Returns (1973-2017)

	$r_{oil,t+1}$	$r_{oil,t+6}$	$r_{oil,t+12}$	$r_{oil,t+18}$	$r_{oil,t+24}$
ClimPol	0.000	0.005	-0.009	-0.035	-0.027
S.E.	(0.005)	(0.013)	(0.020)	(0.026)	(0.036)
# Obs.	534	529	523	517	511
R^2	0.004	0.016	0.014	0.028	0.013

Table 6: Climate Policy Impact on Oil Sector Returns (1996-2017)

	$r_{oil,t+1}$	$r_{oil,t+6}$	$r_{oil,t+12}$	$r_{oil,t+18}$	$r_{oil,t+24}$
ClimPol	-0.002	-0.002	-0.034	-0.072	-0.085
S.E.	(0.007)	(0.015)	(0.024)	(0.032)	(0.038)
# Obs.	260	260	260	260	260
R^2	0.021	0.017	0.016	0.045	0.041

These tables show the impact of climate policy events as measured by the *ClimPol* index on returns for the value-weighted US Oil sector portfolio. The regression specification is given by $r_{i,t+1,t+h} = a_i + b_i X_t + c_i \text{ClimPol}_t + \varepsilon_{i,t}$. $r_{i,t+1,t+h}$ is the k-month cumulative return for the value-weighted US Oil sector portfolio. X_t is a vector of control variables that includes the lagged values for the value-weighted market portfolio return, the value-weighted oil sector portfolio return, lagged oil production innovations, lagged oil spot price returns, and lagged real economic activity innovations. The top panel are estimates using the full time sample of data (1973-2017), and the bottom panel are estimates using the policy-relevant time subsample (1996-2017). I omit the constant and control coefficients from the table. See text for full definition of variables.

the prediction that shocks to climate policy that lead to an oil run also depress oil sector firm values and oil prices. Second, the impact has a dynamic effect on outcomes as the negative returns persist and increase in magnitude over the longer cumulative return horizons explored. Finally, as was seen with oil production, the increased magnitude over the most recent time period as compared to the full time period estimates is consistent with an increased impact of climate policy as temperature increases and the likelihood of significant climate policy occurring increases.

8.3.3 Climate Policy*Temperature Interaction Estimates

To strengthen the validity of the climate policy index analysis and further connect the estimated results to the model, I augment these regressions by using a climate policy index and temperature interaction term. The model specifies that the likelihood of climate policy is tied to increases in temperature, and therefore increases in temperature should amplify the impact of climate policy risk.

Table 7: Climate Policy Impact on Oil Price Returns (1973-2017)

	$r_{oil,t+1}$	$r_{oil,t+6}$	$r_{oil,t+12}$	$r_{oil,t+18}$	$r_{oil,t+24}$
ClimPol	-0.002	0.042	0.050	0.018	0.017
S.E.	(0.008)	(0.026)	(0.042)	(0.045)	(0.047)
# Obs.	534	529	523	517	511
R^2	0.045	0.021	0.022	0.016	0.021

Table 8: Climate Policy Impact on Oil Price Returns (1996-2017)

	$r_{oil,t+1}$	$r_{oil,t+6}$	$r_{oil,t+12}$	$r_{oil,t+18}$	$r_{oil,t+24}$
ClimPol	-0.004	0.021	-0.001	-0.076	-0.099
S.E.	(0.011)	(0.033)	(0.045)	(0.048)	(0.050)
# Obs.	260	260	260	260	260
R^2	0.013	0.020	0.029	0.043	0.047

These tables show the impact of climate policy events as measured by the *ClimPol* index on returns for the WTI spot price of oil. The regression specification is given by $r_{i,t+1,t+h} = a_i + b_i \text{ClimPol}_t + c_i X_t + \varepsilon_{i,t}$. $r_{i,t+1,t+h}$ is the k-month cumulative return for the WTI spot price of oil. X_t is a vector of control variables that includes the lagged values for the value-weighted market portfolio return, the value-weighted oil sector portfolio return, lagged oil production innovations, lagged oil spot price returns, and lagged real economic activity innovations. The top panel are estimates using the full time sample of data (1973-2017), and the bottom panel are estimates using the policy-relevant time subsample (1996-2017). I omit the constant and control coefficients from the table. See text for full definition of variables.

The use of the climate policy index interacted with temperature directly tests this link, while also still testing the impact that policy and climate have on oil production, oil sector returns, and oil price returns. Though I have previously proposed that the increased effects seen when comparing the subsample of recent, policy relevant data are related to increases in temperature and increases in policy concern, these interaction estimates should confirm whether or not this is the case.

I continue to include the same lag values and controls as before in the regression equations. The key difference is that the dependent variable of interest is now an interaction variable of the climate policy index and the one year moving average of global mean temperature, $\text{ClimPol} * \text{Temp}$. More precisely, the regression specifications for the production, oil return, and oil price returns estimates are respectively given by

$$Y_t = \alpha + \beta_Y Y_{t-1} + \phi_{\text{ClimPol} * \text{Temp}} \text{ClimPol}_t * \text{Temp}_t + \varepsilon_t$$

$$r_{t+1,t+h}^e = a + b X_t + c_{\text{ClimPol} * \text{Temp}} \text{ClimPol}_t * \text{Temp}_t + \varepsilon_t$$

$$r_{t+1,t+h}^{spot} = a + bX_t + c_{ClimPol*Temp}ClimPol_t * Temp_t + \varepsilon_t$$

I leave the tables with the estimated coefficients for the appendix and outline briefly here the estimation results. For oil production, an increase in the interaction term leads generally to an increase in oil production as it did before. Also similar to before, the impacts are increasing in magnitude and significance for the more recent subsample of data. However, the interaction term for climate policy and temperature is now positive and statistically significant for the US, Non-OPEC, and World regions in the 1973-2017 sample of date. Thus we see an enhanced effect by accounting for temperature within the impact of the climate policy risk, which is in line with the models prediction of increased effects from increased temperature and the proposed justification for increased effects seen in the more recent, policy-relevant and higher temperature subsample of data.

The estimated effects of the interaction term for oil sector and oil price returns also line up with the previous results. An increase to the interaction term has a negative impact on oil sector and oil price returns. The effect is more negative and is statistically significant for the estimates based on the more recent time period sample at the 18- and 24-month horizons, and the magnitude of the impact and predictability are also increasing with the horizon of the cumulative returns. However, again with these estimates there are key main differences from the previous results. First, the estimated impact of the interaction term on the oil sector returns is now monotonically increasing in magnitude and statistical significance, and are statistically significant for the 18- and 24-month cumulative return horizons for the 1973-2017 data sample estimates. And though the impact on oil price returns is not statistically significant for 1973-2017 data sample estimates, the impacts are now all negative and monotonically increasing in magnitude and significance for longer horizons. This results again validates that there is an enhanced effect by accounting for temperature within the impact of the climate policy risk as the model implies, as now even the full sample of data estimates are significant, and further confirms the impacts on oil production and prices estimated in the previous regressions that are consistent with the model implications for climate policy risk.

8.4 Vector Autoregression Analysis

To further the empirical estimation of the dynamic effects of uncertain climate policy on oil sector quantity and price outcomes, I estimate a structural vector autoregression (VAR) for the global oil market. Augmenting the global oil market VARs proposed and used by Kilian and Park (2009), Baumeister and Hamilton (2017), and others, I estimate:

$$y_t = \nu + \sum_j A_j y_{t-j} + u_t$$

where the vector of endogenous state variable vector y_t is defined by

$$y_t = [ClimPol_t, \Delta prod_t, rea_t, \Delta p_t^{oil}]'$$

ClimPol is the climate policy index measure I mentioned previously. $\Delta prod$ is the percent change in global oil production available from the EIA. REA is a measure of real economic activity given by innovations in the log OECD industrial production index suggested in recent work by James Hamilton. r_t^{mkt} is log differences in the real West Texas Intermediate (WTI) monthly closing price for crude oil.

I use a Cholesky decomposition of the estimated variance-covariance matrix for identification of the structural shocks. This identification strategy imposes a recursive interpretation of the impact of the shocks. The general representation and interpretation of this identification is as follows:

$$u_t = B \left[\epsilon_{\text{climate policy}}, \quad \epsilon_{\text{oil supply}}, \quad \epsilon_{\text{aggregate demand}}, \quad \epsilon_{\text{oil-specific demand}} \right]'$$

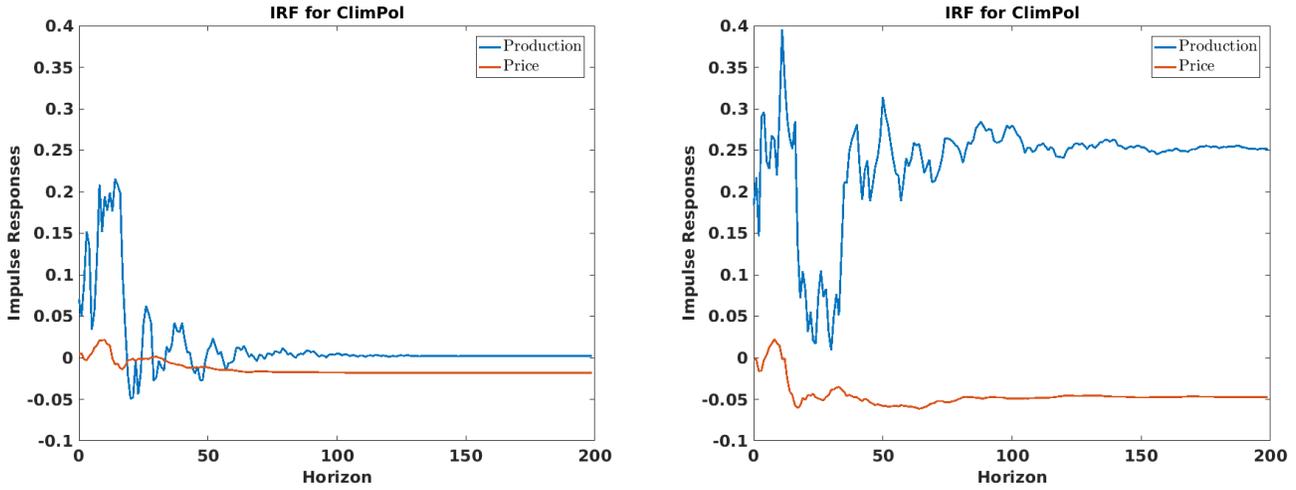
where B is the lower triangular matrix derived from the Cholesky decomposition of the estimated variance covariance matrix $\hat{\Sigma}$, i.e., $BB' = E_t[u_t' u_t] = \hat{\Sigma}$. I outline the specific interpretation and identification of each shock in what follows.

*ClimPol*_t, the focus of this exercise, captures changes in the likelihood of future climate policy that restricts oil use, that is changes in $\lambda(T_t)$ from the model. Although long-run temperature directly maps to the likelihood of significant climate policy action in the model, in practice this link is less precise. Figure 20 in the appendix, which shows the US temperature anomaly time series over the annual *ClimPol* index measure, demonstrates this relationship. The time series for the two variables are positively correlated, but the correlation is obviously not one. For this reason, I use the more direct measure of *ClimPol* to capture changes in the likelihood of future climate policy that restricts the use of oil.

This ordering assumes the likelihood of significant climate policy is contemporaneously predetermined with respect to oil sector shocks and the oil sector is contemporaneously influenced by shifts to the likelihood of future climate policy. This assumption is intuitive and maintains consistency with the model in that the likelihood of significant climate policy responds only with a lag to oil sector shocks as a result of emissions from oil production impacting the climate policy arrival rate. Recent climate science work by Matthews et al. (2009), Ricke and Caldeira (2014), and Zickfeld and Herrington (2015) has shown that impacts on temperature from carbon emissions can take many years or even decades to fully realize, which further validates this restriction. The order for the remaining variables follows the setting of Kilian (2009). Thus, this interpretation of the structure fits this setting as well: 1.) a vertical short-run supply curve and downward sloping demand curve; and 2.) oil demand and supply shocks imply immediate changes in the real oil price.

To further highlight the consistency of the VAR model with the theoretical model, consider the following. The supply shocks and real aggregate demand shocks can be interpreted as the shocks to oil reserves and shocks to capital in the model. As in the model, supply and policy shocks are important determinants of oil supply or production, along with temperature shocks, which are correlated with climate policy likelihood. Combined with capital shocks, and the preferences of agents, these shocks then determine oil prices. These effects together then pin down asset prices in the model. The link

Figure 10: ClimPol Shock IRF - 1973-2017 vs. 1996-2017 Time Samples

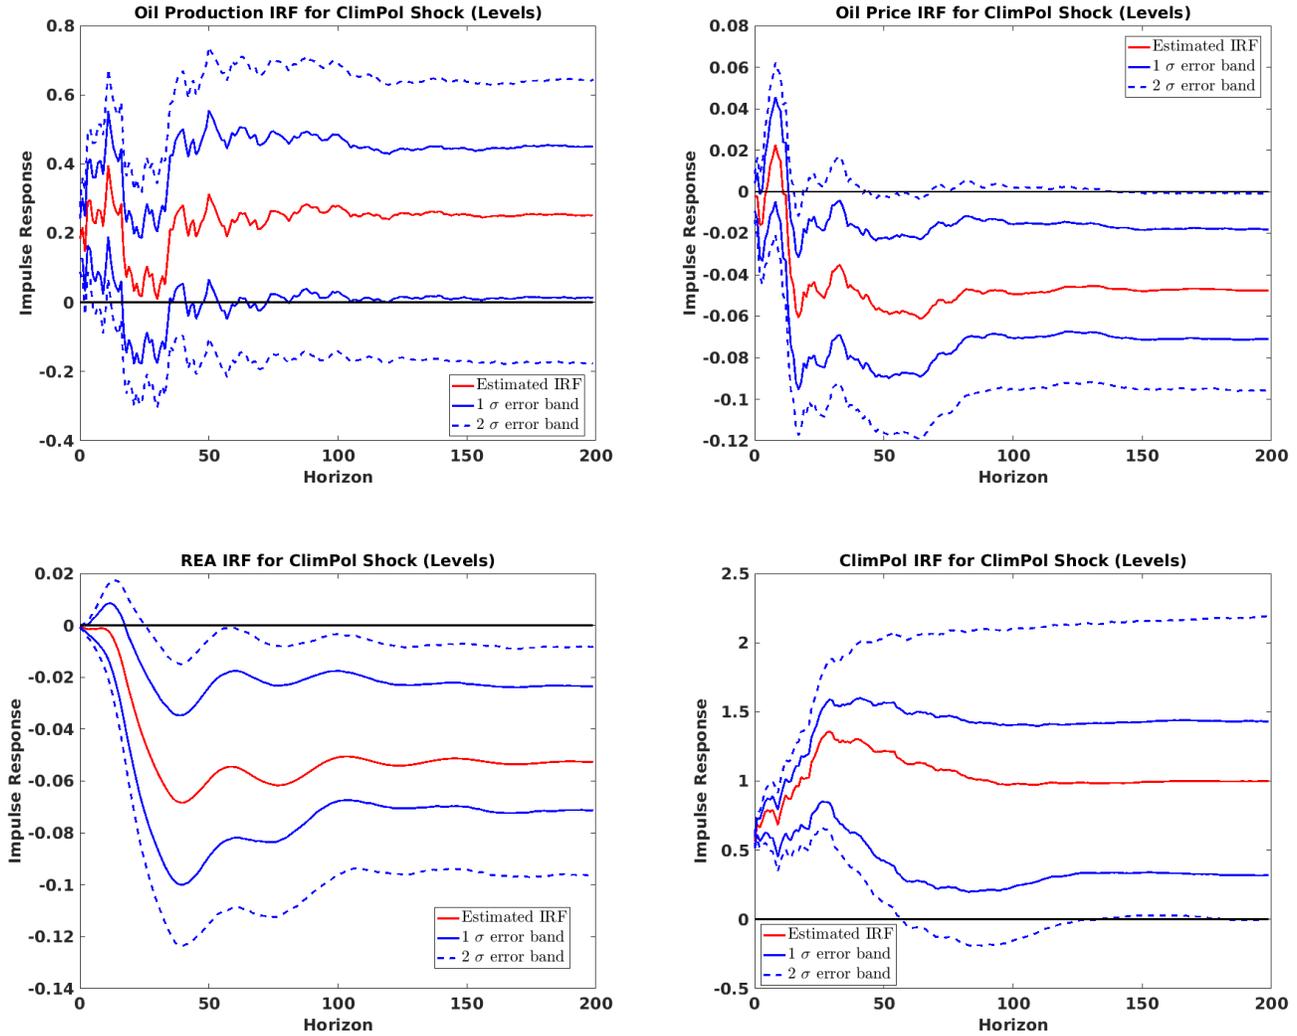


These figures show estimated impulse response functions for global oil production and the WTI spot price of oil for a shock to the *ClimPol* index. The left panel are estimates using the full time sample of data (1973-2017), and the right panel are estimates using the policy-relevant time subsample (1996-2017). See text for the full VAR specification used and definition of variables. The VAR is estimated using 24 lags.

between the climate and the economy in the model comes through how emissions from oil produced impact temperature and then how temperature influences climate damages and the climate policy likelihood, which feed into the determination of the economic and financial outcomes of interest. Thus, the variables included in the VAR and the ordering of the variables in the recursive decomposition is consistent with the theoretical model framework.

From the VAR estimates and the recursive identification structure, I derive impulse response functions (IRFs), or the cumulative responses to a given structural variable shock, which are the results I use to examine the validity of the model mechanism. To understand how the IRFs generated from the VAR estimation can help validate the model, consider first the expected IRFs from alternative model settings. These alternative settings include the case of no uncertain policy arrival or a temperature-independent, constant policy arrival rate of changes in the energy input demand share. In the setting without any uncertain policy, a shock to the climate policy variable would lead to a decrease in oil production and an increase in the spot price of oil if the event increased the carbon tax. There would be no change in outcomes if the event did not lead to a change in the carbon tax. In the uncertain policy setting where the arrival rate is climate-independent and constant, a shock to the climate policy variable should lead to an increase in oil production and decrease in the spot price of oil, but this effect would not persist because of the lack of temperature dependence and the prevailing Hotelling-type forces. However, in the uncertain climate policy setting with a climate-dependent arrival rate, a shock to the likelihood of significant climate policy occurring leads to an increase in oil production and a decrease in the oil spot price. Furthermore, the impacts of a shock to the likelihood of significant

Figure 11: ClimPol Shock IRF - 1996-2017 Time Sample w/ C.I.s



These figures show the estimated impulse response functions for global oil production, the WTI spot price of oil, real economic activity, and the *ClimPol* climate policy index measure for a shock to the likelihood of climate policy measured by the *ClimPol* index. The red line is the estimated IRF, the solid blue lines represent the on-standard deviation error bands, and the blue dashed lines represent the two standard deviation error bands. Error bands are estimated using bootstrapping. The estimates use the policy-relevant time subsample (1996-2017). See text for the full VAR specification used and definition of variables. The VAR is estimated using 24 lags.

climate policy occurring should produce impacts that are persistent and potentially increasing in magnitude dynamically for these outcomes, two defining features of what I have termed a run on oil. Therefore, one can “test” the validity of the model proposed in this paper by looking at the sign and dynamics of the IRFs for oil production and oil spot prices resulting from a shock to the likelihood of significant climate policy occurring.

Figure 10 shows the cumulative level impulse response functions of oil production and oil prices for a shock to the likelihood of significant climate policy occurring. The left plot is for the VAR estimated using the full time sample (1973-2017) and the right plot for the VAR estimated using the more recent, policy-focused time sample (1996-2017). Figure 11 shows the individual impulse response functions with bootstrapped standard errors for the policy focused time sample (1996-2017), where the solid blue lines are for the one-standard deviation confidence interval, the dashed blue lines are for the two-standard deviation confidence interval, and the red line is for the estimated IRF.

The impulse response functions further confirm the results seen previously in the reduced-form estimates. For the impulse response functions generated from the full time sample VAR estimates, the responses are substantially muted and quite close to zero. However, focusing on the impulse responses generated from the VAR estimates using the policy-focused time sample, we see results that correspond to the model predictions. A shock to climate policy leads to an increase in oil production and a decrease in the spot price of oil. Moreover, these impacts grow in magnitude and persist over time. Though only the impact on the spot price of oil is statistically significant for the two variables of interest, as seen in Figure 11, the direction and dynamics of these results taken together are consistent with the key predictions of the model and the other empirical exercises.

8.5 Return-Weighted Climate Policy Index

Finally, I provide an extension of the VAR analysis which uses return-weighted climate policy index measures. The purpose of this extension is to account for the magnitude and dynamics of the effects of the climate policy events on the oil price and production of oil by incorporating the forward-looking information of relevant asset price returns into my climate policy index. This extension ties the asset pricing and production implications into a single analysis to provide a more complete test of the model implications. The results of this extension are not only consistent with the results given above, but in fact identify larger and more significant impacts by exploiting the informational value of asset prices and accounting for policy magnitudes and dynamic implications.

A natural asset price measure of the magnitude of the impact of climate policy events follows from the event study analysis done previously. Each event in the event study analysis required estimating a climate policy risk exposure measure for the cross-section of sector portfolios. These climate policy risk exposure measures provide a weighting scheme for a type of factor-mimicking portfolio that captures not only the realization of climate policy events, but the also the magnitude of the impacts of these events. Furthermore, this method will also capture the dynamic implications of the events

through the return response as we saw with the event study analysis. Thus this extension should allow me to determine not only the effect of a climate policy event realization, but also the magnitude and dynamics of the effects of the climate policy events on the oil price and production of oil.

I implement this method as follows. First, I focus on a particular event to estimate the climate policy risk exposures from exposure to oil price innovations for each sector. Given these estimate exposure measure, I rescale the exposures to sum to one to provide a portfolio weighting scheme. These event-study estimated weights are then used to calculate a factor-mimicking portfolio from the sector portfolios. I focus on the equal weighted portfolios for simplicity. With this factor-mimicking portfolio, I can implement the analysis in two ways. The first is to incorporate the factor mimicking portfolio returns directly into the VAR in place of the climate policy event index as a continuous measure of responses to climate policy events. The second is to use an interaction term of the climate policy event index dummy used originally with the factor-mimicking portfolio that I have constructed. The interaction term will highlight directly the events as given by the climate policy event index, while also incorporating magnitudes through the return measure. This analysis ties the asset pricing and production implications into a single analysis to test the full model implications. I will first focus on the weights generated from the Supreme Court hold on the Clean Power Plan, and for robustness will test portfolio weightings based on different dynamic response times and different events.

Figure 12 shows the cumulative level impulse response functions of oil production and oil prices for a shock to the likelihood of significant climate policy occurring using these alternative climate policy indices. All the plots are for the VAR estimated using the more recent, policy-focused time sample (1996-2017). Each plot include includes the the estimated IRF (red line), the one-standard deviation bootstrapped confidence interval (solid blue line), and the two-standard deviation bootstrapped confidence interval (dashed blue line).

The top two plots show the results for the factor mimicking portfolio returns used as the climate policy risk likelihood measure. Here the results are quite similar to the original ClimPol IRFs. A shock to climate policy leads to an increase in oil production and a decrease in the spot price of oil. Moreover, these impacts grow in magnitude and persist over time. As before, the direction and dynamics of these results are consistent with the key predictions of the model. While again only the impact on the spot price of oil is statistically significant, the significance is even greater in this setting.

The bottom two plots show the results for the factor mimicking portfolio returns interacted with the original ClimPol index used as the climate policy risk likelihood measure. Again, a shock to climate policy leads to an increase in oil production and a decrease in the spot price of oil, and these impacts grow in magnitude and persist over time. There are a number of key differences in this setting. First, the magnitude of the IRFs is greater for the impact on both oil prices and oil production. Second, the impact on the spot price of oil and oil production is statistically significant. This setting shows that not accounting for magnitude likely understated the measured impact from before, as well as show the value of using asset prices to study the impact of climate policy risk.

Figure 12: ClimPol Shock IRF - 1996-2017 Time Sample w/ C.I.s

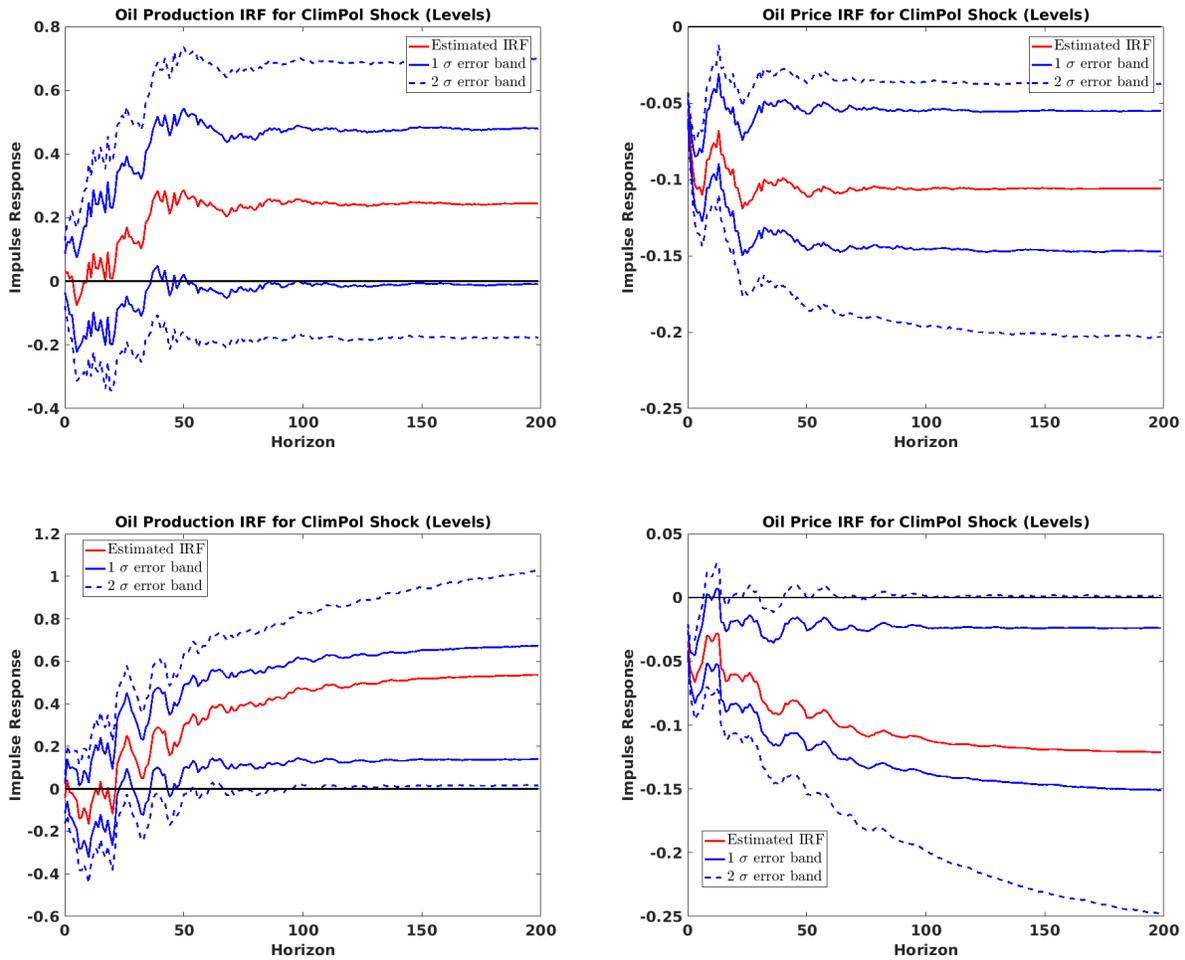


Figure 12 shows the estimated impulse response functions for global oil production, the WTI spot price of oil, real economic activity, and variations of the *ClimPol* climate policy index measure for a shock to the likelihood of climate policy. The first variation uses the returns of the factor mimicking portfolio created from sector portfolios weighted by their normalized climate policy risk exposure estimated value. The second variation uses the returns for this same factor mimicking portfolio, but is interacted with the original *ClimPol* index. The red line is the estimated IRF, the solid blue lines represent the on-standard deviation error bands, and the blue dashed lines represent the two standard deviation error bands. Error bands are estimated using bootstrapping. The estimates use the policy-relevant time subsample (1996-2017). See text for the full VAR specification used and definition of variables. The VAR is estimated using 24 lags.

9 Conclusions

I study the impacts of uncertain climate policy on financial and economic outcomes using a general equilibrium, production-based asset pricing model. Introducing uncertain climate policy causes a model with oil extraction that otherwise closely follows standard Hotelling model-type outcomes to generate a run on oil, meaning firms increase oil extraction and exploration more and more as oil reserves go down and the atmospheric temperature increases. Uncertain climate policy alters the subjective discount rate to be dependent on climate change. Thus, concerns about future climate policy lead to a run on oil production as future profits from oil are increasingly discounted as the likelihood that oil reserve assets become stranded increases. The run on oil induced by uncertain climate policy dynamically pushes down oil firm values and oil spot prices as agents incorporate the concern of stranded assets into their discounting, while also generating a significant shift down in the level of oil firm values. Not accounting for uncertain climate policy in oil firm values would lead to a “carbon bubble,” where oil firm values are overpriced because the price does not incorporate the risk of oil reserves becoming stranded and the run on oil production that stranded assets risk causes.

In addition, I find empirical evidence in support of the theoretical predictions of the model. First, I estimate an event-study analysis around the 2016 US presidential election and find that sectors with large positive exposure to climate policy risk, measured by the model-suggested proxy of exposure to oil price innovations, saw increased cumulative abnormal returns one day after the election. This effect increased up to four weeks after the election, following the predicted dynamic implications of the model. In reduced-form regressions and structural VAR estimates, climate policy related shocks lead to dynamic and persistent increases in oil production and decreases in oil firm values and oil spot prices. The reduced form estimates show the increase in oil production from a climate policy shock is statistically significant for the US and other non-OPEC regions, whereas it is not statistically significant for OPEC. The impacts for both sets of estimates are greater in magnitude when estimated using only the more recent, policy-focused time period sample than the estimates using the entire available time period. This increase in policy impacts with increasing temperature and climate policy concern supports the novel mechanism of my model.

There are a number of interesting areas of research related to my paper that I leave for future work. A quantitative empirical analysis of the impacts of uncertain climate policy on macroeconomic outcomes and asset prices is a valuable extension that would help determine the perceived significance of uncertain climate policy risk. Furthermore, estimating the parameters for the arrival rate of uncertain climate policy would further demonstrate market expectations of the likelihood of a significant climate policy shock and is something which asset prices may potentially help to identify. Another important area to consider empirically is using derivatives such as oil options, oil futures prices, and the term structures of these financial instruments to estimate the long-run expectations about uncertain climate policy and climate change. My model, and extensions of it, could prove valuable in providing testable predictions for these types of assets. Apart from this theoretical extension of alternative,

long-run assets, other important theoretical extensions that can be explored in future work include alternative policy proposals and a multi-country framework of oil production and climate policy with imperfect competition. Each of these are important dimensions building off of this paper that can help us understand the impacts of climate change and uncertain climate policy.

References

- Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr (2016). Transition to clean technology. *Journal of Political Economy* 124(1), 52–104.
- Anderson, E. W. and W. A. Brock (2017). Logarithmic depreciation. *SSRN Working Paper* 1(1), 1.
- Anderson, E. W., W. A. Brock, and A. Sanstad (2016). Robust consumption and energy decisions.
- Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003). A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association* 1(1), 68–123.
- Bansal, R., M. Ochoa, and D. Kiku (2016). Climate change and growth risks. Technical report, National Bureau of Economic Research.
- Barnett, M. (2017). Climate change and uncertainty: An asset pricing perspective. *Working Paper*.
- Barnett, M., W. Brock, and L. P. Hansen (2018). Pricing uncertainty induced by climate change. *Working Paper. In-principle acceptance of RFS Climate Finance registered report, completion in progress..*
- Baumeister, C. J. and J. D. Hamilton (2017). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. Technical report, National Bureau of Economic Research.
- Belo, F., V. D. Gala, and J. Li (2013). Government spending, political cycles, and the cross section of stock returns. *Journal of Financial Economics* 107(2), 305–324.
- Bornstein, G., P. Krusell, and S. Rebelo (2017). Lags, costs, and shocks: An equilibrium model of the oil industry. Technical report, National Bureau of Economic Research.
- Brock, W. A. (1982). Asset prices in a production economy. In *The economics of information and uncertainty*, pp. 1–46. University of Chicago Press.
- Brock, W. A. and L. P. Hansen (2017). Wrestling with uncertainty in climate economic models.
- Cai, Y., K. L. Judd, and T. S. Lontzek (2015). The Social Cost of Carbon with Economic and Climate Risks. *Working Paper*.
- Carlson, M., Z. Khokher, and S. Titman (2007). Equilibrium exhaustible resource price dynamics. *The Journal of Finance* 62(4), 1663–1703.
- Casassus, J., P. Collin-Dufresne, and B. R. Routledge (2009). Equilibrium commodity prices with irreversible investment and non-linear technologies.
- Cochrane, J. H. (1991). Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations. *The Journal of Finance* 46(1), 209.
- Crost, B. and C. P. Traeger (2011). Risk and aversion in the integrated assessment of climate change.
- Dasgupta, P. and G. Heal (1974). The optimal depletion of exhaustible resources. *The review of economic studies* 41, 3–28.

- David, A. (2015). Exploration activity, long run decisions, and the risk premium in energy futures.
- Dell, M., B. F. Jones, and B. A. Olken (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics* 4(3), 66–95.
- Deschenes, O. and M. Greenstone (2007). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather. *The American Economic Review* 97(1), 354–385.
- Dietz, S., C. Gollier, and L. Kessler (2017). The climate beta. *Journal of Environmental Economics and Management*.
- Duffie, D. and C. Skiadas (1994). Continuous-time security pricing: A utility gradient approach. *Journal of Mathematical Economics* 23(2), 107–131.
- Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014). Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica* 82(1), 41–88.
- Gomes, J., L. Kogan, and L. Zhang (2003). Equilibrium Cross Section of Returns. *Journal of Political Economy* 111(4), 693–732.
- Gomes, J. F., L. Kogan, and M. Yogo (2009). Durability of output and expected stock returns. *Journal of Political Economy* 117(5), 941–986.
- Hambel, C., H. Kraft, and E. Schwartz (2015). Optimal carbon abatement in a stochastic equilibrium model with climate change. Technical report, National Bureau of Economic Research.
- Hamilton, J. D. (2005). Oil and the macroeconomy. *The New Palgrave Dictionary of Economics* 2.
- Hamilton, J. D. (2008). Understanding crude oil prices. Technical report, National Bureau of Economic Research.
- Hansen, L. P. and J. Miao (2018). Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proceedings of the National Academy of Sciences* 115(37), 9163–9168.
- Hansen, L. P. and T. J. Sargent (2011). Robustness and ambiguity in continuous time. *Journal of Economic Theory* 146(3), 1195–1223.
- Hansen, L. P., T. J. Sargent, G. Turmuhambetova, and N. Williams (2006). Robust control and model misspecification. *Journal of Economic Theory* 128(1), 45–90.
- Hong, H., F. W. Li, and J. Xu (2016). Climate risks and market efficiency. Technical report, National Bureau of Economic Research.
- Hotelling, H. (1931). The economics of exhaustible resources. *Journal of political Economy* 39(2), 137–175.
- Hsiang, S., R. Kopp, A. Jina, J. Rising, M. Delgado, S. Mohan, D. Rasmussen, R. Muir-Wood, P. Wilson, M. Oppenheimer, et al. (2017). Estimating economic damage from climate change in the united states. *Science* 356(6345), 1362–1369.

- Jermann, U. J. (1998). Asset pricing in production economies. *Journal of Monetary Economics* 41(2), 257–275.
- Kelly, B., L. Pástor, and P. Veronesi (2016). The price of political uncertainty: Theory and evidence from the option market. *The Journal of Finance*.
- Kelly, D. L. and C. D. Kolstad (1999). Bayesian learning, growth, and pollution. *Journal of economic dynamics and control* 23(4), 491–518.
- Kilian, L. (2008). Exogenous oil supply shocks: how big are they and how much do they matter for the us economy? *The Review of Economics and Statistics* 90(2), 216–240.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99(3), 1053–69.
- Kilian, L. and C. Park (2009). The impact of oil price shocks on the us stock market. *International Economic Review* 50(4), 1267–1287.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica* 73(6), 1849–1892.
- Kogan, L., D. Livdan, and A. Yaron (2009). Oil futures prices in a production economy with investment constraints. *The Journal of Finance* 64(3), 1345–1375.
- Kogan, L. and D. Papanikolaou (2014). Growth opportunities, technology shocks, and asset prices. *The Journal of Finance* 69(2), 675–718.
- Koijen, R. S., T. J. Philipson, and H. Uhlig (2016). Financial health economics. *Econometrica* 84(1), 195–242.
- Kotlikoff, L. J., A. Polbin, and A. Zubarev (2016). Will the paris accord accelerate climate change? Technical report, National Bureau of Economic Research.
- Kushner, H. J. and P. Dupuis (2001). *Numerical Methods for Stochastic Control Problems in Continuous Time*, Volume 24. Springer Science & Business Media.
- Lemoine, D. M. and C. P. Traeger (2012). Tipping points and ambiguity in the economics of climate change. Technical report, National Bureau of Economic Research.
- Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 1447–1498.
- MacDougall, A. H. and P. Friedlingstein (2015). The origin and limits of the near proportionality between climate warming and cumulative co2 emissions. *Journal of Climate* 28(10), 4217–4230.
- Matthews, H. D., N. P. Gillett, P. A. Stott, and K. Zickfeld (2009). The proportionality of global warming to cumulative carbon emissions. *Nature* 459(7248), 829–832.
- Matthews, H. D., S. Solomon, and R. Pierrehumbert (2012). Cumulative carbon as a policy framework for achieving climate stabilization. *Phil. Trans. R. Soc. A* 370(1974), 4365–4379.

- McGlade, C. and P. Ekins (2015). The geographical distribution of fossil fuels unused when limiting global warming to 2 degrees c. *Nature* 517(7533), 187–190.
- Nordhaus, W. (2014). *A question of balance: Weighing the options on global warming policies*. Yale University Press.
- Papanikolaou, D. (2011). Investment shocks and asset prices. *Journal of Political Economy* 119(4), 639–685.
- Pástor, L. and P. Veronesi (2009). Technological revolutions and stock prices. *American Economic Review* 99(4), 1451–83.
- Pastor, L. and P. Veronesi (2012). Uncertainty about government policy and stock prices. *The Journal of Finance* 67(4), 1219–1264.
- Pindyck, R. S. and N. Wang (2013). The Economic and Policy Consequences of Catastrophes. *American Economic Journal: Economic Policy* 5(4), 306–339.
- Ready, R. (2015). Oil Consumption, Economic Growth, and Oil Futures: A Fundamental Alternative to Financialization. *Working Paper*.
- Ricke, K. L. and K. Caldeira (2014). Maximum warming occurs about one decade after a carbon dioxide emission. *Environmental Research Letters* 9(12), 124002.
- Salant, S. W. and D. W. Henderson (1978). Market anticipations of government policies and the price of gold. *Journal of political economy* 86(4), 627–648.
- Santa-Clara, P. and R. Valkanov (2003). The presidential puzzle: Political cycles and the stock market. *The Journal of Finance* 58(5), 1841–1872.
- Sialm, C. (2006). Stochastic taxation and asset pricing in dynamic general equilibrium. *Journal of Economic Dynamics and Control* 30(3), 511–540.
- Sinn, H.-W. (2007). Public policies against global warming. Technical report, National Bureau of Economic Research.
- Stern, N. (2007). *The Stern Review: The economics of climate change*. Cambridge University Press.
- Tourre, F. et al. (2017). A macro-finance approach to sovereign debt spreads and returns. *manuscript, University of Chicago*.
- Zickfeld, K. and T. Herrington (2015). The time lag between a carbon dioxide emission and maximum warming increases with the size of the emission. *Environmental Research Letters* 10(3), 031001.

10 Appendix A: Theoretical Derivations

10.1 Macroeconomics Outcomes

10.1.1 No Climate Component

Starting from the Planner's Hamilton-Jacobi-Bellman (HJB) equation

$$\begin{aligned}
 0 &= \max_{L, I, N, i_R} \rho(1 - \xi)V \log(A_C K^\gamma L^\alpha (N - i_R R)^{\nu(1-\alpha-\gamma)} (A_G(1 - L)^\omega)^{(1-\nu)(1-\alpha-\gamma)} - I) \\
 &\quad - \rho V \log((1 - \xi)V) + V_K K (\ln B + \delta_1 \ln I - \delta_2 \ln K) + V_R (-N + \Gamma R_t i_{R,t}^\theta) \\
 &\quad + \frac{1}{2} \sigma_R^2 R^2 V_{RR} + \frac{1}{2} \sigma_K^2 K^2 V_{KK}
 \end{aligned}$$

Then first order conditions (FOC) are

$$\begin{aligned}
 n &= \frac{\rho(1 - \xi)(1 - C_1)^{-1} \nu(1 - \alpha - \gamma)V}{V_R R} + (\Gamma\theta)^{1/(1-\theta)} \\
 i_R &= (\Gamma\theta)^{1/(1-\theta)} \\
 I &= \frac{V_K K \delta_1}{\rho(1 - \xi)V + \delta_1 V_K K} A_C K^\gamma L^\alpha (R(n - i_R))^{\nu(1-\gamma-\alpha)} (A_G(1 - L)^\omega)^{(1-\nu)(1-\gamma-\alpha)} \\
 L &= \frac{\alpha}{\alpha + \omega(1 - \nu)(1 - \gamma - \alpha)}
 \end{aligned}$$

Guess and verify the value function and its coefficients are given by

$$\begin{aligned}
 V &= c_0 K^{c_1} R^{c_2} \\
 c_0 &= \frac{1}{1 - \xi} \exp\left(\frac{1}{\rho} \left\{ \rho(1 - \xi) \log(A_C L^\alpha (n_t - i_R)^{\nu(1-\alpha-\gamma)} ((1 - L)^\omega)^{(1-\nu)(1-\alpha-\gamma)} (1 - C_1)) \right. \right. \\
 &\quad \left. \left. + c_2(-n + \Gamma R_t i_{R,t}^\theta) + c_2(c_2 - 1) \frac{\sigma_R^2}{2} \right. \right. \\
 &\quad \left. \left. + c_1(\log B + \delta_1 \log(A_C L^\alpha (n_t - i_R)^{\nu(1-\alpha-\gamma)} ((1 - L)^\omega)^{(1-\nu)(1-\alpha-\gamma)} C_1)) + \frac{\sigma_K^2}{2} c_1(c_1 - 1) \right\} \right) \\
 c_1 &= \frac{\rho(1 - \xi)\gamma}{\rho - \gamma\delta_1 + \delta_2} \\
 c_2 &= \frac{(1 - \xi)\nu(1 - \alpha - \gamma)(\rho + \delta_2)}{\rho - \gamma\delta_1 + \delta_2} \\
 C_1 &= \frac{c_1 \delta_1}{\rho(1 - \xi) + c_1 \delta_1}
 \end{aligned}$$

10.1.2 Climate and Uncertain Climate Policy

The planner's problem for the pre-policy state can be written as an HJB equation that is given by

$$\begin{aligned}
0 = & \max_{L, I, N, i_R} \rho(1 - \xi)V(\log(\exp(-\eta T)A_C K^\gamma L^\alpha (N - i_R R)^{\nu(1-\alpha-\gamma)}(A_G(1-L)^\omega)^{(1-\nu)(1-\alpha-\gamma)} - I) \\
& - \frac{1}{(1-\xi)} \log((1-\xi)V)) + V_K K(\ln B + \delta_1 \ln I - \delta_2 \ln K) + V_R(-N + \Gamma R_t i_{R,t}^\theta) + \varphi N V_T \\
& + \frac{1}{2}\sigma_R^2 R^2 V_{RR} + \frac{1}{2}\sigma_T^2 V_{TT} + \frac{1}{2}\sigma_K^2 K^2 V_{KK} + \lambda(T)[V_{post} - V_{pre}]
\end{aligned}$$

and the First Order Conditions (FOC) are given by

$$\begin{aligned}
N &= \frac{\rho(1-\xi)V(1-C_1)^{-1}\nu(1-\alpha-\gamma)}{V_R - \varphi_T V_T} + \left(\frac{V_R \Gamma \theta}{(V_R - \varphi_T V_T)}\right)^{1/(1-\theta)} R \\
i_R &= \left(\frac{V_R \Gamma \theta}{(V_R - \varphi_T V_T)}\right)^{1/(1-\theta)} \\
I &= \frac{V_K K \delta_1}{\rho(1-\xi)V + V_K K \delta_1} \exp(-\eta T) A_C K^\gamma L^\alpha (N - I_R)^{\nu(1-\gamma-\alpha)} (A_G(1-L)^\omega)^{(1-\nu)(1-\gamma-\alpha)} \\
L &= \frac{\alpha}{\alpha + \omega(1-\nu)(1-\gamma-\alpha)}
\end{aligned}$$

The planner's post-policy problem can be written as an HJB equation that is given by

$$\begin{aligned}
0 = & \rho(1 - \xi)V(\log(\exp(-\eta T)A_C K^\gamma L^\alpha (A_G(1-L)^\omega)^{(1-\alpha-\gamma)} - I) - \frac{1}{(1-\xi)} \log((1-\xi)V)) \\
& + V_K K(\ln B + \delta_1 \ln I - \delta_2 \ln K) + \frac{1}{2}\sigma_T^2 V_{TT} + \frac{1}{2}\sigma_K^2 K^2 V_{KK}
\end{aligned}$$

and the First Order Conditions (FOC) are given by

$$\begin{aligned}
I &= \frac{V_K K \delta_1}{\rho(1-\xi)V + V_K K \delta_1} \exp(-\eta T) A_C K^\gamma L^\alpha (A_G(1-L)^\omega)^{(1-\gamma-\alpha)} \\
L &= \frac{\alpha}{\alpha + \omega(1-\alpha-\gamma)}
\end{aligned}$$

Guess and verify that the pre- and post-policy value functions are given by

$$V_{pre} = K^{c_1} v(R, T) \quad V_{post} = \tilde{c}_0 K^{c_1} \exp(c_3 T)$$

where the coefficients of the value functions are given by

$$\begin{aligned}\tilde{c}_0 &= \frac{1}{1-\xi} \exp\left(\frac{1}{\rho} \left\{ \rho(1-\xi) \log(A_C L^\alpha (A_G(1-L)^\omega)^{(1-\alpha-\gamma)} (1-C_1)) + c_3^2 \frac{1}{2} \sigma_T^2 \right. \right. \\ &\quad \left. \left. + c_1 (\log B + \delta_1 \log(A_C L^\alpha (A_G(1-L)^\omega)^{(1-\alpha-\gamma)} C_1)) + \frac{\sigma_K^2}{2} c_1 (c_1 - 1) \right\}\right) \\ c_1 &= \frac{\rho(1-\xi)\gamma}{\rho - \gamma\delta_1 + \delta_2} \\ c_3 &= -\frac{\eta(1-\xi)(\rho + \delta_2)}{\rho - \delta_1\gamma + \delta_2} \\ C_1 &= \frac{c_1\delta_1}{\rho(1-\xi) + c_1\delta_1}\end{aligned}$$

and the remaining differential equation v solves

$$\begin{aligned}0 &= \rho(1-\xi)v(\log(\exp(-\eta T) A_C L^\alpha (N - i_R R)^{\nu(1-\alpha-\gamma)} (A_G(1-L)^\omega)^{(1-\nu)(1-\alpha-\gamma)} (1-C_1))) \\ &\quad - \frac{1}{(1-\xi)} \log((1-\xi)v) + v_R(-N + \Gamma R_t i_{R,t}^\theta) + \varphi N v_T \\ &\quad + v c_3 (\ln B + \delta_1 \ln(\exp(-\eta T) A_C L^\alpha (N - i_R R)^{\nu(1-\alpha-\gamma)} (A_G(1-L)^\omega)^{(1-\nu)(1-\alpha-\gamma)} C_1)) \\ &\quad + \frac{1}{2} \sigma_R^2 R^2 v_{RR} + \frac{1}{2} \sigma_T^2 v_{TT} + \frac{1}{2} \sigma_K^2 v c_1 (c_1 - 1) + \lambda(T) [\tilde{c}_0 \exp(c_3 T) - v]\end{aligned}$$

10.1.3 Constant Policy Arrival Rate

Extending the uncertain climate policy setting above, I guess and verify that

$$v = \exp(c_3 T) f(R)$$

and the optimal FOC for extraction and exploration are

$$\begin{aligned}n &= \frac{\rho(1-\xi) f(1-C_1)^{-1} \nu(1-\alpha-\gamma)}{(f_R - \varphi c_3 f)} + \left(\frac{f_R \Gamma \theta}{(f_R - \varphi c_3 f)} \right)^{1/(1-\theta)} \\ i_{R,t} &= \left(\frac{f_R \Gamma \theta}{(f_R - \varphi c_3 f)} \right)^{1/(1-\theta)}\end{aligned}$$

All else remains the same as in the original uncertain policy setting.

10.1.4 No Exploration

Set $\Gamma = 0$ and derive solutions as done in previous settings with exploration.

10.1.5 Oil-Sector Only Policy Setting

For this extension, we alter the final output production function to be of the following form:

$$Y_i = A_C L^\alpha K^\gamma O^{\nu_i(1-\alpha-\gamma)} G^{\beta(1-\alpha-\gamma)}$$

While policy shocks are still given by stochastic shocks to ν_i that follow a finite chain Markov process as before, these shocks no longer alter the energy input demand share of green energy, now given by β . With this, the planner's problem for the pre-policy state can be written as an HJB equation that is given by

$$\begin{aligned} 0 = & \max_{L, I, N, i_R} \rho(1-\xi)V(\log(\exp(-\eta T)A_C K^\gamma L^\alpha (N - i_R R)^{\nu(1-\alpha-\gamma)}(A_G(1-L)^\omega)^{\beta(1-\alpha-\gamma)} - I) \\ & - \frac{1}{(1-\xi)} \log((1-\xi)V)) + V_K K(\ln B + \delta_1 \ln I - \delta_2 \ln K) + V_R(-N + \Gamma R_t i_{R,t}^\theta) + \varphi N V_T \\ & + \frac{1}{2}\sigma_R^2 R^2 V_{RR} + \frac{1}{2}\sigma_T^2 V_{TT} + \frac{1}{2}\sigma_K^2 K^2 V_{KK} + \frac{1}{2}\sigma_{A_C}^2 A_C^2 V_{A_C, A_C} + \frac{1}{2}\sigma_{A_G}^2 A_G^2 V_{A_G, A_G} + \lambda(T)[V_{post} - V_{pre}] \end{aligned}$$

and the First Order Conditions (FOC) are given by

$$\begin{aligned} N &= \frac{\rho(1-\xi)V(1-C_1)^{-1}\nu(1-\alpha-\gamma)}{V_R - \varphi_T V_T} + \left(\frac{V_R \Gamma \theta}{(V_R - \varphi_T V_T)}\right)^{1/(1-\theta)} R \\ i_R &= \left(\frac{V_R \Gamma \theta}{(V_R - \varphi_T V_T)}\right)^{1/(1-\theta)} \\ I &= \frac{V_K K \delta_1}{\rho(1-\xi)V + V_K K \delta_1} \exp(-\eta T) A_C K^\gamma L^\alpha (N - I_R)^{\nu(1-\alpha-\gamma)} (A_G(1-L)^\omega)^{\beta(1-\alpha-\gamma)} \\ L &= \frac{\alpha}{\alpha + \omega\beta(1-\alpha-\gamma)} \end{aligned}$$

The planner's post-policy problem can be written as an HJB equation that is given by

$$\begin{aligned} 0 = & \rho(1-\xi)V(\log(\exp(-\eta T)A_C K^\gamma L^\alpha (A_G(1-L)^\omega)^{\beta(1-\alpha-\gamma)} - I) - \frac{1}{(1-\xi)} \log((1-\xi)V)) \\ & + V_K K(\ln B + \delta_1 \ln I - \delta_2 \ln K) + \frac{1}{2}\sigma_T^2 V_{TT} + \frac{1}{2}\sigma_K^2 K^2 V_{KK} \end{aligned}$$

and the First Order Conditions (FOC) are given by

$$\begin{aligned} I &= \frac{V_K K \delta_1}{\rho(1-\xi)V + V_K K \delta_1} \exp(-\eta T) A_C K^\gamma L^\alpha (A_G(1-L)^\omega)^{\beta(1-\alpha-\gamma)} \\ L &= \frac{\alpha}{\alpha + \omega\beta(1-\alpha-\gamma)} \end{aligned}$$

We can guess and verify that the pre- and post-policy value functions are given by

$$V_{pre} = K^{c_1} v(R, T) \quad V_{post} = \hat{c}_0 K^{c_1} \exp(c_3 T)$$

where the coefficients of the value functions are given by

$$\begin{aligned} \hat{c}_0 &= \frac{1}{1-\xi} \exp\left(\frac{1}{\rho} \left\{ \rho(1-\xi) \log(A_C L^\alpha (A_G(1-L)^\omega)^{\beta(1-\alpha-\gamma)} (1-C_1)) + c_3^2 \frac{1}{2} \sigma_T^2 \right. \right. \\ &\quad \left. \left. + c_1 (\log B + \delta_1 \log(A_C L^\alpha (A_G(1-L)^\omega)^{\beta(1-\alpha-\gamma)} C_1)) + \frac{\sigma_K^2}{2} c_1 (c_1 - 1) \right\}\right) \\ c_1 &= \frac{\rho(1-\xi)\gamma}{\rho - \gamma\delta_1 + \delta_2} \\ c_3 &= -\frac{\eta(1-\xi)(\rho + \delta_2)}{\rho - \delta_1\gamma + \delta_2} \\ C_1 &= \frac{b\delta_1}{\rho(1-\xi) + b\delta_1} \end{aligned}$$

and the remaining differential equation v solves

$$\begin{aligned} 0 &= \rho(1-\xi)v(\log(\exp(-\eta T) A_C L^\alpha (N - i_R R)^{\nu(1-\alpha-\gamma)} (A_G(1-L)^\omega)^{\beta(1-\alpha-\gamma)} (1-C_1))) \\ &\quad - \frac{1}{(1-\xi)} \log((1-\xi)v)) + v_R(-N + \Gamma R_t i_{R,t}^\theta) + \varphi N v_T \\ &\quad + v c_3 (\ln B + \delta_1 \ln(\exp(-\eta T) A_C L^\alpha (N - i_R R)^{\nu(1-\alpha-\gamma)} (A_G(1-L)^\omega)^{\beta(1-\alpha-\gamma)} C_1)) \\ &\quad + \frac{1}{2} \sigma_R^2 R^2 v_{RR} + \frac{1}{2} \sigma_T^2 v_{TT} + \frac{1}{2} \sigma_K^2 v c_1 (c_1 - 1) + \lambda(T) [\hat{c}_0 \exp(c_3 T) - v] \end{aligned}$$

10.1.6 Stochastic TFP and Green Capital Extensions

Using the oil-sector only policy setting, it is straightforward to show how stochastic TFP for the final output sector and green energy sector can be introduced to the model in a straightforward way. Introducing green capital can be done in a similar manner. If we assume that the TFP variables are geometric Brownian motions

$$\begin{aligned} dA_{C,t} &= \mu_{A_C} A_{C,t} dt + \sigma_{A_C} A_{C,t} dB_{A_C} \\ dA_{G,t} &= \mu_{A_G} A_{G,t} dt + \sigma_{A_G} A_{G,t} dB_{A_G} \end{aligned}$$

The stochastic TFP for each sector replace the constant values from before

$$\begin{aligned} Y_t &= A_{C,t} K_t^\gamma L_{C,t}^\alpha E_t^{1-\gamma-\alpha} \\ G_t &= A_{G,t} L_{G,t}^\omega \end{aligned}$$

Then we can guess and verify the value function and its coefficients are given by

$$\begin{aligned}
V &= c_5 A_C^{c_1} A_G^{c_2} K^{c_3} R^{c_4} \\
c_5 &= \frac{1}{1-\xi} \exp\left(\frac{1}{\rho} \left\{ \rho(1-\xi) \log L^\alpha (n_t - i_R)^{\nu(1-\alpha-\gamma)} ((1-L)^\omega)^{\beta(1-\alpha-\gamma)} (1-C_1) \right. \right. \\
&\quad \left. \left. + c_4(-n + \Gamma R_t i_{R,t}^\theta) + c_4(c_4 - 1) \frac{\sigma_R^2}{2} \right. \right. \\
&\quad \left. \left. + c_3(\log B + \delta_1 \log L^\alpha (n_t - i_R)^{\nu(1-\alpha-\gamma)} ((1-L)^\omega)^{\beta(1-\alpha-\gamma)} C_1) + \frac{\sigma_K^2}{2} c_3(c_3 - 1) \right\}\right) \\
c_1 &= \frac{\rho(1-\xi)}{\rho - \gamma\delta_1 + \delta_2} \\
c_2 &= \frac{\rho(1-\xi)\beta(1-\alpha-\gamma)}{\rho - \gamma\delta_1 + \delta_2} \\
c_3 &= \frac{\rho(1-\xi)\gamma}{\rho - \gamma\delta_1 + \delta_2} \\
c_4 &= \frac{(1-\xi)\nu(1-\alpha-\gamma)(\rho + \delta_2)}{\rho - \gamma\delta_1 + \delta_2} \\
C_1 &= \frac{b\delta_1}{\rho(1-\xi) + b\delta_1}
\end{aligned}$$

Everything else remains the same as in the oil-sector only policy setting and asset prices can be derived in the same way as shown here. Additional constant contributions to risk prices and risk premia will result from the inclusion of the stochastic TFP variables. The inclusion of a green capital, whose evolution follows the same structure as the final output capital, would result in a very similar analytical setting as introducing stochastic green sector TFP.

10.1.7 Decentralized Economy

10.1.8 Household

The household optimization problem is given by

$$\begin{aligned}
V &= \max_C E \int \rho(1-\xi)V(\log C - \frac{1}{1-\xi} \log(1-\xi)V)dt \\
s.t. & \quad W_t \geq \int \pi C
\end{aligned}$$

The SDF is given by

$$\pi = \exp\left(\int h_V\right)\rho(1-\xi)VC^{-1}$$

and

$$\begin{aligned} h_C &= \rho(1 - \xi)VC^{-1} \\ h_J &= \rho(1 - \xi) \log C - \rho \log((1 - \xi)V) - \rho \end{aligned}$$

10.1.9 Final Output

The final output firm's profit maximization problem is given by

$$\begin{aligned} V_F &= \max_{L_C, I, O, S} E \int \pi(\exp(-\eta T) A_C K^\gamma L_C^\alpha O^{\nu(1-\alpha-\gamma)} G^{(1-\nu)(1-\alpha-\gamma)} - wL_C - P_I I - P_O O - P_G G) ds \\ \text{s.t. } & dK = K(\ln B + \delta_1 \ln I - \delta_2 \ln K) \end{aligned}$$

The FOC are given by

$$\begin{aligned} P_I &= \lambda_K K \delta_1 \pi^{-1} I^{-1} \\ P_O &= \nu(1 - \alpha - \gamma) \exp(-\eta T) A K^\gamma L_C^\alpha O^{\nu(1-\alpha-\gamma)-1} G^{(1-\nu)(1-\alpha-\gamma)} \\ w &= \alpha \exp(-\eta T) A K^\gamma L_C^{\alpha-1} O^{\nu(1-\alpha-\gamma)} G^{(1-\nu)(1-\alpha-\gamma)} \\ P_G &= (1 - \nu)(1 - \gamma - \alpha) \exp(-\eta T) A K^\gamma L_C^\alpha O^{\nu(1-\alpha-\gamma)} G^{(1-\nu)(1-\alpha-\gamma)-1} \end{aligned}$$

Taking the SDF and value function as given, by definition the Langrangian multiplier λ_K is given by the discounted marginal value of another unit of capital, i.e., $\lambda_K = \exp(\int h_V) V_K$, and so

$$\frac{c_1 \delta_1}{\rho(1 - \xi) + c_1 \delta_1} \tilde{Y} = I$$

given $h_C = \rho(1 - \xi)VC^{-1}$, $V = K^{c_1} v(R, T)$, and $P_I = 1$.

10.1.10 Green Firm

The green firm's profit maximization problem is given by

$$V_G = \max_{L_G} E \int \pi(P_S A_G (L_G)^\omega - w(L_G)) ds$$

The FOC is

$$w = P_G A_G$$

Now, given P_G and w from above and taking the SDF and value function as given, we get

$$\begin{aligned} L_C &= \frac{\alpha}{(1-\nu)(1-\gamma-\alpha)\omega + \alpha} \\ L_G &= 1 - L_C \end{aligned}$$

10.1.11 Oil Firm and Optimal Tax

From the oil firm's profit maximization problem, which includes a tax on the oil extraction piece of output only as that is the only piece contributing to emissions, we see

$$\begin{aligned} V_O &= \max_{n, i_R} E \int \pi(P_O R((1-\tau_{opt})n - i_R)) ds \\ s.t. \quad dR/R &= -n dt + \Gamma i_{R,t}^\theta dt + \sigma_R dB \\ dT &= \varphi_T n R dt + \sigma_T dB \end{aligned}$$

The FOC for extraction and exploration are given by

$$\begin{aligned} P_O &= \lambda_R \Gamma \theta i_R^{\theta-1} R \pi^{-1} \\ P_O &= \lambda_R R (1-\tau)^{-1} \pi^{-1} \end{aligned}$$

Taking P_O , the SDF, and the value function as given previously, and by definition the Langrangian multiplier λ_R is the discounted marginal value of another unit of oil, i.e., $\lambda_R = \exp(\int h_V) V_R$. Plugging in these expressions we find

$$\begin{aligned} n &= \frac{\rho(1-\xi)vY(C)^{-1}\nu(1-\alpha-\gamma)(1-\tau)}{v_R} + i_R R \\ i_R &= (\Gamma\theta(1-\tau))^{1/(1-\theta)} \end{aligned}$$

Note that the Social Planner's FOC derived from the HJB equation are given by

$$\begin{aligned} n &= \frac{\rho(1-\xi)v(1-C_1)^{-1}\nu(1-\alpha-\gamma)}{v_R - \varphi_T v_T} + \left(\frac{v_R \Gamma \theta}{(v_R - \varphi_T v_T)}\right)^{1/(1-\theta)} \\ i_R &= \left(\frac{v_R \Gamma \theta}{(v_R - \varphi_T v_T)}\right)^{1/(1-\theta)} \end{aligned}$$

Equating the SP and decentralized FOCs provides a system of equations from which the optimal tax can be derived as

$$(1 - \tau_{opt}) = \frac{v_R}{v_R - \varphi_T v_T}$$

10.2 Asset Pricing Outcomes

10.2.1 The Stochastic Discount Factor (SDF)

Note that the intertemporal marginal rate of substitution (IMRS) or stochastic discount factor (SDF) following Duffie and Skiadas (1994) is

$$\pi_t = \exp\left(\int_0^t h_J(C, V) ds\right) h_C(C, V)$$

where the utility function h and its derivatives are given by

$$\begin{aligned} h &= \rho(1 - \xi)V \log(C) - \rho V \log((1 - \xi)V) \\ h_C &= \rho(1 - \xi)K^{c_1 - \gamma} v(R, T) \\ &\quad \times \exp(\eta T) A_C^{-1} L_C^{-\alpha} O^{-\nu(1 - \alpha - \gamma)} G^{-(1 - \nu)(1 - \alpha - \gamma)} (1 - C_1)^{-1} \\ h_J &= \rho(1 - \xi) \log(A_C K^\gamma L^\alpha O^{\nu(1 - \alpha - \gamma)} G^{(1 - \nu)(1 - \alpha - \gamma)} (1 - C_1)) \\ &\quad - \rho(1 - \xi)\eta T - \rho \log((1 - \xi)K^{c_1} v(R, T)) - \rho \end{aligned}$$

As shown by Duffie and Skiadas (1994), Ito's Lemma then gives $\frac{d\pi_t}{\pi_t} = h_J dt + \frac{Dh_C}{h_C}$ where

$$\begin{aligned} \frac{dh_C}{h_C} &= (c_1 - \gamma)(\ln B + \delta_1 \ln I - \delta_2 \ln K) dt + \frac{1}{2}(c_1 - \gamma)(c_1 - \gamma - 1)\sigma_K^2 dt \\ &\quad + \left\{ \frac{v_R}{v} - \nu(1 - \alpha - \gamma) \frac{O_R}{O} \right\} R(-N_t + \Gamma R_t i_{R,t}^\theta) dt \\ &\quad + \left\{ \frac{v_T}{v} - \nu(1 - \alpha - \gamma) \frac{O_T}{O} + \eta \right\} (\varphi_T N_t) dt \\ &\quad + \frac{1}{2} \left\{ \frac{v_{RR}}{v} - 2\nu(1 - \alpha - \gamma) \frac{v_R}{v} \frac{O_R}{O} + \nu(1 - \alpha - \gamma) \{ \nu(1 - \alpha - \gamma) + 1 \} \frac{O_R^2}{O^2} \right. \\ &\quad \left. - \nu(1 - \alpha - \gamma) \frac{O_{RR}}{O} \right\} R^2 \sigma_R^2 dt \\ &\quad + \frac{1}{2} \left\{ \frac{v_{TT}}{v} - 2\nu(1 - \alpha - \gamma) \frac{v_T}{v} \frac{O_T}{O} + \nu(1 - \alpha - \gamma) \{ \nu(1 - \alpha - \gamma) + 1 \} \frac{O_T^2}{O^2} \right. \\ &\quad \left. - \nu(1 - \alpha - \gamma) \frac{O_{TT}}{O} \right\} \sigma_T^2 dt + \frac{1}{2} \left\{ 2 \frac{v_T}{v} \eta - 2\nu(1 - \alpha - \gamma) \frac{O_T}{O} + \eta^2 \right\} \sigma_T^2 dt \\ &\quad + \left\{ \frac{c_0 \exp(c_3 T) L_{post}^{-\alpha} G_{post}^{-(1 - \alpha - \gamma)}}{v(R, T) L_{pre}^{-\alpha} G_{pre}^{-(1 - \nu)(1 - \alpha - \gamma)} O^{-\nu(1 - \alpha - \gamma)}} - 1 \right\} dJ + (c_1 - \gamma) dB_K \\ &\quad + \left\{ \frac{v_R}{v} - \nu(1 - \alpha - \gamma) \frac{O_R}{O} \right\} R \sigma_R dB_R + \left\{ \frac{v_T}{v} - \nu(1 - \alpha - \gamma) \frac{O_T}{O} + \eta \right\} \sigma_T dB_t \end{aligned}$$

For the post-policy state, we have

$$\begin{aligned}
h_C &= \rho(1 - \xi)c_0K^{c_1-\gamma}exp(c_3T) \\
&\quad \times exp(\eta T)A_C L_C^{-\alpha} G^{-\omega(1-\alpha-\gamma)}(1 - C_1)^{-1} \\
h_J &= \rho(1 - \xi) \log(A_C K^\gamma L^\alpha G^{(1-\alpha-\gamma)}(1 - C_1)) \\
&\quad - \rho(1 - \xi)\eta T - \rho \log((1 - \xi)c_0K^{c_1} exp(c_3T)) - \rho \\
\frac{dh_C}{h_C} &= (c_1 - \gamma)(\ln B + \delta_1 \ln I - \delta_2 \ln K)dt + \frac{1}{2}(c_1 - \gamma)(c_1 - \gamma - 1)\sigma_K^2 dt \\
&\quad + \frac{1}{2}\{c_3^2 + 2c_3\eta + \eta^2\}\sigma_T^2 dt \\
&\quad + (c_1 - \gamma)dB_K + \{c_3 + \eta\}\sigma_T dB_T
\end{aligned}$$

10.2.2 Risk Prices

The risk prices are the loadings on the Brownians and Jump process for the SDF, so they are

$$\begin{aligned}
\sigma_{\pi,K} &= (\gamma - c_1)\sigma_K \\
\sigma_{\pi,R} &= \left\{ \nu(1 - \alpha - \gamma) \frac{O_R}{O} - \frac{v_R}{v} \right\} \sigma_R R \\
\sigma_{\pi,T} &= \left\{ \nu(1 - \alpha - \gamma) \frac{O_T}{O} - \frac{v_T}{v} - \eta \right\} \sigma_T \\
\Theta_\pi &= \left\{ 1 - \frac{v_{post} \tilde{Y}_{post}^{-1}}{v_{pre} \tilde{Y}_{pre}^{-1}} \right\}
\end{aligned}$$

10.2.3 Firm Prices

To derive firm prices, I apply the envelope theorem to the social planner's Lagrangian. This follows the methodology used by Papanikolaou (2011) for example. Note for the final output firm we have

$$\begin{aligned}
\pi_t S_t^C &= E_t \int_t^\infty \pi_s (exp(-\eta T) A K^\gamma L^\alpha O^{\nu(1-\alpha-\gamma)} G^{(1-\nu)(1-\alpha-\gamma)} - wL - P_O O - P_G G - I) ds \\
&= E_t \int_t^\infty \pi_s (exp(-\eta T) A K^\gamma L^\alpha O^{\nu(1-\alpha-\gamma)} G^{(1-\nu)(1-\alpha-\gamma)} \gamma - i^* K) ds \\
\implies \\
S_t^C &= E_t \int_t^\infty exp\left(\int_t^s h_V\right) \frac{h_{C,s}}{h_{C,t}} (exp(-\eta T) A K^\gamma L^\alpha O^{\nu(1-\alpha-\gamma)} G^{(1-\nu)(1-\alpha-\gamma)} \gamma - i^* K) ds
\end{aligned}$$

For the oil firm, plugging in the socially optimal choices, we have

$$\begin{aligned}
\pi_t S_t^O &= E_t \int_t^\infty \pi_s (P_O R((1 - \tau)n - i_R) ds \\
\implies \\
S_t^O &= E_t \int_t^\infty exp\left(\int_t^s h_V\right) \frac{h_{C,s}}{h_{C,t}} P_O R(n^* - i_R^*) ds
\end{aligned}$$

Note the Lagrangian for the social planner's problem is given by

$$\mathcal{L} = E_t \int_t^\infty \{h(C, V) - \pi_s(C - \exp(-\eta T)AK^\gamma L^\alpha O^{\nu(1-\alpha-\gamma)}G^{(1-\nu)(1-\alpha-\gamma)} + iK) - P_O\pi_s(O - nR + i_R R) - P_G\pi_s(G - A_G L_G)\} ds$$

Therefore, by application of the envelope theorem we know that

$$\frac{\partial \mathcal{L}}{\partial K} = \frac{\partial V}{\partial K} \quad , \quad \frac{\partial \mathcal{L}}{\partial R} = \frac{\partial V}{\partial R}$$

Furthermore, we also know that

$$\frac{\partial K_s}{\partial K_t} K_t = K_s \quad , \quad \frac{\partial R_s}{\partial R_t} R_t = R_s$$

Calculating derivatives of Lagrangian and comparing I find that

$$\begin{aligned} S_t^C &= \frac{1}{h_C} \frac{\partial V}{\partial K} K = c_1 \frac{(1 - C_1)}{\rho(1 - \xi)} \tilde{Y}_t \\ S_t^O &= \frac{1}{h_C} \frac{\partial V}{\partial R} R = \frac{(1 - C_1)}{\rho(1 - \xi)} \frac{v_R R}{v} \tilde{Y}_t \end{aligned}$$

Lastly, for the green energy firm we have

$$\begin{aligned} \pi_t S_t^G &= E_t \int_t^\infty \pi_s (P_G A_G L_G^\omega - w L_G) ds \\ &\implies \\ S_t^G &= E_t \int_t^\infty \exp\left(\int_t^s h_V\right) \frac{h_{C,s}}{h_{C,t}} P_G (1 - \omega) (A_G L_G) ds \end{aligned}$$

Plugging in for P_G from the final output firm's FOC, and through substitution for the optimal choice of investment i , we can rewrite this as

$$\begin{aligned} S_t^G &= (1 - \nu)(1 - \gamma - \alpha)(1 - \omega) E_t \int_t^\infty \exp\left(\int_t^s h_V\right) \frac{h_{C,s}}{h_{C,t}} \tilde{Y}_t ds \\ &= \frac{(1 - \nu)(1 - \gamma - \alpha)(1 - \omega)}{\gamma - C_1} E_t \int_t^\infty \exp\left(\int_t^s h_V\right) \frac{h_{C,s}}{h_{C,t}} (\tilde{Y}_t \gamma - i^* K) ds \\ &= \frac{(1 - \nu)(1 - \gamma - \alpha)(1 - \omega)}{\gamma - C_1} S_t^C \end{aligned}$$

Therefore, the firm prices for each sector in the model are given by

$$S_t^C = c_1 \frac{(1 - C_1)}{\rho(1 - \xi)} \tilde{Y}_t, \quad S_t^G = \frac{(1 - \nu)(1 - \gamma - \alpha)(1 - \omega)}{\gamma - C_1} \frac{c_1(1 - C_1)}{\rho(1 - \xi)} \tilde{Y}_t, \quad S_t^O = \frac{(1 - C_1)}{\rho(1 - \xi)} \frac{v_R R}{v} \tilde{Y}_t$$

10.2.4 Risk Premia

Since the process for returns is defined as $\frac{dP+D}{P}$, the risk premia are given by

$$RP = -cov\left(\frac{dS}{S}, \frac{d\pi}{\pi}\right) = \sigma_{Ret,K}\sigma_{\pi,K} + \sigma_{Ret,R}\sigma_{\pi,R} + \sigma_{Ret,T}\sigma_{\pi,T} + \Theta_{\pi}\Theta_R$$

where $\sigma_{Ret,x}$ comes from the diffusion term of the price equation after applying Ito's to the pricing term. Plugging in appropriate terms gives the risk premium for a given firm (indexed by $X = Mkt, C, O, G$) as

$$RP^X = \gamma(\gamma - c1)\sigma_K^2 + \sum_{\chi=R,T} \left(\left(\frac{\partial}{\partial \chi} S^X \right) / S^X \right) \sigma_{\chi}(\chi) \sigma_{\pi,\chi} + \lambda(T_t) \Theta_{\pi} (S_{post}^X / S_1^X pre - 1)$$

Here $\frac{\partial}{\partial \chi} S^X$ is the partial derivative of price S^X with respect to state χ , $\sigma_{\chi}(\chi)$ are the volatility for the state variables, and $\sigma_{\pi,\chi}$, Θ_{π} are the risk prices derived from the SDF previously.

To calculate these quantities, note that

$$\begin{aligned} S_R/S &= \left\{ \frac{v_{RR}}{vv_R} - \frac{v_R}{v^2} + R^{-1} + \nu(1 - \alpha - \gamma)O^{-1}O_R \right\} \\ S_{RR}/S &= \left\{ \left\{ \frac{v_{RRR}}{v} + 2\frac{v_R^3}{v^3} - 3\frac{v_{RR}v_R}{v^2} \right\} v_R^{-1} + \left\{ \frac{v_{RR}}{v} - \frac{v_R^2}{v^2} \right\} \nu(1 - \alpha - \gamma)v_R^{-1}O^{-1}O_R \right. \\ &\quad \left. + \left\{ v_{RR} + \frac{v_{RR}}{v} - \frac{v_R^2}{v^2} \right\} R^{-1}v_R^{-1} + 2\nu(1 - \alpha - \gamma)O^{-1}O_R \right\} \\ S_T/S &= (-\eta + \nu(1 - \alpha - \gamma)O^{-1}O_T) \\ S_{TT}/S &= (\eta^2 - 2\eta\nu(1 - \alpha - \gamma)O^{-1}O_T + \nu(1 - \alpha - \gamma)\{\nu(1 - \alpha - \gamma) - 1\}O^{-2}O_T^2 \\ &\quad + \nu(1 - \alpha - \gamma)O^{-1}O_{TT}) \\ S_K/S &= \gamma K^{-1} \\ S_{KK}/S &= \gamma(\gamma - 1)K^{-2} \end{aligned}$$

and so the process for returns, given by an application of Ito's lemma, is

$$\begin{aligned} \frac{dS}{S} &= \{\mu_K S_K/S + \mu_R S_R/S + \mu_T S_T/S\}dt \\ &\quad + \frac{1}{2} \{ \sigma_R^2 R^2 S_{RR}/S + \sigma_T^2 S_{TT}/S + \sigma_K^2 K^2 S_{KK}/S + \sigma_{A_G}^2 A_G^2 S_{A_G,A_G}/S + \sigma_{A_C}^2 A_C^2 S_{A_C,A_C}/S \} dt \\ &\quad + \sigma_{A_C} S_{A_C}/S dB_{A_C} + \sigma_{A_G} S_{A_G}/S dB_{A_G} + \sigma_K S_K/S dB_K + \sigma_R S_R/S dB_R + \sigma_T S_T/S dB_T + \Theta_R dJ_t \end{aligned}$$

With this we can plug in and get the expressions for expected returns.

11 Appendix B: Additional Model Extensions

11.1 Ignoring the Climate Externality

A critical assumption to consider within the context of the model is that of using a social planner's problem. In particular, it is important that we understand how the internalization of the climate externality influences the main results of the paper. As we are somewhere between a perfectly decentralized economy that ignores the climate externality and a socially optimizing economy that internalizes the impact of production decisions on climate change. This can be seen by the fact that while numerous countries and producers appear to have no concerns about climate change, there are still significant climate policy actions taking place such as the European Union climate policy actions, including carbon pricing, and California and other states continued efforts to implement the Paris Accord standards even though the US has pulled out of the climate agreement.

To understand this effect, consider the case where we assume an approximately "decentralized" setting with regards to the climate where the planner ignores the climate externality. To do this, I solve what is sometimes known as the second best equilibrium where a planner problem is solved with optimal decisions constrained to satisfy the decentralized economy, i.e., where externalities are ignored. This setting is a more simplified approach than solving the complex setting where the respective firms own the capital stock and oil reserves because this would require jointly solving a system of coupled non-linear partial differential equations jointly characterizing the individual firm problems and the household problem. This "decentralized" setting is closely aligned with the setting where households are assumed to own capital and oil reserves and rent it to firms. Such a setting leads to static firm problems and reduces the problem to a single optimization problem for the household as the shadow prices of capital and oil reserves coincide with the household problem's continuation value marginal values and the climate component influences the marginal value through aggregation and equilibrium adding up but do not influence optimal decisions. With this set-up, the following proposition provides the solution for this case:

Proposition 9. *In the "decentralized" setting where climate externalities are ignored, the value functions are unchanged for the two policy regimes and are given by:*

$$V_{pre}(K_t, R_t, T_t) = K_t^{c_1} v(R_t, T_t) \quad V_{post}(K_t, T_t) = \bar{c}_0 K_t^{c_1} \exp(c_3 T_t)$$

where investment and labor decisions are given by

$$\begin{aligned} I_{pre,t} &= C_1 \tilde{Y}_{pre,t} & L_{pre,C} &= \tilde{L} \\ I_{post,t} &= C_1 \tilde{Y}_{post,t} & L_{post,C} &= \tilde{L} \end{aligned}$$

Exploration and extraction are given by

$$i_{R,t} = (\Gamma\theta)^{1/(1-\theta)} \quad N_t = \frac{v\vartheta}{v_R} + i_{R,t}R_t$$

Note $v(R_t, T_t)$ is the solution to the simplified HJB equation characterizing agent's problem (given in the appendix). The value function constants \bar{c}_0, c_1, c_3 and the FOC constants $\vartheta, C_1, \tilde{L}$ are functions of the model parameters only (also given in the appendix).

While much of the structure of this solution looks quite similar, there are key differences. First, the most obvious difference is for the optimal choices of exploration and extraction. In particular, exploration is now a constant and extraction no longer depends on the marginal cost of climate change. Note first that this is the same level of extraction as when there is no risk of a climate policy shock, which we have seen is significantly higher than the socially optimal exploration when there is climate policy risk. As oil extraction is additively increased by the amount of oil extraction, this leads to higher oil extraction. Second, as the the marginal cost of climate change is positive, i.e., $-\varphi v_T \geq 0$, even with the risk of a climate policy shock, we know that holding all else constant ignoring this impact also increases the level of oil extraction. Thus, the effect of ignoring the climate externality leads to an exacerbated run on oil as long as the same dynamics effects of climate policy risk are in play. Yet, because the risk of a climate policy shock still exists and is increasing in temperature, the same discounting affect that leads to a lower and dynamically decreasing marginal value of oil reserves that was critical in the baselin model setting still exists. As a result, the expected outcome of not internalizing the climate externality is to increase the run on oil, both in levels and dynamically. The intuition is that by ignoring the climate externality the decision maker no longer takes into account the cost of climate change, but they still incorporate into their decision making and discounting the risk of the climate policy shock and thus dynamically run up oil extraction, amplifying the production and pricing impacts seen in the social planner's setting.

11.2 Imperfect Oil Sector Competition

Another important assumption to explore is that of a perfectly competitive oil sector. The existence of OPEC and numerous state-owned oil firms suggests that a model of imperfect competition or market power in the oil sector may more accurately approximate the real world. A simplified approximation of imperfect competition can be derived by assuming a symmetrical oligopolist in the oil sector where firms account for pricing impacts that they have. In particular, I assume J firms use a common pool of reserves, they are homogeneous in production technology, and they internalize their impact on global reserves to ensure a symmetric solution. The evolution of oil reserves and temperature are adjusted

as follows:

$$\begin{aligned}
dR_t &= \left(\sum_{j=1}^J -N_{j,t} \right) dt + \Gamma R_t \left(\sum_{j=1}^J i_{j,R,t} \right)^\theta + \sigma_R R_t dB_R \\
dT_t &= \varphi \sum_{j=1}^J N_{j,t} dt + \sigma_T dB_T
\end{aligned}$$

With this, I solve a constrained planner's problem where oil firms optimize as price setters rather than price takers (as would be socially optimal). However, the impact on climate of oil production is internalized. Thus the constrained optimization is only constrained along the dimension of competition. With these assumptions I derive a symmetric equilibrium solution for this setting. This setting highlights how competition influences equilibrium outcomes when confronting uncertainty climate policy in the model. The following proposition provides the solution for this case of the model:

Proposition 10. *In the oligopolistic oil sector setting, the value functions are unchanged for the two policy regimes and are given by:*

$$V_{pre}(K_t, R_t, T_t) = K_t^{c_1} v(R_t, T_t) \quad V_{post}(K_t, T_t) = \bar{c}_0 K_t^{c_1} \exp(c_3 T_t)$$

where investment and labor decisions are given by

$$\begin{aligned}
I_{pre,t} &= C_1 \tilde{Y}_{pre,t} & L_{pre,C} &= \tilde{L} \\
I_{post,t} &= C_1 \tilde{Y}_{post,t} & L_{post,C} &= \tilde{L}
\end{aligned}$$

Exploration and extraction are given by

$$N_{j,t} = \frac{v \vartheta_J}{J(v_R - v_T \varphi)} \quad i_{j,R,t} = \frac{1}{J} \left(\frac{\Gamma \theta v_R}{v_R - v_T \varphi} \right)^{1/(1-\theta)}$$

Note $v(R_t, T_t)$ is the solution to the simplified HJB equation characterizing agent's problem (given in the appendix). The value function constants \bar{c}_0, c_1, c_3 and the FOC constants C_1, \tilde{L} are functions of the model parameters only (also given in the appendix). The constant ϑ_J is given by

$$\vartheta_J = \rho(1 - \xi)(1 - C_1)^{-1} \nu(1 - \alpha - \gamma) J^{-1} \{ J + \nu(1 - \alpha - \gamma) - 1 \}$$

This proposition provides further valuable intuition about the role of competition in the model and how it interacts with the effects of uncertain climate policy. The first order conditions are almost identical in functional form regardless of whether we are in a perfectly competitive or oligopolistic oil sector. Moreover, as was the case before, the uncertain climate policy has the potential to generate a strong non-linearity in the value of holding reserves by firms and of temperature, thus driving the

potential for a run on oil by the firms.

An important distinction between the optimal extraction in this setting and the competitive setting is the scaling factor ϑ_J that is a function of the number of firms in the oligopoly. Holding the value function constant, the optimal level of oil extracted by each firm decreases as the number of firms goes to infinity. However, there are two other critical effects to keep in mind. As aggregate demand scales extraction by the number of firms, in the limit production would actually go to R_t , the maximum amount of extraction possibly, assuming the value function is held constant. The decrease in production per firm is less than the aggregation scaling, and so increased competition actually amplifies the run through this channel. The other impact of increasing the number of firms is the potential impact it could have on the value function and the marginal value of reserves and temperature, which this simple comparative static ignored. This effect, interacted with the impact of uncertain climate policy, can serve to either amplify or dampen the run on oil effect generated by climate policy. Further exploration of the full numerical solution will allow us to more precisely determine the role that competition has in this model.

11.3 Log Utility Setting

Another important to understand is how the choice of utility function influences the observed outcome. As log utility is represents a special case of the recursive utility specification I use in my analysis, studying the results of the model under log utility helps demonstrate the influence that the recursive structure has on the key outcomes of interest. Therefore, assume in this extension that the utility function is now given by $U(C) = \log(C)$. Then the equilibrium solution is given as follows:

Proposition 11. *With uncertain climate policy where $\nu_t = \nu$ before the policy shock and $\nu_t = 0$ after the policy shock, and where the arrival rate of policy is given by the temperature dependent function $\lambda(T_t)$, and log utility, the value functions for the two policy regimes are given by:*

$$V_{pre} = \hat{c}_1 \log(K_t) + v(R_t, T_t) \quad V_{post} = \log(\hat{c}_0) + \hat{c}_1 \log(K_t) + \hat{c}_3 T_t$$

where investment and labor decisions are given by

$$\begin{aligned} I_{pre,t} &= C_1 \tilde{Y}_{pre,t} & L_{pre,C} &= \bar{L}_{pre} \\ I_{post,t} &= C_1 \tilde{Y}_{post,t} & L_{post,C} &= \bar{L}_{post} \end{aligned}$$

Exploration and extraction are given by

$$i_{R,t} = \left(\frac{\Gamma \theta v_R}{v_R - v_T \varphi} \right)^{1/(1-\theta)} \quad N_t = \frac{\hat{\vartheta}}{v_R - v_T \varphi} + i_{R,t} R_t$$

Note $v(R_t, T_t)$ is the solution to the simplified HJB equation characterizing the planner's problem (given in the appendix). The value function constants \hat{c}_0, c_1, c_3 and the FOC constants $\hat{\vartheta}, C_1, \bar{L}_{pre}, \bar{L}_{post}$ are functions of the model parameters only (also given in the appendix).

There are two significant effects that occur because of the use of log utility. First, the value function becomes additively separable rather than multiplicatively separable. Thus recursive preferences lead to important interactions in the value function that are likely to be for more significant than the additive interactions in this case. The second is that the optimal choice of extraction, though quite similar, no longer has the value function in the numerator of the optimal expression. This highlights again that recursive preferences lead to an amplified effect on the outcomes. As the value function is negative under the recursive utility specification and ϑ is also negative, we saw that increases in temperature and decreases in reserves lead to more negative value function outcomes and thus increased oil extraction as a result of the uncertain climate policy. Here we lose that recursive amplification impact. However, the two critical drivers previously highlighted were the marginal value of reserves and the marginal cost of climate change. Those two components are still present in the optimal choice of oil extraction. Therefore, the same dynamic impacts of the risk of the climate policy shock that strands oil reserves are still in effect. Thus even without recursive utility we would still expect a run on oil and the same dynamic pricing implication. However, as we should expect, we see that the recursive utility specification incorporates an additional amplification effect related to the continuation value and forward looking concerns about the resolution of uncertainty that the log utility setting does not have.

11.4 EZ Preferences where EIS $\neq 1$

While the model with Epstein-Zin type preferences where the EIS is not unitary becomes quite unwieldy, we can still highlight the potential impact that relaxing the model to this setting might have through the risk price characterizations. In particular, I focus on the climate policy jump risk premium. In the case where the EIS is one, we saw that this risk price was given by

$$\Theta_\pi = \left\{ 1 - \frac{V_{post}}{V_{pre}} \left(\frac{C_{post}}{C_{pre}} \right)^{-1} \right\}$$

Without the EIS restriction, and denoting the EIS as ψ^{-1} , preferences are given by

$$h(C, V) = \frac{\rho}{1 - \psi^{-1}} (C^{1-\psi^{-1}} ((1 - \xi)V)^{\frac{\psi^{-1}-\xi}{1-\xi}} - (1 - \xi)V)$$

The stochastic discount factor given by $\pi_t = \exp(\int_0^t h_V)h_C$, but these derivatives are now

$$\begin{aligned} h_V &= -\rho \frac{(\xi - \psi^{-1})}{1 - \psi^{-1}} C^{1-\psi^{-1}} ((1 - \xi)V)^{\frac{\psi^{-1}-1}{1-\xi}} - \rho \frac{(1 - \xi)}{1 - \psi^{-1}} \\ h_C &= \rho C^{-1} C^{1-\psi^{-1}} ((1 - \xi)V)^{\frac{\psi^{-1}-\xi}{1-\xi}} \end{aligned}$$

Therefore, the climate policy jump risk price would therefore be given by

$$\Theta'_\pi = \left\{ 1 - \left(\frac{V_{post}}{V_{pre}} \right)^{\frac{\psi^{-1}-\xi}{1-\xi}} \left(\frac{C_{post}}{C_{pre}} \right)^{-\psi^{-1}} \right\}$$

Note that for the model simulations results the climate policy shock leads to reduced consumption and a more negative continuation value and so $\frac{C_{post}}{C_{pre}} < 1$ and $\frac{V_{post}}{V_{pre}} > 1$. As a result, holding all else constant, when $\psi^{-1} > 1$, we see that

$$\begin{aligned} \left(\frac{C_{post}}{C_{pre}} \right)^{-\psi^{-1}} &> \left(\frac{C_{post}}{C_{pre}} \right)^{-1} \\ \left(\frac{V_{post}}{V_{pre}} \right)^{\frac{\psi^{-1}-\xi}{1-\xi}} &> \left(\frac{V_{post}}{V_{pre}} \right) \end{aligned}$$

On the other hand, holding all else constant, when $\psi^{-1} < 1$, we see that

$$\begin{aligned} \left(\frac{C_{post}}{C_{pre}} \right)^{-\psi^{-1}} &< \left(\frac{C_{post}}{C_{pre}} \right)^{-1} \\ \left(\frac{V_{post}}{V_{pre}} \right)^{\frac{\psi^{-1}-\xi}{1-\xi}} &< \left(\frac{V_{post}}{V_{pre}} \right) \end{aligned}$$

Therefore, the result of relaxing the EIS from being unitary is that when $\psi^{-1} > 1$, holding all else constant, the climate policy jump risk premium is amplified, i.e.,

$$|\Theta'_\pi| > |\Theta_\pi|$$

whereas when $\psi^{-1} < 1$, all else constant, the climate policy jump risk premium is diminished, i.e.,

$$|\Theta'_\pi| < |\Theta_\pi|$$

This comparative static or partial equilibrium analysis highlights the role of the EIS. Consistent with the asset pricing literature, an EIS greater than one leads to increased concern about the resolution of uncertainty and amplifies the magnitude of the risk price of the climate policy jump. However, an EIS less than one leads to decreased concern about the resolution of uncertainty and diminishes the magnitude of the risk price of the climate policy jump. Such an analysis highlights the value of using asset prices in analyzing the impact of climate change and climate policy, and provides insight for

the expected macroeconomic outcomes, where a larger EIS should amplify the run on oil we would expect and a smaller EIS should diminish this effect.

11.5 Reserve Dependent Policy Arrival Rate

Another extension that is important to consider is the setting where the arrival rate of the policy shock is reserve dependent, i.e., $\lambda = \lambda(T_t, R_t)$. Furthermore, I assume that $\frac{\partial \lambda}{\partial R_t} < 0$, consistent with the realized policy outcomes such as increased climate policy compliance for low reserve countries and increase lobbying against climate policy for countries where oil reserves are more significant. France, for example, has proposed a number of substantial climate policy actions as a result of the Paris Climate Accord. These policies include banning oil extraction by 2040 and providing approximately €4 billion for investment in green technology innovation, as well as banning coal generated electricity by 2022 and petroleum and diesel vehicles by 2040. Yet, as France imports 99% of its oil due to the fact that it holds almost no oil reserves of its own, the policy to ban oil extraction and exploration is seen as largely symbolic (The Guardian, December 20, 2017, “France bans fracking and oil extraction in all of its territories”). This response is in direct contrast to the climate policy responses mentioned previously for the US and Norway, two countries with significant oil reserves.

The critical feature of this extension is that now decisions about extraction not only feedback in to temperature, but also into the level of reserves. This modification has two impacts. The first effect is that as reserves go down, there is further amplification of the increased discounting impact of the climate policy risk that would amplify the run on oil. The second effect is that oil firms have another lever by which to impact climate policy risk. By maintaining higher reserves, oil firms can actually help minimize the risk of the climate policy shock taking place. This effect would motivate decreased oil extraction and/or increased oil exploration. The effect is therefore ambiguous ex-ante, and as this setting has a direct, meaningful connection to real world policy settings such as lobbying and market power, important further work will go towards understanding the effect that this alternative specification has on the solution.

11.6 Multi-State Markov Chain

Another variant of the model to consider in terms of alternative policy is the setting where there are intermediate policy shocks, rather than simply a death shock to oil demand. Such policy may seem more implementable, as it requires less severe initial oil restrictions, while still building towards achieving the target of eventually restricting climate change so as not to exceed temperature ceiling thresholds such as the one proposed in the Paris Climate Accord. This equilibrium outcome is in the following proposition:

Proposition 12. *With uncertain climate policy where with M states ($\nu_i > \nu_j$ for $i < j$, $\nu_M = 0$), and where the arrival rate of policy is given by the temperature dependent function $\lambda_j(T_t)$, the intermediate*

and terminal value functions for the planner are given by:

$$V_j(K_t, R_t, T_t) = K_t^{c_1} v(R_t, T_t) \quad V_M(K_t, T_t) = \hat{c}_0 K_t^{c_1} \exp(c_3 T_t)$$

where investment and labor decisions are given by

$$\begin{aligned} I_{j,t} &= C_1 Y_t & L_{j,C} &= \bar{L}_j \\ I_{M,t} &= C_1 Y_t & L_{M,C} &= \bar{L}_M \end{aligned}$$

Exploration and extraction are given by

$$i_{j,R,t} = \left(\frac{v_{j,R} \Gamma \theta}{v_{j,R} - \varphi v_{j,T}} \right)^{1/(1-\theta)} \quad N_{j,t} = \frac{v_j \vartheta}{v_{j,R} - \varphi v_{j,T}} + i_{j,R,t} R_t$$

Note $v_j(R_t, T_t)$ is the solution to the simplified HJB equation characterizing the planner's state j problem (given in the appendix). The value function constants \hat{c}_0, c_1, c_3 and the FOC constants $\vartheta_j, C_1, \bar{L}_j, \bar{L}_M$ are functions of the model parameters only (also given in the appendix).

The difference here compared to the previous uncertain climate policy setting is that the changes in the value function due to uncertain policy are tempered, as are the changes in production and other outcomes of interest. This is seen with the additional, intermediate steps in the value of the demand, and thus additional PDEs to solve jointly as a system. However, the relevant mechanisms highlighted in the previous settings still exist. A run on oil production can still occur as the uncertain climate policy risk alters the discounting of agents and their expectations of future profits.

11.7 Optimal Tax Policy Shock

Another theoretical extension explores an alternative policy shock setting. Instead of a policy shock that shifts the production function, I assume in this setting that when the policy shock occurs the policy that takes place is the implementation of the optimal carbon tax. A carbon tax or carbon pricing is an oft-discussed policy tool, and one policy makers have tried to some degree to implement, as seen in figure 4. Furthermore, taxation is typically the theoretical tool used to achieve the socially optimal equilibrium when faced with an externality. As such, this extension is a particularly relevant framework to consider in terms of comparison with the type of policy I explore throughout the paper.

The structure of the ‘‘climate policy’’ shock in this cases will be similar spirit as before, in that it is meant to mimic the uncertain arrival of a policy trying to impose a temperature or carbon ceiling with a climate-dependent arrival rate of the policy instrument, here an optimal carbon tax. The sense in which agents are concerned about stranded assets is now that the cost of extracting oil for use becomes more costly after the policy arrival. Before the policy shock takes places, I assume that we are in the fully ‘‘decentralized’’ setting as outlined before where the climate externality is ignored

in the first order conditions for the planner. When the policy shock occurs, the economy jumps immediately to having in place the socially optimal carbon tax on oil extraction. Therefore ν , the energy input share of oil, is a fixed constant throughout. The resulting equilibrium for this model is as follows:

Proposition 13. *In the setting where uncertain climate policy is in the form of an optimal tax on oil extraction, the value functions for the two policy regimes are given by:*

$$V_{pre}(K_t, R_t, T_t) = K_t^{c_1} v(R_t, T_t) \quad V_{post}(K_t, R_t, T_t) = K_t^{c_1} \tilde{v}(R_t, T_t)$$

where investment and labor decisions are given by

$$\begin{aligned} I_{pre,t} &= C_1 Y_{pre,t} & L_{pre,C} &= \bar{L} \\ I_{post,t} &= C_1 Y_{post,t} & L_{post,C} &= \bar{L} \end{aligned}$$

In the pre-policy setting, extraction and exploration given by

$$N_t = \frac{v\vartheta}{v_R} \quad i_{R,t} = (\Gamma\theta)^{1/(1-\theta)}$$

In the post-policy setting, extraction and exploration given by

$$N_t = \frac{\tilde{v}\vartheta}{\tilde{v}_R - \tilde{v}_T\varphi} \quad i_{R,t} = \left(\frac{\Gamma\theta\tilde{v}_R}{\tilde{v}_R - \tilde{v}_T\varphi} \right)^{1/(1-\theta)}$$

Note $v(R_t, T_t)$ and $\tilde{v}(R_t, T_t)$ are the solutions to the simplified HJB equations characterizing the planner's pre- and post-policy problems. The value function constant c_1 and the FOC constants ϑ, C_1, \bar{L} are functions of the model parameters only.

While the structure of the solution is quite similar to before, the key difference is the change in how the different forces from uncertain policy identified previously impact oil production and exploration in this new setting. As before, the risk of stranded assets from climate policy still plays an important role in oil production and exploration decisions, and will feed in to the firm values on the asset pricing side. However, because the structure of the production function stays the same and the policy increases the cost of oil extraction through a tax-like policy, there is increased weight in this setting on the desire to avoid policy altogether. As the energy input share of demand for oil still remains high, but the cost of using the input is increased after the policy occurs, the cost of policy in this setting is even higher. Therefore, in numerical results we would expect to see an attenuation on the level of the run on oil as firms will initially want to avoid the policy shock as much as possible. However, because the policy shock still leads to a decrease in the marginal value of oil reserves in the future, we would expect some of the dynamic run on oil impact to occur for the same reasons as given in the baseline model.

12 Appendix C: Numerical Details

12.1 Numerical Solution Method

To solve the HJB equation for the social planner's problem, I use the Markov Chain approximation method developed by Kushner and Dupuis (2001). The key idea behind this method is to solve for the value function similar to value function iteration in discrete time by approximating the probabilities of state transitions appropriately. The state space is discretized, and then one-step transition probabilities are constructed using a Markov chain approximation such that the probabilities consistently approximate the underlying continuous time Brownian motions of the continuous time model. As long as the probabilities are constructed such that they satisfy the necessary conditions, this approximated solution approaches the true solution as the discretization gets small enough.

This method has shown to be quite stable and reliable for my given framework. Also, the method includes proofs of convergence as long as the probabilities constructed to derive the solutions satisfy simple and straightforward to verify conditions. For a full discussion on these methods, and for details on the verification conditions, see Kushner and Dupuis (2001), with examples of implementation found in Papanikolaou (2011) and Tourre et al. (2017). I start by construct the state transition probabilities. Let X_t be the state variable vector. The evolution of X_t , with vector Brownian motion B_t , is

$$dX_t = \mu_X(X_t)dt + \sigma_X(X_t)dB_t$$

Then the Markov chain approximating probabilities are given by

$$\begin{aligned} Pr(X_{n+1}^h = x + h | X_n^h = x) &= \frac{\sigma_X^2(x) + h \max\{0, \mu_X(x)\}}{Q^h(x)} \\ Pr(X_{n+1}^h = x - h | X_n^h = x) &= \frac{\sigma_X^2(x) + h \max\{0, -\mu_X(x)\}}{Q^h(x)} \\ Q^h(x) &= \sigma_X^2(x) + h|\mu_X(x)| \\ \Delta t^h(x) &= \frac{h^2}{Q^h(x)} \end{aligned}$$

where the h is the step-size for the state space discretization. Given these pieces, the value function is solved in an iterative fashion where

$$v^{(i,j+1)}(x) = U(C)\Delta t^h(x) + e^{-\rho\Delta t^h(x)} \times \sum_{x'} Pr\{x'|x\} \times v^{i,j}(x')$$

Convergence is found when $\|v^{(i,j+1)} - v^{(i,j)}\| \leq \epsilon$ for a given tolerance level ϵ .

Furthermore, because my framework generates a nonlinear HJB equation I incorporate an additional

step to ensure stability of the algorithm. For each $v^{(i,j)}$, I do an intermediate iterative loop of

$$v^{(i,j,m+1)}(x) = (U(C) + f(v^{i,j,m}(x'))) \Delta t^h(x) + e^{-\rho \Delta t^h(x)} \times \sum_{x'} Pr\{x'|x\} \times v^{i,j,m}(x')$$

where $f(\cdot)$ is a function of the nonlinear contributions of the value function to the HJB equation. This loop is done for a pre-specified number of iterations to smooth out the impact of the nonlinearities on the iterative convergence algorithm. After the inner loop finishes, the value function and probabilities are updated as explained previously, and then the inner loop is repeated.

Boundary conditions are as follows: For oil reserves R , I set the top and bottom boundary to be a constant derivative of the value function from R_{max-1} to R_{max} and R_{min+1} to R_{min} respectively. For temperature T , I impose a constant derivative of the value function from Y_{max-1} to Y_{max} for the top edge of the grid and from Y_{min+1} to Y_{min} for the bottom edge of the grid. I explore alternatives to test these boundary conditions and the results are stable for other conditions tried.

12.2 Parameter Restrictions

There are a few parameter restrictions required in order for the model to be well defined and for a solution to exist. First, under the assumption that $1 - \xi < 0$, which is a standard assumption for risk aversion and the assumption I use throughout this paper, we need $v < 0$ so that $\log[(1 - \xi)v]$ is defined. Second, I require a non-negative discount factor for the Markov chain approximation method, meaning the following must hold:

$$\begin{aligned} \tilde{\rho} &= -\rho(1 - \xi) \log(A_C L^\alpha (N - i_R R)^{\nu(1-\gamma-\alpha)} (A_G (1 - L)^\omega)^{(1-\nu)(1-\gamma-\alpha)} (1 - C_1)) \\ &\quad - b(\ln B + \delta_1 \ln C_1 A_C L^\alpha (N - i_R R)^{\nu(1-\gamma-\alpha)} (A_G (1 - L)^\omega)^{(1-\nu)(1-\gamma-\alpha)}) \\ &\quad - \frac{1}{2} \sigma_K^2 b(b - 1) + \lambda(T) + \eta(\rho(1 - \xi) + c_1 \delta_1) T > 0 \end{aligned}$$

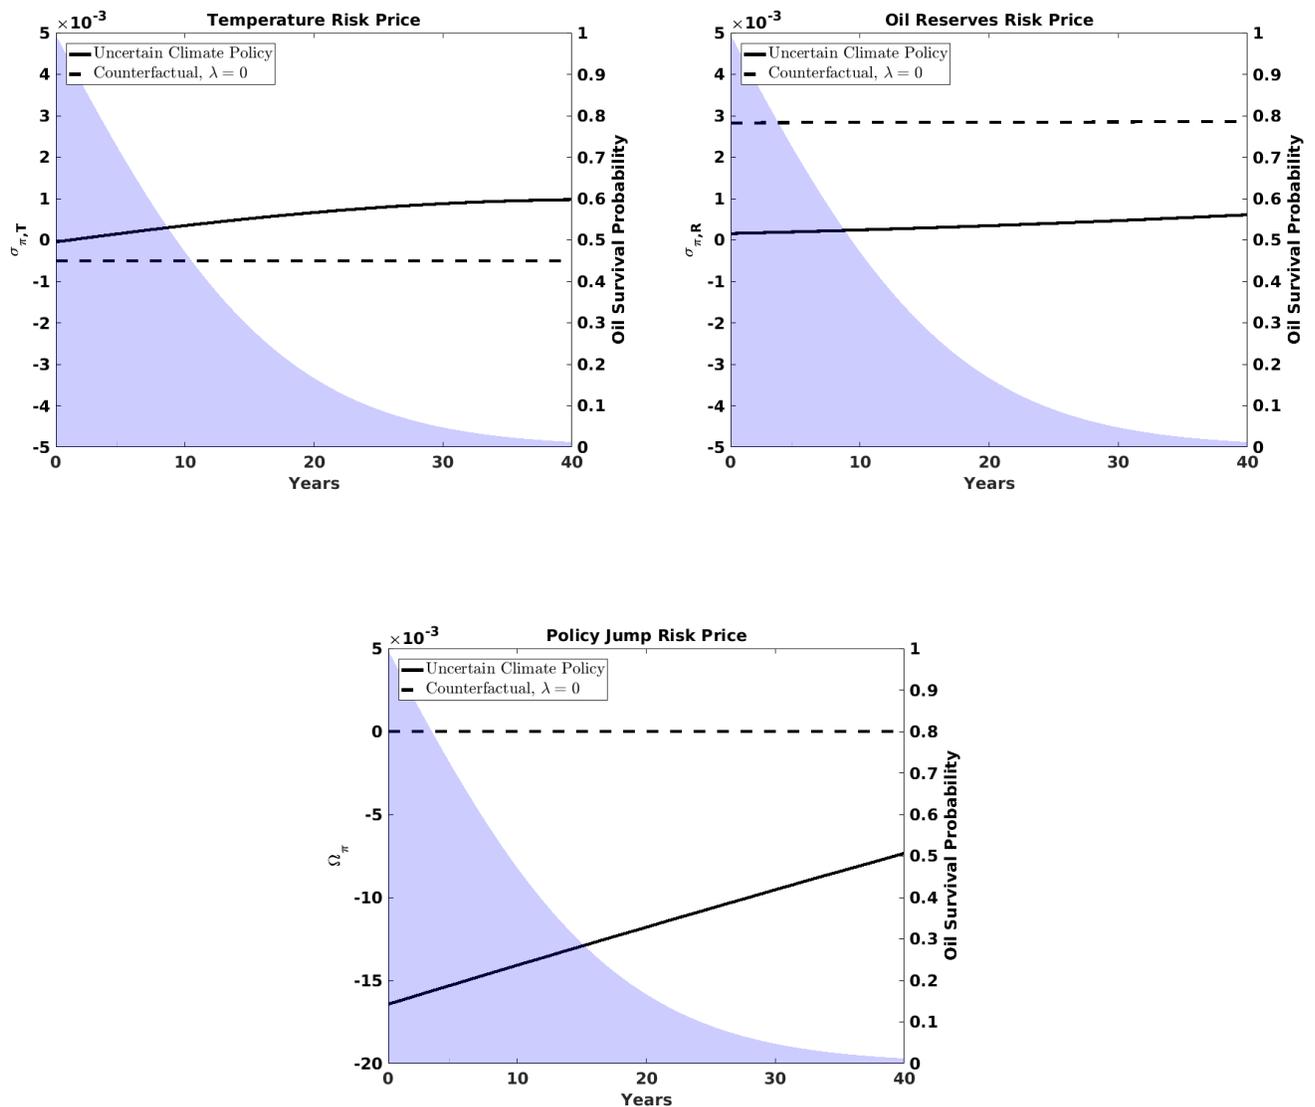
Further restrictions come in the form of the convergence properties for the solution method. The necessary requirement for guaranteed convergence is that the constructed Markov chain approximation probabilities satisfy the local consistency property. This means that

$$\begin{aligned} \mathbb{E}[\Delta R_{\mathbb{Q},n}^h] &= \mu_{\mathbb{Q}}(R, T) \Delta t_{\mathbb{Q}}^h(R, T) \\ var[\Delta R_{\mathbb{Q},n}^h] &= R^2 |\sigma_R|^2 \Delta t_{\mathbb{Q}}^h(R, T) + o(\Delta t_{\mathbb{Q}}^h(R, T)) \\ \mathbb{E}[\Delta T_{\mathbb{Q},n}^h] &= \mu_{\mathbb{Q}}(R, T) \Delta t_{\mathbb{Q}}^h(R, T) \\ var[\Delta T_{\mathbb{Q},n}^h] &= T^2 |\sigma_T|^2 \Delta t_{\mathbb{Q}}^h(R, T) + o(\Delta t_{\mathbb{Q}}^h(R, T)) \\ cov[\Delta R_{\mathbb{Q},n}^h, \Delta T_{\mathbb{Q},n}^h] &= 0 \end{aligned}$$

This is an easily verifiable condition and I choose parameters such that this condition holds.

13 Appendix D: Extended Asset Price Numerical Results

Here I provide the numerical results from the averaged model simulation times series outcomes for the market prices of risk, corresponding to the two main cases highlighted in numerical results section of the paper. The temperature risk price is the volatility loading for the SDF corresponding to temperature shocks, the oil reserves risk price is the volatility loading for the SDF corresponding to oil reserves shocks, and the policy jump risk price is the jump loading for the SDF corresponding to the climate policy jump shock.



14 Appendix E: Numerical Results for Model Extensions

14.1 Comparison with Oil-Sector Only Policy

Figure 13: Uncertain Policy Arrival Comparison - Quantities

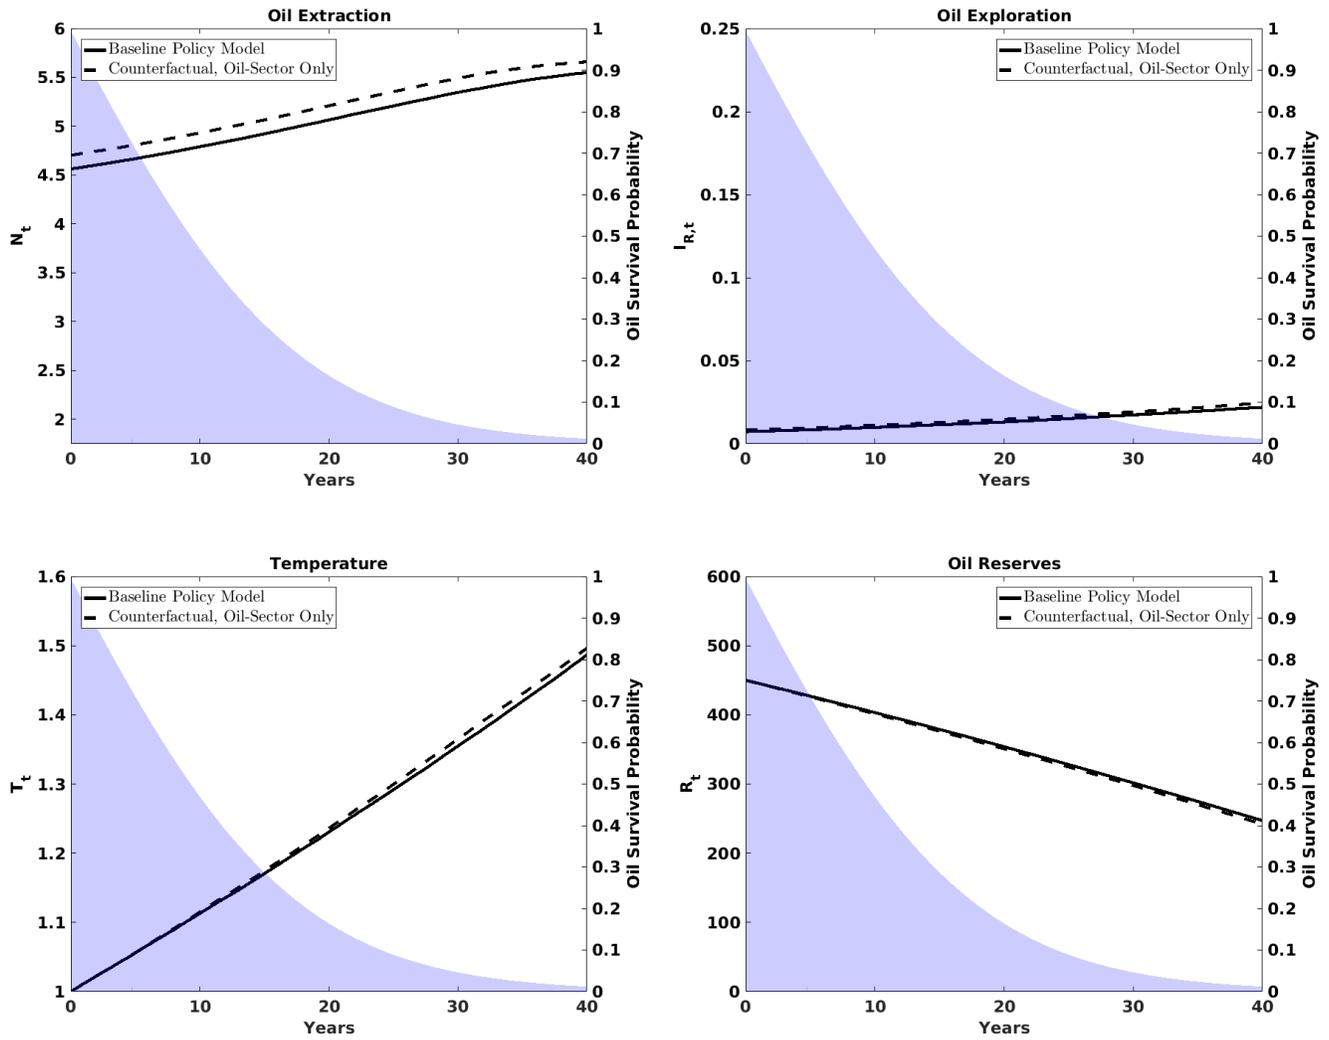
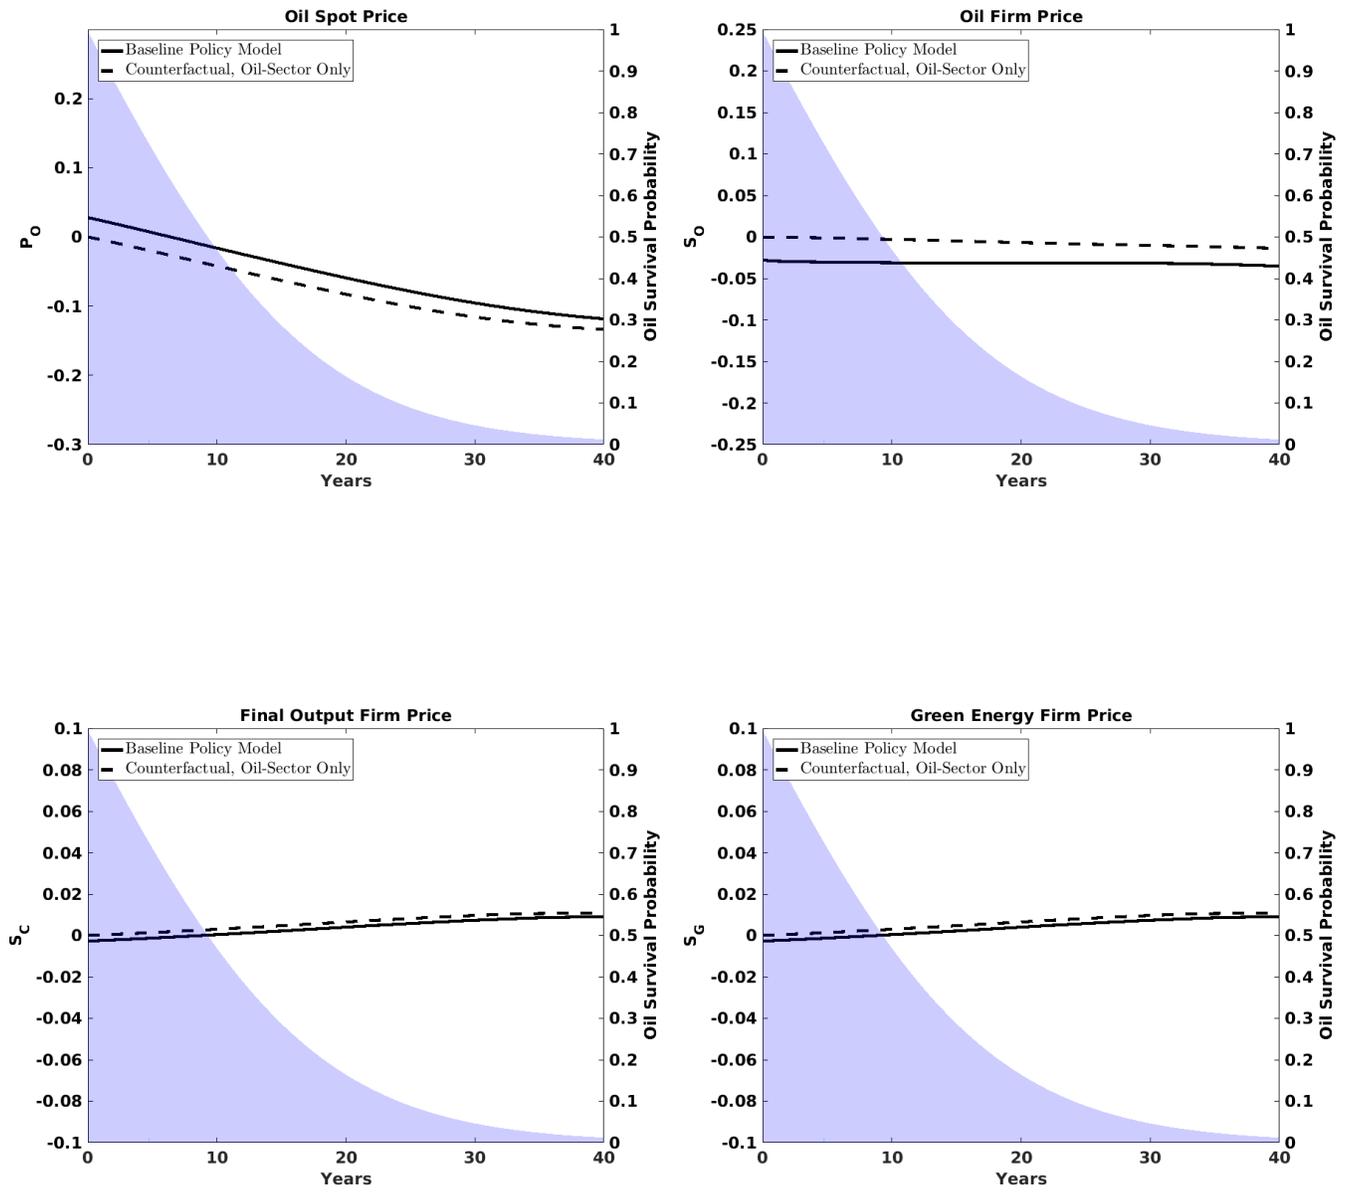


Figure 14: Uncertain Policy Arrival Comparison - Prices



14.2 Comparison with No Exploration Case

Figure 15: Uncertain Policy Arrival Comparison - Quantities

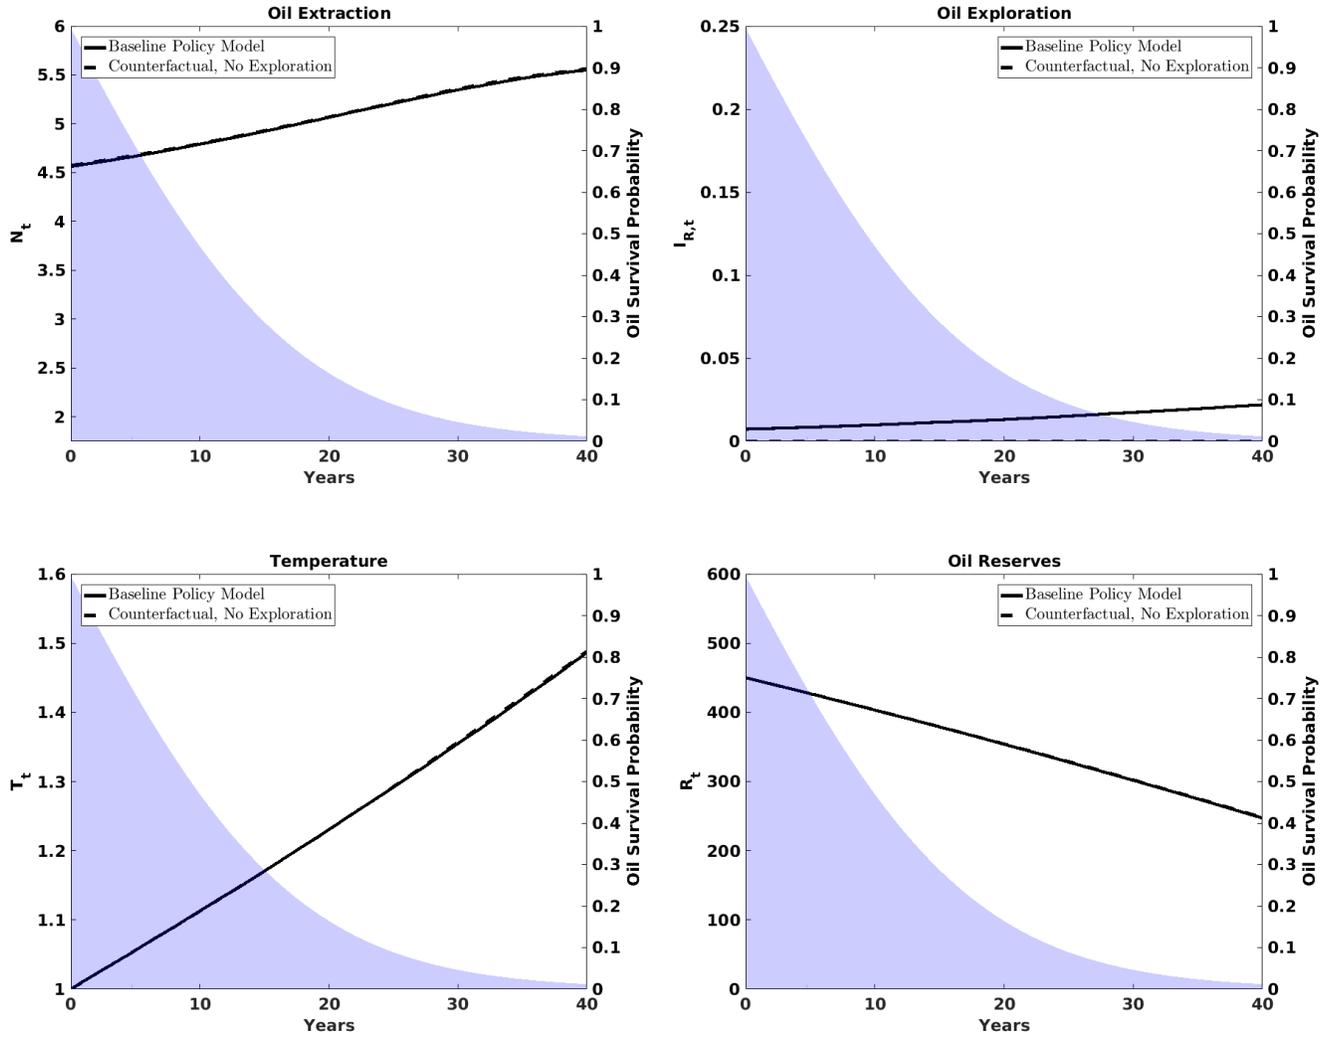
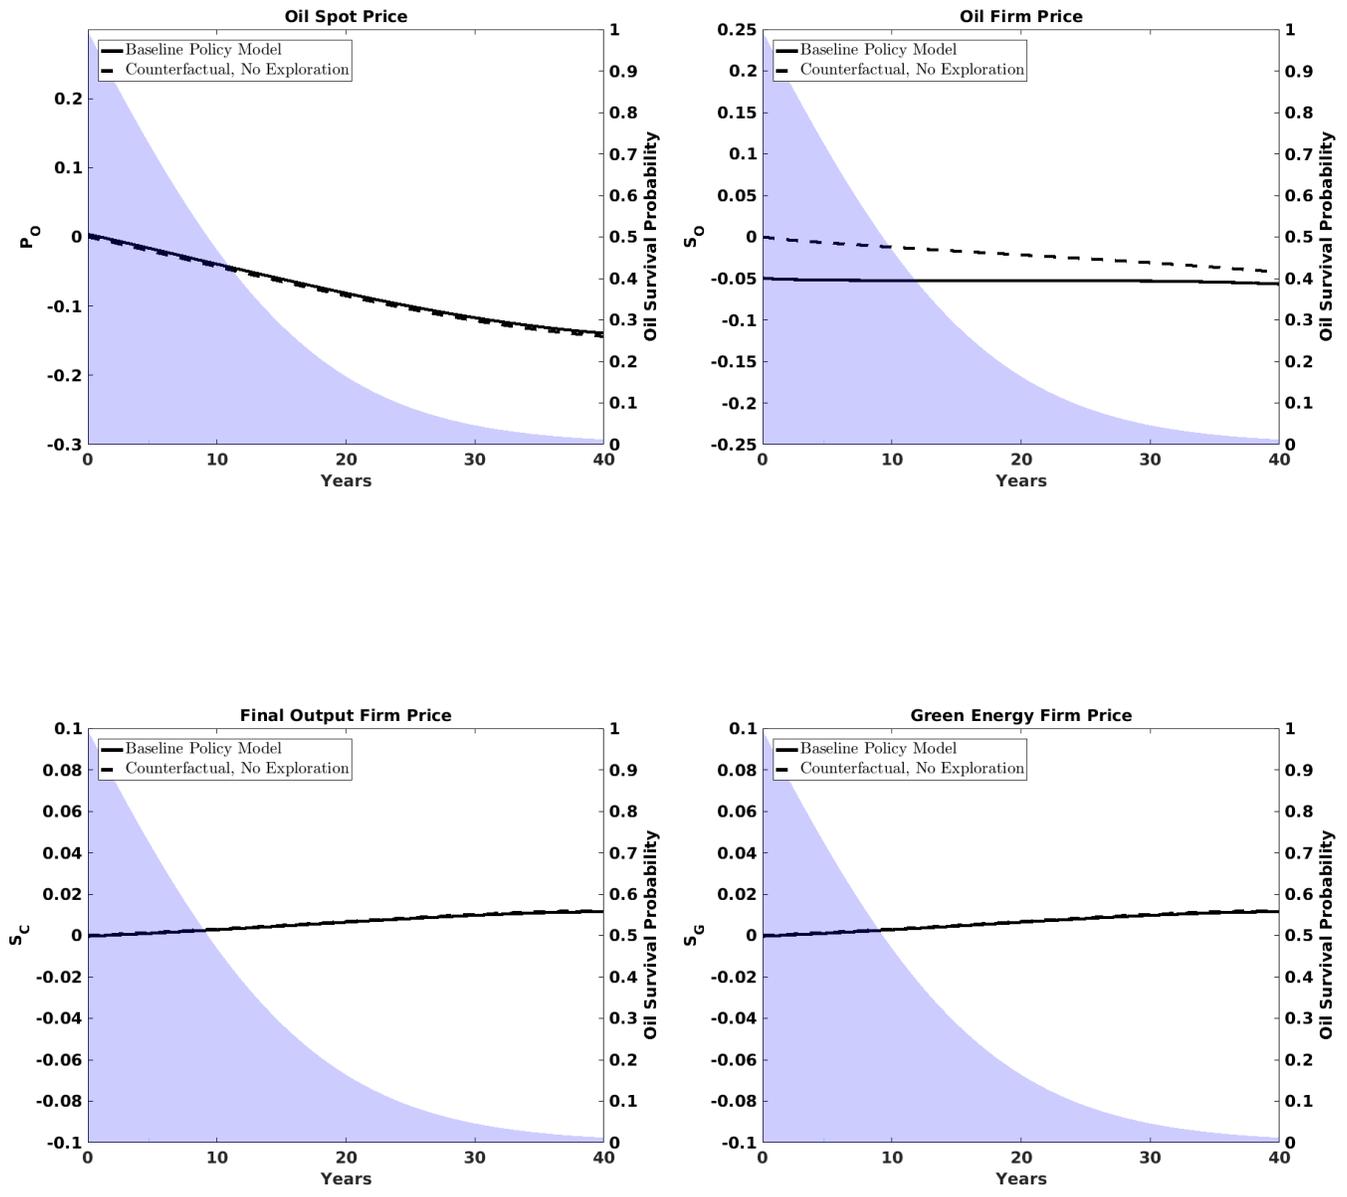


Figure 16: Uncertain Policy Arrival Comparison - Prices



14.3 Policy Welfare Comparison for Model Extensions

Figure 17: Oil-Sector Only Policy Comparison

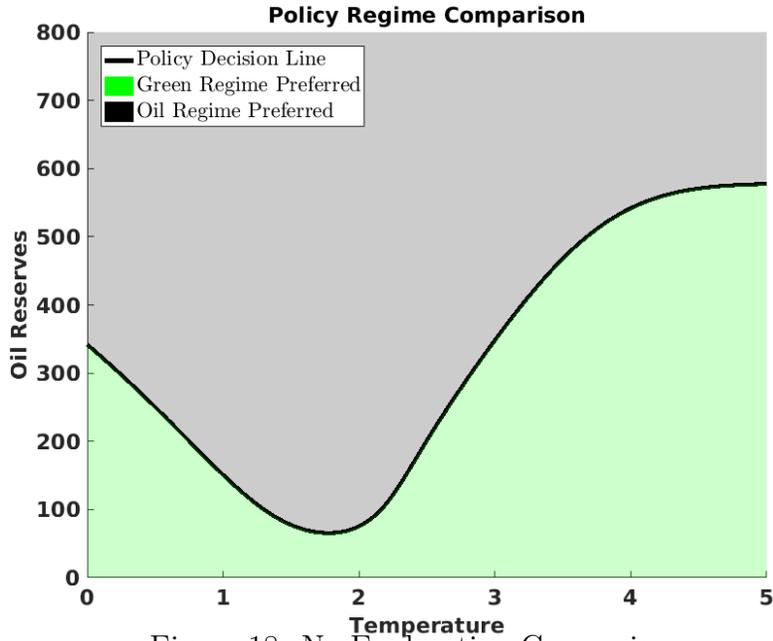
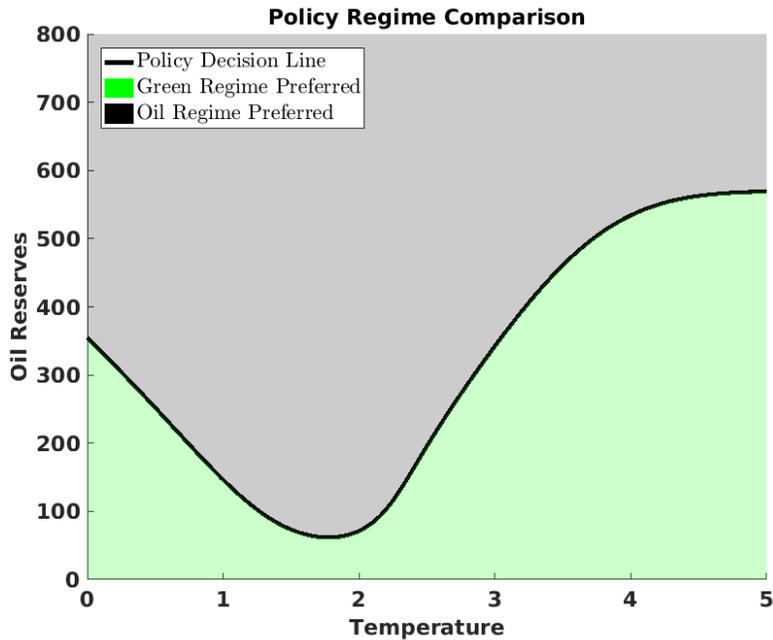


Figure 18: No Exploration Comparison



15 Appendix F: Climate Policy Index Details

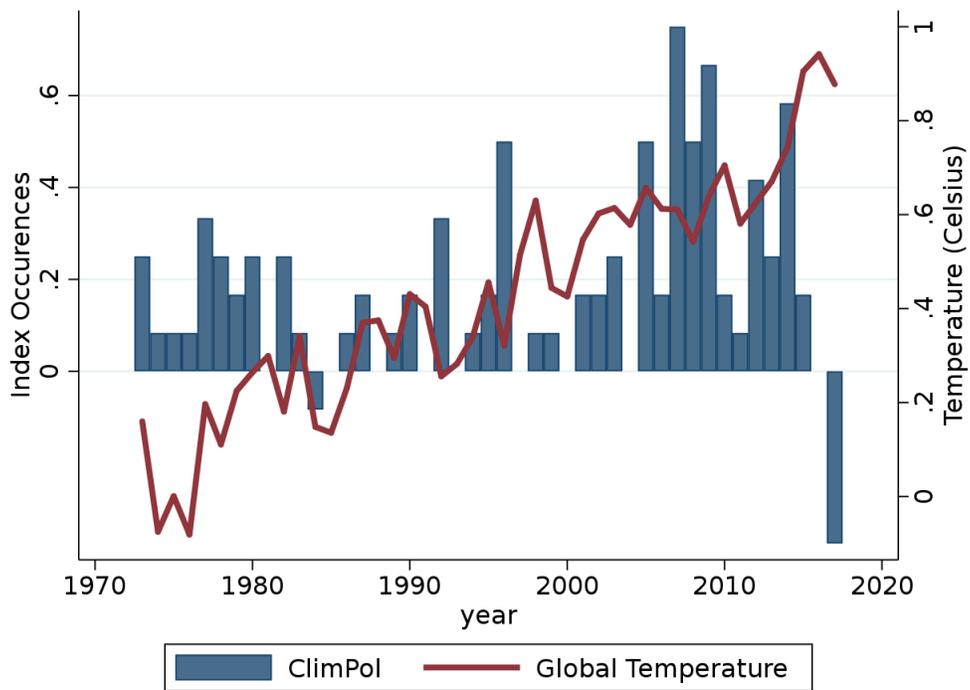
Figure 19: Climate Policy Index List, 1996-Present

Date	Event	Shock Sign	Source
5-Jun-96	Solar Two Plant Demonstrates Low Cost Method of Storing Solar Energy	+	ProCon.org
18-Jul-96	COP 2, Geneva, Switzerland	+	IPCC
9-Oct-96	Hydrogen Future Act of 1996 Is Passed to Further Expand Hydrogen Power Development	+	ProCon.org
29-Oct-96	European Union adopts target of a maximum 2 °C rise in average global temperature	+	Wikipedia
5-Nov-96	Bill Clinton Elected POTUS	+	U.S. Presidential Elections
5-Dec-96	EV1 Electric Car Is Made Available to the Public For Lease; Lease Program and EV1 Later Dismantled by GM	+	ProCon.org
25-Jun-97	US Senate passes Byrd–Hagel Resolution rejecting Kyoto	-	Wikipedia
11-Dec-97	COP 3, The Kyoto Protocol on Climate Change	+	Wikipedia/IPCC
14-Nov-98	COP 4, Buenos Aires, Argentina	+	IPCC
5-Nov-99	COP 5, Bonn, Germany	+	IPCC
7-Nov-00	George W. Bush Elected POTUS	-	U.S. Presidential Elections
25-Nov-00	COP 6, The Hague, Netherlands	+	IPCC
28-Mar-01	President George W. Bush withdraws from the Kyoto negotiations	-	Wikipedia
27-Jul-01	COP 6, Bonn, Germany	+	IPCC
29-Sep-01	IPCC Third assessment report	+	IPCC
10-Nov-01	COP 7, Marrakech, Morocco	+	IPCC
13-May-02	Farm Security and Rural Investment Act	+	Wikipedia
1-Nov-02	COP 8, New Delhi, India	+	IPCC
6-Feb-03	President Bush Unveils the Hydrogen Fuel Initiative	+	ProCon.org
27-Feb-03	Plans Announced to Build Worlds First Zero Emissions Coal Power Plant	+	ProCon.org
12-Dec-03	COP 9, Milan, Italy	+	IPCC
2-Nov-04	George W. Bush Elected POTUS	-	U.S. Presidential Elections
17-Dec-04	COP 10, Buenos Aires, Argentina	+	IPCC
1-Jan-05	EU Emissions Trading Scheme is launched, the first such scheme	+	Wikipedia
16-Feb-05	Kyoto Protocol comes into force (not including the US or Australia)	+	Wikipedia
8-Jul-05	31st G8 summit discusses climate change, relatively little progress made	+	Wikipedia
8-Aug-05	Energy Policy Act	+	Wikipedia
9-Nov-05	US House Prevents Drilling for Oil in the Arctic National Wildlife Refuge	+	ProCon.org
9-Dec-05	COP 11/CMP 1, Montreal, Canada	+	Wikipedia/IPCC
30-Oct-06	The Stern Review is published	+	Wikipedia
17-Nov-06	COP 12/CMP 2, Nairobi, Kenya	+	IPCC
16-Feb-07	February 2007 Washington Declaration	+	IPCC
7-Jun-07	33rd G8 summit	+	IPCC
31-Jul-07	2007 UN General Assembly plenary debate	+	IPCC
3-Aug-07	September 2007 Washington conference	+	IPCC
31-Aug-07	2007 Vienna Climate Change Talks and Agreement	+	IPCC
24-Sep-07	September 2007 United Nations High-Level-Event	+	IPCC
17-Nov-07	IPCC Fourth assessment report	+	IPCC/ProCon.org
17-Dec-07	COP 13/CMP 3, Bali, Indonesia	+	IPCC
19-Dec-07	Energy Independence and Security Act	+	Wikipedia
30-Jan-08	First Commercial Cellulosic Ethanol Plant Goes Into Production	+	ProCon.org
22-May-08	Food, Conservation, and Energy Act	+	Wikipedia
7-Oct-08	National Biofuel Action Plan Unveiled	+	ProCon.org
4-Nov-08	Barack Obama Elected POTUS	+	U.S. Presidential Elections

12-Dec-08	COP 14/CMP 4, Poznań, Poland	+	IPCC
22-Dec-08	Worst Coal Ash Spill in US History in Kingston, Tennessee	+	ProCon.org
17-Feb-09	ARRA (2009) Contains Funding for Renewable Energy	+	ProCon.org/Wikipedia
22-Apr-09	First Framework for Wind Energy Development on the US Outer Continental Shelf Announced	+	ProCon.org
5-May-09	President Obama Issues Presidential Directive to USDA to Expand Access to Biofuels	+	ProCon.org
27-May-09	US Announces Funding in Recovery Act Funding for Solar and Geothermal Energy Development	+	ProCon.org
26-Jun-09	US House of Representatives passes the American Clean Energy and Security Act (Waxman)	+	Wikipedia
22-Sep-09	September 2009 United Nations Secretary General's Summit on Climate Change	+	IPCC
27-Oct-09	US Invests \$3.4 Billion to Modernize Energy Grid	+	ProCon.org
18-Dec-09	COP 15/CMP 5, Copenhagen, Denmark	+	IPCC
20-Apr-10	BP Oil Rig Explodes & Causes Largest Oil Spill in US History	+	ProCon.org
10-Dec-10	COP 16/CMP 6, Cancun, Mexico	+	IPCC
11-Mar-11	Earthquake off Coast of Japan Damages Six Powerplants at Fukushima	+	ProCon.org
1-Sep-11	Solar Power Company Solyndra Declares Bankruptcy	-	ProCon.org
9-Dec-11	COP 17/CMP 7, Durban, South Africa	+	IPCC
9-Feb-12	US Nuclear Regulatory Commission (NRC) Approves New Nuclear Power Plants	+	ProCon.org
27-Mar-12	EPA Announces First Clean Air Act Standard for Carbon Pollution from New Power Plants	+	ProCon.org
17-Apr-12	EPA Issues First Ever Clean Air Rules for Natural Gas Produced by Fracking	+	ProCon.org
6-Nov-12	Barack Obama Elected POTUS	+	U.S. Presidential Elections
7-Dec-12	COP 18/CMP 8, Doha, Qatar	+	IPCC
25-Jun-13	President Obama Releases His Climate Action Plan	+	ProCon.org
20-Sep-13	EPA Issues New Proposed Rule to Cut Greenhouse Gas Emissions from Power Plants	+	ProCon.org
23-Nov-13	COP 19/CMP 9, Warsaw, Poland	+	IPCC
13-Feb-14	Ivanpah, the World's Largest Concentrated Solar Power Generation Plant, Goes Online	+	ProCon.org
9-May-14	President Obama Announces Solar Power Commitments and Executive Actions	+	ProCon.org
2-Jun-14	EPA Proposes First Ever Rules to Reduce Carbon Emissions from Existing Power Plants	+	ProCon.org
22-Sep-14	Rockefellers and over 800 Global Investors Announce Fossil Fuel Divestment	+	ProCon.org
23-Sep-14	Climate Summit 2014	+	IPCC
1-Nov-14	IPCC Fifth assessment report	+	IPCC
12-Dec-14	COP 20/CMP 10, Lima, Peru	+	IPCC
3-Aug-15	President Obama Announces Clean Power Plan (finalized Oct. 23, 2015; Active December, 22, 2015)	+	ProCon.org
12-Dec-15	COP 21/CMP 11, Paris, France	+	Wikipedia/IPCC
8-Nov-16	Donald Trump Elected POTUS	-	U.S. Presidential Elections
18-Nov-16	COP 22/CMP 12/CMA 1, Marrakech, Morocco	+	IPCC
28-Mar-17	President Trump Signs Executive Order to Begin Reversal of President Obama's Clean Power Plan	-	ProCon.org
1-Jun-17	President Donald Trump withdraws the United States from the Paris Agreement	-	Wikipedia
31-Jul-17	Two Nuclear Power Reactors in South Carolina Abandoned Before Construction Completed	-	ProCon.org
22-Dec-17	Tax Bill Opens Arctic National Wildlife Refuge for Oil Drilling	-	ProCon.org
9-May-18	Solar Power to Be Required on All New California Homes by 2020	+	ProCon.org

*Events come from ProCon.org Fossil Fuel and Alternative Energy timeline, IPCC/UNFCCC Meetings, U.S. Presidential Election outcomes, and Wikipedia.org Selective historical timeline of significant climate change political events and List of United States energy acts

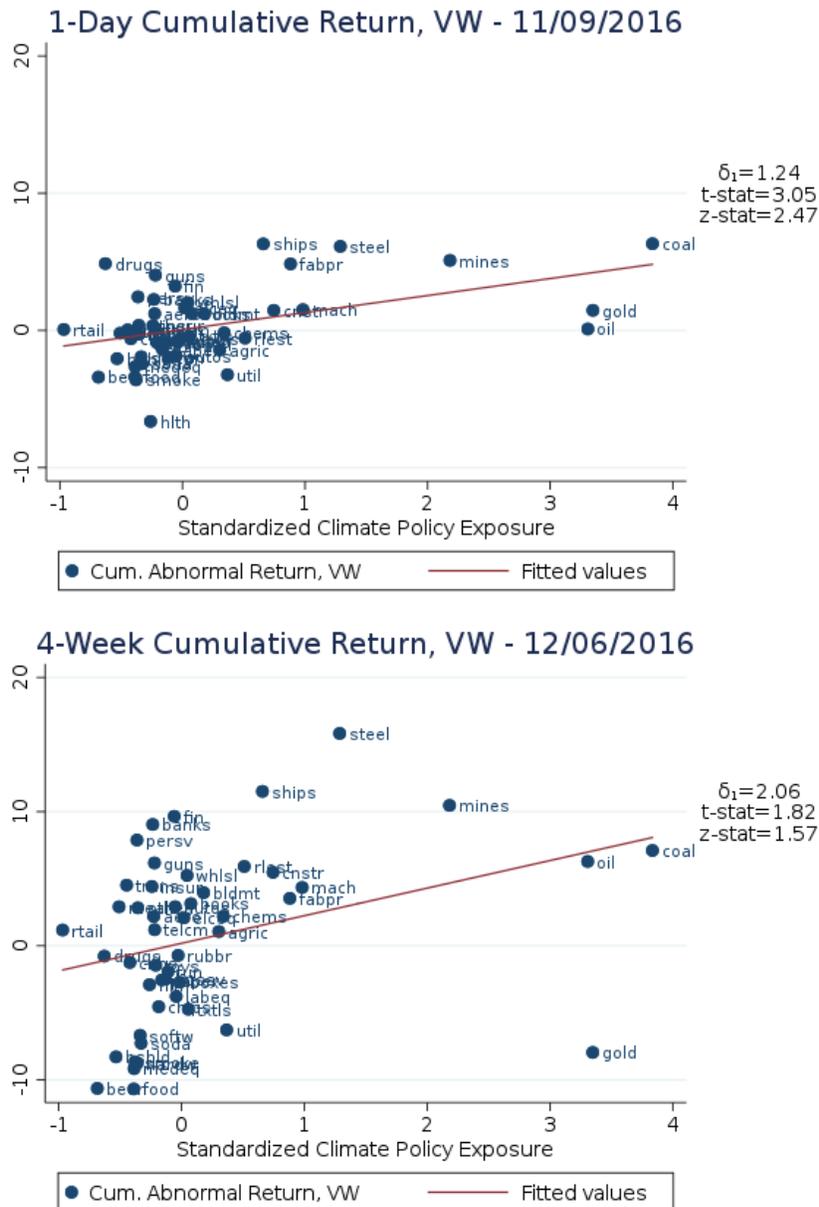
Figure 20: ClimPol Index and Annual Global Mean Temperature



16 Appendix G: Additional Empirical Results

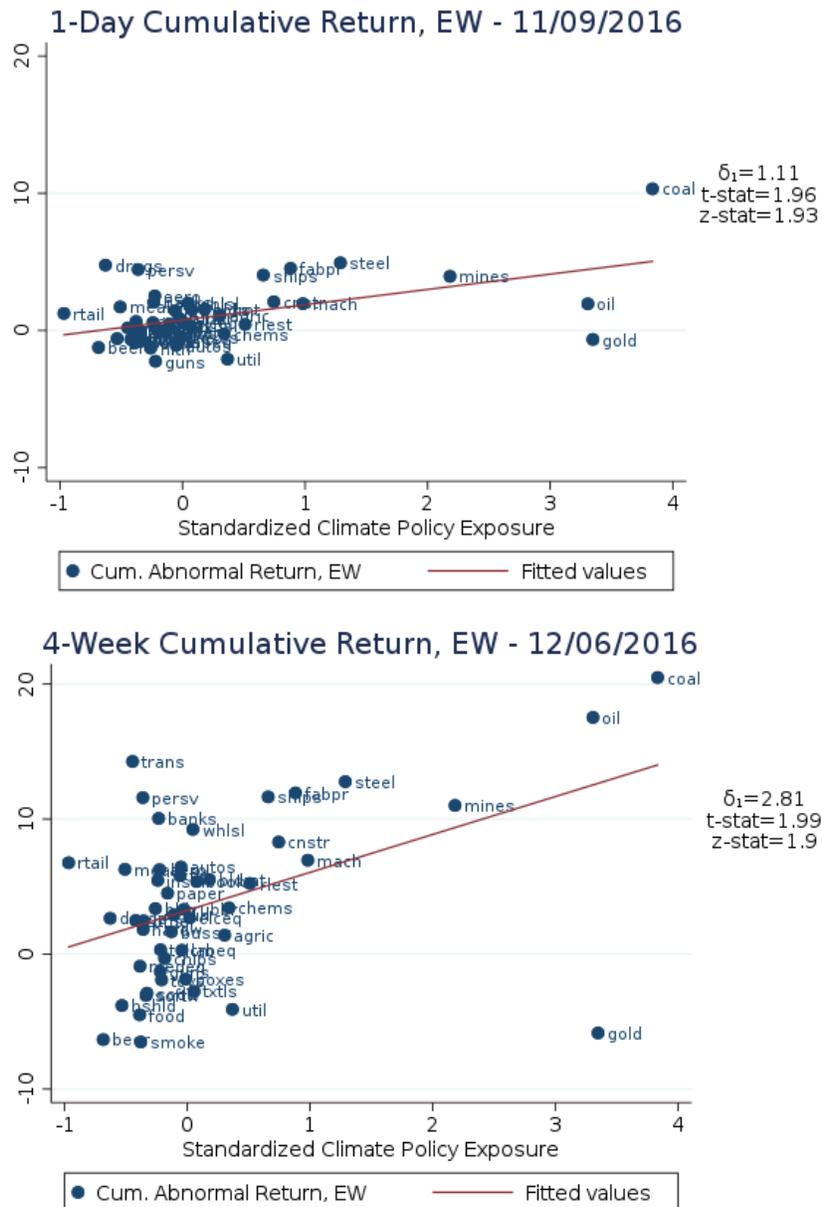
16.1 Election Event Study Results

Figure 21: Election Impact on Returns by Climate Policy Exposure - Value-Weighted



These figures show the relationship between the cumulative abnormal returns of sectors after the 2016 US presidential election and their standardized exposure to climate policy risk. The regression specification is given by $R_{i,t} = \alpha_i + \beta_{i,Mkt}(R_{Mkt,t} - R_{f,t}) + \beta_{i,OilPrice}R_{OilPrice,t} + \varepsilon_{i,t}$. Cumulative abnormal returns are normalized to zero at the election date, and are estimated with respect to the value-weighted excess market return. The top panel are estimates for cumulative abnormal returns one day after the election, and the bottom panel are estimates for cumulative abnormal returns four weeks after the election. See text for full definition of variables.

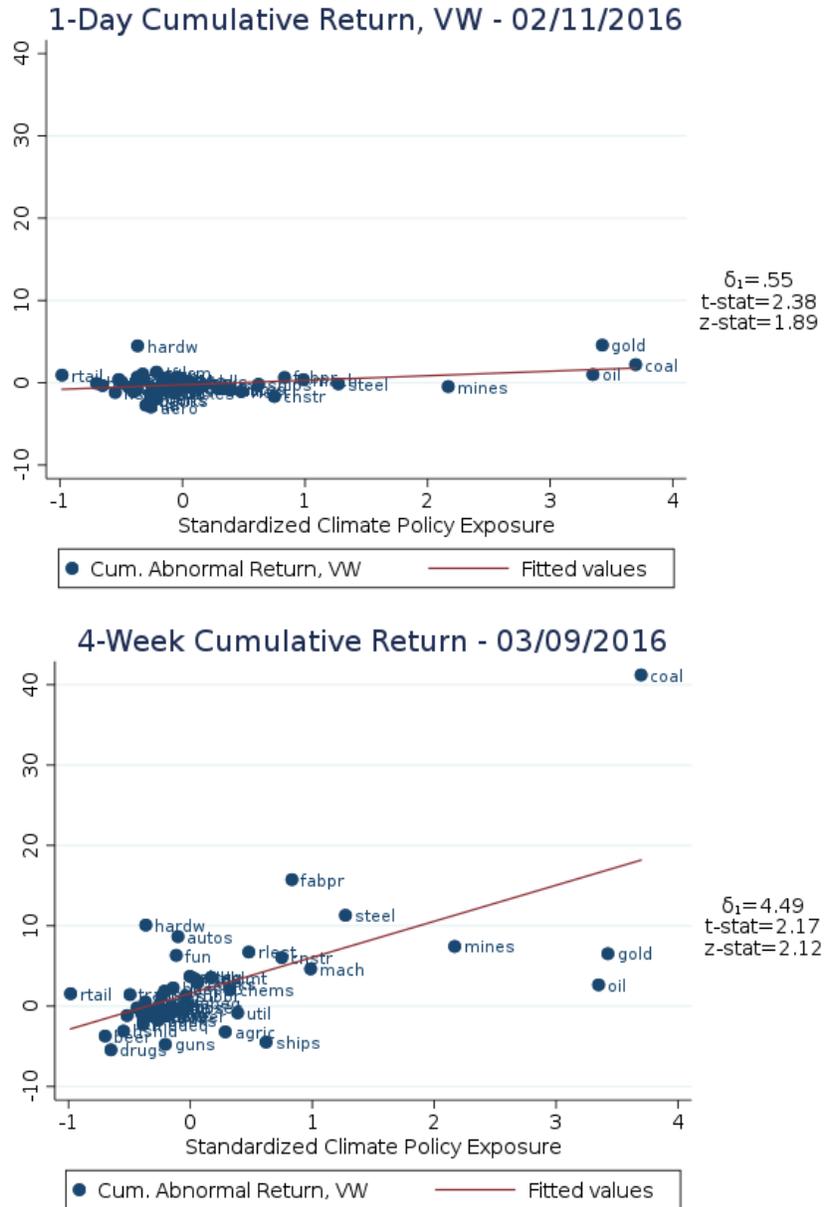
Figure 22: Election Impact on Returns by Climate Policy Exposure - Equal-Weighted



These plots show the relationship between the cumulative abnormal returns of sectors after the 2016 US presidential election and their standardized exposure to climate policy risk. The regression specification is given by $R_{i,t} = \alpha_i + \beta_{i,Mkt}(R_{Mkt,t} - R_{f,t}) + \beta_{i,OilPrice}R_{OilPrice,t} + \varepsilon_{i,t}$. Cumulative abnormal returns are normalized to zero at the election date, and are estimated with respect to the value-weighted excess market return. The top panel are estimates for cumulative abnormal returns one day after the election, and the bottom panel are estimates for cumulative abnormal returns four weeks after the election. See text for full definition of variables.

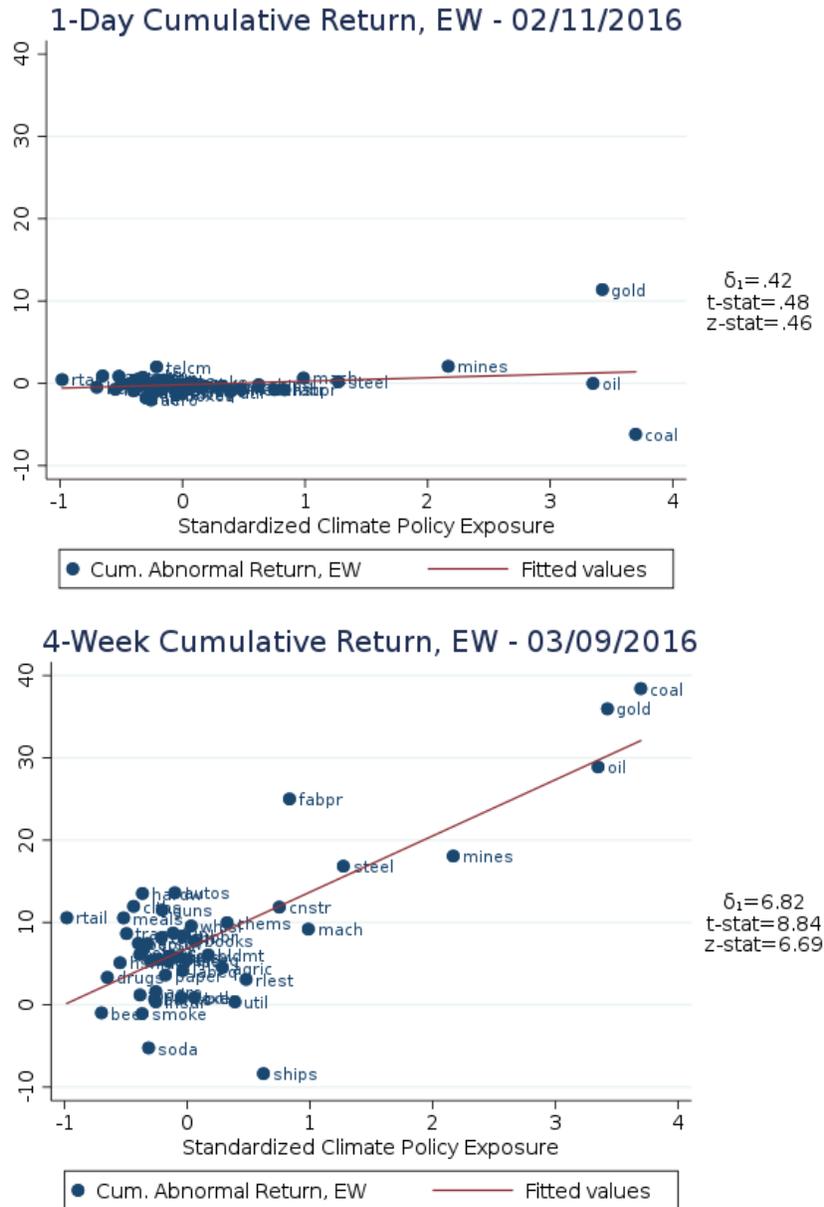
16.2 Supreme Court Event Study Results

Figure 23: Court Impact on Returns by Climate Policy Exposure - Value-Weighted



These plots show the relationship between the cumulative abnormal returns of sectors after the 2016 US Supreme Court decision to put a stay on the Clean Power Plan and their standardized exposure to climate policy risk. The regression specification is given by $R_{i,t} = \alpha_i + \beta_{i,Mkt}(R_{Mkt,t} - R_{f,t}) + \beta_{i,OilPrice}R_{OilPrice,t} + \varepsilon_{i,t}$. Cumulative abnormal returns are normalized to zero at the election date, and are estimated with respect to the value-weighted excess market return. The top panel are estimates for cumulative abnormal returns one day after the election, and the bottom panel are estimates for cumulative abnormal returns four weeks after the election. See text for full definition of variables.

Figure 24: Court Impact on Returns by Climate Policy Exposure - Equal-Weighted



These plots show the relationship between the cumulative abnormal returns of sectors after the 2016 US Supreme Court decision to put a stay on the Clean Power Plan and their standardized exposure to climate policy risk. The regression specification is given by $R_{i,t} = \alpha_i + \beta_{i,Mkt}(R_{Mkt,t} - R_{f,t}) + \beta_{i,OilPrice}R_{OilPrice,t} + \varepsilon_{i,t}$. Cumulative abnormal returns are normalized to zero at the election date, and are estimated with respect to the value-weighted excess market return. The top panel are estimates for cumulative abnormal returns one day after the election, and the bottom panel are estimates for cumulative abnormal returns four weeks after the election. See text for full definition of variables.

16.3 Interaction Estimate of Climate Policy Index

Table 9: Climate Policy Impact on Oil Production (1973-2017)

	OPEC	US	Non-OPEC	World
ClimPol*Temp	0.061	0.053	0.200	0.236
S.E.	(0.095)	(0.023)	(0.070)	(0.117)
# Obs.	536	536	536	536
R^2	0.995	0.979	0.991	0.988

Table 10: Climate Policy Impact on Oil Production (1996-2017)

	OPEC	US	Non-OPEC	World
ClimPol*Temp	0.035	0.062	0.210	0.264
S.E.	(0.055)	(0.024)	(0.075)	(0.116)
# Obs.	260	260	260	260
R^2	0.980	0.973	0.988	0.987

These tables show the impact of climate policy events as measured by the *ClimPol* index on oil production for the Non-OPEC, OPEC, US, and World regions. The top table are estimates using the full time sample of data (1973-2017), and the bottom table are estimates using the policy-relevant time subsample (1996-2017). The regression specification is given by $Y_t = \alpha + \beta_Y Y_{t-1} + \phi_{ClimPol*Temp} ClimPol_t * Temp_t + \epsilon_t$. I omit the constant and lag variable coefficients from the table. See text for full definition of variables.

Table 11: Climate Policy Impact on Oil Sector Returns (1973-2017)

	$r_{oil,t+1}$	$r_{oil,t+6}$	$r_{oil,t+12}$	$r_{oil,t+18}$	$r_{oil,t+24}$
ClimPol*Temp	-0.003	-0.005	-0.058	-0.119	-0.131
S.E.	(0.010)	(0.023)	(0.037)	(0.048)	(0.060)
# Obs.	534	529	523	517	511
R^2	0.004	0.016	0.019	0.037	0.023

Table 12: Climate Policy Impact on Oil Sector Returns (1996-2017)

	$r_{oil,t+1}$	$r_{oil,t+6}$	$r_{oil,t+12}$	$r_{oil,t+18}$	$r_{oil,t+24}$
ClimPol*Temp	-0.003	-0.015	-0.077	-0.153	-0.177
S.E.	(0.011)	(0.025)	(0.040)	(0.051)	(0.063)
# Obs.	260	260	260	260	260
R^2	0.021	0.019	0.023	0.059	0.052

These tables show the impact of climate policy events as measured by the *ClimPol* index on returns for the value-weighted US Oil sector portfolio. The regression specification is given by $r_{t+1,t+h}^e = a + bX_t + c_{ClimPol*Temp}ClimPol_t * Temp_t + \varepsilon_t$. $r_{i,t+1,t+h}$ is the k-month cumulative return for the value-weighted US Oil sector portfolio. X_t is a vector of control variables that includes the lagged values for the value-weighted market portfolio return, the value-weighted oil sector portfolio return, lagged oil production innovations, lagged oil spot price returns, and lagged real economic activity innovations. The top panel are estimates using the full time sample of data (1973-2017), and the bottom panel are estimates using the policy-relevant time subsample (1996-2017). I omit the constant and control coefficients from the table. See text for full definition of variables.

Table 13: Climate Policy Impact on Oil Price Returns (1973-2017)

	$r_{oil,t+1}$	$r_{oil,t+6}$	$r_{oil,t+12}$	$r_{oil,t+18}$	$r_{oil,t+24}$
ClimPol*Temp	-0.009	0.040	0.010	-0.071	-0.083
S.E.	(0.017)	(0.052)	(0.075)	(0.081)	(0.088)
# Obs.	534	529	523	517	511
R^2	0.051	0.018	0.018	0.018	0.023

Table 14: Climate Policy Impact on Oil Price Returns (1996-2017)

	$r_{oil,t+1}$	$r_{oil,t+6}$	$r_{oil,t+12}$	$r_{oil,t+18}$	$r_{oil,t+24}$
ClimPol*Temp	-0.009	0.028	-0.013	-0.145	-0.176
S.E.	(0.017)	(0.028)	(0.078)	(0.079)	(0.081)
# Obs.	260	260	260	260	260
R^2	0.017	0.026	0.030	0.046	0.048

These tables show the impact of climate policy events as measured by the *ClimPol* index on returns for the WTI spot price of oil. The regression specification is given by $r_{t+1,t+h}^{spot} = a + bX_t + c_{ClimPol*Temp}ClimPol_t * Temp_t + \varepsilon_t$ is the k-month cumulative return for the WTI spot price of oil. X_t is a vector of control variables that includes the lagged values for the value-weighted market portfolio return, the value-weighted oil sector portfolio return, lagged oil production innovations, lagged oil spot price returns, and lagged real economic activity innovations. The top panel are estimates using the full time sample of data (1973-2017), and the bottom panel are estimates using the policy-relevant time subsample (1996-2017). I omit the constant and control coefficients from the table. See text for full definition of variables.