

Geopolitical Oil Price Risk and Economic Fluctuations^{*}

Lutz Kilian[†] Michael D. Plante[‡] Alexander W. Richter[§]

May 5, 2025

ABSTRACT

This paper studies the general equilibrium effects of time-varying geopolitical risk in the oil market by simultaneously modeling downside risk from disasters, oil storage, and the endogenous determination of oil price and macroeconomic uncertainty in the global economy. Notwithstanding the attention geopolitical events in oil markets have attracted, we find that geopolitical oil price risk has modest effects on the oil market and the macroeconomy. Even when allowing for the possibility of an unprecedented 20% drop in global oil production, it takes a large increase in the probability of such a disaster to cause a sizable recessionary impact.

Keywords: Geopolitical risk, macroeconomic risk, time-varying uncertainty, rare disasters, oil, endogeneity, shock propagation, economic fluctuations, precautionary savings, inventories

JEL Classifications: E13, E22, E32, Q43

^{*}We thank María Aristizábal-Ramírez and Michael Siemer for discussing our paper at the 2024 System Macro and System Energy meetings. We also thank Mario Crucini, Steven Davis, Paulina Restrepo-Echavarría, Zheng Liu, and numerous conference and seminar participants for helpful comments. This work was supported by computational resources provided by the BigTex High Performance Computing Group at the Federal Reserve Bank of Dallas. The views in this paper do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

[†]Federal Reserve Bank of Dallas, 2200 N Pearl Street, Dallas, TX 75201, and CEPR (lkilian2019@gmail.com).

[‡]Federal Reserve Bank of Dallas, 2200 N Pearl Street, Dallas, TX 75201 (michael.plante@dal.frb.org).

[§]Federal Reserve Bank of Dallas, 2200 N Pearl Street, Dallas, TX 75201 (alex.richter@dal.frb.org).

1 INTRODUCTION

Time-variation in geopolitical risk is widely considered an important determinant of fluctuations in economic activity. The financial press, international organizations, rating agencies and the investment community all vie to assess these risks and their impact on the economy. Clearly, major geopolitical disruptions matter not only when they occur on rare occasions, but also when investors and consumers make decisions in anticipation of the possibility of such events. This fact is nowhere more apparent than when it comes to geopolitical risk in energy markets. For example, many market analysts listed risks to energy supplies as one of the top geopolitical risks in 2024. This assessment was driven in no small part by concerns about OPEC quota decisions, global access to Russian oil amidst Ukrainian attacks on Russian oil infrastructure and efforts to tighten the G7 price cap, dwindling strategic oil reserves, disruptions of oil shipments in the Red Sea and possibly in the Persian Gulf, and concerns about a widening conflict between Israel and Iran.

There is a deep-rooted belief in macroeconomics that higher oil price uncertainty driven by geopolitical risk lowers investment and consumption and hence real GDP.¹ The focus of our paper is to develop a deeper understanding of how time-varying geopolitical risk in oil markets affects oil price uncertainty and economic fluctuations. We introduce a stochastic growth model of the global economy in which geopolitical risk in the oil market arises from rare oil production disasters. The model allows us for the first time to assess the effects of an increase in the probability of a geopolitical oil supply disruption, addressing a key question faced by policymakers and market participants.

Our model includes risk averse economic agents, an oil production sector, oil storage, and limited substitutability between oil and capital. The price of oil is determined endogenously. Since the model is global, we abstract from oil imports and exports and international capital flows. Oil

¹For example, Bernanke (1983), Lee et al. (1995), Ferderer (1996), Edelstein and Kilian (2009), Elder and Serletis (2010), Baumeister and Kilian (2016a), Ready (2018), Gao et al. (2022), and Alfaro et al. (2024) discuss the impact of oil price uncertainty on U.S. real activity, while Kilian (2009), Jo (2014), and Cross et al. (2022) discuss its impact on global real activity. The perception that oil price volatility matters for the transmission of oil price shocks to the economy also helped spawn a large literature on the asymmetric transmission of oil price shocks (see, e.g., Bernanke et al., 1997; Davis and Haltiwanger, 2001; Hooker, 1996, 2002; Kilian and Vigfusson, 2011; Leduc and Sill, 2004; Lee and Ni, 2002; Mork, 1989; Ramey and Vine, 2010).


production disasters are of stochastic length and occur with a time-varying probability. These disasters are modeled after historical events such as the Arab-Israeli War in 1973, the Iranian Revolution in 1979, or the invasion of Kuwait in 1990. The size (a 5% drop in global oil production) and duration (3 quarters on average) of these disasters is set to match the behavior of global oil production during major geopolitical events in the oil market over the past 50 years.²

Our analysis recognizes that, while downside geopolitical risk to oil production raises oil price uncertainty, not all surges in oil price uncertainty are driven by geopolitical events. In particular, a major downturn in the economy or simply the possibility of such a downturn may also cause surges in oil price uncertainty. Thus, the model allows for the endogenous determination of both macroeconomic and oil price uncertainty, in contrast with previous work. Following Gourio (2012), macroeconomic disasters are modeled as sharp declines in economic growth and may be viewed as the result of an economic crisis such as the Great Depression or the Financial Crisis of 2008.

We find that shocks to the probability of an oil production disaster have modest effects on oil inventories, the price of oil, and oil price uncertainty, whether this disaster is realized or not. These shocks also have modest effects on investment and output. They are not a major driver of fluctuations in macroeconomic aggregates because large probability shocks are rare. Nor do they have much of an effect on macroeconomic uncertainty. Shocks to the probability of a growth disaster, which increase macroeconomic uncertainty, in contrast have larger effects on both the price of oil and the macroeconomy. These shocks also play a major role in the determination of oil price uncertainty, which helps explain why higher oil price uncertainty has historically been associated with lower real activity. Our analysis highlights that this association should not be interpreted as evidence of a causal link.

We also explore the possibility that agents might be concerned about larger oil production disasters than observed historically. While this increases the importance of geopolitical oil price risk, even when considering a temporary 20% drop in global oil production, which roughly corresponds to the cessation of oil supplies from the Persian Gulf, it takes a large increase in the probability of

²While a multi-country model would allow us to assess the differential effects of oil production disasters across countries, our focus in this paper is the aggregate effects of oil production disasters on the global economy.

Figure 1: Sources of uncertainty

this disaster or its realization to cause sizable recessionary impacts.

Incorporating downside risk is crucial for our results. While increases in oil price uncertainty may alternatively be explained by stochastic volatility shocks to oil production growth, these shocks have very small effects on the economy because they do not generate risk tilted to the downside. We also show that the ability to store oil plays a central role. Without storage the responses of both oil market variables and macroeconomic aggregates to higher oil production risk tend to be muted, suggesting that models without storage fail to capture the full effects of shifts in oil production risk. Oil storage also matters for the responses of the global economy to growth disaster probability shocks, underscoring the importance of jointly modeling oil production disasters and macroeconomic disasters. Finally, we show that our baseline results are robust to allowing for nominal rigidities and time-varying markups.

Our model differs from previous work in that it draws attention to the fact that changes in oil price uncertainty need not be an indication of exogenous shifts in the uncertainty about future oil supplies. As illustrated in [Figure 1](#), oil price uncertainty also reflects exogenous macroeconomic uncertainty shocks, mirroring the standard result that the real price of oil responds to shifts in the demand for oil. In addition, oil price uncertainty responds to level shocks in the oil market and the macroeconomy, such as realizations of disasters. Thus, not only are level and uncertainty shocks not the same, as implicitly assumed in VAR-GARCH models, but the effects of a level shock are not separable from those of an uncertainty shock, as assumed in VAR models with stochastic

volatility. Similarly, it does not make sense to employ recursive linear VAR models with oil price uncertainty either ordered first or last, since oil price uncertainty is simultaneously determined with macroeconomic aggregates. These results call into question a large body of empirical work that has produced seemingly robust evidence of large recessionary effects of oil price uncertainty shocks and has shaped the policy debate about geopolitical risk.

Implicitly the existing literature has taken the view that oil price uncertainty shocks are entirely driven by geopolitical events in oil markets. This amounts to ignoring the colored lines in [Figure 1](#). Only if oil price uncertainty evolves independently of other shocks in the economy is there a direct causal link from oil price uncertainty shocks to macro aggregates. Once one accounts for the dependence of oil price uncertainty on other shocks, conventional interpretations of the link between oil price uncertainty shocks and the economy break down. Our theoretical model is designed to capture these additional channels of transmission.

Related Literature Our work relates to several strands of the literature. First, it contributes to the large literature on the effects of uncertainty shocks on the macroeconomy (e.g., Berger et al., 2020; Bernstein et al., 2024; Bianchi et al., 2018; Bloom, 2009; Bloom et al., 2018; Fernández-Villaverde et al., 2015, 2011; Gourio, 2012; Jurado et al., 2015; Leduc and Liu, 2016; Ludvigson et al., 2021) by focusing on the interaction between macroeconomic and oil price uncertainty, which has received little attention to date. We show that modeling geopolitical oil production risk and macroeconomic risk jointly is necessary for understanding the evolution of oil price uncertainty.

Second, our analysis contributes to the literature that uses micro-founded models to examine the relationship between oil price uncertainty and the economy. Notable contributions include Başkaya et al. (2013), Gao et al. (2022), Olovsson (2019), and Ready (2018). These studies rely on various simplifications such as stochastic volatility shocks or time-invariant disaster probabilities. Some abstract from macroeconomic risk, while others limit attention to long-run uncertainty. In contrast, we examine this question within a general equilibrium model with endogenous oil price uncertainty. Our analysis accounts for the fact that geopolitical oil price risk is inherently one-sided and reflects the stochastic arrival of oil production disasters driven by geopolitical events.

This allows us to quantify the effects of shifts in probability of an oil production disaster. We show that such shocks have modest effects on real activity and are too infrequent to generate much volatility in macroeconomic aggregates.

Third, our analysis is relevant for a large empirical literature making the case that oil price uncertainty shocks driven by geopolitical events are recessionary (e.g., Bernanke, 1983; Drakos and Konstantinou, 2013; Elder and Serletis, 2010; Ferderer, 1996; Gao et al., 2022; Guo and Kliesen, 2005; Jo, 2014) and affect oil production and storage (e.g., Cross et al., 2022; Kellogg, 2014). We find that oil price uncertainty responds to level and uncertainty shocks to the macroeconomy. This fact complicates the identification of oil price uncertainty shocks and has important implications for empirical work seeking to establish the macroeconomic effects of shocks to oil price uncertainty.

Outline The remainder of the paper is organized as follows. In [Section 2](#), we highlight the importance of geopolitical risk for the economy. We propose an index of the uncertainty in the real price of oil building on Jurado et al. (2015), trace its evolution since the 1970s, and discuss the relationship between downside risk in oil production and oil price uncertainty. [Section 3](#) reviews why many economists expect oil price uncertainty to slow economic activity. [Section 4](#) introduces a calibrated model of the global economy that elucidates the determination of oil price uncertainty and macroeconomic uncertainty. This model contains the key features required for understanding the transmission of shifts in geopolitical oil price risk and should be viewed as a baseline for assessing the quantitative importance of these risks.

In [Section 5](#), we study the relationship between oil price uncertainty and macroeconomic uncertainty, the transmission of uncertainty shocks to the economy, and the ability of these shocks to explain fluctuations in economic growth and the oil price. Our analysis also sheds light on the key economic mechanisms in the model, the importance of modeling downside risk, and the sensitivity of our results to the specification of the oil production disaster. In [Section 6](#), we discuss the relationship between our work and previous models of the transmission of oil price uncertainty shocks. [Section 7](#) discusses implications of our analysis for empirical work on the effects of oil price uncertainty shocks. The concluding remarks are in [Section 8](#).

2 MEASURING OIL PRICE UNCERTAINTY


There has been growing interest in the impact of shifts in geopolitical risk in the oil market in recent years. Historically, increases in oil price risk have been associated, for example, with uncertainty about the implications of the Iranian Revolution in 1979 and the outcome of the invasion of Kuwait in 1990. More recently, there was a surge in uncertainty about the possibility of Russia refusing to sell oil to Europe after the invasion of Ukraine in 2022 and then about the effectiveness of a price cap on Russian oil exports. Other recent sources of oil price uncertainty have included concerns in 2024 about a war between Israel and Iran disrupting global oil exports, changes in OPEC production quotas, and the ability of U.S. shale oil producers to maintain their production increases.

The focus of our paper is to develop a deeper understanding of how time-variation in geopolitical risk in oil markets affects oil price uncertainty and hence economic fluctuations. Our starting point is the downside risk to oil production caused by these events. These downside risks are inherently subjective because they relate to events that have not occurred. In contrast, oil price uncertainty can be quantified using econometric methods. This does not make oil price uncertainty a good indicator of geopolitically driven downside risk in oil production, however, because these two variables need not go hand-in-hand. While downside geopolitical risk to oil production can raise oil price uncertainty, not all surges in oil price uncertainty are driven by geopolitical events.

Figure 2 quantifies the uncertainty about the real price of oil in global oil markets since the modern oil market emerged in the early 1970s. We follow Jurado et al. (2015) in measuring oil price uncertainty (\mathcal{U}_{p^o}) as the one-quarter ahead conditional volatility of the unpredictable component from a predictive model of the real price of oil.³ This definition highlights the fact that what matters for economic decision making is not whether the price of oil has become more or less variable, but whether it has become more or less predictable.⁴ As is standard, the real price of oil is defined as the U.S. refiners' acquisition cost for oil imports deflated by the implicit GDP deflator and is plotted

³Details of the construction of the uncertainty measure can be found in Appendix A.

⁴The definition of uncertainty in Jurado et al. (2015) is closely related to the formal measure of predictability in Diebold and Kilian (2001), since lack of predictability implies uncertainty.

Figure 2: Empirical measures of oil price uncertainty

Notes: All series are indices. The solid line shows the uncertainty about the percent change in the real price of oil obtained by deflating the U.S. refiners' acquisition cost for oil imports by the U.S. CPI for all urban consumers. The method used to quantify this uncertainty is based on Jurado et al. (2015). The dotted line is the option-implied crude oil price volatility index (OVX) published by the Chicago Board Options Exchange. The dash-dotted line shows the (rescaled) text-based oil price uncertainty index in Abiad and Qureshi (2023). The dashed line is the quarterly average of the historical GPR series in Caldara and Iacoviello (2022).

in Appendix C for reference. The predictable component of the growth rate of the real price of oil is approximated using a diffusion index based on largely the same set of variables used by Jurado et al. (2015), augmented by the real price of oil, updated, and aggregated to quarterly frequency.

We estimate the uncertainty about the price of oil from 1974Q4 to 2023Q4.⁵ There are large spikes in 1979, 1986, and 1990 at the time of the Iranian Revolution, the collapse of OPEC, and the

⁵Our analysis of oil price uncertainty cannot be extended back further because the U.S. refiners' acquisition cost for crude oil imports we use as a proxy for the global price of oil is only available starting in early 1974. While the WTI price of oil is available much further back, that price only captures the domestic price of oil in the United States. Since the WTI price remained regulated until the early 1980s and because arbitrage between the U.S. oil market and the global oil market temporarily broke down during the U.S. shale oil boom in the 2010s, the WTI price is not a good proxy for the global price of oil even after 1974. These problems are compounded in the pre-1974 era. The nominal WTI price of oil prior to 1974 was fixed for extended periods, followed by discrete jumps, reflecting the regulation of the U.S. oil market. As a result, not only is there a major structural change in the distribution of the real price of oil in late 1973, but there is a structural break in the predictive correlation between U.S. real GDP growth and the real price of oil. Combining the pre- and post-1974 oil price data thus would be inappropriate (see Alquist et al., 2013).

invasion of Kuwait. Not all geopolitical events are associated with surges in oil price uncertainty, however. For example, neither the outbreak of the Iran-Iraq War in late 1980 nor the outbreak of the Israel-Hamas War in the last quarter of 2023 had a discernible impact on the index.

The largest spike in oil price uncertainty in 2008 was not driven by geopolitical risk, but by macroeconomic risk. Similarly, the surge in oil price uncertainty in 2015 appears to be driven by market forces rather than geopolitics (Baumeister and Kilian, 2016b), as was a smaller spike during the Asian Financial Crisis of the late 1990s. Sometimes, geopolitical events coincide with surges in macroeconomic risk, as was the case in early 2020 when the Covid-19 recession occurred at the same time as the Saudi price war in the oil market or in 1979 when rising geopolitical uncertainty coincided with increased uncertainty about monetary policy.

An alternative way to measure oil price uncertainty is to use the implied volatility index (OVX) published by the Chicago Board Options Exchange. While these data are only available starting in 2007, we find a correlation of 0.71 between our index and the OVX when both series are available.⁶ Figure 2 also shows that our oil price uncertainty index differs systematically from the geopolitical risk (GPR) index of Caldara and Iacoviello (2022), which quantifies the newspaper coverage of geopolitical events not limited to oil markets. The direction of these indices often differ, and their correlation is close to zero. Similarly, the text-based oil price uncertainty index in Abiad and Qureshi (2023) based on the methodology in Baker et al. (2016) differs from our data-based index. The low correlation of 0.3 with our index suggests that the text-based oil price uncertainty index fails to capture changes in the predictability of the real price of oil.

3 WHY ECONOMISTS THINK OIL PRICE UNCERTAINTY MATTERS

Interest in fluctuations in oil price uncertainty dates to the mid-1980s. Economists at the time observed that the economy entered a steep recession after the 1979/80 oil price surge, but a similarly large drop in the price of oil in 1986 did not cause a large economic expansion. This fact is

⁶In related work, Gao et al. (2022) derive an index similar to the OVX series using oil options back to 1990Q1. The relationship between these indices is similar over the extended sample.

consistent with two mutually exclusive narratives. One is that the relationship between oil prices and the U.S. economy is linear, which implies that the effect of oil price shocks on the economy is modest at best and that the recession in the early 1980s is explained in substantial part by other shocks (e.g., Barsky and Kilian, 2002). This explanation is consistent with micro-founded models of the transmission of oil price shocks that predict that a rising price of oil will only modestly slow growth in oil-importing economies, as consumers' income is reduced and firms face higher production costs, and, conversely, falling oil prices will only modestly stimulate growth in oil-importing economies (e.g., Backus and Crucini, 2000).⁷

The other narrative is that this relationship is nonlinear with positive oil price shocks having disproportionately larger effects on the economy. Macroeconomists for many years have been partial to this interpretation. A leading explanation of the nonlinearity required to explain the disproportionately large effect of positive oil price shocks and the negligible effects of negative oil price shocks, is that the rise in oil price uncertainty associated with the 1979/80 oil price surge caused consumer spending and business fixed investment to drop, amplifying the effects of rising oil prices, whereas in 1986 an increase in oil price uncertainty associated with the fall in oil prices largely offset the stimulus from lower oil prices.

3.1 REAL OPTIONS THEORY The most commonly cited reason why oil price uncertainty shocks matter for economic activity was articulated by Bernanke (1983) in a partial equilibrium setting.⁸ Bernanke's point is that—to the extent that the cash flow from an irreversible investment project depends on the price of oil or its derivatives—all else equal, increased uncertainty about the price of oil prompts firms to delay investments. As a result, investment expenditures drop and real output declines. Uncertainty for this purpose may be measured by the expected conditional volatility of the real price of oil over the relevant investment horizon. Exactly the same reasoning applies to purchases of energy-intensive consumer durables such as cars (see Edelstein and Kilian, 2009).

⁷There have been many attempts to design macroeconomic models that amplify the transmission of oil price shocks (e.g., Aguiar-Conraria and Wen, 2007; Atkeson and Kehoe, 1999; Finn, 2000; Rotemberg and Woodford, 1996). These theoretical models are not necessarily supported by the data, however. More importantly, being able to generate a larger recession after 1979/80 makes it even more difficult to explain the absence of an economic expansion in 1986.

⁸For related discussion of the effects of uncertainty shocks more generally see Pindyck (1991) and Bloom (2014).

There are several caveats to this application of real options theory. First, the quantitative importance of this cash flow channel depends on how important the real price of oil is for investment and durable consumption decisions and on the share of such expenditures in aggregate spending. For example, it seems intuitive that uncertainty about the price of oil would be important for decisions about oil drilling in Texas (see Kellogg, 2014). It is less obvious that it would be as important for investment in other sectors of the economy such as textile production or information technology, the expected profitability of which does not depend as much on oil prices.

Second, there is reason to believe that for longer-term investment projects, the variation over time in the uncertainty about the real price of oil is small. Consider an airline purchasing new planes that are expected to fly for 20 years. The cash flow from this investment clearly depends on fuel prices and an increase in expected fuel price uncertainty, all else equal, should cause the airline to delay the investment. However, the predictable component of the variance of the real price of oil quickly reverts to the unconditional variance at longer horizons, so one would not expect variation in the conditional variance at a monthly or quarterly frequency to have a large effect on the investment decision.

Third, Bernanke (1983) takes the real price of oil as exogenously given. This simplifying assumption does not hold in practice, complicating the analysis. The concern is that we may attribute to oil price uncertainty the effects of macroeconomic uncertainty, which is more likely to affect the cash flow from the investment. As discussed in [Section 4.4](#), it is also unclear whether the importance of the real options channel survives in general equilibrium. Thus, the overall importance of this channel for the aggregate economy is open to question.

3.2 PRECAUTIONARY SAVINGS A complementary reason first articulated in Edelstein and Kilian (2009) is that households' increased uncertainty about their future income in the wake of unexpected changes in the real price of oil will cause an increase in precautionary savings. In this interpretation, oil price uncertainty may affect a wide range of consumer expenditures. This argument has subsequently been formalized in Plante and Traum (2012) and Başkaya et al. (2013).

While a response of precautionary savings to oil price uncertainty shocks may seem persuasive

in a partial equilibrium setting, its quantitative importance becomes less obvious when moving to general equilibrium models. For example, in many general equilibrium models with stochastic volatility shocks, consumption barely declines or even increases in response to higher uncertainty (e.g., Born and Pfeifer, 2021; de Groot et al., 2018).

3.3 PRECAUTIONARY INVENTORY DEMAND A third channel by which increased oil price uncertainty can reduce economic activity operates through precautionary demand for oil inventories. This channel was introduced in Kilian (2009) and expanded on in Kilian and Murphy (2014) and Cross et al. (2022). These studies emphasized that higher precautionary demand driven by increases in oil price uncertainty, all else equal, will raise the real price of oil and reduce global economic activity by discouraging oil consumption. This analysis in turn builds on the theoretical insights in Alquist and Kilian (2010) of how mean-preserving shifts in the uncertainty about future oil supply shortfalls affect the real price of oil through inventory accumulation. General equilibrium models incorporating oil storage have been presented in Olovsson (2019) and Gao et al. (2022). The latter study considers stochastic volatility shocks to oil production and finds that firms increase their oil inventory holdings in response to an increase in oil price uncertainty.

4 A MODEL OF THE PROPAGATION OF UNCERTAINTY SHOCKS

In this section, we introduce a model of the global economy designed to elucidate the determinants of oil price uncertainty and to uncover the macroeconomic implications of downside risk. Focusing on the global economy allows us to sidestep the complications involved in modeling oil importing and exporting economies and to concentrate on the key research questions.

4.1 ENVIRONMENT The model is a nonlinear stochastic growth model augmented to include oil production. Oil is used as an intermediate input by a representative firm that produces a final good. The model allows for precautionary savings as well as oil storage, which has been shown to play an important role in driving fluctuations in oil prices (see Kilian and Murphy, 2014). The distinguishing feature of the model is that it includes downside risk to both oil production and

output growth. While downside risk to oil production can be thought of as arising from geopolitical events, downside risk to the macroeconomy involves rare, sharp economic downturns, such as the Great Depression or the Great Recession that are not otherwise captured by the model.

We follow Gourio (2012) in modeling such events as disasters that arrive with a time-varying probability.⁹ Time-variation in the probability of oil production and output growth disasters induces exogenous variation in oil price and macroeconomic uncertainty. While the idea of modeling downside risk in the economy as growth disasters with time-varying probability is not new, we are the first to apply this approach to modelling geopolitical risk in the oil market.¹⁰ One advantage of this approach compared to the more traditional approach of subjecting oil production to a stochastic volatility shock is that it accounts for the fact that the risk agents are typically concerned with in practice is not two-sided. These risks involve a sharp reduction in oil production. Such rare disasters matter not only because of their impact when they occur, but, more importantly, because agents' behavior reflects the anticipation of these disasters even when they are not realized.

Productivity and Growth Disasters The growth rate of productivity, $g_t = a_t/a_{t-1}$, follows

$$\ln g_t = \ln \bar{g} + \sigma_g \varepsilon_{g,t} - \zeta_g (v_t^g - \bar{\pi}_1^g), \quad \varepsilon_{g,t} \sim \mathbb{N}(0, 1),$$

where \bar{g} is the steady-state growth rate. The indicator variable v_t^g equals 1 if a growth disaster occurs and 0 otherwise. The transition matrix for v_t^g is summarized by

$$\Pr(v_{t+1}^g = 1 | v_t^g = 1) = \bar{q}^g,$$

$$\Pr(v_{t+1}^g = 1 | v_t^g = 0) = p_t^g,$$

where the probability of a growth disaster follows

$$\ln p_t^g = (1 - \rho_p^g) \ln \bar{p}^g + \rho_p^g \ln p_{t-1}^g + \sigma_p^g \varepsilon_{p,t}^g, \quad \varepsilon_{p,t}^g \sim \mathbb{N}(0, 1).$$

The size of the disaster is ζ_g , and $\bar{\pi}_1^g = \frac{\bar{p}^g}{1 + \bar{p}^g - \bar{q}^g}$ is the unconditional probability of the disaster.

⁹Applications of disaster risk include Barro and Ursúa (2012), Gourio (2013), Gourio et al. (2013), Wachter (2013), Shen (2015), Farhi and Gabaix (2016), Olovsson (2019), Berger et al. (2020), Kim (2022), and Kilian et al. (2025).

¹⁰Olovsson (2019) uses a related approach, except his model treats the probability of an oil disaster as constant.

Following Gourio (2012), capital is destroyed when the disaster occurs. As in that paper, we use a broad interpretation of capital destruction that can represent a sharp reduction in capital quality due to the loss of intangible capital during economic downturns or the destruction of physical capital during wars, natural disasters, or sectoral reallocations. Let k_t denote the inherited stock of capital and i_t denote investment. The capital stock evolves according to

$$k_{t+1} = e^{-\zeta g^{vg,t+1}}((1 - \delta)k_t + i_t - \phi(i_t/k_t)k_t).$$

The functional form of the adjustment cost follows Jermann (1998) and is given by

$$\phi(i_t/k_t) = i_t/k_t - (\mu_1 + \frac{\mu_2}{1-1/\nu})(i_t/k_t)^{1-1/\nu},$$

where $\mu_1 = (\bar{g} - 1 + \delta)/(1 - \nu)$ and $\mu_2 = (\bar{g} - 1 + \delta)^{1/\nu}$.

Final Goods Firm A representative firm maximizes profits by choosing its investment (i_t), capital (k_{t+1}), labor (n_t), and oil (o_t) inputs. Following the seminal work of Kim and Loungani (1992) and Backus and Crucini (2000), the firm produces a final good y_t using a Cobb-Douglas technology that aggregates labor and capital services, which are produced using a normalized CES production function that aggregates capital and oil. The firm's profit maximization problem is given by

$$V_t = \max_{i_t, k_{t+1}, n_t, o_t} y_t - i_t - p_t^o o_t - w_t n_t + E_t[x_{t+1} V_{t+1}]$$

subject to

$$k_{t+1} = e^{-\zeta g^{vg,t+1}}((1 - \delta)k_t + i_t - \phi(i_t/k_t)k_t),$$

$$y_t = y_0(a_t n_t)^{1-\xi} \left((1 - \alpha)(k_t/k_0)^{1-1/\sigma} + \alpha(o_t/o_0)^{1-1/\sigma} \right)^{\xi/(1-1/\sigma)},$$

where σ is the elasticity of substitution between capital and oil, δ is the depreciation rate of capital, $1-\xi$ is the share of labor in gross output, and α controls the share of oil in the capital services aggregate. The scalars y_0 , k_0 , and o_0 are set so that α is equal to the cost share of oil in the capital services aggregate. These normalizations do not affect the results but simplify the model calibration.¹¹

¹¹A more detailed discussion of normalized CES production functions can be found in Klump et al. (2012). Some more recent studies, such as Başkaya et al. (2013), Hassler et al. (2021), Olovsson (2019), and Ready (2018), favor a specification of the production function that treats oil as complementary to a capital-labor bundle. Our substantive conclusions are robust to this alternative functional form, as shown in Appendix F.

The first-order conditions for the firm's problem are given by

$$\begin{aligned} w_t &= (1 - \xi)y_t/n_t, \\ p_t^o &= \xi\alpha \frac{(o_t/o_0)^{1-1/\sigma}}{(1-\alpha)(k_t/k_0)^{1-1/\sigma} + \alpha(o_t/o_0)^{1-1/\sigma}} \frac{y_t}{o_t}, \\ E_t[x_{t+1}r_{t+1}^i] &= 1, \end{aligned}$$

where

$$\begin{aligned} r_{t+1}^i &\equiv e^{-\zeta_g v_{g,t+1}} (r_{t+1}^k + (1 - \delta + \mu_1 + \frac{\mu_2}{\nu-1} (i_{t+1}/k_{t+1})^{1-1/\nu}) p_{t+1}^k) / p_t^k, \\ r_t^k &\equiv \xi(1 - \alpha) \frac{(k_t/k_0)^{1-1/\sigma}}{(1-\alpha)(k_t/k_0)^{1-1/\sigma} + \alpha(o_t/o_0)^{1-1/\sigma}} \frac{y_t}{k_t}, \\ p_t^k &\equiv \frac{1}{1 - \phi'(i_t/k_t)} = \frac{1}{\mu_2} \left(\frac{i_t}{k_t} \right)^{1/\nu}. \end{aligned}$$

Oil Production and Oil Disasters The production of oil is exogenous and given by $o_t^s = a_t^o e_t$. This assumption is commonly used in micro-founded models of the oil market, given the paucity of data for the oil sector. The permanent component, a_t^o , reflects factors that influence the productive potential of the oil sector, including the evolution of oil reserves and technological progress that increases the ability of the sector to extract oil from current reserves. We include a shock to this permanent component to allow for productivity shocks in the oil sector not related to geopolitical oil supply disruptions. The transitory component reflects temporary changes in the production of oil driven by exogenous geopolitical events. Oil production disasters are modeled as transitory, given evidence that geopolitical supply disruptions historically have not had long-lasting effects on global oil production, as discussed in the calibration section.

The permanent component is cointegrated with productivity in the rest of the economy,

$$a_t^o = \kappa_0 g_t^{\kappa_1} \epsilon_{t-1}^{\kappa_2} a_{t-1}^o \exp(\sigma_{go} \varepsilon_{go,t})$$

where $\epsilon_t = a_t/a_t^o$, κ_1 determines the impact response of a growth shock on a_t^o , and κ_2 affects the speed at which a_t^o converges to a_t . This setup allows for a slow response of oil production to productivity growth shocks in the rest of the economy, which is a key feature of the data.¹²

¹²When $\kappa_1 = 1$ and $\kappa_2 = 0$, $a_t^o = a_t$, so the production of oil responds immediately to changes in productivity elsewhere in the economy. This special case corresponds to the assumption made in Gao et al. (2022). Cointegration

The transitory component of global oil production is given by

$$\ln e_t = \ln \bar{e} - \zeta_e(v_t^e - \bar{\pi}_1^e).$$

The indicator variable v_t^e equals 1 if an oil production disaster occurs and 0 otherwise. The transition matrix for v_t^e is summarized by

$$\Pr(v_{t+1}^e = 1 | v_t^e = 1) = \bar{q}^e,$$

$$\Pr(v_{t+1}^e = 1 | v_t^e = 0) = p_t^e,$$

where the probability of an oil disaster follows

$$\ln p_t^e = (1 - \rho_p^e) \ln \bar{p}^e + \rho_p^e \ln p_{t-1}^e + \sigma_p^e \varepsilon_{p,t}^e, \quad \varepsilon_{p,t}^e \sim \mathbb{N}(0, 1).$$

The size of the disaster is ζ_e , and $\bar{\pi}_1^e = \frac{\bar{p}^e}{1 + \bar{p}^e - \bar{q}^e}$ is the unconditional probability of the disaster.

Oil Storage A representative oil storage firm maximizes profits by choosing inventories, s_{t+1} , and how much oil to supply to the final goods firm, o_t . The firm's maximization problem is given by

$$V_t^o = \max_{o_t, s_{t+1}} p_t^o o_t + E_t[x_{t+1} V_{t+1}^o]$$

subject to

$$s_{t+1} = (1 - \omega)s_t + o_t^s - o_t - \frac{v}{2}(s_t/a_t)^{-2}a_t,$$

where ω is the cost of storage. Following Gao et al. (2022), there is a penalty, v , that prevents stockouts ($s = 0$) from occurring, as they are not observed in the global oil market.

The first-order condition for the storage firm is given by

$$1 = E_t[x_{t+1} r_{t+1}^s],$$

where

$$r_{t+1}^s \equiv ((1 - \omega + v(s_{t+1}/a_{t+1})^{-3})p_{t+1}^o)/p_t^o.$$

is rare in general equilibrium models with oil. One exception is Ready (2018), who models cointegration between oil production and TFP in a setting with long-run risk. Similarly, cointegrated TFP processes have been used in two-country international real business cycle models (e.g., Rabanal et al., 2011).

Household A representative household maximizes the present discounted value of utility by choosing consumption, c_t , hours worked, n_t , bond holdings, b_{t+1} , and equity shares, s_{t+1}^e , which have unit net supply. The household has Epstein-Zin recursive preferences to distinguish between risk aversion, γ , and the intertemporal elasticity of substitution, ψ (see Epstein and Zin, 1989).

The household's maximization problem is given by

$$J_t = \max_{c_t, n_t, s_{t+1}^e, b_{t+1}} \left((1 - \beta) u_t^{1-1/\psi} + \beta (E_t[J_{t+1}^{1-\gamma}])^{\frac{1-1/\psi}{1-\gamma}} \right)^{\frac{1}{1-1/\psi}}$$

subject to

$$\begin{aligned} u_t &= c_t^\chi (a_t(1 - n_t))^{1-\chi}, \\ c_t + p_t^e s_{t+1}^e + b_{t+1}/r_t &= w_t n_t + (p_t^e + d_t^e) s_t^e + b_t, \end{aligned}$$

where β is the discount factor, p_t^e is the equity price, r_t is the risk-free rate, w_t is the wage rate, d_t^e are dividends from firm ownership, and the Frisch elasticity of labor supply $\eta^\lambda = \frac{1-n_t}{n_t} \frac{1-(1-1/\psi)\chi}{1/\psi}$.

The first-order conditions for the household are given by

$$\begin{aligned} \chi w_t (1 - n_t) &= (1 - \chi) c_t, \\ 1 &= E_t[x_{t+1} r_t], \\ 1 &= E_t[x_{t+1} r_{t+1}^e], \end{aligned}$$

where

$$\begin{aligned} r_{t+1}^e &\equiv (p_{t+1}^e + d_{t+1}^e)/p_t^e, \\ x_{t+1} &\equiv \beta (u_{t+1}/u_t)^{1-1/\psi} (c_t/c_{t+1}) (J_{t+1}/z_t)^{1/\psi-\gamma}, \\ z_t &\equiv (E_t[J_{t+1}^{1-\gamma}])^{1/(1-\gamma)}. \end{aligned}$$

The equity risk premium is defined as

$$r_t^{ex} \equiv r_t^e - r_{t-1}.$$

Market Clearing Following Jermann (1998) and Gourio (2012), the final goods firm issues debt to finance its expected asset holdings, where ϑ determines leverage. Since the Modigliani-Miller

theorem holds in our model, the introduction of firm leverage only affects equity returns. There is no effect on household or firm decisions. Aggregate firm dividends are given by

$$d_t^e = d_t^f + d_t^s - \vartheta(E_{t-1}k_t - \frac{1}{r_t}E_t k_{t+1}),$$

where $d_t^f = y_t - i_t - p_t^o o_t - w_t n_t$ and $d_t^s = p_t^o o_t$. Asset market clearing implies that $s_t^e = 1$ and total bond issuance is given by $b_t = \vartheta E_{t-1}k_t$. Market clearing in the goods market implies $c_t + i_t = y_t$.

Due to the stochastic trend in productivity, we detrend the model by defining $\tilde{x}_t \equiv x_t/a_t$. The detrending process introduces the growth terms $g_t = a_t/a_{t-1}$ and $g_{o,t} = a_t^o/a_{t-1}^o$. Appendix E provides the detrended equilibrium system of equations.

Uncertainty We follow Plante et al. (2018) and Bernstein et al. (2024) and define output uncertainty as the conditional volatility of log output growth, which is given by

$$\mathcal{U}_t^y = \sqrt{E_t[(\ln(y_{t+1}/y_t) - E_t[\ln(y_{t+1}/y_t)])^2]}.$$

Oil price uncertainty, $\mathcal{U}_t^{p_o}$, is analogously defined as the uncertainty surrounding $\ln(p_{t+1}^o/p_t^o)$. This definition is equivalent to the measure we used to compute oil price uncertainty in the data.¹³

4.2 SOLUTION METHOD Modeling the oil sector considerably increases the computational cost of solving the model compared to a model that only includes macroeconomic risk. Our model has 7 state variables ($k_t, s_t, v_t^g, v_t^e, \ln p_t^g, \ln p_t^e, \epsilon_{t-1}$), 4 of which are related to the oil market. There are 4 continuous shocks and 2 discrete shocks. In total, the state space contains over 300,000 nodes and 40,000 shock realizations for each node in the state space.

The existence of time-varying disaster risk prevents the use of perturbation methods. We therefore employ a fully nonlinear solution method. Specifically, the model is solved using the policy function iteration algorithm described in Richter et al. (2014), which is based on the theoretical work in Coleman (1991). The algorithm minimizes the Euler equation errors on each node in the state space and computes the maximum change in the policy functions. It then iterates until the

¹³The uncertainty surrounding oil price growth is equivalent to the uncertainty surrounding the log oil price because p_t^o is known at time t and cancels from the definition of $\mathcal{U}_t^{p_o}$. An analogous result holds for output uncertainty.

Table 1: Model calibration at a quarterly frequency

Parameter	Value	Target
Discount Factor (β)	0.997	$E(r)$
Risk Aversion (γ)	10	Gao et al. (2022), Croce (2014)
Intertemporal Elasticity of Substitution (ψ)	2	Gao et al. (2022), Croce (2014)
Frisch Labor Supply Elasticity (η^λ)	2	Peterman (2016), Basu and Bundick (2017)
Capital-Oil Elasticity of Substitution (σ)	0.105	$SD(\Delta p^o)$
Capital Depreciation Rate (δ)	0.025	Depreciation on fixed assets and durables
Capital-Oil Share of Production (ξ)	0.4043	Avg. labor share of income
Investment Adjustment Cost (ν)	3.3	$SD(\Delta i)$
Oil Storage Cost (ω)	0.025	Casassus et al. (2018), Gao et al. (2022)
Average Growth Rate (\bar{g})	1.0043	$E(\Delta y)$
Firm Leverage (ϑ)	0.9	$SD(r^{ex})$
Elasticity of Oil Supply to TFP (κ_1)	0	Newell and Prest (2019)
Oil Supply Adjustment Speed to TFP (κ_2)	0.05	Half life of 3.5 years
Growth Shock SD (σ_g)	0.0095	$SD(\Delta y)$
Oil Production Growth Shock SD (σ_{go})	0.011	$SD(\Delta o^s)$
Growth Disaster Size (ζ_g)	0.018	$E(r^{ex})$
Probability of Entering Growth Disaster (\bar{p}_g)	0.005	Occurs in expectation every 50 years
Probability of Exiting Growth Disaster (\bar{q}_g)	0.9	Gourio (2012)
Growth Disaster Probability Persistence (ρ_{pg})	0.8	$AC(\mathcal{U}_y)$
Growth Disaster Probability SD (σ_{pg})	0.9	$SD(\mathcal{U}_y)$
Oil Production Disaster Size (ζ_e)	0.05	Avg. peak decline in oil production disasters
Probability of Entering Oil Disaster (\bar{p}_e)	0.02	Avg. frequency of oil production disasters
Probability of Exiting Oil Disaster (\bar{q}_e)	0.67	Avg. duration of oil production disasters
Oil Disaster Probability Persistence (ρ_{pe})	0.9	$AC(\mathcal{U}_{p^o})$
Oil Disaster Probability SD (σ_{pe})	1.4	$SD(\mathcal{U}_{p^o})$

maximum change is below a specified tolerance. The algorithm was programmed in Fortran and run on the BigTex supercomputer at the Federal Reserve Bank of Dallas. Appendix D describes the solution method in more detail.

4.3 CALIBRATION Each period in the model is one quarter. The parameters shown in Table 1 are informed by moments in the data and the related literature. For a few parameters, we rely on U.S. data to inform our calibration given the paucity of global data. The moments are computed using data from 1975Q1 to 2019Q4. Excluding the more recent data ensures that the calibration is not unduely influenced by the COVID-19 pandemic. Appendix B provides our data sources.¹⁴

¹⁴Estimation using Bayesian methods or the simulated method of moments is not possible due to the high dimensionality of the model. Even when using a supercomputer with thousands of cores, the model takes several hours to solve.


The discount factor β is set to 0.997 to match the average real interest rate. The relative risk aversion coefficient, γ , and intertemporal elasticity of substitution are set to 10 and 2, respectively, consistent with Gourio (2013), Croce (2014), Gao et al. (2022), and several other recent studies.¹⁵ The Frisch elasticity, which is pinned down by the steady-state labor supply, is set to 2 following Peterman (2016), Basu and Bundick (2017), and many others in the business cycle literature.¹⁶

The elasticity of substitution between capital and oil, σ , is set to 0.105 to match the volatility of oil price growth. Backus and Crucini (2000) adopt the same functional form of the production function and use a similar value (0.09). The Cobb-Douglas weight on capital services (ξ) is set to match the average labor share of income. The investment adjustment cost parameter, ν , is set to match the volatility of per capita investment growth. The capital depreciation rate, δ , matches the annual average rate of depreciation on private fixed assets and durable goods. The oil storage cost, ω , is set to 0.025 following Casassus et al. (2018) and Gao et al. (2022). As in Basu and Bundick (2017), the leverage parameter, ϑ , is set to 0.9 to help match the volatility of the equity premium.

The mean growth rate of productivity, \bar{g} , is set to 1.0043 to match the average growth rate of per capita real GDP. The standard deviation for the growth shock, σ_g , is set to 0.0095 to help match the volatility of real GDP growth. The calibration of the growth disaster parameters is guided by several moments in the data as well as the parameter choices in Gourio (2012). We set the size of the disaster, ζ_g , to 0.018 to match the mean equity premium. The mean probability of entering the disaster state, \bar{p}_g , is set to 0.005, which implies that these disasters happen once every 50 years in expectation given that there were two major events over the last 100 years. The persistence, ρ_{pg} , and standard deviation, σ_{pg} , of this probability are set to 0.8 and 0.9, respectively, to help match the autocorrelation and volatility of output uncertainty. The fixed probability of exiting a growth disaster, \bar{q}_g , is set to 0.9, in line with Gourio (2012). This value implies that growth disasters, on average, last 2.5 years. As shown in Appendix F, the responses to an average growth disaster are

¹⁵Swanson (2018) shows how to compute risk aversion under recursive preferences with an endogenous labor supply. Under our utility kernel, γ corresponds to risk aversion over consumption and leisure.

¹⁶The steady-state price of oil, \bar{p}^o , and steady-state output, \bar{y} , are normalized to 1. We set the steady-state oil expenditure share ($\bar{p}^o \bar{o} / \bar{y}$) to 0.054 so the average across simulations matches the share in the data. Similarly, the steady-state inventory share of oil consumption (\bar{s} / \bar{o}) is set to 1.13 to match the average share in the data.

Figure 3: Shortfall in global oil production during major geopolitical events

Notes: Reproduced from Hamilton (2013) using updated global oil production data. We exclude the 2002/03 episode because the revised data show no evidence of a material shortfall.

very similar to those reported in Gourio (2012), who documents that his responses resemble the empirical estimates in Barro et al. (2013).

The value of κ_1 is set to 0, implying that productivity in the oil sector is unresponsive to changes in productivity in the rest of the economy within the first quarter. This is consistent with the view that oil production in the short run is determined entirely by geological and technological constraints (see, e.g., Newell and Prest, 2019). We set κ_2 to 0.05, so the half-life of the deviation between a_t^o and a_t is 4 years. The standard deviation of the growth shock to oil production, σ_{go} , is set to 0.011 to match the volatility of global oil production.

The parameters controlling the oil production disasters are based on historical oil production data. Following Hamilton (2013), Figure 3 plots global oil production during major geopolitical events, where production is expressed in percent deviations from the level at the beginning of the

Table 2: Data and simulated moments

Moment	Data	Model	Moment	Data	Model
$E(\Delta y)$	0.39	0.39	$SD(\Delta o^s)$	2.01	2.14
$E(s/o)$	0.97	0.97	$SD(\Delta p^o)$	14.39	14.27
$E(p^o o/y)$	0.045	0.046	$SD(r^{ex})$	8.29	5.55
$E(r^{ex})$	2.18	2.11	$SD(\mathcal{U}_y)$	14.51	15.61
$E(r)$	0.22	0.20	$SD(\mathcal{U}_{p^o})$	29.95	30.49
$SD(\Delta y)$	0.74	0.87	$AC(\mathcal{U}_y)$	0.87	0.81
$SD(\Delta i)$	1.95	1.92	$AC(\mathcal{U}_{p^o})$	0.93	0.82
$SD(\Delta s)$	2.30	2.25	$AC(\Delta o^s)$	-0.11	-0.20
$SD(r)$	0.91	0.37	$Corr(\Delta y, \mathcal{U}_y)$	-0.15	-0.48
$AC(\Delta y)$	0.32	0.28	$Corr(\Delta y, \mathcal{U}_{p^o})$	-0.27	-0.44

Notes: Moments above the middle line are targeted while those below it are untargeted. The model is calibrated to data from 1975Q1-2019Q4. $SD(\mathcal{U}_y)$ and $SD(\mathcal{U}_{p^o})$ are normalized by $SD(\Delta y)$ and $SD(\Delta p^o)$, respectively, to be consistent with the normalization in Jurado et al. (2015).

event. We set the size of the oil production disaster, ζ_e , to 0.05 to match the average peak decline in the data. The mean probability of entering the oil disaster state, \bar{p}_e , is set to 0.02 so that disasters occur every 12.5 years in expectation, given that there were four major events over the last 50 years. The persistence, ρ_{pe} , and standard deviation, σ_{pe} , of this probability are set to 0.9 and 1.4, respectively, to help match the autocorrelation and volatility of oil price uncertainty. The fixed probability of exiting an oil production disaster, \bar{q}_e , is set to 0.67 so that a disaster lasts, on average, for 3 quarters, which corresponds to the longest duration observed in the four episodes in [Figure 3](#).¹⁷

To compute the model-implied moments, we simulate the model 10,000 times, each with 180 quarters to match the length of the data used to calibrate the model. We calculate the moments of interest in each simulated data set and then compute the average moments across all simulations. [Table 2](#) compares the data and model-implied moments. The model closely matches most of the targeted moments. This includes moments related to the oil market (e.g., the standard deviations of oil price growth and oil production growth, the oil expenditure share, and the oil inventory-to-oil

¹⁷It may seem that options data could be used to help with the calibration. This is not the case. One challenge is that tail probabilities estimated from equity options as in Barro and Liao (2021) do not help quantify macroeconomic tail risk, but only equity risk. The distinction between financial risk and macroeconomic risk has been emphasized in Gao et al. (2022) and Ludvigson et al. (2021). Likewise, oil disaster probabilities are not recoverable from oil options because these prices reflect both oil and macroeconomic disaster risk in unknown combinations.

consumption share), real activity (e.g., the standard deviations of output and investment growth), asset prices (e.g., the average risk-free rate and equity risk premium), and uncertainty (e.g., the standard deviations and autocorrelations of output uncertainty and oil price uncertainty). Jointly matching all four of these key aspects of the economy gives us confidence that the model provides a good description of oil market, real activity, and uncertainty dynamics.¹⁸

The model also performs well at matching several untargeted moments shown at the bottom of Table 2. These results provide further validation of our model. For example, the volatility of oil inventory growth and the autocorrelations of output growth and oil production growth closely match the data. Output and oil price uncertainty are somewhat more countercyclical in the model than in the data, but this could presumably be addressed by adding exogenous volatility shocks. Finally, allowing for disaster risk raises the volatility in the risk-free rate, which has traditionally been difficult to match in real business cycle models.

4.4 DISCUSSION Our model incorporates precautionary savings by households in response to higher oil price uncertainty as well as storage, the two main economic mechanisms that are thought to propagate oil price uncertainty shocks. While we do not model real options arising from irreversible investment as in Bernanke (1983), our model features limited substitutability between capital and oil. This feature causes the expected return on investment to decline when the probability of an oil production disaster increases, generating recessionary effects in the model.

Although Bernanke’s theoretical analysis is often cited in support of models of oil price volatility shocks, it is not well appreciated that Bernanke was not modeling a monthly or quarterly oil price volatility shock. Rather, he envisioned agents being uncertain about whether the price of oil would permanently move to a higher level or not, which is a different thought experiment. In his model there are two types of capital that differ by their oil efficiency. The irreversibility of the investment decision causes risk averse agents to postpone the acquisition of either type of capi-

¹⁸One potential concern is the oil share calculated using U.S. data may not be representative of the global economy. We also calculated oil shares using supply-use tables from the World Input-Output Database (Timmer et al., 2015). These tables are available for 38 countries, including the U.S., China, Japan, Brazil, India, Indonesia and many countries in Europe. The average across all countries was 0.05, very close to the ergodic mean of the share in our model.

tal. The difficulty in generalizing this model to general equilibrium is that it requires aggregating different types of capital across many firms.

A closely related model that deals with the aggregation of different types of capital in general equilibrium was proposed by Atkeson and Kehoe (1999). In their putty-clay model there is a continuum of capital goods indexed by their oil efficiency. Existing capital goods use oil in fixed proportions, so, in the short run, there is no substitutability between capital and oil. However, firms may invest in new capital with different oil efficiency in response to changes in the price of oil. Although this point is not the focus of Atkeson and Kehoe (1999), their model implies that higher oil price uncertainty would reduce investment, as discussed in Plante and Traum (2012).

The reason we do not incorporate the putty-clay framework within our model is that two key assumptions made by Atkeson and Kehoe (1999) do not hold in our model. One is that the price of oil is exogenous; the other is that under their assumptions oil consumption does not respond to the price of oil on impact. These assumptions not only allow Atkeson and Kehoe to abstract from storage, but they allow them to aggregate across different types of capital without the need to track the distribution of capital types. The fact that the infinite-dimensional state space of capital stocks in their model can be reduced to a one-dimensional space facilitates the solution of their model.

In contrast, in our model the price of oil is endogenously determined. Suppose, for example, that there is an oil supply shock. In that case, we must add storage to the model because otherwise equilibrium in the oil market is unattainable. If oil consumption is predetermined and hence unresponsive to the oil price fluctuations caused by the oil supply shock, oil inventories must absorb any imbalances in the oil market each period. It can be shown that, as a result, the oil inventory moments of the simulated model data differ substantially from the oil inventory moments in the actual data. It may seem that this problem could be addressed by dropping the assumption that oil consumption is unresponsive to the price of oil, but this would render the capital stock intractable, which is why we do not consider the putty-clay framework in our model. However, our model with disaster risk generates investment and output responses that are qualitatively consistent with those in models of irreversible investment. The reason is that in our model risk averse agents are

reluctant to invest given the limited substitutability between capital and oil.

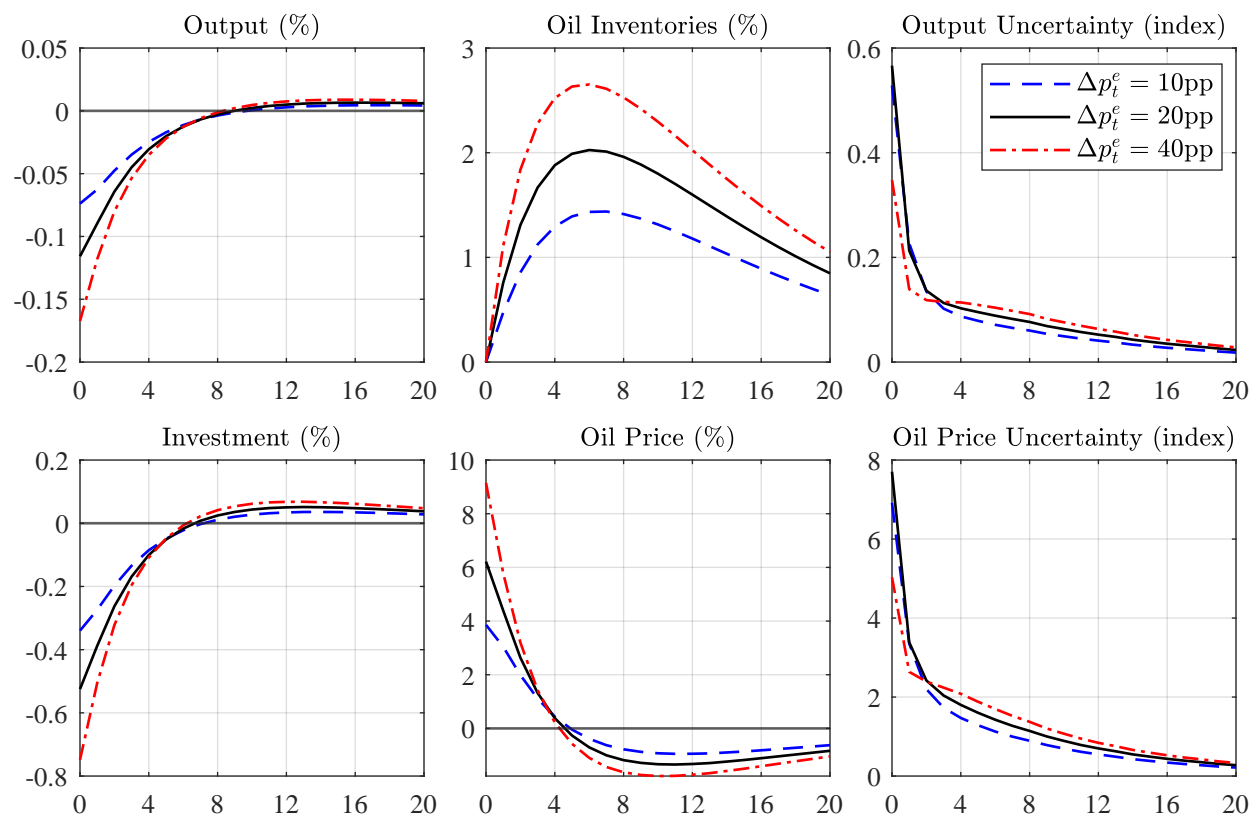
5 WHAT IS THE ROLE OF UNCERTAINTY SHOCKS?

As discussed in [Section 3](#), economists' intuition about the impact of oil price uncertainty shocks continues to be based on insights from partial equilibrium models or from general equilibrium models with exogenous oil price volatility shocks, no fluctuations in macroeconomic uncertainty, and two-sided risk. This section examines the role of uncertainty shocks within the context of a general equilibrium model with endogenous time-varying uncertainty about the price of oil and output driven by macroeconomic and oil production disasters.

The model provides a useful benchmark for what we would expect the relationship between oil price uncertainty and output uncertainty to be in the data. It is capable of generating fluctuations in oil price uncertainty that are qualitatively similar to those in the data, and it does not force output uncertainty and oil price uncertainty to be independent. The model helps us understand to what extent their relationship is driven by exogenous shifts in macroeconomic risk, exogenous shifts in geopolitical risk in oil markets, or by other shocks.

5.1 TRANSMISSION OF UNCERTAINTY SHOCKS Economic agents will witness disasters only infrequently. Nevertheless, disasters matter for economic behavior because there is a probability that they may be realized in the future. For example, the invasion of Ukraine was followed in March 2022 by concerns that Russian oil exports may materially decline. As we highlight in this section, changes in the probability of such disasters, which generate fluctuations in uncertainty, will cause the oil market and the macroeconomy to move even when no disasters actually occur.¹⁹

¹⁹A positive shock to the probability of a geopolitically driven oil production shortfall lowers expected oil production and raises the uncertainty about the future price of oil. Such shocks differ from news about expected oil production or expected oil consumption, as discussed in the previous literature. For example, Arezki et al. (2017) estimate the domestic macroeconomic impact of large oil discoveries rather than the impact on the global economy of a shift in the probability of such a discovery. Moreover, they are not concerned with geopolitical risk nor they do they examine the response of the price of oil or of the uncertainty about the price of oil. Likewise, the literature on the impact of OPEC announcements on oil price expectations is not concerned with uncertainty nor does it focus on the anticipation of rare geopolitical oil supply disruptions (e.g., Känzig, 2021; Kilian, 2024).

Figure 4: Responses to alternative oil production disaster probability shocks

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized.

Oil Disaster Probability Figure 4 shows the responses of key model variables when the exogenous oil disaster probability is increased by 10, 20, and 40 percentage points (pp), respectively. Higher odds of an oil production disaster generate stronger storage demand, reflected in a persistent build-up of oil inventories. This raises the price of oil, with the initial increase ranging from about 4% for a 10pp shock to the disaster probability to about 9% for a 40pp shock.

The probability shock reduces investment, with the effect ranging from a 0.34% drop for a 10pp shock to 0.75% for a 40 pp shock. The negative effect on investment arises for two distinct reasons. First, an oil production disaster, if it were to occur, would reduce the return to capital, since oil and capital are complements in production. Thus, the higher probability of such a disaster lowers the expected return from investing in capital. Second, the return to capital today declines because higher anticipatory demand for oil inventories raises the price of oil. Together, these two

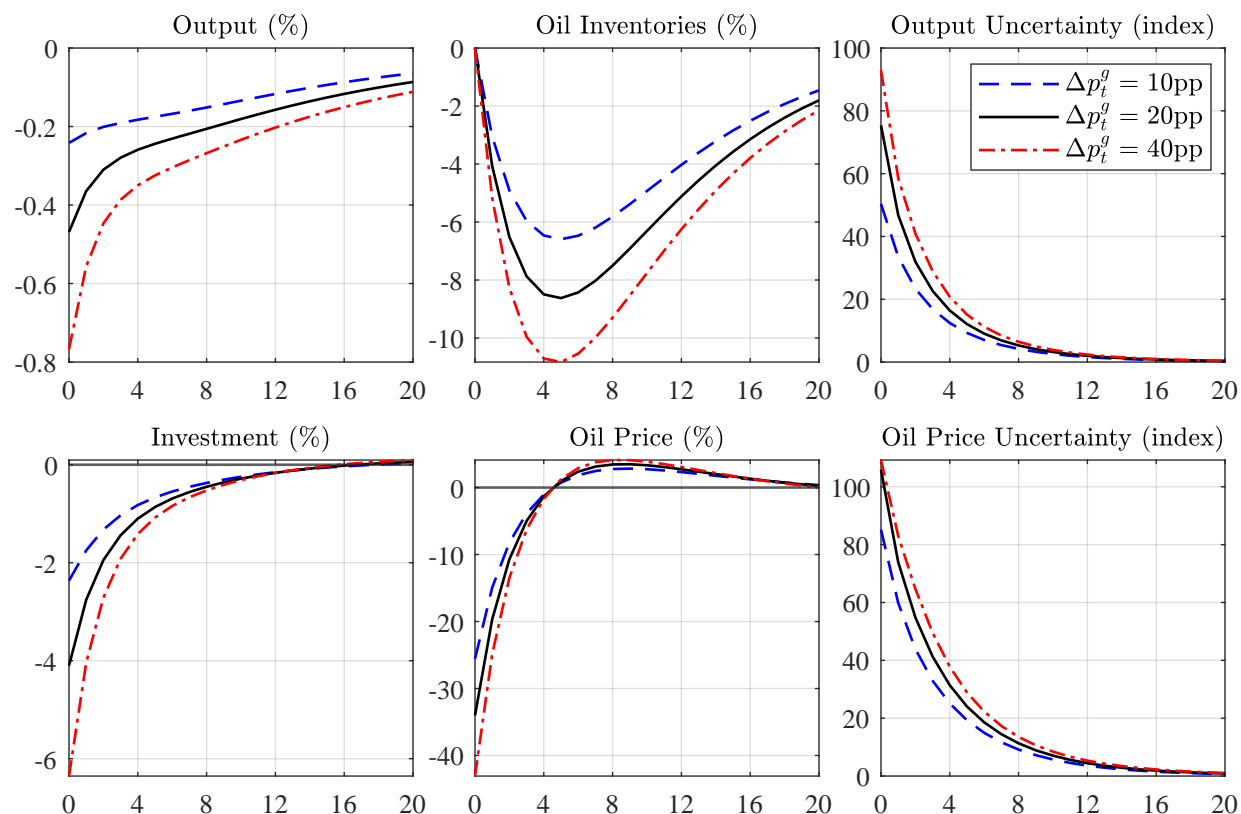
effects push down output, but the overall magnitude is modest.

Higher odds of an oil production disaster raise output uncertainty and oil price uncertainty, but the effects on oil price uncertainty are much larger. For example, a 10pp increase in the probability of an oil disaster raises the output uncertainty index only by 0.35, but the oil price uncertainty index by 5. Thus, the oil disaster probability shock looks in some ways like an exogenous oil price uncertainty shock. However, as discussed later, the downside risk inherent in the oil production disaster leads to very different responses compared to a stochastic volatility shock to oil production.

The model shows that the recessionary effects of the probability shock are reflected in output immediately, but are short-lived. The responses do not change proportionately with the shock size. For example, the responses of output and the price of oil to a 40pp increase in the oil disaster probability are only about 2 times larger than when the probability rises by 10pp. This result highlights that exogenous variation in uncertainty transmits to the macroeconomy nonlinearly.

Raising the disaster probability beyond a certain point lowers oil price uncertainty because the agents in the model perceive the disaster as increasingly likely. At the same time, the effect of the disaster probability shock on the oil price and real activity strengthens and eventually converges to the effects that occur when an average oil production disaster is realized. As shown in Appendix F, an average oil disaster causes a 14% increase in the price of oil and a 0.22% decline in output. However, unlike an oil disaster probability shock, inventories sharply decline as oil supply is cut.

Growth Disaster Probability A growth disaster such as the Financial Crisis of 2008 acts like a negative demand shock in the oil market by reducing real activity and lowering oil demand. This plays a key role in understanding how the oil market responds to an increased probability of an output growth disaster. [Figure 5](#) shows the responses when the growth disaster probability exogenously increases by 10, 20, and 40pp, respectively. All three shocks have substantial, albeit short-lived, effects on the price of oil. For example, a 10pp increase in the probability causes the price of oil to decline by 26% on impact. There are two related but somewhat distinct mechanisms at play. First, as in Gourio (2012), the higher probability directly reduces the expected return to capital, which lowers oil demand today since capital and oil are complements. Second, lower

Figure 5: Responses to alternative growth disaster probability shocks

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized.

current and expected oil demand also reduces the expected return from holding oil inventories. As a result, oil currently held in storage is sold off, pushing down the oil price even further.

Although the reduction in the price of oil is beneficial for the economy, the net effect of this probability shock on output is negative. In fact, the decline in output is much larger than from an oil disaster probability shock of the same magnitude. This is because the growth probability shock transmits directly to output rather than through the share of oil in output, which is small. In addition, the response is more persistent than the response to an oil disaster probability shock.

In related work, Gourio (2012) showed that an increase in the probability of a growth disaster causes the uncertainty about equity prices to rise. Our results show that the same shock also has a major effect on output uncertainty and oil price uncertainty. If the price of oil and oil price uncertainty were exogenous, this interaction between the uncertainty measures would not occur.

As in [Figure 4](#), the responses do not scale proportionately with the increase in the growth disaster probability. For example, a 10pp increase leads to a 26% decline in the price of oil, whereas the price of oil declines by 43% when the probability rises by 40pp. This is true for the other variables as well, once again highlighting the nonlinearity in the transmission of uncertainty.

The key difference between the two disaster probability shocks is that the growth disaster probability shock has substantial effects on both uncertainty variables, whereas the oil disaster probability shock does not. This suggests that the comovement between oil price uncertainty and output uncertainty tends to reflect shifts in macroeconomic risk rather than geopolitical risk.

External validation of the oil market responses An important question is whether the responses of the real price of oil and of oil inventories reported in [Figures 4](#) and [5](#) (and the corresponding responses associated with the realization of a disaster shown in [Appendix F](#)) align with earlier empirical evidence. A natural benchmark are the impulse responses reported in [Zhou \(2020\)](#) based on the workhorse structural VAR model of the global oil market developed in [Kilian and Murphy \(2014\)](#). These response estimates are consistent with those implied by our micro-founded model.

The responses to a positive shock to the oil production disaster probability may be interpreted as an example of a positive storage demand shock in the Kilian-Murphy model, which is expected to raise oil inventories and the price of oil, consistent with the responses in [Figure 4](#). Likewise, the responses to an increase in the growth disaster probability represent an example of a negative storage demand shock that lowers oil inventories and the price of oil, consistent with the responses in [Figure 5](#). Both shocks involve anticipation and hence affect the oil market through storage.

The realization of an oil production disaster ([Appendix F](#), [Figure 2](#)), in contrast, is best thought of as a negative flow supply shock in the Kilian-Murphy model that lowers inventories and raises the price of oil, consistent with the evidence in [Zhou \(2020\)](#). Finally, the realization of a growth disaster ([Appendix F](#), [Figure 3](#)) instead resembles a negative flow demand shock in the Kilian-Murphy framework associated with lower oil prices and inventories, consistent with the evidence in [Zhou \(2020\)](#). Thus, in all cases our DSGE model replicates the sign of the oil price and inventory responses found in the structural VAR literature. Of course, the persistence and magnitudes of the

Table 3: Decomposition of key volatilities

Moment	Data	Baseline Model	No Output Disaster Risk	No Output Disaster Risk or Oil Production Disaster Risk
$SD(\Delta y)$	0.74	0.87	0.65	0.64
$SD(\Delta i)$	1.95	1.92	1.26	1.23
$SD(\Delta o^s)$	2.01	2.14	2.12	1.12
$SD(\Delta p^o)$	14.39	14.27	6.32	5.64
$SD(\mathcal{U}_y)$	14.51	15.61	0.30	0.14
$SD(\mathcal{U}_{p^o})$	29.95	30.49	4.98	2.01

Notes: The models without disaster risk remove both the probability shock and the disaster state. $SD(\mathcal{U}_y)$ and $SD(\mathcal{U}_{p^o})$ are normalized by $SD(\Delta y)$ and $SD(\Delta p^o)$ in the baseline model, respectively, to be consistent with Jurado et al. (2015).

responses need not match, as our model focuses on specific examples of storage demand, flow demand, and flow supply shocks rather than typical shocks of this type. The advantage of working with a fully-specified model is that it makes explicit how oil price expectations are formed and transmitted to the oil market and the economy. This link is left implicit in the structural VAR model. As a result, we are able to quantify how variation in disaster probabilities changes economic outcomes, which would be impossible based on the structural VAR model.²⁰

5.2 HOW MUCH VOLATILITY IS DUE TO OUTPUT RISK AND GEOPOLITICAL RISK? Table 3 shows that the model generally does an excellent job at capturing the volatility in the data. Dropping output disaster risk from the model substantially lowers the ability of the model to explain the volatility in the data. The resulting model not only substantially understates the standard deviation of the two uncertainty series, but it also understates most other volatilities.

Dropping both output and oil production disasters from the model further lowers the volatilities. In particular, it removes almost all variability in the two uncertainty measures and some of the variability in oil price growth and oil production growth but has little additional effect on the

²⁰Studying the evolution of oil inventories and the oil price during geopolitical disasters would not be informative about the realism of our model because the anticipation of an oil production disaster causes inventories to move in the opposite direction from its realization. Moreover, the behavior of oil inventories and prices during geopolitical disasters may also reflect past and current macroeconomic shocks. Thus, the realism of the responses in our model can only be evaluated based on the conditional moments in the data, which requires a structural VAR model.

volatility of macroeconomic aggregates.

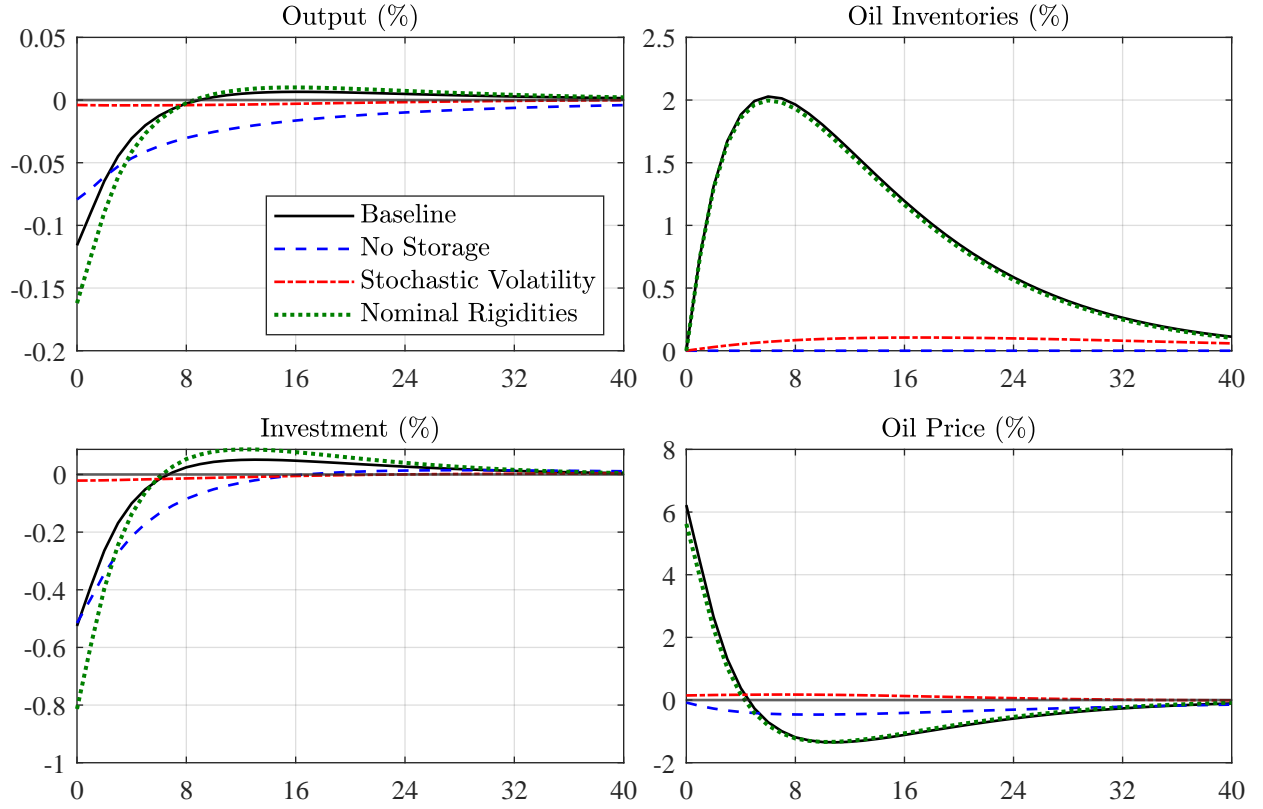
There are two key takeaways from these results. First, output disaster risk is a major driver of fluctuations in oil price uncertainty, highlighting that oil price uncertainty is not exogenous as is often assumed in the literature. Second, oil production disaster risk is not a major driver of fluctuations in macro aggregates or output uncertainty, suggesting that oil price uncertainty does not play a major role in driving business cycles.

5.3 ALTERNATIVE MODEL SPECIFICATIONS In this section, we highlight key features of our model by considering alternative specifications. We first illustrate the central role of oil storage. We then contrast our model with earlier models incorporating stochastic volatility shocks. We also show that the introduction of nominal rigidities does not affect the substance of our results.

Role of storage There are important movements in oil inventories whenever the probability of a disaster increases. These movements affect the price of oil and, therefore, the evolution of macroeconomic aggregates. In this section, we investigate how important storage is for those responses by comparing the baseline results to those from a model without storage.

Figure 6 shows the responses for the oil disaster probability shock. The key difference is that the price of oil slightly declines on impact in the model without storage, whereas it increases substantially in the baseline model. In the absence of storage, the response of the oil price is driven entirely by the expectation of lower output, which reduces the demand for oil and modestly lowers its price. Given the muted response of the price of oil, the impact effect on output is also reduced.

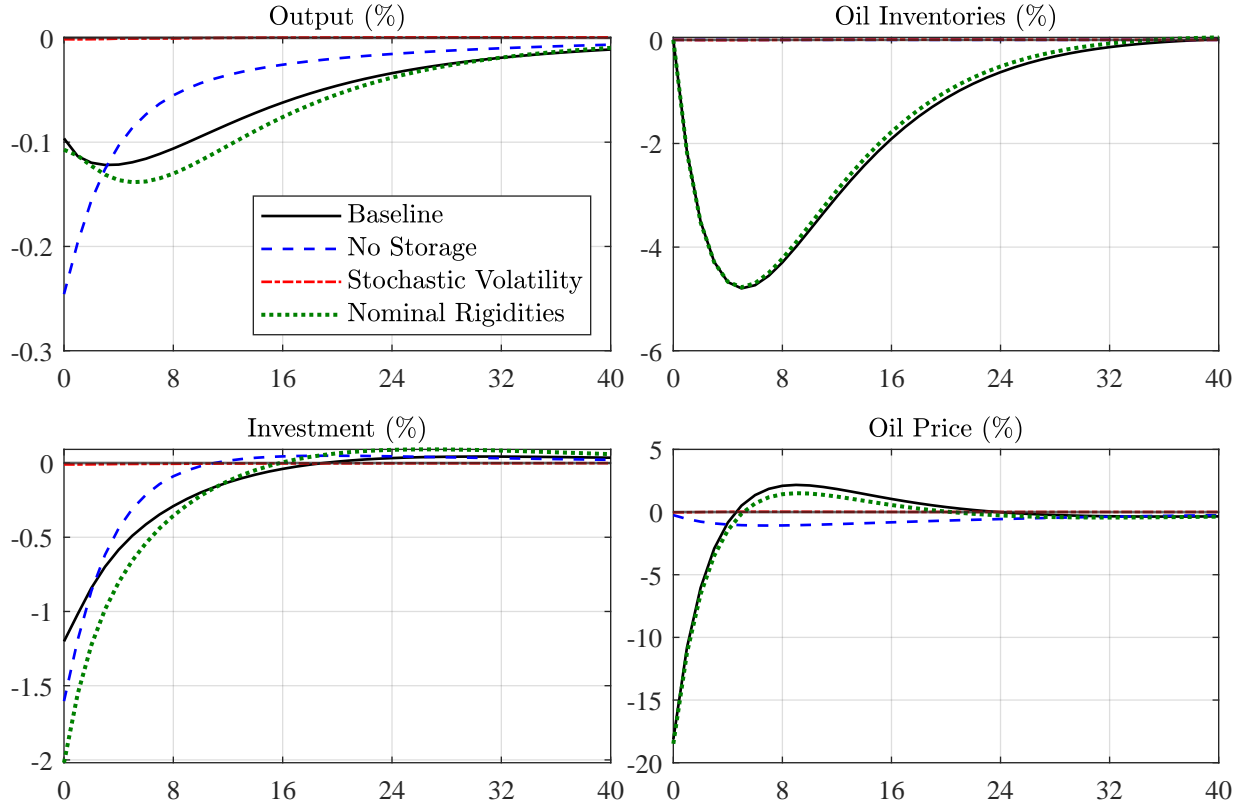
Figure 7 shows that storage also plays a key role in the propagation of a growth disaster probability shock. In the model with storage, a higher probability of a disaster leads to a reduction in oil inventories due to the greater likelihood of a recession. This causes a substantial decline in the price of oil, which does not occur in the no-storage model. Since a lower price of oil offsets some of the negative effects of this shock on the macroeconomy, the impact effect on output is larger when the model does not contain storage. Overall, our results demonstrate that storage is a key ingredient for understanding the effects of uncertainty in both the oil market and the macroeconomy.

Figure 6: Responses to oil production disaster probability and stochastic volatility shocks


Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized. The stochastic volatility shock has been normalized to match the impact response of oil price uncertainty in the baseline model. All disaster probability shocks are 10pp.

Role of downside risk Stochastic volatility (SV) is an alternative way of generating time-varying oil price uncertainty that has been used in previous studies (e.g., Başkaya et al., 2013; Gao et al., 2022; Plante and Traum, 2012). In this section, we compare the results from the baseline model to those from a model where uncertainty is generated by SV. Specifically, we introduce an exogenous volatility shock into productivity growth and oil production growth,

$$\begin{aligned}\ln g_t &= \ln \bar{g} + \sigma_{g,t-1} \varepsilon_{g,t}, \\ \ln g_{o,t} &= \ln \kappa_0 + \kappa_1 \ln g_t + \kappa_2 \ln \epsilon_{t-1} + \sigma_{go,t-1} \varepsilon_{go,t}, \\ \ln \sigma_{g,t} &= (1 - \rho_{sv}^g) \ln \bar{\sigma}_g + \rho_{sv}^g \ln \sigma_{g,t-1} + \sigma_{sv}^g \varepsilon_{sv,t}^g, \\ \ln \sigma_{go,t} &= (1 - \rho_{sv}^{go}) \ln \bar{\sigma}_{go} + \rho_{sv}^{go} \ln \sigma_{go,t-1} + \sigma_{sv}^{go} \varepsilon_{sv,t}^{go},\end{aligned}$$

Figure 7: Responses to growth disaster probability and stochastic volatility shocks

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized. The stochastic volatility shock has been normalized to match the impact response of output uncertainty in the baseline model. All disaster probability shocks are 5pp.

where all shocks are standard normally distributed.²¹ The parameters of the level processes are unchanged. The persistence of both SV processes, ρ_{sv}^g and ρ_{sv}^e , are equal to the persistence of the disaster probability processes. The standard deviation of the growth SV shock, σ_{sv}^g , is set to 0.095 to match the volatility of output uncertainty. Analogously, the standard deviation of the oil production SV shock, σ_{sv}^{go} , is set to 0.145 to match the volatility of oil price uncertainty.

Figure 6 compares the responses to an SV shock in oil production, $\varepsilon_{sv,t}^{go}$, to the responses to our baseline oil disaster probability shock, $\varepsilon_{p,t}^e$. The $\varepsilon_{sv,t}^{go}$ shock is set so the SV specification generates the same impact effect on oil price uncertainty as the oil disaster probability shock. Qualitatively,

²¹Stochastic volatility has also been used to model exogenous uncertainty shocks in a number of other settings including fiscal policy (Fernández-Villaverde et al., 2015), monetary policy (Mumtaz and Zanetti, 2013), household preferences (Basu and Bundick, 2017), and the global interest rate (Fernández-Villaverde et al., 2011).

these shocks move the model variables in the same direction, but there are quantitatively significant differences. While the SV shock naturally generates sizable fluctuations in oil price uncertainty, it has little effect on the macroeconomy and the oil market.

The key difference between the two modeling choices is that an oil disaster introduces a source of downside risk into the economy because it makes a sharp drop in oil production more likely. As a result, when the probability of a disaster increases, it not only increases uncertainty but also shifts the conditional mean of economic outcomes. This generates a stronger precautionary demand motive, which pushes up the price of oil and lowers output. The SV shock, on the other hand, is akin to a mean-preserving spread. It generates a sizable increase in uncertainty but has little effect on the conditional mean. Hence, the responses of the price of oil and output are muted. A similar result holds when replacing the growth disaster probability shock with a SV shock on productivity growth, as shown in [Figure 7](#). We conclude that SV shocks are unable to capture the effects of increases in uncertainty associated with major geopolitical events that affect the oil market.²²

Nominal Rigidities The literature has emphasized that sticky prices and monopolistic competition are potentially important amplifiers of uncertainty shocks (e.g., Basu and Bundick, 2017; Born and Pfeifer, 2014; Fernández-Villaverde et al., 2015; Leduc and Liu, 2016). To explore how these features affect our results, we add a New Keynesian block to our baseline model. This extension replaces the expressions for the factor prices and the goods market clearing condition with

$$\begin{aligned} w_t &= (1 - \xi)mc_t y_t / n_t, \\ p_t^o &= \xi \alpha mc_t \frac{(o_t/o_0)^{1-1/\sigma}}{(1-\alpha)(k_t/k_0)^{1-1/\sigma} + \alpha(o_t/o_0)^{1-1/\sigma}} \frac{y_t}{o_t}, \\ r_t^k &= \xi(1 - \alpha)mc_t \frac{(k_t/k_0)^{1-1/\sigma}}{(1-\alpha)(k_t/k_0)^{1-1/\sigma} + \alpha(o_t/o_0)^{1-1/\sigma}} \frac{y_t}{k_t}, \\ c_t + i_t &= \left(1 - \frac{\varphi}{2} \left(\frac{\pi_t}{\pi} - 1\right)^2\right) y_t, \end{aligned}$$

²²Similar to our responses for the SV specification, Gao et al. (2022) find small impacts of oil production volatility shocks in their baseline model. They show that the responses are amplified when markups are assumed to be time-varying such that the markup falls with oil consumption. The responses are even larger when level and volatility shocks to oil production are also assumed to be negatively correlated. The empirical support for these assumptions is not clear. For example, time-varying markups are a standard feature of New Keynesian models. Plante and Traum (2014) examine such a model where oil is an intermediate input and find that SV shocks to the price of oil have a negligible effect on the macroeconomy.

where mc_t is marginal cost, π_t is the gross inflation rate, and φ scales the price adjustment cost.

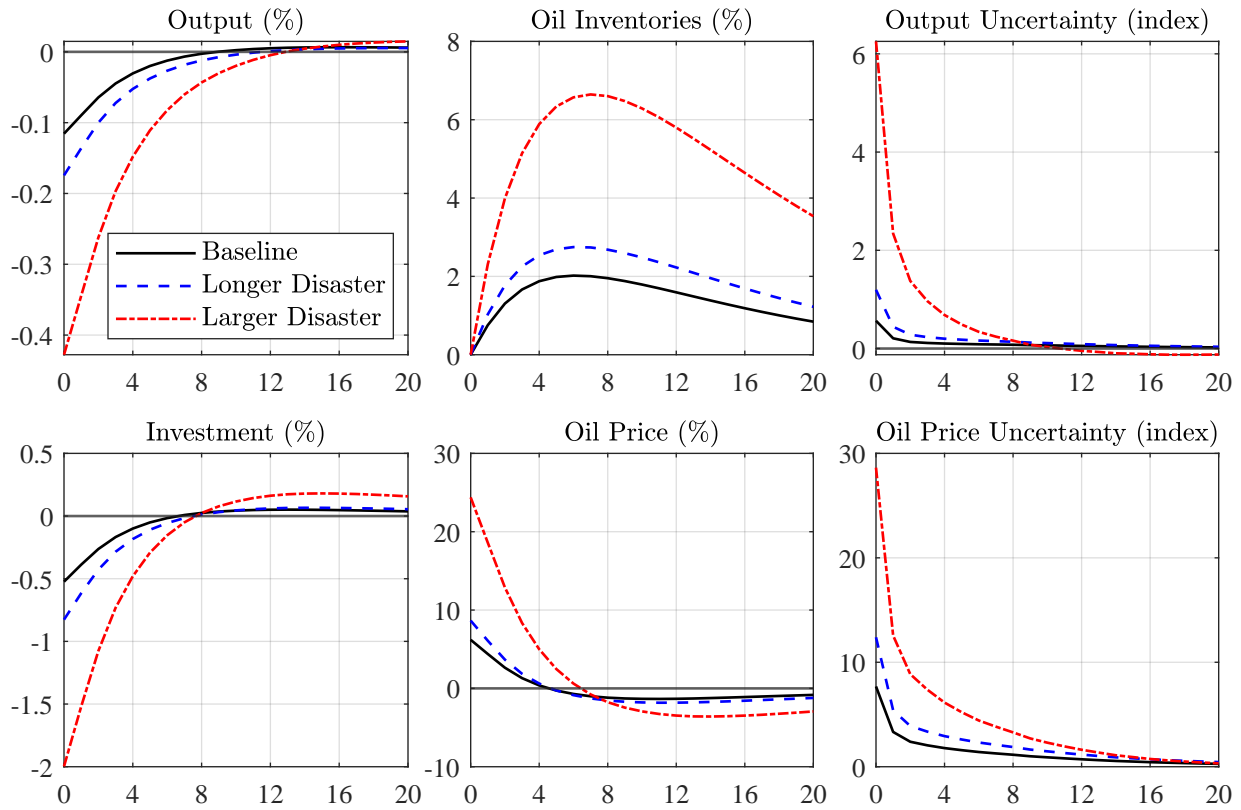
The New Keynesian model also adds a Phillips curve, an Euler equation for the nominal bond, and a Taylor rule given by

$$\begin{aligned}\varphi \left(\frac{\pi_t}{\bar{\pi}} - 1 \right) \frac{\pi_t}{\bar{\pi}} &= 1 - \theta + \theta mc_t + \varphi E_t \left[x_{t,t+1} \left(\frac{\pi_{t+1}}{\bar{\pi}} - 1 \right) \frac{y_{t+1}}{y_t} \frac{\pi_{t+1}}{\bar{\pi}} \right] \\ 1 &= E_t [x_{t+1} r_t^n / \pi_{t+1}] \\ r_t^n &= \bar{r}^n (\pi_t / \bar{\pi})^{\phi_\pi},\end{aligned}$$

where r_t^n is the gross nominal interest rate, θ is the elasticity of substitution between intermediate goods, and ϕ_π is the monetary response to inflation. We set the additional parameters to conventional values in the literature. Specifically, we set $\theta = 6$ to achieve a 20% steady-state markup, $\varphi = 80$ so prices adjust about once a year in a (linear) Calvo setting, $\phi_\pi = 1.5$ in line with a standard Taylor rule, and $\bar{\pi} = 1.005$ consistent with a 2% annual inflation target.

As shown in Figures 6 and 7, the presence of nominal rigidities slightly amplifies the responses of investment and output to both types of disaster probability shocks. In particular, a 20pp oil disaster shock lowers investment by 0.52% and output by 0.12% in our baseline model, compared to 0.81% and 0.16% in the model with nominal rigidities. While these differences confirm the insights from the literature, they do not alter our finding that oil disaster probability shocks have modest effects on the macroeconomy.

5.4 ALTERNATIVE OIL DISASTER SPECIFICATIONS The surge in oil price uncertainty observed after the invasion of Kuwait in August 1990 far exceeded the increase in oil price uncertainty implied by our baseline model, which is calibrated to match the average magnitude and duration of past geopolitically driven oil production disasters. This suggests that the oil market in 1990 was concerned about a much larger or longer lasting disaster than the average realized oil production disaster observed in the data. Figure 8 explores this question by considering alternative specifications of the oil production disaster. We first consider a disaster with the same 5% magnitude as in the baseline model, but with an expected duration of 10 quarters rather than 3 quarters.

Figure 8: Responses to an oil production disaster probability shock under different disasters

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized. The disaster probability shock is 20pp for all three specifications.

We then, alternatively, consider a much larger oil production disaster of 20% of global oil production with the same expected duration as in the baseline model. A 20% shortfall would have been the approximate share of world oil production at risk if Iraq had succeeded in conquering not only Kuwait in 1990, but also Saudi Arabia and its smaller neighbors along the Persian Gulf such as Qatar, Bahrain, and the UAE. It also corresponds to the share of world oil production at risk from Iranian retaliation in the event of an Israeli attack on Iranian oil export facilities in 2024.

While making the oil production disaster longer lasting only modestly raises the effects on the global economy and the uncertainty measures compared to the baseline, making the oil production disaster larger increases the effect on oil price uncertainty six-fold. This brings the response much more in line with the observed spike in oil price uncertainty in 1990 (see [Figure 2](#)). At the same time, the responses of the price of oil and macroeconomic aggregates roughly quadruple in magni-

tude. The increase in the price of oil rises from 6% on impact to 24%, and the reduction in output rises from 0.12% to 0.43%. As the disaster probability increases, these responses increase and converge to what happens if the disaster is realized, which is shown in Appendix F. For example, with a 20pp (40pp) probability shock, the price of oil rises by 24% (34%) and output declines by 0.43% (0.62%). If the disaster is realized, the price of oil rises by 49% and output falls by 0.73%.

The specification of the oil disaster also affects the ability of oil production disaster risk to explain variation in oil price uncertainty. While oil production disaster risk accounts for only 16% of this variability in the baseline model, as shown in [Table 3](#), this share rises to 22% when the disaster is longer lasting and to 69% when the magnitude of the disaster is increased. We conclude that geopolitically driven oil disasters play a more important role when agents expect, at least sometimes, a larger oil production disaster than observed on average in the historical data. Changing the magnitude of the oil disaster, however, does not alter our finding that a large portion of the variability in oil price uncertainty is driven by macroeconomic uncertainty.

6 RELATIONSHIP TO THE EXISTING LITERATURE

We are not the first to study the transmission of oil price uncertainty shocks within a general equilibrium model. For example, Başkaya et al. (2013) analyze the business cycle implications of oil price uncertainty based on a stylized oil-importing small open economy. One important difference is that oil price uncertainty is exogenous in their model. In contrast, our analysis shows that oil price uncertainty in general depends on level shocks as well as macroeconomic uncertainty shocks. Allowing for both macroeconomic and oil price uncertainty is important when assessing their respective roles in explaining economic fluctuations, as illustrated in [Section 5.2](#). Another important difference is their use of SV shocks. Our results in [Section 5.3](#) show that models with SV are unable to capture the responses to macroeconomic and oil production disaster probability shocks.

Our model shares some features with the general equilibrium model in Gao et al. (2022) in that both models allow the price of oil to be endogenously determined and incorporate uncertainty shocks. As with Başkaya et al. (2013), a key difference is that their model features SV shocks to

oil production and productivity growth rather than downside risk. This simplification makes it impossible to explore the relationship between oil price uncertainty and macroeconomic uncertainty and to quantify the effects of increases in uncertainty associated with downside risks.

The study related most closely to ours is perhaps Olovsson (2019) with the important difference that he does not model uncertainty shocks. While his model allows for oil storage and oil production disasters, it does not allow the disaster probability to vary over time, nor does it include macroeconomic disasters. These two features are crucial for understanding the effects of shifts in geopolitical risk and for answering the questions raised by policymakers and market participants.

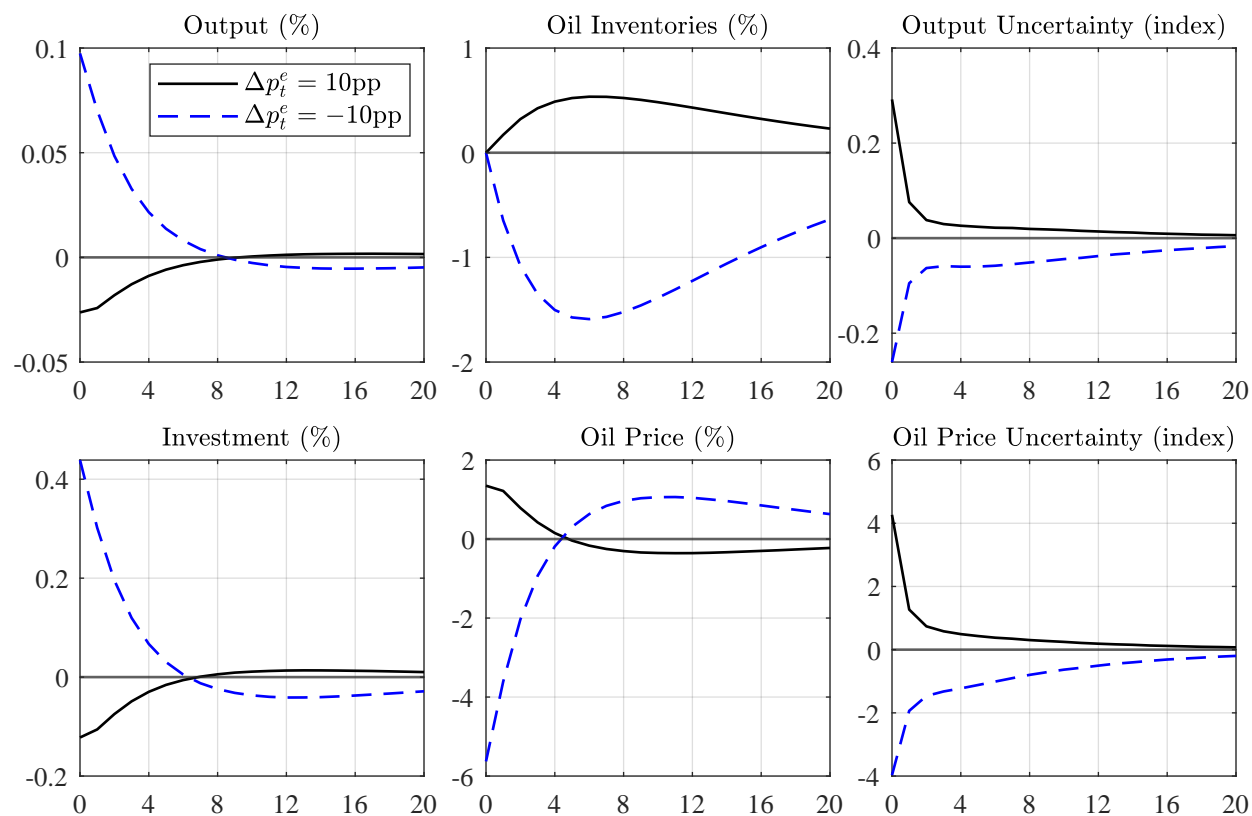
Our model also differs from the general equilibrium model in Ready (2018) who does not study how the economy responds to shifts in the likelihood of a major temporary drop in oil production driven by geopolitical events, but instead examines how discrete shifts in the uncertainty about long-run oil supplies affect asset prices and the economy in a model with long-run risk.²³

Finally, what sets our work apart from all these earlier studies is that we address the implications of the endogeneity of oil price uncertainty for the identification of oil price uncertainty shocks in empirical work. This question, which we turn to next, has not been addressed in prior work.

7 IMPLICATIONS FOR EMPIRICAL WORK

Our analysis highlights that oil price uncertainty endogenously responds not only to exogenous uncertainty about future oil production driven by geopolitical events, but also to exogenous uncertainty about the future path of the economy. Thus, geopolitically driven oil price uncertainty shocks differ in general from shocks to observed oil price uncertainty, as measured by the method of Jurado et al. (2015) or the OVX oil volatility index. This result is consistent with practitioners' understanding that uncertainty about the oil price reflects not only uncertainty about future oil production, but also uncertainty about future oil consumption driven by macroeconomic, financial, and policy uncertainty. For example, market commentators have routinely highlighted the role of uncertainty about the prospects of the Chinese economy and whether the U.S. economy is about

²³In related work, Hamann et al. (2023) examine the effects of sovereign default risk in oil-producing economies.

Figure 9: Responses to positive and negative oil production disaster probability shocks

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized.

to enter a recession in assessing the uncertainty about the price of oil. Perhaps less obviously, the model also shows that oil price uncertainty responds to level shocks in the macroeconomy, as shown in the responses to a growth disaster in Appendix F. These results invalidate the premise of exogenous oil price uncertainty shocks and cast doubt on the ability of standard empirical models to correctly identify exogenous oil price uncertainty shocks.

The endogeneity of oil price uncertainty shocks is not the only concern with these models. For example, VAR models with GARCH errors, as in Elder and Serletis (2010), postulate that every level shock to the price of oil is also an oil price uncertainty shock. In our model, level and uncertainty shocks affect oil price uncertainty differently. Elder and Serletis also impose the restriction that positive and negative oil price shocks both increase oil price uncertainty, which is inconsistent with our model. Figure 9 illustrates this point by comparing the responses to a $\pm 10\text{pp}$

disaster probability shock.²⁴ A decrease in the probability of an oil disaster reduces the price of oil and oil price uncertainty on impact, while an increase in this probability raises the oil price and oil price uncertainty. The same result holds for a growth disaster probability shock, as shown in Appendix F.

Models that break the link between level and uncertainty shocks such as the VAR model with SV in Jo (2014) postulate that empirical measures of oil price uncertainty shocks and shocks to the level of the price of oil are independent. In contrast, we show that oil price uncertainty is endogenous and may be driven by the same shocks as the price of oil. For example, an increase in the probability of a growth disaster not only raises oil price uncertainty, but also causes storage demand to increase, raising the real price of oil.²⁵

The insight that oil price uncertainty is simultaneously determined with macroeconomic aggregates applies not only to GARCH and SV models but also to recursively identified linear VAR models that order oil price uncertainty first as in Gao et al. (2022). Exploring alternative recursive orderings does not address this concern, as shown in Kilian et al. (2025). Thus, the seemingly robust empirical evidence from linear and nonlinear VAR models that oil price uncertainty shocks substantially lower real activity must be viewed with caution.

8 CONCLUDING REMARKS

There has been growing interest in the impact of shifts in geopolitical risk, particularly in the oil market. In this paper, we introduced a stochastic growth model of the global economy that is designed to examine how this risk affects oil price uncertainty and the macroeconomy. Unlike previous studies, we modeled geopolitical risk as downside risk to oil production and allowed oil price and macroeconomic uncertainty to be determined endogenously.

Our analysis highlighted the importance of jointly modeling macroeconomic and geopolitical

²⁴The simulations are initialized at a 15% oil disaster probability to permit a positive and negative shock. The state vector equals the average of periods in the ergodic distribution that are within 1pp of the initial disaster probability.

²⁵Moreover, as shown in [Section 5](#), while an increase in oil price uncertainty may be alternatively generated by an SV shock to oil production, only a disaster probability shock generates meaningful recessionary effects. Thus, one would not expect large recessionary effects of shocks to oil price uncertainty when estimating VAR models with SV.

risk. We showed that oil price uncertainty not only responds to geopolitically driven oil price uncertainty shocks but also to level shocks in the macroeconomy and to macroeconomic uncertainty shocks. In our baseline model, more than half of the observed oil price uncertainty tends to be driven by the macroeconomy, which helps explain why oil price uncertainty has historically been associated with reductions in real activity. This fact, along with other implications of our model discussed in the paper, calls into question standard empirical models of the transmission of oil price uncertainty shocks that provide seemingly robust evidence supporting the conventional wisdom that oil price uncertainty plays a major role in driving the business cycle.

Our results suggest that economists and policymakers need to rethink the role of geopolitical oil price risk in the global economy and be cognizant of the interplay between oil price uncertainty, macroeconomic uncertainty, and the state of the economy. Notwithstanding the attention geopolitical events in oil markets have attracted, we conclude that geopolitical oil price risk is unlikely to generate sizable recessionary effects. Even when considering a geopolitically driven disaster involving a 20% drop in global oil production, which far exceeds the magnitude of historically observed oil production shortfalls, it takes a large increase in the probability of this disaster or its actual realization to cause a sizable recessionary impact.

REFERENCES

- ABIAD, A. AND I. A. QURESHI (2023): “The Macroeconomic Effects of Oil Price Uncertainty,” *Energy Economics*, 125.
- AGUIAR-CONRARIA, L. AND Y. WEN (2007): “Understanding the Large Negative Impact of Oil Shocks,” *Journal of Money, Credit and Banking*, 39, 925–944.
- ALFARO, I., N. BLOOM, AND X. LIN (2024): “The Finance Uncertainty Multiplier,” *Journal of Political Economy*, 132, 577–615.
- ALQUIST, R. AND L. KILIAN (2010): “What Do We Learn From the Price of Crude Oil Futures?” *Journal of Applied Econometrics*, 25, 539–573.
- ALQUIST, R., L. KILIAN, AND R. J. VIGFUSSON (2013): “Forecasting the Price of Oil,” in *Handbook of Economic Forecasting*, ed. by G. Elliott, C. Granger, and A. Timmermann, North-Holland, vol. 2, 427–507.
- AREZKI, R., V. A. RAMEY, AND L. SHENG (2017): “News Shocks in Open Economies: Evidence from Giant Oil Discoveries,” *Quarterly Journal of Economics*, 132, 103–155.
- ATKESON, A. AND P. J. KEHOE (1999): “Models of Energy Use: Putty-Putty versus Putty-Clay,” *American Economic Review*, 89, 1028–1043.
- BACKUS, D. K. AND M. J. CRUCINI (2000): “Oil Prices and the Terms of Trade,” *Journal of International Economics*, 50, 185–213.

- BAŞKAYA, Y. S., T. HÜLAGÜ, AND H. KÜÇÜK (2013): “Oil Price Uncertainty in a Small Open Economy,” *IMF Economic Review*, 61, 168–198.
- BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): “Measuring Economic Policy Uncertainty,” *Quarterly Journal of Economics*, 131, 1593–1636.
- BARRO, R., E. NAKAMURA, J. STEINSSON, AND J. URSÚA (2013): “Crises and Recoveries in an Empirical Model of Consumption Disasters,” *American Economic Journal: Macroeconomics*, 5, 35–74.
- BARRO, R. J. AND G. Y. LIAO (2021): “Rare Disaster Probability and Options Pricing,” *Journal of Financial Economics*, 139, 750–769.
- BARRO, R. J. AND J. F. URSÚA (2012): “Rare Macroeconomic Disasters,” *Annual Review of Economics*, 4, 83–109.
- BARSKY, R. B. AND L. KILIAN (2002): “Do We Really Know That Oil Caused the Great Stagflation? A Monetary Alternative,” in *NBER Macroeconomics Annual 2001, Volume 16*, ed. by M. Eichenbaum, E. Hurst, and V. Ramey, University of Chicago Press, 137–198.
- BASU, S. AND B. BUNDICK (2017): “Uncertainty Shocks in a Model of Effective Demand,” *Econometrica*, 85, 937–958.
- BAUMEISTER, C. AND L. KILIAN (2016a): “Lower Oil Prices and the U.S. Economy: Is This Time Different?” *Brookings Papers on Economic Activity*, 287–357.
- (2016b): “Understanding the Decline in the Price of Oil since June 2014,” *Journal of the Association of Environmental and Resource Economists*, 3, 131–158.
- BERGER, D., I. DEW-BECKER, AND S. GIGLIO (2020): “Uncertainty Shocks as Second-Moment News Shocks,” *Review of Economic Studies*, 87, 40–76.
- BERNANKE, B. S. (1983): “Irreversibility, Uncertainty, and Cyclical Investment,” *Quarterly Journal of Economics*, 98, 85–106.
- BERNANKE, B. S., M. GERTLER, AND M. WATSON (1997): “Systematic Monetary Policy and the Effects of Oil Price Shocks,” *Brookings Papers on Economic Activity*, 28, 91–157.
- BERNSTEIN, J., M. PLANTE, A. W. RICHTER, AND N. A. THROCKMORTON (2024): “A Simple Explanation of Countercyclical Uncertainty,” *American Economic Journal: Macroeconomics*, 16, 143–171.
- BIANCHI, F., C. L. ILUT, AND M. SCHNEIDER (2018): “Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle,” *Review of Economic Studies*, 85, 810–854.
- BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77, 623–685.
- (2014): “Fluctuations in Uncertainty,” *Journal of Economic Perspectives*, 28, 153–176.
- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2018): “Really Uncertain Business Cycles,” *Econometrica*, 86, 1031–1065.
- BORN, B. AND J. PFEIFER (2014): “Policy Risk and the Business Cycle,” *Journal of Monetary Economics*, 68, 68–85.
- (2021): “Uncertainty Driven Business Cycles: Assessing the Markup Channel,” *Quantitative Economics*, 12, 587–623.
- CALDARA, D. AND M. IACOVIELLO (2022): “Measuring Geopolitical Risk,” *American Economic Review*, 112, 1194–1225.
- CASASSUS, J., P. COLLIN-DUFRESNE, AND B. R. ROUTLEDGE (2018): “Equilibrium Commodity Prices with Irreversible Investment and Non-linear Technologies,” *Journal of Banking & Finance*, 95, 128–147.
- COLEMAN, II, W. J. (1991): “Equilibrium in a Production Economy with an Income Tax,” *Econometrica*, 59, 1091–1104.

- CROCE, M. M. (2014): “Long-Run Productivity Risk: A New Hope for Production-Based Asset Pricing?” *Journal of Monetary Economics*, 66, 13–31.
- CROSS, J. L., B. H. NGUYEN, AND T. D. TRAN (2022): “The Role of Precautionary and Speculative Demand in the Global Market for Crude Oil,” *Journal of Applied Econometrics*, 37, 882–895.
- DAVIS, S. J. AND J. HALTIWANGER (2001): “Sectoral Job Creation and Destruction Responses to Oil Price Changes,” *Journal of Monetary Economics*, 48, 465–512.
- DE GROOT, O., A. RICHTER, AND N. THROCKMORTON (2018): “Uncertainty shocks in a model of effective demand: Comment,” *Econometrica*, 86, 1513–1526.
- DIEBOLD, F. X. AND L. KILIAN (2001): “Measuring Predictability: Theory and Macroeconomic Applications,” *Journal of Applied Econometrics*, 16, 657–669.
- DRAKOS, K. AND P. T. KONSTANTINOU (2013): “Investment Decisions in Manufacturing: Assessing the Effects of Real Oil Prices and Their Uncertainty,” *Journal of Applied Econometrics*, 28, 151–165.
- EDELSTEIN, P. AND L. KILIAN (2009): “How Sensitive are Consumer Expenditures to Retail Energy Prices?” *Journal of Monetary Economics*, 56, 766–779.
- ELDER, J. AND A. SERLETIS (2010): “Oil Price Uncertainty,” *Journal of Money, Credit and Banking*, 42, 1137–1159.
- EPSTEIN, L. G. AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937–69.
- FARHI, E. AND X. GABAIX (2016): “Rare Disasters and Exchange Rates,” *Quarterly Journal of Economics*, 131, 1–52.
- FERDERER, P. J. (1996): “Oil Price Volatility and the Macroeconomy,” *Journal of Macroeconomics*, 18, 1–26.
- FERNÁNDEZ-VILLAYERDE, J., P. GUERRÓN-QUINTANA, K. KUESTER, AND J. F. RUBIO-RAMÍREZ (2015): “Fiscal Volatility Shocks and Economic Activity,” *American Economic Review*, 105, 3352–84.
- FERNÁNDEZ-VILLAYERDE, J., P. GUERRÓN-QUINTANA, J. F. RUBIO-RAMÍREZ, AND M. URIBE (2011): “Risk Matters: The Real Effects of Volatility Shocks,” *American Economic Review*, 101, 2530–61.
- FINN, M. G. (2000): “Perfect Competition and the Effects of Energy Price Increases on Economic Activity,” *Journal of Money, Credit and Banking*, 32, 400–416.
- GAO, L., S. HITZEMANN, I. SHALIASTOVICH, AND L. XU (2022): “Oil Volatility Risk,” *Journal of Financial Economics*, 144, 456–491.
- GOURIO, F. (2012): “Disaster Risk and Business Cycles,” *American Economic Review*, 102, 2734–2766.
- (2013): “Credit Risk and Disaster Risk,” *American Economic Journal: Macroeconomics*, 5, 1–34.
- GOURIO, F., M. SIEMER, AND A. VERDELHAN (2013): “International Risk Cycles,” *Journal of International Economics*, 89, 471–484.
- GUO, H. AND K. L. KLIENSEN (2005): “Oil Price Volatility and U.S. Macroeconomic Activity,” *Federal Reserve Bank of St. Louis Review*, 87, 669–684.
- HAMANN, F., J. C. MENDEZ-VIZCAINO, E. G. MENDOZA, AND P. RESTREPO-ECHAVARRIA (2023): “Natural Resources and Sovereign Risk in Emerging Economies: A Curse and a Blessing,” NBER Working Paper 31058.

- HAMILTON, J. D. (2013): “Historical Oil Shocks,” in *Routledge Handbook of Major Events in Economic History*, Routledge, 239–265.
- HASSLER, J., P. KRUSELL, AND C. OLOVSSON (2021): “Directed Technical Change as a Response to Natural Resource Scarcity,” *Journal of Political Economy*, 129, 3039–3072.
- HOOKE, M. A. (1996): “What Happened to the Oil Price-Macroeconomy Relationship?” *Journal of Monetary Economics*, 38, 195–213.
- (2002): “Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime,” *Journal of Money, Credit and Banking*, 34, 540–561.
- JERMANN, U. J. (1998): “Asset Pricing in Production Economies,” *Journal of Monetary Economics*, 41, 257–275.
- JO, S. (2014): “The Effects of Oil Price Uncertainty on Global Real Economic Activity,” *Journal of Money, Credit and Banking*, 46, 1113–1135.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105, 1177–1216.
- KÄNZIG, D. R. (2021): “The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements,” *American Economic Review*, 111, 1092–1125.
- KELLOGG, R. (2014): “The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling,” *American Economic Review*, 104, 1698–1734.
- KILIAN, L. (2009): “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, 99, 1053–69.
- (2024): “How to Construct Monthly VAR Proxies Based on Daily Surprises in Futures Markets,” *Journal of Economic Dynamics and Control*, 168, 104966.
- KILIAN, L. AND D. P. MURPHY (2014): “The Role Of Inventories And Speculative Trading In The Global Market For Crude Oil,” *Journal of Applied Econometrics*, 29, 454–478.
- KILIAN, L., M. D. PLANTE, AND A. W. RICHTER (2025): “Macroeconomic Responses to Uncertainty Shocks: The Perils of Recursive Orderings,” *Journal of Applied Econometrics*, forthcoming.
- KILIAN, L. AND R. J. VIGFUSSON (2011): “Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?” *Quantitative Economics*, 2, 419–453.
- KIM, H. (2022): “Inequality, Disaster risk, and the Great Recession,” *Review of Economic Dynamics*, 45, 187–216.
- KIM, I.-M. AND P. LOUNGANI (1992): “The Role of Energy in Real Business Cycle Models,” *Journal of Monetary Economics*, 29, 173–189.
- KLUMP, R., P. MCADAM, AND A. WILLMAN (2012): “The Normalized CES Production Function: Theory and Empirics,” *Journal of Economic Surveys*, 26, 769–799.
- LEDUC, S. AND Z. LIU (2016): “Uncertainty Shocks are Aggregate Demand Shocks,” *Journal of Monetary Economics*, 82, 20–35.
- LEDUC, S. AND K. SILL (2004): “A Quantitative Analysis of Oil-Price Shocks, Systematic Monetary Policy, and Economic Downturns,” *Journal of Monetary Economics*, 51, 781–808.
- LEE, K. AND S. NI (2002): “On the Dynamic Effects of Oil Price Shocks: A Study Using Industry Level Data,” *Journal of Monetary Economics*, 49, 823–852.
- LEE, K., S. NI, AND R. A. RATTI (1995): “Oil Shocks and the Macroeconomy: The Role of Price Variability,” *Energy Journal*, 16, 39–56.
- LUDVIGSON, S. C., S. MA, AND S. NG (2021): “Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?” *American Economic Journal: Macroeconomics*, 13, 369–410.

- MORK, K. A. (1989): "Oil and Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results," *Journal of Political Economy*, 97, 740–744.
- MUMTAZ, H. AND F. ZANETTI (2013): "The Impact of the Volatility of Monetary Policy Shocks," *Journal of Money, Credit and Banking*, 45, 535–558.
- NEWELL, R. G. AND B. C. PREST (2019): "The Unconventional Oil Supply Boom: Aggregate Price Response from Microdata," *Energy Journal*, 40, 1–30.
- OLOVSSON, C. (2019): "Oil Prices in a General Equilibrium Model with Precautionary Demand for Oil," *Review of Economic Dynamics*, 32, 1–17.
- PETERMAN, W. B. (2016): "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity," *Economic Inquiry*, 54, 100–120.
- PINDYCK, R. S. (1991): "Irreversibility, Uncertainty, and Investment," *Journal of Economic Literature*, 29, 1110–1148.
- PLANTE, M., A. W. RICHTER, AND N. A. THROCKMORTON (2018): "The Zero Lower Bound and Endogenous Uncertainty," *Economic Journal*, 128, 1730–1757.
- PLANTE, M. AND N. TRAUM (2012): "Time-Varying Oil Price Volatility and Macroeconomic Aggregates," Federal Reserve Bank of Dallas Working Paper 1201.
- (2014): "Time-Varying Oil Price Volatility and Macroeconomic Aggregates," Manuscript, HEC Montréal.
- RABANAL, P., J. F. RUBIO-RAMÍ REZ, AND V. TUESTA (2011): "Cointegrated TFP Processes and International Business Cycles," *Journal of Monetary Economics*, 58, 156–171.
- RAMEY, V. A. AND D. J. VINE (2010): "Oil, Automobiles, and the US Economy: How Much Have Things Really Changed?" in *NBER Macroeconomics Annual 2010*, ed. by D. Acemoglu and M. Woodford, University of Chicago Press, vol. 25, 333–368.
- READY, R. C. (2018): "Oil Consumption, Economic Growth, and Oil Futures: The Impact of Long-Run Oil Supply Uncertainty on Asset Prices," *Journal of Monetary Economics*, 94, 1–26.
- RICHTER, A. W., N. A. THROCKMORTON, AND T. B. WALKER (2014): "Accuracy, Speed and Robustness of Policy Function Iteration," *Computational Economics*, 44, 445–476.
- ROTEMBERG, J. J. AND M. WOODFORD (1996): "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money, Credit and Banking*, 28, 550–577.
- SHEN, W. (2015): "News, Disaster Risk, and Time-Varying Uncertainty," *Journal of Economic Dynamics and Control*, 51, 459–479.
- SWANSON, E. (2018): "Risk Aversion, Risk Premia, and the Labor Margin with Generalized Recursive Preferences," *Review of Economic Dynamics*, 28, 290–321.
- TIMMER, M. P., E. DIETZENBACHER, B. LOS, R. STEHRER, AND G. J. VRIES (2015): "An Illustrated User Guide to the World Input Output Database: The Case of Global Automotive Production," *Review of International Economics*, 23, 575–605.
- WACHTER, J. A. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *Journal of Finance*, 68, 987–1035.
- ZHOU, X. (2020): "Refining the Workhorse Oil Market Model," *Journal of Applied Econometrics*, 35, 130–140.