Geopolitical Oil Price Risk and Economic Fluctuations*

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

March 27, 2026

ABSTRACT

Market participants and policymakers are concerned about major oil production shortfalls driven by geopolitical events. Even when such events never materialize, unanticipated increases in the probability of a production shortfall may generate a surge in the price of oil and oil price uncertainty. Our analysis provides the first systematic account of the quantitative importance of time-varying geopolitical risk to oil production for the global economy. We quantify the impact of actual and anticipated global oil production shortfalls on the price of oil and global growth, including the 2026 closure of the Strait of Hormuz, under alternative scenarios.

Keywords: Geopolitical risk, macroeconomic risk, time-varying uncertainty, rare disasters, oil, endogeneity, probability shocks, economic fluctuations, precautionary savings, inventories
JEL Classifications: E13, E22, E32, Q43

---

[*] We thank María Aristizábal-Ramírez, John Elder, Jancy Liu, Casey Petroff, and Michael Siemer for discussing our paper. We also thank Mario Crucini, Steven Davis, Paulina Restrepo-Echavarría, Michel Robe, 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 and the CADRE computing group at the Federal Reserve Bank of Kansas City. The views expressed 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-varying 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.

These concerns are illustrated by the Twelve-Day War between Iran and Israel in June 2025. As the probability of a severe disruption of maritime traffic through the Strait of Hormuz surged to 53% in mid-June, the West Texas Intermediate (WTI) price of oil peaked at $75 per barrel and oil price uncertainty, as measured by implied volatility, surged (see Figure 1). Even though the Strait of Hormuz never closed, the oil market reacted strongly to this shift in the probability of a major geopolitically driven shortfall of oil supplies. The market reaction reflected the fact that a closure of the Strait would have put close to 20% of global oil production at risk. Goldman Sachs warned that such an event would push the Brent price of oil to $110 per barrel and reduce global growth by 0.3 percentage points.[1] Other investment banks and the IMF reached similar conclusions.

Despite the obvious interest among policymakers and market participants in the effects of fluctuations in the probability of a major geopolitically driven oil supply disruption on the global economy, this question has not been addressed in the academic literature.[2] Existing work on the role of oil price uncertainty makes simplifying assumptions that prevent one from analyzing events such as the surge in geopolitical oil production risk in 2025. For example, studies employing stochastic volatility shocks to the price of oil not only assume that oil price uncertainty is exogenous but that the risk in the model is two-sided.

We develop a model of the global economy in which geopolitical risk in the oil market arises from the possibility of rare oil production disasters driven by geopolitical events. Our model includes risk averse agents, an oil production sector, oil storage, and limited substitutability between oil and capital. The price of oil is determined endogenously. Oil production disasters are of stochastic length and occur with a time-varying probability. The baseline size and duration of these disasters—a 5% drop in global oil production that lasts 3 quarters on average—is set to match the behavior of global oil production during major geopolitical events in the oil market over the past 50 years, such as the Arab-Israeli War in 1973, the Iranian Revolution in 1979, and the invasion of Kuwait in 1990. We also explore the possibility that agents might be concerned about larger

---

[1] See Goldman Sachs "Assessing the Economic Impacts of the War in the Middle East," June 23, 2025.

[2] The work most closely related to ours is Olovsson (2019). While Olovsson's model also allows for stochastic oil production disasters, his model is annual and only allows for a two-point *iid* distribution over the disaster probability. Thus, the model is not designed to understand the implications of continuously evolving disaster probabilities, as observed in the data.

Figure 1: Example of geopolitical risk

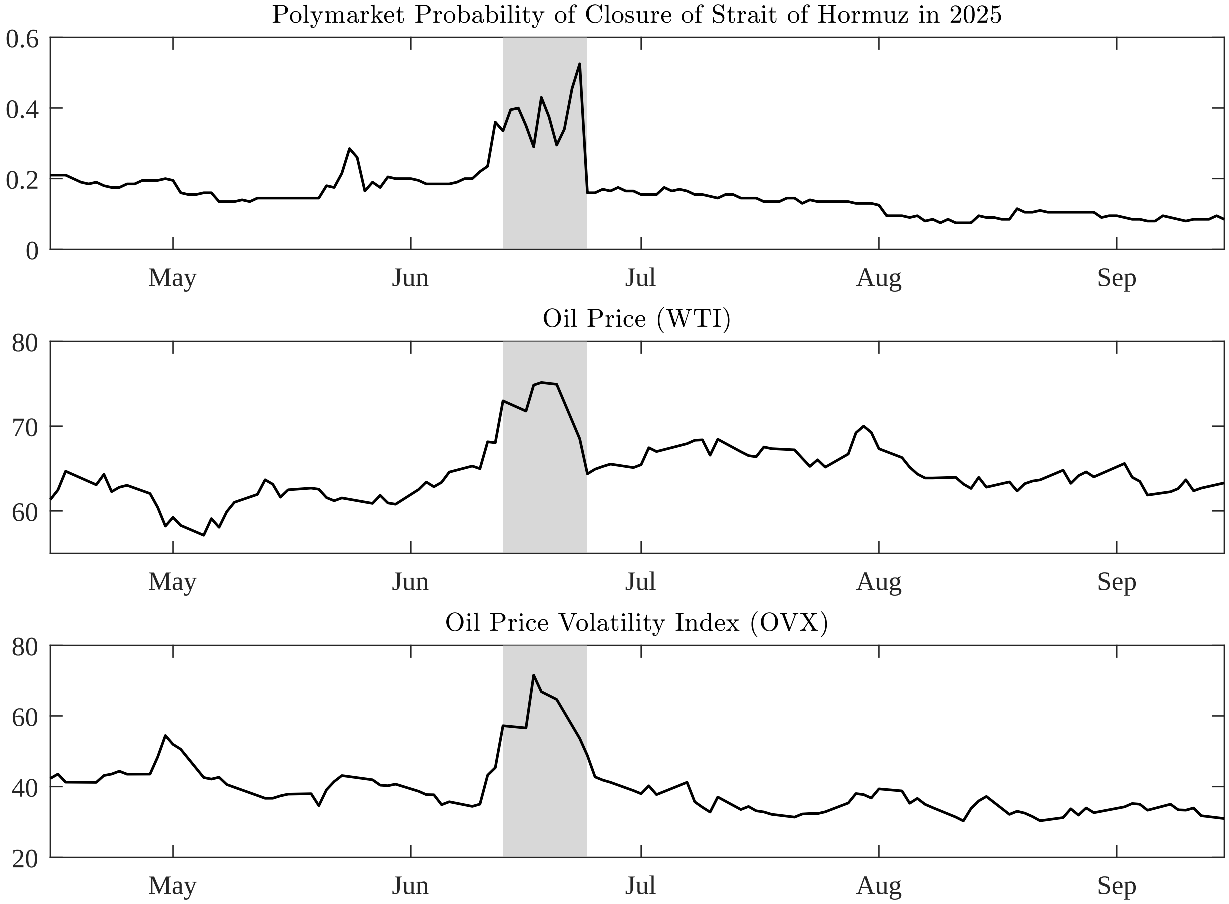

Notes: Daily data from April 15 to September 15, 2025. The shaded area marks the Twelve-Day War between Israel and Iran (June 13–June 24). The probability is from `https://polymarket.com/event/will-iran-close-the-strait-of-hormuz-in-2025`.

oil production disasters than observed historically. The model allows us to assess the effects of an increase in the probability of a major geopolitical oil supply disruption, addressing a key question often faced by policymakers and market participants.[3]

Our model also sheds light on the effects of actual geopolitical oil supply shortfalls. The magnitude and persistence of these effects critically depend on the structure of the underlying model. Our analysis differs from the existing literature employing linearized DSGE models to examine the economic impacts of oil price shocks, such as Backus and Crucini (2000) and Bodenstein et al. (2011), in that we explicitly model geopolitical oil supply disruptions in a nonlinear setting, tak-

---

[3] Since the model is global, we abstract from oil imports and exports and international capital flows. While 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. This reflects data limitations and the need to keep our highly nonlinear model computationally tractable.

ing into account agents' expectations surrounding these events. The importance of revisiting this question is underscored by the 59% surge in the WTI oil price after military conflict in the Persian Gulf created the largest oil supply disruption in history in March 2026.

We first show that shocks to the probability of an oil production disaster cause a simultaneous surge in oil price uncertainty and the price of oil, much like what occurred in June 2025 during the Iran-Israel War. These shocks cause oil inventories to be accumulated before the disaster is realized, as storage demand increases. In addition, they lower aggregate investment and aggregate output. How large these macroeconomic effects are depends on the magnitude of the probability shock and the magnitude of the potential oil production disaster. We show that under the baseline calibration, a 20 percentage point increase in the probability of a 5% shortfall in global oil production causes a 0.12% reduction in output on impact (or a 0.46 percentage point reduction in annualized output growth). When considering a 20% shortfall, which roughly corresponds to the cessation of oil supplies from the Persian Gulf, the drop in output nearly quadruples. As the probability of this event increases, the effect on output approaches 0.7%. The fact that the model jointly matches moments related to the oil market, real activity, asset prices, and uncertainty gives us confidence in the model's predictions. We further show that our model can generate an increase in the price of oil similar to that in Figure 1, given the observed surge in the probability of a closure of the Strait of Hormuz in June 2025, providing further validation of the model.

We then illustrate how our model may be used to shed light on the economic consequences of actual oil supply shortfalls. Our analysis highlights the similarities and differences between actual and anticipated oil production disasters. Whereas in the former case inventories are drawn down in response to lower oil production, in the latter case the price response is driven by oil inventory accumulation. Motivated by recent events, we examine the implications of the 2026 closure of the Strait of Hormuz under a range of scenarios about the duration of this closure. We find that under plausible assumptions this event drives up the oil price by 63% on impact and reduces output by 0.73% (implying a 2.9% decline in annualized growth), with further oil price increases and growth reductions depending on the duration of the closure. For example, if the Strait remains closed for three quarters, the oil price would rise by 118% relative to the baseline and the cumulative impact on annualized growth would rise from 0.2 to 1.3 percentage points at the end of 2026. Moreover, any increase in the expected duration of the closure (or in the probability of future oil supply disruptions) would further amplify the model predictions.

Related Literature Our work relates to several strands of the literature. First, it complements a 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; Jurado et al., 2015; Leduc and Liu, 2016; Ludvigson et al., 2021) by uncovering the underlying sources of time-varying uncertainty, building on Gourio (2012). This allows us to study the interaction between macroeconomic and oil price uncertainty and to make more precise statements about the role of geopolitical oil production risk in the global economy.

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 a surge 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 show that the responses of output uncertainty and oil price uncertainty to a shock to the probability of a macroeconomic disaster are an order of magnitude larger than the responses to a shock to the probability of an oil production disaster of the same magnitude. More than half of the observed oil price uncertainty tends to be driven by the macroeconomy, which helps explain why higher oil price uncertainty has historically been associated with lower real activity. This result highlights that this association should not be interpreted as a causal link.

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), and Olovsson (2019). These studies rely on simplifications such as stochastic volatility shocks or *iid* distributed disaster risk, and all of them abstract from macroeconomic risk.[4] In contrast, we examine this question within a general equilibrium model with endogenous oil price and macroeconomic uncertainty. Our analysis accounts for the fact that geopolitical oil price risk is time-varying and inherently one-sided.

We demonstrate that modeling downside risk is crucial. While increases in oil price uncertainty may alternatively be explained by stochastic volatility shocks to oil production growth, these shocks have little effect on the economy because they do not generate risk tilted to the downside. Likewise, the ability to store oil plays a central role. Without storage, the responses of the price of oil and real activity to higher oil production risk tend to be muted, suggesting that models without storage will fail to capture the full effects of shifts in oil production risk. Oil storage also matters for the responses to macroeconomic disaster probability shocks, underscoring the importance of jointly modeling oil production disasters and macroeconomic disasters.

Third, our analysis is relevant for a large empirical literature making the case that oil price uncertainty shocks 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

---

[4] Our model also differs from the general equilibrium model in Ready (2018) who 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, rather than addressing how the economy responds to shifts in the likelihood of a major temporary drop in oil production driven by geopolitical events.

affect oil production and storage (e.g., Cross et al., 2022; Kellogg, 2014). Our finding that oil price uncertainty responds to macroeconomic level and uncertainty shocks has important implications for empirical work seeking to establish the macroeconomic effects of shocks to oil price uncertainty. We show that level and uncertainty shocks are not the same, as implicitly assumed in VAR-GARCH models, and that 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. 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. 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.

Finally, our analysis complements the literature seeking to understand anticipation in oil markets based on linear models. One key difference is that in our model a 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 shocks to the expected oil price, as discussed in the previous literature. For example, Kilian and Murphy (2014) and more recent structural VAR models of the oil market derived from this study make no attempt to explicitly identify and model the determinants of oil price expectations. Another branch of the literature has focused on the domestic macroeconomic impact of news about large oil discoveries (see Arezki et al., 2017). This literature is not concerned with the impact of oil discoveries on the price of oil or, for that matter, the global economy, and it is not concerned with geopolitical risk. Likewise, the literature on the impact of OPEC announcements on oil price expectations is not concerned with the anticipation of rare major geopolitical oil supply disruptions (e.g., Känzig, 2021; Kilian, 2024; Plagborg-Møller and Wolf, 2022). In fact, recent evidence suggests that OPEC announcements capture changes in expected demand rather than expected supply and that the price effects of OPEC announcements are much smaller than originally reported. More generally, none of these studies consider probability shocks or nonlinearities.

Outline The remainder of the paper is organized as follows. In Section 2, 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 introduces a calibrated model of the global economy that elucidates the determination of oil price uncertainty and macroeconomic uncertainty. We use this model to quantify the impact of geopolitical oil price risk. In Section 4, we study the effects of shocks to the probability of oil production disasters and macroeconomic disasters on the economy and their ability to explain economic fluctuations. Our analysis sheds light on the key economic mechanisms in the

Figure 2: Empirical measures of oil price uncertainty, 1974Q4-2023Q4

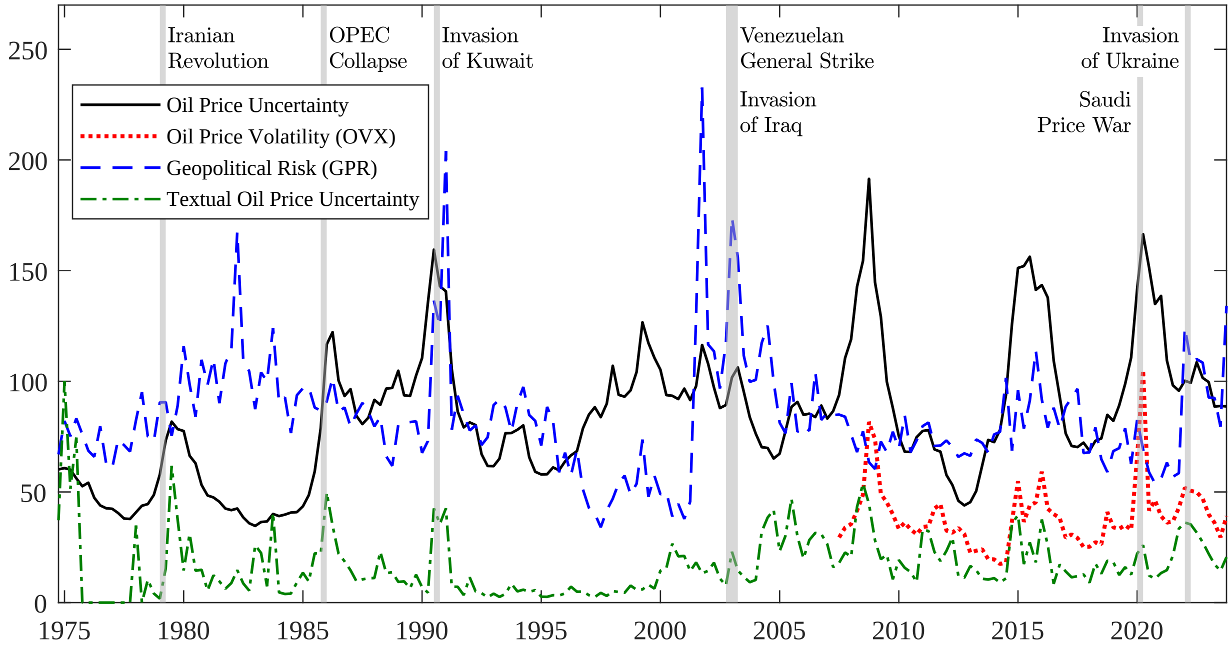

Notes: The solid line shows the uncertainty about the percent change in the real price of oil. The method used to quantify this uncertainty is based on Jurado et al. (2015). The dotted line is the option-implied crude oil price volatility index (OVX) published by the Chicago Board Options Exchange. The 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 geopolitical risk (GPR) series in Caldara and Iacoviello (2022).

model, the importance of modeling downside risk, and the sensitivity of our results to the specification of the model. Section 5 quantifies the impact of the 2026 closure of the Strait of Hormuz under a range of scenarios. Section 6 discusses implications of our analysis for empirical work on the effects of oil price uncertainty shocks. The concluding remarks are in Section 7.

2 MEASURING OIL PRICE UNCERTAINTY

It is useful to develop a deeper understanding of how time-variation in geopolitical risk in oil markets affects oil price uncertainty and global economic fluctuations. Our starting point is the downside risk to oil production caused by major geopolitical 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. 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.[5] 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 (see Appendix F). 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.[6] There are large spikes in 1979, 1986, and 1990 at the time of the Iranian Revolution, the collapse of OPEC, and the 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 evidently 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 (see Baumeister and Kilian, 2016), 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). 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.[7] Figure 2 also shows that our data-based index of oil

---

[5] The 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. Details of the construction of the uncertainty measure can be found in Appendix B.

[6] Our analysis 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 domestic price effectively 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 U.S. price is not a good proxy for the global price of oil after 1974. These problems are compounded in the pre-1974 era. The nominal price of oil prior to 1974 was stable for extended periods except for discrete jumps, reflecting the regulation of the U.S. oil market. As stressed in Hamilton (1985, p. 99), the striking stability of the nominal price of crude oil between the 1950s and the early 1970s "can be attributed to state regulatory commissions' policy of defending posted prices." 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).

[7] In 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.

price uncertainty and the OVX index differ systematically from the text-based oil price uncertainty index proposed in Abiad and Qureshi (2023) based on the methodology in Baker et al. (2016). The correlation of the text-based oil price uncertainty index with the OVX index, when that index is available, is only 0.55. This suggests that text-based measures fail to capture changes in oil price uncertainty as perceived by financial markets. Moreover, the correlation with our index over the full sample is only 0.30. This is not surprising because nothing in the construction of the text-based index ensures that it captures changes in the degree of uncertainty over time, as opposed to merely changes in press coverage. Finally, 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 the changes in these indices often differs, and the correlation between the indices is close to zero. Next, we propose a model of the global economy that captures the empirical relationship between oil price uncertainty and macroeconomic uncertainty, elucidates their determinants, and accounts for the role of uncertainty in the transmission of geopolitical risk shocks. A review of the key mechanisms in the literature linking oil price uncertainty to the economy can be found in Appendix A.

3 A MODEL OF GEOPOLITICAL RISK IN OIL MARKETS

In this section, we introduce a model of the global economy designed to examine the impact of both anticipated and actual geopolitical oil supply disruptions.

3.1 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 distinguishing feature of the model is that it includes downside risk to both oil production and the macroeconomy. 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. Time-variation in the probability of oil production and macroeconomic disasters induces exogenous variation in oil price and macroeconomic uncertainty.[8] 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 faced by agents is not two-sided. Rather it involves a sharp reduction in oil production. Rare oil production disasters matter not only because of their impact when they occur, but because agents' behavior reflects the anticipation of these disasters even when they are not realized.

---

[8] Applications 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).

Productivity and Macroeconomic Disasters Productivity growth $g_t = a_t/a_{t-1}$ 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 macroeconomic disaster occurs and 0 otherwise. The transition matrix for v_t^g is summarized by

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

where the probability of a macroeconomic disaster follows

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

which ensures that p_t^g is bounded between 0 and 1. The size of the disaster is ζ_g, and $\bar{\pi}_1^g$ is the unconditional probability of the disaster, which is evaluated by simulation. Details about how the ergodic probability is computed can be found in Appendix D.

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

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

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

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

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.[9] The firm's profit maximization problem is given by

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

---

[9] It should be noted that production network effects could potentially amplify the effects of geopolitical risk. We do not consider this possibility because production networks tend to be static. Modeling production networks abstracts from decisions related to investment and oil inventories, which are inherently intertemporal and which play a key role in how geopolitical risk affects the economy.

subject to

$$k_{t+1} = e^{-\zeta_g v_{g,t+1}}((1-\delta)k_t + i_t - \phi(i_t/k_t)k_t),$$
$$y_t = y_0(a_t n_t)^{1-\xi}\left((1-\alpha)(k_t/k_0)^{1-1/\sigma} + \alpha(o_t/o_0)^{1-1/\sigma}\right)^{\xi/(1-1/\sigma)},$$

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.[10]

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

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

where

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

Oil Production and Oil Production Disasters The production of oil is given by $o_t^s = a_t^o e_t$. 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. We also allow a_t^o to depend on the state of the economy since productivity in oil production is assumed to be cointegrated with productivity in the rest of the economy. This allows oil production to respond to changes in oil demand. 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 of oil production is given by

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

---

[10] A 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 G.

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.[11]

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

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

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 $\bar{\pi}_1^e$ is the unconditional probability of the disaster, which is evaluated by simulation in the same way as the ergodic probability for the macroeconomic 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

$$s_{t+1} = (1 - \omega)s_t + o_t^s - o_t - \tfrac{\tau}{2}\big(\tfrac{a_t}{s_t}\big)^2 a_t,$$

where ω is the cost of storage. Following Gao et al. (2022), the law of motion for s_t includes a penalty function that prevents stockouts, as they are not observed in the global oil market. The penalty function ensures $s_t > 0$, removing the need for a non-negativity constraint. The advantage of this approach is that it captures the fact that inventories rarely get close to zero in the data.

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

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

where

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

---

[11] When $\kappa_1 = 1$ 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 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^e_{t+1}, 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

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

subject to

$$u_t = c_t^{\chi} (a_t(1-n_t))^{1-\chi},$$
$$c_t + p^e_t s^e_{t+1} + b_{t+1}/r_t = w_t n_t + (p^e_t + d^e_t) s^e_t + b_t,$$

where β is the discount factor, p^e_t is the equity price, r_t is the risk-free rate, w_t is the wage rate, d^e_t 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

$$\chi w_t (1-n_t) = (1-\chi) c_t,$$
$$1 = E_t[x_{t+1} r_t],$$
$$1 = E_t[x_{t+1} r^e_{t+1}],$$

where

$$r^e_{t+1} \equiv (p^e_{t+1} + d^e_{t+1})/p^e_t,$$
$$x_{t+1} \equiv \beta (u_{t+1}/u_t)^{1-1/\psi} (c_t/c_{t+1}) (J_{t+1}/z_t)^{1/\psi-\gamma},$$
$$z_t \equiv (E_t[J_{t+1}^{1-\gamma}])^{1/(1-\gamma)}.$$

The equity risk premium is defined as $r^{ex}_t \equiv r^e_t - r_{t-1}$.

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

$$d^e_t = d^f_t + d^s_t - \vartheta (E_{t-1} k_t - \tfrac{1}{r_t} E_t k_{t+1}),$$

where $d^f_t = y_t - i_t - p^o_t o_t - w_t n_t$ and $d^s_t = p^o_t o_t$. Asset market clearing implies that $s^e_t = 1$ 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 $g_{o,t} = a^o_t/a^o_{t-1}$. 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.[12]

3.2 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 size of the model is at the limit of what is computationally feasible. In particular, versions of this model that allow for different blocks of countries would be computationally intractable.

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

3.3 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 some 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 unduly influenced by the COVID-19 pandemic. Appendix C provides our data sources.[13]

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.[14] The Frisch elasticity, which is pinned down by the steady-state labor supply, is set to 2 following

---

[12] The 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.

[13] Estimation 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, evaluating the model at a given parameterization takes many hours to solve.

[14] Swanson (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.

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 Substitution Elasticity (ψ)	2	Gao et al. (2022), Croce (2014)
Frisch Labor Supply Elasticity (η^λ)	2	Peterman (2016), Basu-Bundick (2017)
Capital-Oil Elasticity of Substitution (σ)	0.105	$SD(\Delta p^o)$
Capital Depreciation Rate (δ)	0.025	Depreciation on fixed assets, 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)
Oil Production Weight (α)	0.134	$E[o/y]$
Oil Inventory Stockout Cost (τ)	0.00001	$E[s/o]$
Average Growth Rate ($\bar{g}$)	1.0039	$E(\Delta y)$
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	$E(r^{ex})$
Prob. of Entering Growth Disaster ($\bar{p}_g$)	0.005	Occurs in expectation every 50 years
Prob. of Remaining in Growth Disaster ($\bar{q}_g$)	0.9	Gourio (2012)
Growth Disaster Prob. Persistence (ρ_{pg})	0.8	$AC(\mathcal{U}_y)$
Growth Disaster Prob. Shock SD (σ_{pg})	0.9	$SD(\mathcal{U}_y)$
Oil Production Disaster Size (ζ_e)	0.05	Avg. peak decline in oil prod. disasters
Prob. of Entering Oil Disaster ($\bar{p}_e$)	0.02	Avg. frequency of oil prod. disasters
Prob. of Remaining in Oil Disaster ($\bar{q}_e$)	0.67	Avg. duration of oil prod. disasters
Oil Disaster Prob. Persistence (ρ_{pe})	0.9	$AC(\mathcal{U}_{p^o})$
Oil Disaster Prob. Shock SD (σ_{pe})	1.4	$SD(\mathcal{U}_{p^o})$

Peterman (2016), Basu and Bundick (2017), and many others in the business cycle literature.

The weight on oil in the production function α is set to 0.134 to match the average oil expenditure share in the data. Similarly, the parameter τ controlling the size of the oil inventory stockout cost is set to match the average inventory share of oil consumption. 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). Using a lower elasticity would cause the model to significantly overstate the volatility of the oil price in the data. 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 1.0039 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 macroeconomic 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 remaining in a macroeconomic disaster, $\bar{q}_g$, is set to 0.9, in line with Gourio (2012). This value implies that macroeconomic disasters, on average, last 2.5 years. As shown in Appendix G, the responses to an average macroeconomic disaster are 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 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 fit of the model is not very sensitive to the value of ρ_{pe}. In the baseline calibration, we focus on $\rho_{pe} = 0.9$, rather than a smaller value, to not limit the role of oil disaster risk *a priori*. However, in Sections 4.6 and 4.7 we explore alternative values of ρ_{pe}, motivated by specific historical episodes.

The fixed probability of remaining in 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. This assumption may be defended given the disproportionate impact of the 1973/74 event on public perceptions about oil production disasters. In Section 4.4, we consider

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

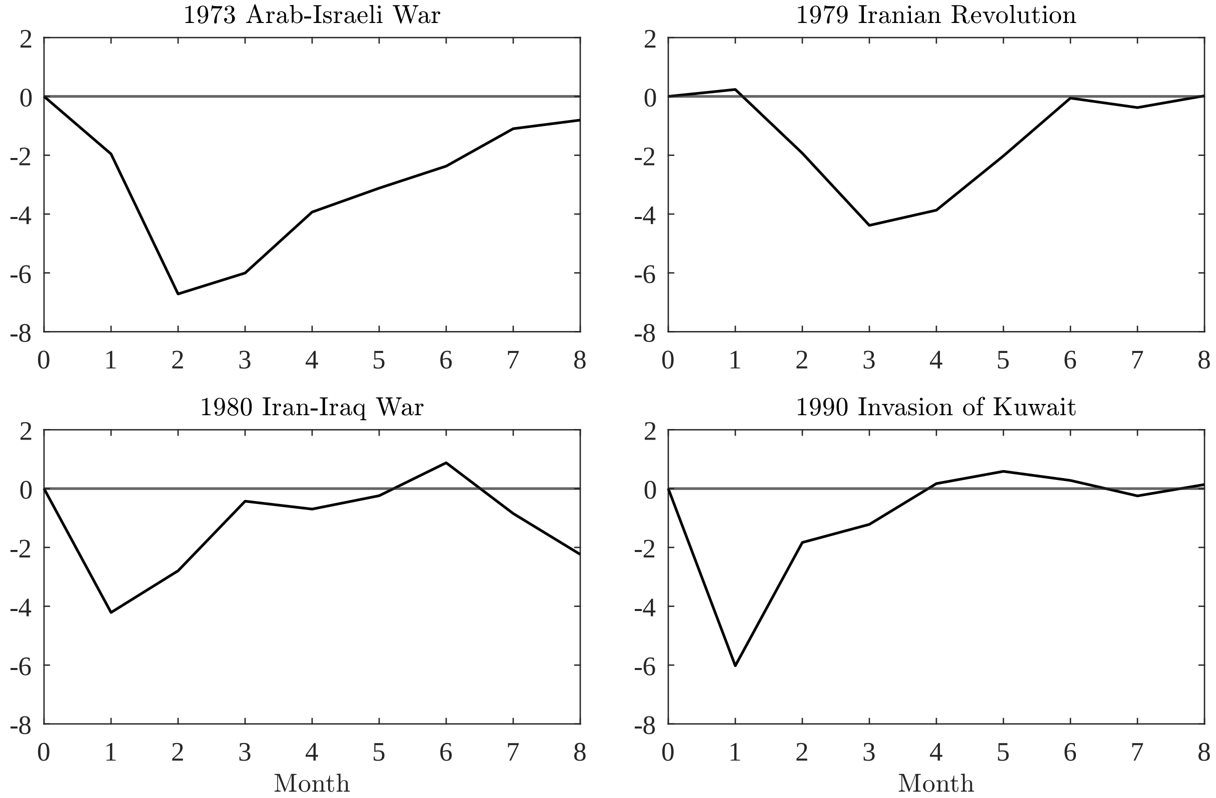

Notes: Reproduced from Hamilton (2013) using updated global oil production data. We exclude the 2002/03 episode because the revised data show no evidence of a material shortfall.

alternative oil disaster specifications.[15]

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 global 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 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

---

[15] It 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. For this reason, we conduct extensive sensitivity analysis on the parameters governing the oil production disaster and we validate the model implications against external evidence.

Table 2: Data and simulated moments

Moment	Data	Model	Moment	Data	Model
$E(\Delta y)$	0.39	0.39	$SD(\Delta o^s)$	2.06	2.14
$E(s/o)$	0.97	0.97	$SD(\Delta p^o)$	14.39	14.29
$E(p^o o/y)$	0.045	0.046	$SD(r^{ex})$	8.29	5.51
$E(r^{ex})$	2.18	2.09	$SD(\mathcal{U}_y)$	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	2.23	$AC(\Delta y)$	0.32	0.28
$SD(r)$	0.91	0.37	$AC(\Delta o^s)$	-0.11	-0.20
$Corr(\Delta y, \mathcal{U}_y)$	-0.15	-0.48	$AC(\Delta p^o)$	0.19	0.00
$Corr(\Delta y, \mathcal{U}_{p^o})$	-0.27	-0.44	$AC(p^o)$	0.95	0.94

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 Jurado et al. (2015). The standard deviations and rates are percents.

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.[16]

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 global 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. The model also reproduces the near random-walk behavior of the price of oil. 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.

3.4 DISCUSSION Our model incorporates precautionary savings by households in response to higher oil price uncertainty as well as oil storage. In addition, 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 that resemble those hypothesized by Bernanke (1983) and others.

Although Bernanke's theoretical analysis is often cited in support of models of oil price volatil-

---

[16] One 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 average share in our model.

ity shocks, it is not well appreciated that Bernanke's framework does not allow for monthly or quarterly energy price volatility shocks. Rather, he envisioned agents being uncertain about whether the price of energy 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 energy efficiency. The irreversibility of the investment decision causes risk averse agents to postpone the acquisition of either type of capital. 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 there is no substitutability between capital and oil in the short run. However, firms may invest in new capital with a 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 allow Atkeson and Kehoe to not only abstract from storage, but also to aggregate across different types of capital without tracking 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.

To illustrate this point, Appendix H evaluates the Atkeson-Kehoe (AK) putty-clay model with exogenous oil prices, augmented to include a stochastic volatility shock to the price of oil. We compare the responses of consumption and investment to an oil price uncertainty shock in this model to the corresponding responses in an alternative model that is identical except that the putty-clay structure has been replaced by the CES production function from our baseline model. We illustrate that the closer σ is to zero, the better the responses implied by the CES production function approximate the responses implied by the Leontief structure embodied in the AK model. For $\sigma = 0.1$, which corresponds to the value used in our baseline model with endogenous oil prices, the approximation is quite good. It is in this sense that our baseline model captures the essence of the results in Bernanke (1983) and Atkeson and Kehoe (1999), while remaining tractable in general equilibrium even when the oil price is endogenous.

4 THE EFFECTS OF TIME-VARYING DISASTER RISK

The model sheds light on the impact of geopolitical oil price risk on the global economy. 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.

4.1 TRANSMISSION OF DISASTER PROBABILITY SHOCKS

While economic agents witness disasters only infrequently, even unrealized disasters matter for economic behavior because there is a probability that they may be realized in the future. For example, the Twelve-Day War between Iran and Israel in June 2025 triggered concerns that the Strait of Hormuz might be closed, disrupting oil exports. As we highlight in this section, changes in the probability of such disasters will cause the oil market and the macroeconomy to move even when no disasters actually occur.

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 10% for a 40pp shock. This pattern is consistent with evidence in Kilian and Murphy (2014) that the oil price increases in 1979 and 1990 were driven in substantial part by higher oil inventory demand.

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.

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

Figure 4: Responses to alternative oil production disaster probability shocks

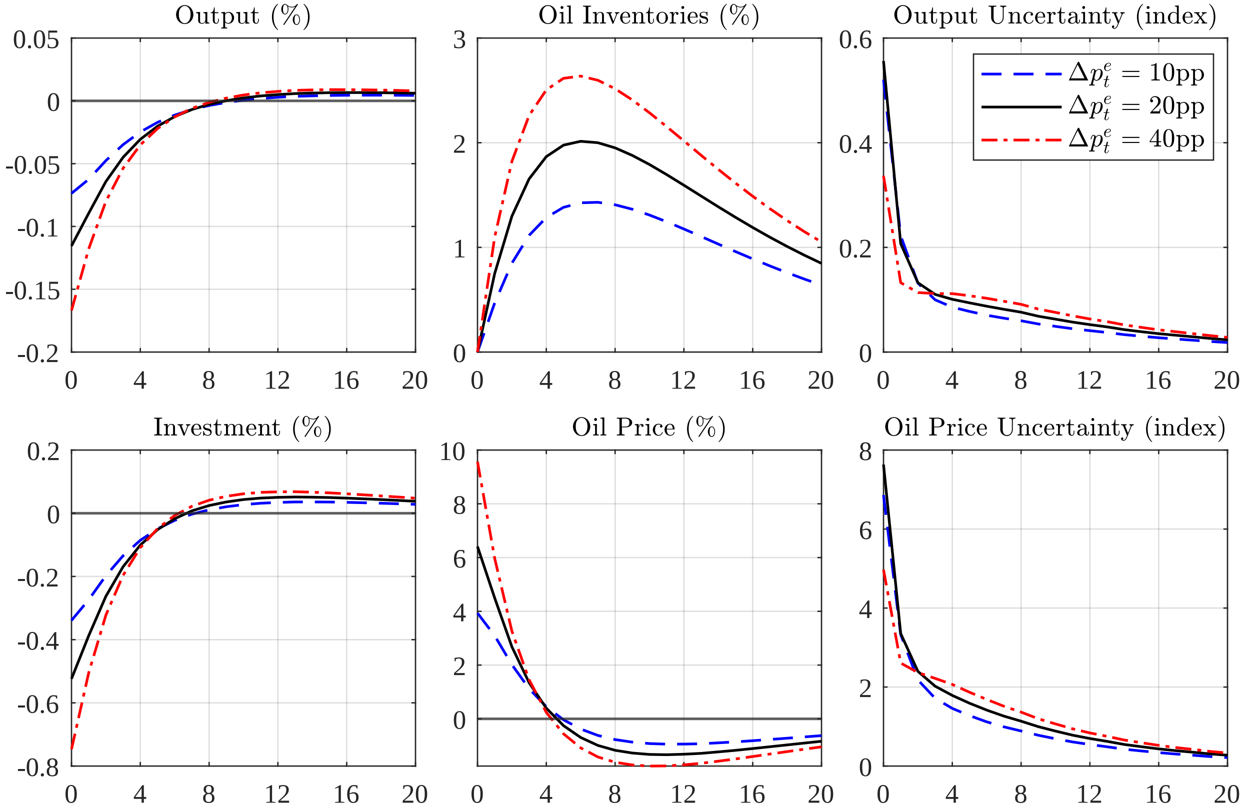

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized.

of an oil disaster raises the output uncertainty index only by 0.5, but the oil price uncertainty index by 6.9. 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.

Macroeconomic Disaster Probability A macroeconomic 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 a macroeconomic disaster. Figure 5 shows the responses when the growth disaster probability exogenously increases by 10, 20, and 40pp, respectively. As in the case of the oil disas-

Figure 5: Responses to alternative macroeconomic disaster probability shocks

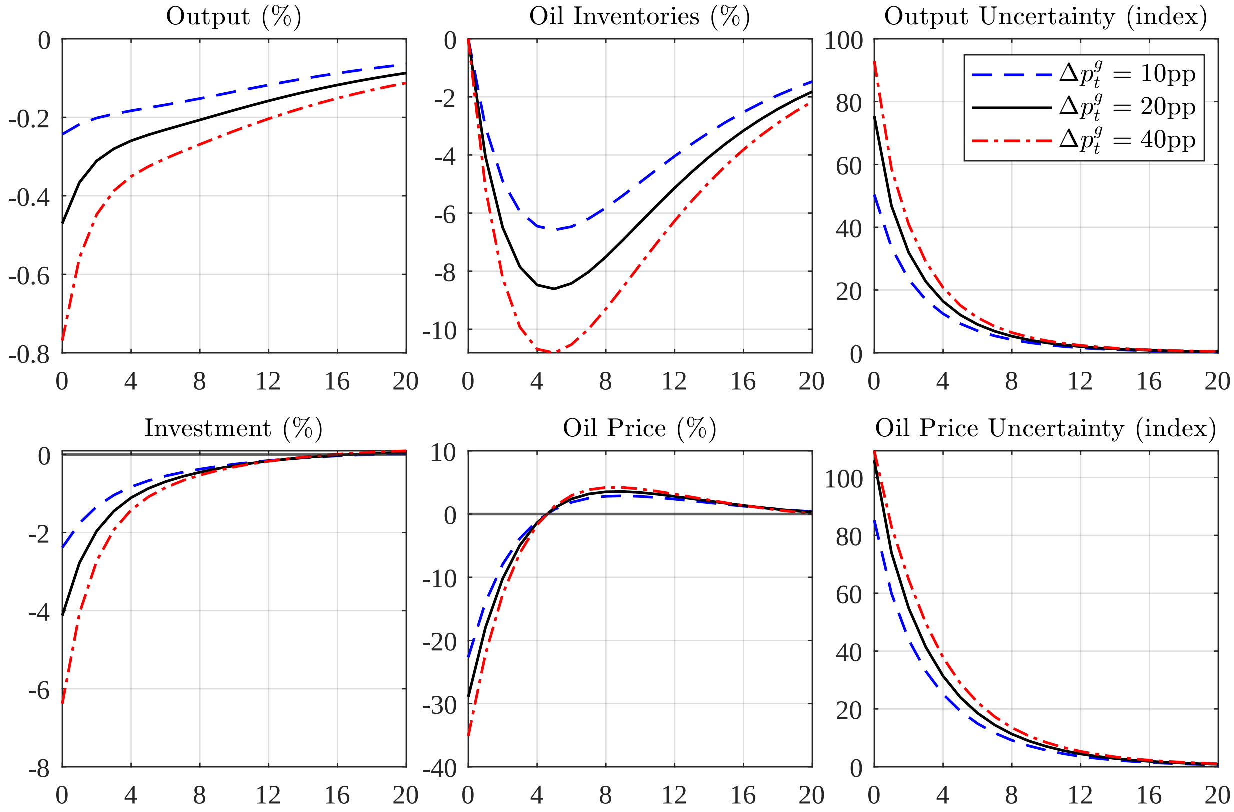

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized.

ter probability shock, the baseline is a probability close to zero. 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 23% 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 current and expected oil demand also reduce 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 disaster probability shock transmits directly to output rather than through the share of oil in output. 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 macroeconomic 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

Table 3: Decomposition of key volatilities

Moment	Data	Baseline Model	No Macro Disaster Risk	No Macro Disaster Risk or Oil Production Disaster Risk
$SD(\Delta y)$	0.74	0.87	0.65	0.65
$SD(\Delta i)$	1.95	1.92	1.27	1.24
$SD(\Delta o^s)$	2.06	2.14	2.12	1.12
$SD(\Delta p^o)$	14.39	14.29	6.30	5.63
$SD(\mathcal{U}_y)$	14.51	15.61	0.30	0.13
$SD(\mathcal{U}_{p^o})$	29.95	30.49	4.97	1.26

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).

price uncertainty were exogenous, this interaction between the uncertainty measures would not occur. This point is important. It formally shows that oil price uncertainty is endogenous with respect to the macroeconomy, consistent with the empirical evidence discussed in Section 2.

As in Figure 4, the responses do not scale proportionately with the increase in the macroeconomic disaster probability. For example, a 10pp increase leads to a 23% decline in the price of oil, whereas the price of oil declines by 35% when the probability rises by 40pp. This is true for the other variables as well.

A key difference between the two disaster probability shocks is that the macroeconomic 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.[17]

4.2 VARIANCE DECOMPOSITION

Table 3 shows that the model generally does an excellent job at capturing the volatility in the data. Dropping macroeconomic 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. Removing both macroeconomic 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

---

[17] Studying 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 nonlinear structural econometric model consistent with the structure of our model. No such model exists in the literature, and it would be inappropriate to compare our nonlinear responses to responses from linear VAR models.

but has little additional effect on the volatility of macroeconomic aggregates.

There are two key takeaways from these results. First, macroeconomic 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 geopolitical risk does not play a major role in driving business cycles. This does not mean, however, that such risk is not important during specific historical episodes, as illustrated by our impulse response analysis.

4.3 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, which introduces uncertainty about marginal costs, does not affect the substance of our results and explore the implications of allowing households to directly consume oil and of a higher average oil expenditure share.

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 macroeconomic 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.

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).[18] In this section, we compare the results from the baseline model to

---

[18] Stochastic 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).

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

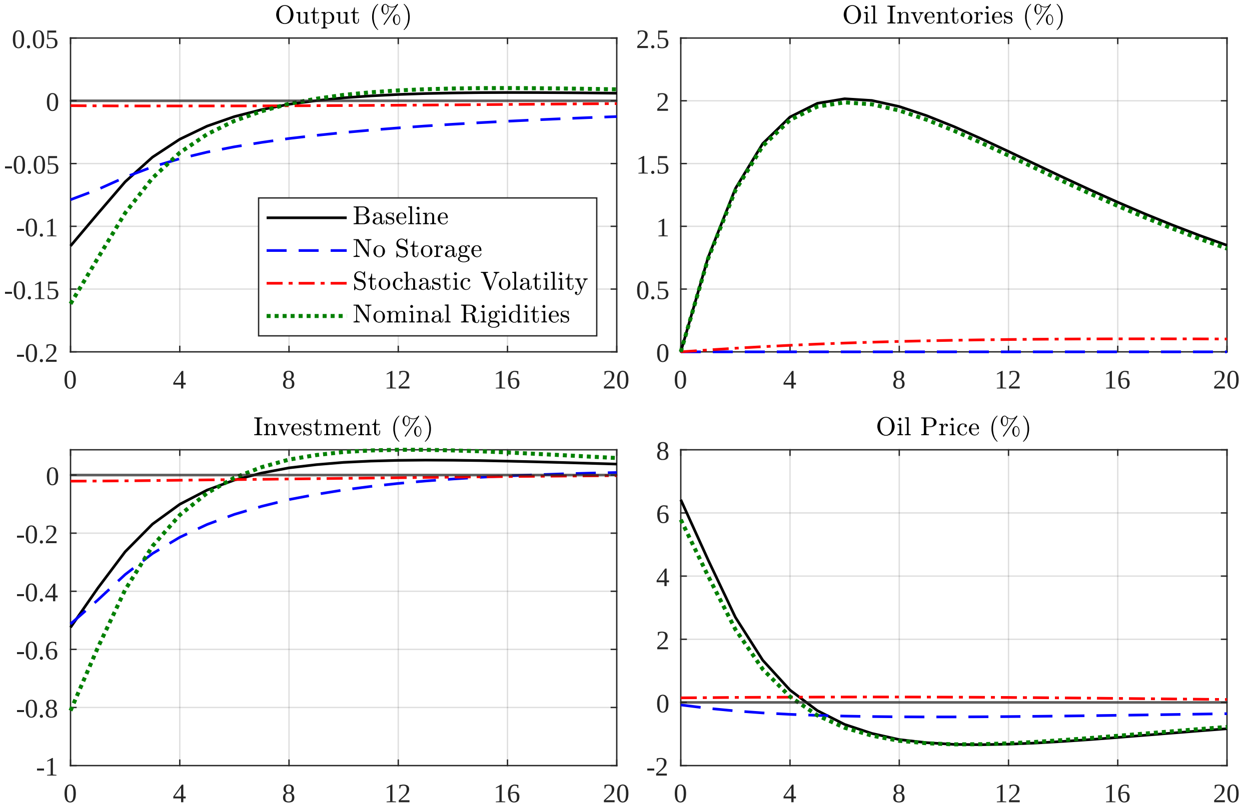

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized. The stochastic volatility shock has been normalized to match the impact response of oil price uncertainty in the baseline model. All disaster probability shocks are 20pp.

those from a model where uncertainty is generated by SV. Specifically, we introduce an exogenous volatility shock into productivity growth and oil production growth,

$$\ln g_t = \ln \bar{g} + \sigma_{g,t-1}\varepsilon_{g,t},$$

$$\ln g_{o,t} = \ln \kappa_0 + \kappa_1 \ln g_t + \kappa_2 \ln \epsilon_{t-1} + \sigma_{go,t-1}\varepsilon_{go,t},$$

$$\ln \sigma_{g,t} = (1-\rho_{sv}^{g}) \ln \bar{\sigma}_g + \rho_{sv}^{g} \ln \sigma_{g,t-1} + \sigma_{sv}^{g}\varepsilon_{sv,t}^{g},$$

$$\ln \sigma_{go,t} = (1-\rho_{sv}^{go}) \ln \bar{\sigma}_{go} + \rho_{sv}^{go} \ln \sigma_{go,t-1} + \sigma_{sv}^{go}\varepsilon_{sv,t}^{go},$$

where all shocks are standard normally distributed. 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

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

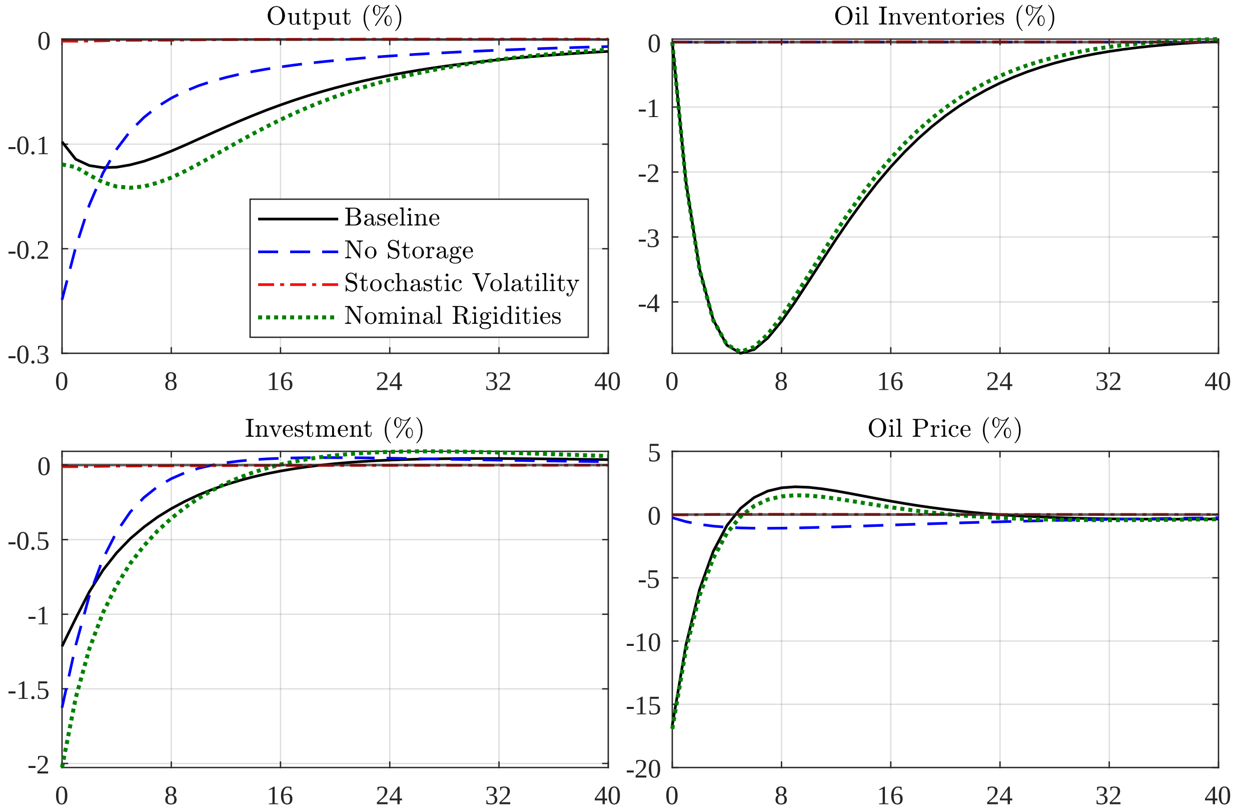

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized. The stochastic volatility shock has been normalized to match the impact response of output uncertainty in the baseline model. All disaster probability shocks are 5pp.

baseline oil disaster probability shock, $\varepsilon^{e}_{p,t}$. The $\varepsilon^{go}_{sv,t}$ shock is set so the SV specification generates the same impact effect on oil price uncertainty as the oil disaster probability shock. Qualitatively, 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.[19]

Nominal Rigidities The literature has emphasized that time-varying markups 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

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

where mc_t is marginal cost, π_t is the gross inflation rate, and φ scales the price adjustment cost. It 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 materially change our findings from the baseline model.

---

[19] Similar 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. When modeling time-varying markups using more conventional assumptions, as is standard in New Keynesian models, the effects are much smaller. For example, 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.

Oil Share and Household Oil Consumption It may seem that our analysis understates the importance of oil production risk because oil was more important for the economy in the 1970s than it is today. However, there is no evidence that the average oil share, as measured in this paper, was substantially higher in the 1970s than we assumed. For example, the actual U.S. oil share was near 6% in the mid-1970s compared with our average share of 4.5%.

This is not to say that there have not been fluctuations in this share about its average value. For example, the U.S. oil share in the early 1980s briefly reached an all-time maximum of 8%, while dropping below 3% in 1998 and 2020. Such fluctuations are largely driven by variation in the price of oil. When the price of oil is unusually high or low, so is the oil share. Our calibrated model replicates this pattern. Nevertheless, it is of interest to know how sensitive our key findings are to a higher average oil share. In Appendix G, we consider an unrealistically high average oil share of 8%. We find that in this case the effect of a 20pp shock to the probability of an oil production disaster on output is only 0.05 percentage points larger than in the baseline model and there is little change in the oil price response. It also remains true that oil production disaster risk is not an important determinant of macroeconomic volatility.

Another potential concern is that our baseline model does not allow oil to be consumed directly by households. In the model, households only purchase consumer goods produced with oil as an intermediate input. This concern, however, is alleviated by the fact that the primary effect of allowing for oil to be consumed directly involves the transfer of income from oil importing countries to oil exporting countries, as households consume imported crude oil in the form of fuels (e.g., Baumeister et al., 2018; Edelstein and Kilian, 2009). As consumer spending on other goods declines in the oil-importing economy, the economy contracts unless the petrodollars are recycled. Since our global model is a closed economy, there is no income transfer by construction, obviating this concern. Additional analysis in Appendix G that allows oil to directly enter the utility function confirms that this change does not materially affect our findings.

4.4 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, given a 20pp probability shock.

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. We then, alternatively, consider a much larger oil production disaster of 20% of global oil production with the same expected duration

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

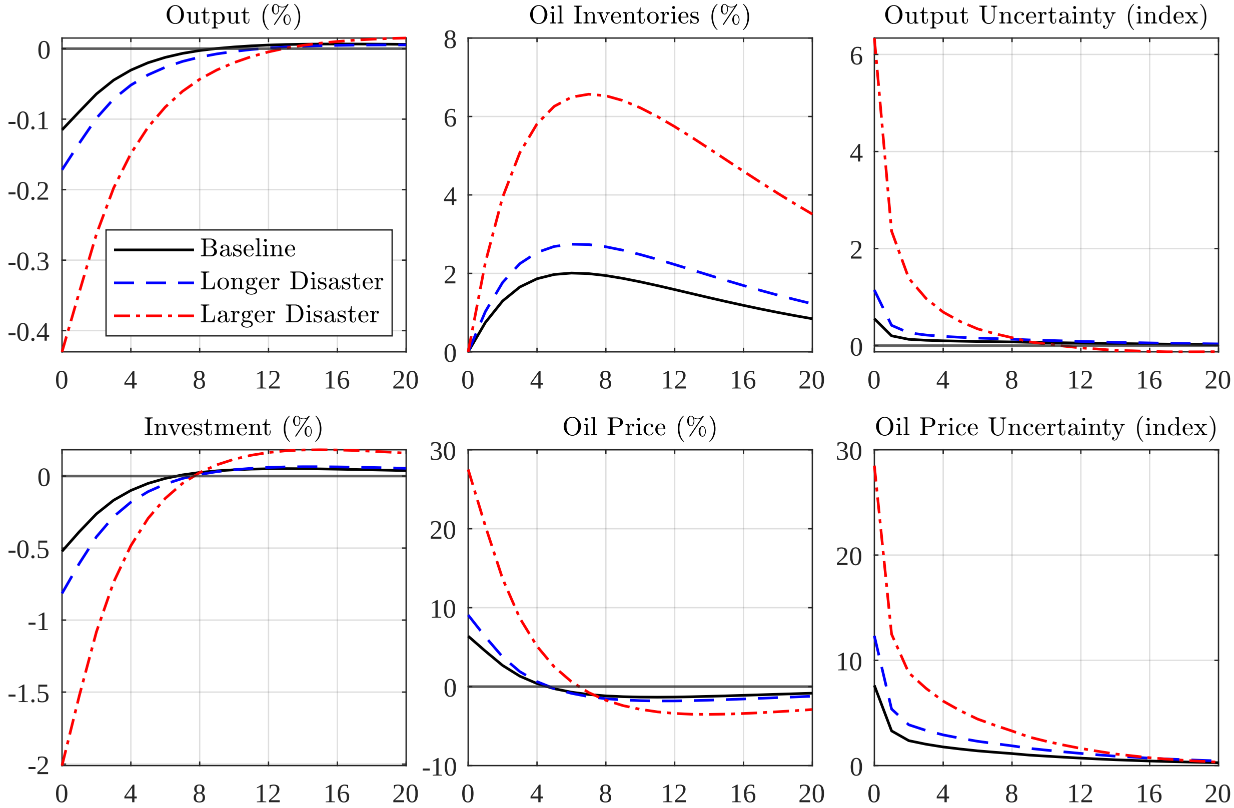

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized. The disaster probability shock is 20pp for all three specifications.

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.

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 nearly four-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 magnitude. The increase in the price of oil rises from 6% on impact to 28%, and the reduction in output rises from 0.12% to 0.43%, implying a 1.71 percentage point reduction in annualized growth.

The specification of the oil disaster also affects the ability of oil production disaster risk to

Figure 9: Responses to an oil production disaster realization

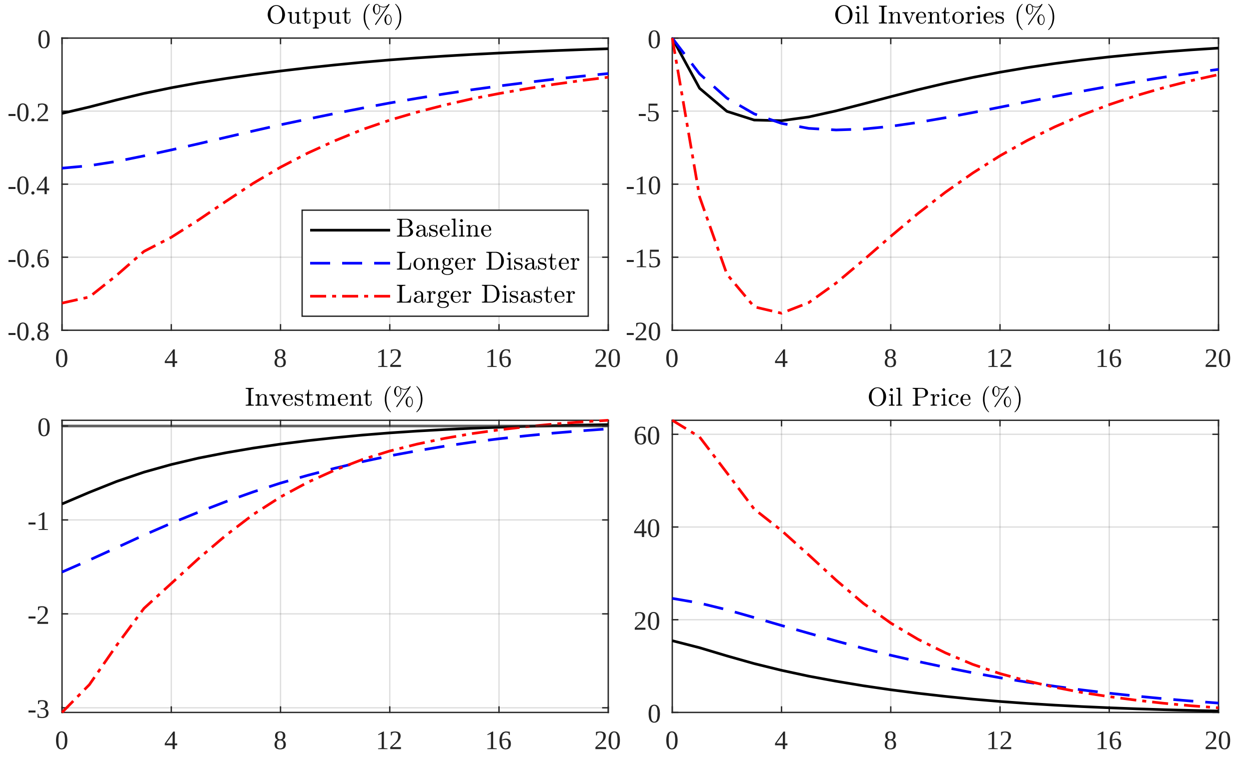

explain variation in oil price uncertainty. While oil production disaster risk accounts for only 16% of this variability in the baseline model, as shown in Table 3, this share rises to 22% when the disaster is longer lasting and to 69% when the magnitude of the disaster is increased. We conclude that 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 substantial portion of the variability in oil price uncertainty is driven by macroeconomic uncertainty.

4.5 OIL DISASTER REALIZATION

So far, we reported the responses to oil production and macroeconomic disaster probability shocks, conditional on no disasters being realized. It is of independent interest to assess the impact of the realization of such disasters. As the probability of an oil production disaster increases, the economic effects tend to converge to those that occur when a disaster is realized, except for oil inventories. The latter response changes sign when the production shortfall is realized, as storage firms respond to the temporarily higher price of oil.

Figure 9 shows responses under the baseline calibration (5% drop in global oil production that is expected to last for 3 quarters) and when allowing for a longer disaster (10 quarters on average) or a larger disaster (20% drop in global oil production). In the baseline model, we find that output declines by 0.21% on impact, but the effects nearly double when allowing the disaster to

Figure 10: Responses to an oil production disaster probability shock of different persistence

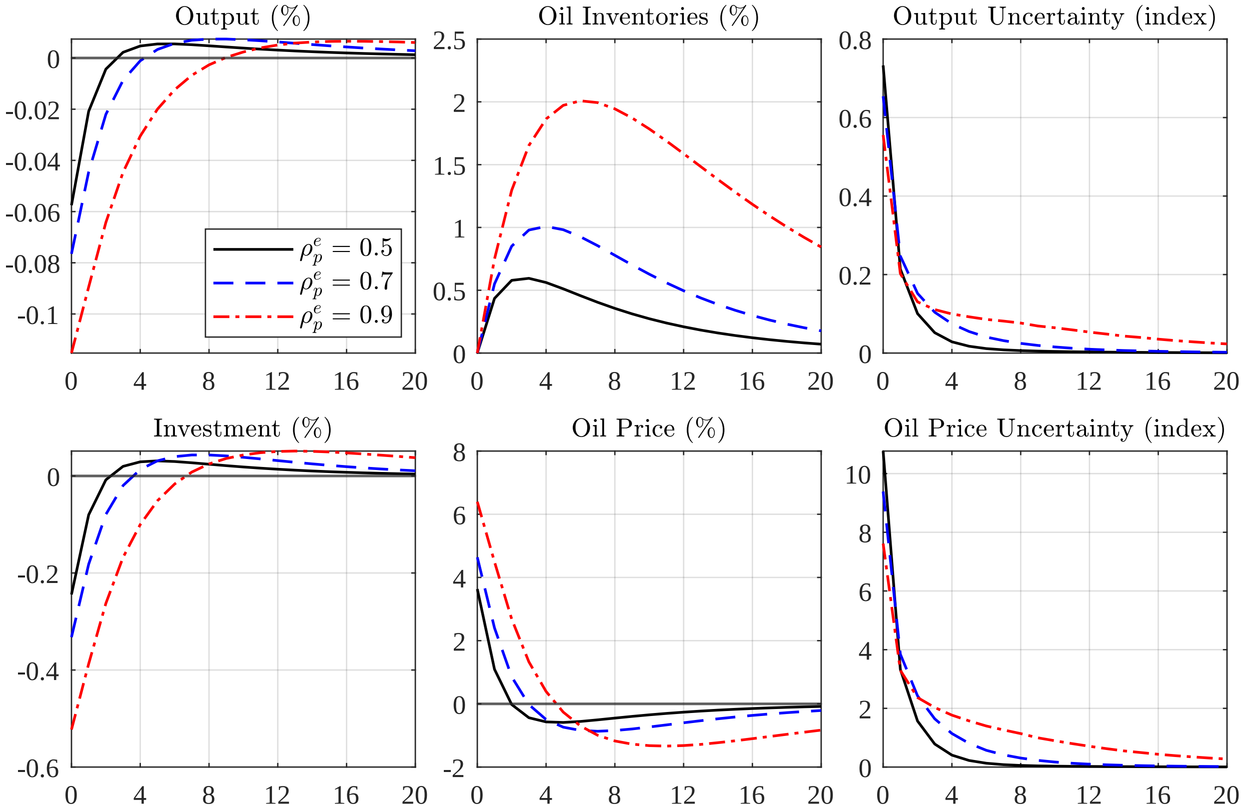

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized. The disaster probability shock is 20pp for all three specifications.

last longer and reach -0.73% on impact when increasing the magnitude of the shortfall to 20% of global oil production, corresponding to a 2.9 percentage point reduction in annualized growth. In the latter case, the oil price jumps by 63% on impact compared to 15% in the baseline model. This evidence is consistent with the view that oil production disasters may have substantial effects on the economy and on the price of oil. Our model helps quantify what these effects are based on a formal model and highlights which model parameters are driving the response.[20]

4.6 PERSISTENCE OF OIL DISASTER PROBABILITY

As noted earlier, the parameter ρ_{pe} governing the persistence of the oil disaster probability is difficult to pin down in the calibration. In the baseline calibration, we work with $\rho_{pe} = 0.9$. This parameter choice is plausible for events such as the invasion of Kuwait in August 1990, which persistently raised the probability of an oil disaster. That probability only started dropping in November of 1990, as discussed in Kilian (2008). This parameter choice may not be representative of all such events, however.

---

[20] The corresponding results for the realization of the macroeconomic disaster are provided in Appendix G.

Table 4: Impact effect of alternative oil production disaster shocks (%)

	5% Oil Disaster				20% Oil Disaster			
	Output		Oil Price		Output		Oil Price	
Shock	0.5	0.9	0.5	0.9	0.5	0.9	0.5	0.9
20pp	−0.06	−0.12	3.6	6.4	−0.22	−0.43	16.3	27.5
40pp	−0.10	−0.17	6.5	9.6	−0.39	−0.62	28.8	40.8
60pp	−0.14	−0.20	9.0	11.8	−0.53	−0.75	39.7	50.3
Realized	−0.21	−0.21	16.0	15.5	−0.81	−0.73	72.2	63.0

Notes: All results are in percent deviations from the baseline. Impact effects of a 20, 40, and 60 percentage point increase in the probability of an oil production disaster and its realization. The results shown are for probability shock persistence 0.5 and 0.9.

One example is the June 2025 war between Iran and Israel. This event temporarily raised the probability of a closure of the Strait of Hormuz, putting about 20% of global oil production at risk, and generated an increase in the price of oil consistent with a much lower persistence parameter (see Figure 1). From June 11, when anticipation of the war started driving up the price of oil, to June 24 when the war ended, there was a 13 percentage point surge in the probability of a closure of the Strait of Hormuz on average and a 9% increase in the price of oil. Our model can match that increase in the price of oil when the persistence parameter ρ_{pe} is set to 0.5.[21] This episode provides motivation for considering alternative values of ρ_{pe}. Figure 10 shows that lowering ρ_{pe} preserves the qualitative pattern of the baseline responses, while reducing the magnitudes.

4.7 IMPLICATIONS FOR POLICYMAKERS

One practical use of our model is to help policymakers gauge the likely effect of geopolitical oil disaster risk on the global economy. Table 4 shows how increased probabilities of a 5% or, alternatively, a 20% global oil supply disruption driven by geopolitical events affect global output and the price of oil. While the model could be used to consider oil disasters of any size, events putting 5% or 20% of global oil production at risk are particularly useful benchmarks. We focus on the impact effects of such risks under alternative assumptions about the persistence of the disaster probability, which may vary in practice as discussed earlier. We also include the impact effect of a realization of these oil disasters.

The impact effects of geopolitical risk are more modest when only 5% of global oil production is at risk. Even for reasonably large probability shocks, the price of oil only rises by between 4% and 16% and global output only falls by between 0.1% and 0.2%. Substantially larger effects arise when considering a potential 20% shortfall in global oil production, corresponding to a closure of the Strait of Hormuz, for example. In that case, the same probability shocks raise the price of oil

---

[21] A similarly good fit is obtained when converting the data into quarterly averages.

by between 16% and 50% and lower global output by between 0.2% and 0.8%.

5 IMPACT OF THE 2026 CLOSURE OF THE STRAIT OF HORMUZ

In this section, we further illustrate how our model may be used in practice. We further investigate the implications of a complete closure of the Strait of Hormuz, not allowing for any diversion of oil exports by pipeline to the Red Sea, which removes close to 20% of global oil supplies from the market. This exercise is motivated by the recent conflict in the Middle East. There are a number of ways in which these export restrictions may be weakened over time, but this scenario serves as a useful benchmark in thinking about the impact of the ongoing conflict on the global economy.

Unlike in Section 4, we condition on scenarios in which the Strait of Hormuz remains closed for one, two, or three quarters. This contrasts with our assumption in Section 4.5 of a disaster happening for one quarter and ending with probability $1-\bar{q}_e$. Suppose the Strait of Hormuz shuts down in 2026Q2 for the duration of the quarter.[22] We do not make any projections about how long the closure will actually last. Instead, we illustrate the potential impact under alternative assumptions. For example, suppose the Strait reopens in 2026Q3, after only one quarter. In that case, the probability of a closure is 100% on impact, before dropping to zero in 2026Q3 and beyond. Alternatively, the Strait may remain closed with 100% probability for two or three quarters before reopening.

The upper panel of Figure 11 illustrates the percent increase in the quarterly average global oil price, relative to the baseline of no geopolitical oil supply disruption, for each of these three scenarios. The lower panels show the corresponding effects on output growth and on the level of output, indicating that these scenarios imply potentially large and persistent effects on the global economy. Table 5 maps these model predictions into the average quarterly WTI price per barrel and the quarterly output growth rate at annualized rates.

Regardless of when the Strait reopens, the model implies that a closure of the Strait of Hormuz that removes close to 20% of global oil supplies from the market during 2026Q2 is expected to raise the average WTI price of oil from \$60 per barrel on average in the three months ending in January 2026 (before the war was anticipated) to \$98 per barrel. It also lowers output growth by an annualized 2.9 percentage points in 2026Q2.

The subsequent effects depend on when oil exports resume. For example, if the Strait reopens after one quarter, the oil price drops to \$68 per barrel and growth increases 2.2 percentage points in 2026Q3. While the oil price drop causes growth to recover, the level of output remains 0.2% below its pre-closure level by year-end 2026 and 0.1 percent below its initial level by year-end 2027.

When the oil supply shortfall lasts longer than one quarter, richer dynamics arise. Extending

---

[22] In reality, the closure occurred in March 2026, so the quarter in the model does not line up perfectly with the calendar quarter. This timing difference is immaterial for understanding the economic effects.

Figure 11: Implications of Closure of Strait of Hormuz

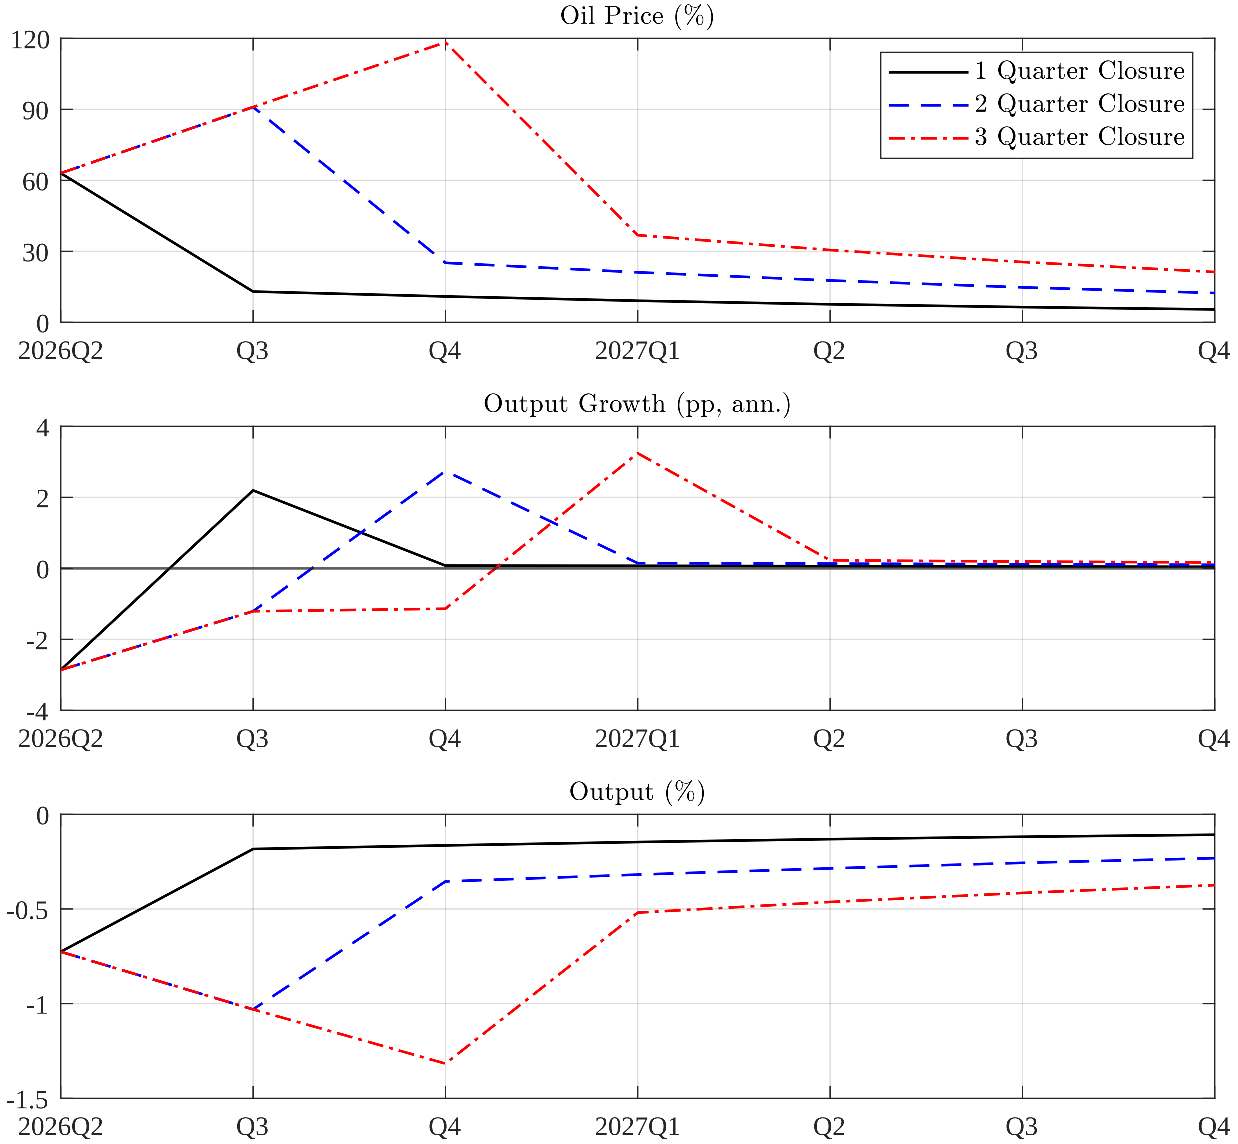

Notes: Results based on a 20% oil production shortfall starting in 2026Q2. All other parameters are set to their baseline values.

the closure to two quarters causes the oil price to rise further to $115 per barrel in 2026Q3 before falling to $76 per barrel in 2026Q4. The impact on output growth only turns positive in 2026Q4. If shipping resumes after three quarters, the oil price will rise even further before declining, reaching $132 per barrel by year-end. The impact on growth remains negative through year-end 2026. Thus, Q4/Q4 output growth in 2026 could fall 0.2 percentage points if the oil supply disruption remains limited to one quarter or 0.3 percentage points if it lasts two quarters. The reduction would sharply increase to 1.3 percentage points if the disruption persists for three quarters. These results

Table 5: Effect on WTI price and output growth by duration of closure

	Oil Price ($/barrel)			**Output Growth (pp, ann.)**		
	1-Quarter	2-Quarters	3-Quarters	1-Quarter	2-Quarters	3-Quarters
2026Q1	60	60	60	0.0	0.0	0.0
Q2	98	98	98	−2.9	−2.9	−2.9
Q3	68	115	115	2.2	−1.2	−1.2
Q4	67	76	132	0.1	2.7	−1.1
2027Q1	66	73	83	0.1	0.1	3.2
Q2	65	71	79	0.1	0.1	0.2
Q3	64	69	76	0.1	0.1	0.2
Q4	64	68	73	0.0	0.1	0.2

Notes: The baseline oil price is the average WTI price from November to January.

are useful for thinking about possible paths the economy might take in response to the ongoing struggle to control the Strait of Hormuz.

These scenarios are based on the assumption that the shortfall has an expected duration of three quarters, corresponding to the duration of the 1973 oil supply disruption. Shifts in the market expectation of how long the Strait will remain closed, reflected in a higher expected duration or a more persistent increase in the probability of future closures in the model, would imply potentially much higher oil prices with implications for economic growth. There are also factors outside of the model that could cause the oil price increase to be more substantial than our model suggests such as dislocations in the oil tanker market and rising insurance rates for oil tankers, as reflected in the rising spread between WTI and Brent prices.

6 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 or the U.S. economy for the uncertainty about the price of oil. Perhaps less obviously, the model also shows that oil price uncertainty re-

Figure 12: Responses to positive and negative oil production disaster probability shocks

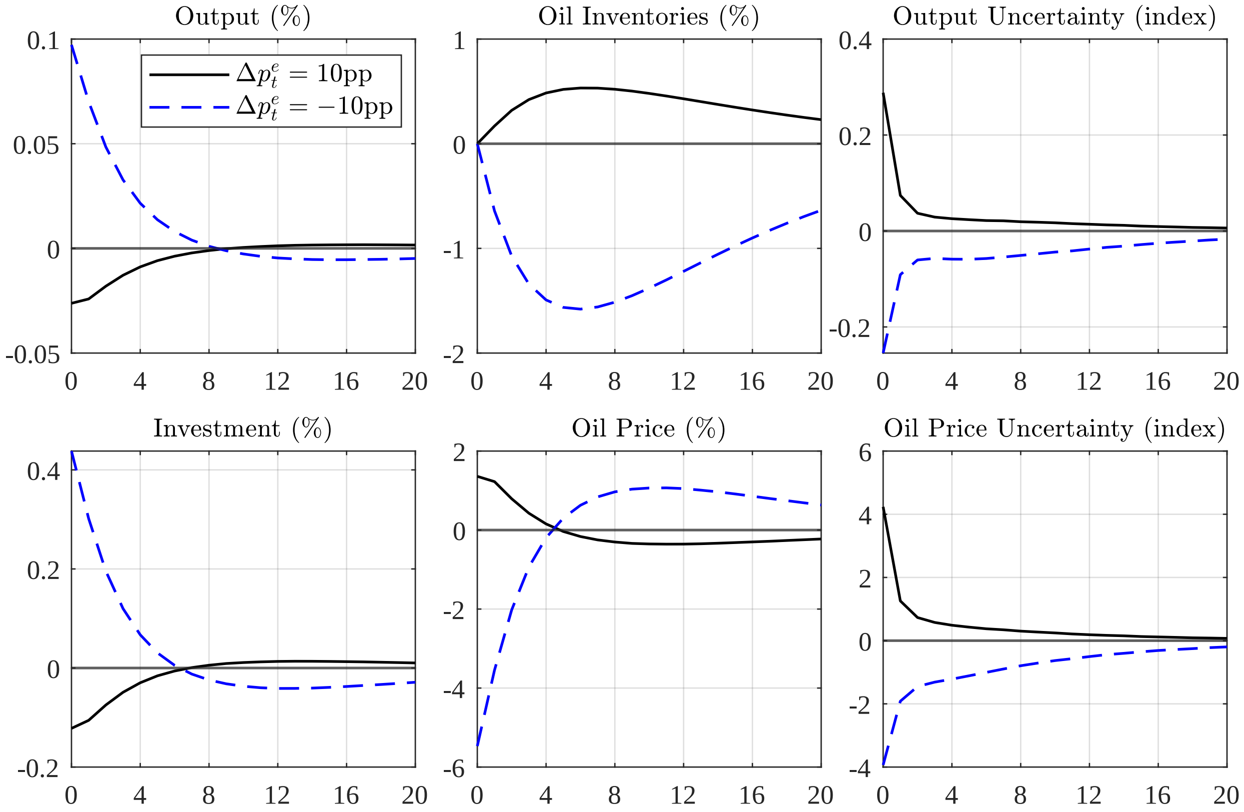

Notes: Responses in deviations from the baseline. Simulations assume no disasters are realized. The simulations are initialized at a 15% oil disaster probability to permit a positive and negative shock. The state equals the average of periods in the ergodic distribution that are within 1pp of the initial disaster probability.

sponds to level shocks in the macroeconomy. 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. Specifically, GARCH models impose the restriction that positive and negative oil price shocks both increase oil price uncertainty, which is inconsistent with our model. Figure 12 illustrates this point by comparing the responses to a ±10pp disaster probability shock under the baseline calibration. 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 macroeconomic disaster probability shock, as shown in Appendix G.

Alternative models that break the link between level and uncertainty shocks such as the VAR model with SV in Jo (2014) instead postulate that oil price uncertainty shocks and shocks to the level of the price of oil are independent. However, our analysis shows 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 macroeconomic disaster not only raises oil price uncertainty, but also causes storage demand to increase, raising the real price of oil.[23]

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.

7 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 the price of oil, oil price uncertainty, and the macroeconomy. Geopolitical risk is modeled as downside risk to oil production, and both the oil price and uncertainty are determined endogenously.

We first illustrated how the effects of anticipated geopolitical oil supply disruptions depend on the magnitude of the shortfall, its expected duration, and the persistence of the probability of a disruption. For example, under our baseline calibration, we find that a 20 percentage point increase in the probability of a 5% shortfall in global oil production causes a 0.12% reduction in global output, implying a 0.46 percentage point reduction in annualized growth. The drop in output nearly quadruples when considering a 20% shortfall, which roughly corresponds to the cessation of oil supplies from the Persian Gulf. As the probability of a 20% shortfall increases, the initial decline in global output approaches 0.7%, indicating that anticipated geopolitical oil supply shortfalls can have important effects on the global economy.

Motivated by recent events in the Middle East, we then further investigated the implications of a complete closure of the Strait of Hormuz, not allowing for any diversion of oil exports by pipeline to the Red Sea, which removes close to 20% of global oil supplies from the market. We consider alternative scenarios in which the Strait of Hormuz remains closed for one, two, or three quarters. The model implies that in all three scenarios the average WTI price of oil is expected to rise from

---

[23] Moreover, as shown in Section 4, 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.

\$60 per barrel on average before the crisis was anticipated to \$98 per barrel. Output growth is expected to drop on impact by 2.9 percentage points at an annualized rate. The subsequent effects depend on when oil exports resume. For example, if shipping only resumes after three quarters, the oil price will rise even further before declining, reaching \$132 per barrel, with the cumulative reduction in growth reaching 1.3 percentage points after three quarters.

These results provide a benchmark for assessing the quantitative impact of geopolitical risk in oil markets on the global economy. Our analysis highlights that this impact crucially depends on the expected duration of the shortfall and on whether the probability of future oil supply disruptions has persistently risen in response to the military conflict in the Persian Gulf. In particular, if one expects the Strait to remain closed much longer, the model implies much larger oil price increases and a sharper contraction.

REFERENCES

ABIAD, A. AND I. A. QURESHI (2023): "The Macroeconomic Effects of Oil Price Uncertainty," *Energy Economics*, 125.

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.

BASU, S. AND B. BUNDICK (2017): "Uncertainty Shocks in a Model of Effective Demand," *Econometrica*, 85, 937–958.

BAUMEISTER, C. AND L. KILIAN (2016): "Understanding the Decline in the Price of Oil since June 2014," *Journal of the Association of Environmental and Resource Economists*, 3, 131–158.

BAUMEISTER, C., L. KILIAN, AND X. ZHOU (2018): "Is the Discretionary Income Effect of Oil Price Shocks a Hoax?" *Energy Journal*, 39, 117–137.

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.

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.

BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2018): "Really Uncertain Business Cycles," *Econometrica*, 86, 1031–1065.

BODENSTEIN, M., C. J. ERCEG, AND L. GUERRIERI (2011): "Oil Shocks and External Adjustment," *Journal of International Economics*, 83, 168–184.

BORN, B. AND J. PFEIFER (2014): "Policy Risk and the Business Cycle," *Journal of Monetary Economics*, 68, 68–85.

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.

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-VILLAVERDE, 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-VILLAVERDE, 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.

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. KLIESEN (2005): "Oil Price Volatility and U.S. Macroeconomic Activity," *Federal Reserve Bank of St. Louis Review*, 87, 669–684.

HAMILTON, J. D. (1985): "Historical Causes of Postwar Oil Shocks and Recessions," *Energy Journal*, 6, 97–116.

——— (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.

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. (2008): "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *Review of Economics and Statistics*, 90, 216–240.

——— (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*, 40, 395–410.

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.

LUDVIGSON, S. C., S. MA, AND S. NG (2021): “Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?” *American Economic Journal: Marcoeconomics*, 13, 369–410.

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.

PLAGBORG-MØLLER, M. AND C. K. WOLF (2022): “Instrumental Variable Identification of Dynamic Variance Decompositions,” *Journal of Political Economy*, 130, 2164–2202.

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.

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.

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.