Journal of Econometrics 247 (2025) 105914

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Natural disasters as macroeconomic tail risks☆

Sulkhan Chavleishvili [a], Emanuel Moench [b,c,*]

[a] Department of Economics and Business Economics, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark
[b] Finance Department, Frankfurt School of Finance & Management, Adickesallee 32-34, 60322 Frankfurt am Main, Germany
[c] CEPR, London, UK

ARTICLE INFO

JEL classification:
C22
Q54

Keywords:
Quantile vector autoregression
Counterfactual distribution forecasts
Disaster scenario analysis

ABSTRACT

We introduce quantile and moment impulse response functions for structural quantile vector autoregressive models. We use them to study how climate-related natural disasters affect the predictive distribution of output growth and inflation. Disasters strongly shift the forecast distribution particularly in the tails. They result in an initial sharp increase of the downside risk for growth, followed by a temporary rebound. Upside risk to inflation increases markedly for a few months and then subsides. As a result, natural disasters have a persistent impact on the conditional variance and skewness of macroeconomic aggregates which standard linear models estimating conditional mean dynamics fail to match. We perform a scenario analysis to evaluate the hypothetical effects of more frequent large disasters on the macroeconomy due to increased atmospheric carbon concentration. Our results indicate a substantially higher conditional volatility of growth and inflation as well as increased upside risk to inflation particularly in a scenario where only currently pledged climate policies are implemented.

1. Introduction

Natural disasters such as hurricanes, wildfires, floods and droughts not only threaten the safety and livelihood of the people directly affected by such events, but they can also disrupt economic activity at a larger scale. While disasters are relatively rare events, they have become more frequent and larger in magnitude in recent decades. This is widely attributed to an increase in global mean temperatures resulting from anthropogenic emissions of carbon dioxide (CO2) and other greenhouse gases. Quantifying the macroeconomic costs of physical risks arising from climate change requires a thorough understanding of the economic dynamics associated with natural disasters. While there is a growing body of work studying the economic effects of disasters, the existing literature has paid surprisingly little attention to the tail response of macroeconomic aggregates to such events.

In this paper, we quantify the impact of natural disasters on macroeconomic aggregates with a particular emphasis on conditional density forecasts. We first make a number of methodological contributions by extending the structural quantile vector autoregressive (QVAR) approach of Chavleishvili and Manganelli (2024). Specifically, we introduce the concept of *quantile impulse response functions* (QIRFs) as the difference between forecasts of regression quantiles with and without a shock to one of the variables in the QVAR. This allows us to trace the entire distribution of potential outcomes conditional on the occurrence of a shock. We then show how to employ these distributional forecasts to compute *moment impulse response functions* (MIRFs). The latter can be used to study the

☆ We thank Sydney Ludvigson, Serena Ng and Sai Ma for kindly sharing the Costly Disaster time series with us. We are particularly grateful to Jean-Paul Renne for useful input about the modeling of rare disasters with the Gamma-zero distribution. We further thank Laura Bakkensen, Matthias Kalkuhl, Kai Lessmann, Oliver Richters, Mark Weth and seminar participants at the IAAE for helpful comments and discussions.

* Corresponding author at: Finance Department, Frankfurt School of Finance & Management, Adickesallee 32-34, 60322 Frankfurt am Main, Germany.
E-mail addresses: sulkhan.chavleishvili@econ.au.dk (S. Chavleishvili), e.moench@fs.de (E. Moench).

https://doi.org/10.1016/j.jeconom.2024.105914
Received 24 February 2023; Received in revised form 8 September 2024; Accepted 10 September 2024
Available online xxxx

dynamic evolution of the moments of a forecast distribution conditional on a shock. We provide a subsampling algorithm to do inference for these QIRFs and MIRFs.

We apply this novel toolkit to trace the predictive distribution of macroeconomic outcomes conditional on the occurrence of natural disasters in the U.S. In our baseline analysis, we measure the intensity of natural disasters using the Costly Disaster (CD) series proposed by Ludvigson et al. (2020), from which we purge disasters that are unrelated to weather events. We start by showing that the distribution of the costs associated with natural disasters over the past several decades features a heavy right tail: large disasters are common, and increasingly so in recent years. As a result, standard econometric models designed to study the conditional mean dynamics following disasters are inadequate to fully capture their effects. Our results suggest that natural disasters substantially shift the distribution of key macroeconomic indicators. They increase the probability mass in the tails of the predictive distribution, making conditional forecasts substantially more uncertain and skewed. Specifically, disaster events immediately shift downward the conditional distribution of real growth, but subsequently give rise to a strong positive rebound effect. In contrast, the conditional density of inflation increases sharply on impact, but then fluctuates markedly in the following months. As a result, the conditional volatility of real growth and inflation is substantially elevated for some time. In response to disaster shocks, the conditional skewness is initially negative and subsequently positive for real growth and persistently positive for inflation. We subject our results to a host of robustness checks and collect them in the Online Appendix.[1]

Our findings shed light on the tradeoffs that monetary and other stabilization policies are likely to face in the context of climate change. As the near-term effects of large natural disasters resemble those of negative supply-shocks, standard monetary policy tools are likely to have limited effects on stabilizing inflation and output. By quantifying the macroeconomic tail effects of climate-related weather events, our findings can also inform the debate on the likely costs associated with different climate policies or their omission. As a first step in this direction, we use our model to study the macroeconomic impact of scenarios where large natural disasters occur more frequently as a result of increased concentrations of carbon in the atmosphere.

Specifically, we first model the disaster frequency and intensity as a function of the global atmospheric carbon concentration relying on the Gamma-zero distribution introduced in Monfort et al. (2017). We then study how the forecast distributions for output growth and inflation would change under two scenarios for the evolution of this concentration. We obtain these from the Network for the Greening of the Financial System (NGFS). The first scenario assumes that only currently pledged climate policies will take effect, the second posits a transition to net zero carbon emissions by 2050. While both are associated with substantially more mass on high disaster costs relative to the sample from 1980 through 2019, we document that the Current Policies scenario implies considerably higher probabilities of large costly disasters than the Net Zero scenario. Our results also show that while the distribution of the conditional mean of IP growth is fairly similar under both scenarios, there is substantially more volatility assuming that only current policies will be implemented. For CPI inflation, both the conditional mean and the conditional volatility are shifted up under the Current Policies scenario, suggesting that inflation will be substantially higher and more volatile if carbon emissions are not substantially reduced.

Our finding that disasters exert significant macroeconomic impact particularly in the far right tail of the disaster distribution is relevant in light of research emphasizing the potential implications of fat tails in the risk of catastrophic climate change. In a seminal paper, Weitzman (2009) has argued that the combination of heavy tails, unlimited exposure to climate events and high risk aversion implies that the expected loss from climate change is infinite. While his "dismal theorem" relies on knife-edge theoretical assumptions and thus may be of limited practical relevance, the insight that heavy tails in the disaster distribution may have substantial effects on the cost–benefit analysis of climate policies is more generally accepted Nordhaus (2011).

We contribute to studying the aggregate economic effects of natural disasters in the United States using macro-econometric techniques.[2] Kim et al. (2022) analyze the impact of extreme weather events in a smooth transition vector autoregressive framework and find that such events have had significant macroeconomic effects only in recent years. Consistent with our results, extreme weather events initially reduce industrial production growth and raise inflation. However, their approach does not allow to study the effects of large disasters on the entire *distribution* of macroeconomic outcomes. Another important difference between our paper and Kim et al. (2022) lies in the measurement of disaster events. They use the Actuaries Climate Index (ACI) to study the occurrence of extreme weather events across the U.S., whereas our baseline analysis is based on the Costly Disaster (CD) index of Ludvigson et al. (2020) which directly captures financial losses associated with disasters. While both series cover partly overlapping disaster events, the CD index focuses only on disasters which have resulted in significant economic costs. We document robustness with respect to using the ACI in the Online Appendix.

Ludvigson et al. (2020) study the impact of disaster shocks on the U.S. economy. Specifically, they order the CD series first in a recursively identified linear VAR and find that disaster shocks exert significant transitory effects on production and employment. They use the model to analyze the hypothetical effects of the COVID-19 pandemic. Importantly, however, their linear VAR model does not allow to study the effects of natural disasters on the conditional distribution of macroeconomic outcomes which is the focus of our paper. A different approach has been taken by Mohaddes et al. (2022). These authors measure deviations of temperature and

[1] First, we show that they are qualitatively unchanged when we rely on alternative measures of the intensity of natural disasters such as maximum sustained wind speeds as in Bakkensen and Barrage (2018) or the Actuaries Climate Index (ACI) of Kim et al. (2022). Second, we find substantial effects of disasters shocks for the Chicago Fed National Activity Index (CFNAI), a broader measure of aggregate economic activity than Industrial Production. Third, we document that natural disasters do not only affect inflation through their impact on food and energy prices, but also have persistent effects on the tails of core inflation. Finally, we report robustness of our results to the exclusion of the largest disaster in our sample, Hurricane Katrina in 2005.

[2] A vast literature analyzes economic effects of climate change using a variety of alternative modeling approaches, including computable general equilibrium models or integrated assessment models. We provide a selective review in Section 2 and the Online Appendix.

precipitation from their long-term moving-average values across U.S. states and use panel econometric techniques to show that these deviations have persistent adverse effects on real output, labor productivity and employment.

A large literature examines the macroeconomic effects of natural disasters globally. For example, Noy (2009) documents a larger short-run adverse impact of disasters on real GDP in developing compared to developed economies. Cavallo et al. (2013) use a similar dataset to show that these effects are mainly driven by the largest disasters.

A few authors have studied the impact of climate-related natural disasters on inflation. Parker (2018) finds that while storms and floods have minimal effects on inflation in developed economies, they have a more significant and persistent impact in developing markets. Similarly, Faccia et al. (2021) show that heat waves strongly affect food price inflation in emerging markets but have little or sometimes even negative effects on headline inflation

Methodologically, our paper is also loosely related to a new literature on structural VARs with infinite variance regressors. Davis and Ng (2022) show that when the endogenous variables have fat tails characterized by a Pareto distribution, the least squares estimator is consistent but standard inference relying on asymptotic normality of the estimators is invalid. They propose an estimator that is robust to infinite variance. Our inference does not assume a specific form of the limiting distribution. Instead, we rely on the accurate approximations of the entire distributions of the endogenous variables including the disaster series by estimating conditional quantile functions on the grid of quantiles.[3]

The paper is organized as follows. Section 2 sets the stage by briefly reviewing the mechanisms through which natural disasters affect economic activity which have been discussed in the theoretical literature. Section 3 presents the methodology. It starts by discussing estimation, inference, and forecasting in the Quantile VAR framework based on prior work. Building on these, it then introduces the novel concepts of quantile and moment impulse response functions. Section 4 applies our methodology to study the effects of natural disaster shocks on the conditional forecast distributions for output growth and inflation. Section 5 then provides an analysis of the effects of different climate policy scenarios on the frequency and intensity of natural disasters and, by implication, on the forecast distribution of macroeconomic aggregates. Section 6 concludes.

2. The economics of natural disasters

In this section, we aim to provide some brief guidance for interpreting our empirical results below. We do so by summarizing the main classes of models through which natural disasters have macroeconomic consequences.[4] Research distinguishes between direct and indirect disaster effects. Directly, disasters cause immediate damages. Indirectly, they interrupt economic activity and prompt positive spillovers from production substitution or reconstruction. These effects capture short- and long-term production and consumption losses. Typically, disasters prompt economic adjustment by causing sudden production factor losses, such as labor and capital, which lead to shifts of the economic equilibrium.

Growth models. The literature has explored the effects of natural disasters using neoclassical growth models in the Ramsey–Solow tradition. A few prominent examples are Albala-Bertrand (1993), Lusardi (1998), Noy and Nualsri (2007), Loayza et al. (2012), Okuyama (2003). In these models natural disasters can push the economy away from the steady-state by destroying part of the capital stock and/or reducing the labor force. Recovery involves increased savings and rebuilding the capital stock. This leads to a temporary output drop immediately after the disaster followed by a recovery. However, if the labor force shrinks, total output may permanently decrease due to insufficient savings. These models predict a short-term negative impact on output, while the longer-term effects on output can be positive or negative. Furthermore, these models commonly focus on real outcomes and cannot speak to the effects of disasters on inflation.

Input–output models. Input–Output (I–O) models analyze the economic impacts of natural disasters by assuming fixed output proportions among sectors, affecting trade and output economy-wide. Shocks, like disasters, can impact demand or supply of one or more sectors. If consumer demand is affected, all intermediate demands and supplies adjust. If supply is affected, final demand restores economic equilibrium. "Inoperability Input–Output" (IIM) models (for an early example see Santos and Haimes, 2004) capture temporary sector inoperability due to labor or capital stock reductions from disaster shocks. Extended I–O models examine how local disasters can affect other regions or the overall economy (see, e.g., Okuyama, 2004, Sohn et al., 2004 or Hallegatte, 2008). While standard I–O models assume constant prices, Hallegatte (2008) introduces price adjustments during reconstruction. Despite their potential relevance, models analyzing the granular origins of aggregate fluctuations (e.g., Gabaix, 2011, Acemoglu et al., 2012) have – to the best of our knowledge – not yet been applied to natural disaster analysis.

Computable general equilibrium models. Computable General Equilibrium (CGE) macroeconomic models simulate how disruptions to the supply of goods and services affect overall output and prices. They are typically more flexible than I–O models because they include demand and supply in various markets in equilibrium and can allow for nonlinear effects such as economies of scale and nonlinear impact functions. They are thought to better capture the longer-run economic consequences of natural disasters (see, e.g., Rose and Liao, 2005). However, their use in disaster analysis is relatively scarce. Exceptions include Carrera et al. (2015), who study the economic impact of floods in Northern Italy and find substantial indirect economic losses and Rose et al. (2016) who integrate a CGE model into an I–O model of the U.S. economy to study the potential impact of tsunamis hitting the Californian Pacific coast.

[3] The limiting distribution of the quantile regression estimator when variables have infinite variance is an important research question yet to be explored (Koenker, 2005, see pages 127–128 for a brief discussion on the topic).

[4] The Online Appendix contains a more detailed discussion of the economics of natural disasters. This closely follows the excellent literature review of in Botzen et al. (2019) to which we refer the reader for more details.

Integrated assessment models. Integrated assessment models (IAMs) are key in climate policy analyses, combining data and modeling elements from various disciplines. They establish baseline scenarios and model how policy interventions such as carbon taxes affect greenhouse gas emissions and by implication global temperatures. IAMs typically rely on "damage functions" to summarize the impact of higher temperatures on the economy. While numerous IAMs exist, many are extensions of the seminal "DICE" model by Nordhaus (1993). For instance, the REMIND-MAgPIE model used by the NGFS for scenario analysis, from which we obtain the atmospheric carbon concentration pathways used in Section 5, is a modern variant of an IAM which combines different elements. At the core is a Ramsey-type optimal growth model. This is complemented with an energy system module that includes a detailed sectoral representation of energy supply and demand. The MAgPIE extension captures land-use dynamics, and the model is linked to the climate model MAGICC to account for changes in global temperatures. For further details, see Hilaire and Bertram (2020).

Note that while many of these models discussed above are inherently nonlinear and as such may give rise to interesting predictions about the dynamics of conditional higher moments of macroeconomic aggregates to disaster shocks, to the best of our knowledge the theoretical literature has thus far not systematically explored these predictions.

3. Quantile VARs: Estimation, inference, and impulse response analysis

This section introduces the novel concepts of quantile impulse response functions (QIRFs) and moment impulse response functions (MIRFs). These form the basis of our empirical analysis in Section 4, where we trace the dynamic distributional effects of natural disasters on the U.S. economy. Since they are defined in terms of quantiles and moments of the forecast distribution, we first introduce the QVAR framework developed in prior work in Section 3.1. We then discuss estimation in Section 3.2, forecasting in Section 3.3, and inference in Section 3.4. Finally, we describe our novel algorithms to compute QIRFs and MIRFs in Section 3.5.

3.1. The quantile vector autoregression model

As shown by Chavleishvili and Manganelli (2024), one can extend the quantile autoregression framework of Koenker and Xiao (2006) to a multivariate setting. Let $\boldsymbol{X}_t = (X_{1,t}, \dots, X_{K,t})'$ denote a vector of observables and $\{u_{i,t}\}$ be a sequence of *i.i.d.* standard uniform random variables for $i = 1, \dots, K$, with the random variables $u_{1,t}, \dots, u_{K,t}$ being independent of each other. Then, the VAR model with random parameters of order $p \geq 1$ is given by

$$\boldsymbol{X}_t = \boldsymbol{\omega}(\boldsymbol{u}_t) + \sum_{j=1}^{p} \boldsymbol{\Phi}_j(\boldsymbol{u}_t)\boldsymbol{X}_{t-j} + \boldsymbol{\Gamma}(\boldsymbol{u}_t)\boldsymbol{X}_t, \tag{1}$$

where $\boldsymbol{\omega}(\boldsymbol{u}_t), \boldsymbol{\Phi}_j(\boldsymbol{u}_t)$ for $j = 1, \dots, p$, and $\boldsymbol{\Gamma}(\boldsymbol{u}_t)$ are $K \times 1$, $K \times K$ and $K \times K$ matrices of unknown parameters given as

$$\boldsymbol{\omega}(\boldsymbol{u}_t) = \left[\omega_1(u_{1,t}), \dots, \omega_K(u_{K,t})\right]', \tag{2}$$

$$\boldsymbol{\Phi}_j(\boldsymbol{u}_t) = \begin{bmatrix} \Phi_{j,1,1}(u_{1,t}) & \dots & \Phi_{j,1,K}(u_{1,t}) \\ \vdots & & \vdots \\ \Phi_{j,K,1}(u_{K,t}) & \dots & \Phi_{j,K,K}(u_{K,t}) \end{bmatrix}, \tag{3}$$

$$\boldsymbol{\Gamma}(\boldsymbol{u}_t) = \begin{bmatrix} 0 & 0 & 0 & \dots & 0 & 0 \\ \Gamma_{2,1}(u_{2,t}) & 0 & 0 & \dots & 0 & 0 \\ \Gamma_{3,1}(u_{3,t}) & \Gamma_{3,2}(u_{3,t}) & 0 & \dots & 0 & 0 \\ \vdots & \vdots & & & & \vdots \\ \Gamma_{K,1}(u_{K,t}) & \Gamma_{K,2}(u_{K,t}) & \Gamma_{K,3}(u_{K,t}) & \dots & \Gamma_{K,K-1}(u_{K,t}) & 0 \end{bmatrix}. \tag{4}$$

Since the matrix $\boldsymbol{\Gamma}(\boldsymbol{u}_t)$ is lower triangular with zeros on the main diagonal, contemporaneous interactions among the variables in the vector $\boldsymbol{X}_t$ are identified through their ordering. This is in the spirit of recursive identification of structural economic shocks in linear VARs following the seminal work of Sims (1980). We discuss conditions for stationarity and geometric ergodicity of the vector $\boldsymbol{X}_t$ in the Online Appendix.

As shown by Koenker and Xiao (2006), an autoregression with random parameters gives rise to a dynamic conditional quantile function with quantile-specific parameters under certain monotonicity conditions. Adopting this to our multivariate framework, let $\boldsymbol{I}_{1,t} = \left(\boldsymbol{X}'_{t-1}, \boldsymbol{X}'_{t-2}, \dots, \boldsymbol{X}'_{t-p}\right)'$ and $\boldsymbol{I}_{i,t} = (\boldsymbol{I}'_{i-1,t}, X_{i-1,t})'$ for $i = 2, \dots, K$ be a recursive information set. The arguments of Koenker and Xiao (2006) and Chavleishvili and Manganelli (2024) imply that for any $i = 1, \dots, K$, if the random variable $X_{i,t}$ is monotonically increasing in $u_{i,t}$ given $\boldsymbol{I}_{i,t}$, then the conditional quantile function of $X_{i,t}$ is given as

$$q_{X_{i,t}}(\theta_i, \boldsymbol{Z}_{i,t}) = \boldsymbol{Z}'_{i,t}\boldsymbol{\beta}_i(\theta_i), \tag{5}$$

where $\boldsymbol{Z}_{i,t} = (1, \boldsymbol{I}'_{i,t})'$ and $\boldsymbol{\beta}_i(\theta_i)$ contain the parameters of equation i of the System (1) evaluated at the quantile $\theta_i \in (0,1)$.[5] In particular $\boldsymbol{\beta}_1(\theta_1) = \left[\omega_1(\theta_1), \boldsymbol{\Phi}_{1,1}(\theta_1), \dots, \boldsymbol{\Phi}_{p,1}(\theta_1)\right]'$, $\boldsymbol{\beta}_2(\theta_2) = \left[\omega_2(\theta_2), \boldsymbol{\Phi}_{1,2}(\theta_2), \dots, \boldsymbol{\Phi}_{p,2}(\theta_2), \Gamma_{2,1}(\theta_2)\right]'$ and for $i = 3, \dots, K$ we have $\boldsymbol{\beta}_i(\theta_i) = \left[\omega_i(\theta_i), \boldsymbol{\Phi}_{1,i}(\theta_i), \dots, \boldsymbol{\Phi}_{p,i}(\theta_i), \Gamma_{i,1}(\theta_i), \dots, \Gamma_{i,i-1}(\theta_i)\right]'$, where $\boldsymbol{\Phi}_{j,i}(\cdot)$ is the ith row of the matrix $\boldsymbol{\Phi}_j(\cdot)$.

This reveals a very flexible feature of the data-generating process (1) as it gives rise to the dynamic conditional quantile function (5) with quantile-specific parameters. Hence, the random variables $\{X_{1,t}, \dots, X_{K,t}\}$ mutually affect the entire shape of each other's conditional distributions and not just their conditional mean as implied by a standard linear VAR model.[6]

[5] Let the random variable $X_{i,t}$ be denoted as $X_{i,t}(u_{i,t})$, indicating its dependence on $u_{i,t}$. Then $q_{X_{i,t}(u_{i,t})}(\theta_i, \boldsymbol{Z}_{i,t}) = X_{i,t}(q_{u_{i,t}}(\theta_i))$ is an immediate consequence of monotonicity. If we suspect that monotonicity fails, one can use for instance a rearrangement method by Chernozhukov et al. (2010).

3.2. Estimation

Note that the function $q_{X_{i,t}}(\theta_i, \boldsymbol{Z}_{i,t})$, defined in Eq. (5), must satisfy the conditional quantile restriction

$$\Pr\left[X_{i,t} < q_{X_{i,t}}(\theta_i, \boldsymbol{Z}_{i,t}) | \boldsymbol{I}_{i,t}\right] = \theta_i, \quad i = 1, \dots, K, \tag{6}$$

and its estimates are defined as $\hat{q}_{X_{i,t}}(\theta_i, \boldsymbol{Z}_{i,t}) = \boldsymbol{Z}'_{i,t}\hat{\boldsymbol{\beta}}_i(\theta_i)$, where

$$\hat{\boldsymbol{\beta}}_i(\theta_i) = \arg\min_{\boldsymbol{\beta}_i \in \mathbb{R}^{pK+i}} \sum_{t=p+1}^{T} \rho_{\theta_i}(X_{i,t} - \boldsymbol{Z}'_{i,t}\boldsymbol{\beta}_i). \tag{7}$$

Here, $\rho_\theta(v) = v(\theta - I(v < 0))$ is an asymmetric loss function introduced by Koenker and Bassett (1978), $I(\cdot)$ is an indicator function and $\theta_i \in (0, 1)$ is a quantile index. In practice, estimation of the parameters is done equation by equation using the interior point algorithm of Koenker and Park (1996).

The asymptotic properties of dynamic quantile regressions have been studied extensively in the literature. Most importantly, White et al. (2015, Theorem 1 and Theorem 2) show that $\hat{\boldsymbol{\beta}}_i(\theta_i)$ is consistent and asymptotically normal. The key input for distributional impulse responses that we introduce below are conditional forecasts. As these depend on a large number of parameters, it is computationally challenging to apply the asymptotic results of White et al. (2015). In order to make feasible inference, we therefore rely on subsampling. We discuss this approach in Section 3.4 below.

3.3. Forecasting

Note that there is an inverse mapping from Eq. (5) to the vector $\boldsymbol{X}_t$ in Eq. (1). Hence, when we utilize the regression quantile framework to estimate the unknown parameters of the model, we can subsequently simulate the data to run counterfactual economic analyses. To illustrate this, consider the case when $K = 2$, and $p = 1$, and suppose $X_{1,T+1}$ is subject to a shock of size κ_1. This will have an instantaneous effect on the entire conditional distribution, not just the mean, of the random variable $X_{2,T+1}$ and thus also on the entire conditional distributions of the future values $(X_{1,T+l}, X_{2,T+l})$, $l = 2, \dots, h$. Hence, we can attribute particular dynamics in the conditional distribution of the endogenous variables to this shock. To investigate this impact, we can simulate a standard uniform random variable $u^{(1)}_{2,1}$ and calculate $X^{(1)}_{2,T+1} = q_{X_{2,T+1}}\left(u^{(1)}_{2,1}, \boldsymbol{Z}_{2,T+1}\right)$ where $\boldsymbol{Z}_{2,T+1} = \left(1, X_{1,T}, X_{2,T}, \kappa_1\right)'$. To simulate further future values, again simulate the independent standard uniform random variables $u^{(1)}_{1,2}$ and $u^{(1)}_{2,2}$ and construct $X^{(1)}_{1,T+2} = q_{X_{1,T+2}}\left(u^{(1)}_{1,2}, \boldsymbol{Z}^{(1)}_{1,T+2}\right)$ and $X^{(1)}_{2,T+2} = q_{X_{2,T+2}}\left(u^{(1)}_{2,2}, \boldsymbol{Z}^{(1)}_{2,T+2}\right)$ where $\boldsymbol{Z}^{(1)}_{1,T+2} = \left(1, \kappa_1, X^{(1)}_{2,T+1}\right)$ and $\boldsymbol{Z}^{(1)}_{2,T+2} = \left(1, \kappa_1, X^{(1)}_{2,T+1}, X^{(1)}_{1,T+2}\right)$.

This procedure can be iterated further and repeated multiple times such that one obtains multiple forecast paths $\left(\kappa_1, X^{(s)}_{2,T+1}\right)$, $\left(X^{(s)}_{1,T+2}, X^{(s)}_{2,T+2}\right), \dots, \left(X^{(s)}_{1,T+h}, X^{(s)}_{2,T+h}\right)$ for $s = 1, \dots, S$. Hence, the conditional distributions of each variable can be studied at any given forecast horizon $l = 1, \dots, h$. This allows us to generate a random variable using its quantile function and standard uniform random variables.

This forecasting technique, which is sometimes referred to as inverse transform sampling, has also been used by Xiao (2017, Section 17.9) in a univariate context, and by Wei (2008, Lemma 1) and Chavleishvili and Manganelli (2024, Section 2.2) in the triangular multivariate case. In practice, the forecasting procedure entails setting up a grid of N quantile indices $0 < \theta_1 < \theta_2 < \cdots < \theta_N < 1$ and estimating the parameters of the model at those quantiles for each equation i.[7] The forecasting is then carried out by randomly and independently drawing from those parameters. The following sampling procedure for forecasting is a generalization to multiple lags of the approach in Chavleishvili and Manganelli (2024, Section 2.2). In Section 3.5, we then extend this idea to design the bootstrap and subsampling procedures for inference purposes.

Algorithm 1 (*Sampling for Forecasting*)**.** The following algorithm constructs the forecast distribution of the vector $\boldsymbol{X}_t$.

(1) Specify a finite grid of quantiles

$$0 < \theta_1 < \theta_2 < \cdots < \theta_N < 1. \tag{8}$$

Estimate the parameters for each equation of the model (1)–(5) at these grid values, construct and store the matrices $\hat{\boldsymbol{\omega}}(\boldsymbol{\theta}), \hat{\boldsymbol{\Phi}}_1(\boldsymbol{\theta}), \dots, \hat{\boldsymbol{\Phi}}_p(\boldsymbol{\theta})$, and $\hat{\boldsymbol{\Gamma}}(\boldsymbol{\theta})$ at each quantile $\boldsymbol{\theta} = (\theta_1, \dots, \theta_K)'$ from the grid (8).

[6] In a standard linear VAR model, the parameter matrices $\boldsymbol{\Phi}_j(\boldsymbol{u}_t) = \boldsymbol{\Phi}_j$, for $j = 1, \dots, p$ and $\boldsymbol{\Gamma}(\boldsymbol{u}_t) = \boldsymbol{\Gamma}$ are constants for any t. Then, a conditional quantile function of the random variable $X_{i,t}$ will only have a quantile specific intercept and will thus only affect its conditional mean.

[7] As discussed by Portnoy (1991), in finite samples the number of distinct quantiles are limited. This justifies the use of a grid.

(2) Simulate independently a sequence of a $K \times 1$ vector of uniform random variables $\left\{ \boldsymbol{u}_l^{(s)} \right\}_{l=1}^{h}$ and use the coefficient matrices corresponding to the respective quantiles given by these draws. Based on the contemporaneous and lagged vector of data $\boldsymbol{X}_T, \boldsymbol{X}_{T-1}, \dots, \boldsymbol{X}_{T-p+1}$, and the estimated parameters for the quantile $\left\{ \boldsymbol{u}_l^{(s)} \right\}_{l=1}^{h}$, calculate the out-of sample values as follows

$$
\begin{aligned}
\boldsymbol{X}_{T+1}^{(s)} &= \hat{\boldsymbol{\omega}}\left(\boldsymbol{u}_1^{(s)}\right) + \hat{\boldsymbol{\Phi}}_1\left(\boldsymbol{u}_1^{(s)}\right)\boldsymbol{X}_T + \hat{\boldsymbol{\Phi}}_2\left(\boldsymbol{u}_1^{(s)}\right)\boldsymbol{X}_{T-1} + \cdots + \hat{\boldsymbol{\Phi}}_p\left(\boldsymbol{u}_1^{(s)}\right)\boldsymbol{X}_{T-p+1} \\
&+ \hat{\boldsymbol{\Gamma}}\left(\boldsymbol{u}_1^{(s)}\right)\boldsymbol{X}_{T+1}^{(s)}, \\
\boldsymbol{X}_{T+2}^{(s)} &= \hat{\boldsymbol{\omega}}\left(\boldsymbol{u}_2^{(s)}\right) + \hat{\boldsymbol{\Phi}}_1\left(\boldsymbol{u}_2^{(s)}\right)\boldsymbol{X}_{T+1}^{(s)} + \hat{\boldsymbol{\Phi}}_2\left(\boldsymbol{u}_2^{(s)}\right)\boldsymbol{X}_T + \cdots + \hat{\boldsymbol{\Phi}}_p\left(\boldsymbol{u}_2^{(s)}\right)\boldsymbol{X}_{T-p+2} \\
&+ \hat{\boldsymbol{\Gamma}}\left(\boldsymbol{u}_2^{(s)}\right)\boldsymbol{X}_{T+2}^{(s)}, \\
&\vdots \\
\boldsymbol{X}_{T+l}^{(s)} &= \hat{\boldsymbol{\omega}}\left(\boldsymbol{u}_l^{(s)}\right) + \hat{\boldsymbol{\Phi}}_1\left(\boldsymbol{u}_l^{(s)}\right)\boldsymbol{X}_{T+l-1}^{(s)} + \hat{\boldsymbol{\Phi}}_2\left(\boldsymbol{u}_l^{(s)}\right)\boldsymbol{X}_{T+l-2}^{(s)} + \cdots + \hat{\boldsymbol{\Phi}}_p\left(\boldsymbol{u}_l^{(s)}\right)\boldsymbol{X}_{T+l-p}^{(s)} \\
&+ \hat{\boldsymbol{\Gamma}}\left(\boldsymbol{u}_l^{(s)}\right)\boldsymbol{X}_{T+l}^{(s)}, \quad l = p+1, \dots, h.
\end{aligned}
$$

(3) Repeat step (2) $S >> 0$ times and collect the forecasts $\boldsymbol{\zeta}_{S,l} = \left\{ \boldsymbol{X}_{T+l}^{(1)}, \dots, \boldsymbol{X}_{T+l}^{(S)} \right\}$ for $l = 1, \dots, h$.

The forecasting algorithm above is based on the assumption that no shock is hitting the variables in $\boldsymbol{X}$. However, we can also use Algorithm 1 to compute forecasts $\boldsymbol{\zeta}_{S,l}^{*} = \left\{ \boldsymbol{X}_{T+l}^{*(1)}, \dots, \boldsymbol{X}_{T+l}^{*(S)} \right\}$ conditional on a shock $X_{i,T+1} = \kappa_i$ for any $l = 1, \dots, h$. Specifically, $\boldsymbol{X}_{i,T+1}$ is replaced by κ_i in each repetition of the sampling algorithm. In Section 3.5 below, we make use of this insight to construct distributional impulse response functions.

3.4. Inference

For inference, we rely on the quantile bootstrap approach proposed by Koenker (1994), which is closely related to the forecasting Algorithm 1. The difference is that the quantile bootstrap regenerates the data in-sample, while Algorithm 1 constructs them out-of-sample. Since it relies on sampling from the estimated quantile functions, as long as the conditional quantile functions are accurately specified, this procedure can be applied to any type of model.[8] The following algorithm adapts his procedure to our model.

Algorithm 2 (*QVAR Bootstrap*)**.**

(1) For a finite grid of quantiles (8), estimate for each equation i the parameters of the model (1)–(5) at these grid values, construct and store the matrices $\hat{\boldsymbol{\omega}}(\theta), \hat{\boldsymbol{\Phi}}_1(\theta), \dots, \hat{\boldsymbol{\Phi}}_p(\theta)$, and $\hat{\boldsymbol{\Gamma}}(\theta)$. Construct estimates of the conditional quantile functions $\hat{q}_{X_{i,t}}(\theta_j, \boldsymbol{Z}_{i,t})$ for each equation $i = 1, \dots, K$ and quantile $j = 1, \dots, N$.

(2) For each $p < t \leq T$ and equation $i = 1, \dots, K$, find the θ_j from the grid (8) that minimizes the distance $|X_{i,t} - \hat{q}_{X_{i,t}}(\theta_j, \boldsymbol{Z}_{i,t})|$. Collect and store those values as $\{\theta_t\}_{t=p+1}^{T}$.

(3) Draw independently a sequence of a $K \times 1$ vector of random variables $\left\{ \boldsymbol{u}_t^{(b)} \right\}_{t=p+1}^{T}$ from the sequence $\{\theta_t\}_{t=p+1}^{T}$ and generate the bootstrap sample recursively as follows

$$
\begin{aligned}
\boldsymbol{X}_{p+1}^{(b)} &= \hat{\boldsymbol{\omega}}\left(\boldsymbol{u}_{p+1}^{(b)}\right) + \hat{\boldsymbol{\Phi}}_1\left(\boldsymbol{u}_{p+1}^{(b)}\right)\boldsymbol{X}_p + \hat{\boldsymbol{\Phi}}_2\left(\boldsymbol{u}_{p+1}^{(b)}\right)\boldsymbol{X}_{p-1} + \cdots + \hat{\boldsymbol{\Phi}}_p\left(\boldsymbol{u}_{p+1}^{(b)}\right)\boldsymbol{X}_1 \\
&+ \hat{\boldsymbol{\Gamma}}\left(\boldsymbol{u}_{p+1}^{(b)}\right)\boldsymbol{X}_{p+1}^{(b)}, \\
\boldsymbol{X}_{p+2}^{(b)} &= \hat{\boldsymbol{\omega}}\left(\boldsymbol{u}_{p+2}^{(b)}\right) + \hat{\boldsymbol{\Phi}}_1\left(\boldsymbol{u}_{p+2}^{(b)}\right)\boldsymbol{X}_{p+1}^{(b)} + \hat{\boldsymbol{\Phi}}_2\left(\boldsymbol{u}_{p+2}^{(b)}\right)\boldsymbol{X}_p + \cdots + \hat{\boldsymbol{\Phi}}_p\left(\boldsymbol{u}_{p+2}^{(b)}\right)\boldsymbol{X}_2 \\
&+ \hat{\boldsymbol{\Gamma}}\left(\boldsymbol{u}_{p+2}^{(b)}\right)\boldsymbol{X}_{p+2}^{(b)}, \\
&\vdots \\
\boldsymbol{X}_t^{(b)} &= \hat{\boldsymbol{\omega}}\left(\boldsymbol{u}_t^{(b)}\right) + \hat{\boldsymbol{\Phi}}_1\left(\boldsymbol{u}_t^{(b)}\right)\boldsymbol{X}_{t-1}^{(b)} + \hat{\boldsymbol{\Phi}}_2\left(\boldsymbol{u}_t^{(b)}\right)\boldsymbol{X}_{t-2}^{(b)} + \cdots + \hat{\boldsymbol{\Phi}}_p\left(\boldsymbol{u}_t^{(b)}\right)\boldsymbol{X}_{t-p}^{(b)} \\
&+ \hat{\boldsymbol{\Gamma}}\left(\boldsymbol{u}_t^{(b)}\right)\boldsymbol{X}_t^{(b)}, \quad t = 2p+1, \dots, T.
\end{aligned}
$$

(4) Repeat step (3) $B >> 0$ times and collect the bootstrap sample $\left\{ \left\{ \boldsymbol{X}_t^{(b)} \right\}_{t=p+1}^{T} \right\}_{b=1}^{B}$.

Note that if the parameter matrices $\boldsymbol{\Gamma}(\theta)$ and $\boldsymbol{\Phi}_j(\theta)$ for $j = 1, \dots, p$ are identical across quantiles, then Algorithm 2 is equivalent to the standard recursive residual-based bootstrap applied to linear dynamic time-series models with constant parameters (Gonçalves and Kilian, 2004).

[8] Note that the quantile bootstrap by Koenker (1994) is directly related to the rearrangement technique proposed by Chernozhukov et al. (2010) for monotonization of quantile functions. In fact, these authors introduce the term *quantile bootstrap*.

3.5. Distributional impulse response analysis

In this section we introduce two different but related concepts of impulse response functions (IRFs). First, we quantify the effect of structural shocks on the quantiles of the forecast distribution. Second, we explore how specific moments of the forecast distribution respond to these shocks.

3.5.1. Quantile IRFs

We assess the statistical significance of a shock to a variable $X_{j,t}$ in the QVAR model (5) on the entire conditional forecast distribution of a variable $X_{i,t}$, for $i, j = 1, \dots, K$. We do so by defining quantile IRFs (QIRFs) as the difference between quantiles of the population conditional forecast distribution of a random variable $X_{i,t}$ with and without the shock

$$\boldsymbol{\Delta}_{i,j}(\theta, l) \equiv q_{X_{i,T+l}}(\theta, X_{j,T+1} = \kappa_j, \boldsymbol{X}'_T, \dots, \boldsymbol{X}'_{T-p+1}) - q_{X_{i,T+l}}(\theta, \boldsymbol{X}'_T, \dots, \boldsymbol{X}'_{T-p+1}), \tag{9}$$

for $l \geq 1$, where κ_j is the size of the shock hitting $X_{j,T+1}$ and $\theta \in (0, 1)$ is the quantile of the forecast distribution. As a result, the QIRFs generalize the standard concept of IRFs which trace the mean responses of the conditional forecast distribution (Lütkepohl, 2005).[9]

In line with Chernozhukov et al. (2013), we label the forecast distribution conditional on the shock the "counterfactual distribution" and that conditional on no shock hitting the "factual distribution". The QIRF compares the same quantiles of these two forecast distributions. In particular, at any horizon $l \geq 1$ we can calculate the empirical quantiles θ of the two sets of forecasts $\boldsymbol{\zeta}^*_{S,l} = \left\{ \boldsymbol{X}^{*(1)}_{T+l}, \dots, \boldsymbol{X}^{*(S)}_{T+l} \right\}$ and $\boldsymbol{\zeta}_{S,l} = \left\{ \boldsymbol{X}^{(1)}_{T+l}, \dots, \boldsymbol{X}^{(S)}_{T+l} \right\}$ calculated using Algorithm 1 for a large number of draws S, where the former is constructed conditional on the shock $X_{j,T+1} = \kappa_j$. We define the estimator $\hat{\boldsymbol{\Delta}}_{i,j}(\theta, l)$ of the quantile impulse response functions in Eq. (9) as the difference between these two empirical quantiles. This allows us to measure the impact of a shock at various forecast horizons and forecast quantiles. Note that our approach is similar in spirit to that for mean IRFs in nonlinear dynamic models which are calculated as empirical averages of predictions obtained through simulations (Koop et al., 1996).

Further note that since they are empirical quantiles of forecasts, these forecast distributions are monotone by construction. However, their difference does not have to be. As a result, the QIRFs defined above can be non-monotone implying that a given quantile impulse response at a given forecast horizon can be larger or smaller than impulse responses for adjacent quantiles at the same horizon.

It is important to stress that the QIRFs defined in Eq. (9) bear similarities with the QIRFs introduced in Chavleishvili and Manganelli (2024, Section 3). Their concept, influenced by a literature on stress testing, compares future values of a variable with and without a shock calculated using a sequence of specific quantile indices. In contrast, we compute the difference of empirical quantiles of forecasts with and without a shock calculated using a large set of draws of *all* quantile indices. As a result, we compare entire forecast distributions rather than specific quantiles, and thus our approach is more general.

Note that the QIRFs in Eq. (9) are calculated using a large number of parameters. To facilitate inference, we suggest the following subsampling algorithm. This procedure does not require re-computation of the parameters and is therefore straightforward to apply.

Algorithm 3 (*Subsampling*)**.**

(1) Apply Algorithm 1 and produce a set of S predicted values $\boldsymbol{\zeta}^*_{S,l} = \left\{ \boldsymbol{X}^{*(1)}_{T+l}, \dots, \boldsymbol{X}^{*(S)}_{T+l} \right\}$ and $\boldsymbol{\zeta}_{S,l} = \left\{ \boldsymbol{X}^{(1)}_{T+l}, \dots, \boldsymbol{X}^{(S)}_{T+l} \right\}$, where $\boldsymbol{\zeta}^*_{S,l}$ is conditional on the shock $X_{j,T+1} = \kappa_j$, $l = 1, \dots, h$, $S >> 0$ and $j = 1, \dots, K$. Define the estimator $\hat{\boldsymbol{\Delta}}_{i,j}(\theta, l)$ of the QIRFs in Equation (9) as the difference between empirical quantiles θ of the sets $\boldsymbol{\zeta}^*_{S,l}$ and $\boldsymbol{\zeta}_{S,l}$ at any l, $\theta \in (0, 1)$ and for $i = 1, \dots, K$.

(2) For a given l from the sets $\boldsymbol{\zeta}^*_{S,l}$ and $\boldsymbol{\zeta}_{S,l}$, randomly construct n subsamples each of size m. Calculate empirical quantiles for each subsample $r = 1, \dots, n$ at a quantile index $\theta \in (0, 1)$ and calculate the QIRFs $\hat{\boldsymbol{\Delta}}^{(r)}_{i,j}(\theta, l)$ as their difference.

(3) For a significance level α, calculate the $(\alpha/2)$ and $(1 - \alpha/2)$ empirical quantiles of the set $\left\{ \hat{\boldsymbol{\Delta}}^{(r)}_{i,j}(\theta, l) - \hat{\boldsymbol{\Delta}}_{i,j}(\theta, l) \right\}$ and denote them as $\hat{t}_{\boldsymbol{\Delta}_{i,j}(\theta,l)}(\alpha/2)$ and $\hat{t}_{\boldsymbol{\Delta}_{i,j}(\theta,l)}(1 - \alpha/2)$. Then define the confidence intervals for the QIRFs as follows:

$$\mathrm{CI}_{\boldsymbol{\Delta}_{i,j}(\theta,l)}(\alpha) \equiv \left[\hat{\boldsymbol{\Delta}}_{i,j}(l, \theta) - \hat{t}_{\boldsymbol{\Delta}_{i,j}(\theta,l)}(1 - \alpha/2), \hat{\boldsymbol{\Delta}}_{i,j}(l, \theta) - \hat{t}_{\boldsymbol{\Delta}_{i,j}(\theta,l)}(\alpha/2) \right].$$

These confidence intervals are the percentile intervals of Hall (1992), as discussed e.g. by Lütkepohl (2005, p. 710). Suppose $T \to \infty$, then for $m, S \to \infty$ and $m^2/S \to 0$ the subsampling procedure with replacement has been shown to be valid in Politis et al. (1999, Theorem 2.2.1 and Corollary 2.3.1.). These authors also discuss various rules for selecting the subsample size m and number n with emphasis either on the minimum variance or stability of the confidence intervals of estimates. In practice, simple rules prove to be effective. In particular, we follow Sakov and Bickel (2000) and use the rule $m = \lfloor \tau S^{2/5} \rfloor$, where $\lfloor \cdot \rfloor$ stands for the integer part, with $\tau = 5$, $S = 10^5$ and $n = 10^4$.[10]

[9] Note that if the data-generating process of $\boldsymbol{X}_t$ is Gaussian, then the median IRF will correspond to a standard impulse response function $\boldsymbol{\Delta}_{i,j}(0.5, l) \approx \mathrm{E}(X_{i,T+l} | X_{j,T+1} = \kappa_j, \boldsymbol{X}'_T, \dots, \boldsymbol{X}'_{T-p+1}) - \mathrm{E}(X_{i,T+l} | \boldsymbol{X}'_T, \dots, \boldsymbol{X}'_{T-p+1}), \quad l \geq 1$.

[10] See for instance Chernozhukov (2002) and Chernozhukov and Fernández-Val (2005) for a related discussion. We also consider $\tau = 4, 6$ for the robustness checks reported in the Online Appendix.

3.5.2. Moment IRFs

We have shown how to trace the effect of structural shocks on the entire forecast distribution of variables in a QVAR. It is instructive to summarize these forecast distributions using the first four moments. We define moment IRFs (MIRFs) as the difference between moments of the conditional forecast distribution of a random variable $X_{i,t}$ with and without a shock:

$$\delta_{i,j}(l) \equiv m_{X_{i,T+l}}(X_{j,T+1} = \kappa_j, \boldsymbol{X}'_T, \dots, \boldsymbol{X}'_{T-p+1}) - m_{X_{i,T+l}}(\boldsymbol{X}'_T, \dots, \boldsymbol{X}'_{T-p+1}), \tag{10}$$

for $l \geq 1$, $i, j = 1, \dots, K$, where κ_j is the size of the shock hitting $X_{j,T+1}$ and where the function $m_{X_{i,T+l}}(\cdot)$ represents one of the first four conditional moments of variable $X_{i,T+l}$ at forecast horizon l.

The following procedure describes the computation of the MIRF via random sampling. First, we generate forecasts $\boldsymbol{\zeta}^*_{S,l} \equiv \left\{ \boldsymbol{X}^{*(1)}_{T+l}, \dots, \boldsymbol{X}^{*(S)}_{T+l} \right\}$ and $\boldsymbol{\zeta}_{S,l} \equiv \left\{ \boldsymbol{X}^{(1)}_{T+l}, \dots, \boldsymbol{X}^{(S)}_{T+l} \right\}$ using Algorithm 1, where $\boldsymbol{\zeta}^*_{S,l}$ is conditional on the shock $X_{j,T+1} = \kappa_j$. Second, we calculate the empirical mean, standard deviation, skewness and kurtosis of each pair of forecast sets $\boldsymbol{\zeta}^*_{S,l}$ and $\boldsymbol{\zeta}_{S,l}$ for any forecast horizon l. We define the estimator $\hat{\boldsymbol{\delta}}_{i,j}(l)$ of the moment impulse response functions in Eq. (10) as the difference between these two empirical moments.

For inference, we again adapt Algorithm 3 to the definition of the MIRFs and calculate the confidence intervals accordingly. In particular, we define the confidence interval at a significance level α as $\widehat{\mathrm{CI}}_{\delta_{i,j}(l)}(\alpha) \equiv \left[\hat{\boldsymbol{\delta}}_{i,j}(l) - \hat{\imath}_{\delta_{i,j}(l)}(1-\alpha/2), \hat{\boldsymbol{\delta}}_{i,j}(l) - \hat{\imath}_{\delta_{i,j}(l)}(\alpha/2) \right]$, where the quantities $\hat{\imath}_{\delta_{i,j}(l)}(1-\alpha/2)$ and $\hat{\imath}_{\delta_{i,j}(l)}(\alpha/2)$ are the $(1-\alpha/2)$ and $(\alpha/2)$ empirical quantiles of the set $\left\{ \hat{\boldsymbol{\delta}}^{(r)}_{i,j}(l) - \hat{\boldsymbol{\delta}}_{i,j}(l) \right\}$, and where $\hat{\boldsymbol{\delta}}^{(r)}_{i,j}(l)$ is constructed as the difference between the empirical moments calculated using $r = 1, \dots, n$ sub-samples each of size m from the sets $\boldsymbol{\zeta}^*_{S,l}$ and $\boldsymbol{\zeta}_{S,l}$.

4. Tracing the effects of disasters with a QVAR

In this section, we use the QVAR methodology laid out in the previous section to quantify the macroeconomic effects of natural disasters in the U.S. We first discuss the properties of the Costly Disaster index which we use for our baseline analysis. We then discuss the model specification and provide evidence on the fit of the conditional quantile function. We next document the in-sample fit of the model for a specific large disaster event and then discuss the quantile and moment impulse responses of disaster shocks more generally. Finally, we discuss our results with a particular focus on explanations why disasters that hit only parts of the U.S. have had nation-wide effects.

4.1. A first look at the data

In this section we take a first look at the data to study the properties of the economic costs of natural disasters in the United States. We measure natural disasters using the costly disaster (CD) index constructed by Ludvigson et al. (2020). This index captures financial losses due to natural disasters based on insurance data for an exhaustive list of disasters in the U.S. The index aggregates information on the costs of disasters from several sources, including the National Oceanic and Atmospheric Administration (NOAA) and the Insurance Information Institute (III), for more details see Ludvigson et al. (2020). Since our focus is on natural disasters only, we restrict our sample to the period from January 1980 through December 2019 and thus exclude the coronavirus epidemic. We also drop the terrorist attacks of 9/11 from our sample.

Fig. 1 plots the financial losses related to natural disasters in the U.S. from 1980 through 2019. The largest three disasters in this sample period were all hurricanes: *Katrina* in August 2005, *Sandy* in October 2012 and *Harvey/Irma/Maria* which occurred in short succession in August and September of 2017. Those events resulted in major damage and loss of lives, and the U.S. economy incurred large financial losses as captured by the CD series. Hurricane Katrina, for example, was associated with staggering economic costs of more than 160 billion in terms of 2019 US dollars. Another observation that follows from Fig. 1 is that very costly disasters occur frequently, particularly in the latter part of the sample. This is in line with evidence that the observed increase in global mean temperatures has raised the probability of natural disasters such as hurricanes and floods in recent decades Masson-Delmotte et al. (2018).

To illustrate that large disasters occur frequently, Fig. 2 shows a Q-Q plot of the costly disaster series. This allows us to compare the quantiles of the costly disaster series with those of a standard normal distribution. If the CD series was normally distributed, the blue dots would lie on the red line. Instead, we see a sharp deviation. This implies that the CD series has a heavy right tail. In other words, very costly disasters occur relatively often. Heavy tails imply that standard econometric tools which estimate conditional mean dynamics may not be well suited to study the full impact of disasters on the macroeconomy.

As an alternative, we propose to use the quantile vector autoregression (QVAR) machinery summarized in Section 3. This allows us to estimate the propagation of disaster shocks to the *entire* conditional distribution of the included endogenous variables. We can thus estimate the tail effects of natural disasters on key macroeconomic aggregates.

To illustrate how a large natural disaster may affect aggregate economic outcomes, Fig. 3 plots the monthly CD series along with real Industrial Production (IP) growth and CPI inflation around Hurricane Katrina. Immediately after the event, IP growth dropped sharply by about two percent and inflation surged by almost 1.5 percent. These initial reactions were followed by a sharp rebound of growth and a decline of inflation in subsequent months. The opposite responses of output and inflation immediately after the disaster suggest that it initially acted as a supply shock to the U.S. economy. This is in line with the notion that natural disasters may lead to disruptions in supply chains and thus result in a temporary shortage of goods and services which drives up prices. The

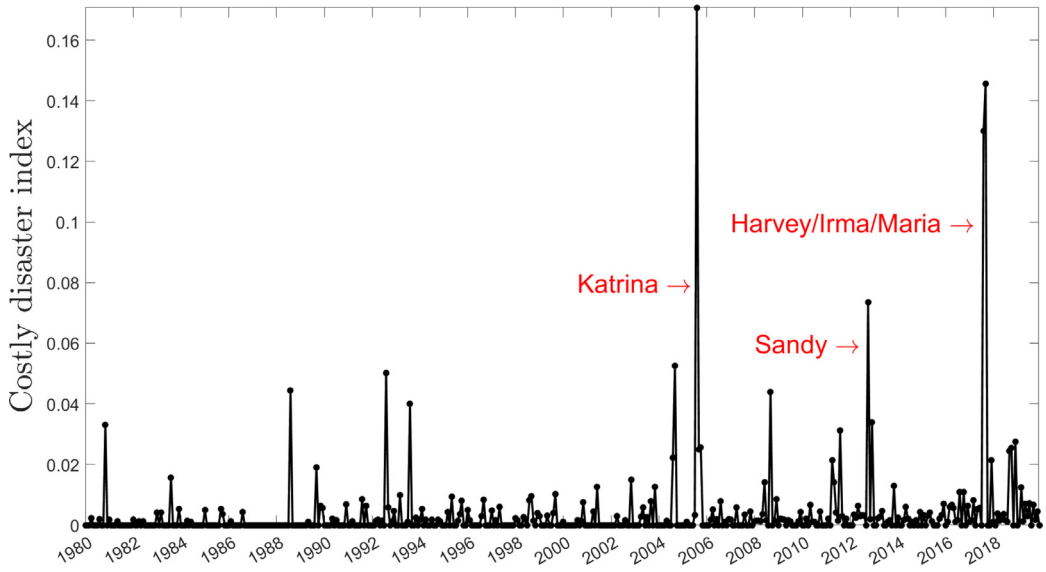

Fig. 1. Costly disaster series (CD). This figure plots financial losses in trillions of 2019 US dollars caused by natural disaster events in the U.S over the period from January 1980 through December 2019. The index is obtained from Ludvigson et al. (2020) and the terrorist attacks of September 2001 are dropped from the sample.

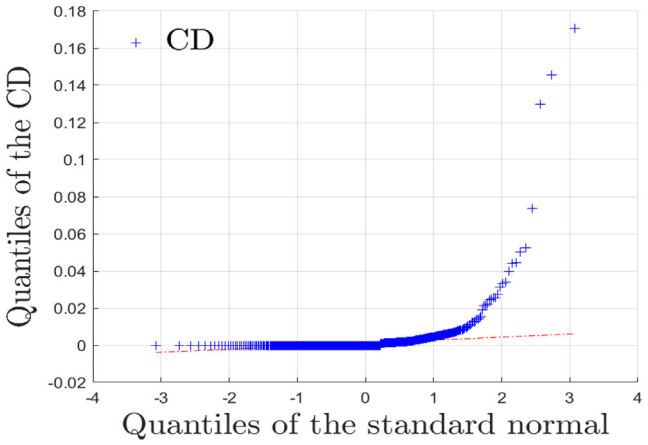

Fig. 2. QQ-Plot of the costly disaster series. This figure compares the quantiles of the costly disaster series with the quantiles of a standard normal distribution.

subsequent rebound of economic activity against the backdrop of continued inflation pressures, in turn, suggests that disasters also affect aggregate demand several months after the initial impact.

The identification of the effects of disaster shocks in Ludvigson et al. (2020) relies on the assumption that natural disasters are genuinely exogenous events. Specifically, they adopt a recursive ordering in their VAR with the CD series ordered first, implying that disasters affect the other variables contemporaneously but the opposite is not true. In our analysis, we adopt the same identification scheme. In contrast to Ludvigson et al. (2020), however, and motivated by the heavy-tailedness of the disaster distribution, we employ a model of the U.S. economy which explicitly allows for disasters to affect the quantiles of the endogenous variables differently.

4.2. Model specification

To trace the economic effects of disasters in the QVAR framework, we rely on the same set of indicators as Ludvigson et al. (2020). In addition to the CD index, which is ordered first, these are initial claims for unemployment insurance, monthly CPI inflation,

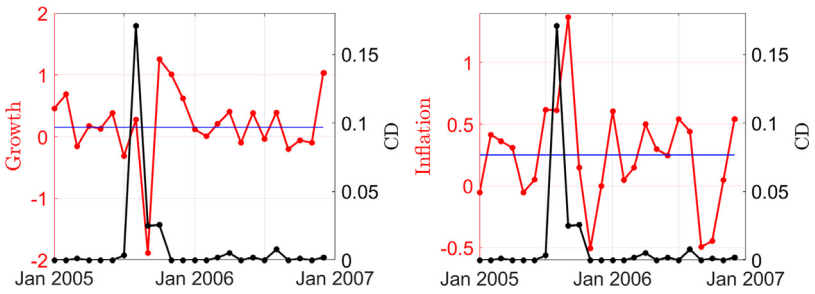

Fig. 3. Macroeconomic indicators and the cost of Hurricane Katrina. This figure shows real industrial production growth and CPI inflation along with the CD index around Hurricane Katrina. The black line represents the CD index, while the red lines correspond to the real IP growth and inflation. The horizontal lines correspond to the sample mean of real IP growth and inflation, respectively.

monthly real industrial production growth, the federal funds rate, and the measure of macroeconomic uncertainty constructed by Jurado et al. (2015).[11]

We use six lags for the QVAR model. This is in line with Ludvigson et al. (2020) and is a common choice in the empirical macroeconomics literature using monthly data.[12] The model selection for regression quantiles, particularly in the context of multiple variables and quantiles, requires additional research. To the best of our knowledge, there is no unified information criterion across all equations and quantiles that we might utilize in our context. We discuss available criteria for a single quantile and equation in the Online Appendix. Before we use the QVAR to analyze the distributional effects of natural disasters, we first ensure that the model is stationary. The results are also given in the Online Appendix.

We next provide an in-sample specification test to evaluate the fit of the regression quantiles. As discussed by Engle and Manganelli (2004), if $q_{X_{i,t}}(\theta_i, \boldsymbol{Z}_{i,t})$ in Eq. (5) is the true conditional quantile function, then Eq. (6) must hold. Equivalently, the hit sequence $\{I\left(X_{i,t} < q_{X_{i,t}}(\theta_i, \boldsymbol{Z}_{i,t})\right) - \theta_i\}$, where $I(\cdot)$ is an indicator function, must be an i.i.d. mean zero process for any $\theta_i \in (0,1)$. Several alternative tests have been proposed to evaluate this condition (Kupiec, 1995; Christoffersen, 1998; Berkowitz et al., 2009).[13]

Here we adopt the test by Engle and Manganelli (2004). Let $\boldsymbol{\Psi}_{i,t} \in \boldsymbol{I}_{i,t}$ be a vector of past observations and $\{h_{i,t}(\theta_i) = I\left(X_{i,t} < q_{X_{i,t}}(\theta_i, \boldsymbol{Z}_{i,t})\right) - \theta_i\}$ the hit sequence. Under the null hypothesis, $h_{i,t}(\theta_i)$ has zero mean and is uncorrelated with past information. Under the alternative, the quantile function in Eq. (5) is a poor approximation of the unknown true conditional quantile function (Gaglianone et al., 2011). Put differently, under the null the parameter vector $\boldsymbol{\pi}_i(\theta_i) = [v_i(\theta_i), \boldsymbol{\gamma}_i(\theta_i)']'$ is zero in the following auxiliary regression

$$h_{i,t}(\theta_i) = v_i(\theta_i) + \boldsymbol{\Psi}'_{i,t}\boldsymbol{\gamma}_i(\theta_i) + \xi_{i,t}(\theta_i), \quad t = p+1, \dots, T, \tag{11}$$

where $\xi_{i,t}(\theta_i)$ is a zero mean i.i.d. residual. We follow Engle and Sheppard (2001) and estimate Eq. (11) via least squares. This yields $\hat{\boldsymbol{\pi}}_i(\theta_i)$ with asymptotic covariance matrix $\hat{\boldsymbol{V}}_i(\theta_i)$. The test statistic is then given by

$$\hat{\boldsymbol{\pi}}_i(\theta_i)'\hat{\boldsymbol{V}}_i(\theta_i)^{-1}\hat{\boldsymbol{\pi}}_i(\theta_i) \sim \chi^2_{\dim(\boldsymbol{\pi}_i(\theta_i))}, \tag{12}$$

where χ^2_τ is a chi-squared distribution with $\tau > 0$ degrees of freedom.[14] The implementation of the test requires a choice of the vector $\boldsymbol{\Psi}_{i,t}$. We use $\boldsymbol{Z}_{i,t}$, which contains all p lags of the endogenous variables $\boldsymbol{X}_t$.[15]

We rely on our Bootstrap Algorithm 2 to compute the test statistic (12). In particular, we re-estimate model (5) for each bootstrap sample, then construct $\hat{h}^{(b)}_{i,t}$, and finally run the least squares regression (11). We calculate $\hat{\boldsymbol{V}}_i(\theta_i)$ as the empirical variance–covariance of the estimated parameters across bootstrap samples $\left\{\hat{\boldsymbol{\pi}}^{(1)}_i(\theta_i), \dots, \hat{\boldsymbol{\pi}}^{(B)}_i(\theta_i)\right\}$. We perform $B = 1000$ bootstrap repetitions and use the grid of quantile indices $[0.05, 0.06, \dots, 0.95]$ as defined in (8). This choice of grid is standard in the literature, e.g. Koenker (2005, Chapter 4.3) and Chernozhukov et al. (2017). The reason is that even with monthly data and a relatively large number of observations as in our case, extreme right or left tail quantiles may not be estimated very precisely.

[11] We thank Sydney Ludvigson for kindly sharing the CD series with us. The uncertainty index is available at: https://www.sydneyludvigson.com/data-and-appendixes. All remaining series have been obtained from FRED and are seasonally adjusted, where applicable. We compute monthly growth rates as log differences.

[12] For instance Kilian and Lütkepohl (2017, section 2.6) and Ivanov and Kilian (2005) conclude that the lag order must be large enough to capture responses of standard macroeconomic aggregates. Various empirical studies reviewed in Ivanov and Kilian (2005, Table 1) suggest at least a lag order of six for monthly data and two for quarterly data for impulse repose analysis. We provide robustness checks for lag ordersp = 3,9 and 12 in the Online Appendix. Our results remain intact with alternative lag orders. A similar strategy is followed by Bakkensen and Barrage (2018), see their Online Appendix.

[13] A detailed discussion of these tests and simulations to evaluate their relative performance in the dynamic regression quantile framework can be found in Gaglianone et al. (2011).

[14] This result derives directly from the properties of the regression quantile estimates $\hat{\boldsymbol{\beta}}_i(\theta_i)$ defined in Eq. (7) and discussed in Engle and Manganelli (2004), as well as the standard properties of quadratic forms discussed in White (2001, Theorem 4.31).

[15] Note that Gaglianone et al. (2011) suggest to use $\boldsymbol{\Psi}_{i,t} = \hat{q}_{X_{i,t}}(\theta_i, \boldsymbol{Z}_{i,t})$ instead. We follow their suggestion in a robustness check in the Online Appendix. The results are essentially unchanged.

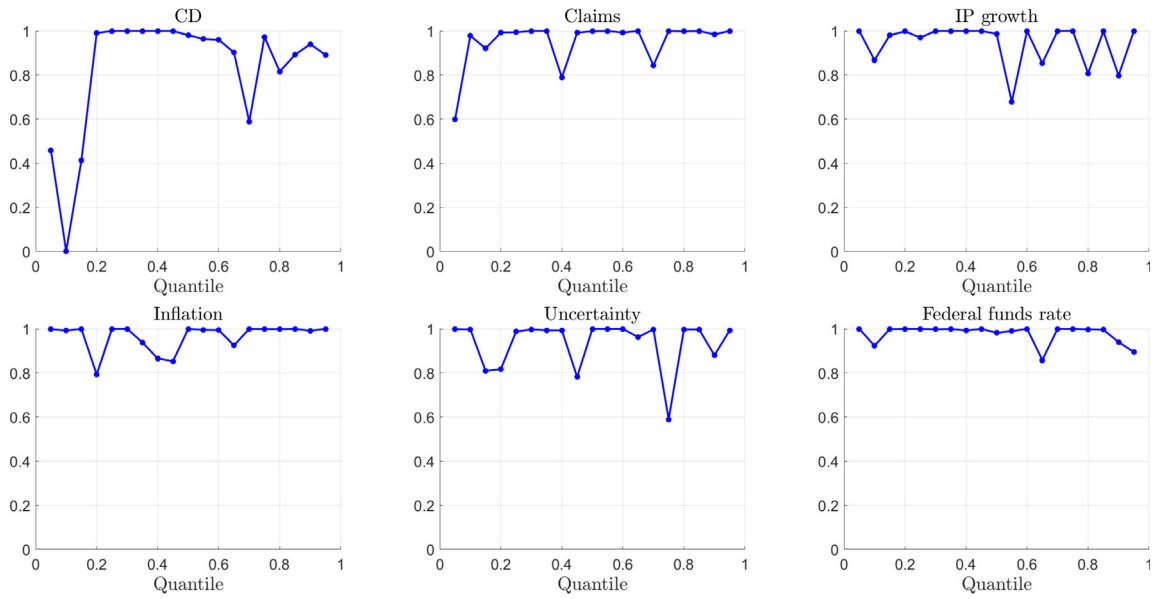

Fig. 4. P-values of the dynamic quantile test: This figure provides *p*-values for the conditional quantile test provided in Eq. (12) for the grid of quantiles $\theta_i \in [0.05, 0.1, \dots, 0.95]$. The test statistic is constructed using the bootstrap Algorithm 2.

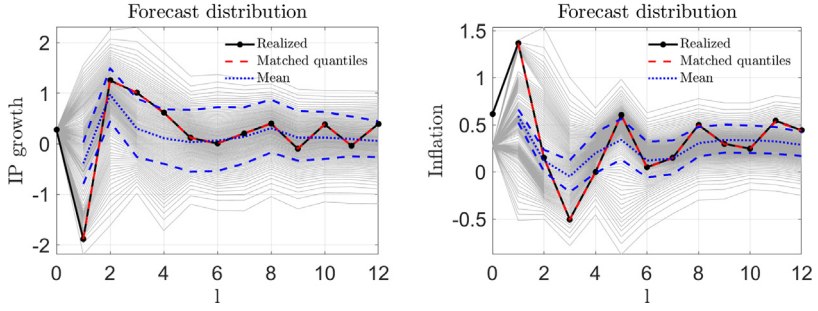

Fig. 5. Matching the realizations with the distribution forecasts. Distributional forecasts of IP growth and CPI inflation following Hurricane Katrina. The figure shows the evolution of the forecast distribution for IP growth and CPI inflation. The solid gray lines correspond to the quantiles $[0.01, 0.02, \dots, 0.99]$ of forecasts. Forecasts are generated by Algorithm 1 which uses the grid of quantiles $[0.25, 0.2578, \dots, 0.95]$ for the CD series and $[0.05, 0.06, \dots, 0.95]$ for the remaining variables.

The results are provided in Fig. 4. They show that the conditional quantile restriction of Eq. (6) cannot be rejected for almost all quantiles and all variables. The exception is the left tail of the CD index which captures the disasters with the lowest economic costs. The empirical density of the CD index is concentrated around zero, implying that while the series has a heavy right tail, it has essentially no far left tail. Hence, to model the dynamics of disaster series, it is not necessary to consider the far left quantiles.[16] In the following analysis, we thus truncate the grid of quantile indices for the CD index at $[0.25, 0.2578, \dots, 0.95]$, whereas for the remaining variables we use the standard grid $[0.05, 0.06 \dots, 0.95]$. As a result, both grids have an equal number of quantiles.

4.3. Evaluating the quantile function in-sample

We can evaluate the ability of our model to capture the observed dynamics of macroeconomic aggregates in the aftermath of large disasters. To this end, we estimate the model and then use Algorithm 1 to recursively construct in-sample forecasts $\zeta_{S,l} = \left\{ \boldsymbol{X}_{t+l}^{(s)}, s = 1, \dots, S \right\}$, $l \geq 1$ using data up to August 2005, the month when hurricane Katrina hit the Gulf coast. We use $S = 10{,}000$ samples. Given the forecasts at horizon l, we then calculate the empirical quantiles of the forecast distribution.

Fig. 5 displays these quantiles for IP growth and CPI inflation. We superimpose the conditional mean forecast, which we obtain using a standard linear VAR also estimated using data through August 2005. The associated 95% confidence interval is obtained using

[16] In unreported results, we see that the estimated conditional quantile function for the CD index for quantile indices below $\theta = 0.25$ are muted.

a standard residual-based bootstrap as in Gonçalves and Kilian (2004). Several results are noteworthy. First, industrial production growth initially dropped sharply and then strongly recovered while inflation rose markedly just after hurricane Katrina hit and then declined in subsequent months. Strikingly, these dynamics are not well captured by the conditional mean forecasts, shown as blue dotted lines. Instead, the realized values coincide with the lowest predicted quantiles for IP growth and the highest quantiles for inflation in the first month after the event. In subsequent months, the realizations are well matched by high predicted quantiles of IP growth and lower quantiles for inflation. A similar observation can be made for CPI inflation on the right-hand side of the figure. The realized inflation readings correspond to some of the highest (lowest) quantiles of the predictive distribution in the first few months after the disaster. These preliminary results suggest that large disasters such as hurricane Katrina imply macroeconomic dynamics that fall into the tails of the conditional forecast distribution.

4.4. Quantile impulse response analysis of natural disasters

In this section, we use the Quantile IRFs introduced in Section 3.5.1 to study the macroeconomic effects of natural disasters in the U.S. for the sample period from 1980 to 2019. We rely on the recursive identification scheme captured in Eq. (4) and discussed in Section 3. We scale the shock to $\kappa_1 = 9 \times \sigma_{CD}$, where σ_{CD} is the empirical standard deviation of the CD index ordered first in our QVAR. This roughly corresponds to the economic cost associated with hurricane Katrina. Our emphasis is thus on how the tails of the forecast distribution evolve in response to very large disasters.

To assess the statistical significance of the shifts in the forecast distribution, we plot the quantile impulse response functions $\hat{\boldsymbol{\Delta}}_{i,1}(\theta, l)$ for the 1st, 5th, 50th, 95th and 99th quantile in Fig. 6. We construct the error bands using subsampling Algorithm 3 with $n = 10{,}000$ repetitions. There are several noteworthy findings. First, while the bottom quantiles of the IP growth distribution (top row) are significantly compressed immediately after the shock, they quickly revert back to their unconditional level. In contrast, both the median and the top quantiles also initially drop but then strongly and significantly overshoot their initial level in the two to six months after the disaster. Hence, while the conditional forecast distribution of IP growth is meaningfully shifted to the left immediately following large disasters, there is a sizeable rebound effect of economic activity which subsequently shifts the growth distribution to the right.

For inflation, the picture looks somewhat different. The bottom quantiles and the median significantly increase immediately after the shock but then quickly recede. In contrast, the effect of the disaster shock is considerably more persistent in the top quantiles of the predictive inflation distribution, suggesting that large disasters come with meaningful upside risks to inflation in the two to six months after the shock.

It is instructive to compare these quantile IRFs with those for the conditional mean implied by a standard linear VAR. To this end, we estimate the linear VAR with the same number of lags $p = 6$ and for a shock of the same magnitude. Specifically, we calculate the following quantity $\boldsymbol{\Delta}_{i,1}(l) \equiv \mathrm{E}(X_{i,T+l} | X_{1,T+1} = 9\sigma_{CD}, \boldsymbol{X}'_T, \ldots, \boldsymbol{X}'_{T-p+1}) - \mathrm{E}(X_{i,T+l} | \boldsymbol{X}'_T, \ldots, \boldsymbol{X}'_{T-p+1})$ for $l = 1, \ldots, h$ and $i = 1, \ldots, K$. We construct confidence intervals using a standard residual-based bootstrap at the 1% significance level. Fig. 7 displays the IRFs for real IP growth and CPI inflation. Both show qualitatively similar but quantitatively much more modest dynamics of the conditional mean in response to a large disaster shock than the tail QIRFs provided in Fig. 6. We interpret this as evidence that the QVAR framework is better able to capture the full conditional distribution of macroeconomic dynamics following disaster shocks.

4.5. Moment impulse response analysis of natural disasters

The QIRFs reported in the previous section provide a granular account of the effects of natural disasters on the forecast distribution of output and inflation. A convenient way to summarize this information is to compute the resulting moments of the conditional forecast distributions with and without disaster shocks, as defined in Eq. (10). The results are provided in Fig. 8. We first focus on the moments of the forecast distribution of IP growth, shown in the left column. The conditional mean sharply contracts in the first month after the shock but then strongly rebounds and remains elevated for about six months. The conditional standard deviation also increases markedly and remains significantly elevated for about six months. The conditional skewness of the predicted IP growth distribution first responds negatively but then briefly moves into positive territory after about four months. The conditional kurtosis only significantly increases in the first month after the disaster shock. Combined, these charts reiterate that conditional on large disasters the forecast distribution of IP growth is initially shifted to the left before seeing a substantial rebound, while the volatility of forecasts is elevated for several months after the disaster.

Turning to the corresponding moments for inflation in the right column, we see that the conditional mean of the predicted inflation distribution is positive for about four months after the shock, while the conditional standard deviation initially drops and then remains elevated for several months. The conditional skewness of the forecast distribution for inflation is initially positive, but then turns negative after about half a year. Finally, as for IP growth the conditional kurtosis is only significantly higher in the first month after the shock. A key message that emerges from this exercise is that to fully capture the rich macroeconomic dynamics following natural disasters, one needs models that allow the entire distribution to shift, not only the mean.

Note that the unconditional mean, standard deviation, skewness and kurtosis of IP growth in our sample equal $0.15, 0.67, -1.16$, and 9.07, respectively. The same moments for CPI inflation are estimated to be 0.25, 0.29, -0.22, and 10.37. This allows us to gauge the economic significance of the moment IRFS above. We see a sharp drop of the conditional mean of IP growth by about -0.4% on impact, followed by a subsequent rebound of more than 0.5% when a larger disaster shock hits compared to a scenario without such a shock. These responses are thus considerably larger than the unconditional mean of IP growth. For CPI inflation, the 0.2% initial increase of the conditional mean subsequent to a disaster shock is also sizeable and of the same magnitude as its unconditional mean.

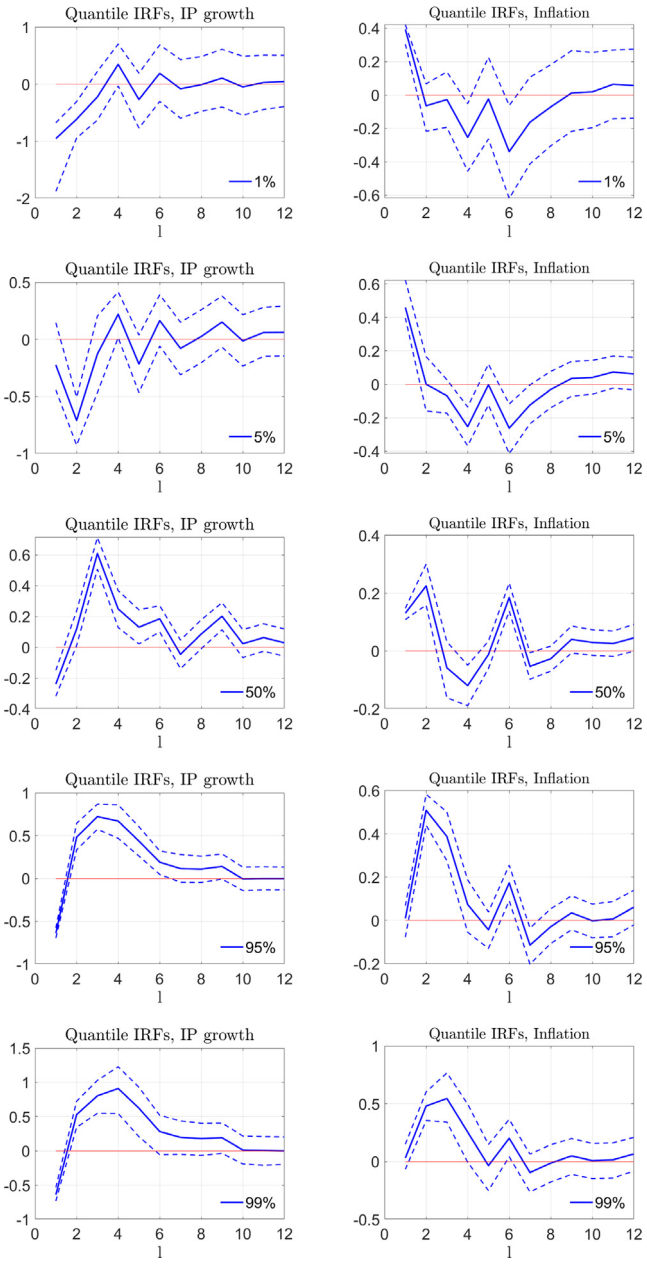

Fig. 6. QIRFs. Quantile impulse response functions for IP growth and CPI inflation as defined in Eq. (9) to a $9 \times \sigma_{CD}$ disaster shock. The point estimates are based on Algorithm 1, which uses a grid of quantile indices $[0.25, 0.2578, \dots, 0.95]$ for the CD index and $[0.05, 0.06, \dots, 0.95]$ for the remaining variables. Confidence intervals are calculated via subsampling Algorithm 3 at the 1% significance level.

4.6. Discussion

The findings presented thus far suggest that natural disasters have a relatively short-lived but sizeable impact on the aggregate U.S. economy. In particular, large disasters meaningfully shift the conditional distribution of IP growth and CPI inflation. This may be surprising in light of the fact that many natural disasters occur locally. That said, as discussed in Section 2, disasters may have aggregate economic effects if they lead to a disruption of supply chains or hit sectors that are central to overall economic activity through input–output linkages.

The largest disasters in our sample are hurricanes. Some of these made landfall in regions that are home to some important industries. For example, according to the National Hurricane Center, hurricane Katrina had a major impact on Florida, Louisiana,

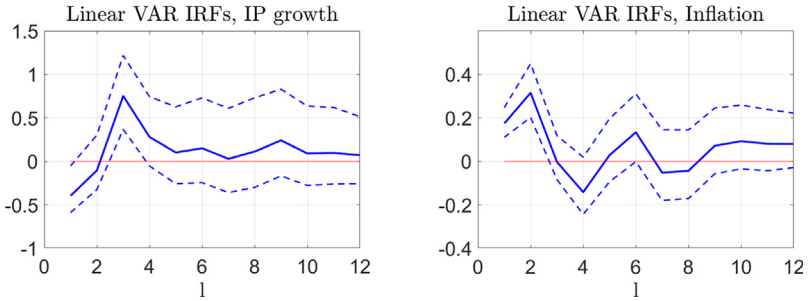

Fig. 7. Linear VAR impulse response functions: Linear VAR impulse response functions for IP growth and CPI inflation to a $9\times\sigma_{CD}$ disaster shock as defined in Section 4.4. Confidence intervals are calculated with the residual-based bootstrap at the 1% significance level.

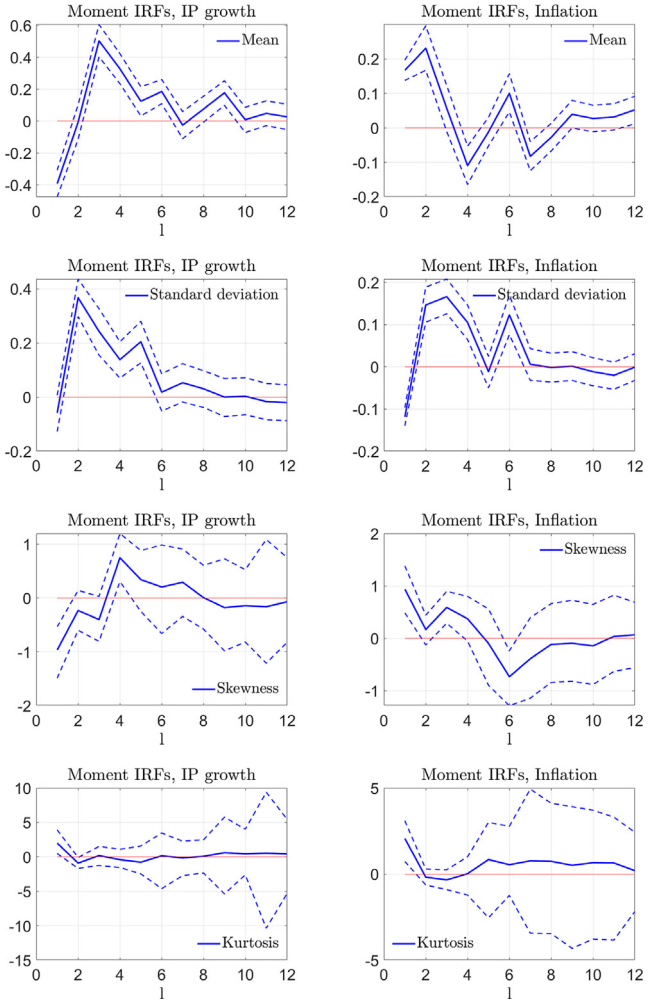

Fig. 8. Moment impulse response functions: Moment impulse response functions for IP growth and CPI inflation to a $9\times\sigma_{CD}$ disaster shock, as defined in Eq. (10). Calculations are based on Algorithm 1, which uses a grid of quantile indices $[0.25, 0.2578, \ldots, 0.95]$ for the CD index and $[0.05, 0.06, \ldots, 0.95]$ for the remaining variables, see the discussion in Section 3.5.2. Confidence intervals are calculated via subsampling Algorithm 3 at the 1% significance level.

Mississippi, and Alabama.[17] As such, it significantly affected industrial production in the Gulf Coast region which is home to important companies active in oil and gas extraction, industrial chemicals manufacturing, and petroleum refining.[18] Moreover, the hurricane was reported to have given rise to significantly higher natural gas prices due to its effects on gas production.[19]

Other major hurricanes primarily affected different U.S. regions of national economic importance. For example, Hurricane Sandy severely impacted production in the Northeast in October 2012, resulting in a 0.4% decline of aggregate industrial production.[20] The main effects were felt in the utilities, food, chemicals, computers and electronic products and transportation equipment sectors. Finally, Hurricane Harvey, which made landfall on the Gulf Coast of Texas in August 2017, led to a decline of national industrial production of 0.9% in that month as this region is home to important oil and gas extraction, petroleum refining, and petrochemical manufacturing industries (Bayard et al., 2017).[21]

5. Scenario analysis

There is unequivocal evidence that global mean temperatures have been increasing as a result of heightened concentration of greenhouse gases in the atmosphere. There is also widespread agreement among climate scientists that the frequency and intensity of extreme weather events will increase in large parts of the world as global mean temperatures continue to rise (IPCC, 2023). In this section, we use our QVAR framework to assess the potential macroeconomic outcomes associated with scenarios based on policies which have differential implications for the path of global carbon emissions and thus the carbon concentration in the atmosphere.

We take the following approach. First, we directly model the historical relationship between the atmospheric carbon concentration and natural disasters in the U.S. Second, we simulate the evolution of disasters over the next twenty years based on carbon concentration levels under two policy scenarios entertained by the Network for the Greening of the Financial System (NGFS) and study the implications for the conditional forecast distributions of IP growth and inflation. We now explain each of these steps in turn.

5.1. Modeling the frequency and intensity of disasters

Climate scientists often rely on Integrated assessment models (IAMs) to simulate complex interactions between the economy, the climate, and the environment, see the discussion in Section 2 and in the Online Appendix. Modeling these complex interactions is beyond the scope of our paper. Instead, we take a simple but, in our view, plausible shortcut. Specifically, we use the Gamma-zero distribution introduced in Monfort et al. (2017) to model the intensity of natural disasters as a function of the atmospheric carbon concentration. In particular we assume

$$CD_t|E_t \sim \gamma_0(\lambda_t, \mu),$$

where $\lambda_t > 0$ for any t, $\mu > 0$, and E_t is the carbon concentration. We measure the global atmospheric carbon concentration using the NOAA Global Monitoring Laboratory data.[22]

The Gamma-zero distribution is a combination of a Gamma and a Poisson distribution. Specifically, let $CD_t|E_t \sim \gamma(\nu_t, \mu)$ be a Gamma distributed random variable with shape ν_t and scale μ. Further assume that the shape ν_t, which governs the number of events modeled with a Gamma distribution, itself follows a Poisson distribution

$$\nu_t \sim \mathcal{P}(\lambda_t).$$

As shown by Monfort et al. (2017), these assumptions imply that the conditional first and second moment of the disaster distribution are given by

$$\mathrm{E}\left(CD_t|E_t\right) = \mu\lambda_t,$$

and

$$\mathrm{Var}\left(CD_t|E_t\right) = 2\mu^2\lambda_t.$$

To ensure that the conditional mean and variance of the costs of disaster distribution are both non-negative at all times, we model λ_t as a time-varying exponential function of the carbon concentration:

$$\lambda_t = \exp(\alpha + \beta E_t).$$

[17] See https://www.aoml.noaa.gov/hrd/hurdat/Data_Storm.html.

[18] The Federal Reserve documented the effects of hurricane Katrina on these industries in September (https://www.federalreserve.gov/releases/g17/20050914/) and October 2005 (https://www.federalreserve.gov/releases/g17/20051014/).

[19] See https://www.minneapolisfed.org/article/2005/gasoline-prices-climb-in-response-to-hurricanes.

[20] See the Federal Reserve statistical release of November 16, 2012: https://www.federalreserve.gov/releases/g17/20121116/.

[21] See also the Federal Reserve statistical release in September 2017: https://www.federalreserve.gov/releases/g17/20170915/.

[22] These data are available at https://doi.org/10.15138/9N0H-ZH07. They are reported in Micromols of carbon per mol of air.

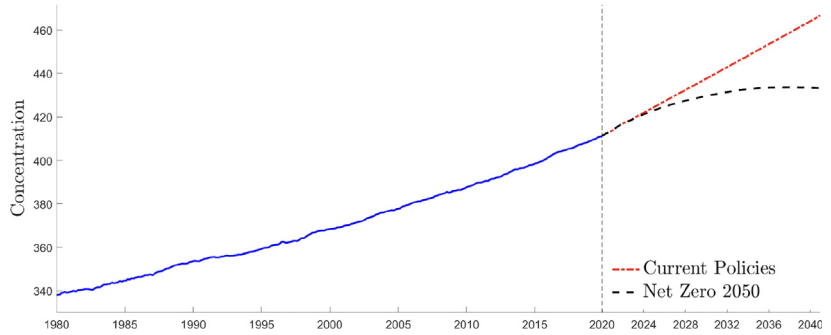

Fig. 9. Atmospheric carbon concentration. The plot shows the global average atmospheric carbon concentration observed since 1980 and projected under the two scenarios "Current Policies" and "Net Zero 2050". The carbon concentration data are measured as micromols of CO2 per mol of air and are obtained from the NOAA Global Monitoring Laboratory. The scenarios are provided by the Network of Central Banks and Supervisors for Greening the Financial System (NGFS).

As a result, the mean and variance of disasters are time-varying and depend on the evolution of the carbon concentration.[23] The two conditional moments can thus be used to estimate the parameters (α, β, μ). Specifically, we minimize the following sum of squared deviations (SSD):[24]

$$\sum_{t=1}^{T} \left(CD_t - \mu\lambda_t\right)^2 + \sum_{t=1}^{T} \left([CD_t - \mu\lambda_t]^2 - 2\mu^2\lambda_t\right)^2.$$

We obtain the following parameter estimates, with heteroskedasticity robust standard errors in parentheses: $\hat{\mu} = 0.028$ (0.007), $\hat{\alpha} = -12.14$ (3.83) and $\hat{\beta} = 0.027$ (0.01) and SSD=0.096. To check that these estimated parameters capture the moments of the observed disaster distribution well, we use them to simulate the CD series 1000 times and then compare the mean and the tails of the observed and the simulated CD index. The mean and 90th, 95th, and 99th quantiles of the simulated data are given by $0.004, 0.007, 0.029$, and 0.080, respectively. The corresponding sample counterparts in the observed CD index take on the values $0.004, 0.007, 0.015$, and 0.074, respectively. Hence, the model represents the properties of the actual CD distribution quite well.

5.2. Simulating natural disasters

Using these estimated parameters, we next simulate hypothetical future realizations of the CD index conditional on different paths for the global carbon concentration. Specifically, we rely on two key scenarios for the atmospheric carbon concentration that have been studied by the NGFS, "Current Policies" and "Net Zero 2050".[25] The Current Policies scenario assumes that only currently implemented policies are preserved, resulting in no substantial decreases of carbon emissions. The Net Zero 2050 scenario, in contrast, is associated with stringent climate policies and innovation so that global CO2 emissions reach net zero around 2050 and global warming is limited to 1.5 °C. Fig. 9 plots the evolution of atmospheric carbon concentration since 1980 and appends the assumed evolution under the two scenarios.[26] As one can see, the carbon concentration continues to rise sharply under the Current Policies Scenario, but starts to decline in the 2030s under the Net Zero 2050 scenario.

Given our model for the disaster intensity discussed above, these assumed carbon concentration pathways give rise to different future disaster distributions. Fig. 10 displays a representative draw extending the observed sample of CD realizations by twenty years of simulated outcomes for the CD index. The chart shows that the frequency and magnitude of disasters is substantially higher under the Current Policies scenario than under the Net Zero 2050 scenario.

To study the impact of natural disasters on the conditional forecast distributions of IP growth and CPI inflation under the two scenarios, we perform the following exercise. First, for each scenario we use the projected evolution of the atmospheric carbon concentration $E_{T+1}, \ldots, E_{T+h}$ over a twenty-year horizon ($h = 240$) starting in January 2020. We use the estimated parameters $\hat{\alpha}, \hat{\beta}, \hat{\mu}$ to simulate $S = 1000$ draws $\{CD_{T+1}^{(s)}, \ldots, CD_{T+h}^{(s)}, s = 1, \ldots, S\}$ from the Gamma-zero distribution which characterizes disasters conditional on the carbon concentration. For each draw s, we annualize the simulated CD series and compute densities for the forecast period from 2020 through the end of 2039. We then average these densities across the 1000 draws.

[23] Note that we abstract from climate adaptation policies which in practice may affect the relationship between the carbon concentration and the economic costs of natural disasters.

[24] To verify the validity of our estimation method, we first performed a Monte Carlo exercise using the following parameters: $T = 480$, $E_t = 334 + 0.15t + 2\varepsilon_{E_t}$, $\varepsilon_{E_t} \sim N(0, 1)$, $\mu = 0.2$, $\alpha = -12$, and $\beta = 0.05$. With this data generating process we reestimate the parameters 1000 times. The empirical means and standard deviation of the Monte Carlo estimates are $\hat{\alpha} = -11.963(.222)$, $\hat{\beta} = .05(.0004)$, and $\hat{\mu} = .198(.025)$. Hence, our estimation approach is well suited to estimate the parameters of the Gamma-zero distribution.

[25] All scenarios considered by the NGFS are summarized here: https://www.ngfs.net/ngfs-scenarios-portal/. These scenarios are based on the most current vintage of the integrated assessment model REMIND-MAgPIE 3.0–4.4 developed by the Potsdam Institute for Climate Impact Research. The model documentation is provided here: https://www.iamcdocumentation.eu/index.php/Model_Documentation_-_REMIND-MAgPIE.

[26] We thank Oliver Richters for sharing these data with us.

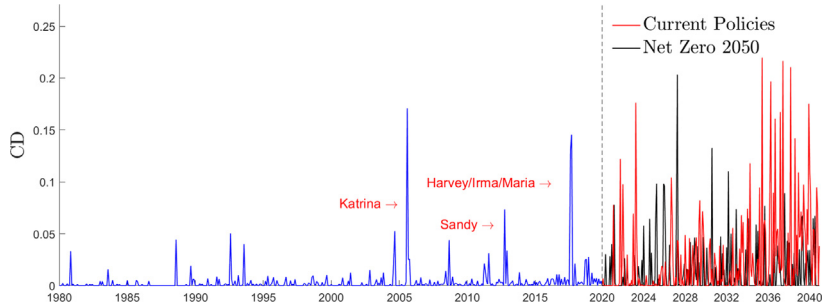

Fig. 10. Representative Draws of the Simulated CD Series. The plot shows the observed evolution of the CD index from 1980 through 2019, extended by representative draws under the two scenarios "Current Policies" and "Net Zero 2050". The cost of disaster series follows a Gamma-zero distribution which depends on the global average carbon concentration, see Section 5.1. The scenarios are provided by the Network of Central Banks and Supervisors for Greening the Financial System (NGFS).

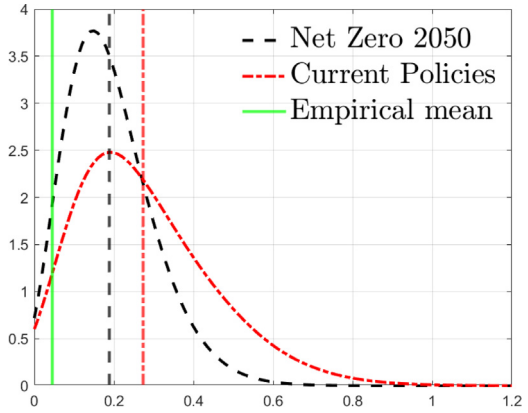

Fig. 11. Kernel Density Estimates of Annualized Costs of Disasters under Two NGFS Scenarios. This figures shows average kernel density estimates of the annualized CD index conditional on the two NGFS scenarios "Current Policies" and "Net-Zero 2050". We superimpose the average of the annualized CD realizations for the period from 1980 through 2019.

The result is shown in Fig. 11. We superimpose the corresponding means as well as the average of the CD index over the sample from 1980 through 2019 as vertical lines. The chart suggests that under both scenarios, the distribution of disaster-related economic costs will be much more heavily tilted to the upside compared to the historically observed average disaster costs. Moreover, the Current Policies scenario generates a considerably heavier right tail for the cost of disasters than the Net-zero 2050 scenario.

Given the simulated costs of disasters over the next twenty years, in the second step we generate conditional distributions for the macroeconomic aggregates of interest. Specifically, for each draw $\left\{CD_{T+1}^{(s)},\dots,CD_{T+240}^{(s)}\right\}$, we use Algorithm 1 to generate $\tilde{S}=1000$ new draws for each remaining endogenous variable in our QVAR for each of the $l=1,\dots,240$ horizons. Based on these simulated paths, we then compute the conditional moments for each horizon $T+l$ across draws $\left\{IP_{T+l}^{(\tilde{s})}\right\}$ and $\left\{CPI_{T+l}^{(\tilde{s})}\right\}$, where $\tilde{s}=1,\dots,1000$. We calculate kernel densities of these conditional moments across forecast horizons $T+l$ for $l=1,\dots,240$. We repeat this exercise for each draw $s=1,\dots,S$ of the simulated CD index. In the end, we average the $S=1000$ kernel densities for each moment. The purpose of this exercise is to trace the distribution of the moments of IP growth and CPI inflation conditional on disaster paths under the two scenarios, taking into account the randomness of the disaster outcomes.

It is instructive to compare these simulated conditional moments to their historical counterparts. To this end, we compute the conditional moments of IP growth and CPI inflation over the 20-year period from 2000 through 2019. Specifically, we use Algorithm 1 to draw $\tilde{S}=1000$ paths for the remaining endogenous variables in our QVAR for each month in the sample period January 2000 through December 2019. Then, we calculate the first four moments for IP growth and CPI inflation and compute their kernel densities for each of the 240 months. We superimpose these densities on the simulated densities of conditional moments under the two scenarios.

Fig. 12 provides the conditional moments of the counterfactual distributions along with their historical counterparts. The conditional mean of the IP growth distribution only minimally differs across the two scenarios. However, both are markedly shifted to the left compared to the last twenty years. The conditional standard deviation of IP growth is substantially tilted upwards under both scenarios compared to historical outcomes, but more strongly so under Current Policies than under the Net-Zero 2050 scenario. The skewness of the IP growth distribution also has substantially more mass at negative values under both scenarios relative to the past 20 years. The kurtosis, in turn, has a somewhat longer right tail under the two scenarios. Combined, these results suggest that

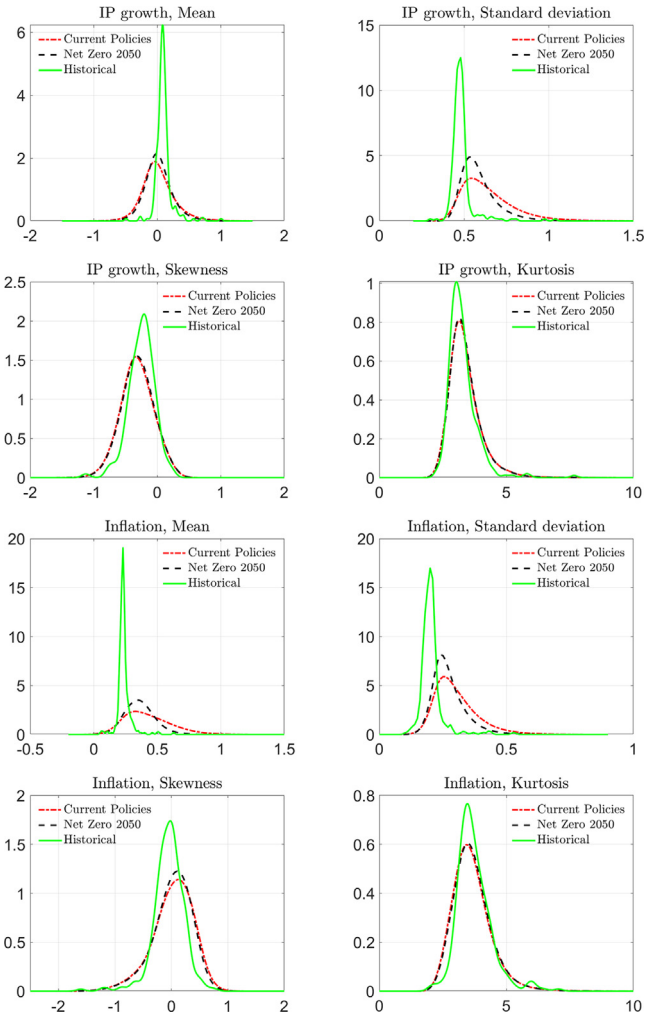

Fig. 12. Kernel Density Estimates of the Conditional Moments of Counterfactual Distributions under two NGFS Scenarios. This figure shows averages of the kernel density estimates of the moments for IP growth (top panel) and CPI inflation (bottom panel) under two NGFS scenarios: "Current Policies" and "Net-Zero 2050". The scenarios and procedure are described in Section 5.2.

the conditional distribution of real output growth will be considerably more tilted towards negative values and be more volatile in the future compared to recent experiences. That said, climate policies achieving net zero global carbon emissions by 2050 would considerably mitigate these distributional shifts.

Turning to CPI inflation in the bottom four panels of Fig. 12, we observe that both the conditional mean and the conditional volatility have substantially heavier right tails under the two scenarios as compared to the historical evidence. Under Current Policies, however, these distributional shifts are again more sizeable. The distributions of the conditional skewness and kurtosis of CPI inflation are not meaningfully different under the two scenarios, but feature somewhat longer right and left tails than the historical distribution.

In sum, our evidence suggests that the increased frequency and severity of natural disasters imply substantially more mass on lower growth and higher inflation outcomes and considerably higher macroeconomic volatility. This is particularly true when carbon emissions continue to rise as in the Current Policies scenario.

6. Conclusion

This paper has analyzed the distributional impact of disasters on real economic activity and inflation in a Quantile VAR framework. Natural disasters are followed by significant shifts in the forecast distribution, especially in the tails. Initially, there is substantial downside risk to real activity, but this is followed by a sizeable rebound. For inflation, upside risks dominate. Disaster shocks result in persistent increases of the conditional volatility and skewness of output and inflation. Our findings highlight that natural disasters significantly increase macroeconomic vulnerability, making it more difficult for policy makers to stabilize output

and inflation in the presence of frequent large natural disasters. Our results also show that to fully capture the economic effects of natural disasters requires a complete account of the conditional distribution of key macroeconomic aggregates in response to physical climate risk events. Arguably, this should have important implications for calculations of the social cost of carbon and the calibration of climate policies. Both depend on the estimated damages from extreme weather events associated with climate change. While it is beyond the scope of our paper to incorporate tail risks in such calculations, it appears to be an important avenue for future research.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jeconom.2024.105914.

References

Acemoglu, D., Carvalho, V., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. Econometrica 80 (5), 1977–2016.

Albala-Bertrand, J.M., 1993. Natural disaster situations and growth: a macroeconomic model for sudden disaster impacts. World Dev. 21 (9), 1417–1434.

Bakkensen, L., Barrage, L., 2018. Climate Shocks, Cyclones, and Economic Growth: Bridging the Micro-Macro Gap. NBER Working Paper 24893, National Bureau of Economic Research, Inc.

Bayard, K., Decker, R., Gilbert, C., 2017. Natural disasters and the measurement of industrial production: hurricane harvey, a case study. In: FEDS Notes. Board of Governors of the Federal Reserve System, Washington, http://dx.doi.org/10.17016/2380-7172.2086.

Berkowitz, J., Christoffersen, P., Pelletier, D., 2009. Evaluating value at-risk models with desk-level data. Manage. Sci. 57 (12), 2213–2227.

Botzen, W.J.W., Deschenes, O., Sanders, M., 2019. The economic impacts of natural disasters: a review of models and empirical studies. Rev. Environ. Econ. Policy 13 (2), 167–188.

Carrera, L., Standardi, G., Bosello, F., Mysiak, J., 2015. Assessing direct and indirect economic impacts of a flood event through the integration of spatial and computable general equilibrium modelling. Environ. Modell. Softw. 63, 109–122.

Cavallo, E., Galiani, S., Noy, I., Pantano, J., 2013. Catastrophic natural disasters and economic growth. Rev. Econ. Stat. 95 (5), 1549–1561.

Chavleishvili, S., Manganelli, S., 2024. Forecasting and stress testing with quantile vector autoregression. J. Appl. Econometrics 39 (1), 66–85.

Chernozhukov, V., 2002. Inference on the quantile regression process, an alternative, Working Paper 02-12, MIT Department of Economics Working Paper Series.

Chernozhukov, V., Fernández-Val, I., 2005. Subsampling inference on quantile regression processes. Sankhyā 67, 253–276.

Chernozhukov, V., Fernández-Val, I., Galichon, A., 2010. Quantile and probability curves without crossing. Econometrica 78, 1093–1125.

Chernozhukov, V., Fernández-Val, I., Kaji, T., 2017. Extremal quantile regression. In: Chernozhukov, V., He, X., Koenker, R., Peng, L. (Eds.), Handbook of Quantile Regression. CRC Chapman-Hall.

Chernozhukov, V., Fernández-Val, I., Melly, B., 2013. Inference on counterfactual distributions. Econometrica 81 (6), 2205–2268.

Christoffersen, P.F., 1998. Evaluating interval forecasts. Internat. Econom. Rev. 39, 841–862.

Davis, R., Ng, S., 2022. Time series estimation of the dynamic effects of disaster-type shocks. J. Econometrics 235, 180–201.

Engle, R.F., Manganelli, S., 2004. CAViaR: Conditional autoregressive value at risk by regression quantile. J. Bus. Econom. Statist. 22, 367–381.

Engle, R., Sheppard, K., 2001. Theoretical and Empirical Properties of Dynamic Conditional Correlation Multivariate GARCH. NBER Working Paper 8554, National Bureau of Economic Research.

Faccia, D., Parker, M., Stracca, L., 2021. Too hot for stable prices? International evidence on climate change and inflation. ECB Working Paper No. 2626.

Gabaix, X., 2011. The granular origins of aggregate fluctuations. Econometrica 79 (3), 733–772.

Gaglianone, W.P., Lima, L.R., Linton, O., Smith, D.R., 2011. Evaluating value-at-risk models via quantile regression. J. Bus. Econom. Statist. 29 (1), 150–160.

Gonçalves, S., Kilian, L., 2004. Bootstrapping autoregressions with conditional heteroskedasticity of unknown form. J. Econometrics 123, 89–120.

Hall, P., 1992. The Bootstrap and Edgeworth Expansion. Springer, New York.

Hallegatte, S., 2008. An adaptive regional input–output model and its application to the assessment of the economic cost of Katrina. Risk Anal. 28, 779–799.

Hilaire, J., Bertram, C., 2020. The REMIND-MAgPIE Model and Scenarios for Transition Risk Analysis. A Report Prepared By PIK for the UNEP-FI Banking Pilot Project (Phase 2). Potsdam-Institut für Klimafolgenforschung, Potsdam, http://dx.doi.org/10.2312/pik.2020.006, 20 p.

Intergovernmental Panel on Climate Change (IPCC), 2023. Summary for policymakers. In: Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, pp. 3–32. http://dx.doi.org/10.1017/9781009157896.001.

Ivanov, V., Kilian, L., 2005. A practitioner's guide to lag order selection for VAR impulse response analysis. Stud. Nonlinear Dyn. Econom. 9, 1–36.

Jurado, K., Ludvigson, S.C., Ng, S., 2015. Measuring uncertainty. Am. Econ. Rev. 105 (3), 1177–1216.

Kilian, L., Lütkepohl, H., 2017. Structural Vector Autoregressive Analysis. Cambridge University Press, Cambridge, UK.

Kim, H.S., Matthes, C., Phan, T., 2022. Severe Weather and the Macroeconomy. Working Paper 21-14, Federal Reserve Bank of Richmond.

Koenker, R., 1994. Confidence intervals for regression quantiles. In: Mandl, P., Huskova, M. (Eds.), Asymptotic Statistics. Springer-Verlag, New York, pp. 349–359.

Koenker, R., 2005. Quantile Regression. Cambridge Univ. Press, Cambridge, UK.

Koenker, R., Bassett, G., 1978. Regression quantiles. Econometrica 46, 33–50.

Koenker, R., Park, B.J., 1996. An interior point algorithm for nonlinear quantile regression. J. Econometrics 71, 265–283.

Koenker, R., Xiao, Z., 2006. Quantile autoregression. J. Amer. Statist. Assoc. 101 (475), 980–1006.

Koop, G., Pesaran, M.H., Potter, S.M., 1996. Impulse response analysis in nonlinear multivariate models. J. Econometrics 74, 119–147.

Kupiec, P., 1995. Techniques for verifying the accuracy of risk measurement models. J. Deriv. 3, 73–84.

Loayza, N., Olaberría, E., Rigolini, J., Christiansen, L., 2012. Natural disasters and growth-going beyond the averages. World Dev. 40 (7), 1317–1336.

Ludvigson, S.C., Ma, S., Ng, S., 2020. Covid19 and the macroeconomic effects of costly disasters. NBER Working Paper No. 26987.

Lusardi, A., 1998. On the importance of the precautionary saving motive. Am. Econ. Rev. Pap. Proc. 88, 449–453.

Lütkepohl, H., 2005. New Introduction to Multiple Time Series Analysis. Springer, Berlin, Heidelberg.

Masson-Delmotte, V., Zhai, P., Pörtner, H., Roberts, D., Skea, J., Shukla, P., Pirani, A., Moufouma-Okia, W., Péan, C., Pidcock, R., 2018. An IPCC special report on the impacts of global warming of 1.5 C.

Mohaddes, K., Ryan, N.C., Ng, M., Pesaran, H., Raissi, M., Yang, J., 2022. Climate change and economic activity: evidence from U.S. states. CESIfo Working Paper No. 9542.

Monfort, A., Pegoraro, F., Renne, J.-P., Roussellet, G., 2017. Staying at zero with affine processes: An application to term structure modelling. J. Econometrics 201 (2), 348–366.

Nordhaus, W.D., 1993. Optimal greenhouse-gas reductions and tax policy in the "DICE" model. Amer. Econ. Rev. 83 (2), 313–317.

Nordhaus, W.D., 2011. The economics of tail events with an application to climate change. Rev. Environ. Econ. Policy 5 (2), 240–257.

Noy, I., 2009. The macroeconomic consequences of disasters. J. Dev. Econ. 88, 221–231.

Noy, I., Nualsri, A., 2007. What Do Exogenous Shocks Tell Us About Growth Theories? Working Paper 07-28, University of Hawaii, Hawaii, USA.
Okuyama, Y., 2003. Economics of natural disasters: a critical review. Regional Research Institute Working Papers, 131.
Okuyama, Y., 2004. Modeling spatial economic impacts of an earthquake: Input–output approaches. Disaster Prev. Manag. 13, 297–306.
Parker, M., 2018. The impact of disasters on inflation. Econ. Disasters Clim. Change 2 (1), 21–48.
Politis, D.N., Romano, J.P., Wolf, M., 1999. Subsampling. Springer-Verlag, New York.
Portnoy, S., 1991. Asymptotic behavior of the number of regression quantile breakpoints. SIAM J. Sci. Stat. Comput. 12 (4), 867–883.
Rose, A., Liao, S.-Y., 2005. Modeling regional economic resilience to disasters: a computable general equilibrium analysis of water service disruptions. J. Reg. Sci. 45, 75–112.
Rose, A., Wing, I.S., Wei, D., Wein, A., 2016. Economic impacts of a California tsunami. Nat. Hazards Rev. 17 (2), 1–12.
Sakov, A., Bickel, P.J., 2000. An edgeworth expansion for the m out of n bootstrapped median. Statist. Probab. Lett. 49, 217–223.
Santos, J.R., Haimes, Y.Y., 2004. Modelling the demand reduction Input-Output (I-O) inoperability due to terrorism of interconnected infrastructures. Risk Anal. 24, 1437–1451.
Sims, C.A., 1980. Macroeconomic and reality. Econometrica 48, 1–48.
Sohn, R.A., Barclay, A.H., Webb, S.C., 2004. Microearthquake patterns following the 1998 eruption of axial volcano, Juan de Fuca ridge: Mechanical relaxation and thermal strain: Axial volcano microearthquakes. J. Geophys. Res. 109, B01101.
Wei, Y., 2008. An approach to multivariate covariate-dependent quantile contours with applications to bivariate conditional growth charts. J. Amer. Statist. Assoc. 103, 397–409.
Weitzman, M.L., 2009. On modeling and interpreting the economics of catastrophic climate change. Rev. Econ. Stat. 91 (1), 1–19.
White, H., 2001. Asymptotic Theory for Econometricians. Academic Press, San Diego, CA.
White, H., Kim, T.-H., Manganelli, S., 2015. VAR for VaR: Measuring tail dependence using multivariate regression quantile. J. Econometrics 187 (1), 169–188.
Xiao, Z., 2017. QAR and quantile time series analysis. In: Chernozhukov, V., He, X., Koenker, R., Peng, L. (Eds.), Handbook of Quantile Regression. CRC Chapman-Hall.