The Nonlinear Case Against Leaning Against the Wind*

Nina Boyarchenko[1,2], Richard K. Crump[1], Keshav Dogra[1], Leonardo Elias[1] and Ignacio Lopez Gaffney[1]

[1]Federal Reserve Bank of New York
[2]CEPR and CesIfo

April 2024

Abstract

We re-examine the relationship between monetary policy and financial stability in a setting that allows for nonlinear, time-varying relationships between monetary policy, financial stability, and macroeconomic outcomes. Using novel machine learning techniques, we estimate a flexible "nonlinear VAR" for the stance of monetary policy, real activity, inflation, and financial conditions, and evaluate counterfactual evolutions of downside risk to real activity under alternative monetary policy paths. We find that a tighter path of monetary policy in 2003 – 2005 would have increased the risk of adverse real outcomes 3 – 4 years ahead, especially if the tightening had been large or rapid. This suggests that there is limited evidence to support "leaning against the wind" even once one allows for rich nonlinearities, intertemporal dependence, and crisis predictability.

*The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. We thank Marco Del Negro, Kinda Hachem, Anna Kovner, Silvia Miranda-Agrippino, and Guillaume Roussellet for helpful comments and suggestions. Emails: nina.boyarchenko@ny.frb.org; richard.crump@ny.frb.org; keshav.dogra@ny.frb.org; leonardo.elias@ny.frb.org; ignacio.lopezgaffney@ny.frb.org.

1 Introduction

In a speech in 2013, Federal Reserve Governor Stein argued that, in the real-world setting of an imperfect ability of supervisory and regulatory tools to promptly address rising financial vulnerabilities, monetary policy has the potential advantage of "getting in all the cracks" and, as such, it may be appropriate to use monetary policy to pursue financial stability objectives. From a theoretical perspective, using monetary policy to "lean against the wind" (LAW) of rising financial vulnerabilities trades off the potential benefit of reducing either the likelihood of a crisis and/or the depth of a recession conditional on a crisis occurring with dampened growth outside of crisis periods. A large empirical literature in the aftermath of the Global Financial Crisis (GFC) of 2008 has evaluated the merits of LAW and has mostly found the net benefits of LAW to be either small or even negative.

One potential criticism of this literature is that it makes a number of simplifying assumptions in quantifying both the potential benefits of intervening to prevent the build-up of financial vulnerabilities, as well as the potential costs of tightening monetary policy in response to a build-up of financial vulnerabilities. In this paper, we take a first step in addressing these shortcomings by considering a fully non-parametric, nonlinear specification for the joint dynamics of monetary policy, financial stability, and macroeconomic outcomes. This approach lets the data inform us about the intertemporal tradeoff (if any) between the central tendency and the left tail of the distribution of real outcomes, and how the stance of monetary policy affects the joint dynamics of financial conditions and downside risk to real outcomes.

We first show that, consistent with the findings in the "outlook-at-risk" literature, all else equal, tighter financial conditions increase downside risk to growth at near horizons but reduce downside risk at the 2 – 3 year horizon. This suggests a potential role for LAW in the nonlinear setting: more restrictive monetary policy tightens financial conditions, increasing near-term risk to growth but reducing medium-term risk to growth. Such dynamics would be consistent with the intuition that monetary policy may sometimes want to trigger a small, short recession to prevent the possibility of a financial crisis down the road.

However, when we evaluate counterfactuals with respect to the stance of monetary policy, we find that tighter monetary policy always increases the medium-term risk of adverse real outcomes, even when starting from a relatively loose level of financial conditions. In particular, a tighter path of policy in 2003 – 2005 would have substantially increased the risk of adverse real outcomes 3 – 4 years ahead, especially so if the tightening had been large or rapid. Examining how the effect of a counterfactual tightening of monetary policy varies over time, we find that a tighter path of policy increases medium-term downside risk to growth throughout the sample. Furthermore, tighter monetary policy modestly improves near-term risks to growth only immediately following recessions.

How can we reconcile the findings from our counterfactuals with respect to monetary policy tightening with the term structure of the counterfactuals with respect to tighter financial conditions? We argue that tighter financial conditions reduce medium-term risks to growth because monetary policy loosens in response to financial conditions. Instead, when we evaluate the counterfactuals with respect to a tighter path of policy, we preclude monetary policy from responding to the resulting financial conditions tightening. This further increases upside risk to financial conditions and therefore downside risk to real activity.

Our results thus suggest that, once we take into account the endogenous response of monetary policy to financial conditions, there is limited evidence to support LAW even in our setting which allows for rich nonlinearities, intertemporal dependence, and crisis predictability. Instead, the term structure of the predictive relationship between current financial conditions and future downside risk to growth suggests a possible role for macroprudential policies in mitigating financial instabilities.

We conclude by evaluating the contribution of monetary policy tightening in the most recent tightening cycle (starting in March 2022) to downside risk to growth. We find that monetary policy tightening meaningfully increased prospective tail risk to economic activity in early 2022 but that the contribution of the tightening cycle to downside risk to growth has since abated. As of 2023Q4, the policy path projected in the December 2023 Summary of Economic Projections (SEP) is close to our model's forecast and has little incremental effect on tail risk.

In our study, we rely on novel machine learning techniques that enable us to conduct counterfactual exercises in a computationally parsimonious manner. These techniques allow us to model jointly the dynamic evolution of monetary policy, financial conditions, inflation, and real activity without imposing parametric assumptions on the nature of the dynamic nonlinearities nor distributional assumptions on underlying shocks in the system. This flexibility comes at the cost of not being able to construct uncertainty bounds around the distributional impulse response functions (IRF) we estimate. We thus view this work as a first step in studying the interplay between monetary policy, financial vulnerabilities, and economic outcomes in a flexible nonlinear setting.

The rest of the paper is organized as follows. We describe our empirical approach in Section 2. Section 3 then presents our main results on the counterfactuals with respect to the net effect of monetary policy tightening. We evaluate the contribution of the current tightening cycle to downside risk to growth in Section 4. We discuss the relationship to the existing literature in Section 5. Section 6 concludes.

2 Empirical Approach

In this paper, we re-evaluate LAW in a setting that allows for nonlinear joint dynamics between inflation, real activity, policy, and financial conditions. More specifically, we consider a dynamic system of the form

$$\vec{y}_{t+1} | \, \text{past} \sim f\left(\cdot | \, \vec{y}_t, \ldots, \vec{y}_{t-L}\right), \tag{1}$$

where $\vec{y}_t$ is the vector of our variables of interest and f is a conditional distribution function that relates the past history of $\vec{y}_t$ to future realizations of $\vec{y}_{t+1}$. The general functional form in equation (1) allows for both nonlinear, state-dependent shock amplification as well as non-Gaussian features, and nests the more familiar linear Gaussian VAR

$$\vec{y}_{t+1} = B\left(L\right)\vec{y}_t + \Sigma^{1/2}\epsilon_{t+1} \quad \Rightarrow \quad \vec{y}_{t+1} \,|\, \vec{y}_t, \ldots, \vec{y}_{t-L} \sim \mathcal{N}\left(B\left(L\right)\vec{y}_t, \Sigma\right). \tag{2}$$

The nonlinear model also nests more complex parametric and semi-parametric dynamics such as a linear VAR with stochastic volatility

$$\vec{y}_{t+1} = B\left(L\right)\vec{y}_t + \Sigma_t^{1/2}\epsilon_{t+1} \quad \Rightarrow \quad \vec{y}_{t+1} \,|\, \vec{y}_t, \ldots, \vec{y}_{t-L} \sim \mathcal{N}\left(B\left(L\right)\vec{y}_t, \Sigma_t\right), \tag{3}$$

and a linear VAR with a non-Gaussian innovation distribution

$$\vec{y}_{t+1} = B\left(L\right)\vec{y}_t + \epsilon_{t+1} \quad \Rightarrow \quad \epsilon_{t+1} \,|\, \vec{y}_t, \ldots, \vec{y}_{t-L} \sim f_{\theta_\epsilon} \tag{4}$$

for some distribution f_{θ_ϵ}. Adrian et al. (2021) show that, even when estimating just the bivariate conditional joint distribution of real GDP growth and financial conditions, a fully nonlinear model outperforms both a linear VAR with normal innovations and a linear VAR with an arbitrary distribution for the innovations ϵ. We thus take the (possible) nonlinearity of the joint dynamics between inflation, real activity, policy, and financial conditions as given and let the data inform us on the shape of that nonlinearity.

While the nonlinear model is appealing from the perspective of allowing flexibility in the dynamic interactions between policy, financial conditions, and real activity, direct estimation of such dynamics is computationally infeasible. Instead, we use novel techniques from the machine learning literature that allow us to easily sample from the non-parametric conditional distribution of interest. We now describe these techniques and how we apply them to evaluating the net benefits of LAW.

2.1 Application to studying leaning against the wind

We are interested in evaluating the impact of counterfactual paths of monetary policy on the conditional distribution of real activity and other macroeconomic variables across different forecast horizons. We will thus estimate the dynamic relation between 4 variables: real activity, inflation, financial conditions, and the stance of monetary policy. We estimate the model using quarterly data from 1971Q1 to 2019Q4 (so that our estimates are not influenced by the evolution of the economy during the COVID-19 pandemic and in the aftermath of the pandemic).

We measure real activity ($ggap_t$) as the gap between the annualized quarterly growth rate of GDP and the Laubach and Williams (2003) (LW) 2-sided estimate of trend growth, and the stance of monetary policy ($rgap_t$) as the gap between the real interest rate and two-sided estimate of r_t^* from LW. Inflation is measured as the annualized quarterly core PCE inflation, π_t. Following Adrian et al. (2019), we measure financial conditions (fci_t) using the Federal Reserve Bank of Chicago National Financial Conditions Index (NFCI). A main advantage of the NFCI is that it includes measures of both quantities of credit and prices of credit (and risky assets). This is particularly important in our setting given that the literature on predictable financial crises (e.g Schularick and Taylor, 2012; Krishnamurthy and Muir, 2017; Greenwood et al., 2022) has emphasized both expansions in the quantity of credit and exuberance in pricing of credit as predictors of sharp economic downturns. Figure A.4 plots the time series of all four variables for reference.

Define $z_t = \{fci_t, ggap_t, \pi_t\}$ and $x_t = \{rgap_t, z_t\}$. We partition the distribution $f(x_t|x_{t-1}, ..., x_{t-4})$ as

$$f(x_t|x_{t-1}, ..., x_{t-4}) = f(rgap_t|z_t, x_{t-1})f(z_t|x_{t-1}, ..., x_{t-4}). \tag{5}$$

Notice that we assume $f(rgap_t|z_t, x_{t-1}) = f(rgap_t|z_t, x_{t-1}, ..., x_{t-4})$, i.e. the distributional "policy rule" only depends on contemporaneous variables and possibly one lag of variables, as in standard specifications of a Taylor rule. This way of partitioning the joint distribution also reflects an assumption we will make in our policy experiments: changes to $rgap_t$ do not contemporaneously affect z_t. We separately estimate $f(rgap_t|z_t, x_{t-1})$ and $f(z_t|x_{t-1}, ..., x_{t-4})$.

Our empirical methodology allows us to sample from the (estimated) conditional distribution $f(x_t|x_{t-1}, ..., x_{t-4})$. To sample a history $\{x_t, x_{t+1}\}$, we exploit the Markovian structure

$$f(x_{t+1}, x_t|x_{t-1}, ..., x_{t-4}) = f(x_{t+1}|x_t, ..., x_{t-3})f(x_t|x_{t-1}, ..., x_{t-4}). \tag{6}$$

That is, we draw x_t from the one-step ahead distribution given $x_{t-1}, ..., x_{t-4}$, then draw x_{t+1} from the one-step ahead distribution given $x_t, ..., x_{t-3}$. In this way, we can sample histories $\{x_t, ..., x_{t+H}\}$ of any length H. This gives us an estimate of the conditional distribution of each variable (e.g. $rgap$) at any horizon h, $f(rgap_{t+h}|x_{t-1}, ..., x_{t-4})$. We summarize these "baseline" or estimated distributions, given initial conditions, by their 10th, 50th and 90th percentiles, which we denote (e.g.) $Q^{rgap}_{t+h|t-1}(0.1)$, $Q^{rgap}_{t+h|t-1}(0.5)$, $Q^{rgap}_{t+h|t-1}(0.9)$.

In our experiments, we perturb this distribution – by changing initial conditions or shifting the conditional distribution of $rgap_t$ – generating an "alternative" conditional distribution of any variable at any horizon h. Again, we summarize these "alternative" distributions by their percentiles $\widetilde{Q}^{rgap}_{t+h|t-1}(0.1)$, $\widetilde{Q}^{rgap}_{t+h|t-1}(0.5)$, $\widetilde{Q}^{rgap}_{t+h|t-1}(0.9)$. Our figures plot "distributional impulse response functions" showing the difference between the corresponding percentiles of the h-quarters ahead distributions of each variable. For example, the difference in medians of $rgap_t$ is defined as

$$\Delta Q^{rgap}_{t+h|t-1}(0.5) \equiv \widetilde{Q}^{rgap}_{t+h|t-1}(0.5) - Q^{rgap}_{t+h|t-1}(0.5),$$

and the corresponding "distributional IRF" plots $\Delta Q^{rgap}_{t+h|t-1}(0.5)$ as a function of forecast horizon h. Note that these differences of percentiles are not percentiles of differences; thus, the *effects* on the various percentiles $\Delta Q^{rgap}_{t+h|t}(0.1), \Delta Q^{rgap}_{t+h|t}(0.5), \Delta Q^{rgap}_{t+h|t}(0.9)$ need not be ordered (i.e. the lines on the graph can cross), even though the percentiles of each underlying distribution are ordered.[1] Also, the effect of any perturbation will depend nonlinearly on both initial conditions and the magnitude of (e.g.) a shift in the distribution of $rgap_t$. Finally, when plotting IRFs of $ggap$ and π, we report the difference in quantiles of the distribution of annualized cumulative $ggap_{t+h}$ and π_{t+h}, e.g. the distribution of $\frac{4}{h+1}\sum_{k=0}^{h} ggap_{t+k}$, rather than the distribution of quarterly growth rates.

2.2 Wasserstein generative adversarial network (WGAN)

To estimate the conditional distributions given in equation (5) and evaluate joint dynamics of variables of interest under policy counterfactuals, we need an empirical approach that (*i*) allows us to flexibly estimate conditional distributions from the data and (*ii*) allows us to easily simulate from the estimated distribution. Adrian et al. (2021) estimate a bivariate nonlinear VAR using a non-parametric kernel approach, combined with discrete space Monte Carlo to simulate multi-period ahead distributions. The discrete space Monte Carlo approach, however, becomes computationally infeasible as more variables are added to the nonlinear VAR. Instead, we rely on a conditional Wasserstein generative adversarial network (WGAN) to estimate the joint dynamics.[2]

Intuitively, a WGAN chooses a distribution that is easy to simulate from (the "noise distribution", e.g., multivariate standard normal) and then passes each draw through a (possibly) nonlinear function with the aim of capturing the key features of the actual data. To achieve the latter goal, the functional form of this nonlinear function is "chosen" (i.e., estimated) in a way such that it is "hard" to statistically differentiate simulated samples from the actual sample. How does it do so? It utilizes a mini-max game between a "generator", which produces simulated samples as discussed with the aim of mimicking the data, and a "critic", which endeavors to successfully discriminate between artificial samples and the actual sample. To allow for a high degree of flexibility, the generator and the critic are

[1] Since we estimate the full conditional distribution, our estimated percentiles do not cross by construction.

[2] See Goodfellow et al. (2014) and Arjovsky et al. (2017) for the unconditional setting and Mirza and Osindero (2014), Odena et al. (2017), Kocaoglu et al. (2017), and Liu et al. (2018) for the conditional setting. For an econometric application, see Athey et al. (2021).

modeled via neural networks which can accommodate general forms of nonlinear behavior. This game iterates back and forth with a user-chosen stopping rule to create the estimated density. Although the density estimate is not available analytically, because of the structure, sampling from it is (essentially) as easy as sampling from the noise distribution itself.

To make things concrete, consider the example where we are interested in the conditional distribution of $\vec{y}$ given $\vec{x}$. The results of the WGAN estimation is a function $g(\cdot, \vec{x}; \widehat{\theta}_g)$ where $\widehat{\theta}_g$ are the estimated parameters. In our case, this is a large dimensional vector of estimated parameters for our neural network.[3] Then, if we draw u from the noise distribution, we have that $g(u, X_0; \widehat{\theta}_g)$ is a draw from our estimate of the conditional distribution of $\vec{y}$ given $\vec{x} = X_0$. We can then repeat this many times. To build intuition, recall the linear Gaussian VAR from the previous section. In this case, $\theta_g = \{B(L), \mathrm{vech}(\Sigma)\}$, and

$$g(u, \vec{x}; \theta_g) = B(L)\,\vec{x} + \Sigma^{1/2} u, \tag{7}$$

so that simulation draws can be obtained using $g\left(u, \vec{x}; \widehat{\theta}_g\right)$ with $u \sim \mathcal{N}(0, I)$.

As with other machine learning methods, implementation requires a number of choices of tuning parameters. Where possible, we use the default choices for tuning parameters from the `wgan` package.[4] We tailor the remaining tuning parameters to our setting, in particular, the modest sample size and large conditioning set. We use a batch size of one which gives the critic a single observation to discriminate from the data each time it updates its parameters. We then choose a generator dropout rate, maximum epoch, and the critic penalty term factor by matching stylized facts about economic variables. In particular, we enforce that stronger activity (higher *ggap*) predicts higher inflation and that monetary policy raises *real* rates in response to stronger activity or higher inflation.[5] Figure A.1 shows how the estimated

[3] The estimated parameters are not, themselves, of interest.

[4] See `https://github.com/gsbDBI/ds-wgan`. The `wgan` package is the replication package for Athey et al. (2021). We use the default tuning parameters for the noise distribution (multivariate standard normal), the neural network architecture (3 hidden layers with 128 nodes in each layer for both the critic and the generator), the critic dropout rate, the critic steps, and the test set size. We also leave the parameters for the stochastic gradient descent algorithm at their default choices.

[5] Since we measure the policy stance in terms of the real-rate gap, then a tightening of real rates requires

model satisfies these relations. The top row of Figure A.1 shows that the model satisfies a conventional Phillips Curve relation between inflation and activity. In Panel (a) we show $Q^{\pi}_{t|t-1}(c)$ for different values of c where we fix all initial conditions $x_{t-1}, \ldots, x_{t-4}$ except for $ggap_{t-1}$ which we allow to vary. All other variables are set to $ggap_{t-s} = rgap_{t-s} = fci_{t-s} = 0$, $\pi_{t-s} = 2$ for $s = 1, ..., 4$. In Panel (b) we again fix all initial conditions except that we vary $ggap_{t-1} = \ldots = ggap_{t-4}$. Both plots show that inflation is monotonically increasing with activity. The bottom row of Figure A.1 shows the response of monetary policy to changes in activity and inflation. Panels (c) and (d) shows $Q^{rgap}_{t|t-1}(c)$ for different values of c where we fix z_t and all initial conditions $x_{t-1}, \ldots, x_{t-4}$ (as above) except we vary $ggap_t = ggap_{t-1}$ (left panel) or $\pi_t = \pi_{t-1}$ (right panel).

Finally, we note that uncertainty quantification for the method we use is currently infeasible, and so we are unable to report standard errors around our estimates of (e.g.) the effect of monetary policy shocks on the 10th percentile of GDP.

3 Main Results

In this section, we report our main results. First, we show the term structure of the predictive relationship between current financial conditions and future downside risk to growth. We then show that the medium-termmedium-term downside risk to growth is *higher* under a counterfactually tighter path of monetary policy during 2003 –2005. Finally, we show that the adverse impact of tighter monetary policy on downside risk to growth is even larger if the tightening were large or rapid.

3.1 Financial conditions and downside risk to GDP

Adrian et al. (2022) find that tight financial conditions reduce the median and 5th percentile of real GDP growth in the near-term; however, at the 2-3 year horizon, tighter FCI predicts a

that nominal rates increase more than one for one (the Taylor principle).

slight increase in the 5th percentile. If interpreted causally, this suggests a tradeoff: tighter financial conditions reduce activity on average, but can potentially lessen the tail risk of a financial crisis in the medium-term. Potentially then, by tightening financial conditions, tighter monetary policy could reduce the tail risk of a crisis, albeit at the cost of lower economic activity on average.

We first verify that Adrian et al. (2022)'s stylized fact, which they estimate using a linear quantile regression methodology in a sample of 11 advanced economies, holds for the U. S. according to our empirical methodology. Figure 1 plots distributional IRFs of all variables to a shock of 0.5 to the NFCI fci_t.[6] Under the baseline, we set all initial conditions as follows: $ggap_{t-s} = rgap_{t-s} = fci_{t-s} = 0$, $\pi_{t-s} = 2$, for all lags $s = 1, ..., 4$, and $ggap_t = fci_t = 0$, $\pi_t = 2$. We then draw r_t, $x_{t+1}, x_{t+2}, ...$ from the estimated conditional distributions. Under the alternative, we modify fci_t to 0.5, keep $ggap_t$, π_t, and all lagged variables the same, and draw $r_t, x_{t+1}, x_{t+2}, ...$ from the estimated distributions.

The bottom left panel of Figure 1 plots the response of the annualized cumulative change in $ggap_t$, h quarters out.[7] Broadly in line with Adrian et al. (2022), tighter financial conditions sharply reduce activity in the near-term, but increase the 10th percentile of GDP, relative to baseline, after the 2-year horizon.

3.2 Effect of tighter monetary policy

Do these results imply that contractionary monetary policy can reduce downside risk to GDP growth? To answer this question, we look at the 2003 – 2005 period. Starting with Taylor (2007), many commentators have argued that loose monetary policy during this period helped accelerate the housing boom and encouraged risk-taking, contributing to the

[6] This is roughly the standard deviation of quarterly changes to NFCI in our sample. Recall that the NFCI is standardized to have a full sample mean of 0 and standard deviation of 1.

[7] For example, the blue line plots the effect of the NFCI shock on the 10th percentile of the distribution of annualized cumulative $ggap_{t+h}$, $\Delta Q^{ggap}_{t+h|t-1}(0.1)$. $\Delta Q^{ggap}_{t+4|t-1}(0.1) \approx -0.6$. This means the 10th percentile of the distribution of GDP, 4 quarters ahead, is around 0.6 percent lower under the "alternative" distribution following an FCI shock, relative to the "baseline" distribution with no shock.

Figure 1. Response to a shock of 0.5 to the NFCI This figure presents distributional IRFs $\Delta Q^{v}_{t+h|t-1}(c)$ to a shock of 0.5 to the NFCI, where $c \in \{0.1, 0.5, 0.9\}$ and $v \in \{fci, rgap, ggap, \pi\}$. We start from initial condition $ggap_{t-s} = rgap_{t-s} = fci_{t-s} = 0$, $\pi_{t-s} = 2$, for all lags $s = 1, ..., 4$, $ggap_t = fci_t = 0$, $\pi_t = 2$. For $ggap$ and π we report the difference in quantiles of the distribution of annualized cumulative $ggap_{t+h}$ and π_{t+h}, e.g. the distribution of $\frac{4}{h+1}\sum_{k=0}^{h} ggap_{t+k}$. The x-axes denote the forecast horizon h.

(a) NFCI

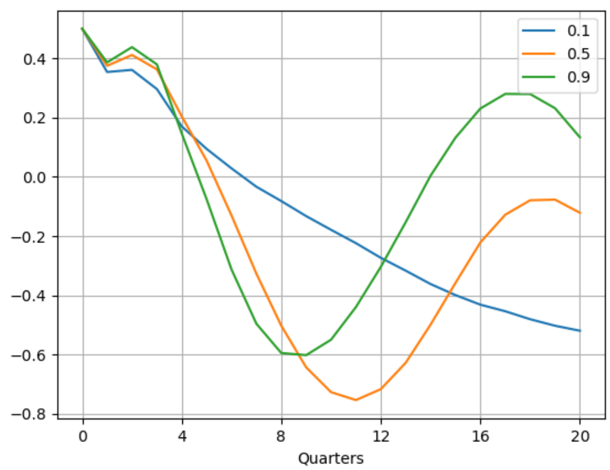

(b) Real rate gap

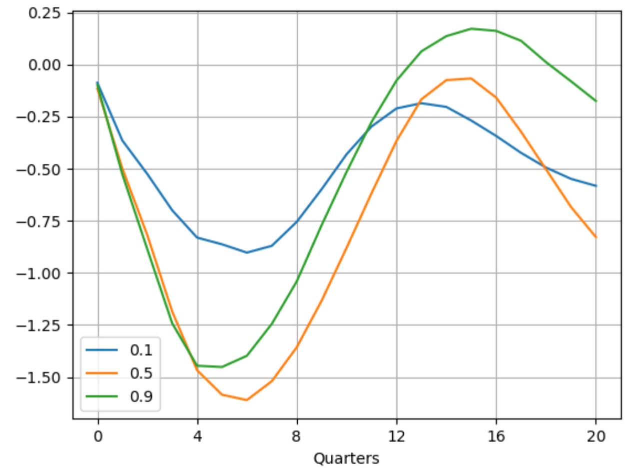

(c) Annualized cumulative $ggap_t$

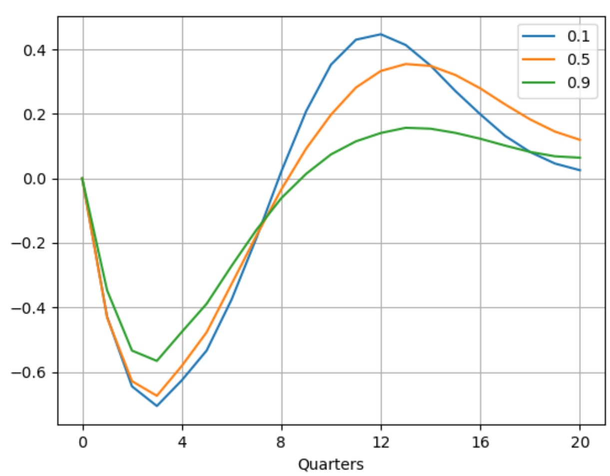

(d) Annualized cumulative core PCE inflation

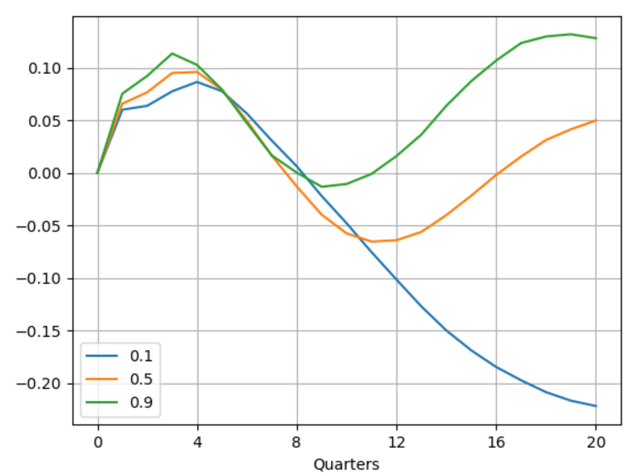

global financial crisis. Arguably then, this period should provide the best case to assess the hypothesis that leaning against the wind can reduce tail risk.[8]

In our baseline, we set t=2004Q1 and use realized values from 2003Q1 through 2004Q1 as initial conditions. In our alternative scenario, we perturb $f(rgap_t|z_t, x_{t-1})$ for 4 quarters

[8] In unreported results, we also conducted the monetary policy tightening experiment for initial conditions set as an average of "normal time" $ggap$, $rgap$, π and fci. The conclusions of that experiment are broadly in line with those for the 2003 – 2005 experiment discussed here.

by adding a series of shocks to $rgap_t$ that cumulatively raise its conditional median value by 50bps per quarter, relative to the baseline distribution, i.e. after 4 quarters, the conditional median of $rgap_t$, $\widetilde{Q}^{rgap}_{t+3|t-1}(0.5)$, is 200bps above its baseline value $Q^{rgap}_{t+3|t-1}(0.5)$ (the orange line in the top left panel of Figure 2 plots this shift in the median, $\Delta Q^{rgap}_{t+h|t-1}(0.5)$).[9] This increases the real interest rate gap from its initial value of -240bps to roughly zero (given that the gap narrows slightly even under the baseline distribution).

Figure 2 plots distributional IRFs to this monetary policy shock. As expected, tighter monetary policy leads to tighter financial conditions (top right panel). However, rather than reducing the tail risk of GDP, this tightening shifts down the whole distribution, particularly the left tail, at the 3-year horizon (blue line in bottom left panel).[10] This is reminiscent of Svensson (2017)'s argument that even if tighter policy reduces the probability of a crisis, leaning against the wind is undesirable since it weakens economic activity *conditional* on a crisis occuring. An important difference is that rather than imposing the tight parametric restrictions in Svensson (2017)'s calibration exercise, we flexibly estimate the effect of policy on the whole distribution of outcomes.

Given that tighter financial conditions predict less GDP tail risk in the medium-term, why does tighter monetary policy fail to have a similar effect? One hypothesis is that, at least according to our model, the relation between NFCI and GDP risk shown in Figure 1 reflects the endogenous response of monetary policy. The top right panel in Figure 1 shows that real rates fall rapidly in response to tighter FCI, bottoming out around 6 quarters out; lower real rates mitigate the effect of tighter FCI and potentially explain the increase in GDP at the 3-year horizon. Thus, while *ceteris paribus* tighter financial conditions reduce tail risk 3 years out – because they elicit *looser* monetary policy – tighter financial conditions induced by *tighter* monetary policy increase tail risk. This endogenous response of monetary

[9] Note that this is a different experiment than adding a series of consecutive 50bps shocks, given that in the alternative scenario, $rgap_t$ endogenously responds to past shocks, in part via their effect on macroeconomic outcomes. That said, the effect of a series of 50bps shocks is qualitatively similar.

[10] The bottom right panel shows that our model, like linear VARs estimated with similar variables over a similar period, suffers from a price puzzle: tighter monetary policy slightly increases inflation.

Figure 2. Response to tighter monetary policy in 2004Q1 This figure presents distributional IRFs $\Delta Q^{v}_{t+h|t-1}(c)$ to a series of $rgap_t$ shocks that raise $rgap_t$ by 50bps per quarter for 4 quarters, where $c \in \{0.1, 0.5, 0.9\}$ and $v \in \{fci, rgap, ggap, \pi\}$. We start from an initial condition t=2004Q1. For $ggap$ and π we report the difference in quantiles of the distribution of annualized cumulative $ggap_{t+h}$ and π_{t+h}, e.g. the distribution of $\frac{4}{h+1}\sum_{k=0}^{h} ggap_{t+k}$. The x-axes denote the forecast horizon h.

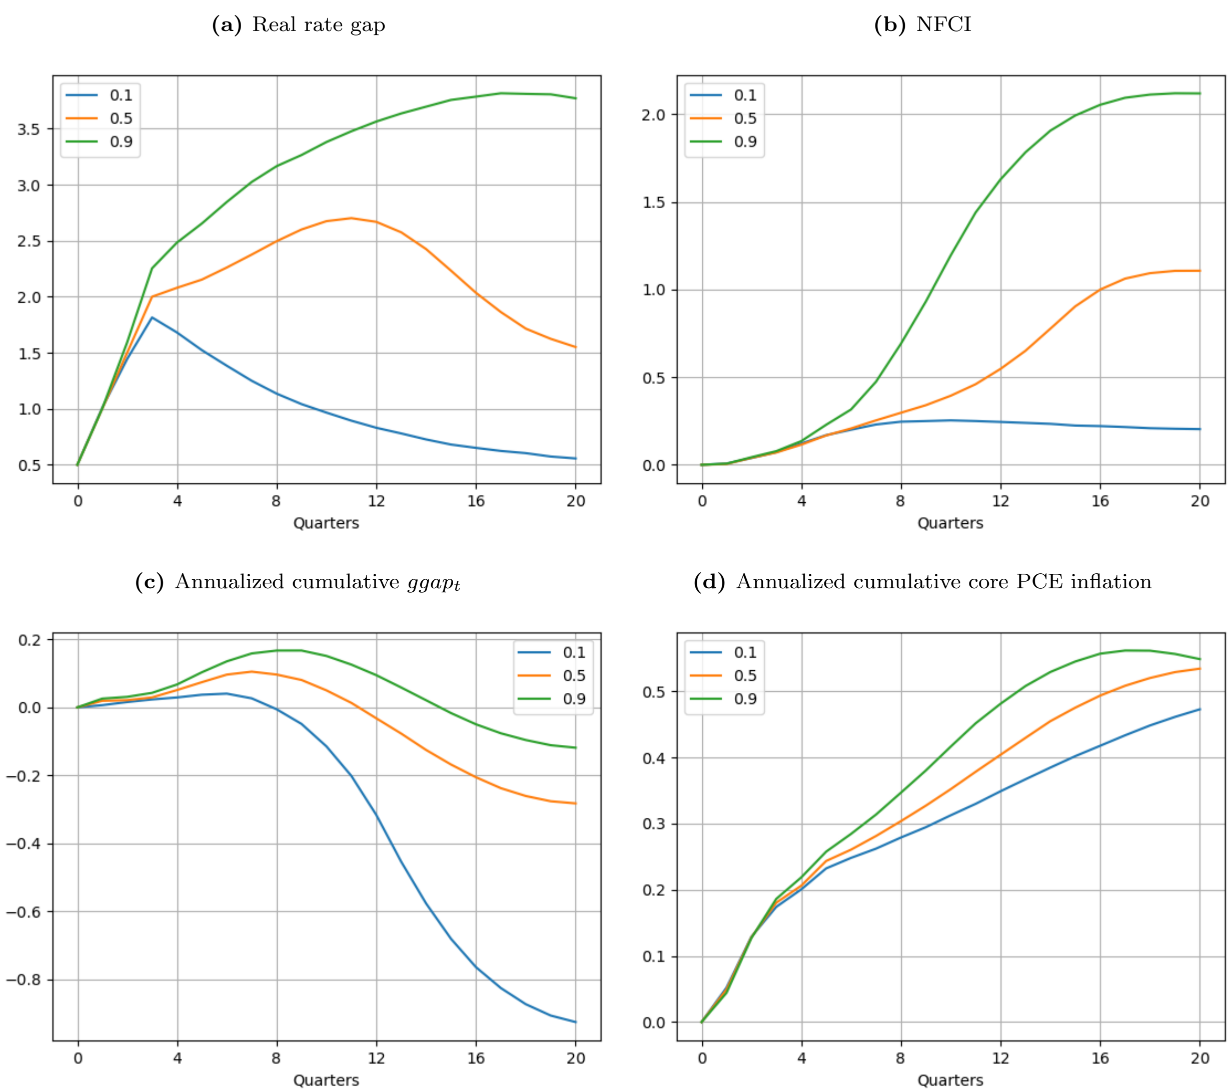

(a) Real rate gap

(b) NFCI

(c) Annualized cumulative $ggap_t$

(d) Annualized cumulative core PCE inflation

policy to tighter financial conditions is also consistent with the findings in Brunnermeier et al. (2021) who argue that the negative predictive relationship between credit growth and future economic outcomes documented in the predictable financial crises literature is the consequence of endogenous responses of monetary policy to looser credit conditions.

A limitation of Figure 1 is that it only plots "marginal" conditional IRFs, and does not show the joint evolution of *ggap*, NFCI, and *rgap*. While the top right panel shows that an FCI shock is generally followed by a loosening of monetary policy, it does not show whether this loosening occurs along histories associated with tail outcomes for *ggap*. Thus, it is not clear whether this endogenous response of monetary policy is responsible for the increase in the 10th percentile of GDP 3 years after an FCI shock. To investigate this further, we perform the following conditioning exercise. We first calculate the 10th percentile path for *ggap* under the baseline (no shock to financial conditions) and alternative (0.5 shock to financial conditions) scenarios. The difference between these two paths is shown by the blue line in the bottom left panel of Figure 1. We then consider all of the draws for the baseline scenario and select only those for which *ggap* is sufficiently close to its 10th percentile path under the baseline. We then calculate the median, 25th, and 75th percentiles of the other variables (e.g., *rgap*) within this subset of draws. The median of *rgap* within this subset, for example, can be interpreted as the median of *rgap* conditional on a tail scenario for growth. We perform the same exercise in the alternative scenario, calculating the median, 25th and 75th percentiles of all variables, conditional on *ggap* being sufficiently close to its 10th percentile path under the alternative scenario. We then have the conditional median, 25th and 75th percentiles under the baseline and alternative scenarios. As in Figure 1, we plot their differences, which are shown in Figure 3.

The bottom left panel shows the differences in percentiles for *ggap*. By construction, all three of these lines are almost exactly equal to the blue line in the bottom left panel of Figure 1. The top right panel shows the difference in *rgap* with and without the FCI shock, again conditioning on outcomes where *ggap* is at its 10th percentile. Consistent with our conjecture that monetary policy plays a key role in the response of growth to tightening financial conditions, the distributional IRF for *rgap* is sharply negative around 4 quarters out. Following a tightening of financial conditions, monetary policy loosens more aggressively when downside outcomes for growth are realized, as compared to the policy response to weak growth in the absence of an FCI shock. Furthermore, the policy response to the FCI shock more than undoes the initial tightening of financial conditions. The top left panel shows that NFCI is initially tighter following the shock (by construction), but at a 2- to 3-year horizon, it is looser relative to the baseline.

Figure 3. Conditional response to a shock of 0.5 to the NFCI This figure presents conditional distributional IRFs to a shock of 0.5 to the NFCI where we restrict to paths which are sufficiently close to the 10th percentile path of $ggap$ under the baseline and alternative scenarios. We plot the difference of the median, 25th and 75th percentiles of each variables in baseline and shocked scenarios, where in both baseline and shocked scenarios we condition on the event that $\sum_h |ggap_{t+h} - Q^{ggap}_{t+h|t-1}(0.1)|^2 < 1$ (that is, we drop all draws for which the evolution of $ggap_t$ does not approximately equal its 10th percentile). We start from initial condition $ggap_{t-s} = rgap_{t-s} = fci_{t-s} = 0$, $\pi_{t-s} = 2$, for all lags $s = 1, ..., 4$, $ggap_t = fci_t = 0$, $\pi_t = 2$. For $ggap$ and π we report the difference in quantiles of the distribution of annualized cumulative $ggap_{t+h}$ and π_{t+h}, e.g. the distribution of $\frac{4}{h+1}\sum_{k=0}^{h} ggap_{t+k}$. The x-axes denote the forecast horizon h.

(a) NFCI

(b) Real rate gap

(c) Annualized cumulative $ggap_t$

(d) Annualized cumulative core PCE inflation

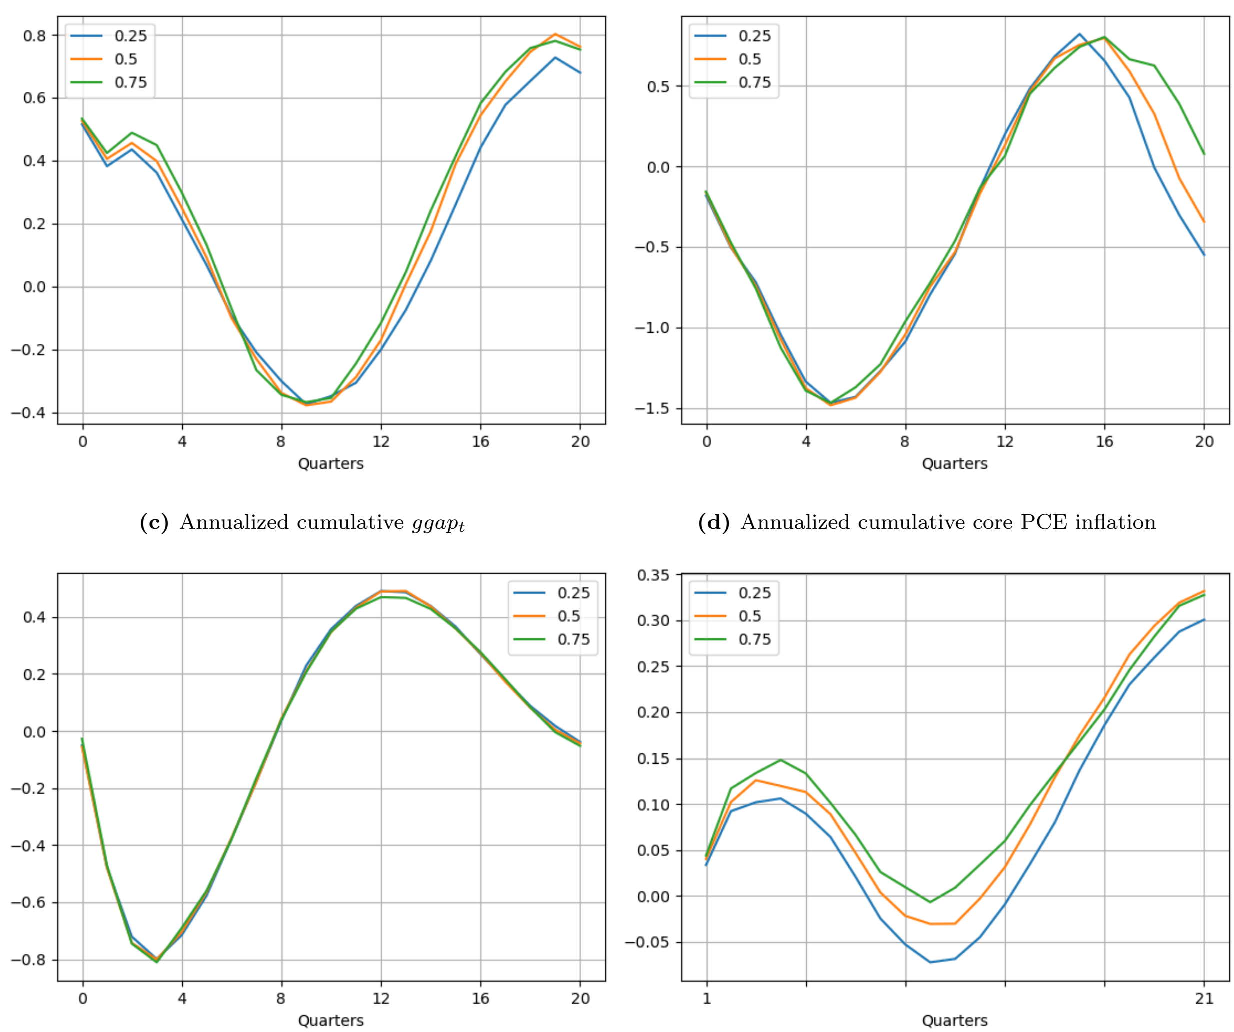

3.3 How do the speed and magnitude of policy tightening matter?

One might worry that the adverse effects of the monetary tightening scenario above arise because of its large magnitude and rapid pace (four consecutive 50bps increases). Perhaps a smaller or more gradual tightening might induce different dynamics for tail risk. In principle, our nonlinear VAR allows for this possibility. We therefore consider four alternative experiments in which we vary the magnitude of tightening (four consecutive 25bps or 100bps hikes, rather than 50bps) and its pace (eight 25bps hikes or two 100bps hikes, rather than four 50bps hikes). The right panel of Figure 4 plots the response of the 10th percentile of annualized cumulative $ggap_t$; the left panel plots the median real rate gap.

Increasing the magnitude of the policy tightening – from 25bps per quarter (blue line in left panel), to 50bps (orange line), to 100bps (green line) – deepens and frontloads the decline in the left tail of GDP (shown by the corresponding lines in the right panel). Increasing the pace of the tightening while keeping its total magnitude (200bps) fixed produces a similar effect – compare the purple (eight 25bps hikes), orange (four 50bps hikes) and red (two 100bps hikes) lines in the right panel. But while a more moderate and gradual pace of hikes has less harmful effects on tail risks, it never *reduces* risk.

We have focussed on the period preceding the GFC in our exercises. Since our model is nonlinear and state-dependent it is possible that these results are unrepresentative of other time periods in our sample. Figure A.2 in the Appendix shows a time series of $\Delta Q^{ggap}_{t+h|t-1}(c)$ for $c \in \{0.1, 0.5\}$ and $h \in \{4, 8, 12\}$, where we vary t over the sample. In Panel (a) each observation corresponds to the effect of a cumulative 200 basis point policy tightening over 4 quarters on median cumulated $ggap$ at various forecast horizons starting at a *different* period t. Panel (b) is constructed in the same way but showing the response of the 10th percentile rather than the median. The observations at $t = 2004Q1$ correspond to the distributional IRFs shown in the bottom left panel in Figure 2 at horizons of 4, 8, and 12 quarters. The charts show that this policy tightening experiment never meaningfully increases either the median or the left tail of GDP, and it often has a substantial negative effect. Thus, for any

Figure 4. Alternative policy tightening experiments This figure plots the response of the median real rate gap $\Delta Q^{rgap}_{t+h|t-1}(0.5)$ and the 10th percentile of cumulative $ggap_t$, $\Delta Q^{ggap}_{t+h|t-1}(0.1)$ under the five tightening experiments described in the text. We start from an initial condition t=2004Q1. For $ggap$ we report the difference in quantiles of the distribution of annualized cumulative $ggap_{t+h}$ and π_{t+h}, e.g. the distribution of $\frac{4}{h+1}\sum_{k=0}^{h} ggap_{t+k}$. The x-axes denote the forecast horizon h.

(a) Median real rate gap $rgap_t$ **(b)** 10th percentile of cumulative $ggap_t$

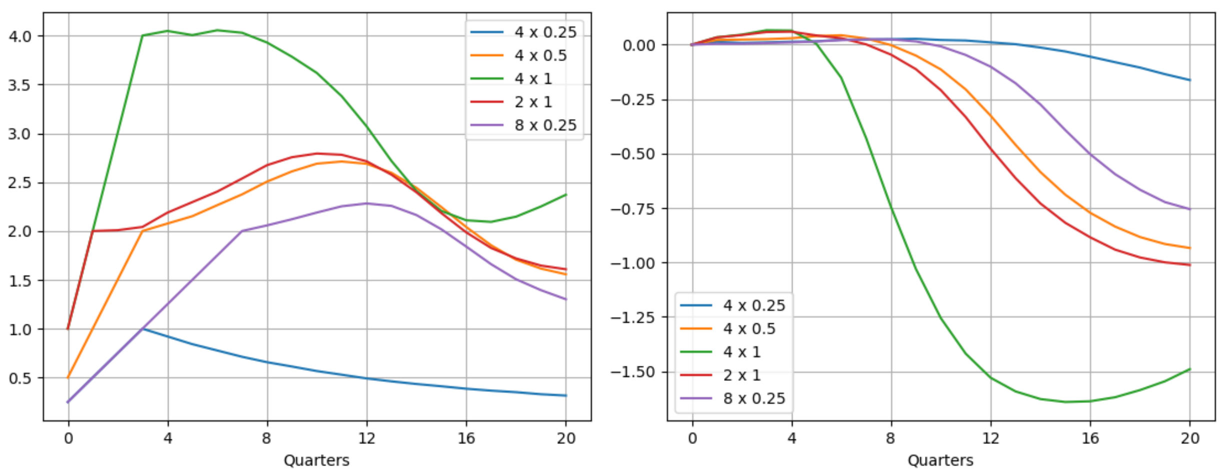

initial conditions in our sample, we fail to find that "leaning against the wind" is beneficial for materially reducing tail risk to activity.

Finally, we have centered our empirical approach on being sufficiently flexible to capture the type of nonlinear and time-varying effects that are implicit in the LAW hypothesis. To that end, we proceeded with only a modest degree of regularization to ensure that the estimated model had the required flexibility to capture these effects if they are indeed present. As a robustness check, in the Appendix we provide results from a more stringently regularized version of the model. As a simple way to summarize the results from this auxiliary model, Figure A.3 presents $\Delta Q^{ggap}_{t+h|t-1}(c)$ for $c \in \{0.1, 0.5\}$ and $h \in \{4, 8, 12\}$, where we vary t over the sample. We can compare directly to Figure A.2 and observe again that monetary policy tightening shocks are generally associated with downward shifts to the $ggap$ distribution at all horizons. The only exceptions are periods immediately following a recession when the monetary policy is already loose.

4 The recent policy cycle through the lens of the model

We have seen that according to our model, the pace and magnitude of policy tightening affect macroeconomic outcomes, particularly tail risks, in a nonlinear fashion. Given the unprecedented pace of the tightening cycle over the past two years, it is natural to ask what effect the path of policy – both the realized path in our sample, and the expected future path as projected in the December 2023 SEP medians – had on the evolution of tail risks. We do so as follows.

To construct a projected path of $rgap_t$ beyond 2023Q4 when our sample ends, we take the median SEP participant's projection for the end of year federal funds rate for 2024Q4, 2025Q4 and 2026Q4. To construct real rates, we subtract SEP median projections of Q4/Q4 core PCE inflation at these horizons. We assume the LW estimate of r^* remains constant at its 2023Q4 value of 1.12. We linearly extrapolate the Q4 values of $r_t - r_t^*$ for intervening quarters. In what follows, for convenience we refer to the realized path of $rgap_t$ extended by these SEP projections as the "FOMC path".

For each quarter t=2022Q2,...,2023Q4, we use realized values through quarter t as initial conditions. The top-left panel in Figure 5 shows the evolution of the h-quarter ahead 10th percentile of annualized cumulative $ggap_t$ under the *baseline* distribution for $h = 4, 8, 12$. This baseline distribution is constructed by conditioning on realized data on z_t through quarter t and $rgap_t$ through quarter $t - 1$, but *not* the subsequent path of $rgap$, and instead drawing from the conditional forecast distribution. Importantly, here "baseline" denotes the model's forecast conditional on information until date t, *not* the rate path that actually realized. In early 2022, the model sees substantial downside risk across all three forecast horizons. As we move through 2022 into 2023, near-term downside risk recedes (the blue line rises), whereas tail risk generally worsens at the three year horizon. This likely reflects two competing forces: the model sees better than expected realizations of growth, which reduces near-term risks, but also sees a tighter than expected path of policy, which increases medium-term risks.

Figure 5. Effect of 2022-23 tightening Panel (a) presents the 10th percentile of the distribution of cumulative $ggap_t$ $h \in \{4, 8, 12\}$ quarters ahead, $Q^{ggap}_{t+h|t-1}(0.1)$, conditional on data through quarter t =2022Q2,...,2023Q4 but *not* the subsequent path of rates. Panels (b) through (d) plot the distributional impulse response of the 10th percentile of cumulative $ggap_t$, $\Delta Q^{ggap}_{t+h|t-1}(0.1)$, the median of $rgap_t$, $\Delta Q^{rgap}_{t+h|t-1}(0.5)$, and the 90th percentile of NFCI, $\Delta Q^{fci}_{t+h|t-1}(0.9)$, to a series of $rgap_t$ shocks, starting in quarter t =2022Q2,...,2023Q4. The sequence of $rgap_t$ shocks is chosen to raise the conditional median of $rgap_t$, $\widetilde{Q}^{rgap}_{t+h|t-1}(0.5)$, to coincide with the FOMC path for dates t through 2026Q4. The x-axes denote quarter t.

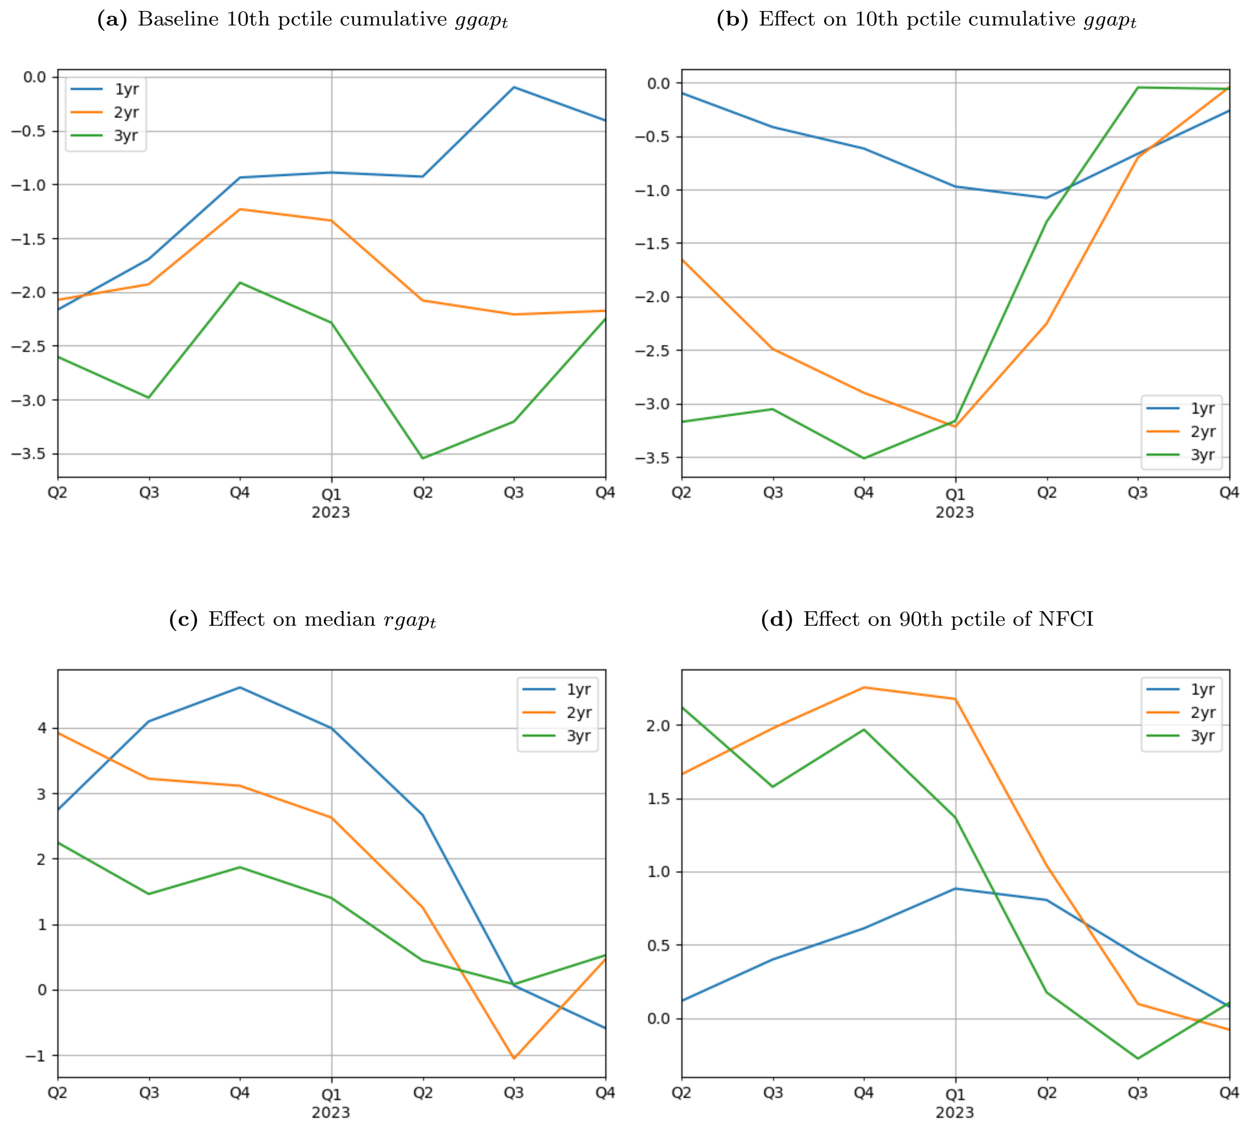

(a) Baseline 10th pctile cumulative $ggap_t$

(b) Effect on 10th pctile cumulative $ggap_t$

(c) Effect on median $rgap_t$

(d) Effect on 90th pctile of NFCI

We compare these baseline distributions to "shocked" distributions, where in each quarter $t + h$ from t to 2026Q4, we add an $rgap_t$ shock that raises the conditional median value of $rgap_t$, $\widetilde{Q}^{rgap}_{t+h|t-1}(0.5)$, to match the realized FOMC path. The bottom-left panel of Figure

5 shows the difference between the median value of $rgap_t$ h quarters out under the FOMC path and the baseline forecast, starting from each quarter t, for $h = 4, 8, 12$. Initially, the model expects substantially less tightening than actually occurred, especially at shorter horizons; that is, relative to the baseline, the FOMC path reflects a substantial and frontloaded tightening of policy. As we move through 2022 and the model sees the tightening that has already occurred, the difference between its future projections and the FOMC path narrows. As of 2023Q4, the median path projected by the model is broadly similar to the FOMC path.

The top-right panel of Figure 5 shows the effect of the FOMC path, relative to baseline, on the 10th percentile of cumulative annualized $ggap_t$, starting from t=2022Q2 to 2023Q4. For example, the green line shows that as of t=2023Q1, $\Delta Q^{ggap}_{t+12|t-1}(0.1) \approx -3.0$. That is, from the perspective of 2023Q1, the additional policy tightening embodied in the subsequent FOMC path, relative to the model's forecast based on 2023Q1 information, reduced the 10th percentile of annualized GDP growth over the subsequent 3 years by 3 percent. Since tighter policy affects activity with long lags according to the model, the substantial policy tightening represented by the FOMC path increases tail risks only modestly at a 1-year horizon. At a 2- and 3-year horizon, however, this tightening substantially increases tail risk from the perspective of early 2022. As we move into 2023 and the amount of future expected tightening diminishes (bottom left panel), it increases tail risk at a 2 or 3 year horizon less and less (orange and green lines in top right panel). By 2023Q4, the near-term FOMC path is broadly similar to the model's baseline median forecast, and so the effect on medium-term tail risk is negligible. The effect of the FOMC path on the 90th percentile of financial conditions (bottom right panel) is a rough mirror image of its effect on the left tail of GDP: upside risk to NFCI is associated with downside risk to GDP. Broadly speaking then, the model sees the recent tightening cycle as an ex-ante risky bet – arguably a bet worth taking, given elevated inflation – which has (so far) paid off ex-post.

5 Discussion and Relation to Existing Literature

Overall, the results in Figures 2–4 suggest that leaning against the wind may have adverse consequences for downside risk to growth even in a nonlinear setting. This stands in contrast with recent findings in Grimm et al. (2023), who argue that persistently loose monetary policy leads to elevated financial crisis risks. There are a number of reasons why our exercise and theirs may come to different conclusions.

First, Grimm et al. (2023) consider much longer periods of "low for long" monetary policy (5 year average deviations of r from r^*) than the policy tightening experiments we study, and find significant effects at a longer horizon (5-10 years out). However, given the results of the experiments reported in Figure 4, it is unlikely that a prolonged (5 year) but shallow monetary policy tightening would reduce downside risk. Second, while Grimm et al. (2023) study the effect of policy on the probability of financial crises (and find no effect on the probability of "normal" recessions), we study its effect on downside risk more broadly, which may arise from either financial crises or normal recessions. Also, given the potentially nonlinear relationship between real activity, policy stance and financial conditions, the effects of monetary policy tightening and loosening may be asymmetric. That is, it is possible in a "nonlinear VAR" for a "Goldilocks" principle to apply, where both excessively tight and excessively loose policy increase financial crisis risk.

Finally, and perhaps most importantly, Grimm et al. (2023) and the majority of papers on predictable financial crises rely on long-history international panel data. While these data have the advantage of covering multiple periods of crises and/or loose policy, the connection between the monetary policy stance and financial conditions in the rest of the world likely differs from that in the U.S. For example, the literature on the global financial cycle (Rey, 2015, and the subsequent literature) emphasizes that there is a large degree of comovement in global asset prices, so that local financial conditions are less affected by local monetary policy than is the case in the U.S. The effect of monetary policy on the evolution of financial conditions and downside risks to growth may thus be fundamentally different in the rest of the world than in the U.S. Boyarchenko and Elias (2024) further show that there is a global component to downside risk to growth, with a tightening in global credit conditions predicting increased downside risk to domestic growth across a number of economies.

More broadly, our paper brings together the literature on leaning against the wind with the literature on outlook-at-risk. In a seminal paper, Svensson (2017) argues that the net benefits of LAW are likely to be small both because tighter monetary policy has a substantial impact on unemployment if a crisis does not materialize and only a small effect on the probability of a crisis in the medium run. In the context of growing evidence on the predictability of financial crises (see e.g. Schularick and Taylor, 2012; Mian et al., 2017; Krishnamurthy and Muir, 2017; Richter et al., 2021; Greenwood et al., 2022, and the subsequent literature), the negative conclusion of Svensson (2017) on LAW seems counterintuitive: if a policy maker can detect the build-up of financial vulnerabilities in real time, how can monetary policy be ineffective in mitigating the rising downside risks to real activity?

One potential criticism of Svensson (2017) is that it considers LAW in a highly simplified setting, with a deterministic impact of monetary policy tightening on both the average level of unemployment and future crisis probability. The subsequent LAW literature has relaxed a number of the assumptions implicit in Svensson (2017). For example, Schularick et al. (2021) consider the effect of discretionary leaning against the wind in response to rising levels of credit in the economy, and find that a one percentage point tightening in the policy rate on average increases crisis probability by 2 percentage points 1 to 2 years out. Similarly, Brandão-Marques et al. (2021) and Ajello and Pike (2022) consider the effect of monetary policy on downside risk to growth in a quantile regression setting and find no to slightly detrimental effect of monetary policy tightening on risks in the economy.

In contrast to the empirical literature on LAW, the theoretical literature is more optimistic on the net benefits of LAW. Gourio et al. (2018) argue that the answer to whether or not LAW is beneficial depends on the source of the shocks that monetary policy responds to. Monetary policy that *systematically* responds to excess credit growth lowers average financial crisis probability, at the cost of larger amplitudes of "normal" business cycles. In the same spirit, Adrian and Duarte (2018) argue that optimal monetary policy responds to financial vulnerabilities in an economy in which financial conditions affect the conditional distribution of future real outcomes, even if financial stability itself is not an objective of monetary policy. Boissay et al. (2021) likewise argue that, when financial crises are endogenous, a central bank that moves from an inflation targeting framework to one that targets also an output gap and an index of financial fragility achieves a lower probability of financial crises. Ajello et al. (2019) also find that, once monetary policy takes into account the uncertainty around the point estimates of the relationship between credit conditions, economic outcomes, and financial crises, the optimal policy responds more aggressively to financial instability.

6 Conclusion

We study the benefits and costs of leaning against the wind – that is, changing the conduct of monetary policy in response to a build-up of financial vulnerabilities – in a flexible, non-parametric model of the dynamic interactions between monetary policy, financial conditions, and macroeconomic outcomes. We find that downside risk to growth increases in response to a counterfactual tightening of the path of monetary policy, suggesting that LAW is detrimental even once one allows for rich nonlinearities, intertemporal dependence, and crisis predictability.

Our conclusions are subject to a number of caveats. First, we are using a reduced form model to predict the effect of changes in policy. In doing so, we are effectively assuming that a tighter path of policy is equivalent to a sequence of contractionary monetary policy "shocks"; our measure of the stance of policy is only one-dimensional, and we neglect the role of expected future policy or explicit forward guidance. Furthermore, we effectively identify the response of our variables of interest to a monetary policy shock using timing restrictions: we assume that within a period, $rgap_t$ reacts to z_t, but not vice versa. This assumption is vulnerable to all the same critiques as in a linear VAR setting, especially since our system includes a financial variable (NFCI) which may respond contemporaneously to policy. Of course, in linear VAR models, the literature has considered various alternative identification approaches to circumvent this and other issues; we view our approach as a first step to studying monetary policy shocks in our rich nonlinear framework.

Second, the counterfactual exercises we conduct are subject to the usual Lucas critique. A monetary policy that systematically reacts to loose financial conditions would likely fundamentally change the relationship between monetary policy and downside risk to growth. For example, Adrian and Duarte (2018) consider the conduct of monetary policy by a central bank facing an economy in which financial conditions affect the conditional distribution of future real outcomes. In that economy, the optimal monetary policy rule always depends on financial vulnerabilities (as well as the output gap, inflation, and the natural rate), even when financial conditions themselves are not a target of monetary policy.

Finally, our estimates are subject to the usual caveats on the distinction between the effects of monetary policy surprises, realized monetary policy stance, and monetary policy rules. The overall conduct of monetary policy may affect the buildup of financial vulnerabilities through its impact on households', firms', and investors' policy expectations and investment and consumption decisions made conditional on those expectations. Thus, for example, households in Grimm et al. (2023) are able to borrow more, with higher house prices, because monetary policy is systematically loose. Estimating the effect of changes in the conduct of monetary policy on the buildup of financial vulnerabilities, however, remains challenging due to the paucity of changes in the conduct of monetary policy, the simultaneous impact of an evolving regulatory environment, and the rare nature of financial crises.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable growth," *American Economic Review*, 109, 1263–1289.

——— (2021): "Multimodality in macrofinancial dynamics," *International Economic Review*, 62, 861–886.

Adrian, T. and F. Duarte (2018): "Financial vulnerability and monetary policy," .

Adrian, T., F. Grinberg, N. Liang, S. Malik, and J. Yu (2022): "The Term Structure of Growth-at-Risk," *American Economic Journal: Macroeconomics*, 14, 283–323.

Ajello, A., T. Laubach, D. López-Salido, and T. Nakata (2019): "Financial Stability and Optimal Interest Rate Policy," *International Journal of Central Banking*, 15, 279–326.

Ajello, A. and T. Pike (2022): "Getting in All the Cracks: Monetary Policy, Financial Vulnerabilities, and Macro Risk," Working Paper, Federal Reserve Board of Governors.

Arjovsky, M., S. Chintala, and L. Bottou (2017): "Wasserstein Generative Adversarial Networks," in *Proceedings of the 34th International Conference on Machine Learning*, ed. by D. Precup and Y. W. Teh, PMLR, vol. 70 of *Proceedings of Machine Learning Research*, 214–223.

Athey, S., G. W. Imbens, J. Metzger, and E. Munro (2021): "Using Wasserstein generative adversarial networks for the design of monte carlo simulations," *Journal of Econometrics*.

Boissay, F., F. Collard, J. Galí, and C. Manea (2021): "Monetary policy and endogenous financial crises," Working paper N. 991, BIS.

Boyarchenko, N. and L. Elias (2024): "The Global Credit Cycle," Staff Report N. 1094, Federal Reserve Bank of New York.

Brandão-Marques, L., G. Gelos, M. Narita, and E. Nier (2021): "Leaning against the wind: An empirical cost-benefit analysis," Discussion Paper No. 15693, CEPR.

Brunnermeier, M., D. Palia, K. A. Sastry, and C. A. Sims (2021): "Feedbacks: financial markets and economic activity," *American Economic Review*, 111, 1845–1879.

Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio (2014): "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems*, ed. by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Curran Associates, Inc., vol. 27.

Gourio, F., A. K. Kashyap, and J. W. Sim (2018): "The trade offs in leaning against the wind," *IMF Economic Review*, 66, 70–115.

Greenwood, R., S. G. Hanson, A. Shleifer, and J. A. Sørensen (2022): "Predictable financial crises," *The Journal of Finance*, 77, 863–921.

Grimm, M., Ò. Jordà, M. Schularick, and A. M. Taylor (2023): "Loose monetary policy and financial instability," WP 30958, National Bureau of Economic Research.

Kocaoglu, M., C. Snyder, A. G. Dimakis, and S. Vishwanath (2017): "CausalGAN: Learning Causal Implicit Generative Models with Adversarial Training," *CoRR*, abs/1709.02023.

Krishnamurthy, A. and T. Muir (2017): "How credit cycles across a financial crisis," Working Paper 23850, National Bureau of Economic Research.

Laubach, T. and J. C. Williams (2003): "Measuring the Natural Rate of Interest," *Review of Economics and Statistics*, 85, 1063–1070.

Liu, Y., Z. Qin, T. Wan, and Z. Luo (2018): "Auto-painter: Cartoon image generation from sketch by using conditional Wasserstein generative adversarial networks," *Neurocomputing*, 311, 78–87.

Mian, A., A. Sufi, and E. Verner (2017): "Household debt and business cycles worldwide," *The Quarterly Journal of Economics*, 132, 1755–1817.

Mirza, M. and S. Osindero (2014): "Conditional Generative Adversarial Nets," *CoRR*, abs/1411.1784.

Odena, A., C. Olah, and J. Shlens (2017): "Conditional Image Synthesis with Auxiliary Classifier GANs," in *Proceedings of the 34th International Conference on Machine Learning*, ed. by D. Precup and Y. W. Teh, PMLR, vol. 70 of *Proceedings of Machine Learning Research*, 2642–2651.

Rey, H. (2015): "Dilemma not trilemma: the global financial cycle and monetary policy independence," Working Paper 21162, National Bureau of Economic Research.

Richter, B., M. Schularick, and P. Wachtel (2021): "When to lean against the wind," *Journal of Money, Credit and Banking*, 53, 5–39.

Schularick, M., L. t. Steege, and F. Ward (2021): "Leaning against the wind and crisis risk," *American Economic Review: Insights*, 3, 199–214.

Schularick, M. and A. M. Taylor (2012): "Credit booms gone bust: monetary policy, leverage cycles, and financial crises, 1870–2008," *American Economic Review*, 102, 1029–1061.

Svensson, L. E. (2017): "Cost-benefit analysis of leaning against the wind," *Journal of Monetary Economics*, 90, 193–213.

Taylor, J. B. (2007): "Housing and monetary policy," Working paper N. 13682, National Bureau of Economic Research.

A Additional Results

A.1 Model Properties

Figure A.1. Model-Implied Macroeconomic Relations This figure presents the model-implied relation between inflation and $ggap$ (top row) along with the model-implied relation between the $rgap$ and $ggap$ or π (bottom row). Panel (a) shows $Q^{\pi}_{t|t-1}(c)$ for $c \in \{0.25, 0.5, 0.75\}$ where we fix all initial conditions $x_{t-1}, \ldots, x_{t-4}$ except $ggap_{t-1}$ which is shown on the x axis. All other variables are set to $ggap_{t-s} = rgap_{t-s} = fci_{t-s} = 0$, $\pi_{t-s} = 2$ for $s = 1, ..., 4$. In Panel (b) we again fix all initial conditions except that we vary $ggap_{t-1} = \ldots = ggap_{t-4}$ which is shown on the x axis. Panel (c) shows $Q^{rgap}_{t|t-1}(c)$ for $c \in \{0.25, 0.5, 0.75\}$ where we fix z_t and all initial conditions $x_{t-1}, \ldots, x_{t-4}$ (as in Panels (a) and (b)) except for $ggap_t = ggap_{t-1}$ which is shown on the x axis. Panel (d) conducts the same exercise except fixing $ggap$ and instead varying $\pi_t = \pi_{t-1}$.

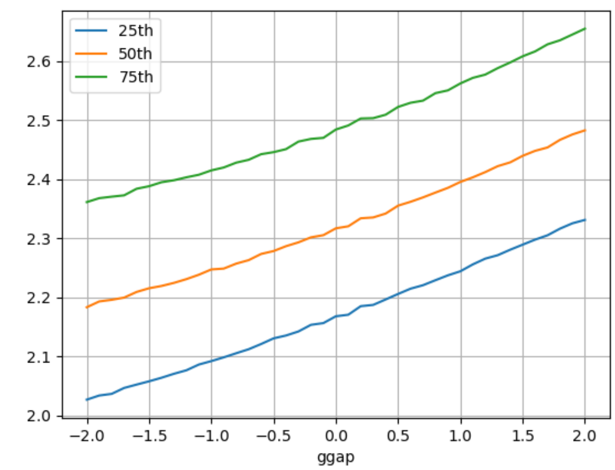

(a) Response of π_t to One Qtr. Change in $ggap$

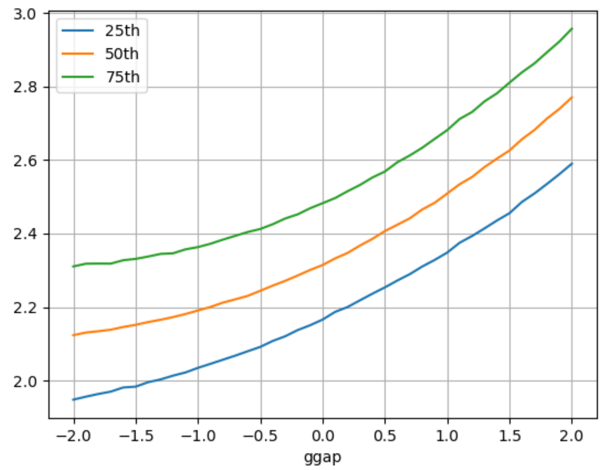

(b) Response of π_t to Four Qtrs. Change in $ggap$

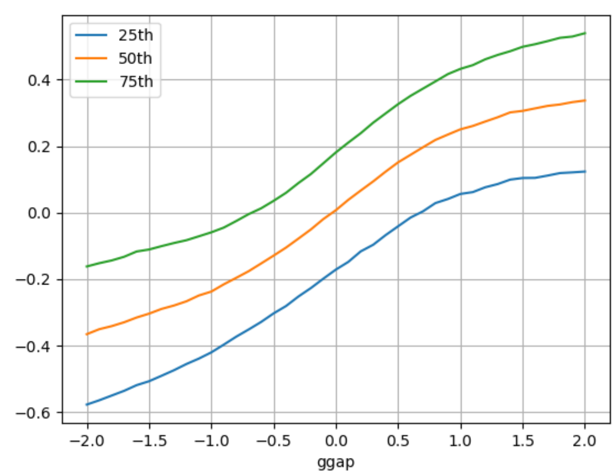

(c) Response of $rgap_t$ to Two Qtrs. Change in $ggap$

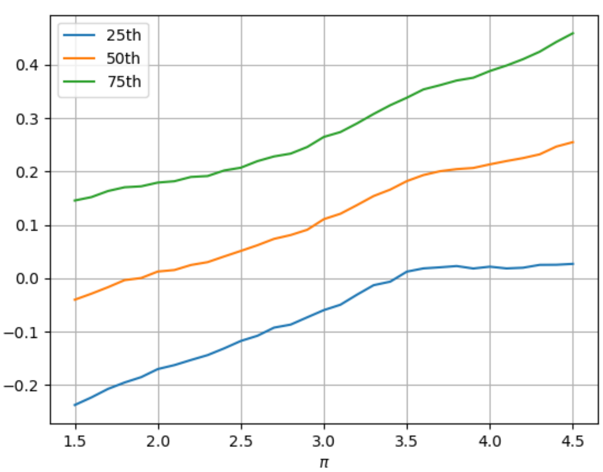

(d) Response of $rgap_t$ to Two Qtrs. Change in π

A.2 Role of Initial Conditions

Figure A.2 plots the time series of IRFs $\Delta Q^{v}_{t+h|t-1}(c)$ to a series of $rgap_t$ shocks that raise $rgap_t$ by 50bps per quarter for 4 quarters for the median ($c = 0.5$) and the 10th percentile ($c = 0.1$) of $ggap$. The near-term effect on the 10th percentile of the cumulative $ggap$ tends to be slightly positive immediately following recessions.

Figure A.2. Response to tighter monetary policy over time This figure presents the distribution IRFs $\Delta Q^{ggap}_{t+h|t-1}(c)$ to a series of $rgap_t$ shocks that raise $rgap_t$ by 50bps per quarter for 4 quarters for the median ($c = 0.5$) and the 10th percentile ($c = 0.1$) of $ggap$. Each observation corresponds to initial conditions that prevailed as of the quarter. We report the quantiles of the distribution of annualized cumulative $ggap_{t+h}$, e.g. the distribution of $\frac{4}{h+1}\sum_{k=0}^{h} ggap_{t+k}$. Each line corresponds to a different forecast horizon h. Grey shading denotes NBER recessions.

(a) Median　　**(b)** 10th percentile

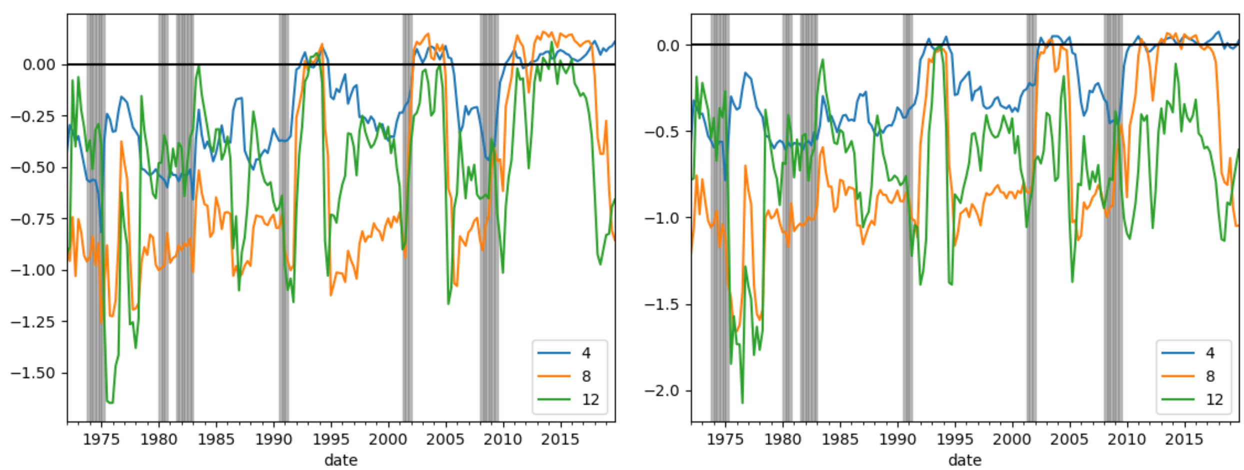

A.3 Alternative Specification

As discussed in the main text, we also consider an implementation which relies on more stringent regularization than our main results. In this auxiliary implementation of the WGAN we increase the generator dropout rate, increase the critic penalty term factor, and, for the estimation of the policy rule, reduce the max epoch. Figure A.3 below provides the same time series plots as in Figure A.2 but for this more regularized model.

Figure A.3. Response to tighter monetary policy over time (auxiliary WGAN implementation) This figure presents the distribution IRFs $\Delta Q^{v}_{t+h|t-1}(c)$ to a series of $rgap_t$ shocks that raise $rgap_t$ by 50bps per quarter for 4 quarters for the median ($c = 0.5$) and the 10th percentile ($c = 0.1$) of $ggap$. Each time series point corresponds to initial conditions as of that quarter. We report the difference quantiles of the distribution of annualized cumulative $ggap_{t+h}$, e.g. the distribution of $\frac{4}{h+1}\sum_{k=0}^{h} ggap_{t+k}$. Each line corresponds to a different forecast horizon h. NBER recession shadings are reported in grey.

(a) Median **(b)** 10th percentile

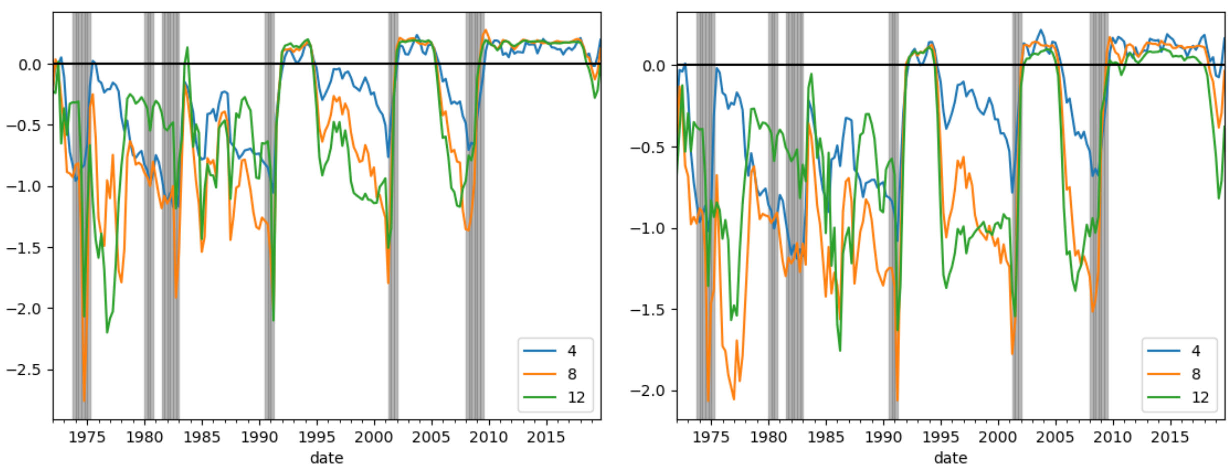

Figure A.4. Time Series of Variables. This figure presents time series plots for our four variables of interest over the common sample 1971Q1–2023Q4.

(a) *Chicago Fed NFCI* (fci_t)

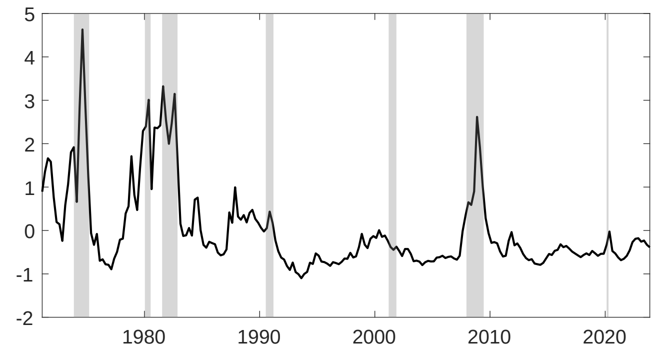

(b) *Real Rate Gap* ($rgap_t$)

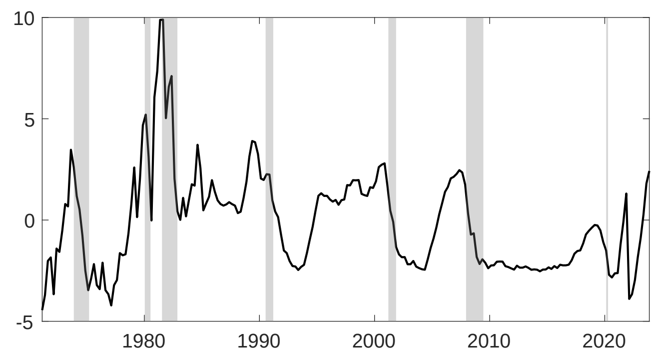

(c) *Real GDP Growth Gap* ($ggap_t$)

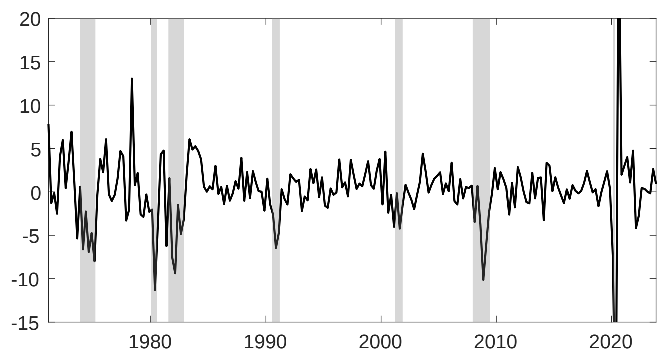

(d) *Core PCE Inflation* (π_t)