The Downside Risk Channel of Monetary Policy

Niklas Schmitz*

First Version: October 24, 2021

This Version: July 18, 2022

Abstract

This paper documents a new transmission channel of monetary policy shocks to equity premia: Monetary shocks can affect the probability of bad outcomes for future macroeconomic growth and hence support equity prices beyond any effect on expected mean growth rates. I estimate a monthly index of downside risks to consumption growth. A loosening in the monetary policy stance can significantly reduce consumption downside risks in crisis times. Increases in downside risk predict higher future equity returns, in line with the disaster risk hypothesis. Consumption downside risk predicts stock markets in the aggregate and across a wide range of industry portfolios.

*Faculty of Economics, University of Cambridge (email: ns751@cam.ac.uk). I thank my supervisor Chryssi Giannitsarou and my advisor Elisa Faraglia. I thank Rhys Bidder, Giancarlo Corsetti, Radu Cristea, Petra Geraats, Alexander Meyer-Gohde, Daniel Ostry, Charles Parry, Alba Patozi, Carolin Pflueger, Giovanni Ricco, as well as participants at the Cambridge Macroeconomic PhD Workshop, T2M 2022 Conference, CERF Lunchtime Talks, the IAAE Conference 2022, and the QCGBF Annual Conference 2022 for helpful comments. I thank Refet S. Gürkaynak, Marek Jarociński, and Jakob Berndt for sharing data with me. I thank CERF for the Cambridge Finance Best Student Paper Award, the IAAE for the Best PhD Student Paper Award, and the QCGBF for the Young Economist Prize. I acknowledge receipt of a student travel grant for the 2022 IAAE Conference.

1 Introduction

Much of the monetary policy communication during crisis times describes efforts to mitigate "downside risks" and avert bad outcomes.[1] If monetary policy can significantly reduce macroeconomic downside risks, its interventions have an important stabilization role for the economy. Changes in downside risk can also have important asset pricing implications: A higher risk of bad macroeconomic outcomes leads risk averse investors to demand higher future returns.[2]

This paper argues that the effect of monetary policy shocks on downside risks to macroeconomic growth can drive equity prices. Even if expectations about mean growth in the economy do not change dramatically following a monetary policy shock, stock prices can rise strongly if investors believe the policy action has mitigated downside risk. This could particularly explain sudden stock market rebounds following central bank interventions in crisis times. Monetary policy can have a powerful effect on equity markets via the downside risk channel that goes unseen by focusing on mean effects.

I study the downside risk channel empirically in three steps. First, I estimate a model-free measure of downside risk to aggregate consumption growth. Consumption most closely reflects the demand side of the economy and is intimately linked to equity returns in consumption-based asset pricing models. In contrast to GDP, consumption data is available at a monthly frequency, which allows for a detailed analysis of the link between downside risk and monetary policy as well as equity returns. Second, I study the effect of monetary policy shocks on downside risk to consumption using vector autoregressions. Third, I analyze the predictive ability of changes in downside risk for stock market returns.

The monthly index of downside risks to U.S. aggregate consumption growth is

[1] In a January 2009 article in *The Economist*, Olivier Blanchard famously explained how policymakers should respond to the financial crisis: "First and foremost, reduce uncertainty. Do so by removing tail risks, and perceptions of tail risks."

[2] This intuition is formalized in consumption-based asset pricing models with disaster risk, e.g. Barro (2006), Gabaix (2012), Gourio (2012), and Nakamura et al. (2013).

constructed based on an approximation of the conditional distribution of expected U.S. consumption growth over the next year. To accurately estimate the growth distribution, I pre-select variables from a large set of predictors and aggregate their information via factor estimation to estimate conditional quantiles of the growth distribution. Then, I fit a flexible skewed t-distribution over the estimated quantiles to approximate the full conditional distribution of future consumption growth. Based on the conditional distribution, I construct a measure for downside risks to consumption growth. Loosely speaking, higher values in the downside risk index indicate a higher probability of below-average outcomes for aggregate consumption growth over the next year. Analogously, I estimate an index of upside potential to consumption growth, which reflects the probability of above-average future growth outcomes.

The index of consumption downside risk is the first contribution of my paper. Downside risk exhibits substantial variation over time. Increases in consumption downside risk coincide with major U.S. recessions. Downside risk and upside potential have positive correlation but overall different time series behaviour. Downside risk is more volatile over time and rises more strongly around recessions. During recessions, the downside risk index comoves closely with common measures of economic risk or uncertainty such as the VIX, the Economic Policy Uncertainty Index (Baker et al. (2016)), the Jurado et al. (2015) macroeconomic uncertainty index, or credit spreads. However, the series behaves markedly different in recovery periods, episodes of pure financial stress, and non-crisis times.

Overall, the downside risk index contains information distinct from other indicators of risk, uncertainty, or activity. The downside risk index does not conflate risk and risk aversion into a composite index but provides a clean measure of macroeconomic risk on the consumer side. In contrast to existing risk measures based on financial market data (e.g. Kelly & Jiang (2014)), it is not affected by pure financial stress that does not spill over into the macroeconomy. Further, the index allows a distinction between upside and downside risks, which is an improvement over symmetric measures of risk or uncertainty.

Importantly, the downside risk index is not designed to focus on extreme events in the tail, but instead captures the much more frequent potential for good versus bad consumption growth rates. This makes the downside risk index easier to estimate and avoids any critique about the difficulty of estimating extreme tails with small samples of macroeconomic data. An important result of this paper is that an asymmetry between upside and downside risk exists, and that the transmission of monetary policy to equity premia can be explained without having to concentrate on worst- and best-case scenarios.

Given the index for consumption downside risk, I analyze the effects of monetary policy shocks on downside risk. I estimate a vector autoregression (VAR) including the downside risk index, a monetary policy indicator, and a set of controls. The analysis is complicated by a simultaneity issue: While monetary policy can affect consumption downside risk, policymakers may also take the most recent developments in downside risk into account when making their policy decision.[3] Imposing a naive recursive structure on the system via short-run restrictions would not allow for this simultaneous interaction. Instead, I identify the system using an external instruments approach as in Gertler & Karadi (2015). This method uses high-frequency changes in Fed funds futures around FOMC announcements to identify monetary policy surprises. The approach yields a sufficient number of restrictions to estimate the impulse responses to a policy shock without restricting any contemporaneous relations between the variables. The instrument is adjusted to remove Fed information effects using Jarociński & Karadi (2020)'s "poor man's sign restrictions".

The VAR analysis indicates a significant effect of monetary policy on downside

[3] For example, Fed Chairman Powell noted during the November 2020 FOMC press conference: "In terms of the tail risks, I mean, I think clearly the tail risks we were worried about have, have subsided. And, you know, we were worried about very negative potential outcomes. ... We don't just look at the most likely case. We, we ask 'How do you make policy in light of the risks?' and often it's downside risks ... " (FOMC (2020), p.17).

risk to consumption. Following a 25bp surprise loosening in the monetary policy stance, downside risk falls on impact and keeps decreasing, experiencing the strongest decrease of around one third of a standard deviation after eight months. The effect on downside risk reverts afterwards. A threshold VAR analysis confirms that these effects are driven by the effect of monetary policy interventions on consumption risk during recession periods. In contrast, monetary policy has little to no measurable impact on downside risk during normal times. Further, monetary policy has no significant effect on upside potential of consumption growth. This suggests a loosening in the monetary policy stance is a successful crisis mitigation tool but has not reduced risks to economic growth during normal times over the sample period. The effect of monetary policy on downside risk remains after controlling for measures of risk aversion and macroeconomic uncertainty, supporting the idea of a distinct downside risk channel.[4]

Lastly, I study the relation of downside risk and equity returns using predictive regressions. Changes in the index of consumption downside risk have strong predictive power for aggregate U.S. equity returns over a horizon of three to twelve months. A one standard deviation rise in the downside risk index predicts equity returns to increase by about 1.5 percentage points over the next three months, and by about 3 percentage points over the next six months. These results are strongly driven by the predictive power of downside risk for returns in crisis times. The coefficients can increase severely to 4 (three months) to 8.5 (six months) percentage points during recessions, in line with strong stock market volatility of such periods. The predicted effects are in line with the disaster risk hypothesis: A rise in downside risk causes risk averse investors to require higher returns to hold equities. In contrast, downside risk shows no predictive ability during normal times. These results are robust to including controls for risk aversion and uncertainty in financial markets, again

[4]In the Online Appendix, I show that including controls mean consumption, median consumption, or the variance of consumption does not change the results. Downside risk captures information about future consumption growth prospects that is different from the information content of any of these measures.

suggesting the existence of a downside risk channel distinct from overall uncertainty and risk aversion concerns.

Downside risk not only predicts stock returns in the aggregate, but also across a wide range of industry portfolios. In all cases, the predictive ability of downside risk is concentrated in U.S. recessions. Industries which are conventionally thought of as being procylical, for example manufacturing or durable goods consumption, show a stronger sensitivity to changes in downside risk than less procyclical industries such as utilities or healthcare. The differences across industries are quantitatively large. A one standard deviation rise in consumption downside risk predicts an increase in six-month returns by 1.9 percentage points for the non-durable goods sector, but an increase by 5.6 percentage points for the durable goods industry.

The rest of the paper is structured as follows. Section 2 reviews the related literature. I describe the methodology to construct the downside risk measure in Section 3 and present results on its time series properties. Section 4 analyzes the effect of monetary policy surprises on downside risk using vector autoregressions. The predictive regressions to study the relation between downside risk and equity returns are in Section 5. Section 6 concludes.

2 Related Literature

The theoretical motivation of the downside risk channel and the focus on risks to aggregate consumption is rooted in the vast literature on consumption-based asset pricing, see Cochrane (2017) for a review. More specifically, the downside risk channel is connected to ideas from the disaster risk literature, which posits that the Mehra-Prescott equity premium puzzle can be explained by the risk of rare but very large declines in stock prices. This idea was originally proposed by Rietz (1988) and subsequently reconsidered by Barro (2006), who estimates the probability and severity of rare disasters from a panel of 20th century macroeconomic crises. Gabaix (2012), Gourio (2012) and Wachter (2013) extend this literature by considering time-varying disaster risk in different model frameworks. These papers give rise to the disaster risk hypothesis: An increase in the expected probability of a bad economic outcome in the near future should decrease equity prices as investors require larger compensation to hold risky assets. Nakamura et al. (2013) consider consumption disasters with partial recoveries, finding smaller equity premia as a significant share of the initial drop in consumption in a crisis is subsequently reversed.

In contrast to the disaster risk literature, my paper does not consider extreme worst-case outcomes, but instead focuses on “moderate”disasters. This is in line with recent empirical evidence (Beason & Schreindorfer (2022)), which suggests that investors are mostly concerned with large but not huge declines in equity prices (between -10 and -30 percent). Based on this evidence, disaster risk models rely too much on events in the extreme left tail to achieve a realistic equity premium.

The empirical literature on the link between risk and equity premia finds that different measures of (tail) risks or risk aversion can predict equity returns. Bollerslev et al. (2009) and Bekaert & Hoerova (2014) demonstrate that the variance risk premium – a measure of risk aversion – has predictive power for future stock returns at short horizons.[5] Kelly & Jiang (2014) show that tail risks can predict excess stock market returns, where higher tail risk predicts higher stock returns at horizons between one month and five years. Bollerslev & Todorov (2011*b*) and Bollerslev et al. (2015) find that fear of negative tail events accounts for up to five percentage points of the equity premium and explains much of the predictive ability of the variance risk premium.

The existing literature usually derives risk (aversion) measures from financial market data such as options, high-frequency stock returns, or individual firm returns. This makes it hard to link the risk measure to macroeconomic fundamentals. Barro & Liao (2021) is a recent attempt to build a consumption-based option pricing

[5] The variance risk premium is obtained from a decomposition of the VIX index into proxies for risk and risk aversion, see for example Carr & Wu (2009). The VIX measures 30-day option-implied expected volatility on the S&P 500 stock market index.

model for far out-of-the-money options. Nevertheless, the evidence of Backus et al. (2011) points towards difficulties in reconciling the empirical evidence from option prices with existing consumption-based asset pricing models. To avoid concerns about misspecification, my paper takes a different route and starts from a model-free measure of consumption risk, which is then linked to equity premia. While the downside risk index does not focus on tail events and hence avoids the usual estimation problems associated with rare events, I find that it only exhibits predictive power for future stock returns during crisis times, when downside risks are the strongest. This lends support to the disaster risk hypothesis.

This paper relates to a growing literature on the effects of monetary policy on risks to economic fundamentals or equity returns. Some papers indicate that monetary policy decisions can reduce risks in equity markets (Bekaert et al. (2013), Hattori et al. (2016), Cortes et al. (2020)) or corporate bond markets (Haddad et al. (2022)). The main difference to these papers is that my analysis of the effect of monetary policy shocks on downside risks uses a risk measure based on macroeconomic fundamentals instead of financial market data, which also allows to cover a longer sample. Duprey & Ueberfeldt (2020) and Ajello & Pike (2022) study the effect of monetary policy on GDP tail risk and argue that loose policy reduces risks in the short run but increases them over the medium term. In contrast, my work studies the downside risk channel in the transmission of monetary policy shocks to equity markets at a short horizon.[6]

To measure (tail) risks in equity markets, the finance literature has relied on using high-frequency returns (Bollerslev & Todorov (2011*a*)), firm-level stock returns (Kelly & Jiang (2014)), or option prices (Backus et al. (2011)). While using financial market data to construct risk measures directly reflects how investors value these risks and may be most closely related to asset pricing, this approach also has disadvantages for

[6]Another difference between the papers is the empirical methodology: Duprey & Ueberfeldt (2020) and Ajello & Pike (2022) use quantile regression and vector autoregression but their approach rests on using coefficients estimated from quantile regressions to infer quantile impulse responses from the mean impulse responses estimated in a VAR. In contrast, my paper stresses the importance of using a very large dataset to estimate downside risks accurately, before using the downside risk index as an input directly in a VAR.

my application. First, high-frequency return data is often not available before the late 1990s, which restricts the sample size and complicates studying the behaviour of tail risk in a historical context. Second, the relation of market-based risk measures to macroeconomic fundamentals is not necessarily clear, which makes it harder to relate these measures to the intuition of consumption-based asset pricing models. The goal of this paper is to construct a pure measure of macroeconomic risk, whereas market-based measures my be driven by liquidity concerns or idiosyncratic firm effects.

The approach of this paper is inspired by a recent literature in macroeconomics, which aims to estimate quantiles of the underlying distribution of the variable of interest and then construct risk measures based on the quantile estimates. Giglio et al. (2016) use partial quantile regression to study systemic risk indices in their ability to forecast negative shocks to macroeconomic outcomes out of sample. Adrian et al. (2019) estimate the conditional distribution of GDP growth using a two-step procedure. First, they predict certain quantiles of GDP growth using quantile regression. Then, they recover the full distribution at each point in time by fitting a skewed t-distribution across these quantiles. The authors find asymmetries between upside and downside risks. My paper combines methodologies from both of these works to improve the accuracy of the downside risk estimates. I am not aware of any papers using this methodology to obtain a measure of consumption risk for the purpose of studying the predictions of consumption-based asset pricing models.

3 An Index of Consumption Downside Risk

This section introduces expected downside risk to U.S. aggregate consumption growth as a measure of macroeconomic risk. The index for consumption downside risk is constructed based on the entire conditional distribution of expected future consumption growth at a given point in time. The methodology follows the three-step procedure of Adrian et al. (2019) to estimate the consumption growth distribution: Given time series data on consumption growth and a set of predictors, first forecast several quantiles of the consumption growth distribution at each point in time. In a second step, fit a flexible density function across the estimated quantiles to approximate the full distribution at each point in time. Third, based on the conditional distribution, construct a time series for the downside risk index.

Working with consumption growth data has several advantages. Since consumption growth data is available at a monthly frequency, the downside risk index is monthly as well. This facilitates the analysis of the effect of monetary policy shocks on downside risk. It also allows to study the relation between downside risk and stock returns at the monthly frequency, which yields more detailed insights than quarterly regressions. Further, consumption risk directly captures risks to the demand side of the economy, in contrast to many other measures of macroeconomic or financial activity. Lastly, consumption-based asset pricing models suggest a close link between consumption and asset prices. From this perspective, studying consumption risk is the most natural way to identify a downside risk channel of monetary policy.

3.1 Forecasting U.S. consumption growth

Starting with a description of the forecasting model, let c_t be the aggregate real consumption growth in the U.S. and denote the k-dimensional vector of predictors at time t by x_t. An out-of-sample prediction of a quantile $Q(\theta)$ of consumption growth for time $t+h$ via quantile regression with data available at time t yields the following (consistent) estimate for the regression coefficients:

$$\hat{\beta}_\theta = \underset{\beta_\theta \in R^k}{argmin} \sum_{s=h+1}^{t} \left(\theta \cdot \mathbb{I}_{(c_s \geq x'_{s-h}\beta_\theta)} |c_s - x'_{s-h}\beta_\theta| + (1-\theta) \cdot \mathbb{I}_{(c_s < x'_{s-h}\beta_\theta)} |c_s - x'_{s-h}\beta_\theta| \right) \tag{1}$$

$$= \underset{\beta_\theta \in R^k}{argmin} \, \widehat{\mathcal{C}}_\theta \left(\beta_\theta\right), \tag{2}$$

where $\mathbb{I}$ denotes the indicator function and $\widehat{\mathcal{C}}_\theta$ is the sample analog of the criterion function which the quantile regression estimator of β aims to minimize.

In contrast to OLS, the coefficient estimates vary across estimated quantiles $Q(\theta)$. This is because the objective function used in quantile regression differs from the standard OLS case. First, we minimize absolute deviations between the predicted value and the actual value instead of squared deviations. Second, the absolute deviations are weighted depending on the quantile θ we are estimating. More weight is applied to errors close to the quantile of interest.

Given the estimated coefficients $\hat{\beta}_\theta$ and the set of predictors at time t, the predicted value for a θ-quantile of real consumption growth c at time $t+h$ is:

$$\hat{Q}_{c_{t+h}|x_t}(\theta|x_t) = {x_t}'\hat{\beta}_\theta. \tag{3}$$

This yields a time series of quantile forecasts for each θ_j-th quantile, where θ_j lies in a set of targeted quantile indices $\Theta = \{\theta_1, \theta_2, ..., \theta_J\}$.

To obtain good predictions for the consumption growth quantiles, it is desirable to combine a wide range of information about various parts of the economy to increase the forecasting performance. The set of predictors used here comes from the FRED-MD database by McCracken & Ng (2016). The dataset contains 128 macroeconomic time series at a monthly frequency starting in January 1959. The variables cover information about output and income, the labor market, consumption and orders, inventories, money and credit, interest and exchange rates, prices, and the stock market. Apart from comprising a large set of possibly useful predictors, the dataset is also updated continuously to take care of changing variable definitions or data revisions. I use the April 2021 vintage, as most data for December 2020 only becomes available with a lag or may be subject to revisions.[7]

[7]The values for April 2020 were missing for the 3-month AA Financial Commercial Paper Rate (CP3Mx) and 3-month Commercial Paper Rate minus Fed Funds rate (COMPAPFFx). This is because there was insufficient trading data in that month. I replaced the values by linearly interpolating between March and May 2020.

However, the standard quantile regression estimator in equation (1) may no longer be consistent if the number of regressors is large relative to the sample size, see e.g. Belloni & Chernozhukov (2011). Including a large set of predictors may also overfit the data and weaken the forecasting performance. To address these problems, I use factor estimation to reduce the number of predictors while keeping their information content rich.

The construction of the independent variables proceeds in two steps. First, an algorithm pre-selects a subset of regressors from the FRED-MD dataset. Since some variables in the dataset are highly correlated (for example, the sample includes 13 different measures of industrial production), directly estimating factors from the entire set of predictors may overemphasize certain features of the data. In addition, shrinking the number of predictors each period provides the factor estimation procedure with more flexibility to adapt to changes in parameters over time (Bai & Ng (2008*a*), p. 314). Pre-selecting a subset of variables from the entire sample before estimating the factors can therefore improve forecasting performance.

I pre-select a subset of predictors using ℓ_1-penalized quantile regression (Belloni & Chernozhukov (2011)). All predictor series are standardized to have zero mean and unit variance. For a given quantile index θ, the ℓ_1-penalized quantile regression coefficients $\tilde{\beta}_\theta$ satisfy

$$\tilde{\beta}_\theta = \underset{\beta_\theta \in R^k}{argmin} \quad \widehat{\mathcal{C}}_\theta \left(\beta_\theta\right) + \frac{\lambda\sqrt{\theta(1-\theta)}}{n} \sum_{i=1}^{k} \hat{\sigma}_i |\beta_i|, \tag{4}$$

where $\widehat{\mathcal{C}}_\theta$ is the sample analog of the criterion function and n is the sample size. $\hat{\sigma}_i$ is the sample variance of predictor x_i. The penalty term $\lambda\sqrt{\theta(1-\theta)}$ depends on the quantile index θ and a penalty level λ.[8] ℓ_1-penalized quantile regression applies the logic of LASSO regression to a quantile regression setting. The penalty restricts the sum of absolute values of the coefficients to be below a fixed value, thus shrinking the

[8]The penalty level is chosen as described in Belloni & Chernozhukov (2011), p.86.

coefficient values for variables with little explanatory power towards zero. For each quantile in the targeted set Θ, the algorithm selects the $\tilde{k} = 40$ predictors with the largest absolute coefficients. The number of selected predictors is based on results by Boivin & Ng (2006), who find that factors constructed from 40 pre-screened variables can yield better forecasting performance than factors estimated from a large dataset with over 100 macroeconomic time series.[9]

In a second step, the information from the subset of regressors is condensed into a factor estimate. Given the subset of $\tilde{k}$ predictors, I combine those predictors using partial quantile regression (Giglio et al. (2016)). This approach constructs a single factor to predict the dependent variable. In constructing the factor, the individual predictors are weighted depending on their predictive power for the dependent variable. This contrasts partial quantile regression (PQR) from factor estimation via principal components, which aims to construct factors that describe the most variation in the set of predictors. Instead, PQR aims to construct a factor with the best predictive power for the variable of interest.

Assume that a scalar f_t contains all relevant information for a conditional quantile of consumption growth. Constructing a factor for each quantile via PQR follows a two-step procedure. First, run univariate quantile regressions of c_{t+h} on each predictor $x_{i,t}$ for $i = 1, ..., \tilde{k}$. The slope estimates are $\hat{\gamma}_{\theta,i}$. Second, compute the cross-sectional covariance of $x_{i,t}$ with i's first stage slope estimate in each period t: $\hat{f}_{\theta,t} = \sum_{i=1}^{\tilde{k}} (x_{i,t} - \bar{x}_{i,t}) \left(\hat{\gamma}_{\theta,i} - \bar{\hat{\gamma}}_{\theta,i}\right)$, where $\bar{\cdot}$ denotes the sample average. The covariance estimate $\hat{f}_{\theta,t}$ serves as an estimate of the latent factor realization. It is a weighted average of individual predictors with weights determined by their predictive strength for $Q_{c_{t+h}}(\theta)$ from the first step.

Given these factors, we obtain the predicted quantiles from the quantile regressions of c_{t+h} as given in equation (1), where the predictors are the second stage factor estimates $\hat{f}_{\theta,t}$. Appendix A contains further information on an out-of-sample

[9] I generally find that more than 40 predictors have non-zero coefficients, hence selecting 40 variables does not include any predictors which would otherwise have been excluded by the ℓ_1-penalty.

forecasting comparison against nineteen alternative factor model specifications. The results demonstrate that the forecasting approach presented here strongly outperforms its competitors.

The goal is to forecast the 5, 25, 50, 75 and 95 percent quantiles of consumption growth over the next twelve months at a monthly frequency.[10] For a given month, I define year-over-year consumption growth as the percent change in U.S. aggregate real personal consumption expenditures (excluding durables) over the past twelve months. The sample period covers January 1960 to December 2020. At a given point in time, I pre-select the predictors and apply PQR to forecast out-of-sample with a horizon of 12 months. The first out-of-sample prediction is for January 1980, and I extend the sample each month with the new data available.

Figure 1 plots the predicted median (black) as well as $5, 25, 75$ and 95% quantiles (grey) versus the realized year-over-year growth rate of U.S. aggregate consumption (blue dashed). Precisely, at a given point in time, the blue dashed line is the realized growth rate of consumption over the *past* twelve months. The values for the quantiles are the values predicted over the *next* twelve months, using the information available until the given point in time. While one-year ahead forecasts are generally hard, the model does a decent job of capturing the main business cycle movements of the consumption growth rate. Most realizations of the growth rate fall within the predicted 5 to 95% range. The Covid crisis was virtually impossible to forecast one year in advance, but given the new information arriving in March and April of 2020 the model quickly adjust its predictions for future consumption growth, albeit with a slight lag.

[10] Predictive quantile regressions can lead to quantile crossings (e.g. the 5% quantile being predicted above the 10% quantile). To address this issue, I estimate a fine grid of quantiles with $\theta \in \{0.05, 0.06, ..., 0.94, 0.95\}$ and use the Chernozhukov et al. (2010) rearrangement method to ensure monotonicity of the quantile function. In evaluating the forecasting performance, I only focus on the five targeted quantiles. To save computation time, I pre-select the predictors via ℓ_1-penalized regression only for $\theta \in \{0.05, 0.15, 0.35, 0.5, 0.65, 0.85, 0.95\}$.

Figure 1: Predicted Quantiles vs Realized Consumption Growth

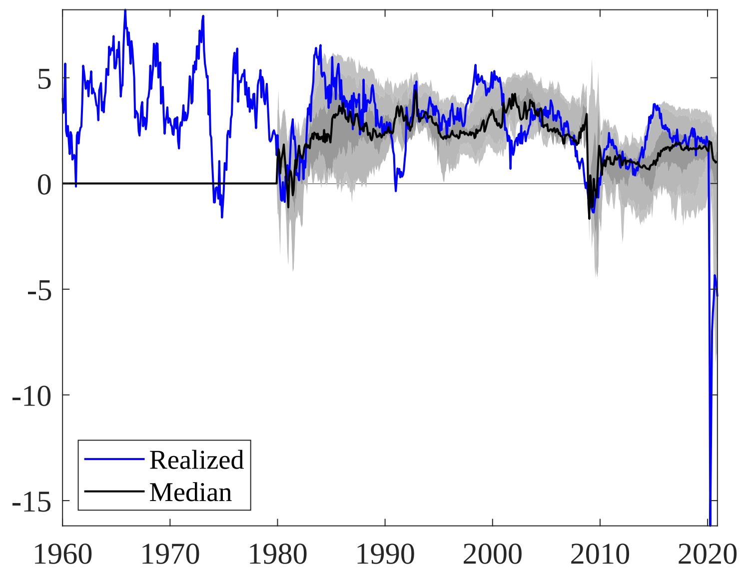

Note: The blue dashed line is the realized year-over-year rate of consumption growth. The black line is the estimated median. The grey areas correspond to the quantiles estimated for 5, 25, 75, and 95%. The sample period is 1:1980-12:2020.

3.2 Approximating the consumption growth distribution

Given the estimated quantiles, the goal is to estimate the downside risk to consumption growth based on the full conditional distribution of real consumption growth for each point in time. I approximate the conditional distribution by fitting a skewed t-distribution by Azzalini & Capitanio (2003) across the targeted quantiles in a given period:

$$f(c;\mu,\sigma,\alpha,\nu) = \frac{2}{\sigma} t\left(\frac{c-\mu}{\sigma};\nu\right) T\left(\alpha\frac{c-\mu}{\sigma}\sqrt{\frac{\nu+1}{\nu+\left(\frac{c-\mu}{\sigma}\right)^2}};\nu+1\right), \tag{5}$$

where $t(\cdot)$ and $T(\cdot)$ are the probability density function and cumulative distribution function of the Student-t distribution, respectively. Denote the PDF of the skewed t-distribution as $f(\cdot)$ and its CDF as $F(\cdot)$.

Figure 2: Conditional Distribution of Consumption Growth

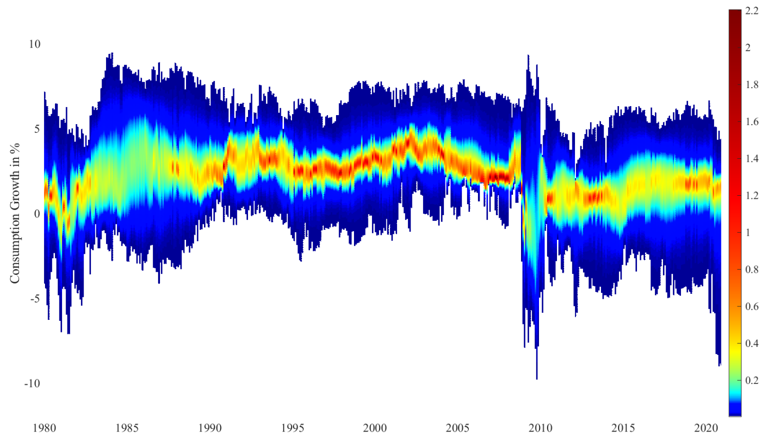

Note: The figure plots the density of the skewed t-distribution fitted across the estimated quantiles of consumption growth for each month. Low density values have a blue color, high values are red. The color scale is nonlinear to reflect that only few density values lie above one. The sample period is 1:1980-12:2020.

Figure 3: Moments of Conditional Consumption Growth Distribution

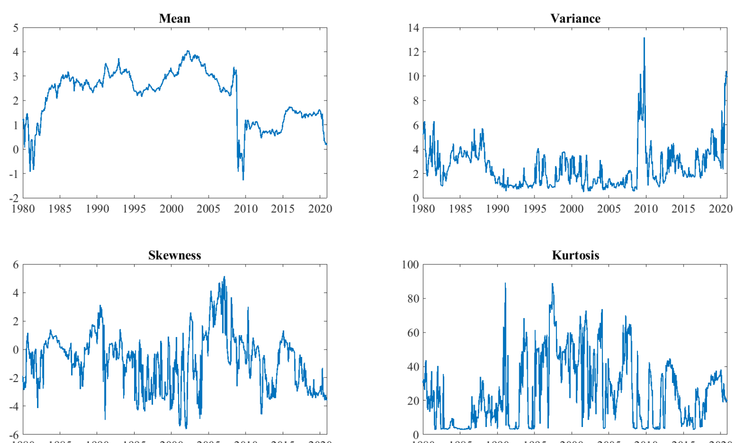

Note: The moments are derived from the fitted skewed t-distribution. The sample period is 1:1980-12:2020.

The skewed t-distribution allows to capture key properties of the consumption growth distribution without imposing too much structure on the data. It is governed by the parameters μ (mean), σ (variance), α (shape) and ν (thickness). As in the case of the standard t-distribution, we can account for changes in the location, scale, and size of the tails of the distribution. However, the standard t-distribution is symmetric. In contrast, the skewed t-distribution can capture asymmetries by using the cumulative distribution function to shape the probability density function. The skewed t-distribution can therefore capture different behaviour of upper and lower quantiles, which can reflect differences between risks on the upside and the downside of the conditional distribution. The shape parameter α governs to which extent the CDF is used to skew the PDF. For $\alpha = 0$, the distribution is symmetric. For $\nu \to \infty$, the skewed t-distribution converges to a skewed Gaussian distribution.

The skewed t-distribution in a period t is fitted by choosing $\{\mu_t, \sigma_t, \alpha_t, \nu_t\}$ to minimize the sum of squared differences between the estimated quantiles for $\{\theta_1, \theta_2, ..., \theta_J\}$ and the values implied by the skewed t-distribution:

$$\{\hat{\mu}_t, \hat{\sigma}_t, \hat{\alpha}_t, \hat{\nu}_t\} = \underset{\mu,\sigma,\alpha,\nu}{argmin} \sum_{\theta} \left(\hat{Q}_{c_{t+h}|x_t}(\theta|x_t) - F^{-1}(\theta; \mu, \sigma, \alpha, \nu) \right)^2 . \tag{6}$$

Figure 2 shows the conditional distribution of consumption growth over time. The color scale indicates the level of the density curve at a given point in time. Blue colours indicate low values, and red indicates high density. The distribution exhibits substantial variation over time. Upper quantiles are more stable over time than lower quantiles, suggesting that downside risks are varying more over time than upside risk. Especially recessions can be characterized by strong asymmetries between upper and lower quantiles, whereas the conditional distribution is more symmetric during normal times. The Volcker disinflation of the early 1980s and the 2008 financial crisis show decreases in the median and long left tails in the conditional distribution. The distribution is markedly slimmer during the 1990s and early 2000s, reflecting the Great Moderation. After the financial crisis, the growth distribution keeps a lower mean and larger variance, in line with lower growth prospects after the crisis. The Covid recession is associated with long left tails but only a moderate shift in the median of the expected distribution: While tail risks were perceived to be large, the most likely outcome for growth over the next year was not expected to be as severe as for the Great Recession.

Figure 2 also stresses the importance of using a flexible density function to capture the complex changes in the conditional distribution function. While crisis periods are often characterized by a lower mean and higher variance, skewness and kurtosis adjust in non-trivial ways. Figure 3 plots the first four moments of the conditional consumption growth distribution over time. The behaviours of mean and variance are easy to interpret. The mean of the consumption growth distribution falls sharply during the early 1980s, the recession of 2008, and the Covid shock. Following the global financial crisis, the mean does not fully recover to past levels, indicating slower growth dynamics after the crisis. Recessions are associated with increases in the variance of consumption growth. Particularly the global financial crisis and the Covid shock show strong spikes in variance. The time series for skewness and kurtosis are harder to interpret and we cannot identify a clear pattern related to the business cycle. This in line with Plagborg-Møller et al. (2020), who make similar findings for the GDP growth distribution.

3.3 Estimating the downside risk measure

Given the conditional distribution of consumption growth, I estimate downside risk taking the entire probability mass below the median into account. This poses an advantage over value at risk approaches, which are constrained to a certain quantile (e.g. 5%) and thus only consider a certain point on the distribution. As in Adrian et al. (2019), I use relative entropy to describe downside and upside risks to the economy.

Relative entropy defines downside risk as the excess probability mass below a certain π-quantile of the conditional distribution relative to the probability mass that the unconditional density $\hat{g}_{c_{t+h}}(c)$ assigns to those same values. Relative entropy therefore summarizes the probability that the conditional distribution assigns to

Figure 4: A Measure of Consumption Downside Risk

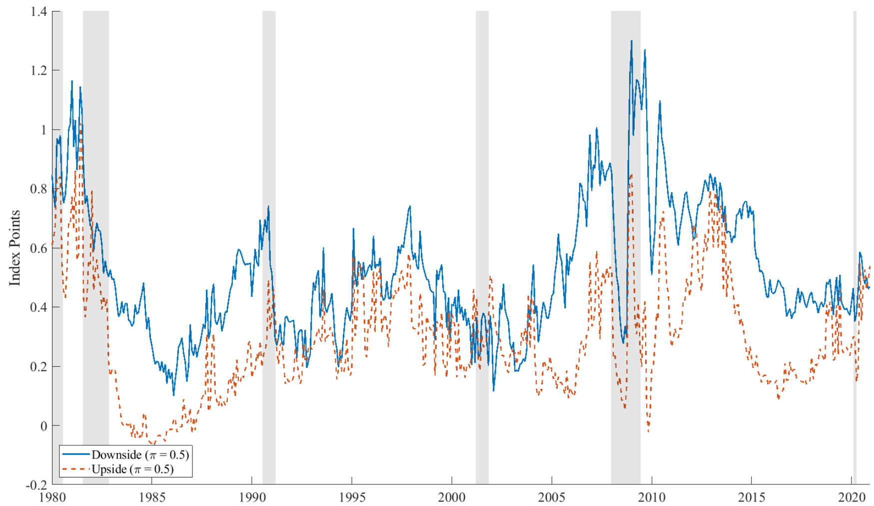

Note: The blue solid line is the downside entropy of the predicted distribution of consumption growth over the next twelve months relative to the unconditional distribution. The red dashed line shows the upside entropy. The shaded grey areas indicate NBER recessions. The sample period is 1:1980-12:2020.

"downside events" in relative terms. If the conditional distribution assigns a high probability to its low realizations of the consumption growth rate relative to the probability that the unconditional distribution assigns to those values, relative entropy is high. Skewness in the unconditional distribution affects relative entropy as it changes the difference in probability masses that the conditional and unconditional distribution attribute to downside events.

Relative entropy is a one-sided measure and can be computed for both the downside and the upside. *Downside entropy* is:

$$\mathcal{L}_t^D(\pi) = -\int_{-\infty}^{\hat{F}^{-1}_{c_{t+h}|x_t}(\pi|x_t)} \left(log\hat{g}_{c_{t+h}}(c) - log\hat{f}_{c_{t+h}|x_t}(c|x_t) \right) \hat{f}_{c_{t+h}|x_t}(c|x_t)dc. \tag{7}$$

Downside entropy measures the probability mass starting from a quantile with index π all the way into the left tail of the conditional distribution, relative to the probability mass that the unconditional distribution assigns to those values. During a recession, we should expect stronger downside risks to growth reflected in more probability mass in the left side of the conditional distribution, which raises downside entropy. The analog for the right side is called *upside entropy*, where we integrate from a quantile with index π to infinity:

$$\mathcal{L}_t^U(\pi) = -\int_{\hat{F}^{-1}_{c_{t+h}|x_t}(\pi|x_t)}^{\infty} \left(log\hat{g}_{c_{t+h}}(c) - log\hat{f}_{c_{t+h}|x_t}(c|x_t)\right)\hat{f}_{c_{t+h}|x_t}(c|x_t)dc. \tag{8}$$

For upside and downside risk, I choose a value of $\pi = 0.5$. Downside entropy then describes the excess probability in the conditional distribution below the median relative to the unconditional distribution. This is a natural choice to study asymmetries in downside versus upside risk as it splits the distribution of possible outcomes for consumption growth exactly in half.[11]

For the purpose of this paper, constructing the index of consumption downside risk based on the conditional distribution of consumption growth has several advantages over existing measures of risk or uncertainty. First, it does not conflate risk and risk aversion into a composite index. Instead, it clearly identifies risk without any risk aversion component. This stands in contrast to some newspaper-based uncertainty indices (e.g. Baker et al. (2016)) or the widely used stock market volatility index VIX. Second, it allows to distinguish between upside and downside risk − in contrast to symmetric measures such as the aforementioned uncertainty indices. Third, unlike risk measures constructed from financial returns, the index of consumption downside risk isolates risk associated with a macroeconomic variable. Consumption downside risk reflects risks to the consumer side of the macroeconomy and allows to distinguish macroeconomic recessions from pure financial stress. The downside risk index has a clear interpretation in the context of a consumption-based asset pricing model,

[11]Downside risk is different from negative tail risk as it considers all the probability mass below the median instead of focusing on the extreme tail. Estimating extreme tails of macroeconomic variables suffers from large estimation uncertainty since macroeconomic disasters are rare. In contrast, the measure of downside risk is easier to compute and not reliant on extreme properties of the consumption growth distribution.

whereas market-based risk measures may be affected by idiosyncratic firm effects (Kelly & Jiang (2014)) or are based on high-frequency market movements with unclear relations to macroeconomic fundamentals (Bollerslev & Todorov (2011*a*)). Fourth, since the downside risk index does not rely on estimates for the extreme tails, it avoids common issues of high estimation uncertainty for rare events with limited macroeconomic data available.[12]

The time series behaviour of the downside risk measure is shown in Figure 4 as the blue solid line. Downside risk was elevated during the Volcker disinflation period, then declined during the 1980s. The index peaks again in September 1990, in the middle of the next U.S. recession, and reaches a trough with the end of the recession period. The subsequent peak in consumption downside risk coincides with the height of the Asian financial crisis of 1997. Downside risk did not rise considerably during the burst of the dot-com bubble and the following recession in the early 2000s. While stock markets suffered heavy losses, the macroeconomic implications were relatively gentle. In contrast, the global financial crisis of 2008/2009 had severe macroeconomic implications and is associated with an all-time high in the downside risk index. The index provides a granular picture of the Great Recession: Downside risk started to increase in the mid-2000s and experienced a first peak in February 2007, well before the recession officially started. The index then sharply declines in line with the Fed interest rate cuts and Bush tax stimulus of the time. Following the collapse of Lehman Brothers in September 2008, the measure starts a steep climb and peaks in January 2009. After the Great Recession, downside risk remains elevated the start of 2015, before declining even further. This decline follows the end of the Fed's large-scale asset purchases under QE3.

The Covid crisis caused only a moderate increase in the downside risk measure. While the initial impact of the Covid shock was large, the recession was short-lived and one-year ahead predictions were not as dramatic as during the 2008

[12]Further, in contrast to risk measures based on high-frequency market data, the downside risk index can be computed using a widely available macroeconomic dataset such as FRED-MD.

financial crisis. If I compute the downside risk measure based on one-quarter ahead predictions of consumption growth instead, there is a strong sudden rise in downside risk surmounting even the peak of the Great Recession, but an equally fast decline back to normal levels towards the end of 2020, which is in line with the results reported here. This result is reported in the Online Appendix. The weak rise in downside risk following the Covid shock is not because the model is unable to pick up existing downside risks, but instead because the predictions suggest the extreme downside risks will have subsided by the end of the twelve month forecasting horizon.

The dashed red line in Figure 4 is upside entropy. Downside and upside entropy move in the same direction during some times such as the early 1980s and much of the 1990s. Upside entropy can rise during recessions, reflecting the potential for economic recovery after a shock. However, both measures have overall different time series behaviour and do move in opposite directions during important times such as the lead-up to the Great Recession. The correlation between downside risk and upside potential is 0.51. Downside risk shows stronger variation over time than the upside counterpart and rises more strongly during crisis times. This reflects that upper quantiles of the consumption growth distribution are more stable than lower quantiles, see Figure 2.

3.4 Further results on consumption downside risk

How does the downside risk index compare to other measures of risk, uncertainty, or economic and financial activity? Figure 5 plots the downside risk index (blue dashed line, left scale) against six common indicators of risk/activity (orange line, right scale). The top left panel compares the downside risk index against the U.S. unemployment rate. Rises in downside risk are usually accompanied by rises in unemployment, although unemployment rises slower and often peaks only after the peak in downside risk. The decline in unemployment after a recession is usually also slower. The top right panel shows the Baker et al. (2016) index of economic policy

Figure 5: Downside Risk Index vs Other Variables

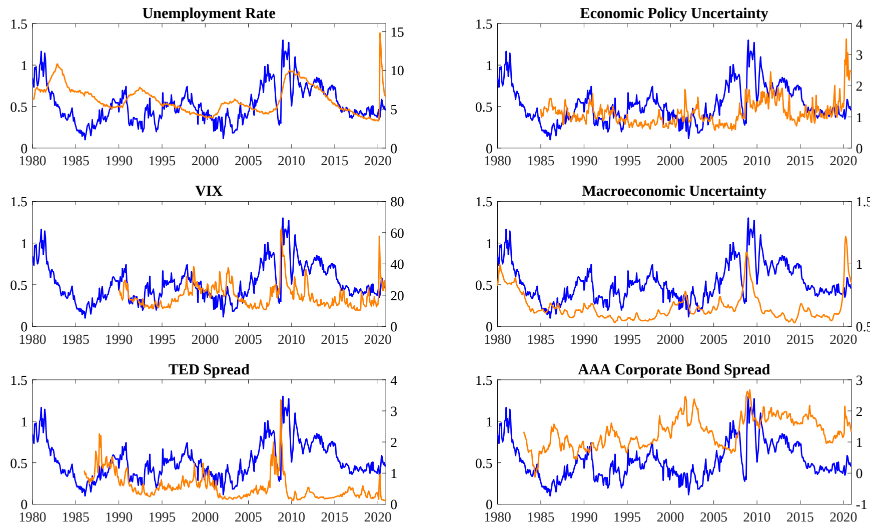

Note: The sample period for the downside risk index is 1:1980-12:2020. The unemployment rate and macroeconomic uncertainty index are plotted over the same period. The data for the Economic Policy Uncertainty index starts in 1:1985. The VIX data starts in 1:1990. Data for the TED spread is available from 1:1986, and data on the AAA corporate bond spread goes back to 1:1983.

uncertainty. Both series have a similar pattern during the late 1980s and 1990s. Economic policy uncertainty is more susceptible to policy-related events, however, and shows occasional upticks that are not reflected in downside risk. Examples include the 9/11 terrorist attacks and the Brexit referendum. Interestingly, policy uncertainty declines around the start of the Great Recession, supporting the idea that policy interventions around that time reduced economic uncertainty. The index also remains elevated for several years after the global financial crisis, similar to the downside risk measure.

The middle left panel plots the VIX stock market volatility index versus downside risk. While the VIX index rises during economic crises, the responses are often sharp and short-lived. The VIX is also more susceptible to stress originating in the financial markets, even if these do not (fully) spill over into the real economy. Examples include the Long Term Capital Management collapse and the series of U.S.

corporate scandals in the early 2000s. The middle right panel contains the Jurado et al. (2015) macroeconomic uncertainty index for a forecasting horizon of twelve months. The index averages the volatility in the unforecastable component of future values across a range of macroeconomic variables. The uncertainty index shows low to modest levels of economic uncertainty for most periods, and only rises significantly in the recessions of the early 1980s, the global financial crisis of 2007-2009, and the Covid crisis. The comparison between both measures following the Covid shock demonstrates that while one-year ahead forecasting uncertainty was high (as also indicated by the predicted variance in Figure 3), one-year ahead macroeconomic *downside* risks increased considerably less than during other recessions.

The bottom panels compare the downside risk index with two yield spreads commonly used to gauge stress in the bond market. The left panel shows the TED spread, which is the difference between the 3-month LIBOR rate and 3-month Treasury yields. During times of economic distress, lending to commercial banks becomes more risky and the TED spread increases. The TED spread is strongly affected by financial stress and peaked both during Black Monday in October 1987 and at the height of the credit crunch in October 2008. The TED spread rose significantly less during the Covid crisis than during the previous financial crisis, in line with the observation that the Covid turmoil in March 2020 did not lead to widespread financial stress. The right panel compares downside risk to the spread between the yield on Moody's AAA-rated seasoned corporate bonds and the yield on 10-year Treasuries. The spread rises mostly following financial turmoil in the corporate bond sector, for example during the early 2000s. The measure also rises sharply during the global financial crisis and remains at elevated levels for several years thereafter, similar to the measure of downside risk.

While the downside risk index moves similarly to the other measures during times of economic stress, the time series can be quite different outside of peak crisis times. This is reflected in low contemporaneous correlations between downside risk and the other variables. The correlation with the TED spread is slightly negative at a value of -0.05. The correlations with the VIX and Economic Policy Uncertainty are 0.17 and 0.10, respectively. Downside risk also has low correlation with the AAA corporate bond spread (0.25). The downside risk index shows the highest correlations with the unemployment rate (0.26) and macroeconomic uncertainty (0.31). This is encouraging since these two measures are those most closely reflecting macroeconomic conditions, whereas the other variables either partly or fully reflect financial conditions. Nevertheless, these correlations remain low. A regression of downside risk on the six variables yields a R-squared of 17.3 percent. Downside risk is a distinct measure of macroeconomic risk and not well explained by existing risk, uncertainty, or activity indicators.

In summary, the index of downside risk shares certain commonalities with many of the existing measures of risk, uncertainty, or economic/financial stress. All measures rise during recessions periods before returning to their long-term average. However, the speed and size of the rise and recovery can differ strongly between the different variables, as well as between recessions. Downside risk has a time series behaviour that is not easily explained by any of the other variables discussed. Its relation to monetary policy and equity markets may therefore be distinct from that of existing measures of risk or uncertainty.

4 Monetary Policy and Consumption Downside Risk

I study the effect of a monetary policy shock on the index of consumption downside risk with vector autoregressions. The first subsection discusses the identification of the monetary policy shock using high-frequency federal funds futures data as an external instrument. The subsequent subsection discusses the main results. The section next discusses state-dependent effects of the monetary policy shock. I close with an analysis of the effects of policy shocks to downside risk when controlling both risk aversion and uncertainty.

4.1 External instrument VAR methodology

To study how monetary policy surprises affect macroeconomic downside risk, I estimate the effect of a monetary policy shock on the risk measure using a structural vector autoregression. Let Y_t be the vector of (macroeconomic and financial) variables, A_j the matrix of coefficients for the vector autoregessive lags, and Φ the matrix of contemporaneous coefficients. I denote the vector of structural shocks as ε_t to get:

$$\Phi Y_t = \mu^* + \sum_{j=1}^{p} A_j Y_{t-j} + \varepsilon_t. \tag{9}$$

The monetary shock is identified using a high-frequency instrument as proposed by Gertler & Karadi (2015). This approach circumvents the simultaneity problem between the monetary policy indicator and the downside risk index: While downside risk may respond contemporaneously to changes in monetary policy, central bankers may also take current developments in downside risk into account when setting the policy rate. The shock is identified by looking at changes in a monetary policy instrument within a tight (usually intra-day) window around monetary policy shifts. The identifying assumption is that no other shocks affect the monetary instrument within this window. This allows to identify monetary policy surprises using an *external instrument* and does not impose any further restrictions on the contemporaneous relation between the variables. Since we are only interested in the effect of a monetary policy shock, this yields the required number of restrictions to estimate dynamic responses of all variables in our VAR to a monetary shock.[13]

The monetary policy instrument is constructed from changes in the interest rate implied by three-month federal funds futures around FOMC announcements. These assets measure the expected average federal funds rate over the third calendar month from the contract.[14] Since the futures are forward-looking, they do not

[13]See Online Appendix D for details on the approach.

[14]For example, a three-month federal funds future in June 2021 measures market expectations about the average effective federal funds rate in September 2021.

only reflect changes in the current short rate but also capture short-term forward guidance.[15] This is especially useful in a low-rate environment in which forward guidance has become an important tool to steer the monetary policy stance. While short-rate surprises may be small during the ZLB period, considering a slightly longer horizon guarantees enough variation in the surprises to allow for a meaningful analysis of the policy shocks. I follow the standard approach in the literature and consider the change in three-month federal funds futures prices 10 minutes before the announcement relative to 20 minutes after the announcement for all FOMC meetings between February 1990 and December 2020 (see e.g. Gürkaynak et al. (2005), Nakamura & Steinsson (2018)).

The raw monetary policy shock series may be contaminated by Fed information effects. Market participants may interpret a change in the monetary policy stance not only as a pure policy shock, but also as new information about the Fed's assessment of the future state of the economy. To separate the policy shock from the information shock, I use the "poor man's sign restrictions" of Jarociński & Karadi (2020). This approach sets the Fed funds rate change to zero if the stock market and the interest rate do not move in opposite directions over the event window. If the stock market interprets a surprise loosening of the policy stance as expansionary, we would expect stock prices to rise. If, however, the loosening in the policy stance is interpreted as a sign for a bad economic outlook, stock prices fall. The goal of the sign restrictions is to only keep the surprise changes from FOMC decisions which were not interpreted as information events. Online Appendix E details the construction of the shock series.[16] To obtain a monthly shock series, I sum up all high-frequency Fed funds rate surprises in a given month after having applied the sign restrictions.

[15] Jarociński & Karadi (2020) point out that usually six weeks elapse between two subsequent FOMC meetings such that the response of the three-month federal funds future price on an FOMC meeting day reflects the expected change in the federal funds rate following the *next* policy meeting.

[16] This appendix also describes how I construct the monetary policy shock series of Miranda-Agrippino & Ricco (2021*b*), which aims to remove information effects using Greenbook projection data. I use this alternative shock series for a robustness exercise.

4.2 Main results

I estimate the VAR for monthly data covering the period February 1984 until December 2019.[17] The Covid period is not part of the baseline specification since there is no clear guidance on how to interpret the Covid shock in a VAR setting.[18] Since the data on the monetary policy instrument only dates back to February 1990, I estimate the lag coefficients of the VAR for the longer sample and obtain the residuals. I then identify the effect of monetary policy shocks using the high-frequency data and residuals from February 1990 until the end of the sample, which yields an estimate of vector s. I choose 12 lags for the baseline specification.[19]

The baseline specification includes industrial production as an indicator of real economic activity, the consumer price index to measure prices, the excess bond premium (EBP) by Gilchrist & Zakrajšek (2012), the one-year U.S. Treasury yield as a monetary policy indicator, and the indices of consumption downside as well as upside risk.[20] The excess bond premium measures investors' attitude towards risk in the U.S. corporate credit market. It outperforms many other financial variables in terms of forecasting power for real activity and can hence serve as a proxy for the financial information relevant to predict macroeconomic indicators.[21]

[17] Jarociński & Karadi (2020) choose February 1984 as the starting date for their VAR since it coincides with the end of the Volcker disinflation period.

[18] Lenza & Primiceri (2022) and Ng (2021) make two different proposals on how to deal with the Covid period in a VAR setting.

[19] Across all of my model specifications, the corrected Akaike Information Criterion (AICc) usually suggests an optimal lag length between six and 12 months. A robustness check shows the results are very similar for six lags.

[20] To ensure stationarity, industrial production and the CPI are in year-over-year growth rates. The monetary policy indicator is in monthly changes, and the downside and upside risk index are in year-over-year absolute changes. The EBP is in levels. Data for the EBP is from an updated version of the Favara et al. (2016) data, the other macro time series are from FRED. I formally test for stationarity using the Augmented Dickey-Fuller test with 12 lags (results not reported).

[21] I choose the one-year U.S. government bond yield as the monetary policy indicator since, relative to the federal funds rate, the one-year rate contains some information about the effect of forward guidance. It can therefore serve as a measure of the monetary policy stance even when the federal funds rate is constrained by the zero lower bound. The empirical evidence of Gertler & Karadi (2015) suggests that combining the one-year Treasury yield as a policy indicator and three-month federal funds futures as a policy instrument allows to identify monetary policy shocks that affect market interest rates and spreads in a sensible way.

Figure 6 shows the impulse responses to a monetary policy shock that lowers the 1-year government bond yield by 25bp, which equals about a one standard deviation impact. In line with economic theory, the policy loosening leads to a rise in industrial production and the inflation rate. The excess bond premium decreases on impact and remains below its steady state value for several months, indicating a loosening in financial conditions. The downside risk index falls on impact, and the effect is strongest after eight months, when the index falls by almost 0.1 points. This is in line with the idea of transmission lags in the effect of monetary policy. The significant effect on changes in downside risk persists for about one year, before returning to zero. In contrast, the effect on downside risk is consistenly insignificant. Monetary policy changes appear able to reduce downside risks, but are not found to be effective in improving upside risks to the economy.

The heteroskedasticity-robust F-statistic from the first-stage IV regression

Figure 6: Main VAR results

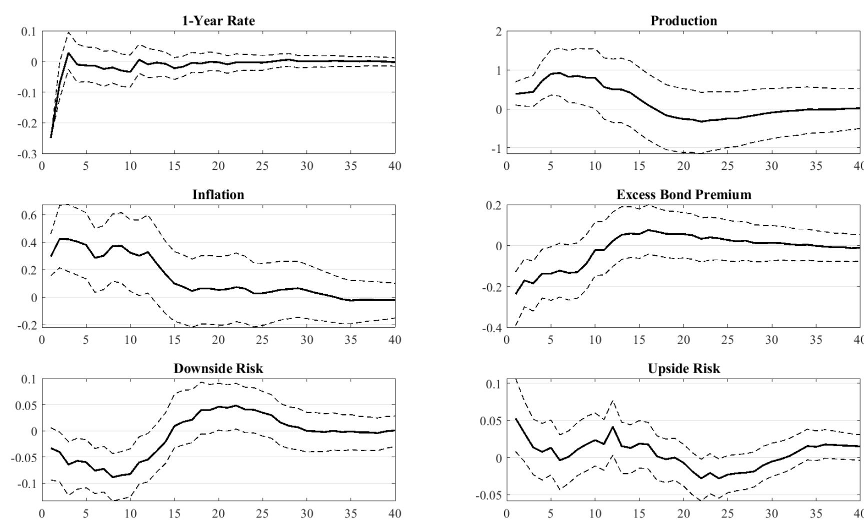

The sample period is 1984:2-2019:12. The controls included are industrial production, the CPI, and the excess bond premium (EBP). The monetary policy indicator is the first-differenced one-year government bond yield, instrumented by the Jarociński & Karadi (2020) shocks. All data at a monthly frequency, lag length is 12 months. Dashed lines are 95% confidence bands from a wild bootstrap with 5,000 replications.

is 10.03. Following the Staiger & Stock (1997) rule of thumb, this indicates the instrument for the monetary policy shock is potentially weak as the F-statistic is only marginally above ten. The impulse responses may be biased. Appendix G reports results from a VAR estimated with the federal funds rate as the policy indicator and high-frequency changes in the 3-month federal funds futures rate as the policy instrument (see Figure 19a). This yields a strong instrument (F-statistic of 35.3). The results are very similar, suggesting that the baseline results are not driven by a weak policy instrument.

Appendix G also shows that the results are robust to using different combinations of policy indicator and instrument, including the Covid period in the sample, using a different monetary policy indicator, to only using six lags instead of twelve, to including various other controls, and to using alternative measures of macroeconomic downside risk. Lastly, the results are neither sensitive to non-invertibility concerns nor potential model misspecification (see the discussion in Appendix F).

4.3 State-dependent results

Do monetary policy shocks have different effects on macroeconomic downside risk depending on the state of the economy? During a recession, we may expect monetary policy to have a stronger effect on downside risk than during normal times, when the realization of negative macroeconomic outcomes is not a major concern. To test this hypothesis, I allow for state dependence by estimating a threshold VAR:

$$Y_t = \mathbb{I}\left(rec_{t-1}\right)\left[\mu_{rec} + \sum_{j=1}^{p} B_{rec,j} Y_{t-j}\right] + \mathbb{I}\left(1 - rec_{t-1}\right)\left[\mu_{norec} + \sum_{j=1}^{p} B_{norec,j} Y_{t-j}\right] + u_t, \tag{10}$$

where $\mathbb{I}\left(rec_{t-1}\right)$ is an indicator function indicating whether the economy is in a recession when the monetary shock hits the economy at the start of period t. All coefficients are allowed to be state-dependent, including the impact effects s of the monetary shock.

The start of a recession is defined following the NBER's Business Cycle Dates. Since only using the official NBER recession months would yield too few recession states for the VAR, I additionally define all months with an unemployment rate above 6.5% as recession months. Instead of only defining recession periods based on the unemployment rate, this combined approach allows to more accurately capture the start of a recession. Since the unemployment rate may only rise with a lag following the start of a recession but monetary policy actions aimed at crisis mitigation often occur at the onset of a downturn, only using a high unemployment rate to define recessions could miss important policy events.

Figures 7 and 8 show the results for recession states and non-recession states, respectively. Monetary policy has significant effects on consumption downside risk for several months during a recession. The impulse response of downside risk is similar in shape and larger in size to the main result. The measurable effect of monetary policy on downside risk are particularly strong in recessions. The confidence bands are wider than for the baseline specification. This is both due to the smaller sample size and the larger volatility in recession periods. Nevertheless, the response of downside risk is still borderline significant in the recession state, even at a 95% confidence level. In contrast, there is no clear effect on upside risk.

In non-recession states, policy shocks have no clear effect on downside risk. The results from the previous section are therefore largely driven by the effect of monetary policy interventions in recessions. Loosening the monetary policy stance can have an important short-term macroeconomic stabilization role during crisis times but has little to no measurable effect on downside risk in other times.[22]

[22] Estimating state-dependent effects complicates the identification of impact responses due to a weak instrument. Since strong monetary policy shocks occur mostly during crisis times, this can especially harm the identification in the non-recession state. The robust F-statistic is 16.9 for the recession state, but only 4.8 for the non-recession state. The results in Figure 8 should therefore be interpreted with care. However, the results are almost idential when using a combination of policy indicator and instrument that yields a strong instrument during normal times (F-statistic of 17.6), see Figure 19b in Appendix G.

Figure 7: State-dependent VAR results (Recession states)

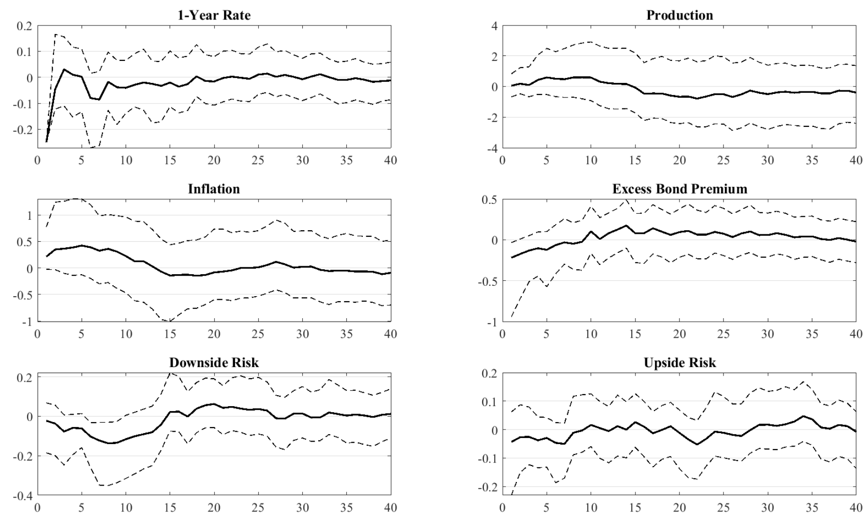

The sample period is 1984:2-2019:12. The controls included are industrial production, the CPI, and the excess bond premium (EBP). The monetary policy indicator is the first-differenced one-year government bond yield, instrumented by the Jarociński & Karadi (2020) shocks. All data is at a monthly frequency, the lag length is 12 months. Dashed lines are 95% confidence bands from a wild bootstrap with 5,000 replications.

4.4 Risk Aversion and Uncertainty

This section demonstrates that the effect of monetary policy on downside risk is not explained by a risk aversion or uncertainty channel. The downside risk index is an asymmetric index of macroeconomic risk and is designed to not capture risk aversion or symmetric uncertainty. Therefore, the results of this chapter should be robust to controlling for measures of risk aversion or uncertainty. To make this point clear, this section adds measures of risk aversion and uncertainty to the vector of variables from the main VAR. The uncertainty measure is the Jurado et al. (2015) index of macroeconomic uncertainty for a forecasting horizon of twelve months. The risk aversion measure comes from Bekaert et al. (2021).[23] Figure 9 shows that even

[23]The authors estimate the risk aversion coefficient for a representative agent in a no-arbitrage asset pricing model. Since asset prices and risk premia are functions of the risk aversion coefficient and other model parameters, risk aversion can be backed out from a set of observed financial variables and the model restrictions. The authors provide data for risk aversion from 1986:6 onwards.

Figure 8: State-dependent VAR results (Non-recession states)

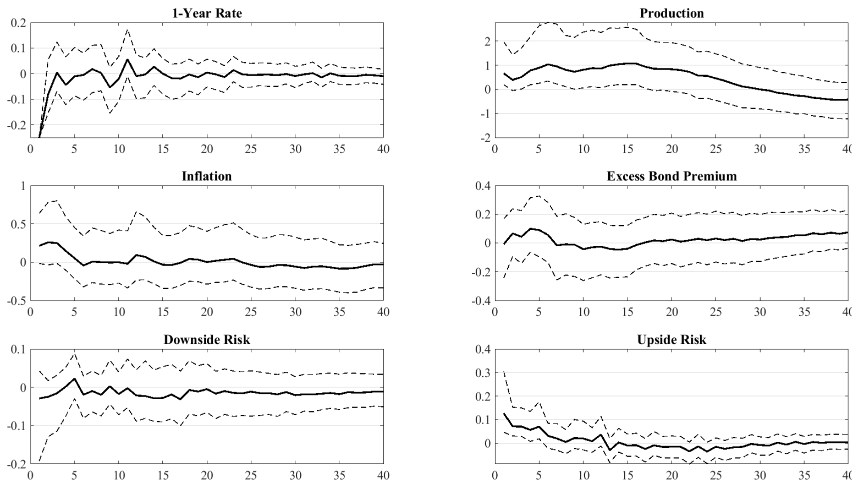

The sample period is 1984:2-2019:12. The controls included are industrial production, the CPI, and the excess bond premium (EBP). The monetary policy indicator is the first-differenced one-year government bond yield, instrumented by the Jarociński & Karadi (2020) shocks. All data is at a monthly frequency, the lag length is 12 months. Dashed lines are 95% confidence bands from a wild bootstrap with 5,000 replications.

after controlling for both of these variables, the effect of a monetary policy shock on downside risk remains very similar.

In summary, this section shows that monetary policy has a significant impact on downside risk to consumption growth. Following a 25bp loosening in the policy stance, downside risk declines significantly for about one year. This effect of monetary policy shocks is driven by crisis periods, whereas monetary policy has little to no measurable effect on downside risk in normal times. Policy shocks do not change upside potential to consumption growth. These results persist when controlling for risk aversion and macroeconomic uncertainty since downside risk is a distinct measure of macroeconomic risk.

This restricts the VAR sample to 1986:6 - 2019:12.

Figure 9: VAR Results - Including Risk Aversion and Uncertainty

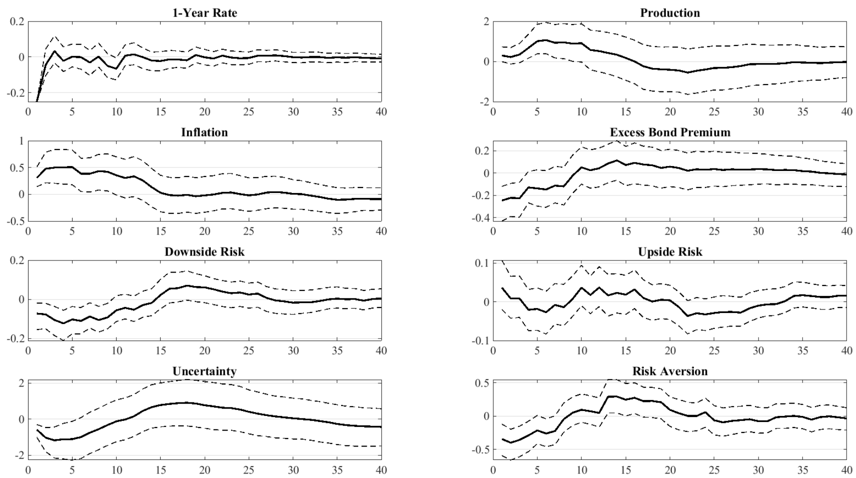

The sample period is 1987:6-2019:12. The controls included are industrial production, the CPI, the excess bond premium (EBP), the Jurado et al. (2015) macro uncertainty index, and the Bekaert et al. (2021) risk aversion index. The monetary policy indicator is the first-differenced one-year government bond yield, instrumented by the Jarociński & Karadi (2020) shocks. All data is at a monthly frequency, the lag length is 12 months. Dashed lines are 95% confidence bands from a wild bootstrap with 5,000 replications.

5 Consumption Risk and Equity Returns

This section studies if downside consumption risk predicts stock market returns in the aggregate and across industries. Let $r_{t,t+s}$ be the aggregate stock market return between period t and $t+s$. For a given forecasting horizon s, the predictive return regression then is

$$r_{t,t+s} = a_s + b_s' x_t + u_{t,t+s}, \tag{11}$$

where x_t is the vector of predictors. The dependent variable is the excess return on the Fama-French market portfolio, where the safe rate is the one-month Treasury bill return. The possible predictors are the consumption downside risk index, the VIX volatility index, the variance risk premium (VRP), and the log price-dividend ratio.[24] As in the VAR analysis, the risk measure is in year-on-year changes. The other variables are in levels, as is common in the empirical literature.

Dividend yields have been extensively studied as a predictor of stock returns,

Table 1: Market return predictability regressions - Baseline results

	Horizon h = 3					Horizon h = 6				
Risk	1.61	1.51	1.61	1.37	1.25	3.04	2.84	3.04	2.78	2.55
	(0.80)	(0.83)	(0.83)	(0.64)	(0.68)	(1.71)	(1.61)	(1.64)	(1.54)	(1.49)
	[0.72]	[0.72]	[0.72]	[0.71]	[0.71]	[1.25]	[1.25]	[1.25]	[1.23]	[1.23]
log(PD)		-1.13			-1.32		-2.42			-2.64
		(0.70)			(0.71)		(1.29)			(1.35)
		[0.77]			[0.76]		[1.39]			[1.38]
VIX			0.33		0.06			1.05		0.80
			(0.84)		(0.58)			(0.95)		(0.74)
			[0.84]		[0.90]			[1.13]		[1.13]
VRP				1.99	2.10				2.18	2.29
				(0.34)	(0.31)				(0.74)	(0.62)
				[0.72]	[0.78]				[0.86]	[0.83]
Constant	2.07	18.49	1.21	0.5	19.31	4.08	38.91	1.37	2.39	38.26
	(0.67)	(10.04)	(2.01)	(0.58)	(10.16)	(1.41)	(18.26)	(2.59)	(1.20)	(20.14)
	[0.66]	[10.98]	[2.03]	[0.86]	[10.75]	[1.32]	[19.71]	[2.69]	[1.38]	[20.07]
R^2	4.3	6.3	4.2	11.0	13.5	7.6	12.3	8.3	11.4	17.2

The dependent variable is the excess return on the S&P 500 index over the next 3 or 6 months. Risk is the downside entropy of the consumption growth distribution, in year-over-year growth rates. The log price-dividend ratio is taken from Robert Shiller's website. The VIX is the monthly level of the VIX index, constructed as the within-month average of daily adjusted closing prices. The variance risk premium is from Hao Zhou's website. The sample period is 01/1990 until 12/2019. All data is at a monthly frequency. Hodrick (1992) and Newey-West standard errors in parentheses and brackets, respectively. R^2 is adjusted for the number of predictors.

especially over longer horizons (e.g. Fama & French (1988)). Crucially, consumption-based asset pricing models such as Campbell & Cochrane (1999) suggest a close link between price-dividend ratios and aggregate consumption. To the extent that these two measures co-move, the price-dividend ratio serves as a gauge of the forecasting power of aggregate consumption for stock returns. If the consumption downside risk

[24] Data on the S&P 500 and the VIX is from Yahoo Finance. The time series for the safe rate is from Kenneth French's Data Library. Data on the variance risk premium is from Hao Zhou's personal website (`https://sites.google.com/site/haozhouspersonalhomepage/`). Data on the price-dividend ratio is from Robert Shiller's website (`http://www.econ.yale.edu/~shiller/data.htm`).

Table 2: Market return predictability regressions - State-dependent

	Horizon h = 3					Horizon h = 6				
Risk	0.32	0.33	0.39	0.35	0.45	0.52	0.53	0.68	0.54	0.79
	(0.43)	(0.40)	(0.39)	(0.43)	(0.42)	(0.81)	(0.63)	(0.66)	(0.71)	(0.58)
	[0.67]	[0.67]	[0.68]	[0.67]	[0.68]	[1.15]	[1.15]	[1.17]	[1.15]	[1.18]
Recession	-3.24	-4.07	-4.2	-3.58	-5.97	-5.06	-6.75	-7.40	-5.34	-10.71
	(2.49)	(2.29)	(2.56)	(2.12)	(1.94)	(3.95)	(4.39)	(5.69)	(3.48)	(5.65)
	[3.04]	[3.03]	[3.14]	[3.08]	[3.17]	[5.37]	[5.42]	[5.84]	[5.38]	[5.83]
Rec*Risk	4.74	4.35	4.48	3.90	2.95	9.31	8.52	8.67	8.60	6.58
	(0.96)	(0.94)	(0.94)	(1.07)	(0.89)	(2.73)	(2.21)	(2.43)	(2.61)	(1.92)
	[1.94]	[1.96]	[1.95]	[1.96]	[1.97]	[2.90]	[2.94]	[2.90]	[2.94]	[2.99]
log(PD)		-1.12			-1.53		-2.32			-2.92
		(0.61)			(0.63)		(1.12)			(1.09)
		[0.78]			[0.80]		[1.45]			[1.42]
VIX			0.59		0.76			1.45		1.97
			(0.79)		(0.57)			(1.08)		(1.02)
			[0.90]		[0.97]			[1.26]		[1.19]
VRP				1.66	1.81				1.41	1.66
				(0.55)	(0.37)				(1.22)	(0.80)
				[0.74]	[0.76]				[0.88]	[0.84]
Constant	2.42	18.78	0.99	1.15	21.38	4.65	38.18	1.15	3.58	40.91
	(0.56)	(8.76)	(1.88)	(0.65)	(9.08)	(1.10)	(16.05)	(2.44)	(1.28)	(15.82)
	[0.65]	[11.14]	[2.13]	[0.80]	[11.08]	[1.27]	[20.51]	[2.87]	[1.28]	[20.40]
R^2	14.7	15.4	13.7	18.9	21.3	24.1	28.1	25.1	25.4	32.5

The dependent variable is the excess return on the S&P 500 index over the next 3,6,12,36 months. Risk is the downside entropy of the consumption growth distribution, in year-over-year growth rates. Recessions are as defined by the NBER. The log price-dividend ratio is taken from Robert Shiller's website. The VIX is the monthly level of the VIX index, constructed as the within-month average of daily adjusted closing prices. The variance risk premium is from Hao Zhou's website. The sample period is 01/1990 until 12/2019. All data is at a monthly frequency. Hodrick (1992) and Newey-West standard errors in parentheses and brackets, respectively. R^2 is adjusted for the number of predictors.

index reflects information other than that contained in aggregate consumption, we should expect the predictive power of the risk index to be robust to the inclusion of the price-dividend ratio.

The VIX is a measure of option-implied expected volatility of the S&P 500 index. While the index formally contains information about risk and investors' risk attitude (risk aversion), it is often considered a measure of uncertainty more generally (e.g. Bloom (2009)). Since the consumption downside risk index is not constructed as a

Table 3: Market return predictability regressions - Industry portfolios

	Linear		State-Dependent					
			Risk		Recession		Risk ∗ Recession	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Non-Durable Goods	1.91	1.01	0.00	0.99	1.10	4.28	7.12	2.08
Utilities	1.97	1.04	0.18	1.00	-7.54	4.59	6.54	2.45
Healthcare	2.26	1.21	0.86	1.25	-1.04	4.96	5.22	2.55
Wholesale	2.44	1.19	0.76	1.21	1.59	4.95	6.30	2.54
Energy	2.53	1.45	0.44	1.60	-8.70	5.28	7.65	2.80
Chemicals	2.95	1.20	0.48	1.12	-1.20	5.47	9.16	2.98
Computers	2.99	1.55	-0.26	1.56	-5.47	7.09	11.99	3.35
Telecommunication	3.32	1.35	1.15	1.23	-10.23	5.99	7.89	3.14
Manufacturing	3.99	1.66	0.69	1.39	-5.08	7.18	12.21	4.24
Finance	4.13	1.72	0.88	1.48	-5.38	7.63	11.99	4.49
Durable Goods	5.55	2.44	1.19	1.88	-2.34	10.66	16.19	6.80

The dependent variable is the return on the industry portfolio over the next 6 months. Risk is the downside entropy of the consumption growth distribution, in year-over-year growth rates. Recessions are as defined by the NBER. All regressions include a constant but the coefficient on the constant is not reported to conserve space. The sample period is 01/1990 until 12/2019. All data is at a monthly frequency. The standard errors are based on Hodrick (1992).

measure of uncertainty in financial markets but of downside risks to the macroeconomy, we should expect both measures to contain different information for predicting stock market returns.

The variance risk premium aims to proxy investor risk aversion and can be obtained from a decomposition of the VIX index into a risk component and a residual associated with risk aversion. The VRP has been shown to be a good short-horizon predictor of aggregate stock market returns (Bollerslev et al. (2009)). However, a large part of its predictive ability may be due to investor aversion of tail risk (extreme negative outcomes) instead of overall risk (Bollerslev et al. (2015)). In this sense, the VRP also serves as a (noisy) measure of tail risk aversion. Since the consumption downside risk index does not contain information about (tail) risk aversion, we should expect its predictive ability to be robust to controlling for the VRP.

The sample period for the baseline regressions is January 1990 until December 2019. All regressions are estimated at a monthly frequency with a forecasting horizon of 3 or 6 months. For these forecasting horizons, the monthly returns are overlapping. I use Hodrick (1992) standard errors to account for this feature of the data. Since Hodrick (1992) standard errors are constructed for the null hypothesis of no predictability by any predictor, using them in regressions with multiple predictors makes them hard to interpret and formally invalid (see Bollerslev et al. (2015)). To address this issue, the main tables report Newey-West standard errors as well.

Table 1 contains the results. All predictors are standardized such that the coefficients indicate the change in excess market returns following a one standard deviation change in the predictor. Changes in downside risk are associated with significant stock market movements, even after controlling for the alternative predictors. A one-standard deviation increase in the downside risk measure predicts an increase in aggregate excess returns of around 1.5 percentage points for the three-month horizon, and of around 3 percentage points for the six-month horizon. The positive sign on the coefficient is in line with the prediction of disaster risk models: Since a rise in downside risk poses an additional source of risk for risk averse stock investors, the risk premium required to hold stocks increases. The multivariate regressions indicate that the downside risk index captures information distinct from the information about risk and risk aversion contained in the VIX or VRP. The risk index also contains information about aggregate consumption beyond the information about the mean contained in the price-dividend ratio.

Is the predictive power of the downside risk index for stock returns largely driven by crisis times? I introduce a recession indicator $\mathbb{I}\left(rec\right)$, which equals one during NBER recession months. This allows to consider state-dependent effects of the downside risk index in the predictive regressions:

$$r_{t,t+s} = a_s + b_s \mathbb{I}\left(rec_t\right) + c_s Risk_t + d_s \mathbb{I}\left(rec_t\right) * Risk_t + e_s' x_t + u_{t,t+s}, \tag{12}$$

where x_t is now the vector of other predictors.

The results for the state-dependent regressions are in Table 2. Again, the coefficients indicate the expected change in stock returns following a one standard deviation rise in the predictor. The predictive ability of downside risk for excess returns is concentrated in recession periods. During recessions, the downside risk indicator predicts future stock market returns for both horizons. A one standard deviation rise in the downside risk measure predicts excess return increases around 4 percentage points for the three-month horizon, and around 8.5 percentage points for the six-month horizon. These magnitudes are economically meaningful when considering that stock markets can easily move by more than 20 percentage points over the course of six months during recessions. For example, the excess return on the market portfolio between September 2008 and February 2009 was -50.8%, while the downside risk index increased by multiple standard deviations. The results are again robust to including the price-dividend ratio, the VIX, and the VRP as additional predictors. Appendix H shows that the results are robust to allowing for state dependence of the other predictors, using different sample periods, controlling for median or realized consumption growth.

We can also expect downside risk to have varying predictive power across industries. If investors are averse to downside risk, sectors with a stronger exposure to downside risk should command higher returns following an increase in risk. To test this hypothesis, I run predictive regressions for the excess returns of the twelve Fama-French SIC industry portfolios.[25]

Table 3 reports the results for the regressions with a forecasting horizon of 6 months. The left column contains the results for univariate regressions with only the downside risk measure as a predictor. Industries that can be thought to have little exposure to downside risk such as non-durable goods, utilities, and healthcare show lower coefficients, between 1.91 and 2.26. In contrast, sectors more closely following the business cycle such as manufacturing, financial services, and durable goods consumption all have higher coefficients. For example, a one standard deviation

[25] The return data comes from Kenneth French's website. I omit the portfolio "Other", leaving eleven portfolios.

rise in downside risk predicts 4.13% higher returns in the financial services sector over the next six months. While downside risk is a significant predictor across many industries, it has quantitatively stronger predictions for sectors with higher downside risk exposure.

The differences across industries remain largely unchanged if we allow for state-dependent effects as given in equation (12). Healthcare and the wholesale sector still have the low sensitivity to downside risk, whereas manufacturing and durable goods show the highest sensitivities. As for the aggregate market return, the predictive power of the downside risk index is virtually entirely concentrated in recession times, but the coefficient sizes for the interaction term Risk $*$ Recession now vary from 5.22 (healthcare) to 16.19 (durable goods). In summary, downside risk predicts excess stock returns for the aggregate market and across a wide range of industries. More procyclical industries have higher sensitivities to downside risk.

6 Conclusion

This paper documents the existence of a downside risk channel of monetary policy. I start by estimating a model-free index of downside risks to aggregate consumption growth. Higher index values indicate an increased probability of low realizations of consumption growth. Rises in downside risk coincide with major U.S. recessions, but downside risk does not reflect pure financial stress. The index of downside consumption risk is directly related to the series of consumption growth and admits a natural interpretation in the context of a consumption-based asset pricing model. The index shows similar but nevertheless distinct time series behaviour relative to several existing measures of risk or uncertainty.

The main result is that monetary policy shocks have asymmetric effects on consumption growth risk. A loosening of the monetary policy stance reduces macroeconomic downside risks during crisis times. This effect persists for about one year. Monetary policy has little to no impact on downside risk during normal times.

The results also document that policy shocks yield no significant change in upside potential. Importantly, these results are not driven by Fed information effects, the Covid period, shifts in the median of the consumption growth distribution, or by risk aversion and uncertainty.

Changes in downside risk are associated with sizeable changes in future stock market returns. Increases in downside risk can predict high future returns especially during crisis times, when the downside risk index can move by multiple standard deviations in a short time. Industry-sorted portfolios have varying sensitivities to changes in downside risk. Sectors such as healthcare and utilities are less sensitive to changes in downside risk than manufacturing or finance.

My findings support the idea that central banks can avert macroeconomic disasters in crisis times. By loosening the policy stance, the probability of negative outcomes can decrease significantly. However, the results suggest monetary policy has little opportunity to increase the growth potential of the economy in normal times. The downside risk channel of monetary policy may also explain the strong reaction of stock markets following policy interventions in recessions. If investors perceive that a rate cut has lowered macroeconomic risks, valuations should increase following the decision. This gives monetary policy a powerful lever on stock prices that goes beyond its effects on realized growth rates of macroeconomic variables.

References

Adrian, T., Boyarchenko, N. & Giannone, D. (2019), 'Vulnerable growth', *American Economic Review* **109**(4), 1263–89.

Ajello, A. & Pike, T. (2022), 'Getting in all the cracks: Monetary policy, financial vulnerabilities, and macro risk', *mimeo* .

Azzalini, A. & Capitanio, A. (2003), 'Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**(2), 367–389.

Backus, D., Chernov, M. & Martin, I. (2011), 'Disasters implied by equity index options', *The Journal of Finance* **66**(6), 1969–2012.

Bai, J. & Ng, S. (2002), 'Determining the number of factors in approximate factor models', *Econometrica* **70**(1), 191–221.

Bai, J. & Ng, S. (2008*a*), 'Forecasting economic time series using targeted predictors', *Journal of Econometrics* **146**(2), 304–317.

Bai, J. & Ng, S. (2008*b*), 'Large dimensional factor analysis', *Foundations and Trends in Econometrics* **3**(2), 89–163.

Baker, S. R., Bloom, N. & Davis, S. J. (2016), 'Measuring economic policy uncertainty', *The Quarterly Journal of Economics* **131**(4), 1593–1636.

Barro, R. J. (2006), 'Rare disasters and asset markets in the twentieth century', *The Quarterly Journal of Economics* **121**(3), 823–866.

Barro, R. J. & Liao, G. Y. (2021), 'Rare disaster probability and options pricing', *Journal of Financial Economics* **139**(3), 750–769.

Beason, T. & Schreindorfer, D. (2022), 'Dissecting the equity premium', *Journal of Political Economy* **130**(8), 000–000.

Bekaert, G., Engstrom, E. C. & Xu, N. R. (2021), 'The time variation in risk appetite and uncertainty', *Management Science* **68**(6), 3975–4753.

Bekaert, G. & Hoerova, M. (2014), 'The vix, the variance premium and stock market volatility', *Journal of Econometrics* **183**(2), 181–192.

Bekaert, G., Hoerova, M. & Duca, M. L. (2013), 'Risk, uncertainty and monetary policy', *Journal of Monetary Economics* **60**(7), 771–788.

Belloni, A. & Chernozhukov, V. (2011), 'l1-penalized quantile regression in high-dimensional sparse models', *The Annals of Statistics* **39**(1), 82–130.

Bloom, N. (2009), 'The impact of uncertainty shocks', *Econometrica* **77**(3), 623–685.

Boivin, J. & Ng, S. (2006), 'Are more data always better for factor analysis?', *Journal of Econometrics* **132**(1), 169–194.

Bollerslev, T., Tauchen, G. & Zhou, H. (2009), 'Expected stock returns and variance risk premia', *The Review of Financial Studies* **22**(11), 4463–4492.

Bollerslev, T. & Todorov, V. (2011*a*), 'Estimation of jump tails', *Econometrica* **79**(6), 1727–1783.

Bollerslev, T. & Todorov, V. (2011*b*), 'Tails, fears, and risk premia', *The Journal of Finance* **66**(6), 2165–2211.

Bollerslev, T., Todorov, V. & Xu, L. (2015), 'Tail risk premia and return predictability', *Journal of Financial Economics* **118**(1), 113–134.

Campbell, J. Y. & Cochrane, J. H. (1999), 'By force of habit: A consumption-based explanation of aggregate stock market behavior', *Journal of Political Economy* **107**(2), 205–251.

Carr, P. & Wu, L. (2009), 'Variance risk premiums', *The Review of Financial Studies* **22**(3), 1311–1341.

Chernozhukov, V., Fernández-Val, I. & Galichon, A. (2010), 'Quantile and probability curves without crossing', *Econometrica* **78**(3), 1093–1125.

Christiano, L. J., Eichenbaum, M. & Evans, C. L. (2005), 'Nominal rigidities and the dynamic effects of a shock to monetary policy', *Journal of Political Economy* **113**(1), 1–45.

Cochrane, J. H. (2017), 'Macro-finance', *Review of Finance* **21**(3), 945–985.

Cortes, G., Gao, G., Silva, F. B. G. & Song, Z. (2020), 'Unconventional monetary policy and disaster risk: Evidence from the subprime and covid-19 crises', *Available at SSRN: https://ssrn.com/abstract=3642970* .

Duprey, T. & Ueberfeldt, A. (2020), 'Managing gdp tail risk', *Bank of Canada Staff Working Paper No. 2020-03* .

Fama, E. F. & French, K. R. (1988), 'Dividend yields and expected stock returns', *Journal of Financial Economics* **22**(2), 3–25.

Favara, G., Gilchrist, S., Lewis, K. F. & Zakrajsek, E. (2016), 'Recession risk and the excess bond premium', *FEDS notes No. 2016-04-08* .

FOMC (2020), 'Chair powell's press conference on 5 november 2020'.
URL: *https://www.federalreserve.gov/mediacenter/files/FOMCpresconf20201105.pdf, Accessed on: 09/09/2021*

Gabaix, X. (2012), 'Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance', *The Quarterly Journal of Economics* **127**(2), 645–700.

Gertler, M. & Karadi, P. (2015), 'Monetary policy surprises, credit costs, and economic activity', *American Economic Journal: Macroeconomics* **7**(1), 44–76.

Giglio, S., Kelly, B. & Pruitt, S. (2016), 'Systemic risk and the macroeconomy: An empirical evaluation', *Journal of Financial Economics* **119**(3), 457–471.

Gilchrist, S. & Zakrajšek, E. (2012), 'Credit spreads and business cycle fluctuations', *American Economic Review* **102**(4), 1692–1720.

Gourio, F. (2012), 'Disaster risk and business cycles', *American Economic Review* **102**(6), 2734–2766.

Gürkaynak, R. S., Sack, B. P. & Swanson, E. T. (2005), 'Do actions speak louder than words? the response of asset prices to monetary policy actions and statements', *International Journal of Central Banking* **1**, 55–93.

Haddad, V., Moreira, A. & Muir, T. (2022), 'Whatever it takes? the impact of conditional policy promises', *Working Paper* .

Hattori, M., Schrimpf, A. & Sushko, V. (2016), 'The response of tail risk perceptions to unconventional monetary policy', *American Economic Journal: Macroeconomics* **8**(2), 111–136.

Hodrick, R. J. (1992), 'Dividend yields and expected stock returns: Alternative procedures for inference and measurement', *The Review of Financial Studies* **5**(3), 357–386.

Jarociński, M. & Karadi, P. (2020), 'Deconstructing monetary policy surprises: The role of information shocks', *American Economic Journal: Macroeconomics* **12**(2), 1–43.

Jurado, K., Ludvigson, S. C. & Ng, S. (2015), 'Measuring uncertainty', *American Economic Review* **105**(3), 1177–1216.

Kelly, B. & Jiang, H. (2014), 'Tail risk and asset prices', *The Review of Financial Studies* **27**(10), 2841–2871.

Koenker, R. & Machado, J. A. (1999), 'Goodness of fit and related inference processes for quantile regression', *Journal of the American Statistical Association* **94**(448), 1296–1310.

Lenza, M. & Primiceri, G. E. (2022), 'How to estimate a var after march 2020', *Journal of Applied Econometrics* **37**(4), 688–699.

Li, D., Plagborg-Møller, M. & Wolf, C. K. (2022), 'Local projections vs. vars: Lessons from thousands of dgps', *Working Paper* .

McCracken, M. W. & Ng, S. (2016), 'Fred-md: A monthly database for macroeconomic research', *Journal of Business & Economic Statistics* **34**(4), 574–589.

Miranda-Agrippino, S. & Ricco, G. (2021*a*), 'Bayesian local projections', *Sciences Po OFCE Working Paper No. 13/2021* .

Miranda-Agrippino, S. & Ricco, G. (2021*b*), 'The transmission of monetary policy shocks', *American Economic Journal: Macroeconomics* **13**(3), 74–107.

Nakamura, E. & Steinsson, J. (2018), 'High-frequency identification of monetary non-neutrality: the information effect', *The Quarterly Journal of Economics* **133**(3), 1283–1330.

Nakamura, E., Steinsson, J., Barro, R. & Ursúa, J. (2013), 'Crises and recoveries in an empirical model of consumption disasters', *American Economic Journal: Macroeconomics* **5**(3), 35–74.

Ng, S. (2021), 'Modeling macroeconomic variations after covid-19', *NBER Working Paper No. 29060* .

Plagborg-Møller, M., Reichlin, L., Ricco, G. & Hasenzagl, T. (2020), 'When is growth at risk?', *Brookings Papers on Economic Activity, Spring 2020* .

Plagborg-Møller, M. & Wolf, C. K. (2021), 'Local projections and vars estimate the same impulse responses', *Econometrica* **89**(2), 955–980.

Plagborg-Møller, M. & Wolf, C. K. (forthcoming), 'Instrumental variable identification of dynamic variance decompositions', *Journal of Political Economy* .

Rietz, T. A. (1988), 'The equity risk premium a solution', *Journal of Monetary Economics* **22**(1), 117–131.

Romer, C. D. & Romer, D. H. (2004), 'A new measure of monetary shocks: Derivation and implications', *American Economic Review* **94**(4), 1055–1084.

Staiger, D. & Stock, J. H. (1997), 'Instrumental variables regression with weak instruments', *Econometrica* **65**(3), 557–586.

Stock, J. H. & Watson, M. W. (2018), 'Identification and estimation of dynamic causal effects in macroeconomics using external instruments', *The Economic Journal* **128**(610), 917–948.

Wachter, J. A. (2013), 'Can time-varying risk of rare disasters explain aggregate stock market volatility?', *The Journal of Finance* **68**(3), 987–1035.

Appendix A Consumption forecasting comparison

To find the best out-of-sample forecasting model for consumption growth quantiles, I compare the out-of-sample forecasting accuracy across a range of different model alternatives. The goal is to forecast the 5, 25, 50, 75 and 95 percent quantiles of the distribution of consumption growth over the next year at a monthly frequency.[26] In a given month, I define year-over-year consumption growth as the percent change in U.S. aggregate real personal consumption expenditures (excluding durables) over the past twelve months.

I consider twenty different forecasting models. All models condense the information from a large set of predictors into factor estimates and use these factors as predictors in quantile regressions as in equation (1). The data comes from the FRED-MD database by McCracken & Ng (2016), which contains 128 monthly macroeconomic and financial variables. The sample period covers January 1960 to December 2020. The first out-of-sample forecast is made in January 1970 for year-over-year consumption growth in January 1971. With each new month, I extend the sample by the additional information from the new month and re-estimate the factors.

The first eight models use as predictors the first one, two, ..., eight principal components of the data.[27] I call these models PC1, PC2, ..., PC8. The method of principal components estimates factors to explain the variation among the set

[26]Predictive quantile regressions can lead to quantile crossings (e.g. the 5% quantile being predicted above the 10% quantile). To address this issue, I estimate a fine grid of quantiles with $\theta \in \{0.05, 0.06, ..., 0.94, 0.95\}$ and use the Chernozhukov et al. (2010) rearrangement method to ensure monotonicity of the quantile function. In evaluating the forecasting performance, I only focus on the five quantiles mentioned.

[27]The Bai & Ng (2002) PC_{p2} criterion indicates eight significant factors for the vast majority of periods.

of predictors.[28] Following Bai & Ng (2008*b*), we write for observations a_{it} with $i = 1, ...N$ and $t = 1, ..., T$

$$a_{it} = \lambda_i' F_t + e_{it}, \tag{13}$$

where λ_i is a $w \times 1$ vector of factor loadings, F_t is a $w \times 1$ vector of (static) factors, and e_{it} is the idiosyncratic error term. Let $A_t = (a_{1t}, a_{2t}, ..., a_{Nt})'$ be a $N \times 1$ vector, $\Lambda = (\lambda_1, ..., \lambda_N)'$ a $N \times w$ matrix, and $e_t = (e_{1t}, ..., e_{Nt})'$. Then, we have

$$A_t = \Lambda F_t + e_t. \tag{14}$$

Further, if we define the $T \times N$ matrices $A = (A_1, ..., A_T)'$ and $e = (e_1, ..., e_T)'$, and the $T \times w$ matrix $F = (F_1, ..., F_T)'$, we obtain the matrix expression

$$A = F\Lambda' + e. \tag{15}$$

Let $\tilde{F}^w$ be a $T \times w$ matrix collecting the estimated factors for all periods, and $\tilde{\Lambda}^w$ be a $N \times w$ matrix of corresponding factor loadings. As described in Bai & Ng (2008*b*), I obtain $\tilde{\Lambda}^w = \sqrt{N}\text{eig}^w\left(A'A\right)$, where $\text{eig}^w(\cdot)$ collects the eigenvectors corresponding to the w largest eigenvalues of the $N \times N$ matrix $A'A$. Under the normalization $\tilde{\Lambda}^{w\prime}\tilde{\Lambda}^w/N = I_w$, we estimate the factors as $\tilde{F}^w = A\tilde{\Lambda}^w/N$. The factors can then be used as predictors in the quantile forecasting problem.

The method of principal components ensures that the w factors selected are those that describe the most variation of the data A. This allows to capture much of the information in the set of predictors with a considerably smaller number of factors. However, when constructing the factors, this method does not take into account the extent to which individual predictors have predictive power for the variable of interest. In forecasting, we may instead wish to construct factors by explicitly

[28] Before estimating the factors from the FRED-MD data set, I use the transformations suggested in McCracken & Ng (2016) to make all series stationary. This loses the first two observations since some series are second differenced. I also remove outliers and remove all series with missing values, which loses up to five variables.

acknowledging that certain variables are good predictors for the variable of interest while others are not, and apply different weights accordingly.

As described in the main text, partial quantile regression (PQR) by Giglio et al. (2016) constructs a single factor to predict the dependent variable. The factor is a weighted average of the predictors with weights determined by their predictive power from a univariate quantile regression of the dependent variable on the predictor in question. Model 9 applies PQR on the entire FRED-MD sample. Since using so many predictors in the PQR may be too noisy, an alternative model estimates a factor via PQR using as input data the eight factors estimated via the method of principal components. I call this model PQR-PC.

Since some variables in the FRED-MD dataset are highly correlated, estimating factors from the entire set of predictors may lead to an unwanted overweighting of certain features of the data. Further, with a large number of 'noisy' time series in the data, the average size of the common component can be small and the cross-correlation of idiosyncratic error terms can be large. Since the asymptotic theory of the method of principal components requires a large size of the common component and weak cross-correlation in idiosyncratic errors, pre-selecting the number of variables can improve the factor estimation and lead to a more precise forecasting model. Boivin & Ng (2006) find that factors constructed from 40 pre-screened variables can yield better forecasting performance than factors estimated from a large dataset with over 100 macroeconomic time series.

Models 11 to 18 pre-select 40 variables from the entire sample before estimating the $1, 2, ..., 8$ factors. I denote these models as PC1$*$, PC2$*$, ..., PC8$*$. The variables are chosen via ℓ_1-penalized quantile regression, see equation (4) in the main text. Model 19 estimates a single factor using partial quantile regression from the pre-selected set of 40 predictors (PQR$*$). Model 20 estimates a single factor from the eight factors of model 18 (PQR-PC$*$).

To compare the out-of-sample forecasting performance across the different models, I compute the Koenker & Machado (1999) R^1, which transfers the logic of the standard R^2 measure used in OLS to quantile regression. For the θ-th quantile, $R^1(\theta)$ compares the sum of weighted prediction errors between an unconstrained model (our forecasting model) and a constrained model (the historical sample average):

$$R^1(\theta) = 1 - \frac{\sum_t \rho_\theta \left(c_t - x_t' \hat{\beta}_{\theta,t}\right)}{\sum_t \rho_\theta \left(c_t - \mathbf{1}' \bar{\beta}_{\theta,t}\right)}, \tag{16}$$

where $\rho_\theta(z) = z(\theta - \mathbb{I}(z < 0))$ is the quantile loss function, $\hat{\beta}_{\theta,t}$ are the coefficients estimated using the data available in period t, and $\mathbf{1}'\bar{\beta}_{\theta,t}$ is the historical sample average of the θ-th quantile. For in-sample estimations, R^1 is bound to lie between 0 and 1. For out-of-sample estimations, R^1 will take a negative value if the forecasting model performs worse than the historical sample average.

The results from the out-of-sample forecasting exercise are in Table 4. The value reported for R^1 is the average across the targeted quantiles $\Theta = \{0.05, 0.25, 0.5, 0.75, 0.9\}$. The first column contains the results for the entire sample, from January 1980 until December 2020. The principal component models generally perform poorly and have negative R^1, except for the model with only one factor. Including additional factors as predictors seems to overfit the data. The PQR model has a considerably higher R^1 value of 17.5. The best model is the PQR with pre-selection (R^1 of 21.1). The second column considers the sample starting in February 1984, which coincides with the start of the VAR sample and excludes the turbulent 1970s and early 1980s. The PQR model with pre-selection also outperforms all other models over this shorter sample.

The third column compares the model fit using only the 5, 25 and 50% quantiles, which are most relevant for constructing the measure of downside risk. The PQR model again demonstrates the best fit when combined with pre-selecting the predictors. All models generally perform poorer for the lower quantiles than for the upper quantiles, which is evident from lower R^1 values in the third column. This could reflect the higher variability and existence of stronger feedback loops/nonlinearities for the lower quantiles, and stresses the importance of choosing an accurate forecasting model. The diffusion index models (PC1 to PC8 and PC1∗ to PC8∗), which are popular choices for mean forecasts, can dramatically underperform relative to the sample average and lead to misguided conclusions.

Table 4: Forecasting Model Comparison

	Full sample	Shorter sample	Lower quantiles
PC1	5.0	4.9	7.2
PC2	-6.4	-4.8	-10.5
PC3	-16.1	-15.6	-25.0
PC4	-14.1	-13.6	-23.2
PC5	-10.9	-11.9	-18.9
PC6	-9.7	-8.8	-19.2
PC7	-7.2	-6.0	-17.0
PC8	-8.1	-6.4	-18.4
PQR	17.5	15.6	7.8
PQR-PC	-4.5	-2.3	-11.3
PC1∗	6.1	3.9	8.2
PC2∗	-9.1	-10.0	-14.9
PC3∗	-10.5	-11.4	-17.0
PC4∗	-8.7	-9.2	-14.9
PC5∗	-6.1	-7.1	-11.8
PC6∗	-5.5	-6.3	-11.3
PC7∗	-4.8	-5.9	-11.1
PC8∗	-3.5	-4.7	-9.5
PQR∗	**21.1**	**19.7**	**13.0**
PQR-PC∗	-2.0	-2.4	-6.3

The value in the table is the average value of R^1 across the targeted quantiles, in percent. The highest value of each column is in bold face. The full sample is 01/1980-12/2020. The shorter sample starts in 02/1984 to coincide with the start of the VAR sample. The third column contains results for the average fit over the 5, 25 and 50% quantiles.

Appendix B Selected predictors for consumption forecasting

Which predictors are particularly useful in predicting quantiles of the consumption growth distribution? Table 5 lists the ten most frequently selected predictors from the pre-selection via ℓ_1-penalized quantile regression. The left column considers the most frequently chosen predictors for the 5% quantile. Real personal income (RPI), the ratio of inventories to sales among all business (ISRATIOx), and the short-term Treasury spread (TB3SMFFM) are included in the set of predictors most often. Three labor market measures also seem to contain useful information for the lower consumption quantiles: the U.S. unemployment rate (UNRATE), the number of civilians unemployed 27 weeks or more (UEMP27OV), and the average weekly overtime hours in the manufacturing industry (AWOTMAN). The list is completed by housing starts in the South (HOUSTS), the US-Canadian exchange rate (EXCAUSx), and two credit spreads. The first is the spread between the 3-month commercial paper rate and the federal funds rate (COMPAPFFx), and the second is Moody's Baa corporate bond yield over the federal funds rate (BAAFFM).

The results for the median and the 95% quantile include similar variables, although the ranking can differ and some additional variables are among the most frequent predictors. For example, the S&P 500 price-earning ratio and several indicators of housing starts (PERMITS, PERMITMW, PERMITW) make the list for the median. The top ten for the 95% quantile includes three foreign exchange rates (US-Canada, US-Japan, and US-Switzerland). Additional variables are the ratio of nonrevolving consumer credit to personal income (CONSPI), the amount of consumer motor vehicle loans outstanding (DTCOLNVHFNM), and the amount of unfilled orders for durable goods (AMDMUOx).

Two qualifications are in order. First, the ℓ_1-penalty usually selects only one variable out of a group of highly correlated variables and disregards the rest. However, the estimator does not care which variable exactly is chosen (Bai & Ng (2008*a*), p. 307). Therefore, taking the list of selected predictors literally is not advisable. Instead, each predictor should be viewed as being replaceable by another highly correlated variable representing similar information about the economy. Second, the frequency with which a predictor is included in the subset of 40 targeted predictors only provides indirect information about the importance of that predictor for a consumption quantile. The strength of a predictor is ultimately given by the weight it obtains in the factor estimation step of the partial quantile regression.

To provide a more thorough picture of the type of variables useful for predicting different quantiles of consumption growth, I first categorize the set of predictors into eight groups following McCracken & Ng (2016). The groups are 1) output and income, 2) labour market, 3) housing, 4) consumption, orders, and inventories, 5) money and credit, 6) interest and exchange rates, 7) prices, and 8) stock market. Given these groups, for each month I compute the share of the 40 selected predictors which fall into each of the eight groups. The results are shown in Figures 10, 11 and 12 for the 5% quantile, median, and 95% quantile, respectively.

Interest and exchange rates as well as indicators of money and credit play a prominent role as predictors for the 5% quantile (especially in the first half of the sample), which is in line with this paper's results on the effect of monetary policy on downside risk to consumption. Labour market variables also feature frequently as selected predictors. The overall composition of the selected subset of predictors varies strongly over time for the 5% quantile but is more stable for the median and the 95% quantile. Interest and exchange rates are selected less frequently as predictors for the higher quantiles, whereas output and income play an increasing role. Housing market variables are picked rarely as predictors for the 95% quantile but more often for the median and 5% quantile. In contrast, prices are picked rarely as predictors for the 5% quantile but more often for the median and the 95% quantile. Variables representing consumption, orders, and inventories represent only a minor share of predictors across all quantiles. This casts doubt on any approach trying to predict the consumption growth distribution using exclusively information about consumption itself. However, note that individual variables out of any of these groups may still be selected frequently as a predictor even if the group only represents a small share of the selected predictors. For example, while consumption, orders, and inventories only account for a small share of predictors, the inventories-to-sales ratio features prominently as a predictor for all quantiles, see Table 5.

Overall, the results from this chapter demonstrate that the consumption growth distribution is best predicted using a wide range of predictors capturing information about different aspects of economic conditions. This validates the estimation approach taken in this paper, which starts from a large data set to consider as many potential predictors as possible. It also means that the transmission channel from monetary policy to downside risk is not immediately clear: While interest rates feature strongly as predictors of the consumption growth distribution, the analysis also demonstrates that many other factors, with which monetary policy may interact, forecast future consumption growth quantiles. More research is necessary to disentangle these components.

Table 5: Most frequently selected predictors

	5% Quantile		Median		95% Quantile	
	Name	Frequency	Name	Frequency	Name	Frequency
1	ISRATIOx	99.8	ISRATIOx	94.3	RPI	84.4
2	RPI	98.4	EXCAUSx	92.5	CONSPI	83.7
3	TB3SMFFM	94.1	PERMITMW	84.5	EXCAUSx	82.5
4	EXCAUSx	92.9	PERMITS	83.5	DTCOLNVHFNM	77.9
5	COMPAPFFx	84.2	S&P PE ratio	81.7	EXJPUSx	76.0
6	UNRATE	83.7	HWIURATIO	81.3	EXSZUSx	75.4
7	HOUSTS	78.7	CES1021000001	78.7	ISRATIOx	73.1
8	AWOTMAN	75.2	PPICMM	77.2	CES1021000001	72.0
9	UEMP27OV	73.8	CONSPI	77.0	AMDMUOx	70.1
10	BAAFFM	68.5	PERMITW	75.8	PERMITS	67.9

Legend: AMDMUOx – Unfilled orders for Durable Goods. AWOTMAN – Average weekly overtime hours: manufacturing. BAAFFM – Moody's Baa Corporate Bond Minus Fed Funds. CES1021000001 – All employees: mining and logging: mining. COMPAPFFx – 3-month commercial paper minus federal funds rate. CONSPI – Nonrevolving consumer credit to personal income. DTCOLNVHFNM – Consumer motor vehicle loans outstanding. EXCAUSx – Canada/US foreign exchange rate. EXJPUSx – Japan/US foreign exchange rate. EXSZUSx – Switzerland/US foreign exchange rate. HOUSTS – Housing Starts, South. HWIURATIO – Help wanted to number of unemployed. ISRATIOx – Total business: inventories to sales. PERMITMW – New private housing permits, Midwest (SAAR). PERMITW – New private housing permits, West (SAAR). PERMITS – New private housing permits, South (SAAR). PPICMM – PPI: Metals and metal products. RPI – Real personal income. S&P P/E Ratio - P/E ratio on S&P composite index. TB3SMFFM – 3-month treasury yield minus federal funds rate. UEMP27OV – Civilians Unemployed for 27 weeks and over. UNRATE – Civilian unemployment rate.

Figure 10: Composition of selected predictors - 5% Quantile

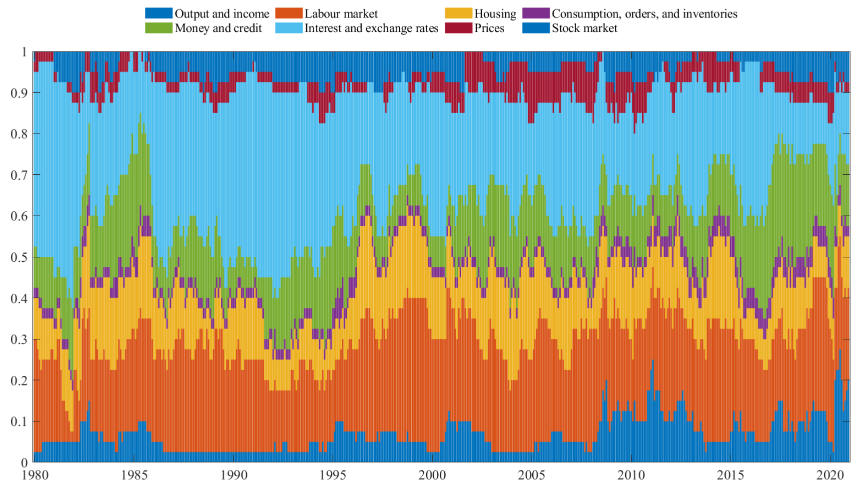

Note: For each month, the figure plots the share of the 40 selected predictors for the 5% quantile coming from different variable groups. The group definitions are as in the appendix of McCracken & Ng (2016).

Figure 11: Composition of selected predictors - Median

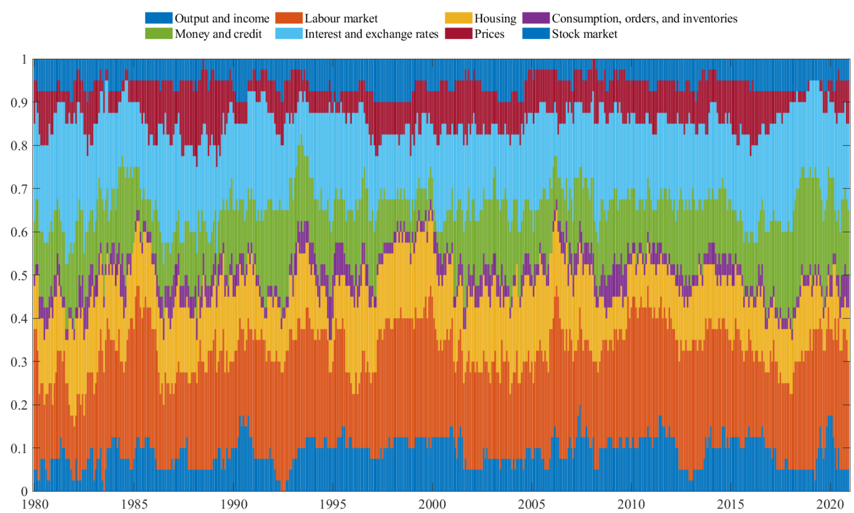

Note: For each month, the figure plots the share of the 40 selected predictors for the median coming from different variable groups. The group definitions are as in the appendix of McCracken & Ng (2016).

Figure 12: Composition of selected predictors - 95% Quantile

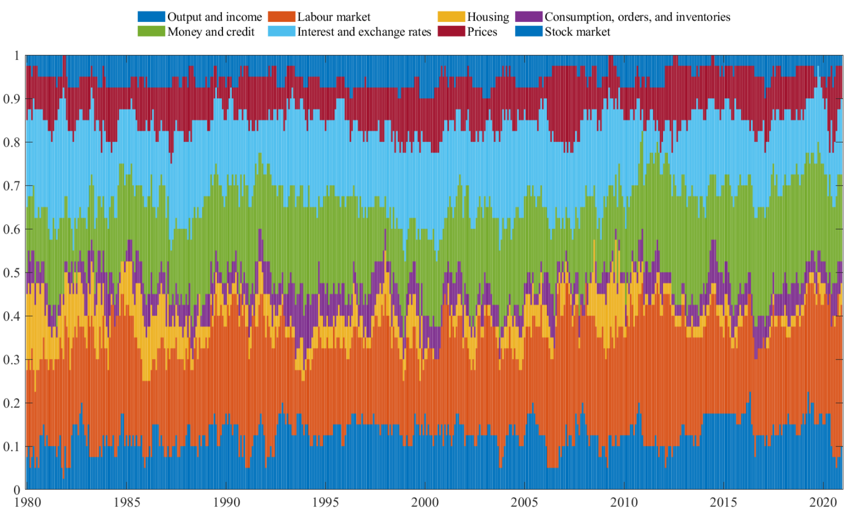

Note: For each month, the figure plots the share of the 40 selected predictors for the 95% quantile coming from different variable groups. The group definitions are as in the appendix of McCracken & Ng (2016).

Appendix C Details on the consumption growth distribution and downside risk

Figure 13 provides further evidence on the accuracy with which the skewed t-distribution approximates the estimated quantiles. The solid blue lines are the 5, 50, and 95% quantiles estimated in the out-of-sample forecasting exercise. The red dashed lines are the corresponding values for those quantiles implied by the skewed t-distribution fitted to the estimated quantiles at a given point in time. As described in the main text, I fit the distribution to the estimated quantiles based on the 5, 25, 50, 75, and 95% quantiles. I only plot the 5, 50, and 95% quantiles here to keep the figure simple. The implied quantiles closely track the estimated quantiles. The skewed t-distribution is flexible enough to fit the different quantiles closely in each period, irrespective of the behaviour of the quantiles. This supports the use of the skewed t-distribution to approximate the conditional distribution of expected consumption growth.

Figure 14 shows the change in the downside and upside risk index relative to the value 12 months ago. Since the risk index in levels does not always yield stationarity in the VAR, I use these 12-month differences in the VAR analysis. For consistency, this transformation is also used in the predictive regressions.

Figure 15 shows the downside and upside entropy for a forecasting horizon of 3 months as opposed to 12 months. Similar to the downside risk index for a 12-month forecasting horizon, the 3-month index spikes during the early 1980s recession and the Great Recession. The index is mostly muted between the mid-1980s and the mid-2000s, reflecting the Great Moderation. The index starts to rise before the Great Recession and then falls with the start of the Great Recession after the Fed rate cuts, even though this increase and subsequent decline are less pronounced than for the 12-month index. The index spikes following the Lehman collapse and then starts a persistent decline with the end of the Great Recession. The biggest contrast with the 12-month index is the steep rise and decline of the 3-month downside risk index around the Covid recession. The steep and sudden increase in short-term downside risk reflects the fear about lockdown measures and the initial drop in aggregate consumption growth. The forecasting model is flexible enough to pick up this dramatic change in economic outlook. At the same time, short-term downside risk falls strongly after its peak in April 2020, supporting the information contained in the 12-month index: While short-term downside risk was very high, the one-year outlook was not as bad as during other recessions.

Figure 13: Estimated vs Implied Quantiles

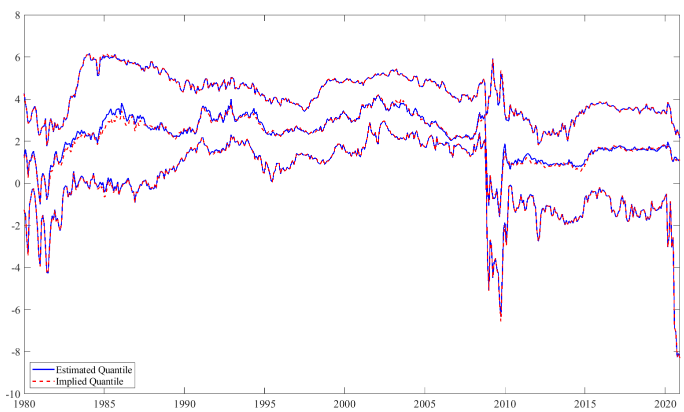

Note: The blue lines are the 5, 50 and 95% quantiles of year-over-year consumption growth from the one-year ahead forecasting exercise. The red dashed lines are the corresponding values for the quantiles implied by the skewed t-distribution fitted across the estimated quantiles. The sample period is 1:1980-12:2020.

Figure 14: 12-month changes of Downside/Upside Risk

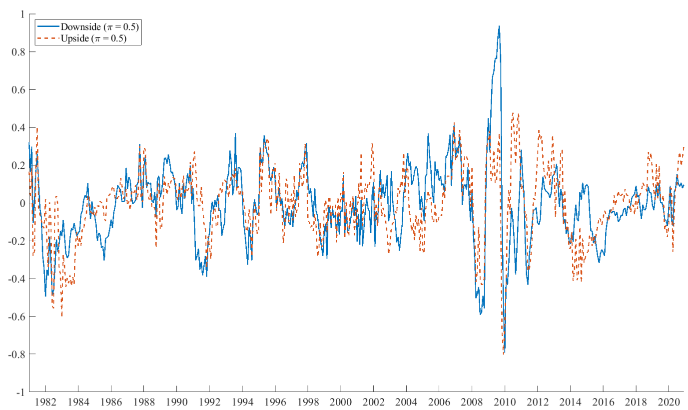

Note: The blue line shows the year-over-year difference in downside entropy for U.S. consumption growth. The red line shows the year-over-year difference in upside entropy. The sample period is 1:1981-12:2020.

Figure 15: Downside/Upside Risk for three-month horizon

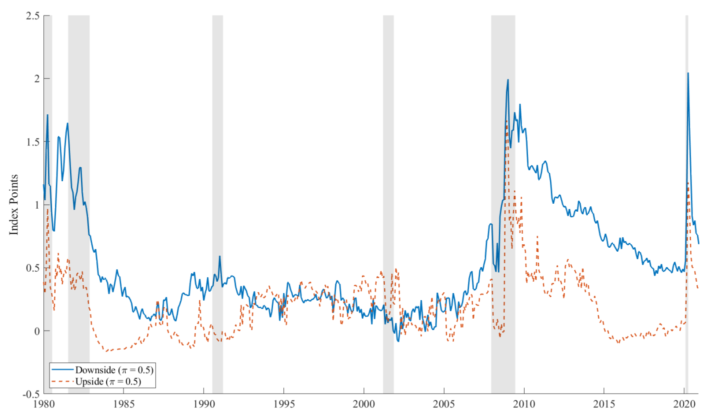

Note: The blue line shows the downside entropy for U.S. consumption growth with a forecasting horizon of three months. The red line shows the corresponding upside entropy. The sample period is 1:1980-12:2020.

Appendix D Details on external instrument VAR

Given an instrument for the monetary policy shocks, we can estimate the effect of a monetary policy shock on downside risk using the external instruments VAR approach of Gertler & Karadi (2015). In the following, I provide an overview of their methodology. Consider a reduced form VAR(p) with shock $u_t = S\varepsilon_t$, where ε_t is the structural shock and $S \equiv \Phi^{-1}$. We let $B_j = SA_j$ be the matrix of coefficients and $\mu = S\mu^*$ such that

$$Y_t = \mu + \sum_{j=1}^{p} B_j Y_{t-j} + u_t. \tag{17}$$

We decompose the structural shock into the monetary policy shock ε_t^p and all other shocks ε_t^q such that $\varepsilon_t = \left[\varepsilon_t^p, \varepsilon_t^{q\prime}\right]'$. Analogously, we can decompose the reduced-form errors $u_t = \left[u_t^p, u_t^{q\prime}\right]'$. Focusing on the impulse responses to a monetary shock ε_t^p, we write

$$Y_t = \mu + \sum_{j=1}^{p} B_j Y_{t-j} + s\varepsilon_t^p \tag{18}$$

where s is the column of matrix S that determines the effect of ε_t^p on u_t. Let Z_t be a vector of instrumental variables such that $\mathbb{E}\left[Z_t \varepsilon_t^p\right] = \Omega$ (instrument relevance) and $\mathbb{E}\left[Z_t \varepsilon_t^{q\prime}\right] = \mathbf{0}$ (instrument exogeneity). In my case, the instrumental variable used is the high-frequency change in the implied interest rate from three-month federal funds futures around FOMC announcements.

We obtain the impulse responses as follows. First, estimate the reduced form VAR via OLS to get estimates for all B_j and the residuals u_t^p and u_t^q. Let s^q be the response of u_t^q to monetary policy shock ε_t^p such that $u_t^q = s^q \varepsilon_t^q$. Then, perform a two-stage least squares regression of u_t^q on u_t^p using instrument Z_t:

$$u_t^q = c + \frac{s^q}{s^p}\hat{u}_t^p + \xi_t, \tag{19}$$

where $\hat{u}_t^p$ is the fitted value from the first stage regression of u_t^p on Z_t. As explained in Gertler & Karadi (2015), the first stage isolates the variation in the reduced form residual for the policy indicator u_t^p that is due to the structural policy shock. The second stage splits u_t^q into a part driven by the policy shock ($s^q\varepsilon_t^p$) and a residual (ξ_t), and replaces ε_t^p with the fitted value from the first stage ($\hat{u}_t^p \approx s^p\varepsilon_t^p \Rightarrow \varepsilon_t^p \approx \frac{\hat{u}_t^p}{s^p}$). Gertler & Karadi (2015) show how to isolate s^q and s^p from equation (19) using the reduced-form variance-covariance matrix. Given estimates for s^p, s^q and B_j, we obtain the impulse responses to a monetary policy shock using the reduced-form VAR in equation (18).

Appendix E Details on monetary policy shocks

To construct the monetary policy shock series, I follow the standard approach in the literature and consider the change in three-month federal funds futures prices 10 minutes before the announcement relative to 20 minutes after the announcement.[29] I consider two alternative measures: The "poor man's sign restriction" shocks of Jarociński & Karadi (2020) and the shocks of Miranda-Agrippino & Ricco (2021*b*). Both approaches aim to identify the policy shock after having removed any Fed information effects that may pollute the information contained in high-frequency Fed funds futures returns.

For the Jarociński & Karadi (2020) approach, I obtain a series of high-frequency changes in federal funds futures and the S&P 500 for the period February 1990 until December 2017, which was kindly provided to me by Refet S. Gürkaynak. From January 2018 until December 2020, I construct the high-frequency changes myself using 1-minute price data on electronically traded futures from FirstRateData. I obtain the FOMC meeting dates and times from a file kindly shared with me by Marek Jarociński.[30] The meeting dates include intermeeting decisions such as the announcements of asset purchases.

For each announcement, I identify the change in the price of the federal funds futures contract that expires in the third month after the FOMC announcement. Denote the price of this contract 10 minutes before the announcement as $f_{t,t+3}$ and 20 minutes after the announcement as $f_{t+\Delta,t+3}$. The implied rate (in percent) of a Fed funds futures contract is 100 minus the current price. The implied interest rate change around the announcement is then $\Delta i_t = (100 - f_{t+\Delta,t+3}) - (100 - f_{t,t+3}) = f_{t,t+3} - f_{t+\Delta,t+3}$.

To account for Fed information effects, Jarociński & Karadi (2020) propose to use sign restrictions: If interest rate cuts are perceived as expansionary policy shocks,

[29] Three-month federal funds futures measure the expected average federal funds rate over the third calendar month ahead. This horizon captures surprises to the short rate and forward guidance.

[30] The FOMC meeting time is the time of the press release. For the period 2018-2020 I take the press release times from the Federal Reserve Board website.

they should be associated with rises in stock markets such that the interest rate and stock prices move in opposite directions. The "poor man's sign restrictions" set policy surprises equal to zero if the implied rate change and stock prices do not move in opposite directions around FOMC announcements. To implement this approach, I obtain the changes in the S&P 500 index value in the same 30-minute window around FOMC announcements as for federal funds futures, and apply the sign restrictions accordingly. The data for S&P 500 prices until 2017 comes from the dataset provided by Refet S. Gürkaynak. The data for 2018-2020 is at a 1-minute frequency and comes from FirstRateData. Lastly, to obtain a time series of policy shocks, I sum up the interest rate changes within the same month. Months without a FOMC announcement have a value of zero. This yields a sample covering February 1990 until December 2020.[31]

For the Miranda-Agrippino & Ricco (2021*b*) shocks, I start from the same high-frequency changes in federal funds futures rates around FOMC announcements. Instead of using stock market information, the authors propose to cleanse the high-frequency shocks from information effects by projecting it on the Fed's private assessment of the macroeconomic outlook and the lags of the shock series. Similar to Romer & Romer (2004), the Fed's private information is proxied by the latest Greenbook projections before the FOMC decision.[32] Data on the Greenbook projec-

[31] Jarociński & Karadi (2020) aggregate the individual shocks to a monthly frequency first and then apply the sign restrictions, whereas I apply the sign restriction at the level of each FOMC meeting and then aggregate to a monthly frequency. I choose the latter approach since it allows for a more granular distinction between "true" monetary policy shocks and Fed information effects. Especially during months with multiple Fed decisions my approach allows for the presence of both a monetary policy and a Fed information effect, whereas the original Jarociński & Karadi (2020) effect restricts one of the effects to be zero. My approach also proves to be more flexible during the recent Covid-19 crisis: While I get a strongly expansionary effect using my approach, the Jarociński & Karadi (2020) approach restricts the monetary shock to be zero since the S&P 500 overall fell following the March 2020 Fed announcements.

[32] While Romer & Romer (2004) only focus on scheduled FOMC meetings, Miranda-Agrippino & Ricco (2021*b*) also consider intermeeting decisions. They acknowledge that the Greenbook projections do not fully capture the Fed's information before intermeeting decisions because Greenbook forecasts are only made prior to scheduled FOMC meetings, whereas the arrival of *new* information between scheduled meetings is likely to have caused the intermeeting session. I do not resolve this issue in my work.

tions is from the Philadelphia Fed's Greenbook Data Set. To obtain the component of the high-frequency federal funds futures rate change around a certain FOMC meeting (Δi_m) that cannot be forecast by the Fed's private information, I regress the high-frequency change on the latest Fed's Greenbook projections:

$$\begin{aligned}\Delta i_m = \alpha &+ \sum_{k=-1}^{3} \beta_k \Delta\tilde{y}_{m,k} + \sum_{k=-1}^{2} \gamma_k \left(\Delta\tilde{y}_{m,k} - \Delta\tilde{y}_{m-1,k}\right) \\ &+ \sum_{k=-1}^{3} \lambda_k \tilde{\pi}_{m,k} + \sum_{k=-1}^{2} \varphi_k \left(\tilde{\pi}_{m,k} - \tilde{\pi}_{m-1,k}\right) \\ &+ \sum_{k=-1}^{3} \theta_k \tilde{u}_{m,k} + \sum_{k=-1}^{2} \psi_k \left(\tilde{u}_{m,k} - \tilde{u}_{m-1,k}\right) + MPI_m,\end{aligned} \tag{20}$$

where $\Delta\tilde{y}_{m,k}$ is the Greenbook projection before FOMC meeting m for real GDP growth k quarters ahead. Analogously, $\tilde{\pi}_{m,k}$ denotes the forecast for the inflation rate and $\tilde{u}_{m,k}$ for the unemployment rate. Note that the unit of observation for this regression is FOMC meetings, not months or days. The residual MPI_m is the unforecastable change in the federal funds futures rate around FOMC meeting m. I sum all residuals for FOMC meetings that occur within the same month to obtain a monthly series. All months without a FOMC decision are assigned a value of zero. Let $\overline{MPI}_t$ denote the value of this new time series in a given month t.

With imperfect information, markets can be slow in absorbing new information from policy shocks. Therefore, changes in market prices can be autocorrelated as prices not only respond to information about the current, but also to past shocks. To purge the shock series from this autocorrelation, Miranda-Agrippino & Ricco (2021*b*) regress the policy shock in a given period on its past lags:

$$\overline{MPI}_t = \phi_0 + \sum_{j=1}^{12} \phi_j \overline{MPI}_{t-j} + MPI_t \tag{21}$$

The regression uses only observations corresponding to non-zero values of the dependent variable $\overline{MPI}_t$. The residual MPI_t is the final policy shock. Since the

Greenbook projections are published with a lag of five years, the shock series is only available until 2015. Correcting for slow information absorption loses the first twelve non-zero observations of the series since we use twelve lags in the autoregression. The sample period is then February 1991 until December 2015. Figure 16 compares the Jarociński & Karadi (2020) and the Miranda-Agrippino & Ricco (2021*b*) shock series.

Figure 16: Monetary Policy Shock Series over Time

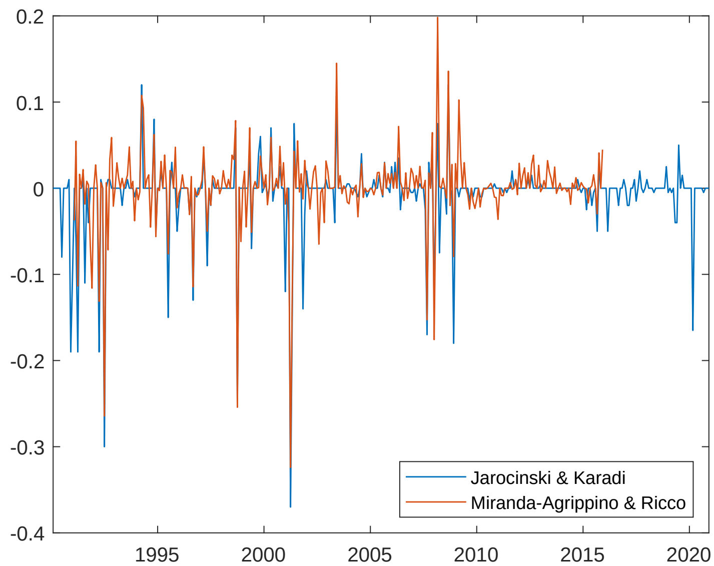

Note: The sample period for the Jarociński & Karadi (2020) shocks is February 1990 until December 2020. The sample period for the Miranda-Agrippino & Ricco (2021*b*) shocks is February 1991 until December 2015.

To check the quality of the monetary policy shock series, I analyze them along several dimensions. First, I check the quality of the raw high-frequency changes in federal funds futures. For the data until 2017, my paper uses the high-frequency changes from an updated version of the Gürkaynak et al. (2005) dataset, which is also used by Jarociński & Karadi (2020). These changes are constructed using tick-by-tick data. However, for the period 2018-2020, I use minute-by-minute data to construct the changes myself. To see if the minute-by-minute data yields comparable results to tick-by-tick data, I construct high-frequency changes from the minute-by-minute data for the period October 2008 until December 2017, which goes as far into the past as my data allows. I then compare the series to the original Gürkaynak et al. (2005) data. The series are similar and I cannot reject the null hypothesis that the mean difference between my shocks and the Gürkaynak et al. (2005) shocks is different from zero for the period October 2008 until December 2017.

I also compare my version of the Miranda-Agrippino & Ricco (2021*b*) shocks to the original shock series constructed by the authors. When comparing the overlapping sample from February 1991 until December 2009, I find no statistically significant average difference between the two series. Both shock series are very similar, even though I ran the autoregressive regression in the last step of constructing the shocks using data from 1991 until 2015, whereas the authors' sample stopped in 2009.

Next, I verify that my updated versions of the Jarociński & Karadi (2020) and Miranda-Agrippino & Ricco (2021*b*) series have zero mean and no autocorrelation for up to twelve lags. For the Jarociński & Karadi (2020) shocks I also check whether they can be predicted by the same Greenbook variables used in constructing the Miranda-Agrippino & Ricco (2021*b*) shocks. That is, I run equation 20 with my version of the Jarociński & Karadi (2020) shocks as the dependent variable. I can reject the joint null hypothesis that all coefficients are equal to zero at the 5% level. However, if I re-run the regression using only data for scheduled FOMC meetings, I cannot reject the joint null hypothesis that all coefficients are zero, even at the 10% level. Since Greenbook forecasts are only published prior to scheduled meetings, this provides additional support for the poor man's sign restrictions in removing information effects.

Appendix F Non-invertibility and misspecification

SVAR analysis implicitly makes the assumption of (partial) invertibility. Since we are only interested in a monetary policy shock, partial invertibility is a sufficient condition to identify the correct dynamic responses to a monetary shock. Under partial invertibility, the shock of interest can be recovered from current and past macro variables such that knowing the true shock series would not provide additional information to the researcher (Plagborg-Møller & Wolf (2021), p.966 and Stock & Watson (2018), p.919). If this assumption is violated, the estimated IRFs are biased across all horizons (Plagborg-Møller & Wolf (forthcoming), Appendix B.4).

As a robustness check, I estimate an internal instrument VAR, which identifies the correct impulse responses even under non-invertibility (Plagborg-Møller & Wolf (2021)). The internal instrument SVAR orders the policy instrument first in a recursively identified SVAR.[33] The ordering of the variables is (Instrument, CPI, Production, Policy Rate, EBP, Downside Risk, Upside Risk). With this order, the monetary policy indicator has contemporaneous effects on the excess bond premium and the downside risk index. The policy instrument is not contemporaneously affected by any of the other variables but has impact effects on all other variables.

The results are in Figure 17. The impact effect on downside risk is almost identical, despite the different identification scheme. After that, downside risk falls even lower than in the baseline regression and declines by up to 0.2 points, which is almost one standard deviation. This effect may be partly influenced by the overall slightly different dynamics of the internal IV VAR system relative to the baseline VAR. By construction, a shock to the monetary policy instrument lowers the 1-year rate by 25bp on impact. However, in the internal IV VAR the response of the 1-year rate does not revert to zero as quickly as in the external instrument approach, and the rate also loosens again by over 10bp after 6 months. This additional loosening

[33] In population, this approach estimates the same impulse responses as the Local Projection Instrumental Variable (LP-IV) approach estimated via two-stage least squares, see Plagborg-Møller & Wolf (2021).

is associated with a stronger rise in industrial production and the CPI, as well as a stronger fall in the excess bond premium. It may also explain the additional fall in downside risk. In any case, the internal IV approach suggests an even stronger response of downside risk to a monetary shock than suggested by the main results.

To address potential concerns about model misspecification, I re-estimate the

Figure 17: VAR Robustness Exercise - Estimating Internal IV VAR

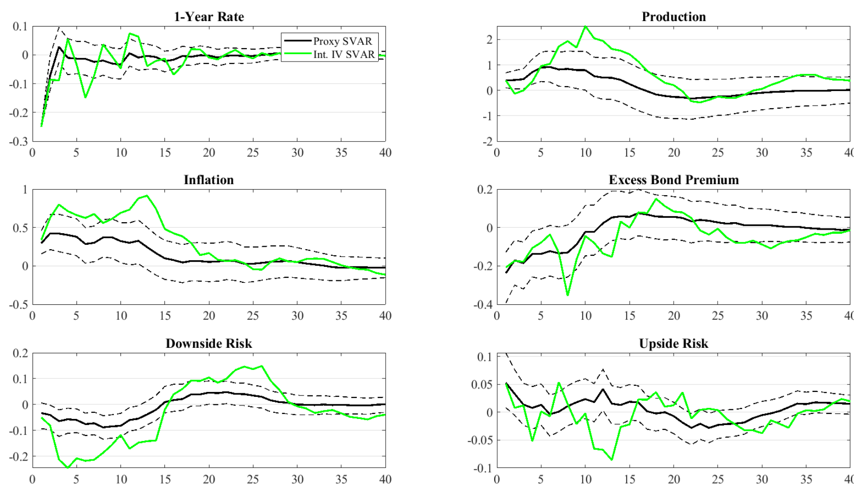

The sample period is 1990:1-2019:12. The controls included are industrial production, the CPI, and the excess bond premium (EBP). The monetary policy indicator is the first-differenced one-year government bond yield, instrumented by the Jarociński & Karadi (2020) shocks. All data is at a monthly frequency, the lag length is 12 months. Dashed lines are 95% confidence bands for the proxy SVAR from a wild bootstrap with 5,000 replications.

impulse responses of the baseline model from local projections identified via the external instrument approach.[34] While local projection and vector autoregression estimate the sample impulse responses in population (Plagborg-Møller & Wolf (2021)), LPs can have lower bias in empirical applications (Li et al. (2022)).[35] Even if the VAR is misspecified, we should expect very similar IRFs up to a horizon equal to lag

[34]Since VAR and LP estimate the same system for a one-step ahead forecast, the identification of the two systems is identical when using external instruments (Miranda-Agrippino & Ricco (2021*a*)).

[35]However, the simulation study by Li et al. (2022) also shows that LPs usually have higher variance. The authors conclude that researchers will generally prefer the VAR approach unless they put a strong weight on bias reduction in the bias-variance trade off.

length p. For horizons larger p, the IRFs may differ markedly. Figure 18 compares the IRFs from the baseline specification with those estimated using local projection. I find that the IRFs are overall very similar, even after horizon p.

Figure 18: VAR Robustness Exercise - Estimating Local Projections

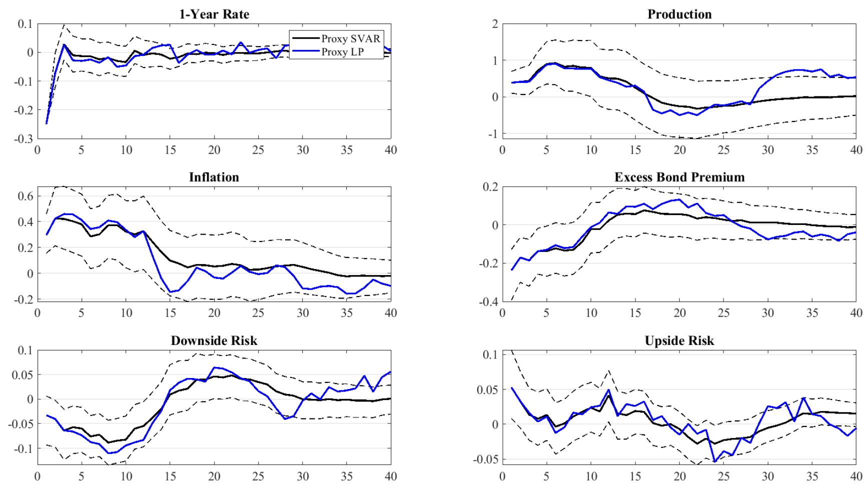

The sample period is 1984:2-2019:12. The controls included are industrial production, the CPI, and the excess bond premium (EBP). The monetary policy indicator is the first-differenced one-year government bond yield, instrument by the Jarociński & Karadi (2020) shocks. All data is at a monthly frequency, the lag length is 12 months. Dashed lines are 95% confidence bands for the proxy SVAR from a wild bootstrap with 5,000 replications.

Appendix G VAR robustness exercises

This section presents several robustness exercises in support of the main VAR result. To save space, I often only report the impulse response of the downside entropy measure and omit the impulse responses for the standard set of controls.

To address concerns about a weak instrument in the main specification, I replace the monetary policy indicator and instrument to obtain a strong instrument. The new policy indicator is the first difference of the effective federal funds rate. The instrument is the high-frequency change in the 3-month federal funds rate, without any adjustments to remove information effects. The first-stage F-statistic is 23.8, well above the rule-of-thumb threshold of 10. Figure 19a shows the results. Figure 19b uses the same combination of indicator and instrument and estimates the impulse responses for the non-recession state. The first-stage F-statistic is 17.6. This confirms that monetary policy shocks have little to no measurable impact on downside consumption risk during normal times.

Figure 19c uses year-over-year differences in the one-year Treasury yield instead of month-over-month differences. The impulse response for downside entropy is similar to the baseline case. However, the robust F-statistic with a value of 6.0 indicates a weak instrument.

Instead of using the Jarociński & Karadi (2020) policy shocks, Figure 19d shows the impulse responses to a policy shock identified via the Miranda-Agrippino & Ricco (2021*b*) shocks. The construction of these shocks is detailed in Appendix E. The results are again similar to the main result. Overall, the choice of policy indicator and instrument therefore seems to make no difference to the main message of the VAR analysis: A loosening monetary policy shock reduces downside consumption risk over the short to medium run.

Figure 20a uses the Jarociński & Karadi (2020) instrument but restricts the sample used for identifying the impact effects to end in December 2007. This excludes the ZLB period, during which the short rate showed little variation. Identifying the policy shock during the ZLB period is complicated, and excluding the period may change the impact effects. However, the impulse responses are unchanged. The first-stage robust F-statistic is 10.7.

Next, I consider two alternative sample periods for the VAR. Figure 20b extends the sample for both the macroeconomic variables and the policy instrument until December 2020, which includes the Covid shock. Including the Covid period does not change the conclusions about the effect of monetary on consumption downside risk. Figure 20c uses a sample from January 1990 until December 2019 to estimate the VAR coefficients. The sample for the instrument is February 1990 until December 2019, which is the same as in the main VAR. Using this shorter sample excludes the entire 1980s, during which a structural break may have occured at some point. The results support the conclusions from the main specification.

Figure 20d uses six instead of twelve lags in the VAR. The impact effects are very similar. The impulse responses now have a smoother path back to zero, but the policy shock still has a significant negative effect on downside risk for about ten months.

Figure 21 shows the impulse responses to a monetary shock from a large-scale VAR with eleven variables. The macro variables are the monetary policy indicator, industrial production, the U.S. employment rate, aggregate real personal consumption expenditures, the PCE deflator, real orders, the real wage, average hours worked, the S&P 500, money stock M2, and the downside risk index of consumption. This is a variation of the macroeconomic VAR in Christiano et al. (2005) inspired by Jurado et al. (2015), who study the effect of uncertainty shocks in a recursively identified macroeconomic VAR. In contrast, I use the VAR to study the effect of a monetary policy shock, which I identify using the Jarociński & Karadi (2020) shocks. This imposes no restrictions on the contemporaneous responses of the variables to a policy shock. The policy shock still has a significant effect on downside risk upon impact, even though the confidence bands are now wider. This may be because of the large number of coefficients to be estimated relative to the sample size. Further, the identification may suffer from a weak instrument as the first-stage F-statistic is only 8.15.

The next four robustness checks consider alternative controls for consumption risk. Figure 22 uses the predicted median from the fitted skewed t-distribution for consumption growth with a forecasting horizon of twelve months. Since the downside entropy measure is computed for all below-median values of the conditional distribution, changes in the median of the conditional relative to the unconditional distribution will be reflected in downside entropy. The goal of the robustness check is to demonstrate that the effect of monetary policy on downside risk is not simply driven by the median of the distribution. Similarly, Figure 22b shows the impulse response of downside entropy when controlling for the expected mean of consumption growth (again taken from the fitted skewed t-distribution). Figure 22c reports the results when controlling for the expected variance. This is particularly relevant since one aim of this paper is to demonstrate the importance of asymmetries between upside and downside risk, which cannot be captured by the variance. In all three cases, the effect of monetary policy on downside consumption risk prevails and is very similar to the main specification.

Figure 22d uses a different measure of downside risk to illustrate the results are not specific to the choice of measuring risk via relative entropy. Given the fitted skewed t-distribution described in the main text, I estimate downside risk as the expected value of all potential outcomes below the median. This risk measure is called *expected shortfall* and given by

$$SF_{t+h}(\pi) = \frac{1}{\pi} \int_0^{\pi} \hat{F}^{-1}_{c_{t+h}|x_t}(\theta|x_t) d\theta, \tag{22}$$

where $\pi = 0.5$ in our case and $\hat{F}_{c_{t+h}|x_t}$ is the conditional distribution of 12-month ahead consumption growth c_{t+h} given the set of predictors x_t. While relative entropy measures the asymmetry of the conditional distribution in excess of the asymmetry exhibited by the unconditional distribution, expected shortfall summarizes downside risk in absolute terms. During a recession, we should expect a higher probability of negative events such that expected shortfall rises. The VAR uses the year-over-year difference in expected shortfall to make the time series stationary. The impulse response for expected shortfall confirms that a loosening in the monetary policy stance raises expected shortfall, thereby lowering downside risk. This effect reverts after about twelve months.

Figure 19: VAR Robustness Exercise - Different monetary policy indicators and instruments

(a) Strong instrument

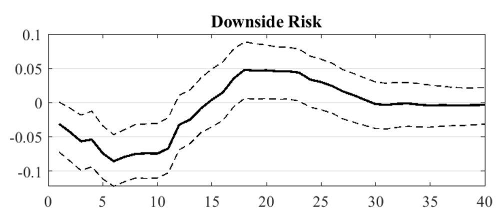

(b) Strong instrument (non-recession state)

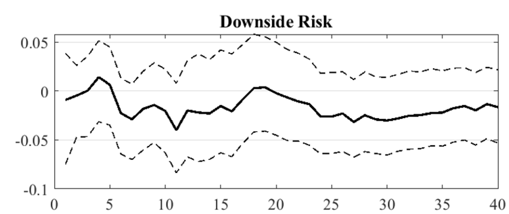

(c) Different policy indicator

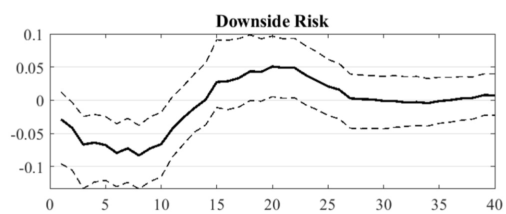

(d) Different policy instrument

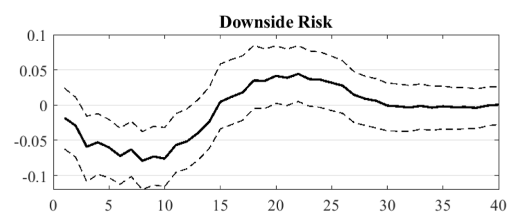

For panels (a) and (b), the monetary policy indicator is the first difference of the effective federal funds rate, instrumented by the raw high-frequency change in the 3-month federal funds rate. For panel (c), the monetary policy indicator is the year-over-year difference in the one-year government bond yield, instrumented by the Jarociński & Karadi (2020) shocks. In panel (d), the monetary policy indicator is the first-differenced one-year government bond yield, instrumented by the Miranda-Agrippino & Ricco (2021*b*) shocks. The sample period is 1984:2-2019:12. The controls included are industrial production, the CPI, and the excess bond premium (EBP). All data is at a monthly frequency, the lag length is 12 months. Dashed lines are 95% confidence bands from a wild bootstrap with 5,000 replications.

Figure 20: VAR Robustness Exercise - Different sample periods and lag lengths

(a) Excluding ZLB period in IV regression

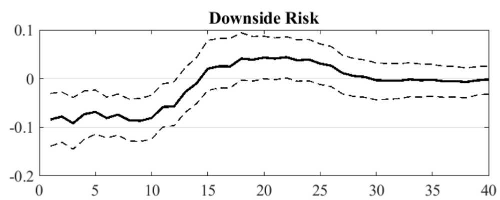

(b) Including the Covid period

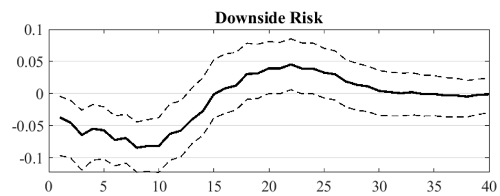

(c) Shorter VAR sample

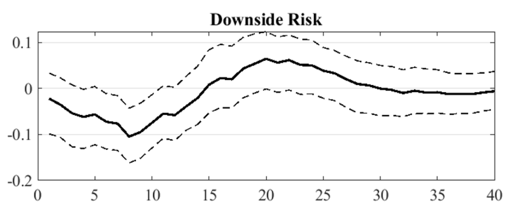

(d) Using only six lags

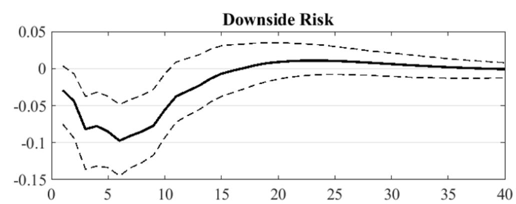

For panel (a), the sample period for the IV regression is 1990:2-2007:12. In panel (b), the sample period for the VAR is 1984:2-2020:12. For panel (c), the sample period is 1990:1-2019:12. For panel (d), the lag length is six months. Unless stated otherwise, the sample period for the VAR is 1984:2-2019:12 and the sample period for the IV regression is 1990:2-2019:12. The controls included are industrial production, the CPI, and the excess bond premium (EBP). All data is at a monthly frequency. Unless stated otherwise, the lag length is 12 months. Dashed lines are 95% confidence bands from a wild bootstrap with 5,000 replications.

Figure 21: VAR Robustness Exercise - 11-variable VAR

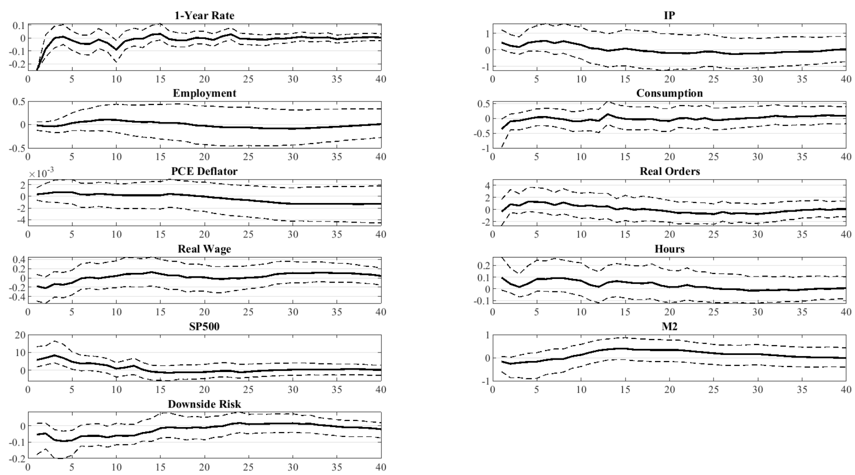

The sample period is 1984:2-2019:12. The controls included are industrial production, the U.S. employment rate, aggregate real personal consumption expenditures, the PCE deflator, real orders, the real wage, average hours worked, the S&P 500, and money stock M2. The monetary policy indicator is the first-differenced one-year government bond yield, instrumented by the Jarociński & Karadi (2020) shocks. All data is at a monthly frequency, the lag length is 12 months. Dashed lines are 95% confidence bands from a wild bootstrap with 5,000 replications.

Figure 22: VAR Robustness Exercise - Consumption risk controls

(a) Controlling for expected median

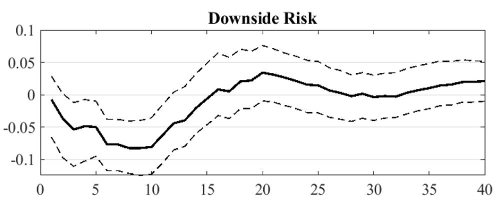

(b) Controlling for expected mean

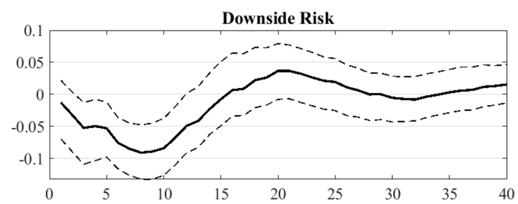

(c) Controlling for expected variance

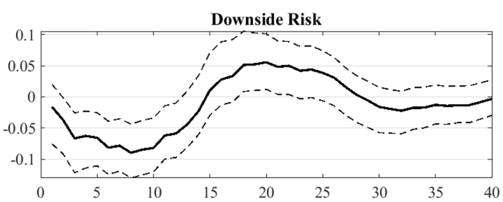

(d) Using expected shortfall

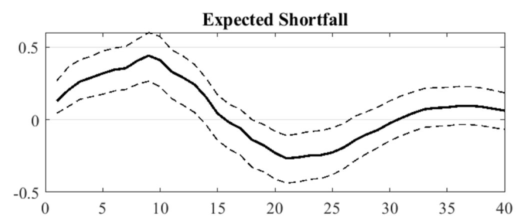

In panel (a), the VAR includes the expected median of consumption growth as a control. Panel (b) controls for the expected mean of consumption growth. In panel (c), the additional control is the expected variance of consumption growth. For panel (d), the downside risk measure is the expected shortfall considering all values below the median. The sample period is 1984:2-2019:12. The variables included are the monetary policy indicator, industrial production, the CPI, and the excess bond premium (EBP). All data is at a monthly frequency, the lag length is 12 months. Dashed lines are 95% confidence bands from a wild bootstrap with 5,000 replications.

Appendix H Return prediction robustness exercises

This appendix shows that the results from the predictive stock market return regressions are robust to several other specifications. To conserve space, I focus on six month-ahead excess market returns unless stated otherwise. I first allow for state dependence of the VIX, variance risk premium and price-dividend ratio. The results are in the left column of Table 6. None of these controls except the variance risk premium has significant predictive power once we include the interaction term of downside risk and the recession indicator. Interestingly, the interaction term between the VRP and the recession indicator has a negative sign, precisely opposite to economic intuition: Investors should demand higher risk premia when risk aversion rises.

The results remain similar when we extend the sample until December 2020, which includes the Covid turmoil from March 2020 (middle column). Market returns show an even stronger sensitivity to downside risk during recession times. Conditional on being in a recession, a one standard deviation rise in downside risk predicts an increase in excess market returns of 7.53 percentage points over the next six months. The variance risk premium also has significant state-dependent effects, although these are quantitatively weaker and again appear with a negative sign. The right column considers a sample going back to February 1984, which is equivalent to the baseline VAR sample. Since the data for the VIX and the VRP does not extend that far into the past, the regression only contains the price-dividend ratio and downside risk as predictors. Again, the downside risk measure is strongly significant.

Table 7 considers alternative measures of consumption growth. The top panel considers the linear case without state dependence. The bottom panel reports the regression results under state-dependent effects. The left column adds realized year-over-year consumption growth as an additional control. Realized consumption growth is a poor predictor of six-month ahead stock market returns over the sample period. Therefore, the results for downside risk are not driven by its association with realized consumption growth.

The middle column adds expected median consumption as a control. Expected median consumption is defined as the 50% quantile of the estimated conditional consumption growth distribution. Since downside entropy is constructed as the excess probability mass below the median of the conditional relative to the unconditional distribution, shifts in the median could be correlated with the downside risk measure. In this case, the effects for downside entropy might turn out to simply capture changes in the median of the consumption growth distribution, which would not support the downside risk hypothesis. The results indicate that the predictive ability of downside risk is robust to controlling for the expected median.

The right column reports results when controlling for the expected variance. Since the downside risk hypothesis stresses asymmetry and the importance of downside risk versus symmetric uncertainty, we should expect the results to go through when including the expected variance. The expected variance is estimated based on the fitted skewed t-distributions in each month. The regression results show that the coefficient on downside risk is still significant during recessions, which is where its predictive ability is concentrated. While the coefficient on the interaction term between the recession indicator and the expected variance is also statistically significant, its sign points in the opposite direction: Conditional on a given level of downside risk, increases in the expected variance are associated with lower risk premia, not higher risk premia.

Table 8 studies if upside entropy has similar predictive ability as downside entropy. The results indicate that both upside and downside risk can predict future excess returns. However, in terms of the downside risk channel of monetary policy, it is important to keep in mind that monetary policy has a strong effect on downside risk during recessions, but not on upside risk. The transmission of monetary policy shocks to equity premia is therefore occuring via the reduction in downside risk, even though upside risk may serve as a predictor for returns as well.

Table 6: Market return predictability regressions - Different Samples

	Baseline Sample		Sample: 1990-2020		Sample: 1984-2019	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Risk	0.70	0.56	0.79	0.58	0.12	0.62
log(PD)	-2.65	1.75	-2.73	1.75	-1.77	1.16
VIX	1.06	1.19	1.59	1.22	–	–
VRP	2.37	0.79	3.26	1.22	–	–
Recession	-11.59	67.32	-14.01	67.60	23.98	45.76
Recession $*$ Risk	6.08	2.26	7.53	2.23	8.00	2.11
Recession $*$ log(PD)	-0.44	4.44	0.35	4.30	-2.67	4.16
Recession $*$ VIX	2.24	1.76	0.77	2.05	–	–
Recession $*$ VRP	-0.54	1.36	-3.76	1.33	–	–
Constant	38.63	24.67	39.37	25.14	24.84	12.87
R^2	32.4		29.9		24.0	

The dependent variable is the excess return on the Fama-French market portfolio over the next 6 months. Risk is the downside entropy of the consumption growth distribution, in year-over-year growth rates. Recessions are as defined by the NBER. The log price-dividend ratio is taken from Robert Shiller's website. The VIX is the monthly level of the VIX index, constructed as the within-month average of daily adjusted closing prices. The variance risk premium is from Hao Zhou's website. The sample period for the baseline case is 01/1990 until 12/2019. All data is at a monthly frequency. Standard errors are Newey-West. R^2 is adjusted for the number of predictors.

Table 7: Market return predictability regressions - Different consumption measures

	Realized Consumption		Median Consumption		Consumption Variance	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Linear						
Risk	2.98	1.75	2.29	1.34	2.37	1.52
Consumption	-0.76	1.53	-2.70	1.02	1.88	0.71
Constant	5.48	3.48	9.80	1.78	1.55	1.81
State Dependent						
Risk	0.48	0.50	0.25	0.50	0.30	0.73
Consumption	-1.58	1.24	-2.21	1.19	1.28	0.73
Recession	-9.27	5.14	-22.18	4.64	4.75	5.92
Recession ∗ Risk	9.86	1.87	13.65	2.96	14.05	4.39
Recession ∗ Cons.	3.50	2.36	7.97	2.77	-6.49	3.35
Constant	7.77	2.35	9.31	2.13	2.95	1.60

The dependent variable is the excess return on the Fama-French market portfolio over the next 6 months. Risk is the downside entropy of the consumption growth distribution, in year-over-year growth rates. Recessions are as defined by the NBER. Realized consumption is the year-over-year growth rate of U.S. aggregate real personal consumption expenditures excluding durables. Median consumption is the predicted 50% quantile of the consumption growth distribution. The last column uses as a risk measure the variance estimated from the fitted distribution of consumption growth. The sample period is 01/1990 until 12/2019. All data is at a monthly frequency. Standard errors are Newey-West.

Table 8: Market return predictability regressions - Including upside risk

	$h = 3$		$h = 6$	
	Coeff.	Std. Err.	Coeff.	Std. Err.
Linear				
Downside	1.44	0.69	2.69	1.52
Upside	0.67	0.65	1.40	1.17
Constant	2.05	0.66	4.05	1.38
State Dependent				
Downside	0.29	0.46	0.44	0.86
Upside	0.15	0.53	0.37	0.90
Recession	-4.53	2.12	-7.47	2.46
Recession $*$ Down	2.96	0.58	5.97	1.49
Recession $*$ Up	4.35	1.15	8.08	1.79
Constant	2.42	0.55	4.65	1.05

The dependent variable is the excess return on the Fama-French market portfolio over the next 3 or 6 months. *Downside* is the downside entropy of the consumption growth distribution, and *Upside* is the upside entropy. Both variables are in year-over-year changes. Recessions are as defined by the NBER. The sample period is 01/1990 until 12/2019. All data is at a monthly frequency. Standard errors are Newey-West.