ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS

WORKING PAPER SERIES

Growth-at-Risk is Investment-at-Risk

Authors	Aaron Amburgey, and Michael W. McCracken
Working Paper Number	2023-020C
Revision Date	August 2025
Citable Link	https://doi.org/10.20955/wp.2023.020
Suggested Citation	Amburgey, A., McCracken, M.W., 2025; Growth-at-Risk is Investment-at-Risk, Federal Reserve Bank of St. Louis Working Paper 2023-020. URL https://doi.org/10.20955/wp.2023.020

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

Growth-at-Risk is Investment-at-Risk*

Aaron J. Amburgey† and Michael W. McCracken‡

August 14, 2025

Abstract

We investigate the role financial conditions play in the composition of U.S. growth-at-risk. We document that, by a wide margin, growth-at-risk is investment-at-risk. That is, if financial conditions indicate U.S. real GDP growth will be in the lower tail of its conditional distribution, we know that quantitatively, the main contributor is a decline in investment. Consumption contributes under extreme financial stress. Government spending and net exports do not play a role. We show that leverage plays a key role in determining both consumption- and investment-at-risk, which provides support to the financial accelerator mechanism proposed by Bernanke et al. (1999).

JEL Nos.: C32, C52, C53, E27, E44

Keywords: growth-at-risk, real-time data, quantiles, expected shortfall

*We are grateful to Scott Brave, Todd Clark, Ana Beatriz-Galvão, and Paulina Restrepo-Echavarria for helpful comments as well as participants at the MacroFor workshop, Dolomite Macro workshop, Midwest Econometrics Group, and seminar participants at the Federal Reserve Bank of St. Louis and the Booth School of Business. We thank Ngân Trần for excellent research assistance. The views expressed here are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

†Northwestern University. Email: amburgey@u.northwestern.edu

‡Federal Reserve Bank of St. Louis. Address: P.O. Box 442. St. Louis, MO 63166, U.S.A. Tel: 314-444-8594. Email: michael.w.mccracken@stls.frb.org

1 Introduction

There is now a substantial empirical literature on downside risks to real GDP growth – often described as growth-at-risk. Much of the work focuses on the role of financial conditions. For example, Adrian et al. (2019; ABG) use quantile regressions to examine the link between the Chicago Fed's National Financial Conditions Index (NFCI) and downside risks to U.S. real GDP growth. Adrian et al. (2022) provide comparable evidence for other advanced economies. Building on this work, alternative sources of tail risk have also been considered, including climate change (Kiley, 2024), natural disasters (Bayoumi et al., 2021), and uncertainty (Hengge, 2019). In addition, alternative methods have been developed for modeling tail risk to GDP, including Bayesian additive regression trees (Clark et al., 2022), GARCH-based location-scale models (Brownlees and Souza, 2021), and generalized autoregressive score models (Delle Monache et al., 2023).

One of the main empirical results in ABG is that tighter current financial conditions (z) is associated with greater downside risk to U.S. real GDP growth (y). One way they do so is by analyzing its expected shortfall. Succinctly, for some level of tail risk α (e.g., 5%), forecast origin t, forecast horizon τ, and conditional quantile for real GDP growth $Q_\alpha(y_{t+\tau}|y_t, z_t)$, expected shortfall is $E(y_{t+\tau}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|y_t, z_t))$. This measure indicates what GDP growth is expected to be conditional on being in the lower tail of potential future outcomes. An example of this is provided in Figure 9 of ABG, which plots the estimated expected shortfall and longrise (upside risk) at the 1- and 4-quarter horizons.[1] Briefly, they show that significant declines in the expected shortfall are almost always associated with large increases in the NFCI and thus drastically tightening financial conditions. In contrast, NFCI values below or near zero do not have any clear effect on growth-at-risk at either horizon.

In this paper, we revisit the original work of ABG, but from a different perspective. We investigate the role financial conditions have on the composition of U.S. growth-at-risk. That is, if

[1]Note that they provide in-sample estimates of expected shortfall. In contrast, we estimate the out-of-sample expected shortfall to emphasize the real-time importance of financial conditions for real GDP growth.

current financial conditions indicate that U.S. real GDP growth is going to be in the lower tail of its conditional distribution next quarter, what does this imply for growth of consumption, investment, government spending, or net exports next quarter? In doing so, we provide a collection of stylized facts that gives a more nuanced view of the economy than when we say (aggregate) growth is at risk. In spirit, our exercise is analogous to Adrian and Brunnermeier (2016), who decompose systemic risk by assessing how the characteristics of individual banks (e.g., leverage, size) contribute to aggregate financial vulnerability.[2] Though we do not explicitly embed our findings within a structural model, we view them as potentially informative for calibrating or validating macroeconomic frameworks that feature financial frictions (e.g., Brunnermeier and Sannikov, 2014) or rare disaster risk (e.g., Gourio, 2012).

Our approach to answering this question is related to Acharya et al.'s (2017) marginal expected shortfall but is closer to Banulescu and Dumitrescu's (2015) component expected shortfall. Let $w \in \{c, i, g, nx\}$ denote the growth rate of consumption, investment, government spending, or net exports. Using copula-based methods to link the conditional distributions of $y_{t+\tau}$ and $w_{t+\tau}$, we form estimates of $E(w_{t+\tau}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|y_t, z_t))$, which we refer to as an approximate marginal expected shortfall (aMES). Each of these is then scaled by the contribution of the level of the component to the level of GDP to form what we refer to as an approximate component expected shortfall (aCES). Methodological details are provided in Section 3.

Figure 1 provides our main result. At the 1- and 4-quarter horizons we plot real-time estimates of each aCES. By a wide margin, investment accounts for the majority of variation in growth-at-risk, as its aCES is almost always the only notable negative contributor. Consumption-at-risk only contributes to downside risks in cases of the most extreme financial distress, as was the case at the height of the Great Recession. Unsurprisingly for U.S. data, contributions from government spending and net exports are negligible.

Inspired by Chari et al.'s (2007) treatment of durables consumption as a form of investment

[2] The key distinction is that while they examine what individual components reveal about aggregate risk, we reverse the logic: we ask what aggregate financial conditions imply for the distribution of risks across GDP components.

in their "Business Cycle Accounting," we also conduct a comparable exercise using the subcomponents of consumption and investment - the details of which we provide later. As expected, durables account for the bulk of consumption's limited negative contribution to growth-at-risk. Non-durables contribute a little, while services have no impact. In contrast, both residential and non-residential investment always contribute negatively.

Finally, with an eye towards understanding the mechanism in which financial conditions affect growth-at-risk – and in particular investment-at-risk – we implement LASSO penalized quantile regressions of GDP and its components on the leverage, risk, and credit subindices of the NFCI. We find that leverage is always the most important source of financial stress. Our novel empirical connection between leverage and investment-at-risk provides fresh support for the financial accelerator framework laid out in Bernanke et al. (1999).

The remainder of the paper proceeds as follows. Section 2 describes the data and evaluates the usefulness of the NFCI for forecasting tails risks associated with the growth of each component of GDP. Section 3 delineates our methodology for forming the composition of GDP growth-at-risk and presents our main empirical results as well as the detailed results on investment, consumption, and net exports. Section 4 shows that the leverage subindex of the NFCI is strongly associated with consumption- and investment-at-risk and connects our findings to the financial accelerator framework. Section 5 concludes.

2 Financial Conditions and the Components of Growth-at-Risk

Our decomposition of growth-at-risk builds on quantile forecasts. As in much of the literature, these forecasts are conditional on one lag of GDP growth and one lag of the quarterly average of the NFCI. The usefulness of financial conditions indices for predicting downside risk to growth in GDP is well established (e.g., Adrian, et al., 2022) despite being the subject of some criticism (e.g., Reichlin et al., 2020). Even so, their usefulness for forecasting the individual components of GDP growth has not been established, and hence we begin our analysis by doing so in the following.

2.1 Quantile Regression Framework

Our evidence is based on fully real-time out-of-sample exercises. These entail sequences of τ-step-ahead quantile forecasts formed sequentially across quarterly forecast origins $t = R, ..., T - \tau$, where R is the minimum estimation window size. In each instance they are formed using current vintage data $\{w_s(t), x_s'(t)\}_{s=1}^{t}$. These data consist of a scalar predictand $w_s(t)$ and vector of predictors $x_s(t)$ associated with observations $s = 1, ..., t$. As an example, when $\tau = 1$, ABG set $w_s(t)$ to current quarter GDP growth (i.e., $y_s(t)$) and set $x_s(t)$ to $(1, y_{s-1}(t), z_{s-1}(t))'$ where z denotes the NFCI.[3] Later we will see other cases in which the quantile regressions treat $w_s(t)$ as current quarter consumption growth (or some other component) while maintaining the same definition of $x_s(t)$ so that lagged GDP growth remains a predictor. We also consider cases in which $\tau = 4$, and in these cases $w_s(t)$ is defined as average quarterly growth over four consecutive quarters.

For a chosen $\alpha \in (0, 1)$, the τ-step-ahead conditional $\alpha-$quantile forecasts are based on linear models $x_s'(t)\beta^{(\alpha)}$. The parameter vector $\beta^{(\alpha)}$ is re-estimated as we move across forecast origins by minimizing the relevant tick loss function and takes the form

$$\hat{\beta}_t^{(\alpha)} = \arg\min_{\beta} \sum\nolimits_{s=1}^{t-\tau} (\alpha - 1(w_{s+\tau}(t) \leq x_s'(t)\beta))(w_{s+\tau}(t) - x_s'(t)\beta).$$

The quantile forecasts $x_t'(t)\hat{\beta}_t^{(\alpha)}$ are evaluated against the future realization $w_{t+\tau}(t')$ for some vintage t' such that $t' \geq t + \tau$. This yields $P = T - \tau - R + 1$ forecast errors, denoted $\hat{u}_{t+\tau}^{(\alpha)}(t') = w_{t+\tau}(t') - x_t'(t)\hat{\beta}_t^{(\alpha)}$ and subsequent loss $L(\hat{u}_{t+\tau}^{(\alpha)}(t')) = (\alpha - 1(\hat{u}_{t+\tau}^{(\alpha)}(t') \leq 0))\hat{u}_{t+\tau}^{(\alpha)}(t')$.

[3]Throughout y and w refer to percentage changes, while z is the level of the NFCI.

2.2 Data

We obtain GNP and GDP (including components) data from the Federal Reserve Bank of Philadelphia's Real-Time Data Set for Macroeconomists (RTDSM).[4] We obtain NFCI data from the ALFRED database. For NFCI vintages before May 25, 2011, we use the unofficial NFCI vintages constructed by Amburgey and McCracken (2023). The data date back to 1973Q1, and we obtain vintages from January 1988 through October 2023. Throughout, we generalize by using the term GDP despite using GNP in vintages before December 1991.

Forecast origins take place at the start of the second month of a given quarter, and we evaluate forecasts for the horizons $\tau = 1, 4$. All forecasts are constructed using a recursive estimation scheme with an initial estimation sample of fifteen years (R = 60). As a result, forecast origins do not necessarily start at the first available vintage, due to an insufficient number of observations for the direct multi-step estimation. Put together, when $\tau = 1(4)$ the first forecast origins are May 1, 1988 (February 1, 1989), with target dates 1988Q2 (1989Q4). This procedure continues up to 2019Q4. In the main text, we abstract from evaluating any forecasts with horizons after 2019Q4 to avoid the atypical macroeconomic data releases during the COVID-19 pandemic.[5] For the NFCI, we use non-calendar 3-month averages—that is, averages that use the most recent 3 months of observations at a given origin. To evaluate tick loss, we use the initial release of each series.

2.3 Evaluating Forecast Accuracy

As noted in section 2.1, we are interested in evaluating the usefulness of the NFCI for forecasting quantiles of each GDP component. To do so, we compare a NFCI-containing model to a baseline model with the NFCI omitted. Formally, the unrestricted model takes the form $\hat{w}_{t,\tau}^{(\alpha)} = x_t(t)'\hat{\beta}_t^{(\alpha)}$ for $x_t(t) = (1, y_t(t), z_t(t))'$, where $\hat{w}_{t,\tau}^{(\alpha)}$ is the τ-quarter-ahead component forecast. The restricted

[4]The only exception is aggregate Private Domestic Investment for which we obtain data from the ALFRED database.

[5]See Online Appendix C for an analysis of the COVID-19 period.

model omits $z_t(t)$ as a predictor.[6] To evaluate whether the NFCI is a useful predictor, we use the OOS-t test for nested quantile forecasting models developed in Amburgey and McCracken (2023), which takes the form

$$OOS - t = \frac{P^{-1/2} \sum_{t=R}^{T-1} L(\hat{u}_{1,t+\tau}^{(\alpha)}) - L(\hat{u}_{2,t+\tau}^{(\alpha)})}{\sqrt{\hat{\Omega}}}, \quad (1)$$

where $L(\hat{u}_{1,t+\tau}^{(\alpha)})$ and $L(\hat{u}_{2,t+\tau}^{(\alpha)})$ represent the tick loss of the restricted and unrestricted models for the forecast target date, $t + \tau$, and the conditional quantile, α.[7]

Table 1 displays the results from this exercise. The first column gives the predicted series, and the second row gives the α-quantile, with $\alpha \in (0.05, 0.1)$ representing the lower tail and $\alpha \in (0.9, 0.95)$ representing the upper tail. For each exercise we display the percentage change in average tick loss associated with including lagged NFCI as a predictor (e.g., a value of -0.460 means the model including the NFCI had 46% lower tick loss on average compared to the baseline model). In parenthesis, we report the p-values associated with a two-sided test of the null using the OOS-t test delineated above. For ease of viewing, we split results for the 1- and 4-quarter-ahead forecasts into two panels.

The results for GDP largely mirror the existing literature. In the lower tail, the model that includes lagged NFCI performs significantly better than the baseline model. Conversely, in the upper tail the baseline model performs better, albeit not at a significant level. The consumption results mimic those of GDP. Perhaps more noteworthy are the investment results. Here we observe a good example of how financial conditions have distinct effects on the components. In contrast to GDP and consumption, the NFCI model performs better at forecasting investment growth across all quantiles (upper and lower) at the shortest horizon and significantly so for all but the largest

[6]In unreported results we also include the predictors of Christiano et al. (2005) (less profits due to its delayed release), and find limited gains over our parsimonious specification.

[7]The test is an adaptation of West (1996) that permits non-differentiability of the tick loss function as well as the presence of data revisions. See Amburgey and McCracken (2023) for details on estimation of the asymptotic variance, $\hat{\Omega}$.

quantile. In addition, at the longer horizon, the NFCI model does better in the lower tail as well as the median, while doing worse in the upper tail. The government spending results are a mixed bag, with the NFCI model performing worse (better) than the baseline model for $\tau = 1(4)$in the lower tail. Finally, in all cases except for one, the net export forecasts are improved by the NFCI, usually significantly so. Amalgamating the results in Table 1, it is generally clear that the inclusion of NFCI is helpful for forecasting growth-at-risk, as well as the components of growth-at-risk – absent government spending. It's also clear that the degree of predictive content for the components differs from that for GDP – especially for investment, government spending, and net exports.

3 Decomposing Growth-at-Risk

In the previous section we established that the NFCI is a useful predictor of downside risk to growth in consumption, investment, and net exports. In this section we build on these observations and decompose growth-at-risk using an approximate component-wise expected shortfall. This requires a method for linking the downside risks between real GDP growth and each of its components, and here we opt for a parametric copula to do so. Other applications of copula-based methods in the growth-at-risk literature include Coe and Vahey (2020), who study the relationship between financial conditions and the distribution of future U.S. real GDP growth across a long historical sample dating back to 1875, and Chernis et al. (2023), who revisit the results in ABG using a copula-based quantile regression framework designed to limit the influence of outliers like those associated with the COVID-19 period.

3.1 Our Methodology

We begin by using real-time conditional quantile forecasts to construct marginal conditional distributions of GDP growth as well as growth of the components. At each forecast origin and for each α= 0.05, 0.1, ..., 0.95, we use one lag of both the NFCI and GDP growth as predictors when forming these quantile forecasts. Within a comparable framework, ABG use these forecasts to inform the parameters of an assumed skewed-t marginal conditional distribution. Instead, we use a semi-parametric method developed in Mitchell et al. (2024, MPZ), which builds on Koenker and Zhao (1996). By doing so we impose fewer restrictions on the shapes of the marginal conditional distributions.[8]

In all cases we follow Algorithm 1 of MPZ. At each forecast origin, their approach treats the forecasted quantiles as the actual quantiles of the marginal distribution. Values between these quantiles are simulated using 1000 uniform draws. Beyond the lowest (highest) quantile, the 0.05 and 0.10 (0.90 and 0.95) quantile forecasts are used to estimate the parameters of a normal distribution, and then 1000 values are drawn from the implied tail of this normal distribution. This collection of quantiles and simulated values are then used to form a kernel-based estimate of the marginal distribution.[9] From here we can easily estimate $E(y_{t+\tau}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|y_t, z_t))$, i.e., the expected shortfall seen in ABG and MPZ.[10]

Given the sequence of marginal distributions, we now construct a sequence of bivariate joint conditional distributions between GDP growth and growth in each of the individual components. To do so, we link the corresponding marginal distributions using a Student's t copula, rather than a Gaussian copula, as it permits greater flexibility of the bivariate tail risks. Specifically, we consider the object:

$$C_{\Theta_w}(u_y, u_w) = \mathbf{t}_{\nu_w,\Sigma_w}(t_{\nu_y}^{-1}(u_y), t_{\nu_w}^{-1}(u_w)),$$

where $\mathbf{t}_{\nu_w,\Sigma_w}$ is the CDF of the bivariate Student's t distribution with correlation matrix Σ_w and degrees of freedom ν_w, t_ν^{-1} denotes the inverse CDF of the univariate Student's t-distribution with degrees of freedom ν, and u_y and u_w are the probability integral transforms of the marginal dis-

[8]In unreported results we follow the methodology of ABG and obtain similar findings.

[9]We use a normal kernel with the theoretically optimal bandwidth from Bowman and Azzalini (1997).

[10]In their figures, rather than the out-of-sample forecasts, they plot the in-sample estimates obtained at the last horizon, i.e., 2019Q4. The in-sample estimates can be obtained by simply following the same set of steps outlined here, but using the fitted quantiles rather than the out-of-sample quantile forecasts.

tributions of GDP growth and component w, respectively. For each $w \in \{c, i, g, nx\}$, we estimate $\Theta_w = \{(\nu_w, \Sigma_w) : \nu \in (1, \infty), \Sigma_w \in \mathbb{R}^{2\times 2}\}$ separately using maximum likelihood over the full sample.[11] Appendix Table A1 reports the estimated hyperparameters for each component and horizon (i.e., for 1- and 4-quarter percentage changes).

The need for joint distributions follows immediately from the aMES, $E(w_{t+\tau}|y_{t+\tau} < Q_{.05}(y_{t+\tau}|x_t))$. To estimate this object we need the ability to generate draws from the joint distributions of future GDP growth and each of its components, conditional on current NFCI and GDP growth, i.e., $F(y_{t+\tau}, w_{t+\tau}|x_t)$ $t = R, ..., T-\tau$. Specifically we need draws from the slice of these distributions for which $y_{t+\tau} < Q_{.05}(y_{t+\tau}|x_t)$. But having constructed the joint distributions, drawing from them is straightforward and the average of these draws forms our estimate.

Figure 2 plots the 1- and 4-quarter-ahead aMES forecasts for each component. By a large margin, growth in investment exhibits the greatest and most consistent risk. Across both horizons, its aMES is nearly always negative, and when financial conditions are tight, it experiences substantial declines. For the 1-quarter-ahead forecasts, leading into the 1990-91, 2001, and 2008 recessions, investment growth is predicted to decline roughly -35%, -25%, and -70% (!), respectively. At the 4-quarter horizon the predicted declines are lower but are still substantial at -15%, -13%, and -30%, respectively. In contrast, the forecasts for consumption are nearly always positive. The exception occurs when financial conditions are at their tightest leading into the Great Recession period. Here we observe a substantial decline of roughly -7% at both horizons.

The third row of panels indicates that, unsurprisingly, government spending is largely unaffected by financial conditions. At both horizons there appears to be a clear, but insignificant, downward trend in the expected shortfall since the 1990s. Finally, the last row of panels shows that the downside risk to net exports tends to move in the positive direction during economic downturns — or rather the risk to imports is greater than the risk to exports. Outside of recessions the expected shortfall of net exports remains near zero, however.

[11] In unreported results we re-estimate the hyperparameters in real-time using the available data at each forecast origin. The main results remain largely unchanged.

3.2 The Composition of GDP Growth-at-Risk

Now that we have the individual aMES components, we can finalize our decomposition of growth-at-risk, adjusting for the fact that these components contribute to GDP by very different magnitudes. On average over our sample, consumption is roughly 70% of GDP while the remaining pieces are roughly 15% (noting that net exports is negative). To account for these features we rescale the results in Figure 2 relative to each component's total proportion of GDP. We denote a component's rescaled aMES as its aCES. Formally, a component's aCES is given by

$$\frac{W_t}{Y_t} E(w_{t+\tau} | y_{t+\tau} < Q_\alpha(y_{t+\tau}|x_t)), \tag{2}$$

where Y_t and W_t represent the levels of GDP and one of its components, respectively. With this adjustment, we return to our main result in Figure 1. First consider that the sum of all four components' estimated aCES (i.e., the total area of the plot) yields an estimate of the expected shortfall. The patterns displayed by this approximation largely mimic the in-sample estimates from ABG and MPZ. Importantly, periods associated with tight financial conditions tend to coincide with markedly negative expected shortfalls.

Quantitatively, investment is the sole consistent contributor to the downside risk to growth, as it is the only element that is consistently negative. Specifically, at the 1- and 4-quarter horizons, investment makes up an astounding 91% and 87% of total negative contributions in our sample, respectively. On the other hand, consumption contributes only in the presence of extreme financial distress, as we saw during the Great Recession. Across our sample, consumption makes up a mere 7% and 13% of total negative contributions at the 1- and 4-quarter horizons, respectively, with the majority of these contributions coming from the Great Recession period at both horizons. Downside risks from government spending and net exports are negligible.

The relative magnitude of investment's contributions to downside risk are large even after considering its high volatility relative to the other components. First consider that investment makes up only approximately 15% of GDP (compared with consumption's 70%). Second, note that the standard deviations of 1-quarter ahead consumption and investment growth in our sample are 2.48 and 15.31, respectively. Thus, investment is roughly 6 times more volatile than consumption while contributing roughly 13 times more to the total negative contributions to downside risk. Due to the established link between financial frictions and an investment wedge, one possibility is that the importance of investment is driven by our inclusion of NFCI as a predictor. To verify this is not the case, Figure A1 plots the aCES when the NFCI is excluded. Investment remains the dominant contributor.

When interpreting these results, it is important to distinguish between statistical contributors and fundamental economic drivers of growth-at-risk. Figure 1 indicates that, when GDP growth is in its lower tail, the majority of the change to GDP comes from a change to investment. However, because each of the GDP components are jointly determined, this result does not necessitate a causal relationship between investment and growth-at-risk. In fact, none of the GDP components are included as predictors in our baseline model. Despite this nuance, our findings indicate that investment makes up the quantitative majority of growth-at-risk.

Finally, while we find our decomposition instructive, as noted earlier, it is still an approximation. Loosely speaking, it is an approximation because while the component shares are relatively stable, they are not constant, and yet we treat them as approximately so. A more detailed explanation is provided in Online Appendix B. Nevertheless, on average across all forecast origins in Figure 1, the average deviation between the sum of the aCES values and the actual expected shortfall for GDP growth at the 1-quarter (4-quarter) horizon is roughly 11 bps (9 bps), albeit with a standard deviation of 1.1% (1.4%).[12]

3.3 A Detailed Composition of GDP Growth-at-Risk

In our main analysis we made a choice about the granularity of GDP's components. Here we revisit this choice with an eye toward the robustness of our findings and any insights gained from a more thorough breakdown of GDP. For example, while growth-at-risk from consumption and investment

[12]Figure A2 plots the comparison.

make up the bulk of our findings, it may be that these categories remain too broad to understand their underlying contributions to growth-at-risk. As an example, Mankiw (1985) emphasizes that consumption of durables is more sensitive to interest rates than consumption of nondurables or services. In addition, Leamer (2007) argues that residential investment plays a dominant role in GDP weakness with non-residential investment playing a lesser role.

In that spirit, our more detailed analysis of the composition of growth-at-risk applies the same modeling choices outlined in Sections 3.1 and 3.2, but to a different set of series, w. We analyze how consumption of nondurable goods (CNDUR), consumption of durable goods (CDUR), consumption of services (CSER), nonresidential investment (INRES), residential investment (IRES), exports (X), and imports (M) contribute to growth-at-risk.[13] [14]

Figure 3 plots the detailed composition of growth-at-risk at the 1-quarter- and 4-quarter ahead horizons. The top panel focuses on consumption, which typically did not contribute negatively to growth-at-risk in Figure 1. Even so, our detailed decomposition reveals that consumer durables is a regular downside risk, as it displays significant negative contributions that tend to arise leading into, and during, recessions. This aligns with our expectation given Chari et al.'s (2007) treatment of durables consumption as a form of investment. Nondurables provide substantial contributions to growth-at-risk only during the height of the financial crisis, reinforcing Houthakker and Taylor's (1970) assertion that certain non-durables can behave as durables in the right economic environment. Across the sample, the contribution from services consumption remains positive and is generally invariant to financial conditions, with the exception that during the height of the Great Recession it dips sharply to zero.

The middle panel gives us a more detailed understanding of investment-at-risk. In contrast to the components of consumption, both nonresidential and residential investment contribute to growth-at-risk as they remain below zero for most of the sample. The balanced contributions to

[13]For imports we multiply equation (2) by -1 since imports contribute negatively to the level of GDP.

[14]We choose to omit change in inventories since it makes up less than 1% of investment. We also omit government spending due to its insignificance in our initial decomposition.

growth-at-risk contrast with Leamer's (2007) assessment above. It's important to note however, that Leamer (2007) is emphasizing the timing of contributions to GDP relative to business cycle peaks, while we're considering risks to GDP at any stage of the business cycle.

Finally, the bottom panel indicates that exports and imports contribute to growth-at-risk more than Figure 1 reveals. During expansions both exports and imports contribute little to growth-at-risk, but during recessions they deliver significant negative and positive contributions, respectively. This finding supports evidence from the international trade literature showing that trade significantly declines during global recessions, e.g., Eaton et al. (2016). Overall, throughout the entire sample there is a clear trend of changes to exports and imports canceling each other out. In other words, trade tends to decrease in both directions during downturns rather than only exports or only imports. Since exports (imports) contribute positively (negatively) to GDP, significant risk to trade is masked by minor risk to net exports.

4 The Determinants of Growth-at-Risk

In Section 3 we established that the vast majority of growth-at-risk is investment-at-risk. We now turn to the underlying mechanisms that explain this overwhelming contribution. To do so, rather than use the NFCI to predict tail risk, we use its three subindices: leverage, credit, and risk. Based on in-sample LASSO quantile regressions, we find that in all cases, leverage is the strongest predictor of downside risk. Finally, we connect the main results from Sections 3 and 4 to the financial accelerator framework.

4.1 Leverage and Growth-at-Risk

For each series in the NFCI, Brave and Butters (2012) assign a label of "leverage," "risk," or "credit," which reflects the theme of each series. Each subindex is estimated in a similar fashion to the NFCI, albeit with a subset of the series. Specifically, the authors simply take the estimated loadings from the NFCI, set the loadings of variables not in the subindex to zero, and reapply the

Kalman filter and smoother. The leverage subindex includes financial sector stock price indices, open interest, debt issuance, and asset-to-GDP ratios; the risk subindex includes repo market volumes, volatility indices, and a large set of yield spreads; and the credit subindex includes household survey questions about credit conditions, senior loan officer opinion survey questions about standards for loans, delinquent credit to loan ratios, and mortgage treasury spreads. Appendix Figure A3 plots the NFCI and each of its subindices. See Brave and Butters (2012) for a more detailed discussion of the series contained in each NFCI subindex as well as their estimation. For the remainder of the paper we utilize the most recent vintage in our dataset (i.e., October 2023) for the NFCI subindices as well as the GDP components. Similarly to the previous exercises, we abstract from the COVID-19 period and use the sample from 1971Q1 to 2019Q4.

As noted above, we use LASSO quantile regressions as a device to determine which of the NFCI subindices are most important for predicting growth-at-risk and its components. To do so we consider an alteration of the unrestricted model discussed in Section 2.3 that takes the form $\hat{w}_{t+\tau}^{(\alpha)} = x_t'\hat{\beta}^{(\alpha)}$ for $x_t = (1, y_t, l_t, r_t, c_t)'$, where the last three terms denote quarterly averages of the leverage, risk, and credit subindices of the NFCI. As in Section 3, we focus on downside risk by setting $\alpha = 0.05$. Furthermore, we now narrow our focus to consumption-at-risk, and investment-at-risk, as we have shown government spending and net exports to be inconsequential to total downside risk. For ease of interpretation, all series are demeaned and standardized to have a variance of 1.

To estimate $\hat{\beta}^{(\alpha)}$ we apply the LASSO quantile regression of Li and Zhu (2008) using, specifically, the algorithm derived in Belloni and Chernozhukov (2011). For a given quantile, α, and noting that for this exercise the intercept term $\beta_0^{(\alpha)}$ is always included, the regression is the solution to the following optimization problem:

$$\min_{\beta^{(\alpha)}} \sum\nolimits_{t=1}^{T-\tau} \left[L(\hat{u}_{t+\tau}^{(\alpha)}) \right] + \lambda \sum_{j=1}^{4} \left| \beta_j^{(\alpha)} \right|, \tag{3}$$

where $L(\hat{u}_{t+\tau}^{(\alpha)})$ is the tick loss of the τ-quarter ahead forecast of the α-quantile (i.e., as in section

2.1) and $\lambda > 0$ is a tuning parameter. Analogous to LASSO regressions in the OLS context, the first term of equation (3) characterizes the minimization problem associated with ordinary quantile regression while the second term adds a ℓ_1-norm penalty whose importance is controlled via the hyperparameter λ.

While a choice of λ is a required to conduct the optimization, the "optimal" choice of λ is not relevant for this exercise.[15] We are using the penalty term strictly as a variable selection device that highlights which of the predictors are most and least important for predicting downside risk. For example, if λ is large enough one would expect all of the slope coefficients to be set to zero because the marginal predictive content of each predictor is insufficient to warrant a non-zero value. As we lower the value of λ, the predictor with the strongest marginal predictive content eventually warrants a non-zero coefficient. Continuing the process of lowering λ eventually leads to the case in which even the least relevant predictor will warrant a non-zero coefficient. The implied ranking of the strength of the predictors is what we are interested in. Note that for each distinct λ we can record the average tick loss, the coefficients $\hat{\beta}^{(\alpha)}$, which predictors have estimated coefficients that are non-zero, and the value of the ℓ_1-norm penalty.

Figure 4 plots the average tick loss associated with penalized quantile regressions of 1-quarter- and 4-quarter-ahead GDP, consumption and investment on x_t. The x-axis displays the ℓ_1-norm associated with the value of λ used for each estimation. Larger values of the ℓ_1-norm always coincide with smaller values of λ. In summary, moving rightward on the plot implies a lower penalty, and thus a larger ℓ_1-norm and lower associated tick loss. Vertical lines for each series in x_t are included to indicate the cutoff at which that series' coefficients become non-zero, i.e., lines further to the right indicate a smaller penalty required for a series to be selected by the model. Strikingly, GDP and leverage are the most important predictors in all cases. For each horizon and component, GDP is included regardless of the penalty-level, while leverage is included for all but very large λ. Moreover, leverage is always the first of the subindices to be included, with a sizeable

[15]There are several proposed criterion in the literature for choosing λ each of which considers a distinct data or modeling environment. See Belloni and Chernozhukov (2011), Li and Zhu (2012), and Lee et al. (2014) for examples.

gap existing between the ℓ_1-norm cutoff of leverage and the next most important predictor. Credit is included for moderate values of the penalty, and risk is included only in extreme cases, i.e., when λ is very small. Appendix Figure A4 plots the coefficients associated with each regression and confirms both that the tightening of each subindex is associated with more tail risk and that leverage has the most negative contribution in all cases.

4.2 Discussion

These results support existing findings in the empirical growth-at-risk literature that show leverage has an important relationship with vulnerabilities to financial markets and GDP growth (Adrian and Shin, 2010; Brave and Butters, 2012; and Reichlin et al. 2020). However, by considering the composition of growth-at-risk directly, we are able to pinpoint the role leverage has on growth-at-risk through both consumption- and investment-at-risk. This is an important distinction, because as we have shown, consumption only contributes to growth-at-risk under extreme financial conditions, while investment always plays a role. These new results then suggest that expanding leverage is associated with significant risk to GDP growth through vulnerabilities to future investment, *even when the economy is doing well.*

This refinement provides a direct link between the empirical growth-at-risk literature and the "financial accelerator" mechanism proposed by Bernanke et al. (1999), which notes that higher amounts of leverage amplify the effect of unexpected asset returns to entrepreneurs' equity. In this canonical model financial frictions do not directly affect consumption, rather, financially constrained firms' investment spending deviates much more from steady-state levels in response to macroeconomic shocks. As a result, all else equal, the growth of more leveraged economies are more vulnerable to shocks that affect the asset position of firms. In the language of the empirical growth-at-risk literature: deepening leverage creates more investment-at-risk. Christensen and Dib (2008) argue that this mechanism greatly amplifies the negative investment response to monetary and demand shocks, and *dampens* the upside response to technology shocks. In line with this latter set of results, Appendix Figures A5 and A6 show that in our framework leverage is the most important dampener of upside risk to investment.

However, our findings also raise questions about the broader transmission of financial conditions. Specifically, the NFCI leverage subindex captures not just firm leverage, but also indirect effects of tightening credit markets and balance sheet stress among financial intermediaries.[16] While standard financial accelerator models focus on borrower-side frictions, our results suggest that a complete account of downside risk should also consider supply-side constraints in the financial sector, as emphasized by Brunnermeier and Sannikov (2014) and Gertler and Kiyotaki (2010). These models show how intermediary net worth affects credit supply and macro volatility, but have largely been developed in isolation from the firm-side accelerator. Our results suggest that joint modeling of both firm leverage and intermediary constraints could better account for the emergence of systemic downside risk.

5 Conclusion

In this paper we address a previously unexplored aspect of growth-at-risk. Using the concept of component-wise expected shortfall, we provide evidence on its composition. By a wide margin we find that U.S growth-at-risk is quantitatively dominated by investment-at-risk, be it residential or nonresidential. Consumption, and especially consumption of durables, contributes to growth-at-risk but only in the most extreme cases of financial stress. Almost no contribution comes from net exports or government spending. Finally, we show that leverage plays a vital role in both consumption- and investment-at-risk as is argued by the financial accelerator literature.

It is worth noting that while we expect investment to be a primary contributor to GDP growth-at-risk, the relative importance of each component may vary significantly across countries. For example, in contrast to exports' moderate role in the U.S. economy, they make up approximately half of Germany's GDP. Thus, in Germany's case, one might expect exports-at-risk to make up a

[16] Aside from firm and household debt ratios, it includes measures of commercial bank loan exposures, unused credit lines, broker-dealer balance sheets, and securitization activity (e.g., ABS and MBS issuance). In unreported results we also include the non-financial leverage subindex, and find that leverage remains the dominant predictor.

more substantial proportion of growth-at-risk. While beyond the scope of this paper, investigating the composition of growth-at-risk in other countries seems a sensible avenue for future research.

References

Acharya, Viral V., Pederson, Lasse H., Phillippon, Thomas, and Matthew Richardson (2017), "Measuring Systemic Risk," *The Review of Financial Studies,* 30, 2-47.

Adrian, Tobias, Boyarchenko, Nina, and Domenico Giannone (2019), "Vulnerable Growth," *American Economic Review*, 109, 1263-1289.

Adrian, Tobias, and Markus K. Brunnermeier (2016), "CoVaR," *American Economic Review*, 106, 1705-1741.

Adrian, Tobias, Grinberg, Federico, Liang, Nellie, Malik, Sheheryar, and Jie Yu (2022), "The Term Structure of Growth-at-Risk," *American Economic Journal: Macroeconomics*, 14, 283-323.

Adrian, Tobias and Hyun Song Shin (2010). "Liquidity and Leverage," *Journal of Financial Intermediation*, 19(3), 418-437.

Amburgey, J. Aaron and Michael W. McCracken (2023), "On the Real-Time Predictive Content of Financial Conditions Indices for Growth," *Journal of Applied Econometrics*, 38, 137-163.

Banulescu, Georgiana-Denisa and Elena-Ivona Dumitrescu (2015), "Which are the SIFIs? A Component Expected Shortfall approach to systemic risk," *Journal of Banking and Finance*, 50, 575-588.

Bayoumi, Tamim, Quayyum, Saad Noor, and Sibabrata Das (2021), "Growth at Risk and Natural Disasters," IMF Working Paper 21/234.

Belloni, Alexandre, and Victor Chernozhukov (2011), "ℓ_1-Penalized Quantile Regression in High-Dimensional Sparse Models," *Annals of Statistics*, 39, 82-130.

Bernanke, Ben S., Gertler, Mark, and Simon Gilchrist (1999). "The Financial Accelerator in a

Quantitative Business Cycle Framework," *Handbook of Macroeconomics*, 1(1), chapter 21, 1341-1393, Elsevier.

Bowman, A. W., and A. Azzalini. Applied Smoothing Techniques for Data Analysis. New York: Oxford University Press Inc., 1997.

Brave, Scott, and R. Andrew Butters (2012), "Diagnosing the Financial System: Financial Conditions and Financial Stress," *International Journal of Central Banking* June, 191-239.

Brownlees, Christian, and André Souza (2021), "Backtesting Global Growth-at-Risk," *Journal of Monetary Economics*, 118, 312-30.

Brunnermeier, Markus K., and Yuliy Sannikov (2014), "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104, 379-421.

Chari, V.V, Kehoe, Patrick J., and Ellen R. McGrattan (2007), "Business Cycle Accounting," *Econometrica*, 75, 781-836.

Chernis, Tony, Coe, Patrick J., and Shaun P. Vahey (2023), "Reassesing the Dependence Between Economic Growth and Financial Conditions since 1973," *Journal of Applied Econometrics*, 38, 260-267.

Christensen, Ian and Ali Dib (2008), "The Financial Accelerator in an Estimated New Keynesian model," *Review of Economic Dynamics*, 11, 155-178.

Christiano, Lawrence J., Eichenbaum, Martin, and Charles L. Evans (2005), "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1-45.

Clark, Todd E., Huber, Florian, Koop, Gary, Marcellino, Massimiliano, and Michael Pfarrhofer (2022), "Tail Forecasting with Multivariate Bayesian Additive Regression Trees," *International Economic Review*, 64, 979-1022.

Coe, Patrick J., and Shaun P. Vahey (2020), "Financial Conditions and the Risks to Economic Growth in the United States since 1875," CAMA Working Papers 2020-36.

Delle Monache, Davide, De Polis, Andrea, and Ivan Petrella (2023), "Modeling and Forecasting Macroeconomic Downside Risk," *Journal of Business and Economic Statistics*, 42, 1010-1025.

Eaton, Jonathan, Kortum, Samuel, Neiman, Brent, and John Romalis (2016), "Trade and the Global Recession," *American Economic Review*, 106, 3401-3438.

Gertler, Mark, and Nobuhiro Kiyotaki (2010). "Financial Intermediation and Credit Policy in Business Cycle Analysis," *Handbook of Monetary Economics*, 3, chapter 11, 547-599, Elsevier.

Gourio, Francois (2012), "Disaster Risk and Business Cycles," *American Economic Review*, 102, 2734-2766.

Hengge, Martina (2019), "Uncertainty as a Predictor of Economic Activity," IHEID Working Paper 19-2019.

Houthakker, Hendrik S., and Lester D. Taylor, Consumer Demand in the United States, 2nd ed. (Cambridge: Harvard University Press, 1970).

Kiley, Michael T. (2024), "Growth at Risk from Climate Change," *Economic Inquiry*, 62, 1134-1151.

Koenker, Roger and Quanshui Zhao (1996), "Conditional Quantile Estimation and Inference for ARCH Models," *Econometric Theory*, 12, 793-813.

Leamer, Edward E. (2007), "Housing is the Business Cycle," NBER Working Paper 13428.

Lee, Eun Ryung, Noh, Hohsuk, and Byeong U. Park (2014), "Model Selection via Bayesian Information Criterion for Quantile Regression Models," *Journal of the American Statistical Association*, 109, 216-229.

Lenza, Michele and Giorgio E. Primiceri (2022), "How to Estimate a Vector Autoregression after March 2020," *Journal of Applied Econometrics*, 37, 688-699.

Li, Youjuan, and Ji Zhu (2008), "L_1-Norm Quantile Regression," *Journal of Computational and Graphical Statistics*, 17, 163-185.

Li, Youjuan, and Ji Zhu (2012), "Semiparametric Quantile Regression with High-Dimensional Covariates," *Stat Sin*, 22(4), 1379-1401.

Mankiw, Gregory (1985), "Consumer Durables and the Real Interest Rate," *The Review of Economics and Statistics*, 67, 353-362.

Mitchell, James, Poon, Aubrey, and Dan Zhu (2024), "Constructing Density Forecasts from Quantile Regressions: Multimodality in Macro-Financial Dynamics," *Journal of Applied Econometrics*, forthcoming.

Reichlin, Lucrezia, Ricco, Giovanni, and Thomas Hasenzagl (2020), "Financial Variables as Predictors of Real Growth Vulnerability," Deutsche Bundesbank Discussion Papers 05/2020.

West, Kenneth D. (1996), "Asymptotic Inference About Predictive Ability,"*Econometrica* 64, 1067-1084.

Table 1: CM Tests of Equal Tick Loss

	$\tau = 1$					$\tau = 4$				
Series\\α	0.05	0.1	0.5	0.9	0.95	0.05	0.1	0.5	0.9	0.95
GDP	-0.460*	-0.357*	0.095*	0.029	0.027	-0.457*	-0.443*	0.094	0.080	0.048
	(0.00)	(0.00)	(0.01)	(0.31)	(0.20)	(0.00)	(0.00)	(0.11)	(0.22)	(0.38)
C	-0.169*	-0.063*	0.052*	0.069	0.008	-0.301	-0.404*	0.092	0.033	-0.010
	(0.00)	(0.03)	(0.02)	(0.23)	(0.45)	(0.21)	(0.00)	(0.17)	(0.2)	(0.46)
I	-0.329*	-0.266*	-0.042*	-0.0281*	-0.034	-0.342*	-0.304*	-0.174*	0.095	0.024
	(0.00)	(0.00)	(0.00)	(0.00)	(0.11)	(0.00)	(0.00)	(0.00)	(0.10)	(0.33)
G	0.030*	0.036*	-0.025*	-0.004	0.069*	-0.102*	-0.081	-0.006	0.0228	0.002
	(0.02)	(0.01)	(0.05)	(0.48)	(0.05)	(0.02)	(0.26)	(0.44)	(0.44)	(0.49)
NX	-0.066*	0.012	-0.018	-0.087*	-0.101*	-0.154*	-0.050	-0.072	-0.270*	-0.248 *
	(0.00)	(0.42)	(0.14)	(0.00)	(0.03)	(0.00)	(0.14)	(0.11)	(0.00)	(0.00)

Notes: This table gives the percentage change in the average tick loss when lagged NFCI is added to a baseline model that only includes an intercept and lagged real GDP growth as a predictor (e.g., -0.460 denotes a 46% lower average tick loss in the unrestricted model than in the restricted model). *P*-values associated with two-sided OOS-t tests are provided in parentheses. Asterisks indicate statistical significance at the 5% level. Results are reported for the horizons $\tau = 1$ and $\tau = 4$, and the quantiles $\alpha \in \{0.05, 0.1, 0.5, 0.9, 0.95\}$. All models are estimated with a recursive scheme.

Figure 1: Real-time composition of growth-at-risk

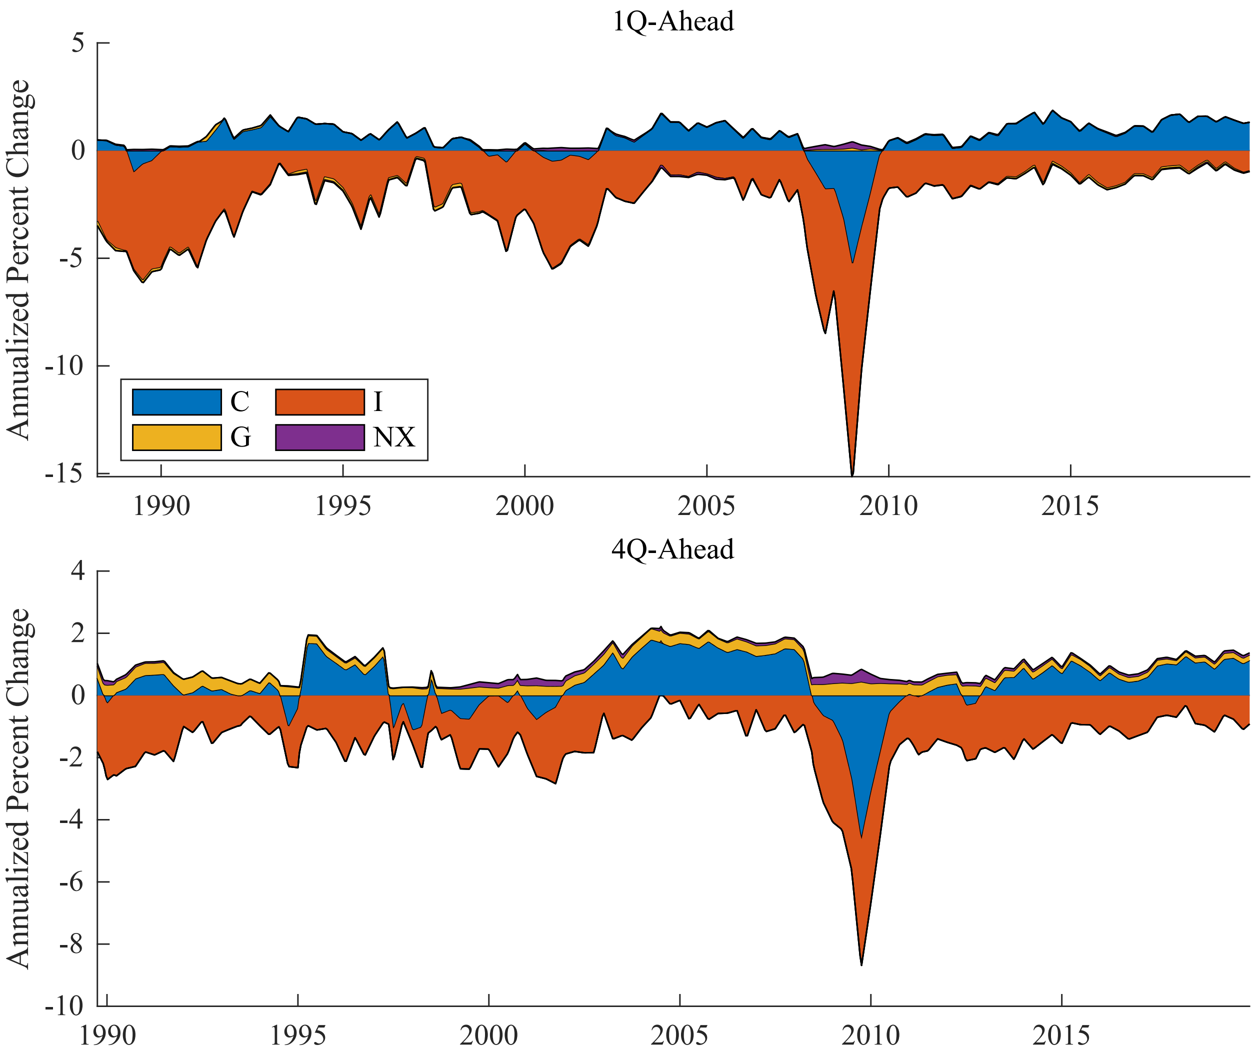

Notes: This figure plots the 1-quarter- and 4-quarter-ahead expected shortfall of GDP, weighted by each of its components (aCES). Specifically, C, I, G, and NX, give the contributions of consumption, investment, government spending, and net exports to growth-at-risk. Forecasts are out-of-sample and use a recursive estimation scheme.

Figure 2: Marginal Component Expected Shortfall

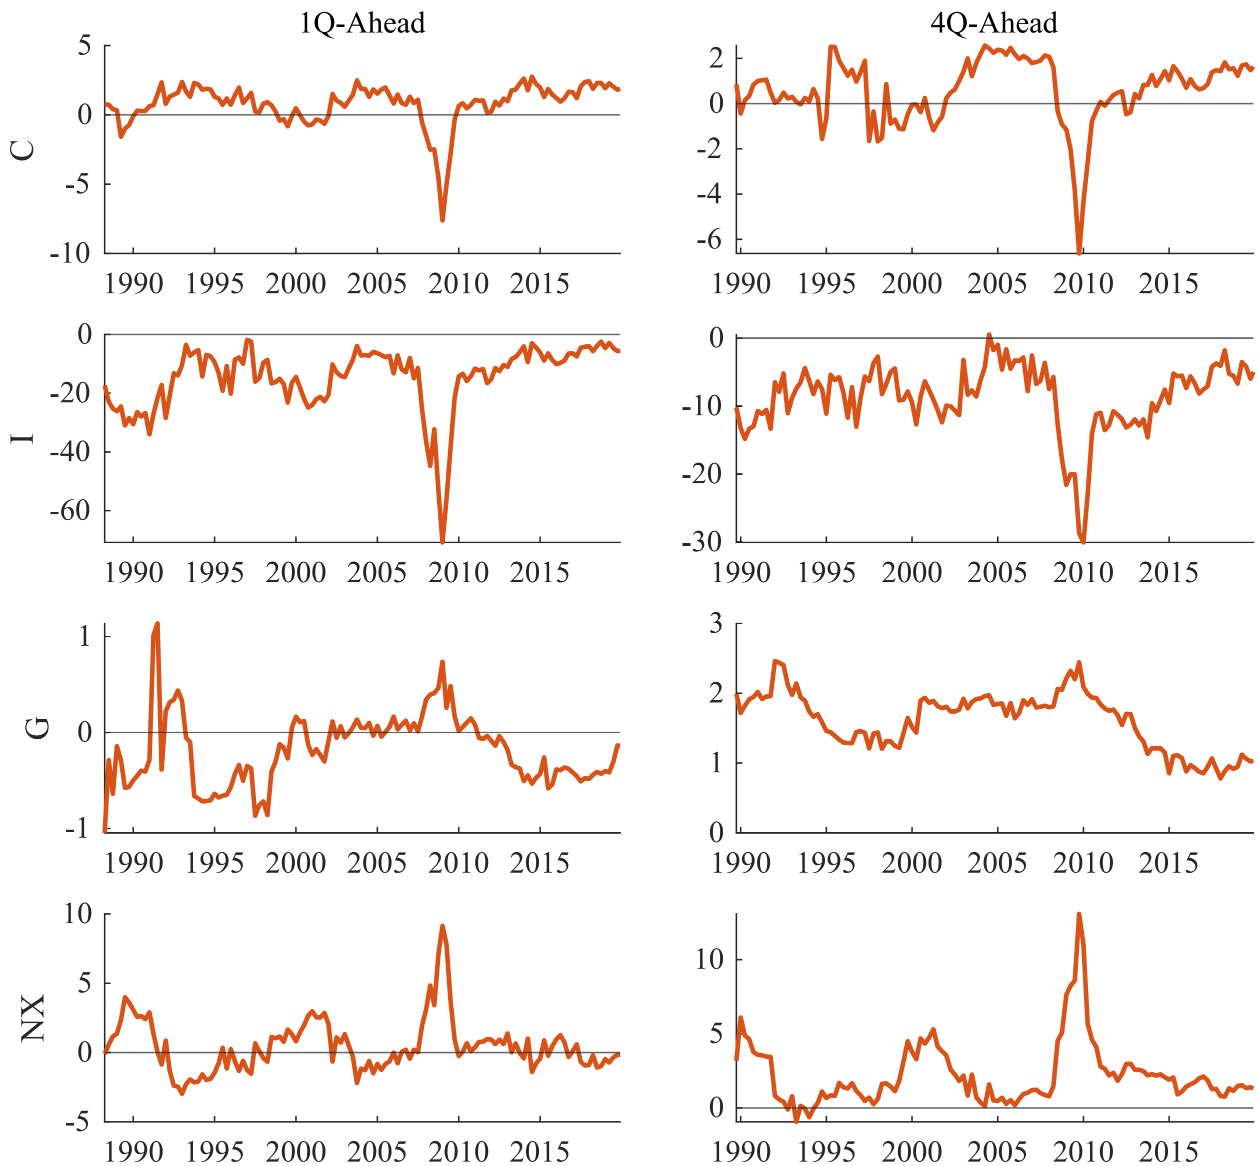

Notes: This figure plots the 1-quarter- and 4-quarter-ahead aMES of each GDP component. Forecasts are out-of-sample and use a recursive estimation scheme.

Figure 3: Real-time composition of growth-at-risk: detailed

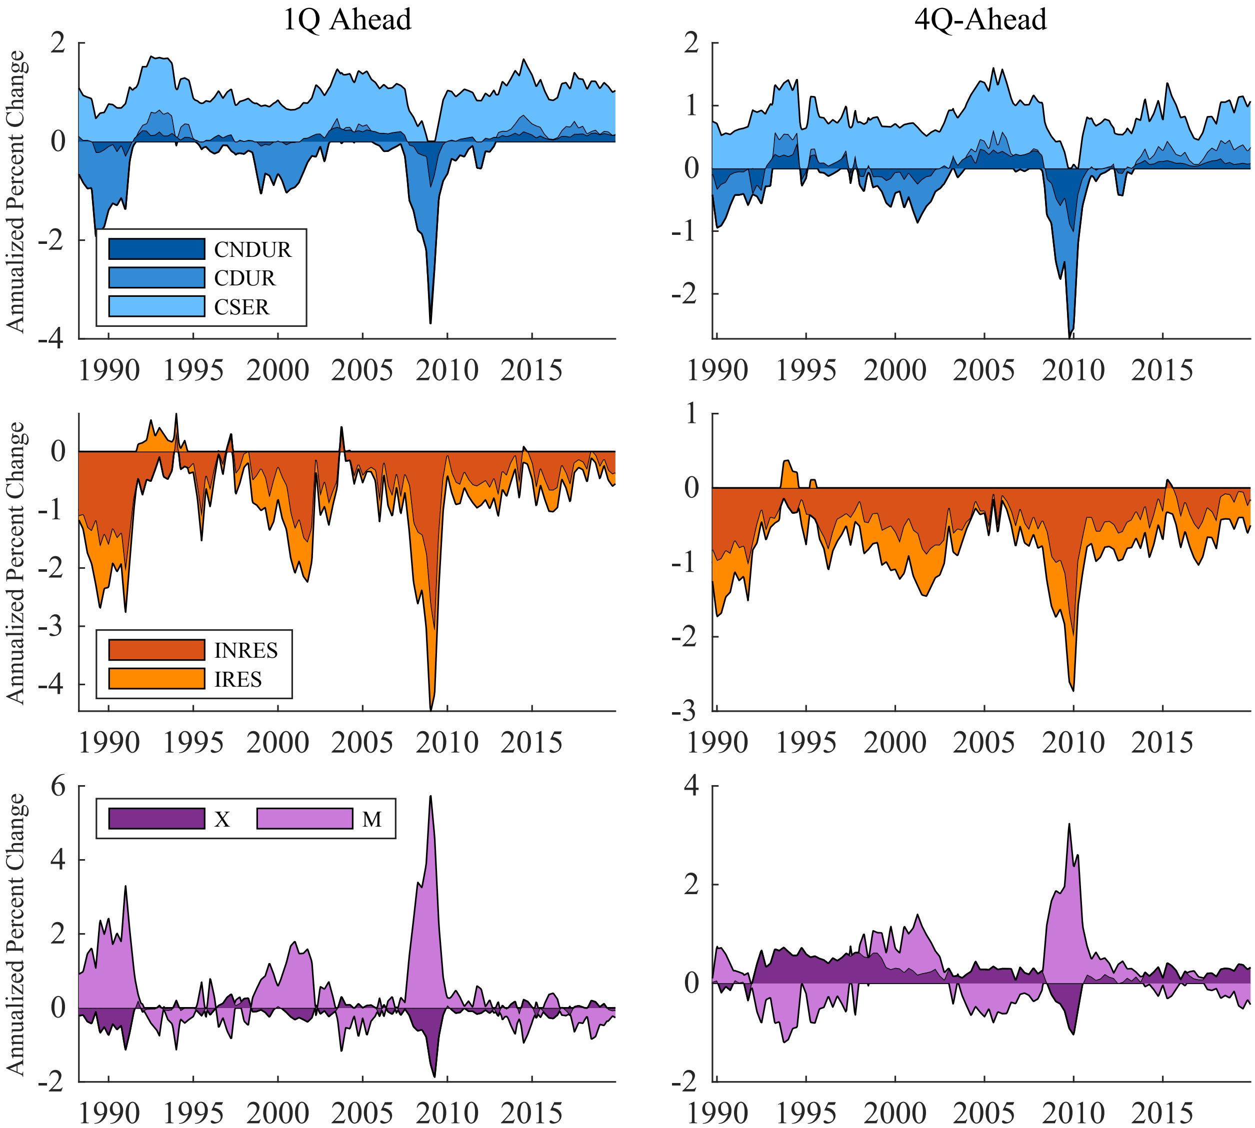

Notes: This figure plots the 1-quarter- and 4-quarter-ahead detailed composition of the expected shortfall. The top, middle, and bottom panels plot the detailed aMES of consumption, investment, and net exports, respectively. Each component is broken down into its subcomponents. For brevity, we exclude government spending. Forecasts are out-of-sample and use a recursive estimation scheme.

Figure 4: ℓ_1-norm Penalized QR with NFCI Subindices

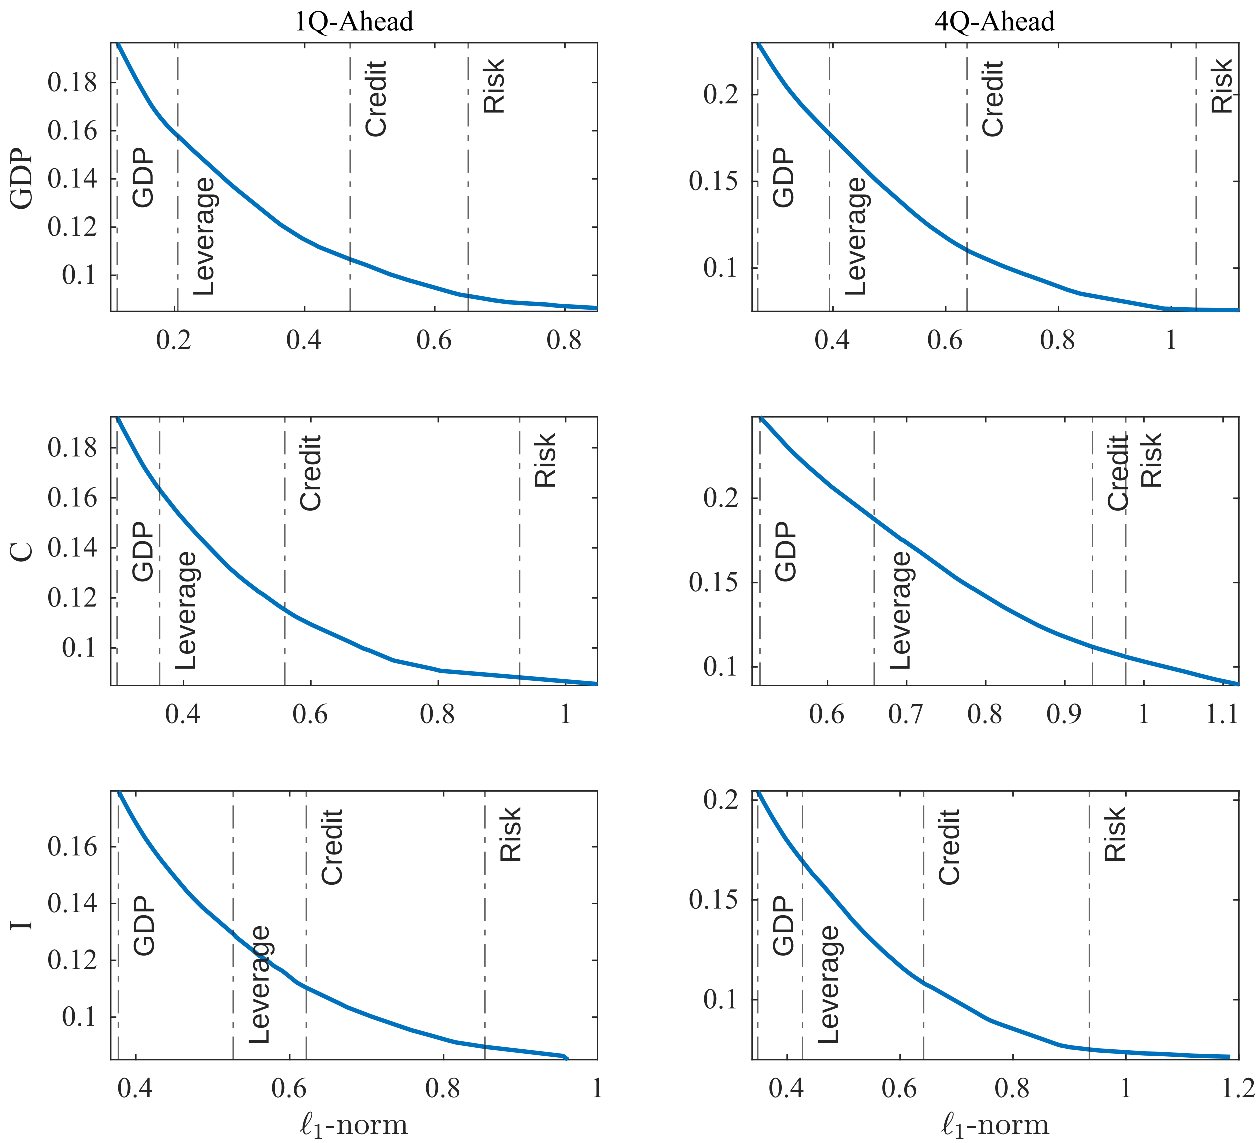

Notes: This figure plots the average tick loss associated with the ℓ_1-norm penalized 0.05-quantile regressions of 1-quarter- and 4-quarter-ahead GDP, consumption, and investment on contemporaneous GDP growth and the leverage, credit, and risk NFCI subindices. Regressions are run for a large set of tuning parameters, λ, and the x-axis displays the ℓ_1-norm associated with a given regression. Larger ℓ_1-norms are associated with smaller penalties. Vertical lines indicate at what penalty level a series' coefficient becomes well separated from zero, i.e., lines further to the right indicate a smaller penalty required for a series to be selected by the model.

Online Appendix

A: Extra Figures and Tables

Figure A1: Real-time composition of growth-at-risk, NFCI excluded

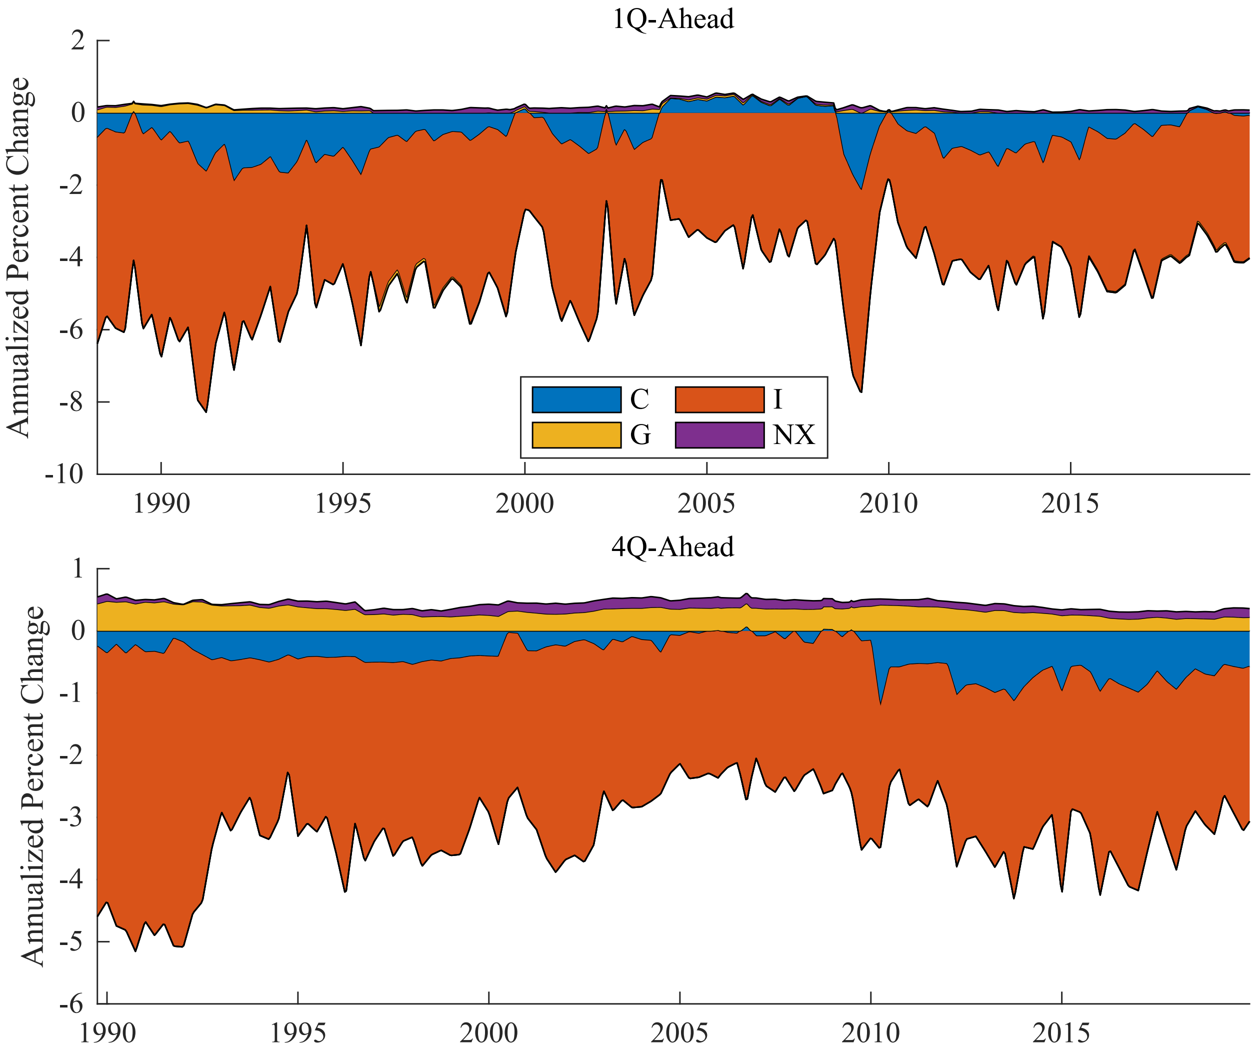

Notes: See Figure 1 for details. The NFCI is excluded as a predictor.

Figure A2: Expected Shortfall and Sum of aCES.

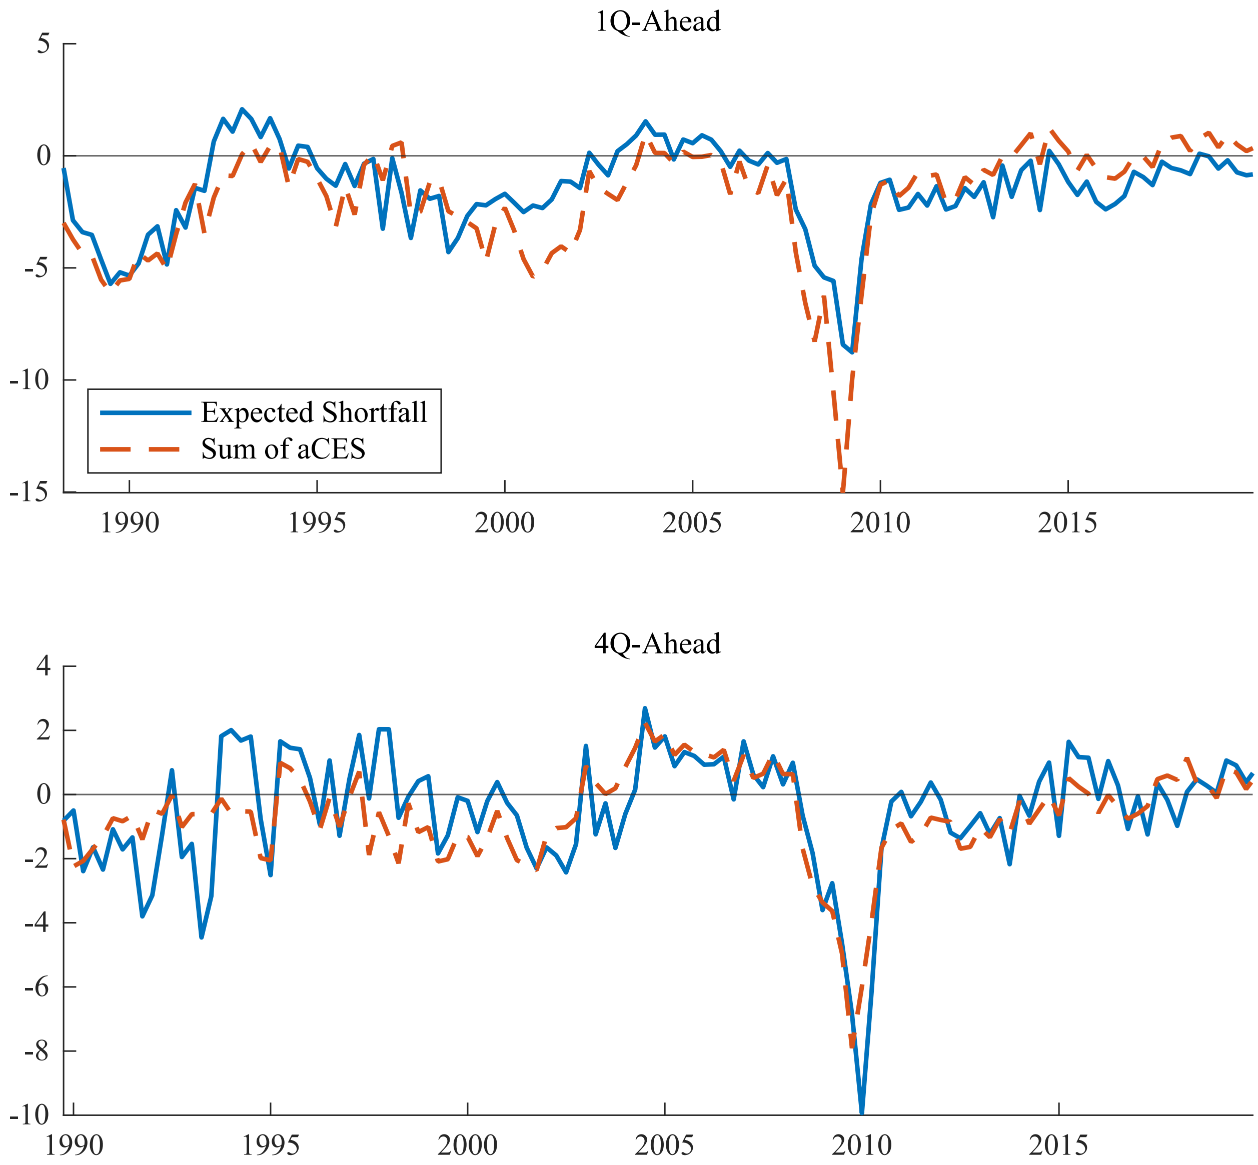

Notes: This figure plots the out-of-sample expected shortfall of GDP growth against the sum of the aCES at the 1- and 4-quarter horizons.

Figure A3: NFCI Subindices

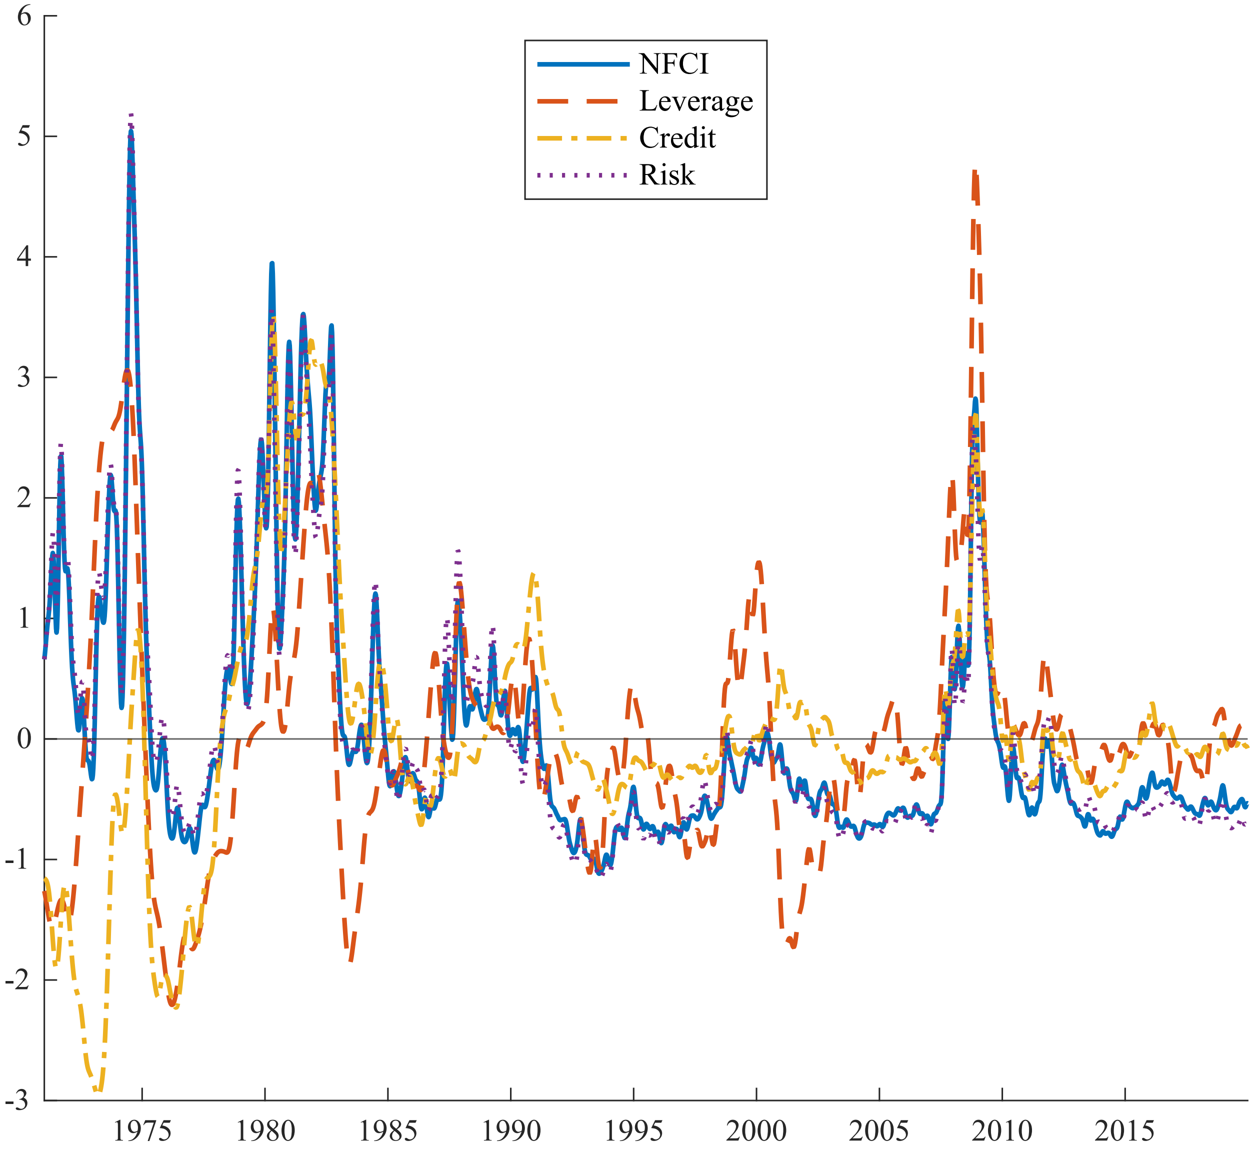

Notes: This figure plots the NFCI as well as its leverage, credit, and risk subindices.

Figure A4: ℓ_1-norm Penalized QR with NFCI Subindices, Coefficients

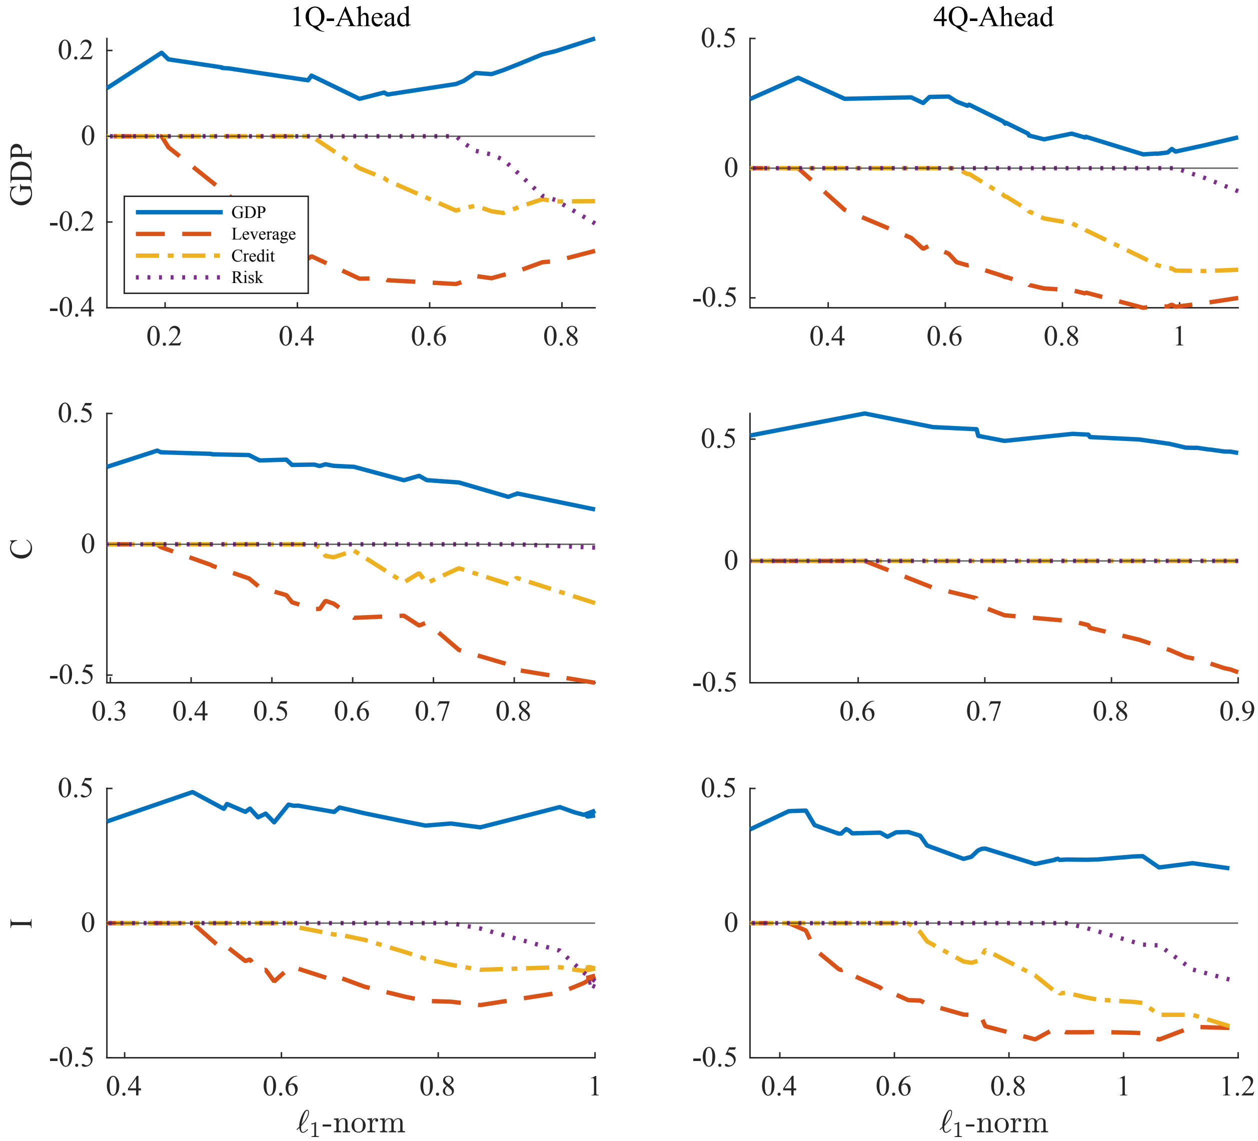

Notes: This figure plots the coefficients associated with the ℓ_1-norm penalized 0.05-quantile regressions of 1-quarter- and 4-quarter-ahead GDP, consumption, and investment on contemporaneous GDP growth and the leverage, credit, and risk NFCI subindices. Regressions are run for a large set of tuning parameters, λ, and the x-axis displays the ℓ_1-norm associated with a given regression. Larger ℓ_1-norms are associated with smaller penalties.

Figure A5: ℓ_1-norm Penalized QR with NFCI Subindices, 0.95-Quantile

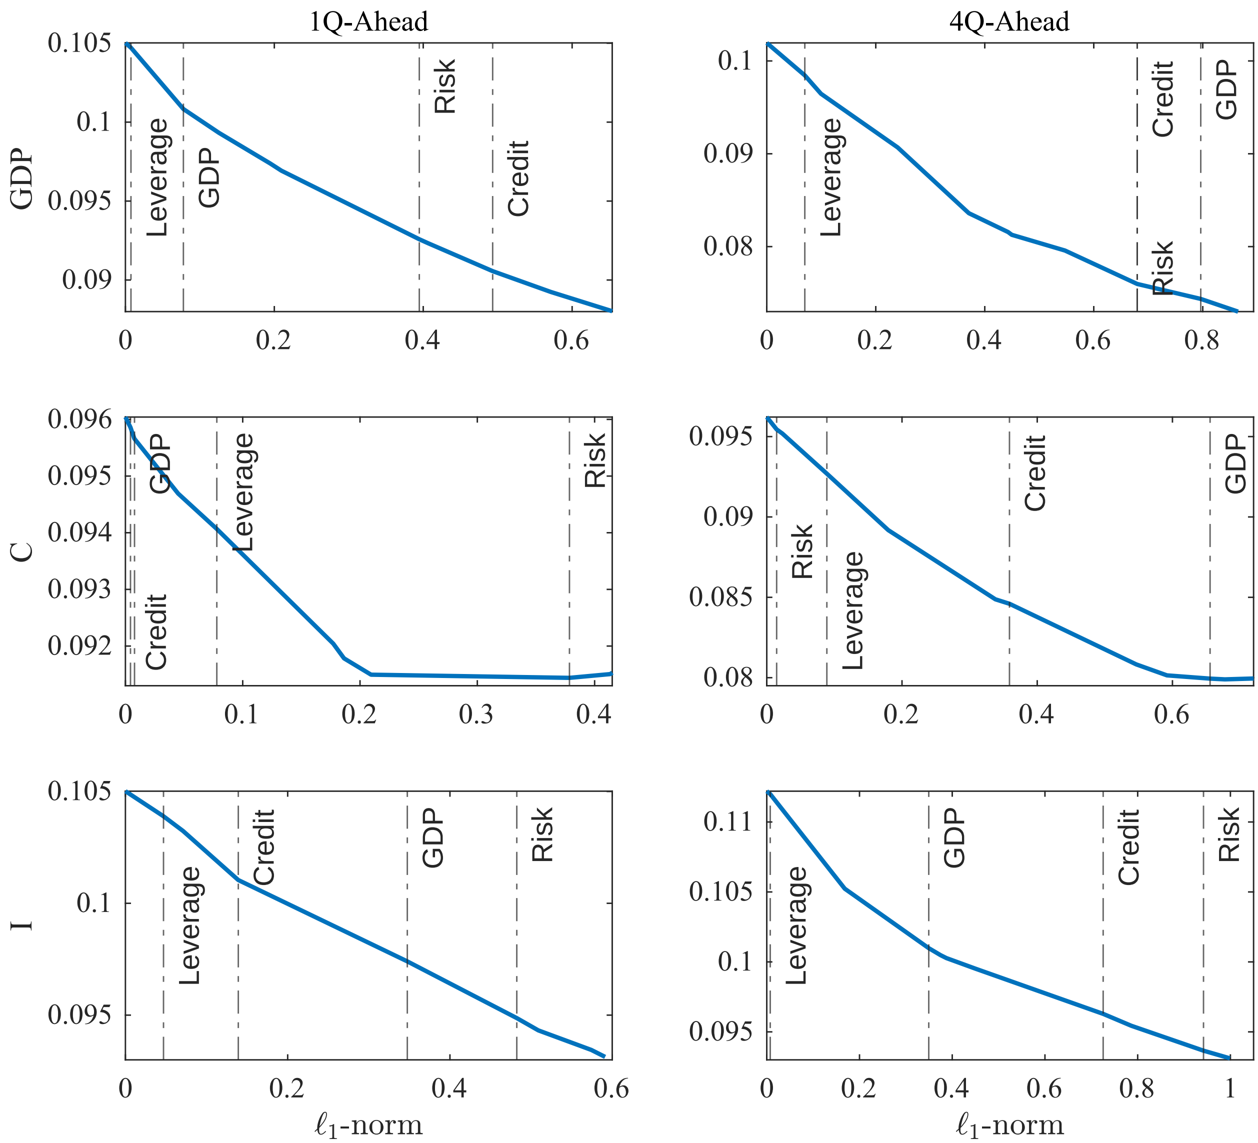

Notes: See Figure 4 for details. All regressions are for the 0.95-quantile.

Figure A6: ℓ_1-norm Penalized QR with NFCI Subindices, Coefficients, 0.95-Quantile

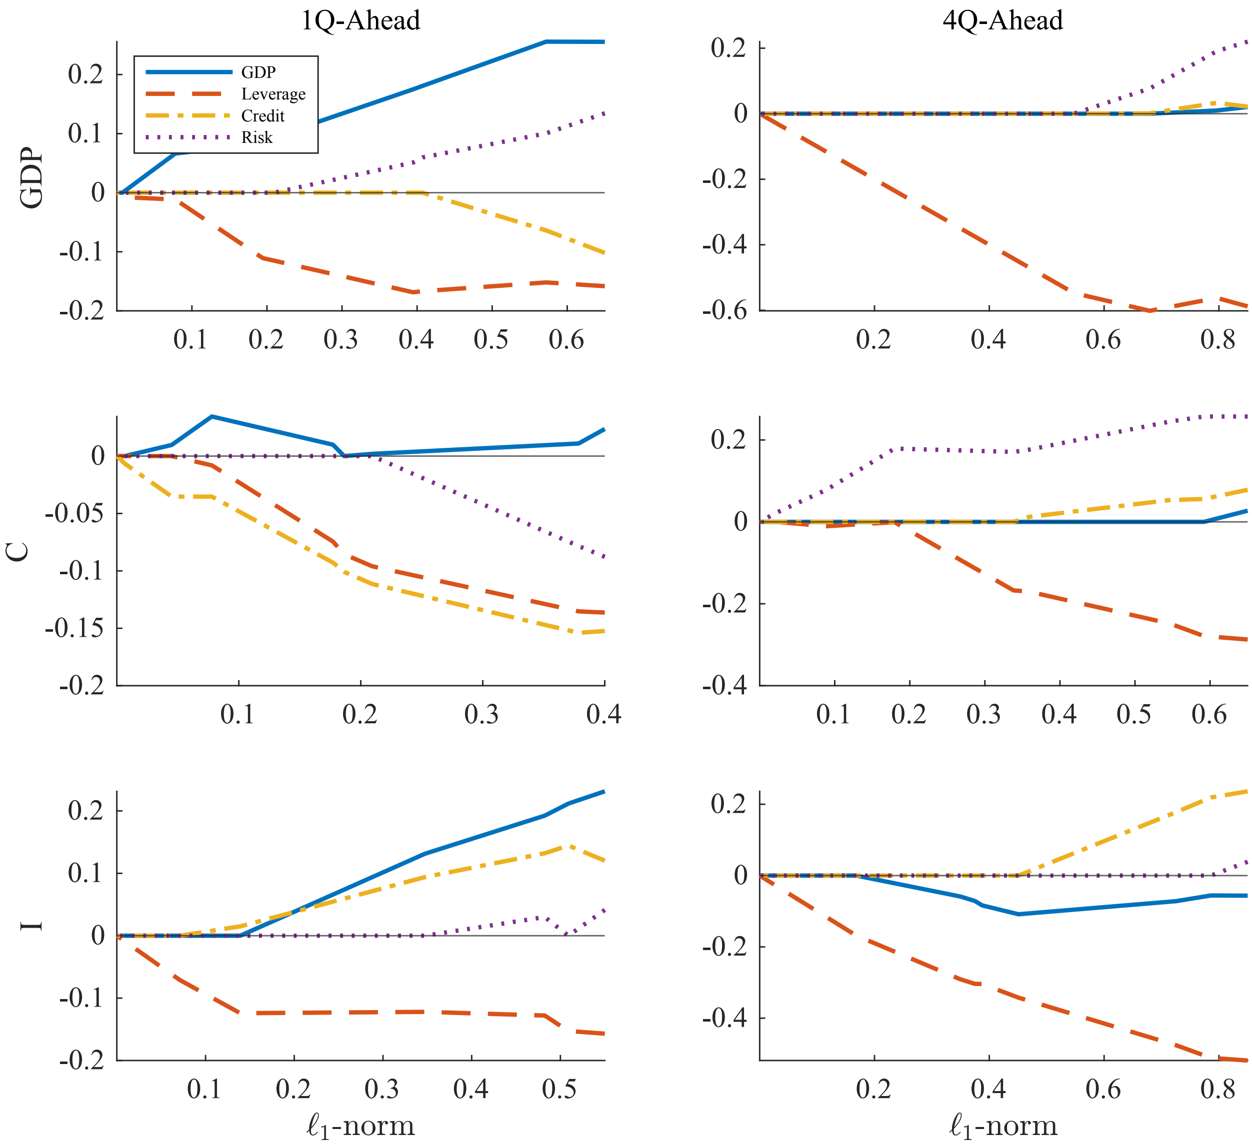

Notes: See Figure A4 for details. All regressions are for the 0.95-quantile.

Table A1: Copula Parameters

	$\tau = 1$		$\tau = 4$	
	ν_w	σ_w	ν_w	σ_w
C	2.368	0.610	9.385	0.880
I	14.056	0.815	2.457	0.838
G	6.043	0.284	1.585e7	0.082
NX	2.889	-0.204	1.860	-0.536

Notes: This table reports the estimated hyperparameters of the Student's t copula between GDP and each of its components at the 1- and 4-quarter ahead horizons. Estimates are obtained over the full sample using maximum likelihood.

B: aCES Derivation

As noted in the main text, the sum of the aCES is only an approximation of ES. We can see this if we breakdown the expected shortfall of GDP growth into component shares of growth using the following steps.

$$
\begin{aligned}
\text{ES} &= E(y_{t+\tau}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|x_t)) \\
&= E(\frac{Y_{t+\tau} - Y_t}{Y_t}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|x_t)) \\
&= E\left(\sum_{W\in(C,I,G,NX)} \frac{W_{t+\tau} - W_t}{Y_t}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|x_t)\right) \\
&= E\left(\sum_{W\in(C,I,G,NX)} \frac{W_{t+\tau} - W_t}{W_t} * \frac{W_t}{Y_t}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|x_t)\right) \\
&= \sum_{W\in(C,I,G,NX)} E\left(\frac{W_{t+\tau} - W_t}{W_t} * \frac{W_t}{Y_t}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|x_t)\right) \\
&\stackrel{?}{=} \sum_{W\in(C,I,G,NX)} \frac{W_t}{Y_t} * E\left(\frac{W_{t+\tau} - W_t}{W_t}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|x_t)\right) \\
&= \sum_{W\in(C,I,G,NX)} \frac{W_t}{Y_t} * E\left(w_{t+\tau}|y_{t+\tau} < Q_\alpha(y_{t+\tau}|x_t)\right) \\
&= \sum_{w\in(c,i,g,nx)} \text{aCES}_w,
\end{aligned}
$$

All but one of the steps is by definition, notation, or a simple rearrangement of terms. The 6th equality requires an assumption. It requires that the information set induced by x_t be equivalent to the information set induced by x_t and $\frac{W_t}{Y_t}$ for each of the four possible W_t. If this is true, as it would if the the shares are constant across time, $\frac{W_t}{Y_t}$ can be pulled out of the expectations operator and we obtain an equality. Since $\frac{W_t}{Y_t}$ remains relatively stable over the sample, we choose to treat them as approximately constant, with the expectation this assumption is inconsequential. Indeed, Figure A2 indicates that the weighted sum of estimated aCES closely resembles the estimated ES. Nonetheless, for transparency we choose to refer to our estimate as aCES rather than CES, and aMES rather than MES by implication.

C: COVID-19

In the main analysis we omitted forecasts with horizons after 2019Q4. As noted in Amburgey and McCracken (2023), the enormous swings in GDP growth led to wildly different parameter estimates from quantile regressions when using that period for estimation.[17] These in turn led to highly inaccurate tail forecasts that were not in-line with the accuracy seen in earlier periods. Nonetheless, COVID-19 caused an unprecedented shock to the economy, and there is value to identifying how the components-at-risk evolved throughout that period.

To do so, we apply the methodology described in section 3.1 when constructing the conditional marginal and joint distributions – with one modification. Rather than re-estimate the quantile regression parameters recursively across the COVID-19 period, for each forecast origin after 2019Q4, we fix the parameter estimates at that value. By doing so we maintain most of the "real-time" nature of our exercises (i.e., by using vintage data), albeit with a look-ahead bias due to how we manage the parameter estimates.

[17] Lenza and Primiceri (2022) and Chernis et al. (2023) note comparable issues in the context of VARs and quantile regressions, respectively.

With that in mind, Figure C1 provides the composition of growth-at-risk from 2020Q1 through 2022Q2. Aside from the forecasts associated with origins 2020Q3 and 2020Q4, the composition remains by and large the same as that in Figure 1. Investment-at-risk is omnipresent with consumption-at-risk remaining positive throughout. On the other hand, the forecasts associated with origins 2020Q3 and 2020Q4 (i.e, 2020Q3 and 2020Q4 for 1-quarter ahead, and 2021Q2 and 2021Q3 for 4-quarter ahead) display wild swings due to the unstable nature of GDP growth during these periods. Specifically, in 2020Q2 and 2020Q3, GDP growth was -35% and 30%, respectively.

Focusing on the 1-quarter horizon, the approximate expected shortfall (i.e., the total area of the aCES) drops to a record low of -16% in 2020Q3 before jumping to a record high of 6% in 2020Q4. Perhaps most notably, for the first and only time across our sample, investment-at-risk contributes positively to downside-risk in 2020Q4. Patterns in the 4-quarter horizon forecasts mirror those outlined above.

Figure C2 provides the detailed composition of growth-at-risk over the same period. At the peak of the recession, for the first time in our sample, services meaningfully contribute negative values to growth-at-risk. This reflects the large drop in service-related activity caused by government-mandated closures of businesses in the industry. Consumer durables and nondurables display similar patterns to that of the Great Recession, albeit with nondurables making up a more substantial portion of downside risk at the height of the downturn. The middle panel of Figure C2 tells us that non-residential investment was the driver of investment's swing from negative to positive contributions during this period. For example, at the 1-quarter horizon, nonresidential investment's aCES swung from -9% to 4.5% between 2020Q3 and 2020Q4. Even more so than our main results, this stands in stark contrast to Leamer (2007), who argues that residential investment plays a more prominent role in recessions.

Figure C1: Real-time composition of growth-at-risk, COVID-19

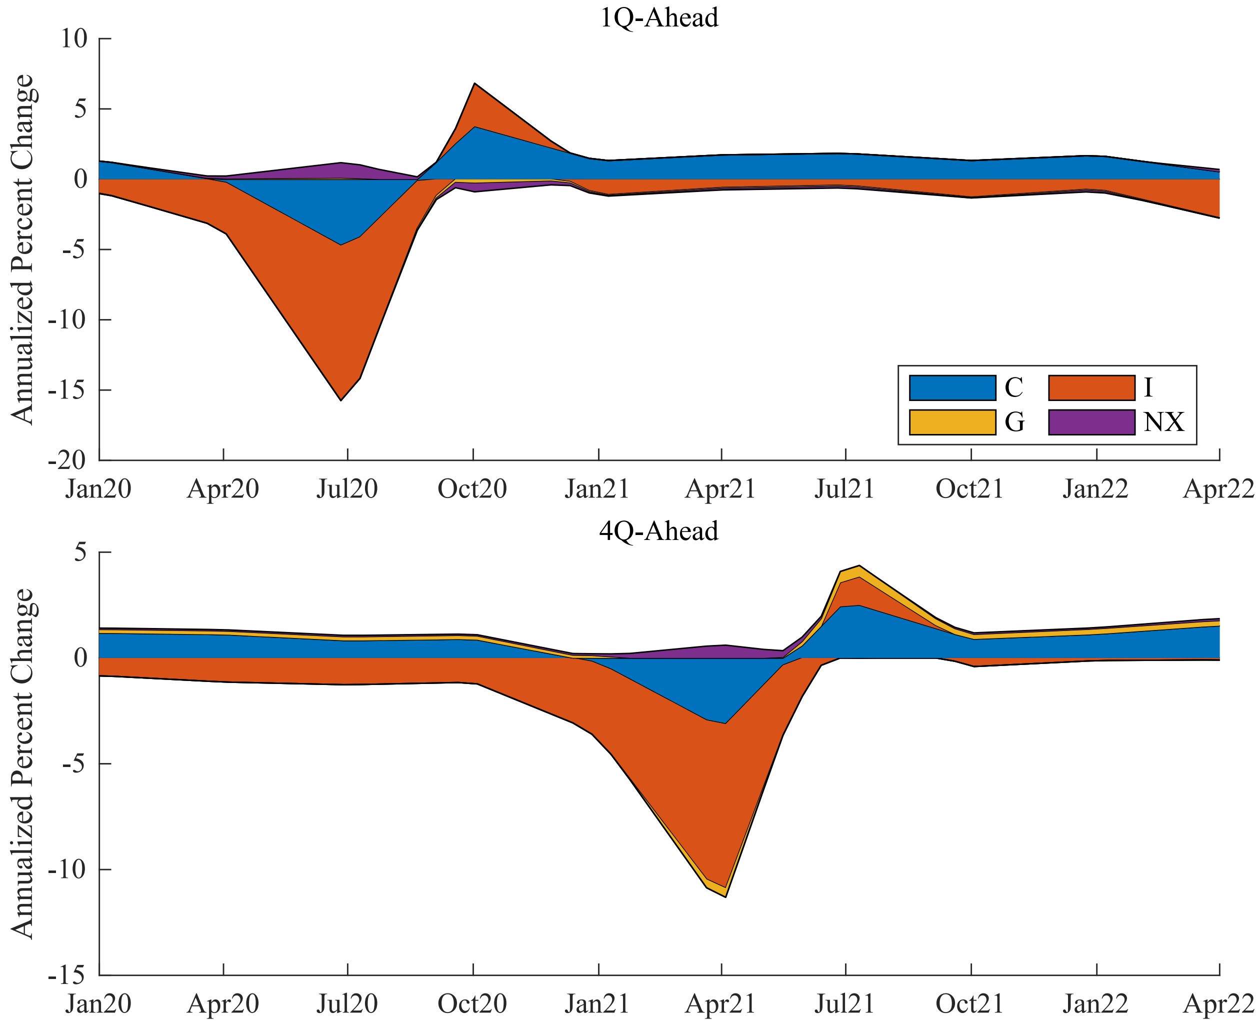

Notes: See Figure 1 for details.

Figure C2: Real-time composition of growth-at-risk: detailed, COVID-19

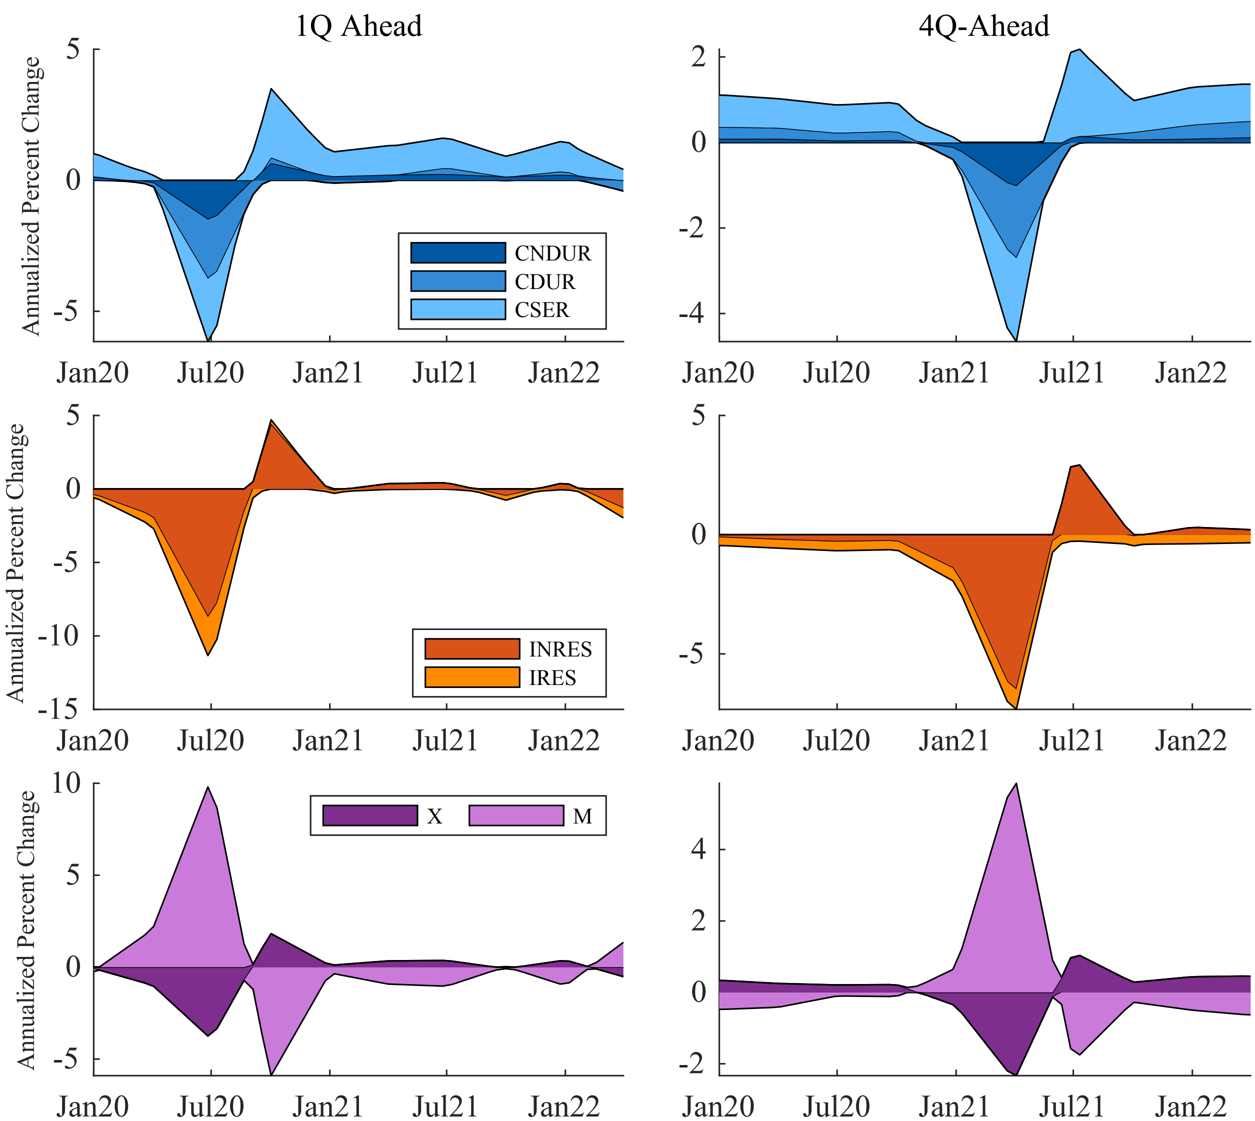

Notes: See Figure 3 for details.