


ARTICLE

A tale of two tails: 130 years of growth at risk*

Martin Gächter^{1,2}, Elias Hasler^{1,2} , and Florian Huber³

¹Liechtenstein Financial Market Authority, Vaduz, Liechtenstein

²University of Innsbruck, Innsbruck, Austria

³University of Salzburg, Salzburg, Austria

Corresponding author: Elias Hasler; Email: hasler.elias@hotmail.com

Abstract

We extend the growth-at-risk (GaR) literature by examining US growth risks over 130 years using a time-varying parameter stochastic volatility regression model. This model effectively captures the distribution of GDP growth over long samples, accommodating changing relationships across variables and structural breaks. Our analysis offers several key insights for policymakers. We identify significant temporal variation in both the level and determinants of GaR. The stability of upside risks to GDP growth, as seen in previous research, is largely confined to the Great Moderation period, with a more balanced risk distribution prior to the 1970s. Additionally, the distribution of GDP growth has narrowed significantly since the end of the Bretton Woods system. Financial stress is consistently associated with higher downside risks, without affecting upside risks. Moreover, indicators such as credit growth and house prices influence both downside and upside risks during economic booms. Our findings also contribute to the financial cycle literature by providing a comprehensive view of the drivers and risks associated with economic booms and recessions over time.

Keywords: Growth at risk; financial crises; business cycles; tail forecasting

JEL classifications: C11; C53; E32; E44; G01; N10

1. Introduction

The empirical growth-at-risk (GaR) concept introduced by Adrian et al. (2019) suggests that deteriorating financial conditions are associated with increased downside risks to economic growth. While standard forecasts focus on the expected value of future GDP growth, the GaR approach places a particular emphasis on the probability and magnitude of potentially adverse outcomes. Similar to the value-at-risk concept in finance, the GaR of an economy for a given time horizon is defined as a specific low quantile of the distribution of the projected GDP growth rate for the respective horizon. In this context, Adrian et al. (2019) show that the left tail of the distribution of (projected) GDP growth is less stable and more affected by financial conditions than the upper quantiles of the distribution. Against this background, the GaR concept is a useful and

*We would like to thank participants of the 11th European Seminar on Bayesian Econometrics and the Annual Meeting of the Austrian Economic Association for comments and suggestions. Huber gratefully acknowledge financial support from the Austrian Science Fund (FWF, grant no. ZK 35) and the Oesterreichische Nationalbank (OeNB, Anniversary Fund, project no. 18304). The views expressed in this paper are those of the authors and do not necessarily represent those of the Liechtenstein Financial Market Authority.

intuitive policy tool to identify and quantify systemic risk and has therefore gained traction among policymakers in recent years.

In the last few years, the GaR idea has been extended in various directions, for example, adding various risk indicators from the financial cycle literature (see, for instance, Aikman *et al.*, 2019) or by examining the term structure of GaR (Adrian *et al.*, 2022). In this context, the GaR framework is used as a composite indicator for systemic risk at the country level and can, therefore, also be taken as an indicator when to activate macroprudential policy measures. Consequently, recent research has also taken into account the impact of various policy instruments on GaR (e.g., Galán, 2024). Remarkably, the entire strand of literature has relied entirely on data samples back to the 1970s without taking into account earlier developments in the last century.¹ This is insofar surprising, as the main objective of macroprudential policy is the prevention of financial crises (or alternatively, the reduction of the costs of such crises if they occur). Since financial crises appear infrequently, as also shown by the widely cited financial cycle literature (Schularick and Taylor, 2012; Jordà *et al.*, 2017), a long time series of the underlying drivers is crucial to capture the tail risks of the variables of interest.

Using long time series, however, raises additional econometric issues. For instance, these time series might be subject to structural breaks in the conditional mean. Such a behavior might reflect changes in key relationships between the conditional distribution of output growth on macro-financial determinants. When conditioned on a particular quantile, standard quantile regressions (QRs) assume that the regression relationship between a given covariate and the corresponding quantile function is constant, and hence constancy of the parameter is assumed throughout. In this paper, we build on recent papers that show that simple heteroskedastic models perform well (or even better) than QRs (Carriero *et al.*, 2024; Brownlees and Souza, 2021) and propose using a time-varying parameter stochastic volatility regression model (TVP-SV) to capture changes in the relevance of different drivers of upside and downside risks to output growth. We use this framework to analyze an extended data set covering 130 years of macroeconomic data and show that it works well for predicting downside risks to GDP.² After providing evidence that such a model is competitive to standard QRs, we back out the contributions of individual variables over time using state-of-the-art techniques from statistics and machine learning (see, e.g., Crawford *et al.*, 2019; Woody *et al.*, 2021).³ Thus, by applying novel methods to historical data, we are able to draw important policy implications for today's policymakers.

Our findings provide a comprehensive historical context and several key policy implications. First, we demonstrate that the stability of upside risks to GDP growth observed in previous research is time-specific, to the period of the Great Moderation, with a more symmetric risk distribution prior to the 1970s. Both upside and downside risks to GDP have substantially decreased over time, remaining relatively low since the start of the Great Moderation. We also find that financial stress consistently correlates with higher downside risks, although the magnitude of this effect fluctuates over time, while upside risks remain unaffected by financial stress. Notably, the impact of credit and house price growth exhibit significant temporal variation. Specifically, the effect of credit growth during the Great Depression and the Great Financial Crisis show remarkable similarities, and the negative impact of house prices on growth risks during the Global Financial Crisis was unprecedented. Interestingly, credit growth and house prices drive both downside and upside risks, leading to a wider distribution of expected GDP growth. Our findings enrich the financial cycle literature by examining the entire distribution of expected GDP, thus providing a more comprehensive picture of the drivers and risks of economic booms and recessions over time. This enhanced understanding is crucial for the effective calibration and timing of macroprudential and monetary policy measures.

The remainder of the paper is structured in the following way. The next section fleshes out the data set and econometric methods adopted. Section 3 includes the empirical findings and consists of an out-of-sample tail forecasting exercise, provides quantitative evidence on the predictive distributions of GDP growth, and then proceeds by discussing the key drivers of GaR. The final section concludes the paper.

2. Data and methods

2.1. Data

Our analysis is based on a US data set stretching from 1893Q1 to 2016Q4. Using three different sources, our data includes annualized real GDP growth, a financial stress indicator, the 3-year average growth rate of the credit-to-GDP ratio, and the 3-year average growth rate of real house prices. The real GDP time series is taken from D’Agostino and Surico (2012) and updated using FRED data.⁴ The dependent variable is constructed by using the logarithm of real GDP, Y_t , and converting it into the annualized growth rate h periods ahead, $y_{t+h} = \frac{Y_{t+h} - Y_t}{h/4}$.

In line with previous literature, we include a measure of financial stress as an explanatory variable. However, unlike other financial stress indicators, which are typically based on financial data and start in the 1970s, we use a historical newspaper-based financial stress indicator built by Püttmann (2018). As Püttmann (2018) notes, the index exhibits a long-run trend; hence, we detrend the time series using a slow-moving Hodrick–Prescott filter with a λ of 5×10^6 .

While financial stress measures are highly relevant for short-term GaR estimations—at least since the 1970s—credit-to-GDP growth and house prices are frequently used as a signal for medium-term financial imbalances (Aikman et al., 2019; Galán, 2024). Unfortunately, consistent quarterly data for credit (loans to the nonfinancial private sector) and house prices are not available for such a long time horizon. Therefore, we use data from Jordà et al. (2017) and convert the annual data into quarterly by using the quadratic spline of Forsythe et al. (1977). Subsequently, annualized 3-year averages of the log differences of the credit-to-GDP ratio as a measure of credit growth and the 3-year average of the log differences of real house prices as a measure of house price growth are used.⁵

2.2. Growth at risk through the lens of TVP regressions

Downside risks are typically analyzed through QRs (see Adrian et al., 2019). However, one key shortcoming of QRs is that *within a given quantile*, they are not capable of capturing structural breaks in the regression coefficients, a feature that is crucial given the length of the time series we analyze. A natural way of capturing time variation in the parameter is through time-varying parameter (TVP) regression models. These models assume that the coefficients evolve smoothly over time and thus capture changes in transmission mechanisms but, conditional on appropriate modeling assumptions, also allow for rapid movements in the underlying parameters. To capture changing volatilities of the shocks, we also allow for heteroskedasticity in the regression model through a standard stochastic volatility specification. Our simple, yet flexible model enables us to capture differences in the relations between the determinants of GaR but also allows for situations where large, unobserved shocks are the main drivers of tail risks.

In its general form, we consider predictive equations with drifting parameters that take the following form:

$$y_{t+h} = \beta'_{t+h}x_t + \varepsilon_{t+h}, \quad \varepsilon_{t+h} \sim \mathcal{N}(0, \sigma^2_{t+h}) \tag{1}$$

where β_{t+h} is a vector of TVPs, which link y_{t+h} to our set of K macrofinancial covariates in x_t . We follow much of the literature (see, e.g., Primiceri, 2005; Cogley and Sargent, 2005) and assume that these evolve according to random walk processes. Moreover, the logarithm of the error variance is assumed to follow an AR(1) process.⁶

These assumptions give rise to a system of state evolution equations:

$$\begin{aligned} \beta_t &= \beta_{t-1} + \eta_t, & \eta_t &\sim \mathcal{N}(\mathbf{0}_K, V_\beta), \\ \log \sigma_t^2 &= \mu_\sigma + \rho_\sigma (\log \sigma_{t-1}^2 - \mu_\sigma) + w_t, & w_t &\sim \mathcal{N}(0, \vartheta^2), \\ \log \sigma_0^2 &\sim \mathcal{N}\left(\mu_\sigma, \frac{\vartheta^2}{1 - \rho_\sigma^2}\right) \end{aligned}$$

where $V_\beta = \text{diag}(v_1^2, \dots, v_K^2)$ is a diagonal matrix with variances v_j^2 . These variances control the amount of time variation in the regression coefficients. If $v_j^2 = 0$, β_{jt} , the j^{th} element of β_t , would be constant over time since $\beta_{jt} = \beta_{jt-1}$ for all t . Hence, the corresponding effect of x_{jt} on y_{t+h} is time invariant. By contrast, setting v_j^2 to a large value implies substantial variation in the corresponding coefficient, giving rise to overfitting concerns.

The coefficients associated with the law of motion of the log-volatilities are the long-run unconditional mean μ_σ , the persistence parameter ρ_σ , and the innovation variance ϑ^2 . If ρ_σ is close to one, the corresponding volatility estimate will be smooth. The variance parameter ϑ^2 controls the amount of time variation in the error variances. This model has two features necessary to produce asymmetries in the predictive distribution. The TVPs provide sufficient flexibility to capture movements in the conditional mean, and the stochastic volatility process allows the variance of the predictive density to shift flexibly.

Deciding on whether we need time variation in the parameters is a nonstandard statistical problem and boils down to testing whether $v_j^2 = 0$ or $v_j^2 > 0$. The Bayesian literature offers several solutions. Shrinkage priors that force v_j^2 toward zero if the data suggests that β_{jt} is time invariant are the most prominent way (see, e.g., Belmonte *et al.*, 2014; Bitto and Frühwirth-Schnatter, 2019; Huber *et al.*, 2021, for recent papers that develop and apply different shrinkage techniques). All these priors, commonly referred to as global-local shrinkage priors (Polson and Scott, 2010), can be written as:

$$\pm\sqrt{v_j^2} \sim \mathcal{N}(0, \lambda\tau_j), \quad \tau_j \sim p_\tau(\tau_j), \quad \lambda \sim p_\lambda(\lambda),$$

where λ is a so-called global shrinkage factor. Everything else equal, a value of λ very close to zero implies that all regression coefficients in (1) are time invariant. To allow for time-varying coefficients if λ is close to zero, the τ_j 's are local shrinkage parameters that pull $\pm v_j$ away from zero, if supported by the data. Both the global and local shrinkage parameters feature their own priors, and by choosing them in a suitable manner, we end up with a wide variety of different shrinkage priors. In this paper, we use the triple Gamma shrinkage prior (Cadonna *et al.*, 2020) as implemented in the R package `shrinkTVP` (Knaus *et al.*, 2021). This choice nests many of the priors commonly used in the literature and has been shown to work well in Cadonna *et al.* (2020). More details on the priors and the posterior simulator are provided in Appendix A.⁷

In this paper, we will take a forecasting stance and consider the predictive distribution of the model in (1). The predictive density is given by:

$$p(y_{T+h} | \text{Data}_{1:t}) = \int p(y_{T+h} | \Xi, \text{Data}_{1:T}) p(\Xi | \text{Data}_{1:T}) d\Xi, \tag{2}$$

where $\text{Data}_{1:T}$ denotes the available information up to time T and Ξ is a generic object that collects coefficients and latent states.

The predictive density is not available in closed form and obtained through simulation-based techniques using the output from the Markov chain Monte Carlo (MCMC) sampler. Notice that $p(y_{T+h} | \Xi, \text{Data}_{1:T})$ is a Gaussian distribution:

$$y_{T+h} | \Xi, \text{Data}_{1:T} \sim \mathcal{N}(\beta'_{T+h} x_T, \sigma_{T+h}^2),$$

and thus symmetric. To see how our model generates heavy tails, consider a simplified version that sets $\rho_\sigma = 0$. The resulting state evolution equation for $\log \sigma_t^2$ reduces to $\log \sigma_t^2 \sim \mathcal{N}(\mu_\sigma, \vartheta^2)$, and hence, ε_{T+h} is a Gaussian term with a variance that follows a log-normal distribution. The kurtosis, in this case, will be given by $3 \cdot \exp(\vartheta^2)$ and is thus strictly larger than the kurtosis of a Gaussian distribution for $\vartheta > 0$. In general, Carnero *et al.* (2004) show that the kurtosis of ε_{T+h} is given by $k_\varepsilon = 3 \cdot \exp(\vartheta^2 / (1 - \rho_\sigma^2))$. This implies that if $\log \sigma_t^2$ is persistent, ρ_σ would be large, and kurtosis would increase, making larger shocks more likely.

Table 1. Out-of-sample model evaluation

Horizon	Pre-WWII			Post-WWII		
	TVP	TVP+	QR+	TVP	TVP+	QR+
$h = 1$	1.063	1.074	1.229	0.856	0.853	1.010
$h = 4$	1.887	1.715	1.223	0.937	0.886	1.117

Note: This table reports the out-of-sample model evaluation for the forecast horizon of one and four quarters. The quantile scores are reported relative to the QR model.

It is worth stressing, however, that other features (such as asymmetries or multimodalities) can arise since, once we integrate over Ξ , the corresponding predictive distribution takes a non-standard form and can accommodate features such as heavy tails or downside asymmetries as reported in, for example, Adrian et al. (2019). This is because (2) can be interpreted as a mixture of Gaussians with a highly complicated mixing distribution $p(\Xi | \text{Data}_{1:T})$ that takes no well known form either.

One additional advantage of our TVP specification relative to a standard QR is that the parameters vary over time. This enables us to investigate whether different elements in x_t vary in importance for explaining growth at risk over time.

3. Empirical findings

This section first demonstrates that our TVP-SV model outperforms quantile regression models in out-of-sample forecasting, underscoring its superior predictive capabilities. Subsequently, we shed light on the predictive densities over time and decompose both downside and upside risks to identify their driving factors, providing a comprehensive overview of how financial risk indicators have influenced tail risks to economic growth over the past 130 years. Finally, we conduct a comparative analysis of the two most significant economic crises in our sample: the Great Depression and the Great Recession. This comparison highlights the shifting drivers of downside risks before and during crises and is of particular interest for policymakers.

3.1. Model evaluation

In a first step, we evaluate the out-of-sample forecasting performance of different TVP and quantile regression models. For both the quantile regression and the TVP model, we estimate two specifications: (i) the baseline model including a constant, the financial stress indicator, and lagged GDP growth as predictors (we will use the abbreviation QR and TVP to refer to these models) and (ii) the extended model, which additionally includes credit growth and house price growth as regressors (hereafter referred to as QR+ and TVP+).

We evaluate the accuracy of the predicted downside risks with the help of quantile scores (see, e.g., Giacomini and Komunjer, 2005; Brownlees and Souza, 2021). Our focus is on the performance in the left tail. Hence, we consider the quantile score at the 5% quantile.

Table 1 reports the quantile scores relative to the QR model for the pre- and post-World War II (WWII) subsample and the forecasting horizons $h = 1$ and $h = 4$.⁸ In general, we find that TVP models perform particularly well during the post-WWII period, especially at the one-step ahead horizon. For both periods and horizons considered, the TVP models produce quantile scores that are superior to those obtained from the extended quantile regression. However, it is worth noting that the small QR model produces slightly more precise tail forecasts during the pre-WWII period.

Once we focus on four-quarter ahead GDP growth forecasts, this pattern becomes slightly less pronounced. In this case, both quantile regressions (QR and QR+) improve upon the TVP models

when the period prior to WWII is considered. When we focus on post-WWII data, the TVP regressions are again outperforming the simple quantile regression model by appreciable margins. The reason for this rather weak performance in the pre-WWII period is driven by the fact that the QRs produce predictive densities with wide credible intervals (see Figure A.7). This helps in periods characterized by sharp breaks in GDP growth but harms forecasting accuracy in tranquil times. Since the pre-WWII features several crises, a model which produces wide forecast intervals during that specific points in time yields favorable overall tail forecasts.

This brief discussion has shown that TVP models can outperform quantile regressions, especially at the one-quarter ahead horizon, but also for multistep ahead forecasts and post-WWII data. Since we focus on two variants of the models (i.e., the baseline and the extended versions), we can also analyze whether including credit and house price growth pays off for obtaining more precise tail forecasts. The results in Table 1 suggest that for one-quarter ahead tail forecasts, additional information does not translate into more precise predictions for the pre-WWII sample and only slightly improves the predictive fit over the post-WWII period. This pattern reverses if our interest is on four-step ahead forecasts. In that case, using more information yields more precise forecasts (which are still inferior to the QR benchmark predictions pre-WWII) for both holdout periods considered. This finding is most likely driven by the fact that one-quarter ahead predictions are dominated by high frequency shocks, which are notoriously difficult to predict, whereas for longer-run forecasts, short- and medium-lived trends become more important. Such short- and medium-lived trends include asset, credit, and house price cycles, as reported by Potjagailo and Wolters (2023), who analyze these cycles over a comparable historical sample.

In sum, our small forecasting exercise shows that TVP regressions are capable of producing competitive tail forecasts without explicitly focusing on the corresponding quantile under scrutiny. This indicates that the increased flexibility provided by TVP models (i.e. allowing for drifts in β_t and changing error variances in σ_t^2) enables us to adequately model the distribution of GDP growth over long samples. This conclusion is further supported by the additional model evaluations presented in Appendix B.

3.2. Features of the predictive densities

Next, we examine the characteristics of the predictive densities. Figure 1 presents the annualized one- and four-quarter ahead GDP growth along with their predicted lower (5th percentile) and upper (95th percentile) bounds. The results for the post-1970s period are consistent with previous research, which shows that lower bounds vary significantly over time while upper bounds are relatively stable (see, e.g., Adrian *et al.*, 2019; Aikman *et al.*, 2019). However, this pattern appears to have emerged only since the start of the Great Moderation period in the 1980s. Prior to that, upper bounds were just as volatile as lower bounds. Table 2 supports this observation by showing the standard deviation of up- and downside risks over four different periods. The table also highlights the significant reduction in overall tail volatility after WWII and the even lower variation in the tails since the Great Moderation. This is not surprising, as the Great Depression led to a change in policymaking, with a greater focus on stabilizing the business cycle and reducing volatility.

3.3. Main results: decomposing up- and downside risks

Next, we focus on the drivers of tail risks over time. Our TVP model with stochastic volatility is able to handle asymmetries in up- and downside risks, which is a necessary feature for this type of analysis. The asymmetries in tail risks imply asymmetries in the unconditional distributions but do not necessarily require asymmetries in conditional predictive distributions (Carriero *et al.*, 2024). Hence, only focusing on the time series evolution of the coefficients (see Figure A.3 and A.6) does not provide a complete picture of how tail risks occur. To shed light on which indicator drives the predictive quantiles of GDP growth, we rely on linear posterior summaries (see, e.g.,

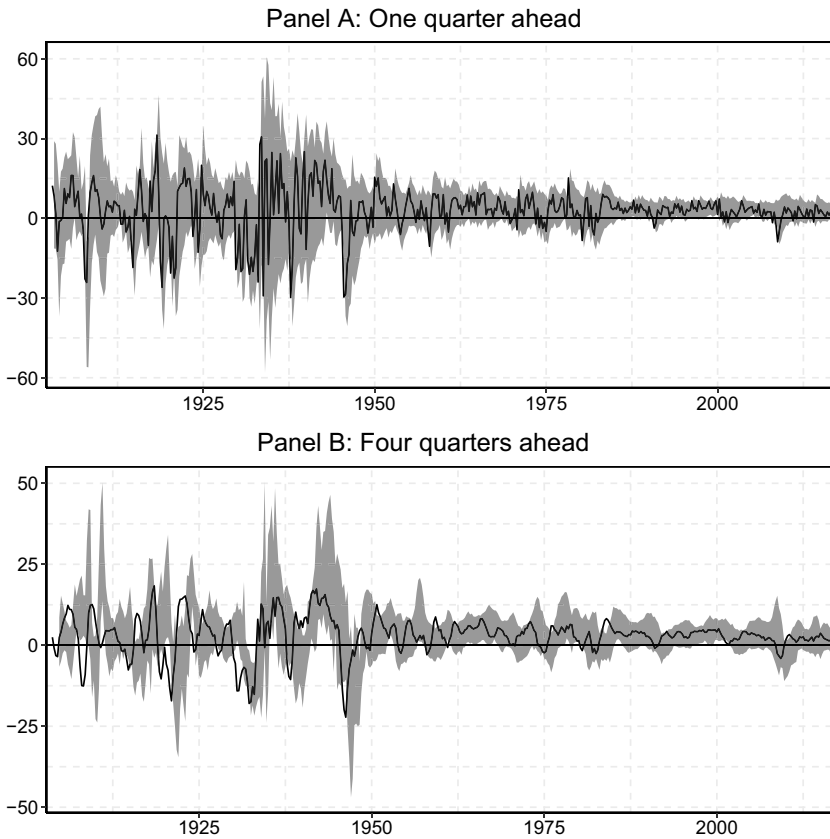


Figure 1. Time series evolution of the predicted tail risks.
 Note: This figure shows the out-of-sample one-quarter ahead (Panel A) and four-quarter ahead (Panel B) forecast of the 5th and 95th percentile together with the realized growth rate.

Woody et al., 2021). This is achieved as follows. Based on the h -step ahead predictive distribution of the model (see Eq. (2)) we, compute a sequence of quantiles $Q_{t+h,p}$. Each of these estimated quantiles is then used as the dependent variable in the following linear regression model:

$$Q_{t+h,p} = \alpha'_p x_t + e_t, \tag{3}$$

where α_p is a quantile-specific set of linear coefficients and e_t is a Gaussian shock with constant error variance. The ordinary least squares (OLS) estimator $\hat{\alpha}_p$ then provides a linear quantile-specific approximation to the predictive density of the flexible TVP regression. The key advantage of this approach is that it improves interpretability and allows for a straightforward decomposition of the driving forces of right and left-tail forecasts of GDP growth.

Figures 2 and 3 illustrate the decomposed tail risks in Panels A and B. The dashed black line shows the level of tail risks as predicted by our TVP model at the time of the prediction. The bars show the decomposition of the predicted 5th and 95th percentile approximated by the 10-year rolling linear regression. Panels C and D each show the linear posterior summary's corresponding coefficients to assess the risk indicators' marginal effect. A circle indicates whether a coefficient is significant at the 5% significance level. Subsequently, we do not discuss the tail risks around WWII because we cannot distinguish between the impact of macrofinancial variables and war time effects.

Table 2. Standard deviation of up- and downside risks

Horizon	Upside risks				Downside risks			
	Pre-WWI	Interwar	Pre-GM	Since GM	Pre-WWI	Interwar	Pre-GM	Since GM
$h = 1$	11.03	11.74	3.09	1.82	12.78	11.62	2.87	2.60
$h = 4$	11.96	11.66	3.62	2.22	7.47	9.29	3.47	2.63

Note: This table reports the standard deviation of the predicted 5th and 95th percentile. The sample is split into four periods: pre-WWI, the interwar era, the time before the Great Moderation (pre-GM), and since the Great Moderation (since GM). After both world wars, 2 years are left out to not capture any wartime effects in the standard deviations.

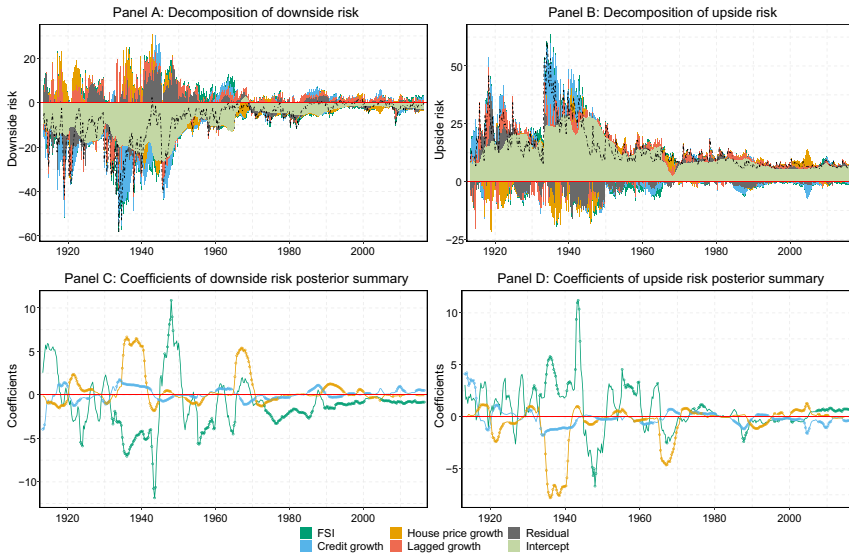


Figure 2. Decomposition of tail risks one-quarter ahead.

Note: Panels A and B show the predictive variable relevance for the predicted 5th and 95th percentile over time. We approximate the predicted 5th and 95th percentile using a 10-year rolling window linear regression model. Panels C and D show the coefficients of the linear approximation over time. A circle indicates a significant coefficient (5% significance level).

Subsequently, we present the results of this analysis by describing the stylized facts about the determinants of tail risks over the past 130 years, their implications for policymakers, and how they connect to the broader literature.

Conditional average. Starting with the intercept, we observe a staggering reduction of the intercept post-WWII (see Figures 2 and 3), which is the conditional average downside risk of the 10-year window of the rolling linear posterior summary. Therefore, the conditional average tail risks nowadays are significantly lower compared to the start of the sample. While this result is unsurprising since the overall tail volatility has substantially decreased since WW II, the sharp reduction in the conditional average growth risk is still eye-catching.

Financial stress. Adrian *et al.* (2022) identify financial stress as the primary short-term predictor of left tail risks. This is particularly evident during significant historical events such as the oil crises, the dot-com bubble, and the Great Recession, where financial stress markedly contributes to downside risks. However, during the Great Recession, credit and house price growth emerged as the primary factors influencing downside risks in the four-quarter ahead forecast (see Figure 3, Panel A). Similarly, prior to WWII, these variables were crucial in forecasting growth at risk,

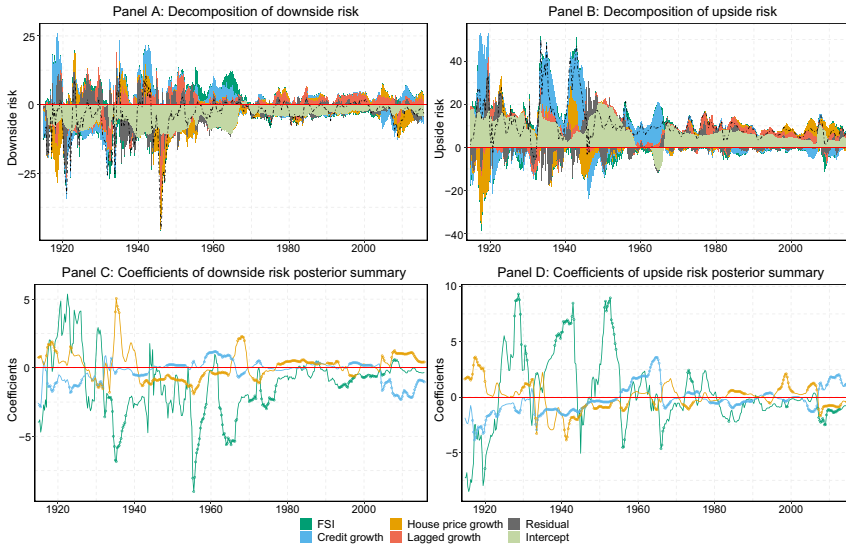


Figure 3. Decomposition of tail risks four quarter ahead.

Note: Panels A and B show the predictive variable relevance for the predicted 5th and 95th percentile over time. We approximate the predicted 5th and 95th percentile using a 10-year rolling window linear regression model. Panels C and D show the coefficients of the linear approximation over time. A circle indicates a significant coefficient (5% significance level).

notably during the early stages of the Great Depression. In fact, during the depths of the Great Depression, credit growth was the predominant contributor to downside risks.

Interestingly, while financial stress significantly affects the left tail of the predicted distribution, it has a negligible impact on the upper tail. This asymmetry is evident in the coefficients, which show a significant negative effect on left-tail risks but a less pronounced effect on upside risks, as depicted in Figures 2 and 3, Panels C and D. Specifically, financial stress exhibits a significant negative marginal effect on left tail risks (Panel C), whereas its marginal effect on upside risks is both economically and statistically insignificant.⁹

Given these findings, policymakers should prioritize mitigating financial stress, as it has exacerbated growth risks for over a century without offering corresponding benefits to the upper tail of the distribution.

Credit growth. Credit growth typically affects downside risks negatively in the medium and long term but less so at shorter time horizons (Adrian et al., 2022; Galán, 2024). While our findings broadly confirm this, we are also able to shed light on temporary and transitional effects over time. In our one-quarter ahead predictions, we find a negative impact of credit growth both after the Global Financial Crisis and the first oil shock, but for completely different reasons. While rapid credit expansion drives downside risks during the first oil crisis, a credit crunch—that is, too little credit—increased downside risks following the Global Financial Crisis (see Figure 2, Panel C). This stands in stark contrast to the effect of credit growth in the four-quarter ahead predictions, where the sharp decline in the coefficient in the mid-2000s results in a large negative overall effect.¹⁰

Before WWII, credit growth was the most important driver of tail risks, contributing to the post-World War I recession, the Depression of 1920–1921, and the deterioration in growth risks during the Great Depression. Again, when looking at Panel C in Figures 2 and 3, we see the importance of using a flexible approach, as the coefficients vary substantially over time both in terms of sign and magnitude.

The stark difference between the Great Recession and the Great Depression lies in the duration and magnitude of the credit crunch, highlighting the effectiveness of central banks’ interventions

in 2008 and the following years. Notably, while credit growth did not significantly increase downside risks in the latter half of the 20th century, it did so in the lead-up to the Global Financial Crisis, comparable to the Great Depression in both absolute and marginal effects. These findings align with the seminal work of Jordà *et al.* (2015), which shows that credit boom-fueled stock bubbles are much costlier compared to noncredit boom bubbles. Additionally, Gorton and Ordóñez (2020) find that the average credit boom in a sample of 34 countries from 1960 to 2010 lasts 11 years and is not necessarily bad. Our analysis supports this perspective; for instance, the credit boom after WWII had a very desirable outcome by decreasing downside risks and increasing upside risks.

While financial stress has only a minimal impact on upside risks, credit growth significantly influences both upper and lower tails, depending on the respective time period. Therefore, policymakers should consider the role of credit and its impact on downside risks more nuancedly compared to financial stress.

House price growth. The effect of financial stress and credit growth on downside risks to economic growth since the 1970s aligns well with previous literature. For house price growth, prior studies suggest little to no effect on downside risks (Aikman *et al.*, 2019; Galán, 2024). However, our findings reveal a significant impact on short-term tail risks during the 1950s, 1960s, early 1990s, and 2000s. Figures 2 and 3, Panel C, illustrate these patterns. Fast-rising house prices in the 1950s, partly due to the GI Bill, and increasing subprime lending in the 2000s heightened economic growth risks by threatening a real estate bubble. Conversely, declining house prices in the early 1990s and late 1960s also negatively impacted downside risks. The most pronounced negative effect is observed in the mid-2000s four-quarter ahead prediction.

These findings contrast with previous studies that suggest house prices have little to no effect on short-term tail risks, highlighting our empirical model's ability to capture time-varying effects. Notably, house price growth also impacts the upside risks of economic growth, consistent with the boom-bust pattern in financial cycle literature (Borio, 2014).

In the pre-WWII period, the effect of house price growth on tail risks was relatively small, except for the house price crash after WWI. Our results align with literature on housing and financial cycles, which concludes that over the last century, increased mortgage lending and total credit to households have amplified feedback effects on the economy, influencing tail risks to economic growth. Since the 1970s, we observe an almost exclusively positive marginal effect of house prices on growth risks. This trend, linked to a rising home ownership ratio, suggests an increasing wealth effect on consumption when house prices collapse (Mian *et al.*, 2013; Jordà *et al.*, 2017; Graham and Makridis, 2023).

The correlation between house prices and credit growth has also increased significantly since the 1960s and has remained high since then. Consequently, house prices have become more relevant from a financial stability perspective, as financial imbalances have likely increased with rising leverage in recent decades compared to the pre-WWII period (see, for instance, Jordà *et al.*, 2015, 2016; Mian *et al.*, 2017). This mortgage-financed housing boom may explain why the effect of house price growth on tail risks was unprecedented during the Global Financial Crisis, aligning with broader financial cycle literature (Borio, 2014).

3.4. Common patterns in times of crises

To gain an in-depth understanding of our empirical results and go beyond the broader picture presented in the previous section, it is useful to explain the findings by referencing the two major economic crises in our sample. For this purpose, we illustrate the results on the basis of TVP local projections during the Great Depression on the one hand and the Global Financial Crisis on the other. These TVP local projections are derived from the posterior summary, as described in Subsection 3.3 and equation 3. Essentially, they provide a different representation of the coefficients shown in Figures 2 and 3. Hence, Figure 4 shows the effect of financial stress, credit growth,

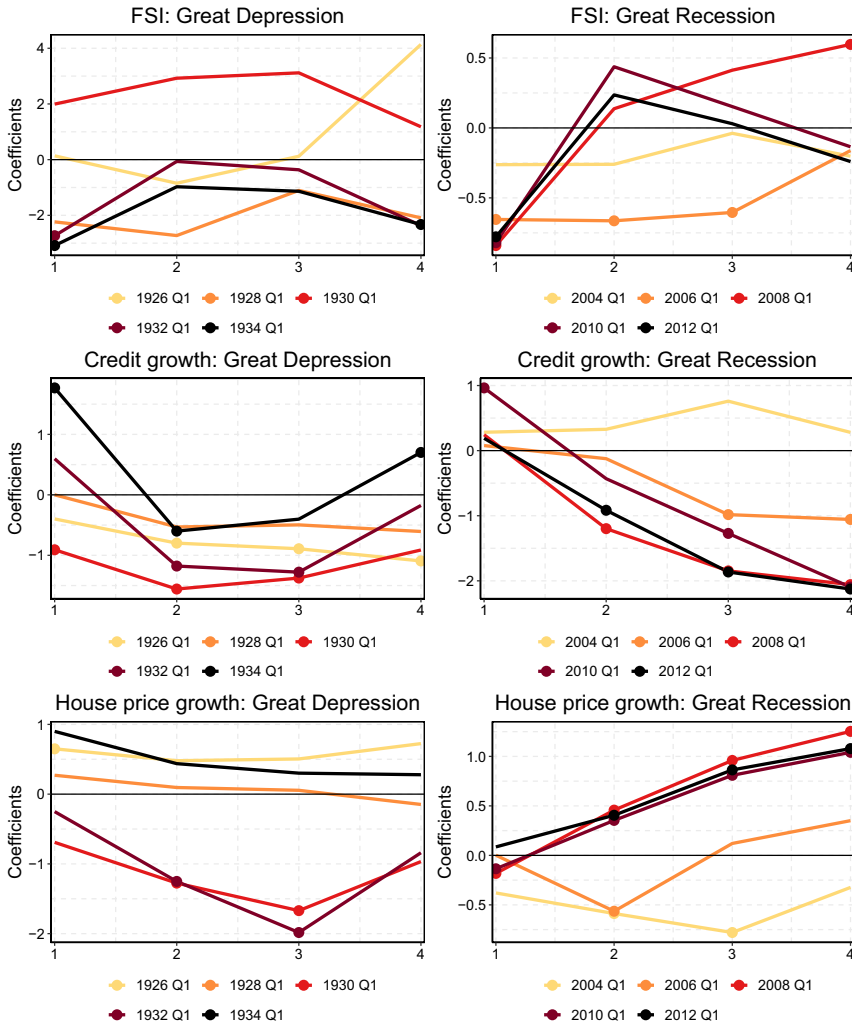


Figure 4. TVP local projections.

Note: This figure shows TVP local projections of the linear approximation at different time points. A circle indicates a significant coefficient (5% significance level).

and house price growth on the predicted 5th percentile of future GDP growth one to four quarters ahead at different points in time.

The results show significant fluctuations in coefficients during both crises, with an even more pronounced effect observed during the latter crisis. Interestingly, there are both similarities and differences in the impact of financial risk indicators on downside risks during the two crisis periods, with particularly strong differences in the effect of house prices.

Both during the Great Depression and the Global Financial Crisis, financial stress amplified downside growth risks at short time horizons, in the case of the Great Depression partly also at longer time horizons.¹¹ In contrast, the effect is somewhat different in the years preceding the crisis, as pre-crisis growth risks are only affected by financial stress prior to the Global Financial Crisis. In both cases, as the crisis unfolded, the effect on longer term horizons disappeared and the impact on short-term (one-quarter ahead) predictions intensified. Thus, financial stress elevated

downside risks to economic growth in both crises; however, it only influenced pre-crisis growth risks before the Global Financial Crisis.

For credit growth, we observe comparable patterns in both crises, with higher credit growth being associated with heightened growth risks both prior to and during the crisis. During the Great Depression, the coefficients for all horizons shifted upward, suggesting that a credit crunch, rather than excessive credit, heightened growth risks in later years. On the contrary, this upward shift in coefficients was only observed for the one-quarter ahead horizons during the Global Financial Crisis, potentially indicating the success of central bank intervention in avoiding a credit crunch.

While the results are comparable for financial stress and credit growth, the relationship between house price growth and downside risks shows strong differences between the two crises. In the Great Depression, rising house prices had a negative impact on growth risks, whereas during the Global Financial Crisis, falling house prices had the same effect. These results for the Global Financial Crisis are well in line with recent findings on housing wealth effects (e.g., Mian *et al.*, 2013). The literature finds that a decline in house prices can also reduce consumption, especially when housing wealth constitutes a large proportion of households' overall wealth.

The fact that the relationship between tail risks and financial risk indicators varies not only over long time horizons but also during crises has significant implications for policymakers. When calibrating models for optimal macroprudential policy, such as the one proposed by Suarez (2022), using non-TVPs (e.g., quantile regression estimates) can lead to misleading policy recommendations, especially during financial crises. Additionally, the effects of both risk indicators and macroprudential policies change over time (see, e.g., Jiménez *et al.*, 2017; Cerutti *et al.*, 2017). Therefore, instead of attempting to model all interactions between risk indicators (and possibly macroprudential policies),¹² it is more effective to use a flexible empirical approach like the one proposed in this paper. This approach provides more reliable warning signals about potential sharp recessions and facilitates a more efficient use of macroprudential policies.

4. Conclusion

Predicting growth risks has become a critical aspect of policymaking since the Global Financial Crisis, and it is particularly important for macroprudential policymakers. While previous research has identified financial stress, credit growth, and house price growth as key indicators for predicting growth risks, these studies often rely on short time series and inflexible models. In contrast, our research uses a flexible empirical approach and a 130-year historical time series to improve forecasting accuracy and to gain a better understanding with respect to the changing relationships between tail risks and financial risk indicators over time.

Our findings suggest that certain risk indicators, such as credit growth, have a time-varying effect on tail risks but tend to follow similar patterns during financial crises. On the other hand, the impact of house prices on tail risks is time-varying for both upside and downside risks, and its impact on growth risks during the Global Financial Crisis was also exceptional compared to historical trends. Additionally, we find that financial stress is consistently associated with increased growth risks (i.e., lower values of GaR) over the entire sample period.

This paper also serves as a note of caution in drawing policy implications from short time series or single historical events, as the observed effects of risk indicators may be time dependent. In addition to financial risk indicators, policy variables may also have a time-varying impact on tail risks throughout the credit and business cycle. Further research is necessary to assess the effectiveness and importance of macroprudential policy in crisis situations and the potential trade-offs that may occur during normal times.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/S1365100524000476>.

Notes

- 1 One notable exception that relies on copula models is Coe and Vahey (2020). As opposed to this paper, they focus on the role of financial conditions for tail risks to economic growth.
- 2 It is important to note that for the QR model to adjust to the structural break occurring in the data after the Second World War, we need to split the sample. This is not necessary for our TVP-SV model. Furthermore, if such a structural break or a significant shift in variable relationships were to appear at the recent end of the sample, a sample split would not be feasible, and the QR model could not adapt. However, this poses no problem for the TVP-SV model, which can naturally adjust to such changes without requiring a sample split.
- 3 Clark et al. (2024) use a similar approach to summarize the effect of different predictors on inflation within a nonparametric model.
- 4 D'Agostino and Surico (2012) construct a time series for quarterly real GDP for pre-WWII by using the real GNP time series of Balke and Gordon (1989).
- 5 Taking the average growth rate of house price and credit growth over several years is standard in the GaR literature (see, e.g., Aikman et al. (2019), Galán (2024), Gächter et al. (2023)). By focusing on the 3-year average growth rate, the intention is to measure a buildup of financial imbalances and vulnerabilities over time, similar to Borio and Lowe (2002). This approach also mitigates some concerns regarding data interpolation. Notably, our findings remain consistent even when using annual credit data for every quarter of the year rather than interpolated quarterly data.
- 6 We use an AR(1) process since, as we detail below, the kurtosis is a function of the unconditional distribution of the log-volatility process. If we assume random walk dynamics, this unconditional distribution is not well defined. Moreover, in theory, assuming a random walk law of motion would imply that the volatility process hits $\pm\infty$ with probability 1 if $T \rightarrow \infty$. We feel that an AR(1) specification is thus more suitable given the data set we aim to analyze.
- 7 It is noteworthy that papers such as Kowal et al. (2019), Prüser (2021), and Hauzenberger et al. (2024) assume time-varying shrinkage factors. These models, however, are computationally more demanding, and the triple gamma prior produces posterior distributions of v_j^2 , which are heavy tailed and hence is also capable of covering situations where parameters change abruptly.
- 8 We need to split the sample because the quantile regression models cannot adequately handle the large drop in GDP volatility post-WWII. Furthermore, Amir-Ahmadi et al. (2016) show that correlations between variables, forecasts, and other statistics change considerably at certain points in time when using long time series, making an a priori choice of subsamples hard to defend. Moreover, Gächter et al. (2023) show that the average downside risks and the magnitude of effects of financial risk indicators depend on the structural characteristics of a country. These country characteristics change over time and therefore should be taken into account when not using a TVP approach.
- 9 Although the OLS coefficient for the one-quarter ahead 95th percentile posterior summary has been consistently significantly positive since the Global Financial Crisis, the magnitude of this effect is minimal.
- 10 This once again confirms the advantages of our flexible empirical approach. In a similar vein, Aikman et al. (2019) show, in their 3-year ahead quantile regression GaR estimation, that the sign of the credit growth coefficient depends on the respective subsample. For example, using subsamples until 1997 or 2002, they find a positive effect of credit growth on the 3-year ahead GaR. Only when extending the subsample to 2007, do they find a negative effect of credit growth. Their—at first view surprising—results are therefore consistent with our findings.
- 11 Please note that the illustrated positive coefficients are not statistically significant, except for 2008Q1, when the positive coefficient four quarters ahead may reflect a partial reversion of the strongly negative effect for the one-quarter ahead tail forecast.
- 12 For example, whether bubbles are leveraged or not, see Jordà et al. (2015).

References

- Adrian, T., F. Grinberg, N. Liang, S. Malik and J. Yu (2022) The term structure of growth-at-risk. *American Economic Journal: Macroeconomics* 14(3), 283–323.
- Adrian, T., N. Boyarchenko and D. Giannone (2019) Vulnerable growth. *American Economic Review* 109(4), 1263–1289.
- Aikman, D., J. Bridges, S. Hacioglu Hoke, C. O'Neill and A. Raja (2019) Credit, capital and crises: A gdp-at-risk approach.
- Amir-Ahmadi, P., C. Matthes and M.-C. Wang (2016) Drifts and volatilities under measurement error: Assessing monetary policy shocks over the last century. *Quantitative Economics* 7(2), 591–611.
- Balke, N. S. and R. J. Gordon (1989) The estimation of prewar gross national product: Methodology and new evidence. *Journal of Political Economy* 97(1), 38–92.
- Belmonte, M. A. G., G. Koop and D. Korobilis (2014) Hierarchical shrinkage in time-varying parameter models. *Journal of Forecasting* 33(1), 80–94.
- Bitto, A. and S. Frühwirth-Schnatter (2019) Achieving shrinkage in a time-varying parameter model framework. *Journal of Econometrics* 210(1), 75–97.
- Borio, C. (2014) The financial cycle and macroeconomics: What have we learnt? *Journal of Banking & Finance* 45, 182–198.

- Borio, C. and P. Lowe (2002) Assessing the risk of banking crises. *BIS Quarterly Review* 7(1), 43–54.
- Brownlees, C. and A. B. M. Souza (2021) Backtesting global growth-at-risk. *Journal of Monetary Economics* 118, 312–330.
- Cadonna, A., S. Frühwirth-Schnatter and P. Knaus (2020) Triple the gamma-a unifying shrinkage prior for variance and variable selection in sparse state space and TVP models. *Econometrics* 8(2), 20.
- Carnero, M. A., D. Peña and E. Ruiz (2004) Persistence and kurtosis in GARCH and stochastic volatility models. *Journal of Financial Econometrics* 2(2), 319–342.
- Carrero, A., T. E. Clark and M. Marcellino (2024) Capturing macro-economic tail risks with bayesian vector autoregressions. *Journal of Money Credit and Banking* 56(5), 1099–1127.
- Cerutti, E., S. Claessens and L. Laeven (2017) The use and effectiveness of macroprudential policies: New evidence. *Journal of Financial Stability* 28, 203–224.
- Clark, T. E., F. Huber, G. Koop and M. Marcellino (2024) Forecasting us inflation using bayesian nonparametric models. *The Annals of Applied Statistics* 18(2), 1421–1444.
- Coe, P. J. and S. P. Vahey (2020) Financial conditions and the risks to economic growth in the united states since 1875.
- Cogley, T. and T. J. Sargent (2005) Drifts and volatilities: Monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8(2), 262–302.
- Crawford, L., S. R. Flaxman, D. E. Runcie and M. West (2019) Variable prioritization in nonlinear black box methods: A genetic association case study. *The Annals of Applied Statistics* 13(2), 958.
- D’Agostino, A. and P. Surico (2012) A century of inflation forecasts. *Review of Economics and Statistics* 94(4), 1097–1106.
- Forsythe, G. E., M. A. Malcolm and C. B. Moler (1977) Computer methods for mathematical computations.
- Gächter, M., M. Geiger and E. Hasler (2023) On the structural determinants of growth-at-risk. *International Journal of Central Banking* 19(2), 251–293.
- Galán, J. E. (2024) The benefits are at the tail: Uncovering the impact of macroprudential policy on growth-at-risk. *Journal of Financial Stability* 74, 100831.
- Giacomini, R. and I. Komunjer (2005) Evaluation and combination of conditional quantile forecasts. *Journal of Business & Economic Statistics* 23(4), 416–431.
- Gorton, G. and G. Ordonez (2020) Good booms, bad booms. *Journal of the European Economic Association* 18(2), 618–665.
- Graham, J. and C. A. Makridis (2023) House prices and consumption: A new instrumental variables approach. *American Economic Journal: Macroeconomics* 15(1), 411–443.
- Hauzenberger, N., F. Huber and G. Koop (2024) Dynamic shrinkage priors for large time-varying parameter regressions using scalable markov chain monte carlo methods. *Studies in Nonlinear Dynamics & Econometrics* 28(2), 201–225.
- Huber, F., G. Koop and L. Onorante (2021) Inducing sparsity and shrinkage in time-varying parameter models. *Journal of Business & Economic Statistics* 39(3), 669–683.
- Jiménez, G., S. Ongena, J.-L. Peydró and J. Saurina (2017) Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the spanish dynamic provisioning experiments. *Journal of Political Economy* 125(6), 2126–2177.
- Jordà, Ò., M. Schularick and A. M. Taylor (2015) Leveraged bubbles. *Journal of Monetary Economics* 76, S1–S20.
- Jordà, Ò., M. Schularick and A. M. Taylor (2016) The great mortgaging: Housing finance, crises and business cycles. *Economic policy* 31(85), 107–152.
- Jordà, Ò., M. Schularick and A. M. Taylor (2017) Macrofinancial history and the new business cycle facts. *NBER macroeconomics annual* 31(1), 213–263.
- Knaus, P., A. Bitto-Nemling, A. Cadonna and S. Frühwirth-Schnatter (2021) Shrinkage in the time-varying parameter model framework using the R package shrinkTVP. *Journal of Statistical Software* 100(13), 1–32.
- Kowal, D. R., D. S. Matteson and D. Ruppert (2019) Dynamic shrinkage processes. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 81(4), 781–804.
- Mian, A., A. Sufi and E. Verner (2017) Household debt and business cycles worldwide. *The Quarterly Journal of Economics* 132(4), 1755–1817.
- Mian, A., K. Rao and A. Sufi (2013) Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics* 128(4), 1687–1726.
- Polson, N. G. and J. G. Scott (2010) Shrink globally, act locally: Sparse bayesian regularization and prediction. *Bayesian Statistics* 9(501-538), 105.
- Potjagailo, G. and M. H. Wolters (2023) Global financial cycles since 1880. *Journal of International Money and Finance* 131, 102801.
- Primiceri, G. E. (2005) Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72(3), 821–852.
- Prüser, J. (2021) The horseshoe prior for time-varying parameter VARs and monetary policy. *Journal of Economic Dynamics and Control* 129, 104188.
- L., Püttmann (2018) Patterns of panic: Financial crisis language in historical newspapers. Available at SSRN 3156287.
- M., Schularick and A. M. Taylor (2012) Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870–2008. *American Economic Review* 102(2), 1029–1061.

- Suarez, J. (2022) Growth-at-risk and macroprudential policy design. *Journal of Financial Stability* 60, 101008.
- Woody, S., C. M. Carvalho and J. S. Murray (2021) Model interpretation through lower-dimensional posterior summarization. *Journal of Computational and Graphical Statistics* 30(1), 144–161.