

**Time-Varying Dynamics of
the German Business Cycle:
A Comprehensive
Investigation**

Magnus Reif

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Time-Varying Dynamics of the German Business Cycle: A Comprehensive Investigation

Abstract

This paper provides insights into the time-varying dynamics of the German business cycle over the last five decades. To do so, I employ an open-economy time-varying parameter VAR with stochastic volatility, which I estimate by quasi-Bayesian techniques. The reduced-form analysis reveals substantial shifts in the variables' long-run growth rates and shock volatilities over time. German trend inflation has strongly decreased and settled at a historically low level. GDP growth volatility exhibits marked fluctuations over time and has dropped to historically low levels only after the global financial crisis. The structural analysis employs externally identified oil supply shocks along with a recursive identification scheme to identify key macroeconomic shocks. The analysis reveals strong fluctuations in both the impact responses of macroeconomic aggregates to these shocks and the shock propagation processes. Thus, I conclude that business cycle stabilization in Germany is driven by both good policy and good luck.

JEL-Codes: E310, E320, E520, E580.

Keywords: time-varying parameters, Bayesian vector autoregression, counterfactuals, stochastic volatility, Great Moderation.

Magnus Reif
Deutsche Bundesbank
Frankfurt / Germany
magnus.reif@bundesbank.de

†I am thankful to Jens Boysen-Hogrefe, Johannes Hoffmann, Robert Lehmann, Svetlana Rujin, Maik Wolters, and Timo Wollmershäuser for helpful comments and suggestions. The views expressed in this paper are those of the author and do not necessarily coincide with the views of the Deutsche Bundesbank or the Eurosystem.

1 Introduction

During the last five decades, the German economy was subject to enormous structural changes. It consummated its reunification, integrated into the global and in particular European economy, and transferred its monetary authority from the Bundesbank to the European Central Bank (ECB). These changes not only came along with considerable shifts in the composition of German GDP over time.¹ The literature also documents a decline in German output growth volatility. However, the timing, the extent, and the sources of these structural changes are not beyond dispute.

This paper provides a comprehensive analysis of this issue and contributes to the literature along two dimensions. First, I jointly model domestic forces along with a global activity proxy using a time-varying parameter VAR with stochastic volatility (TVP-SV-VAR), which I estimate using quasi-Bayesian methods. By this means, I can model rich dynamics while being agnostic about the degree and nature of parameter instability.² Second, I conduct a structural analysis and provide novel evidence on the question whether structural changes are rather driven by smaller shocks over time (good luck) or by changes in the systematic response of the economy to these shocks (good policy).

The reduced-form analysis studies how the structural transformations affected the time series properties of the German economy. Specifically, I examine whether the variables' trends and volatilities are time-varying and document that both statistics are subject to substantial shifts. In particular, the trends of the interest and inflation rates and trend output growth exhibit a strong decline over time. Moreover, I find a strong reduction in business cycle volatility, which appears to be rather gradual than discrete. The latter can be traced back to strong drops in the magnitude of reduced-form shocks hitting inflation, wage growth, and the exchange rate change. The magnitude of shocks to output growth displays marked fluctuations over time and drops to historically low levels only after the global financial crisis. Additionally, counterfactual calculations suggest that smaller shocks are not the sole driver of business cycle stabilization in Germany and that parts of the drop in volatility can be attributed to changes in the response of the German economy to shocks.

The structural analysis provides an economic interpretation of the changes in the time series properties. I use an externally identified oil supply shock series along with a recursive identification scheme to identify key macroeconomic shocks and investigate how their propagation evolves over time. I find strong time-variation in the effects of these shocks.

¹On the expenditure side, the share of exports increased from about 20% in 1970 to more than 40% in 2019. On the production side, the share of the service sector increased from roughly 50% in 1970 to 70% in 2019, whereas the share of manufacturing dropped from 37% to 21% in the same period.

²Recently, Markov Chain Monte Carlo (MCMC) algorithms that can process large dataset and sample from the exact posterior were proposed. A brief overview about these developments is provided in Section 2.1.

In particular, the responsiveness of the German economy to sudden disruptions in the oil market has substantially declined; at the end of the sample, the German economy is virtually non-responsive to oil supply shocks. Moreover, I provide evidence for changes in the shocks' propagation, suggesting that the German economy has improved absorbing exogenous disturbances. Finally, counterfactual simulations indicate that shocks originating from abroad do only account for a minor part of the reduction in business cycle volatility, while domestic shocks and oil supply shocks play a more important role.

I estimate the model using the quasi-Bayesian local likelihood (QBLL) approach introduced by Petrova (2019). Compared to the commonly applied state-space approach of Cogley and Sargent (2005) and Primiceri (2005), the QBLL method features three main advantages, which makes it particularly suited for the analysis at hand: (i) it models the coefficients' time-variation nonparametrically and hence, flexibly allows for a suite of possible parameter processes, including the popular random walk process; (ii) it provides closed-form expressions for the quasi-posterior distributions of the VAR coefficients, which, on the one hand, substantially reduces computational burden compared to the state-space approach. On the other hand, it admits to directly draw the time-varying covariance matrix without decomposing it using some kind of triangularization. Thus, according to the QBLL approach, the ordering of the variables does not influence the posterior distribution; (iii) it allows to impose Bayesian shrinkage directly on the VAR coefficients and hence facilitates estimating large systems.

This paper contributes to the literature on structural change in the German economy, particularly regarding the decline in business cycle volatility. Stock and Watson (2005) document a near-monotonic decline of GDP growth volatility since the 1960s, driven by a decrease of the residual variances. This gradual decline is confirmed by Fritsche and Kuzin (2005), who attribute it, however, to an increasing persistence of the GDP growth process caused by a change in the conduct of monetary policy. In contrast, Buch, Doepke, and Pierdzioch (2004) and Aßmann, Högrefe, and Liesenfeld (2009) present evidence in favor of a discrete transition to a lower volatility state in the early 1990s. While Buch et al. (2004) also ascribe the declining volatility to a change of monetary policy, Aßmann et al. (2009) highlight the importance of shifts in the composition of GDP. Finally, Mills and Wang (2003) and Summers (2005) find a single structural break in the residual variances of the growth process taking place already in the mid-1970s.

Compared to the studies highlighted above, which are either based on small-scale linear multivariate models, univariate models with discrete breaks, or univariate models with gradual parameter change, this paper employs the TVP-SV-VAR that has three advantages. First, the researcher can refrain from taking any stance on whether there is abrupt, gradual, or no structural change at all. Second, the TVP-SV-VAR allows for both drifting VAR coefficients and drifting volatilities. Thus, it can capture time-variation in

the high- and low-frequency domain of the variables considered. Third, the TVP-SV-VAR allows to simultaneously identify structural shocks and their evolution over time.

The remainder of this paper is organized as follows. Section 2 outlines the model. Sections 3 and 4 present the results from the reduced-form and the structural analysis, respectively. Section 5 concludes.

2 Empirical methodology

2.1 Modeling time-variation

Since the seminal contribution of Primiceri (2005), the state-space approach for TVP-SV-VARs has been successfully applied for both structural analysis and forecasting macroeconomic time series (see, among others, Galí and Gambetti, 2009; D’Agostino, Gambetti, and Giannone, 2013; Mumtaz and Zanetti, 2015). However, due to the high computational burden and the risk of overfitting the data, these studies are usually confined to small-scale systems. To overcome the “curse of dimensionality”, several modifications to the original approach have been proposed. Korobilis (2013) suggests to extract a low-dimensional vector of latent factors from a large dataset and to use these factors as additional regressors in small-scale VARs. While this factor structure allows to process more information, it is restrictive in how the additional information is linked to the dependent variables. More recently, MCMC algorithms have been developed, which allow to abstract from this factor structure. Bitto and Frühwirth-Schnatter (2019) use shrinkage priors to estimate whether parameters are time-varying or constant and by this means avoid overfitting. Huber, Koop, and Onorante (2020) add another layer by first shrinking the parameters to zero and second sparsifying the draw to control for model uncertainty. While these models are able to process large datasets without restricting linkages between the variables, they are somewhat restrictive on how the parameters evolve over time. In fact, these approaches either assume a random walk or a stationary autoregressive process. Moreover, tightly parametrized inverse Gamma priors that favor gradual change are usually employed.³

However, from a theoretical point of view, structural change can induce both discrete breaks in the data generating process and gradual parameter change. On the one hand, given that agents do not update their beliefs at the same time, aggregation across agents would rather lead to gradual changes. On the other hand, abrupt and rapid policy shifts might lead to discrete structural change. The key advantage of the QBLL methodology of Petrova (2019) is that it flexibly allows for both kinds of parameter instability. Thus, the

³To circumvent this issue, Prüser (2020) suggests to use horseshoe priors.

QBLL method is less prone to misspecification and allows the researcher to stay agnostic about the precise law of motion for the coefficients.

2.2 The model

To investigate how the structural and institutional changes affected the German economy since the early 1970s, I employ an open-economy TVP-SV-VAR and estimate it using the QBLL methodology introduced by Petrova (2019). The model reads as follows:

$$y_t = c_t + \sum_{i=1}^p B_{i,t} y_{t-i} + u_t, \quad u_t = R_t^{-1/2} \varepsilon_t, \quad \varepsilon_t \sim N(0, I_N), \quad (1)$$

where the $N \times 1$ vector c_t and the $N \times N$ matrices $B_{i,t}$ contain the time-varying intercept terms and the autoregressive coefficients, respectively. $R_t^{-1/2}$ is the positive definite time-varying covariance matrix. p denotes the lag length. The vector y_t consists of nine quarterly variables:

$$y_t = \left[os_t \quad gip_t \quad gdp_t \quad p_t \quad w_t \quad h_t \quad reer_t \quad i_t \quad rac_t \right]', \quad (2)$$

where os_t denotes an oil supply shock series (for a detailed description see Section 4.1), gip_t is a proxy for global industrial production growth, gdp_t is real GDP growth, p_t is inflation measured in terms of CPI growth, w_t is real wage growth, h_t is hours growth, $reer_t$ is the real effective exchange rate change, and rac_t refers to real oil price inflation in terms of the change of US refiners' acquisition costs. The sample covers the period from 1975:Q1 to 2019:Q4. Variables, which are sampled at a higher frequency than quarterly, enter the model as quarterly averages.⁴ To account for the effective lower bound (ELB), I estimate two model specifications: Model A includes the actual short-term interest rate as policy instrument, model B uses the shadow rate for the euro area to capture additional features of monetary policy that do not directly affect the actual short-term interest rate. A detailed description of the dataset is provided in Appendix A.

Defining $x_t = (1, y'_{t-1}, \dots, y'_{t-p})$ and $B_t = (c_t, B_{1,t}, \dots, B_{p,t})$ allows to write the model as:

$$y_t = (I_N \otimes x_t) \beta_t + R_t^{-1/2} \varepsilon_t, \quad (3)$$

⁴In the case of the oil supply shock series this amounts to averaging monthly shocks, which introduces serial correlation. However, the empirical correlation is so low that one can safely treat them as being uncorrelated. The latter is also reflected by the fact that the posterior means (or medians) of the coefficients of the first equation are virtually zero.

where β_t contains the VAR coefficients stacked in a vector. The QBLL methodology estimates the time-varying parameters nonparametrically.⁵ It is based on the following conditions for the vector of time-varying parameters $\theta_t = \left[\beta_t \text{ vech}(R_t^{-1/2}) \right]'$:

- (i) θ_t is a deterministic process $\theta_t = \theta(t|T)$, where $\theta(\cdot)$ is a piecewise differentiable function.
- (ii) θ_t is a stochastic process satisfying: $\sup_{j:|j-t|<h} \|\theta_t - \theta_j\|^2 = O_p(h|t)$ for $1 \leq h \leq t$ for $t \rightarrow \infty$.

If θ_t satisfies either (i) or (ii), the sequence of time-varying parameters moves slowly over time, which is a sufficient property for consistent estimation. Moreover, these conditions allow for a wide array of possible parameter processes, including the random walk (Petrova, 2019). Given (i) and (ii), the local likelihood function of model (1) for each period j is given by:

$$\varphi_{T,j}(\theta_j) = \sum_{t=1}^T \vartheta_{j,t} l_t(y_t | y^{t-1}, \theta_j), \quad \text{for } j, t \in \{1, \dots, T\}, \quad (4)$$

where $l_t(y_t | y^{t-1}, \theta_j)$ is the conditional log-density for observation t and $\vartheta_{j,t}$ reweighs the likelihood of the observations (y_1, \dots, y_T) . For $j, t \in \{1, \dots, T\}$, these weights are computed using a kernel function:

$$\vartheta_{j,t} = \varkappa_{j,t} \omega_{j,t}, \quad \omega_{j,t} = \tilde{\omega}_{j,t} \sum_{t=T}^T \tilde{\omega}_{j,t}, \quad \tilde{\omega}_{j,t} = K\left(\frac{j-t}{H}\right), \quad \varkappa_{j,t} = \left(\sum_{t=1}^T \omega_{j,t}^2\right)^{-1} \quad (5)$$

$K(\cdot)$ is a non-negative, continuous, and bounded kernel function with bandwidth parameter H , satisfying $H \rightarrow \infty$ and $H = o(T/\log T)$.⁶ This kernel function reweighs the model's log-likelihood function at each period j such that observations close to j receive a high weight whereas distant observations receive a low weight. The rate of decay of the weights is govern by the bandwidth parameter H . The higher the value for H is set, the slower the weights decay and the smoother the estimates become. Combining the local likelihood function with a Normal-Wishart prior for β_j and R_j :

$$p(\beta_j | R_j) \sim N(\underline{\beta}_j, (R_j \otimes \underline{\kappa}_j)^{-1}), \quad p(R_j) \sim W(\underline{\alpha}_j, \underline{\gamma}_j^{-1}), \quad \text{for } j \in \{1, \dots, T\}, \quad (6)$$

⁵The QBLL approach extends the frequentist nonparametric approach of Giraitis, Kapetanios, and Yates (2014) to a Bayesian framework. For a detailed description see Petrova (2019).

⁶In the case of a Normal kernel, the weights are given by: $\omega_{j,t} = \frac{1}{\sqrt{2\phi}} \exp((\frac{-1}{2})(j-t)/H^2)$ for $j, t \in \{1, \dots, T\}$.

gives rise to a Normal-Wishart quasi-posterior for β_j and R_j :

$$p(\beta_j|R_j, X, Y) \sim N(\bar{\beta}_j, (R_j \otimes \bar{\kappa})^{-1}), \quad p(R_j|Y, X) \sim W(\bar{\alpha}_j, \bar{\gamma}_j^{-1}), \quad \text{for } j \in \{1, \dots, T\}, \quad (7)$$

where, for $j \in \{1, \dots, T\}$, the posterior parameters are defined as follows:

$$\bar{\beta}_j = (I_N \otimes \bar{\kappa}_j^{-1})[(I_N \otimes X'D_jX)\hat{\beta}_j + (I_N \otimes \underline{\kappa}_j)\underline{\beta}_j], \quad (8)$$

$$\bar{\kappa}_j = \underline{\kappa}_j + X'D_jX, \quad (9)$$

$$\bar{\alpha}_j = \underline{\alpha}_j + \sum_{t=1}^T \vartheta_{j,t}, \quad (10)$$

$$\bar{\gamma}_j = \underline{\gamma}_j + Y'D_jY + c_j\underline{\kappa}_j c_j' - \bar{B}_j \bar{\kappa}_j \bar{B}_j', \quad (11)$$

$$D_j = \text{diag}(\vartheta_{j,1}, \dots, \vartheta_{j,T}), \quad (12)$$

and $\hat{\beta}_j = (I_N \otimes X'D_jX)^{-1}(I_N \otimes X'D_jX)y$ is the local likelihood estimator for β_j derived by Giraitis et al. (2014). Thus, I obtain draws for the time-varying parameters for each period by drawing from a Normal-Wishart posterior, which is computationally feasible even for large systems.

For the subsequent analysis, I cast the VAR in (1) in companion form:

$$Y_t = \mu_t + F_t Y_{t-1} + V_t, \quad V \sim N(0, \Omega^*), \quad (13)$$

where μ_t contains the VAR intercepts, F_t contains the AR-coefficients, and the first $n \times n$ elements of Ω^* correspond to Ω . I compute the time-varying impulse response functions (IRFs) using the model's MA(∞) representation (Giraitis, Kapetanios, and Yates, 2018):

$$Y_t = (I_N - F_t)^{-1} \mu_t + \sum_{h=0}^{\infty} F_t^h V_{t-h} + o_p(1). \quad (14)$$

2.3 Model and prior specifications

For the estimation of the model, I set $p = 4$ and impose Minnesota-style priors for $j \in \{1, \dots, T\}$ as in Kadiyala and Karlsson (1997). Since the model is estimated in growth rates, the prior for each coefficient is centered around zero, implying that each element of $\underline{\beta}_j$ is zero. The priors' hyperparameters are set to standard values. For autoregressive coefficients, $(\underline{\kappa}_j)_{ik}$ is set to $\lambda \sigma_i^2 / l^2$, where σ_i^2 denotes the residual variance from an AR(p) regression via OLS for variable i , and l is the respective lag of the coefficients. For intercept coefficients, $\underline{\kappa}_j$ is set to 10^2 . The overall tightness λ is set to 0.1. The scale and degree of freedom parameters of the Wishart prior $\underline{\alpha}_j$ and $\underline{\gamma}_j$ are set to σ_i^2 and $N + 2$,

which is the minimum number to ensure existence of the prior variance. The bandwidth parameter H is set to \sqrt{T} , which is asymptotically optimal (in terms of MSE minimizing, see Giraitis et al., 2014). In addition, I impose a stability constraint, ensuring that only draws leading to a stable system are retained.

I have conducted several robustness checks and find that the results presented in the subsequent sections are largely stable. First, the results are robust to different degrees of overall tightness ($\lambda \in \{0.01, 0.05, 0.25, 0.5\}$). Only for loose priors, the estimates suggest that the model suffers from overfitting, which leads to economically implausible results. Second, changing the degrees of freedom of the Wishart prior leaves the results virtually unchanged. Third, the results are stable with regard to different values for the bandwidth parameter ($H \in \{T^{0.4}, T^{0.6}\}$). Fourth, changing the lag length to 2 and 6 does not change the overall pattern of the results. Some of these results are presented in Appendix B; the remaining results are available upon request.

3 Reduced-form analysis

This section provides reduced-form evidence for changes of the time series properties of the German economy. I investigate fluctuations in both the low- and high-frequency domain of the variables. While the low-frequency analysis focuses on the long-run trends, the high-frequency analysis examines the variables' volatility.

3.1 Long-run means

I compute long-run trends along the lines of the Beveridge and Nelson (1981) concept, which defines a trend as an infinite horizon forecast:

$$\tau_t = (I_N - F_t)^{-1}\mu_t, \quad \tau_t = \lim_{h \rightarrow \infty} E_t y_{t+h}. \quad (15)$$


Figure 1 graphs the evolution of the time-varying trends for models A (solid line) and B (dashed line). To ease comparison, the trends are expressed in terms of annualized rates.⁷ Overall, both models yield similar results; only after the global financial crisis in 2008/2009 the means diverge. Moreover, the long-run means point at considerable time-variation in the VAR coefficients, which appears to be more pronounced during the first half of the sample.⁸

The trend of the short-term interest rate shows the well-documented decline over time (see, for instance, Summers, 2014). Moreover, this decline appears to gain momentum in

⁷Figure 8 in Appendix B illustrates the long-run means along with the actual series.

⁸Inspecting the determinant of the companion matrix F_t as an overall measure of time-variation underpins this finding.

Figure 1: Evolution of the time-varying trends



Notes: Posterior median of time-varying trends for models A (solid line) and B (dashed line). Trends are expressed in terms of annualized rates. Shaded areas and dotted lines refer to 68% equal-tailed point-wise posterior probability bands. Dashed vertical lines refer to historic event; for description see top left panel. Model A uses actual the short-term interest rate; Model B uses the shadow rate.

the mid-1990s. Prior to the German reunification in 1990, the interest rate trends were rather stable around seven percent. At the end of the sample, the trends drop to zero percent (model A) and even below (model B). This decline implies that, as a consequence of adverse shocks, the ELB may become binding more often in the future.

For the evolution of trend inflation, I provide two major findings. First, trend inflation strongly drops over time—similar to other advanced economies (Garnier, Mertens, and Nelson, 2015). Second, this drop emerged well before the Economic and Monetary Union (EMU) has been established in 1992; roughly 60% of the overall drop in trend inflation has occurred already between 1975 and 1985. Thus, contrary to other major economies, the cost-push shocks of the 1970s and 1980s are not reflected in higher trend inflation. After Germany gave up its monetary authority, trend inflation remained stable and settled at about 1.3%. This implies that the median estimate for the trend real interest rate, which can be related to a measure of the natural interest rate, is below zero percent from 2010 onward.⁹

⁹This results is broadly in line with the findings of Fries et al. (2018). However, fluctuations in the trend real interest rate can be due to shifts in the natural interest rate or shifts in the inflation target. Since the TVP-SV-VAR can not differentiate between both sources of variation, its results should be taken with some caution.


Trend GDP and wage growth evolve less steady. After a temporary increase during the late 1980s, they peak just before the reunification. Afterwards, they exhibit a steady decline and bottom out around 2005. One reason for the decline in the long-run GDP growth might be the lower productivity level of Eastern versus Western Germany, with consequences for the overall productivity level. Major sources of the productivity differential are a strong focus on domestic markets, missing economies of scale, and lower human capital in Eastern Germany (see, for instance, Burda and Severgnini, 2018).

The decline in real wage growth after the reunification is even more pronounced. In particular, changes in the institutional setup of the German labor market—for instance, lower union coverage or a decreasing share of industry-wide union agreements—have led to a strong decentralization in wage setting in Germany. The latter allowed wages to react more flexibly in Germany compared to other advanced economies and led to a persistent increase in competitiveness (Dustmann et al., 2014). Accordingly, the trends of GDP growth and wage growth pick up again around 2005. While trend wage growth reaches its pre-reunification level, GDP trend growth, however, remains on a substantially lower level. Moreover, the variability of these trends appears to have decreased and particularly trend wage growth does not display notable changes during the last years of the sample. The trend in hours growth displays a smooth increase until the late 1980s. During the 1990s, trend hours growth is roughly constant, while it starts to increase in the early 2000s. Compared to trend GDP growth, trend hours growth seems to have picked up somewhat earlier, stronger, and more persistent, suggesting a slowdown in trend labor productivity growth in Germany since the early 2000s. The trend of the real exchange rate change shows strong variations in the first half of the sample. After the build-up of the EMU, however, it is close to zero, implying that there is no systematic appreciation or depreciation.

In addition, the global financial crisis has only a minor impact on the trend estimates, suggesting that the models interpret it as a temporary shock that mainly affects the residuals' volatility. Similar to Ball (2014), this outcome assigns a limited role to the financial crisis on German potential output estimates. One explanation might be the so-called German labor market miracle (Burda and Hunt, 2011). The latter refers to the fact that while Germany experienced a strong drop in real GDP growth in 2008/09, the increase in unemployment was surprisingly moderate. In contrast, the US, France, and the UK experienced strong reactions of both GDP growth and the unemployment rate.¹⁰ A likely rationale behind these differences is the German short-time working scheme, which was gradually made more attractive for firms during the crisis (Brenke,

¹⁰Annual GDP growth in 2009 in Germany was -5.6% in 2009. The figures for the US, France, and the UK are, respectively: -2.5%, -2.9%, and -4.2%. The unemployment rate increased in the same period by 0.3 percentage points in Germany, while it increased in the US, France, and the UK by, respectively, 3.5, 1.6, and 1.9 percentage points.

Figure 2: Evolution of the total prediction variance



Notes: Posterior median of total prediction variance ($\log |\Omega_t|$) along with 68% equal-tailed point-wise posterior probability bands for models A (solid line) and B (dashed line).

Rinne, and Zimmermann, 2013) and allowed them to maintain their level of employment by reducing the hours worked per worker.¹¹ Particularly export-orientated firms from the manufacturing sector—suffering the most from the crisis—strongly benefited from this possibility (Rinne and Zimmermann, 2012). When global demand was recovering, these firms could quickly adapt and increase production. Hence, a hysteresis effect with regard to the unemployment rate did not emerge and GDP trend growth remained largely unaffected.

3.2 Volatility

Figure 2 plots the evolution of the total prediction variance, which can be expressed as the log determinant of the VAR’s residual covariance matrix (Whittle, 1953). This measure is interpreted as the total size of shocks hitting the economy in each period and serves as a proxy of the system’s short-run uncertainty (Cogley and Sargent, 2005). Evidently, the total prediction variance steadily declines over time, indicating a substantial decrease in short-run uncertainty since the 1970s. However, it appears that business cycle volatility is rather stable from the late 1990s until the onset of the global financial crisis. After the crisis, the total prediction variance continues its downward trend and reaches a historically low level at the end of the sample. Consequently, the German economy is currently less exposed to shocks. Similar to the US case (see, for example, Gadea Rivas, Gómez-Loscos, and Pérez-Quirós, 2014), also in Germany the global financial crisis seems to have only disrupted, but not ceased, the Great Moderation.

The results for models A (solid line) and B (dashed line) are virtually identical only until the pre-global financial crisis boom. Afterwards, according to model B, short-run uncertainty is somewhat higher compared to model A, but still on a very low level. Thus,

¹¹Balleer et al. (2016) find that the increase in the unemployment rate in Germany during the financial crisis was dampened by 1.29 percentage points due to short-time working, which amounts to roughly 466000 saved jobs.

the results suggest that focusing solely on actual interest rates underestimates the actual uncertainty of the system, because it ignores the impact of unconventional monetary policy since the financial crisis.

To gauge the sources behind this development, Figure 3 provides a closer look at both the unconditional standard deviations of each variable (left column) and the standard deviations of the reduced-form shocks to the variables (right column), that is the remaining elements of Ω_t . The interest rate (according to model A), inflation, wages, and the exchange rate display an almost monotonic decline in unconditional standard deviations, reaching historically low levels at the end of the sample. For GDP and hours growth, the fluctuations are more cyclical, but also downward trending. Moreover, it appears that the cyclical behavior of hours growth decreases stronger than those of output growth, suggesting that shocks to output go along with smaller shocks to hours worked. At the end of the sample, both GDP growth and hours growth volatilities are on historically low levels.


Relative to the remaining variables, exchange rate volatility displays the strongest decline and approaches two percent at the end of the sample. The major part of this decline, though, occurs already in the 1970s and 1980s, suggesting that the implementation of EMU is not the major driver of the reduction of exchange rate volatility in Germany. Interestingly, the exchange rate volatility is unaffected by the European Monetary System (EMS) crisis in 1992.

The unconditional standard deviations from models A and B coincide for each variable except for the interest rate. In the latter case, the volatility of model B, which includes the shadow rate, is on a relatively high level of about 0.5 percent in 2019. By contrast, the actual short-term rate's volatility is close to zero. Thus, the increase in the noise hitting the German economy can be traced back to a strong increase in interest rate variability. At the same time, the unconditional volatilities of the remaining variables are, according to both models, on a remarkably low level.

The declines in the volatilities of the series over time go along with strong reductions of the volatilities of the reduced-form shocks to these variables (right column). Thus, the largest share of variations in the variables' volatilities can be attributed to variations in the magnitude of the reduced-form shocks. Notably, in the case of the interest rate, the innovation standard deviations are smaller than the unconditional standard deviations, implying that changes in the volatility of shocks to the interest rate only account for a fraction of the change in the unconditional volatility. A likely explanation for this result is that the interest rate reacts to changes in GDP growth and inflation (according to a Taylor-rule), while exogenous fluctuations in the instrument itself are avoided.

In sum, the reduced-form analysis points at important changes in the German economy and a stabilization of the business cycle during the last five decades. To provide a first

Figure 3: Evolution of the covariance matrix



Notes: Posterior median of the unconditional standard deviations (left column) and the innovation standard deviations (right column) along with 68% equal-tailed point-wise posterior probability bands. The solid lines refer to model A; the dashed lines refer to model B.

assessment whether this stabilization is rather driven by changes in the dynamics of the system, that is the structure of the economy, or by a drop in the magnitude of the reduced-form shocks, I compute counterfactual unconditional standard deviations along the lines of Stock and Watson (2003). Specifically, I assume that the coefficients are fixed on their average over the first ten years of the sample (t_1) while the covariance matrix is fixed to its average of the last ten years of the sample (t_2) and vice versa.

Table 1 reports the results from this exercise. Columns two and three additionally depict the sample standard deviations for the periods t_1 and t_2 . These are fairly similar

Table 1: Counterfactual unconditional standard deviations

Variable	Sample std. dev.		Unconditional std. dev.			
	$t_1=1975-1985$	$t_2=2010-2019$	$\sigma(\hat{B}_{t_1}, \hat{R}_{t_1})$	$\sigma(\hat{B}_{t_2}, \hat{R}_{t_2})$	$\sigma(\hat{B}_{t_2}, \hat{R}_{t_1})$	$\sigma(\hat{B}_{t_1}, \hat{R}_{t_2})$
Short-term interest rate	0.64	0.25	0.55	0.19	0.54	0.38
Inflation	0.43	0.14	0.45	0.26	0.35	0.30
GDP growth	0.98	0.60	0.94	0.71	0.84	0.75
Wage growth	1.13	0.33	1.02	0.38	0.87	0.46
Hours growth	0.58	0.37	0.55	0.39	0.51	0.41
Exchange rate change	1.91	1.12	1.07	1.70	1.07	1.72

Notes: Columns two and three depict the series' average sample standard deviations over the specified windows. Columns four and five depict the average over the median unconditional standard deviations. Columns six and seven depict the average over the median counterfactual unconditional standard deviations.

to the estimated standard deviations from the model (columns four and five). Using the shocks from t_2 along with the coefficients from t_1 (last column) leads, in most cases, to a substantial decline in volatility; for GDP growth, for instance, the standard deviations drop from 0.94 to 0.75, which is close to the actual unconditional standard deviation from the model for this period (0.71). Similar patterns hold for the remaining variables, suggesting that the shocks account for a great deal of the reduction in volatility. Conversely, using the coefficients from t_2 along with the shocks from t_1 (column six) leads, in most cases, only to a moderate yet sizable decline in the magnitude of the unconditional standard deviations. A notable exception is the exchange rate change. In this case, the drop in volatility is almost entirely driven by changes in the coefficients. Overall, the results hence suggest that smaller shocks are not the sole driver of business cycle stabilization in Germany and that parts of the decline in volatility can be attributed to changes in the way the German economy responds to these shocks.

4 Structural analysis

To shed light on the sources behind the results presented in the previous sections, this section provides a structural analysis based on the TVP-SV-VAR and proceeds as follows.¹² First, I present the shocks' average effects computed as the average of the median responses over time. Second, I provide evidence on how the impact reactions to the shocks and the shock propagation have changed over time and by this means provide an economic interpretation of the results presented in Section 3. Finally, I compute counterfactual simulations to assess how the identified shocks have contributed to the decline of the business cycle volatility in Germany.

¹²Since the reduced-form results from models A and B are virtually identical until the short-term interest rate reaches the ELB, this section reports only the results for model B that uses the shadow rate as policy instrument.

4.1 Identification

I study the response of the German economy to three key macroeconomic shocks: an oil supply shock, a global activity shock, and a domestic activity shock. The following factors motivate the choice of the shocks: (i) oil supply shocks are found to induce strong effects on the German economy (Blanchard and Galí, 2007; Kilian, 2008); (ii) high export orientation and energy intensity of German economic activity render oil supply and global activity shocks particularly important (see Peersman and van Robays, 2012, for a comparison of the role of oil across countries); (iii) while there is strong evidence for changes in the effects of oil supply shocks over time for several countries (see, for instance, Herrera and Pesavento, 2009; Peersman and van Robays, 2012; Baumeister and Peersman, 2013), a detailed time-varying analysis for Germany is hitherto missing.

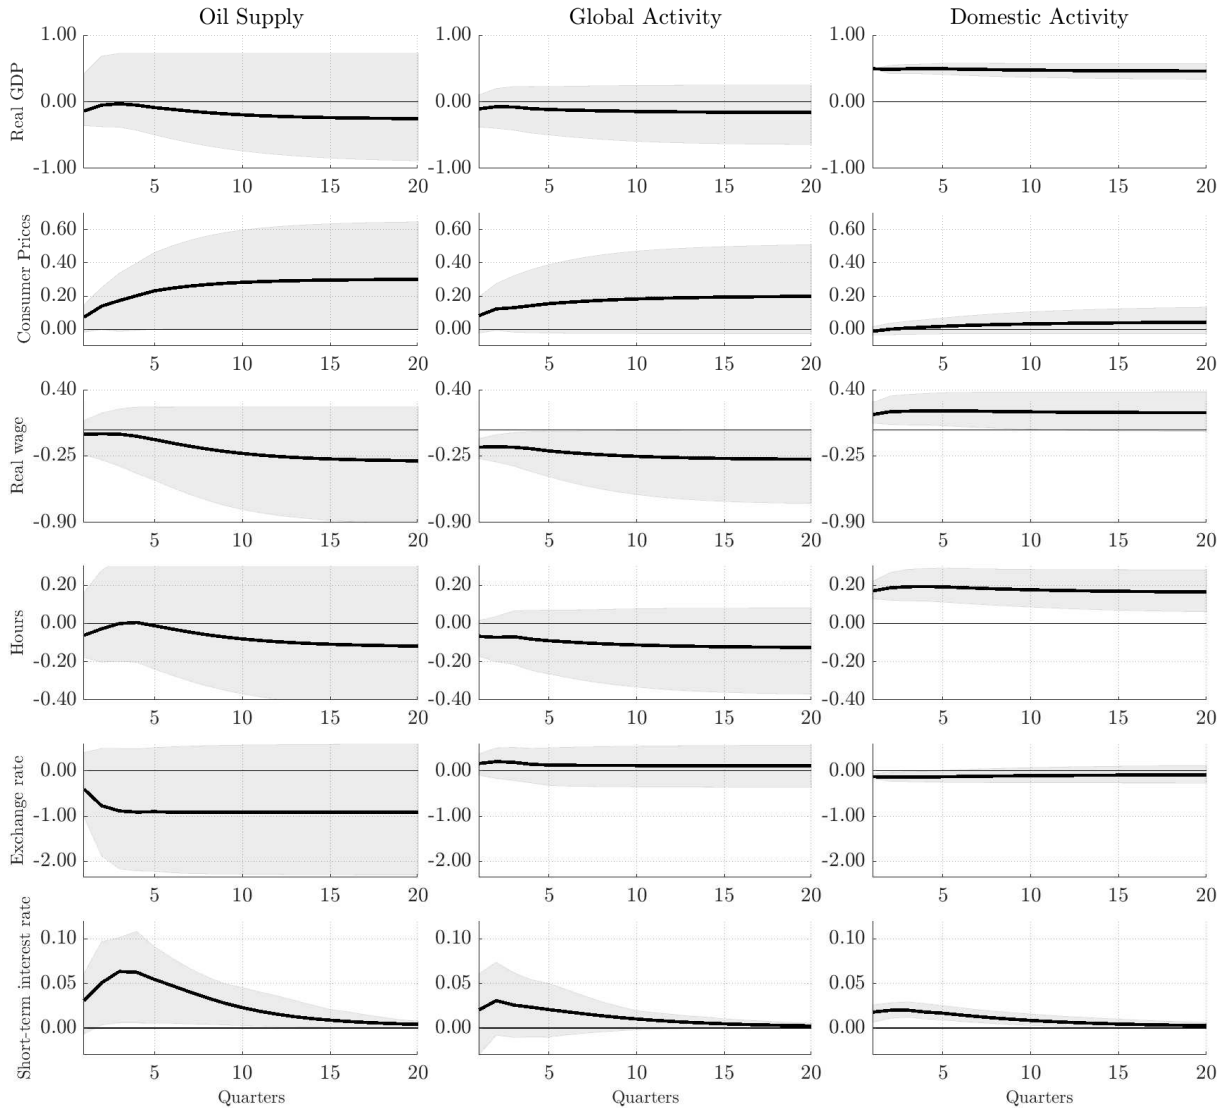
As mentioned in Section 2, the oil supply shock is not internally identified but directly enters the model as an external instrument. Specifically, I use the oil supply news shocks series constructed by Känzig (2021), which applies the high-frequency identification methodology for monetary policy shocks (see, for example, Kuttner, 2001) to the oil market. Accordingly, oil supply surprises are defined as the (log) difference between an oil future price on the day of an OPEC announcement and its price on the last trading day prior to the announcement. In the following, I use the composite measure of Känzig (2021), which is the first principal component of surprises referring to future maturities ranging from one to twelve months. I choose this measure for three reasons: (i) it is shown to be a strong instrument for the price of oil; (ii) it is shown to have statistically and economically significant effects on the oil market; (iii) it is readily available for a long period of time.

To assess the impact of these shocks, I use a recursive identification scheme as in, for instance, Kilian (2006). Given the exogeneity of the oil shock series, it is ordered first, which allows for valid structural estimation of the responses to the instrument (Plagborg-Møller and Wolf, 2021). The innovation in the growth rate of global industrial production, which is ordered second, is labeled global activity shock. The innovation in German GDP growth, which is ordered third, is labeled domestic activity shock. This scheme implies that shocks originating in the German economy do not affect global production within a quarter. Since Germany can be classified as a small open economy, this assumption is innocuous. Albeit concededly simple, this identification allows for assessing the time-varying response of the German economy to key shocks.

4.2 The average effects of the shocks

Figure 4 shows the variables' average responses to the shocks (depicted by the mean over the median responses per period) and the heterogeneity of the responses over time

Figure 4: Average impulse responses over time



Notes: Solid lines depict average impulse response over the entire sample, i.e. the average of the median responses per period. Shaded areas are 68% confidence bands. Responses have been accumulated and are reported in levels.

(depicted by 16th and 84th percentiles). Thus, wide confidence bands indicate strong heterogeneity in the responses over time. For each period, the activity shocks are normalized to a 50 basis point (bp) impact increase in global industrial production growth and German GDP growth, respectively. The oil supply shock is normalized to a 10% increase of the real oil price. The estimated IRFs, except for the interest rate, have been accumulated and are reported in levels.

The IRFs exhibit reasonable patterns and confirm findings from previous research, suggesting that the shocks are correctly identified. Moreover, wide confidence bands suggest that the responses exhibit substantial time-variation. On average over all periods, the impact response of German GDP to an unfavorable oil supply shock (first column)

is near-zero, but slowly accumulates over time—a finding which has been documented by Hamilton (1983) for the US and subsequently has been confirmed for the US and other advanced economies, including Germany, by several studies (see, for instance, Carstensen, Elstner, and Paula, 2013; Peersman and van Robays, 2014). The drop in GDP also extends to the labor market and leads to both decreasing real wages and hours. The long-run impact on wages is about twice the size of the long-run impact on hours. The real exchange rate depreciates on impact and hence prevents a stronger drop in domestic economic activity. Unsurprisingly, consumer prices persistently increase, reflecting both direct (since energy is part of CPI) and indirect effects (for example, higher input costs for firms). Consistent with this finding, monetary policy tightens to stabilize inflation.¹³

The global activity shock (second column) operates, on average over all periods, as a negative supply shock, dampening output and boosting inflation. A likely rationale for this result is that the indirect contractionary effect from the oil price rise, which is induced by the sudden hike in global production, outweighs the direct expansionary effect of higher global demand. The labor market responses resemble those of the oil supply shock, while the exchange rate moves in the opposite direction. Accordingly, this external appreciation suppresses the rise in CPI, which is somewhat lower for the global activity shock and causes a weaker reaction of monetary policy.

The domestic activity shock (third column) resembles, on impact, a positive technology shock, moving output and inflation in opposite directions and increasing real wages. These reactions are consistent with a wide array of theoretical models (see Peersman and Straub, 2009, among others). After some quarters, however, the response of CPI turns positive. Compared to the other shocks, the price reaction is only small, though. Further, the domestic activity shock has an expansionary labor input effect, as evidenced by a rise in hours. As suggested by Rujin (2019), given strict German labor market institutions and high costs associated with hiring and firing workers, the response of hours worked is likely to be driven by the adjustments along the intensive margin (hours per worker) rather than by adjustments along the extensive margin (employment). The exchange and interest rates are virtually non-responsive to the domestic activity shock.

¹³The model's implicit monetary policy rule is probably misspecified from 2000 onward since it only includes German and global variables, whereas ECB conducts monetary policy for the euro area and thus reacts to changes in euro area variables. Since I use the euro area short-term rate as policy variable and the German business cycle exhibits a strong co-movement with the euro area cycle (Eickmeier and Breitung, 2006; Lehwald, 2013), the identified monetary policy shocks after implementation of the EMU can nevertheless serve as a proxy for the actual non-systematic monetary policy in the euro area (Bijsterbosch and Falagiarda (2015)).

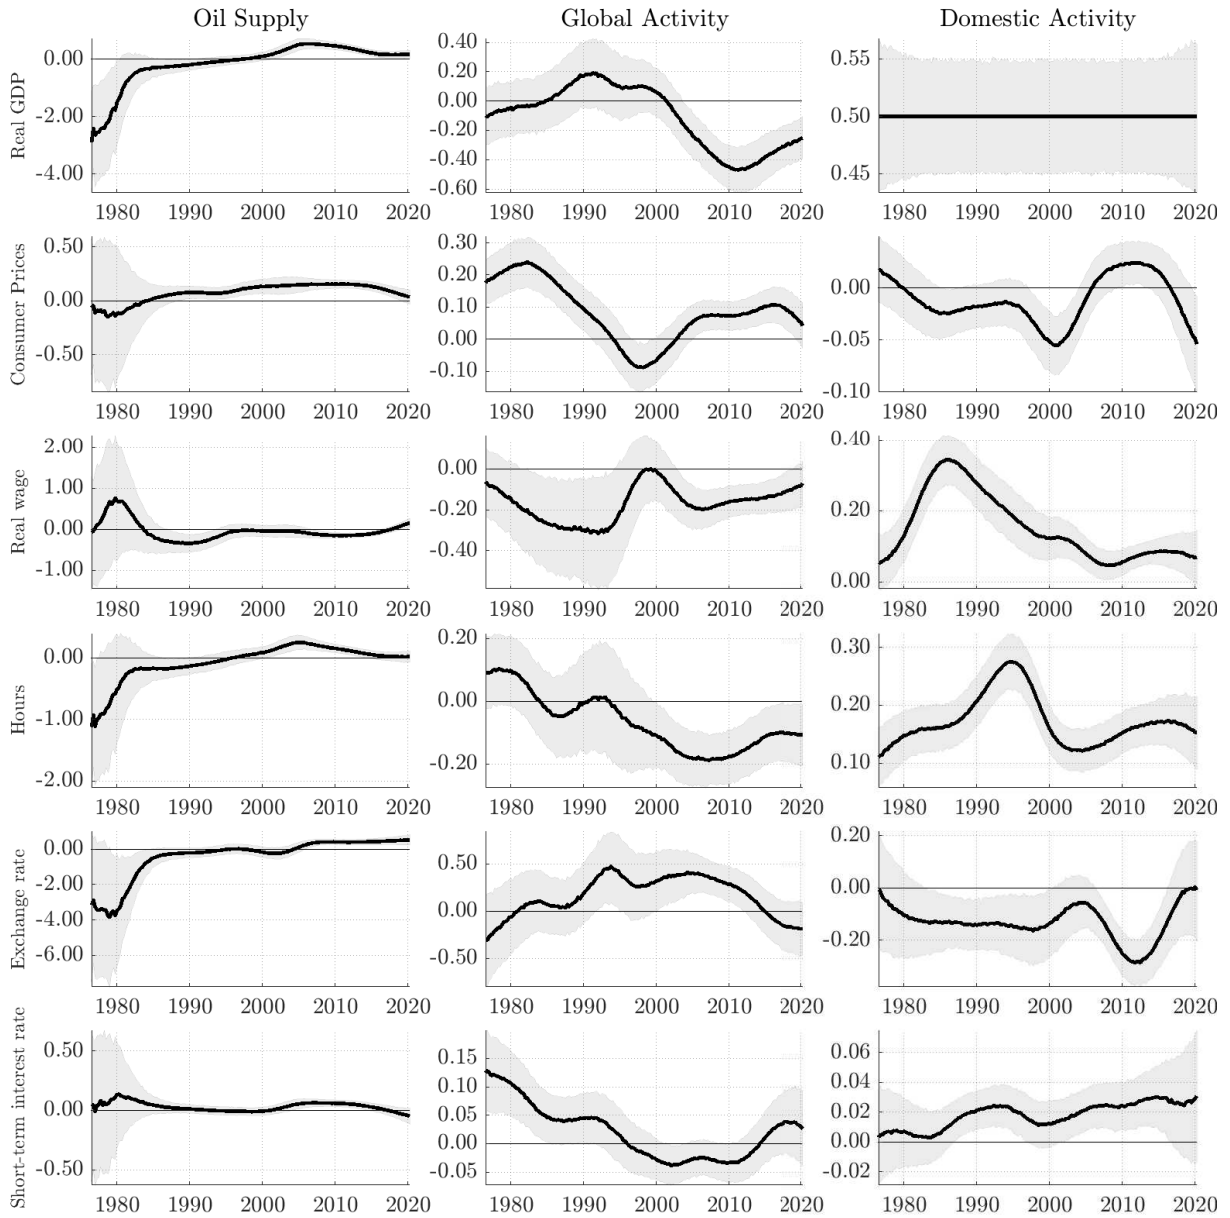
4.3 The time-varying effects of the shocks

To arrive at a better description of the responses' time-variation, Figure 5 plots the impact responses to the shocks for each period. The impact responses to the oil supply shocks display remarkable shifts over time. During the 1970s, an unfavorable oil supply shock causes German GDP to drop by about 2.5%. Subsequently, the magnitude of the impact response quickly diminishes. Around 1995, it switches the sign and becomes positive, but remains—on a historical scale—small. The declining responsiveness of GDP also transmits to the remaining variables—in 2019, the German economy is virtually non-responsive to sudden oil supply disruptions. The latter is consistent with Peersman and van Robays (2012), who document a considerable decline in the responsiveness of German GDP to oil supply shocks in the post-1986 period. Moreover, the muted response of German GDP resembles the findings for US GDP of Baumeister and Peersman (2013). A likely rationale for this pattern is the decreasing relative importance of the energy-intensive manufacturing sector, whose share in total gross value added dropped by about 17 percentage points in the sample period.

The impact responses to the global activity shock exhibit marked fluctuations over time, too. In most periods, the global activity shock dampens German GDP. Only during the late 1980s and 1990s, it induces a rise in GDP. The magnitude of the impact response is only mild until the mid-2000s, while it peaks during the global financial crisis with a decline of about 0.45%. Subsequently, it converges to zero. Its effect on German GDP is again relatively mild in 2019. The size of the effect on the real wage culminates in the mid 1990s. Afterwards, the magnitude substantially decreases. Global activity shocks exert only minor impact on real wages, particularly after the global financial crisis. Conversely, the size of the impact on hours peaks during the global financial crisis and subsequently stabilizes on a historically high level. The impact response of the interest rate is, in most periods, as expected—a rise in CPI goes along with a hike in the interest rate. However, from about 2005 to 2015, the interest rate drops despite an increase in CPI. A reasonable explanation for this result is that Germany responded differently to global activity shocks than the euro area during this period—in particular it recovered much faster from the global financial crisis. The latter is broadly consistent with Bobeica and Jarociński (2019), who stress the importance of global shocks for euro area inflation during this period.

For the domestic activity shock, I obtain particularly pronounced fluctuations in the impact responses to the labor market variables. During the 1980s, an expansionary domestic activity shock induces a decline in the real wages by up to 0.4% on impact. The response of hours culminates about ten years later during the post-reunification boom. Subsequently, the impact responses of both variables considerably decrease. Moreover, while the domestic activity shock operates, on average over time, as a technology shock, it resembles an expansionary demand shock around the global financial crisis.

Figure 5: Impact responses over time




Notes: Impact responses to identified shocks. Shaded areas are 68% confidence bands.

While Figure 5 indicates strong time-variation in the impact responses, there is also evidence for changes in the shock propagation. Figure 6 depicts the averages of the median responses for four time periods, which all consist of at least one entire stylized business cycle (see Schirwitz, 2009; Carstensen et al., 2020, for German business cycle chronologies). Overall, Figure 6 provides two key findings.

First, the average responses for the period 1975–1985 (solid line) are markedly different from the remaining responses. At the same time, the differences between the remaining responses are, in most cases, less pronounced, suggesting that large parts of the structural

Figure 6: Average impulse responses for selected periods



Notes: Figure depicts average over the median responses for 1975–1985 (solid line), 1992–1996 (dashed line), 2000–2009 (dotted line), and 2014–2019 (dashed-dotted line). Responses have been accumulated and are reported in levels.

change have occurred already before 1992. The shocks induce stronger effects in both the short and the long run in 1975–1985. Moreover, the shocks are substantially more persistent in the latter period.

Second, an impact response of a similar magnitude can lead to substantially different long-run effects. For example, an instantaneous rise in GDP by 0.5% leads to a long-run effect that varies between 0.4% and 0.55%. Conversely, different impact responses can exert similar long-run effects, see, for example, the response of hours to global activity

shocks. Thus, Figure 6 suggest that there is sizable change in the shock propagation process.

For 1992–1996 (dashed line), three observations stand out. First, an expansionary global activity shock induces a rise in domestic GDP, whereas GDP drops in the remaining periods. Second, the positive domestic activity shock leads to strong and persistent CPI inflation and wage growth with long-run effects that are substantially larger than in remaining periods. Third, the domestic activity shock leads to an appreciation of the exchange rate in the long run, while the exchange rate depreciates in the long run for the remaining periods.

In 2000–2009 (dotted line), the responses are, in most cases, characterized by a persistent hump-shaped pattern. For example, the response of GDP to the domestic activity shock shows a larger effect in the short run, but a smaller effect in the long run compared to 1992–1996 and 2014–2019. Moreover, the interest rate response to the global activity shock is quite distinct in 2000–2009—the drop in GDP and rise in CPI go along with a monetary loosening.

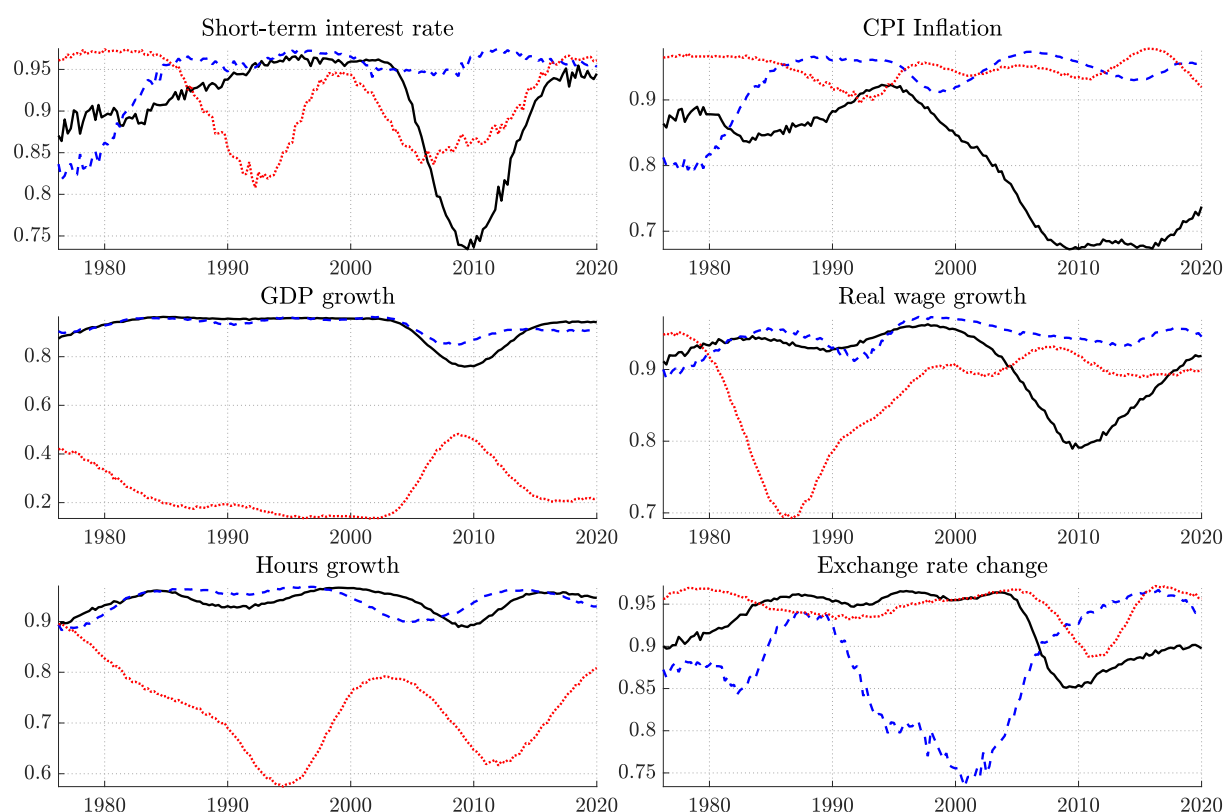
For 2014–2019 (dashed-dotted line), the shocks appear to be rather unpersistent. Following an impact reaction, the variables quickly settle at a new steady state, suggesting that the German economy has improved in absorbing exogenous shocks. The latter provides evidence, that the decline in GDP growth volatility (see Figure 3) is not solely driven by good luck.

In summary, the structural analysis underpins the findings of the reduced-form analysis and reveals substantial change in the German business cycle. However, the results suggest that it is unlikely that good luck is the only explanation for the reduction of business cycle volatility in Germany. In fact, I document that the variables’ impact responses to shocks of constant size have, in most cases, substantially decreased over time. Moreover, the results indicate that the shock propagation has changed considerably. In particular, the shocks’ persistence appears to have declined. Thus, the results indicate that the German economy has become more resilient to exogenous disturbances.

4.4 Counterfactual Analysis

In this section, I assess the role played by the structural shocks for the reduction in business cycle volatility documented in Section 3.2. To this end, I follow Akram and Mumtaz (2019) by computing counterfactual unconditional variances under the assumption that the volatility of the shocks are set to zero one at a time. If the differences between the baseline unconditional volatility of a variable (see Figure 3) and its counterfactual volatility is large, it is suggested that the excluded shock was the driving force of the variable’s volatility.

Figure 7: Counterfactual unconditional volatilities



Notes: Counterfactual median estimates for unconditional volatilities according to Model B. Solid line: oil supply shock volatility is zero; dashed line: global activity shock volatility is zero; dotted line: domestic activity shock volatility is zero. Figures are expressed as percentages relative to the baseline unconditional volatilities.

Figure 7 depicts these counterfactual volatilities in percentage terms relative to the baseline unconditional volatility. Setting the volatility of the oil supply shock to zero (solid line) appears to be rather inconsequential for the variables' volatility until the mid 2000s. Afterwards, the series' volatility would be substantially lower in most cases. For example, the unconditional volatility of CPI inflation would be up to 30% lower around 2010. Setting the volatility of the global activity shock to zero (dashed line) also reduces the variables' volatility, albeit in most cases to a smaller extent compared to the oil supply shock. Except for the exchange rate change, the reduction in volatility fluctuates around 10%, suggesting that the volatility of the global activity shock played a minor role for the decline in German business cycle volatility. Conversely, the volatility of the domestic activity shock strongly affects each variables' unconditional variances. In particular, GDP growth volatility would be substantially (up to 85%) lower when the volatility of domestic activity shocks is switched off. Thus, German GDP growth volatility is primarily driven by domestic activity shocks, which underpins the findings of previous studies (see, for instance, Carstensen and Salzman, 2017). However, also the labor market variables'

volatility would be strongly reduced if the volatility of domestic activity shocks is zero. Hence, these results provide evidence that the reduction in business cycle volatility in Germany is primarily shaped by a reduction of the volatility of domestic activity shocks while the evolving volatility of the oil supply and global activity shocks played only a minor role.

5 Concluding remarks

The reduction of business cycle volatility has been found for several countries, including Germany. This paper provides a comprehensive view on this issue by means of a time-varying parameter VAR with stochastic volatility. I conduct both a reduced-form and a structural analysis. The former demonstrates that not only the volatility of output growth has substantially declined over time, but also the volatility of other key macroeconomic variables. These reductions were mainly driven by smaller reduced-form shocks hitting the variables over time. A closer look at the time profile of the impact effects of identified structural shocks reveals important changes in the systematic responses to these shocks. In particular, the results suggest that the German economy has become more resilient to exogenous disturbances. Moreover, inspecting the shock propagation reveals that the shocks' persistence has declined and the variables' adjustment to a new steady state has fasten. Thus, the results suggest that both good luck and good policy have contributed to business cycle stabilization in Germany. Finally, counterfactual simulations highlight the importance of fluctuations in the volatility of shocks originating in the German economy for the reduction in business cycle volatility.

References

- Akram, Q. F. and H. Mumtaz (2019). Time-varying dynamics of the Norwegian economy. *The Scandinavian Journal of Economics* 121(1), 407–434.
- Aßmann, C., J. Hogrefe, and R. Liesenfeld (2009). The decline in German output volatility: a Bayesian analysis. *Empirical Economics* 37, 653–679.
- Ball, L. (2014). Long-term damage from the Great Recession in OECD countries. *European Journal of Economics and Economic Policies: Intervention* 11(2), 149–160.
- Balleer, A., B. Gehrke, W. Lechthaler, and C. Merkl (2016). Does short-time work save jobs? A business cycle analysis. *European Economic Review* 84, 99–122.
- Baumeister, C. and J. D. Hamilton (2019). Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks. *American Economic Review* 109(5), 1873–1910.
- Baumeister, C. and G. Peersman (2013). Time-Varying Effects of Oil Supply Shocks on the US Economy. *American Economic Journal: Macroeconomics* 5(4), 1–28.
- Beveridge, S. and C. R. Nelson (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. *Journal of Monetary Economics* 7(2), 151–174.
- Bijsterbosch, M. and M. Falagiarda (2015). The macroeconomic impact of financial fragmentation in the euro area: Which role for credit supply? *Journal of International Money and Finance* 54, 93–115.
- Bitto, A. and S. Frühwirth-Schnatter (2019). Achieving shrinkage in a time-varying parameter model framework. *Journal of Econometrics* 210(1), 75–97.
- Blanchard, O. J. and J. Galí (2007). The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s? In *International Dimensions of Monetary Policy*, NBER Chapters, pp. 373–421. National Bureau of Economic Research, Inc.
- Bobeica, E. and M. Jarociński (2019). Missing Disinflation and Missing Inflation: A VAR Perspective. *International Journal of Central Banking* 15(1), 199–232.
- Brenke, K., U. Rinne, and K. F. Zimmermann (2013). Short-time work: The German answer to the Great Recession. *International Labour Review* 152, 287–305.
- Buch, C. M., J. Doepke, and C. Pierdzioch (2004). Business cycle volatility in germany. *German Economic Review* 5(4), 451–479.

- Burda, M. C. and J. Hunt (2011). What Explains the German Labor Market Miracle in the Great Recession. *Brookings Papers on Economic Activity* 42(1), 273–335.
- Burda, M. C. and B. Severgnini (2018). Total factor productivity convergence in German states since reunification: Evidence and explanations. *Journal of Comparative Economics* 46(1), 192–211.
- Carstensen, K., S. Elstner, and G. Paula (2013). How Much Did Oil Market Developments Contribute to the 2009 Recession in Germany? *The Scandinavian Journal of Economics* 115(3), 695–721.
- Carstensen, K., M. Heinrich, M. Reif, and M. H. Wolters (2020). Predicting Ordinary and Severe Recessions with a Three-State Markov Switching Dynamic Factor Model. An Application to the German Business Cycle. *International Journal of Forecasting* 36(3), 829–850.
- Carstensen, K. and L. Salzmann (2017). The G7 business cycle in a globalized world. *Journal of International Money and Finance* 73, 134–161.
- Cogley, T. and T. J. Sargent (2005). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics* 8(3), 262–302.
- D’Agostino, A., L. Gambetti, and D. Giannone (2013). Macroeconomic forecasting and structural change. *Journal of Applied Econometrics* 28(1), 82–101.
- Dustmann, C., B. Fitzenberger, U. Schönberg, and A. Spitz-Oener (2014). From Sick Man of Europe to Economic Superstar: Germany’s Resurgent Economy. *Journal of Economic Perspectives* 28(1), 167–88.
- Eickmeier, S. and J. Breitung (2006). How synchronized are new EU member states with the euro area? Evidence from a structural factor model. *Journal of Comparative Economics* 34(3), 538–563.
- Fries, S., J.-S. Mésonnier, S. Mouabbi, and J.-P. Renne (2018). National natural rates of interest and the single monetary policy in the euro area. *Journal of Applied Econometrics* 3(6), 763–779.
- Fritsche, U. and V. Kuzin (2005). Declining output volatility in Germany: impulses, propagation, and the role of monetary policy. *Applied Economics* 37(21), 2445–2457.
- Gadea Rivas, M., A. Gómez-Loscos, and G. Pérez-Quirós (2014). The two greatest. Great Recession vs. Great Moderation. Working Paper 1423, Banco de Espana.

- Galí, J. and L. Gambetti (2009). On the Sources of the Great Moderation. *American Economic Journal: Macroeconomics* 1(1), 26–57.
- Garnier, C., E. Mertens, and E. Nelson (2015). Trend Inflation in Advanced Economies. *International Journal of Central Banking* 11(4), 65–136.
- Giraitis, L., G. Kapetanios, and T. Yates (2014). Inference on stochastic time-varying coefficient models. *Journal of Econometrics* 179(1), 46–65.
- Giraitis, L., G. Kapetanios, and T. Yates (2018). Inference on Multivariate Heteroscedastic Time Varying Random Coefficient Models. *Journal of Time Series Analysis* 39(2), 129–149.
- Hamilton, J. D. (1983). Oil and the Macroeconomy since World War II. *Journal of Political Economy* 91(2), 228–248.
- Herrera, A. M. and E. Pesavento (2009). Oil Price Shocks, Systematic Monetary Policy, and the “Great Moderation”. *Macroeconomic Dynamics* 13(1), 107–137.
- Huber, F., G. Koop, and L. Onorante (2020). Inducing Sparsity and Shrinkage in Time-Varying Parameter Models. *Journal of Business & Economic Statistics* forthcoming.
- Kadiyala, K. R. and S. Karlsson (1997). Numerical Methods for Estimation and Inference in Bayesian VAR-Models. *Journal of Applied Econometrics* 12(2), 99–132.
- Känzig, D. R. (2021). The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements. *American Economic Review* 111(4), 1092–1125.
- Kilian, L. (2006). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. CEPR Discussion Papers 5994, C.E.P.R. Discussion Papers.
- Kilian, L. (2008). A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries. *Journal of the European Economic Association* 6(1), 78–121.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99(3), 1053–69.
- Korobilis, D. (2013). Assessing the Transmission of Monetary Policy Using Time-varying Parameter Dynamic Factor Models. *Oxford Bulletin of Economics and Statistics* 75(2), 157–179.
- Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics* 47(3), 523–544.

- Lehwald, S. (2013). Has the Euro changed business cycle synchronization? Evidence from the core and the periphery. *Empirica* 40(4), 655–684.
- Mills, T. C. and P. Wang (2003). Have output growth rates stabilised? evidence from the g-7 economies. *Scottish Journal of Political Economy* 50(3), 232–246.
- Mumtaz, H. and F. Zanetti (2015). Labor Market Dynamics: A Time-Varying Analysis. *Oxford Bulletin of Economics and Statistics* 77(3), 319–338.
- Peersman, G. and R. Straub (2009). Technology shocks and robust sign restrictions in a euro area SVAR. *International Economic Review* 50(3), 727–750.
- Peersman, G. and I. van Robays (2012). Cross-country differences in the effects of oil shocks. *Energy Economics* 34(5), 1532–1547.
- Peersman, G. and I. van Robays (2014). Oil and the Euro area economy. *Economic Policy* 24(60), 603–651.
- Petrova, K. (2019). A quasi-Bayesian local likelihood approach to time varying parameter VAR models. *Journal of Econometrics* 212(1), 286–306.
- Plagborg-Møller, M. and C. K. Wolf (2021). Local Projections and VARs Estimate the Same Impulse Responses. *Econometrica* 89(2), 955–980.
- Primiceri, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies* 72(3), 821–852.
- Prüser, J. (2020). A global-local prior for time-varying parameter VARs and monetary policy. Discussion Paper 20/2020, TU Dortmund.
- Rinne, U. and K. F. Zimmermann (2012). Another economic miracle? The German labor market and the Great Recession. *IZA Journal of Labor Policy* 1(1), 3.
- Rujin, S. (2019). What are the effects of technology shocks on international labor markets? Ruhr Economic Papers 806, RWI - Leibniz-Institut für Wirtschaftsforschung, Ruhr-University Bochum, TU Dortmund University, University of Duisburg-Essen.
- Schirwitz, B. (2009). A comprehensive German business cycle chronology. *Empirical Economics* 37(2), 287–301.
- Stock, J. H. and M. W. Watson (2003). *Has the Business Cycle Changed and Why?*, pp. 159–230. MIT Press.
- Stock, J. H. and M. W. Watson (2005). Understanding changes in international business cycle dynamics. *Journal of the European Economic Association* 3(5), 968–1006.

- Summers, L. H. (2014). Reflections on the ‘New Secular Stagnation Hypothesis’. In C. Teulings and R. Baldwin (Eds.), *Secular Stagnation: Facts, Causes and Cures*. CEPR Press.
- Summers, P. M. (2005). What caused the Great Moderation? Some cross-country evidence. *Economic Review* (Q III), 5–32.
- Whittle, P. (1953). The analysis of multiple stationary time series. *Journal of the Royal Statistical Society: Series B (Methodological)* 15(1), 125–139.
- Wu, J. C. and F. D. Xia (2017). Time-Varying Lower Bound of Interest Rates in Europe. Research Paper 17-06, Chicago Booth.

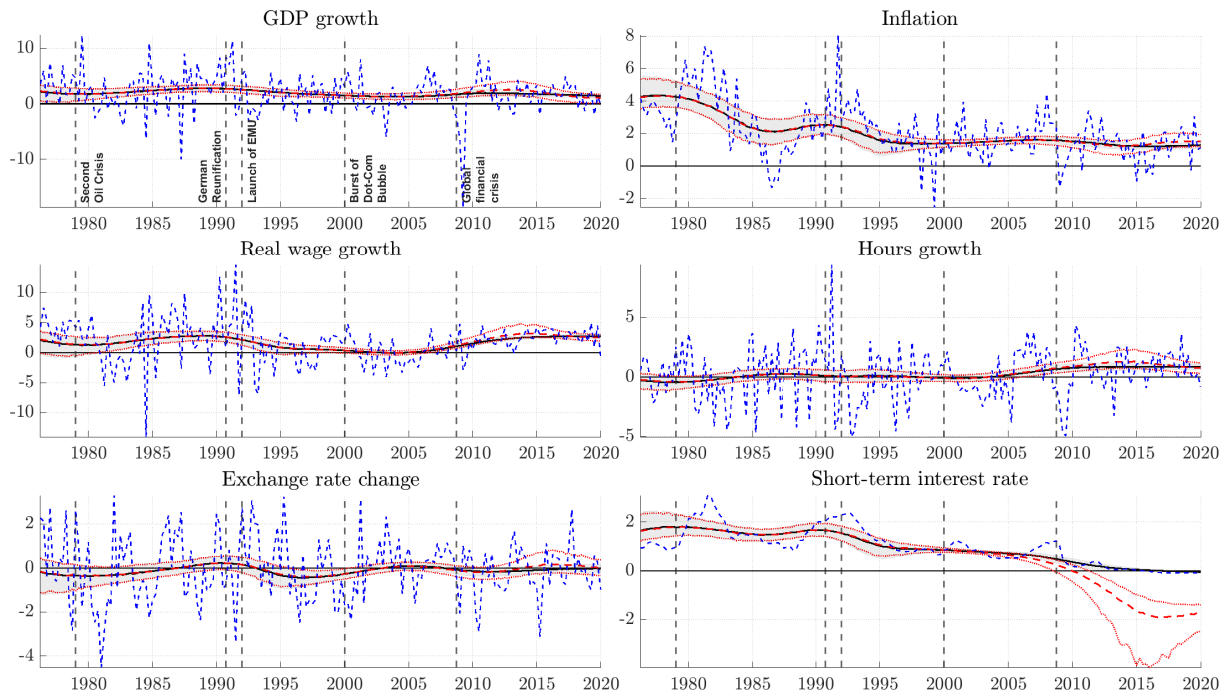
Appendix

A Data

I use quarterly data covering the period from 1975:Q1 to 2019:Q4. Real GDP is measured in terms of the seasonally adjusted series provided by the OECD quarterly national accounts, which refer to West Germany until 1991 and afterwards to reunified Germany. The seasonally adjusted consumer price index (CPI) is obtained from the German Bundesbank. The real effective exchange rate is provided by the Bank for International Settlements and refers to the weighted average of the nominal exchange rates of Germany vis-à-vis its 23 most important trading partners (narrow index) deflated with CPI. Hours worked and nominal wages are provided by the Federal Statistical Office. I obtain real wages by deflating nominal wages with CPI. As a proxy for global activity, I employ the global industrial production index provided by Baumeister and Hamilton (2019), which is an extended version of the OECD's index of monthly industrial production in the OECD and six major other countries. The included countries account for about 75% of global GDP. Following the recommendation of Kilian (2009), I use US refiners' acquisition cost of crude oil imports deflated with US CPI as a proxy for the global real price of crude oil. Refiners' acquisition cost are provided by the US Energy Information Administration (EIA); US CPI is obtained from the FRED database of the St. Louis FED. Finally, I use quarterly averages of the FIONIA until the end of 1999 and afterwards I switch to EONIA for model A. Regarding model B, EONIA is replaced by the shadow rate for the euro area provided by Wu and Xia (2017) from 2009 onward. Each series except the interest rates enter the model in quarter-on-quarter percentages changes. To make the figures for the interest rates commensurable with the remaining series, I compute the quarterly effective interest rate as $r_t = ((1 + r_t^A)^{0.25} - 1) \cdot 100$, where r_t^A denotes the annualized quarter-on-quarter interest rate.


B Additional Figures

Figure 8: Evolution of the time-varying trends along with actual series




Notes: Posterior median of time-varying trends for models A (solid line) and B (dashed line). Trends are expressed in terms of annualized rates. Shaded areas and dotted lines refer to 68% equal-tailed point-wise posterior probability bands. Dashed vertical lines refer to historic event; for description see top left panel. Dashed blue lines refer to actuals.

Figure 9: Evolution of total prediction variance for for different bandwidth parameters



Notes: Posterior mean of $\log |\Omega_t|$ for models A.

Figure 10: Evolution of the time-varying trends for different bandwidth parameters



Notes: Posterior mean of time-varying trends for models A. Trends are expressed in terms of annualized rates.