Chapter 3

Saving Shocks and Business Cycle Fluctuations

Abstract

A common modeling tool – and a popular narrative – used to explain the financial crisis of 2008-2010 is a sudden increase in the desire to save. Such "marginal propensity to save" (MPS) shocks can be triggered by, for instance, a rise in uncertainty surrounding the economic climate, and depress interest rates, inflation, and generally cause an economic contraction. This paper uses the long-run properties arising from MPS shocks in both exogenous- and endogenous growth models with sticky prices in order to identify their causal effect on output. We find that time series data from the United States is strongly supportive of the notion that MPS shocks indeed have a causal, and contractionary, effect on economic activity, lending support to the most common approach of studying the financial crisis.[1]

Keywords: Long-run restrictions; endogenous growth; marginal propensity to save.
JEL Classification: E21, E32, O33, O42.

[1]This chapter is based on joint work with Pontus Rendahl.

3.1 Introduction

There is growing consensus on a particular narrative about what triggered the financial crisis of 2008-2010. This narrative is essentially based on three concurring factors: First, housing prices started to decline in the summer of 2007. Second, the financial system was heavily invested in housing-related assets and mortgage-backed securities of low quality and third, the shadow banking system was highly leveraged in housing assets and highly vulnerable to bank runs (Christiano et al., 2017). The fall in housing prices led to a decline in asset values which forced the shadow banking system to sell them immediately, reinforcing the decline and damaging the whole banking system. The subsequent credit crunch made this decrease in asset values even more severe. The reduction in household wealth due to the interplay of these factors and the rise in uncertainty triggered households to spend less and save more. This increase in the saving rate is depicted in the left panel of figure 3.1[2]. While the savings rate has been declining for decades, there is a clear structural break around the time of the financial crisis sending the saving rate on an upward trend. As households cut back on spending, firms reduced investments and hiring due to falling sales. Hence, it can be argued that the increase in the marginal propensity to save (MPS) reinforced the effect of the three factors above, pushing the economy deeper into recession. Consequently, a common modelling tool to explain the financial crisis has been the introduction of MPS shocks as in Eggertsson et al. (2003) and Christiano et al. (2011).

Building on this narrative, this paper investigates the causal link between saving shocks and business cycle fluctuations. The right panel of figure 3.1 illustrates the need for identification as there is no clear relationship between personal savings and output growth. Moreover, identification approaches using short-run restrictions are hard to justify given the simultaneity issue between output and the saving rate, with the saving rate being defined as savings over income. Therefore, we use the long-run properties arising from MPS shocks in both exogenous- and endogenous growth models with sticky prices in order to identify their causal effect on output. Standard textbook growth models as the Solow model or endogenous versions such as the AK model and Schumpeterian growth models would predict an increase in growth following an increase in the saving rate.

[2] The plot excludes the period of the Covid crisis as an outlier for legibility reasons. Since this paper was written before the Covid crisis and to make sure results are not driven by the spike in saving rates, baseline results also exclude the Covid period, but we do include it for robustness checks.

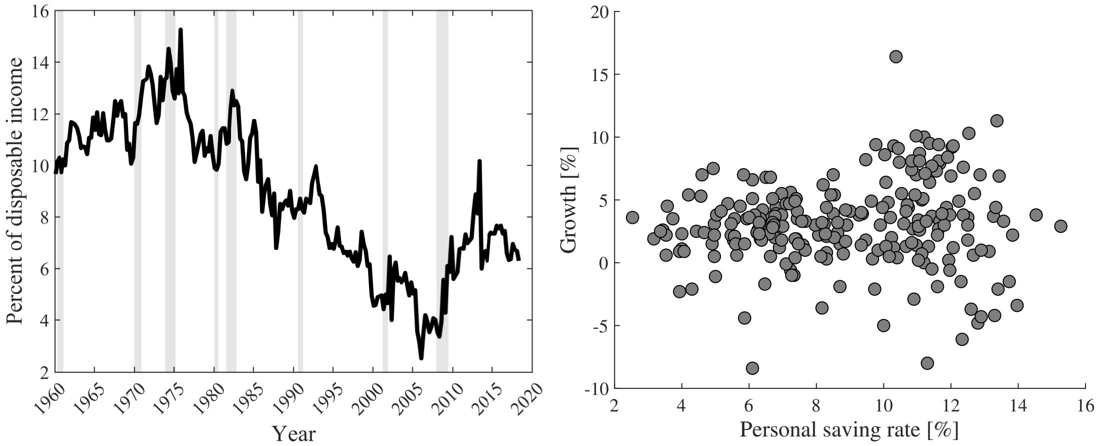

Fig. 3.1 Personal saving rate and growth

Notes. The left panel depicts the personal saving rate in the U.S. from Q1 1960 to Q2 2018. Grey shaded areas indicate the starting and ending dates for recession periods, as determined by the National Bureau of Economic Research (NBER). The right panel depicts the relationship between growth and changes in the saving rate.

However, they only offer a long-run perspective. When taking into account price rigidities, a rise in the saving rate can have a contractionary effect on output in the short-run[3].

Using time series data from the United States, our empirical findings are threefold: First, MPS shocks indeed have a causal, and contractionary, effect on economic activity, lending support to the most common approach of studying the financial crisis. Secondly, we find evidence of the paradox-of-thrift, namely that a rise in the saving rate can lead to a decrease in total savings as output contracts. This corroborates the narrative of the financial crisis being ultimately the response to a negative shock to the demands for goods all across the board and underlines the claim by Christiano et al. (2017) to reconsider the theory of the paradox-of-thrift in modern macroeconomics. Thirdly, combining long- and short-run restrictions we find evidence that saving shocks can have a negative impact orthogonal to uncertainty shocks. In fact, uncertainty shocks appear to be transmitted almost entirely via changes in the saving rate.

Literature. The link between uncertainty surrounding the economic climate and precautionary savings is well documented in theoretical and empirical contributions to the

[3] Guerrieri & Lorenzoni (2017) show that an increase in precautionary savings can generate an output drop even if prices are flexible in a model with incomplete markets and borrowing constraints. Under sticky prices, the output drop will be larger if the economy is in a liquidity trap.

literature[4]. Greater uncertainty increases the incentive for households to save in order to protect themselves against the higher likelihood of adverse economic outcomes in the future. Theoretical contributions on the topic of uncertainty and precautionary savings include Leland (1978), Skinner (1988), Zeldes (1989), Deaton (1989), Caballero (1991) and Carroll et al. (1992) among others. Empirical work studying the importance of precautionary savings at the household level include Carroll (1997), Engen & Gruber (2001) and Gourinchas & Parker (2002). The effect of uncertainty on business cycles is of more recent interest in response to the financial crisis. Gourio (2012) for instance shows that an increase in uncertainty modeled as a rise in the perceived probability of disaster leads to a collapse of investment and output. When inferring the probability of disaster from observed asset prices, the variation in this probability measure can explain a significant fraction of business cycle dynamics, especially for the duration of the financial crisis. Related to this, Christiano et al. (2014) find that fluctuations in risk are the most important shock driving business cycle fluctuations. Furthermore, Basu & Bundick (2017) argue that higher uncertainty has even more negative effects if monetary policy can no longer perform its usual stabilizing function because of the zero lower bound and consequently, increased uncertainty about the future played a key role in worsening the financial crisis. In a recent paper, Bloom et al. (2018) find theoretical evidence that uncertainty shocks can generate drops in gross domestic product of up to 2.5%.

Outlook. The rest of the paper is structured as follows. First, we present a version of the quality ladder model by Aghion & Howitt (1992) with sticky prices. An exogenous growth model is nested if we shut down the R&D channel. Section 3.3 introduces the empirical approach used to identify saving rate shocks and their effect on output in the data. Results for U.S. time series data are presented in section 3.4 and section 3.5 concludes.

3.2 Theoretical model

In this section we present the equilibrium conditions of a growth model with sticky prices that are utilized to derive the long-run restrictions imposed in the empirical model. We consider the simplest version of a Schumpeterian growth model where endogenous growth is induced through improvements in the quality of goods as introduced by Aghion & Howitt (1992) and Grossman & Helpman (1991). The main idea of these so-called quality-ladder models is that refinements of existing intermediate goods and techniques

[4] See Lugilde et al. (2019) for a comprehensive review on the empirical literature on precautionary savings.

increase their productivity in producing the final output. Quality improvements occur stochastically and at different rates across intermediate good sectors. An important feature of these models is that, when a product is improved, it tends to displace the old one given its higher productivity and substitutability in the production process of the final good. Hence, successful research along the quality dimension eliminates the monopoly rents of the preceding intermediate product. This process of endogenous growth is commonly referred to as "creative destruction". An exogenous version of the growth model is nested when shutting down the option to invest into R&D.

The model departs from the textbook quality ladder model as laid out in Barro & Martin (2004) by assuming that prices are sticky. This allows us to identify short run effects on output following a shock in the households' MPS.

3.2.1 Setup

Household. There is an infinitely lived representative household which aims to maximise its expected present value of utility over its lifetime using a discount factor β. The household obtains utility from consumption according to CRRA preferences and disutility from supplying labour which are additively separable

$$U_t = \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\gamma} - 1}{1-\gamma} - \psi \frac{L_t^{1+\eta}}{1+\eta} \right). \tag{3.1}$$

The household maximizes equation (1) subject to its budget constraint choosing optimal consumption, bond purchases and how much labour to supply

$$P_t C_t + B_t = P_t w_t L_t + r_t B_{t-1} + \Pi_t. \tag{3.2}$$

First order conditions with respect to consumption and bond holdings lead to the intertemporal Euler equation of the household

$$U_{C_t} = \beta \mathbb{E}_t \left[r_{t+1} \frac{P_t}{P_{t+1}} U_{C_{t+1}} \right]. \tag{3.3}$$

Shocks in the marginal propensity to save are modelled as shocks to the discount factor β. Following a rise in β it is optimal for the household to reduce its contemporaneous consumption in order to save for future periods. The first order condition with respect to labour leads to the intratemporal labour choice

$$\psi L_t^{\eta} = C_t^{-\gamma} w_t. \tag{3.4}$$

Final goods sector. In the perfectly competitive final goods sector, the final good Y is produced using technology A, labour L and a continuum of varieties of imperfectly substitutable intermediate goods. Intermediate goods are augmented by their quality rung q^{k_i} and combined according to a constant elasticity of substitution (CES) aggregator.

$$Y_t = A_t L_t^{1/\sigma} \int_{i=0}^{1} \left(q^{k_{i,t}} y_{i,t} \right)^{\frac{\sigma-1}{\sigma}} di, \tag{3.5}$$

where $y_{i,t}$ is the amount of sector i's intermediate input used at the final production stage, and q^{k_i} denotes its quality rung. Elasticity of substitution between intermediate goods σ is larger than unity to allow for endogenous growth through quality improvements. Increases in the level of technology A can be viewed as a source of exogenous growth. For the growth accounting exercise further below it is convenient to define the economy's aggregate quality index as the weighted sum of quality levels associated with each sector's intermediate input

$$Q_t = \int_{i=0}^{1} \left(q^{k_{i,t}} \right)^{\sigma-1} di. \tag{3.6}$$

Growing Q_t through successful research raises the quality of final consumption as much as it raises its quantity. From the representative firm's profits which are given as follows

$$\Pi_t^Y = P_t Y_t - \int_0^1 p_{i,t} y_{i,t} di - w_t P_t L_t, \tag{3.7}$$

we can derive the firm's demand for labour and the optimal production of the intermediate good y_i, respectively,

$$\frac{1}{\sigma} \frac{Y_t}{L_t} = w_t, \tag{3.8}$$

and

$$y_{i,t} = \left(\frac{p_{i,t}}{P_t} \right)^{-\sigma} \left(\frac{\sigma-1}{\sigma} A_t \right)^{\sigma} L_t q^{(\sigma-1)k_{i,t}}. \tag{3.9}$$

Intermediate goods producers. There is a continuum of intermediate firms producing the intermediate input using output as the sole factor of production. They produce according to the latest technology acquired from the research sector. Innovations are assumed to be drastic such that the intermediate monopolist is unconstrained by potential competition from previous patents. As intermediate goods are imperfect substitutes for the final good, producers compete in an environment of monopolistic competition and choose their optimal price taking into account the sector's demand function. However, price setting comes with a quadratic adjustment cost following Rotemberg (1982). Due to nominal rigidities the firms face an intertemporal pricing problem

$$\max_{p_{i,t+s}} \mathbb{E}_t \left[\sum_{s=0}^{\infty} m_{t,t+s} \left(p_{i,t+s} - P_{t+s} - \frac{\theta}{2} \left(\frac{p_{i,t+s}}{p_{i,t+s+1}} - 1 \right)^2 p_{i,t+s} \right) y_{i,t+s} \right], \tag{3.10}$$

where $m_{t,t+s}$ is a stochastic discount factor capturing the general discount factor β, the marginal rate of substitution between future and present consumption and the survival rate of the firm. The survival rate is defined as $(1-\mu)$, where μ is the probability of a jump on the quality ladder, hence the likelihood for creative destruction in this sector. The pricing problem is symmetric across sectors as the optimal price and the likelihood for being replaced is independent of the rung on the quality ladder as we will see below. This results in the following mark-up expression for the optimal price

$$p_{i,t}^{*} = \psi_t P_t, \tag{3.11}$$

with

$$\psi_t = \frac{\sigma}{(\sigma-1)(1-\frac{\theta}{2}(\pi_t-1)^2) + \theta(\pi_t-1)\pi_t - \theta\mathbb{E}_t\left[m_{t,t+1}(\pi_{t+1})^2(\pi_{t+1}-1)\frac{y_{i,t+1}}{y_{i,t}}\right]}, \tag{3.12}$$

where π is the inflation rate. Normalizing the aggregate price level to unity we have that intermediate goods' prices equal the markup. For flexible prices we get the standard markup expression in terms of the elasticity of substitution between intermediate input, $\psi_t = \frac{\sigma}{\sigma-1}$.

The market value of the intermediate firm using the latest patent is defined as the discounted expected future profits, which can be written recursively as follows

$$V_{k_i} = \left(p_{k_i,t} - P_t - \frac{\theta}{2}(\pi_t - 1)^2 p_{k_i,t} \right) y_{k_i,t} + m_{t,t+1} V'_{k_i} \tag{3.13}$$

$$= \left(\psi_t - (1 + \frac{\theta}{2}(\pi_t-1)^2) \right) y_{k_i,t} + m_{t,t+1} V'_{k_i}. \tag{3.14}$$

R&D and patents. For a given quality rung k_i the probability of success depends on the amount of research resources $z_{i,t}$ in form of the final good used

$$\mu_{k_{i,t}} = \phi(k_{i,t}) z_{i,t}, \tag{3.15}$$

where $\phi(k_{i,t})$ captures the fact that for any given level of research, the probability of success decreases in the number of innovations already made in this sector (thus $\phi'(k_{i,t}) < 0$). Following standard specifications of $\phi(k_{i,t})$ in the literature, we define it as follows

$$\phi(k_{i,t}) = \frac{1}{\lambda}(q^{k_i+1})^{(1-\sigma)}, \tag{3.16}$$

where $1/\lambda$ is the productivity of goods invested in research and development. Note that the specification is chosen such as to offset the positive effect of a sector's position on the quality ladder on profits. Recall, that instantaneous profits depend linearly on q^{k_i+1} via $y_{k_{i,t}}$.

Given that intermediate firms make profits, there will be competition for the latest patents. Ultimately, a single firm is willing to pay the equivalent to its present value at market entrance. Free entry into the R&D sector leads to a zero profit condition

$$\phi(k_{i,t}) V_{k_{i,t}} - P_t = 0, \tag{3.17}$$

which implies that expected profits from research per researcher equal the wage. The zero profit condition pins down the optimal level of research and hence the probability of success μ which turns out to be constant over time. Using the expression of a firm's present value producing on the latest quality rung, we can see that the free entry condition is independent of the specific sector as q^{k_i+1} drops out

$$\frac{1}{\lambda}(q^{k_i+1})^{(1-\sigma)}\mathbb{E}_t \sum_{s=0}^{\infty} m_{t,t+s}\left(p_{i,t+s} - P_{t+s} - \frac{\theta}{2}(\pi_{t+s}-1)^2 p_{i,t+s}\right) \tag{3.18}$$

$$\left(\frac{p_{i,t}}{P_t}\right)^{-\sigma}\left(\frac{\sigma-1}{\sigma}A\right)^{\sigma} L_t q^{(\sigma-1)(k_i+1)} - P_t, \tag{3.19}$$

$$= \frac{1}{\lambda}\mathbb{E}_t \sum_{s=0}^{\infty} m_{t,t+s}\left(p_{i,t+s} - P_{t+s} - \frac{\theta}{2}(\pi_{t+s}-1)^2 P_{i,t+s}\right) \tag{3.20}$$

$$\left(\frac{p_{i,t}}{P_t}\right)^{-\sigma}\left(\frac{\sigma-1}{\sigma}A\right)^{\sigma} L_t - P_t = 0. \tag{3.21}$$

Monetary authority. Lastly, there is a central bank setting the nominal interest rate following a Taylor rule. The central bank takes into account the deviation from price stability as captured by inflation and the deviation of output from its potential as captured by its steady-state measure:

$$\log(R_t/R_{ss}) = \rho_R \log(R_{t-1}/R_{ss}) + (1-\rho_R)[\theta_\pi \log(\pi_t/\pi_{ss}) + \theta_Y \log(Y_t/Y_{ss})], \tag{3.22}$$

where π_t is gross inflation of the final good price, θ_π is the weight on inflation in the reaction function. Analogously, θ_Y is the weight on the output gap, measured as deviation from steady state output, Y_{ss}. ρ_R is a smoothing parameter.

3.2.2 Model predictions

Table 3.1 Non-estimated parameters and calibrated values

Parameter	Value	Description	Source
		Panel A: Parameters	
β	0.994	Discount factor	Basu & Bundick (2017)
γ	1	Risk aversion	Log utility
η	2	Frisch elasticity	Basu & Bundick (2017)
ψ	0.95	Weight on leisure preference	Basu & Bundick (2017)
ρ_r	0.8	Taylor rule smoothing parameter	Standard
θ_π	1.5	Taylor rule inflation coefficient	Basu & Bundick (2017)
θ_y	0.2	Taylor rule output gap coefficient	Basu & Bundick (2017)
θ_p	100	Rotemberg coefficient of price adjustment cost	Basu & Bundick (2017)
A	1	Economy wide TFP	Normalized
Q	1	Quality index	Normalized
σ	6	CES parameter	Basu & Bundick (2017)
		Panel B: Steady state values	
g_Y	0.5%	Quarterly growth rate of final good	
Y	1	Output of final good	
ψ	1.2	Markup	

Both, the endogenous and exogenous growth model are calibrated such that they result in the same quarterly growth rate of 0.5%. The discount factor is chosen such that it corresponds to a quarterly interest rate of 0.5% which roughly amounts to a 2% annual interest rate. The only shock in this model is a shock to the saving behaviour of households. A shock in the MPS is modeled as a one percent shock to the discount factor. The constant technological progress A in case all growth comes through quality refinements and the constant quality index Q in the case of exogenous growth are set to one such that final output is the same in both models for the same amount of labour utilized. The constant elasticity of substitution between intermediate inputs is chosen in order to obtain a steady state markup of 20%. Based on this markup, the Rotemberg (1982) coefficient of price adjustment is chosen to mimic a Calvo frequency re-optimizing firms of around 25% each quarter following Keen & Wang (2007)[5]. In the following analysis, this is compared to a model with a flexible-price calibration ($\theta_p = 0$).

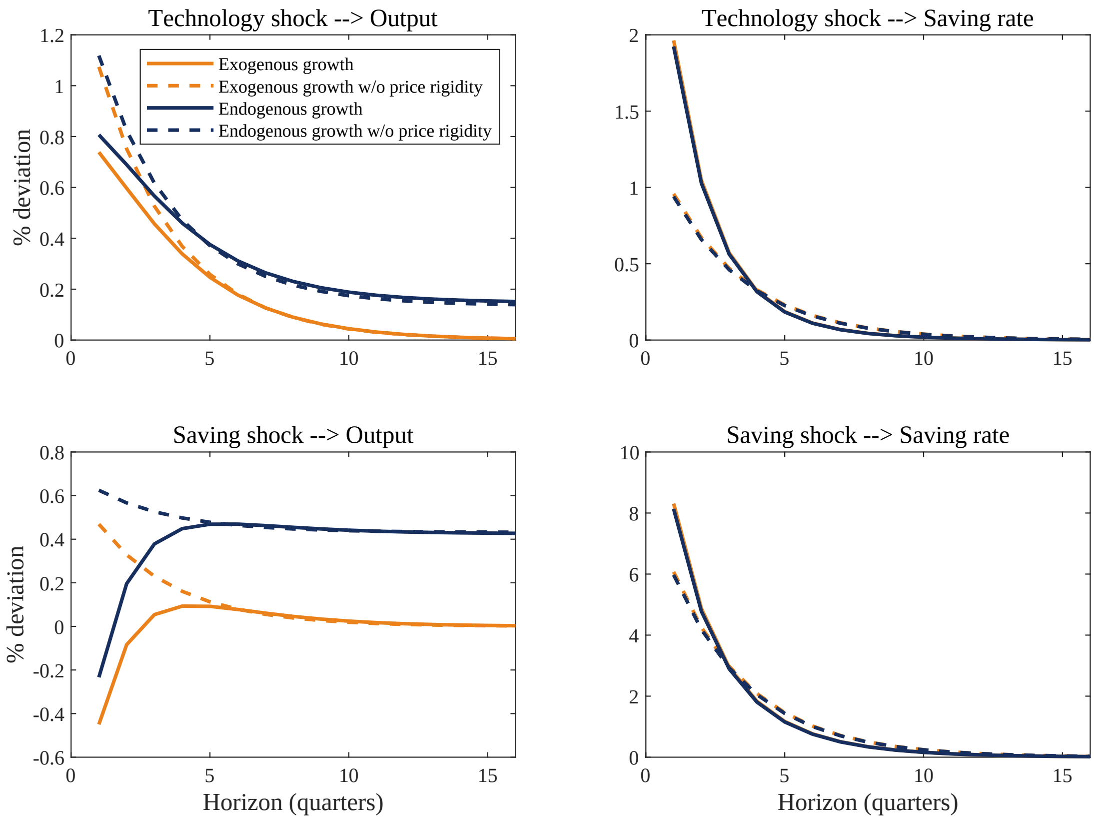

Fig. 3.2 Impulse response functions for the exogenous and endogenous growth model.

Notes.

Figure 3.2 plots the impulse responses of the endogenous and exogenous growth model for both, the sticky-price and flexible-price calibration in case of a temporary shock[6]. Following a productivity shock, output increases for all specifications. In the

[5]This is also broadly in line with the calibration in Basu & Bundick (2017) who use a slightly higher quadratic adjustment cost due to their inflation target of 2%.

[6]Impulse responses for a permanent shock to the discount factor are shown in figure C.3 in the appendix

presence of price stickiness the response is lower in magnitude, due to the quadratic adjustment cost faced by intermediate good producers. While, for the exogenous growth model, output goes back to its steady state level, extra savings lead to more R&D activity in the endogenous growth model, ultimately resulting in more growth in the quality index and hence output. In case of a shock to the discount factor, the response depends on whether prices are flexible or rigid. In a flexible-price regime, prices adjust to compensate for the lack in demand and higher savings are increasing output, both in the exogenous and endogenous growth model. If prices are sticky however, the demand shock has a contractionary effect on output. Firms face lower demand, but are unable to flexibly change their prices, ultimately having to reduce production to still maximize profits. Eventually, output will go back to its steady state level in the exogenous growth model, while it reaches a higher steady state in the endogenous growth model due to innovations in the R&D sector.

Those model predictions allow us to derive several identification restrictions for our empirical strategy to estimate the effect of saving shocks on the business cycle. These are lined out in the next section.

3.3 Empirical strategy

This sections introduces the bivariate[7] structural VAR specifications and long-run exclusion restrictions used to identify saving shocks in the data.[8] The empirical model interprets the variations in differences in (log) output, Δy_t, and the saving rate, sr_t, as originating from two types of exogenous disturbances, namely saving shocks and technology shocks which are by definition orthogonal to each other. The propagation mechanisms of those shocks over time are unspecified. We can formalize this idea by writing it as a structural vector autoregression (SVAR) model:

$$\boldsymbol{A}\begin{bmatrix} \Delta y_t \\ sr_t \end{bmatrix} = \boldsymbol{B}(L)\begin{bmatrix} \Delta y_{t-1} \\ sr_{t-1} \end{bmatrix} + \begin{bmatrix} \epsilon_t^y \\ \epsilon_t^s \end{bmatrix}, \tag{3.23}$$

where ϵ_t^y and ϵ_t^s denote structural technology shocks and saving shocks, respectively. The orthogonality assumption combined with the normalization that $\epsilon_t = (\epsilon_t^y, \epsilon_t^s)' \sim (0, I_2)$

[7]We also considered a trivariate VAR including the federal funds rate to disentangle MPS shocks from saving shocks induced by monetary policy, yet results were very similar to the bivariate VAR.

[8]For a detailed explanation of SVARs and long-run restrictions please refer to Kilian & Lütkepohl (2017).

implies that $\mathbb{E}[\epsilon_t \epsilon_t'] = I$. Given equation 3.23 and assuming invertability, the reduced form VAR is given by

$$\begin{bmatrix} \Delta y_t \\ sr_t \end{bmatrix} = \boldsymbol{F}(L) \begin{bmatrix} \Delta y_{t-1} \\ sr_{t-1} \end{bmatrix} + \boldsymbol{Q} \begin{bmatrix} \epsilon_t^y \\ \epsilon_t^s \end{bmatrix}, \tag{3.24}$$

where $\boldsymbol{F}(L) = \boldsymbol{A}^{-1}\boldsymbol{B}(L)$, and $\boldsymbol{Q} = \boldsymbol{A}^{-1}$. Using the mapping from structural shocks to reduced form errors

$$u_t = \boldsymbol{Q} \begin{bmatrix} \epsilon_t^y \\ \epsilon_t^s \end{bmatrix}, \tag{3.25}$$

we can simplify the reduced form VAR to

$$Y_t = \boldsymbol{F}(L) Y_{t-1} + u_t. \tag{3.26}$$

Informed by the model predictions, we derive several long-run restrictions for saving shocks, while technology shocks can have a permanent effect as in Gali (1999). The long-run exclusion restrictions are imposed on the moving average representation of the SVAR model, which is given by

$$Y_\infty = (I - F(L))^{-1} Q \begin{bmatrix} \epsilon_t^y \\ \epsilon_t^s \end{bmatrix} = \Theta \begin{bmatrix} \epsilon_t^y \\ \epsilon_t^s \end{bmatrix}, \tag{3.27}$$

where Θ is the long-run impact matrix.

In order to identify the structural shocks from the residuals of reduced form estimations of the system of variables, we impose the long-run restrictions implied by the theoretical model following the identification approach originally proposed by Blanchard & Quah (1989). They show that even if the economy is subject to multiple demand shocks, the identification conditions are satisfied as long as only the supply shock has lasting effects on output and the other demand disturbances do not alter the relationship between output and the disturbance of interest, here saving shocks. We argue that this is the case as sudden changes to consumption choices by households are at the very core of the definition of demand disturbances.

Identification restrictions. Based on the theoretical results of temporary and permanent saving shocks in an exogenous and endogenous growth model as shown in the appendix, we derive three exclusion restrictions[9].

Exclusion Restriction 1 *In an exogenous growth model, a temporary shock to the marginal propensity to save has no long-term effect on output.*

In this case $Y_t = [\Delta y_t, sr_t]'$ and the effect of any of the structural shocks on Y_t will approach zero as the horizon increases. The effect of a structural shock on the level of real GDP y_t is the cumulative sum of its effects on Δy_t. The long-run cumulative effects are given by the matrix Θ. Hence, the first exclusion restriction corresponds to $\Theta_{12} = 0$. In other words, the restriction demands that a saving shock is such that when cumulating the impulse response function of growth, there will be eventually no effect on real GDP. The other entries of the long-run effect matrix remain unrestricted as supply shocks such as technology shocks affect the level of real GDP in the long run. Moreover, there are no restrictions on the second row of Θ because the cumulative responses of a stationary series are different from zero.

Exclusion Restriction 2 *In an endogenous growth model, a temporary shock to the marginal propensity to save has no long-term effect on output growth.*

In an endogenous growth model, a reduction in contemporaneous consumption in combination with sticky prices will act as a demand shock to output but will also increase investment into R&D as output decreases by less than consumption. Higher investment in R&D will lead to a higher probability of research success and subsequently a temporary rise in TFP growth. Therefore, under the assumption of an endogenous growth model, a temporary saving shock can act as a supply shock, leading to a permanent change in the level of output in the long-run. However, it has no impact on the long-term growth rate. Implementing the second exclusion restriction, the vector of stationary processes is defined as $Y_t = [\Delta^2 y_t, sr_t]'$. The restriction to the long-run cumulative effects matrix is the same as above, i.e. $\Theta_{12} = 0$ such that a saving shock has no long-run effect on the growth rate of output.

Exclusion Restriction 3 *In an endogenous growth model, a permanent shock to the marginal propensity to save has no long-term effect on the change in output growth.*

[9] Please note that we do not consider the restriction of a permanent shock to the saving rate having no long-run effect on growth as this is true by assumption in an exogenous growth model and is not an outcome of the model.

A permanent saving shock increases the steady state investment in R&D and thereby fosters higher TFP growth in the long-run. While the growth rate will increase to reach its higher steady state, it is constant in the long-run. Hence, in case of an endogenous growth model, a permanent positive saving shock leads to a rise in the growth rate but no change in the slope of the growth rate thereafter. This leads us to the third exclusion restriction, where we define the vector of stationary processes as $Y_t = [\Delta^3 y_t, \Delta sr_t]'$. Here, we assume the saving rate to be I(1) such that a temporary increase in the first difference leads to a permanent increase in the level of the saving rate. The restriction to the long-run cumulative effects matrix is the same as above, i.e. $\Theta_{12} = 0$ such that a saving shock has no long-run effect on the slope of the growth rate.

Reduced-form specification. Consistent estimates of the coefficients of $F(L)$ are obtained as functions of the estimated parameters of a reduced-form VAR. The baseline specification of the reduced-form estimation is as follows

$$Y_t = c_t + \sum_{i=1}^{p} \beta_i Y_{t-i} + u_t, \tag{3.28}$$

where Y_t is defined according to the specific exclusion restriction utilized, c_t is a (2×2) matrix containing a constant term and a linear trend.[10] β_i is a (2×1) vector of parameter estimates for the i-th lag of Y_t with a maximum of p lags. For our baseline specification we use eight lags, i.e. two years for quarterly data, but consider alternative choices as robustness checks. The last term u_t denotes the residuals of the reduced-form estimation.

From the reduced form estimates we obtain the structural estimates by following the standard implementation of long-run restrictions. First, we calculate the companion form of the beta estimates F, then we find the restricted long-run impact matrix, D, from a Cholesky decomposition of $LR \sum LR'$, where $\sum$ is the covariance matrix of the residuals and LR is a matrix containing the first two row and column entries of $(I-F)^{-1}$. Using the long-run impact matrix, we can back out the contemporaneous impact matrix as $Q = (I-F)D$.

[10] We also included dummies for several potential structural breaks but they did not alter the results in any meaningful way.

3.4 Empirical results and robustness

3.4.1 Baseline results

In this section, we present the results of the SVAR models using quarterly U.S. time series data from 1960-Q1 until 2019-Q4. All data comes from the Federal Reserve Economic Database (FRED). Growth is defined as the seasonally adjusted percentage change in real gross domestic product from the preceding period. We use two measures for the saving rate. The first is the personal saving rate as the ratio between personal savings and disposable income of households. The second measures the aggregate saving rate as the ratio between personal savings and the real gross domestic product. Just by the definition of the saving rate, the simultaneity issue becomes apparent as output directly enters the definition of the saving rate. Hence, any short-run identification strategy would be flawed. Long-run restrictions, on the other hand, aim to disentangle this simultaneity by differentiating between technology shocks shifting output permanently and temporary demand shifts in the form of saving shocks. Using the identified temporary response in output after an increase in the desire to save, the aggregate definition allows us to back out the response in the level of personal savings to a saving shock.

Figures 3.3 to 3.5 show the resulting impulse response functions when estimating the baseline specification with eight lags for the first two long-run exclusion restrictions over the sample from 1960-Q1 to 2019-Q4.[11] These results are obtained for aggregate saving rate. Estimates using the personal saving rate are very similar and are reported in the appendix. Imposing the first exclusion restriction, i.e. that a temporary shock in the saving rate cannot have a long-run effect on output, results in a negative short-term impact on economic activity as depcited in figure 3.3. Hence, a sudden cut back in consumption appears to have a contractionary effect which underlines the hypothesis that increased savings made the financial crisis more severe. Only after roughly two years, economic activity is picking up again slowly approaching the level at which it would be in the absence of a saving shock. Interestingly, when computing the response in total personal savings as the sum of the cumulated impulse response of growth in real GDP and the impulse response in the saving rate, we obtain the paradox-of-thrift result as shown in Figure 3.4. Although, households increase their saving rate, the level of savings decreases as the fall in real GDP dominates the increase in the saving rate.

[11] Plots for exclusion restriction 3 are in the appendix. Given that this excl. restriction assumes GDP to be an I(3) process we refrain from giving the result much weight.

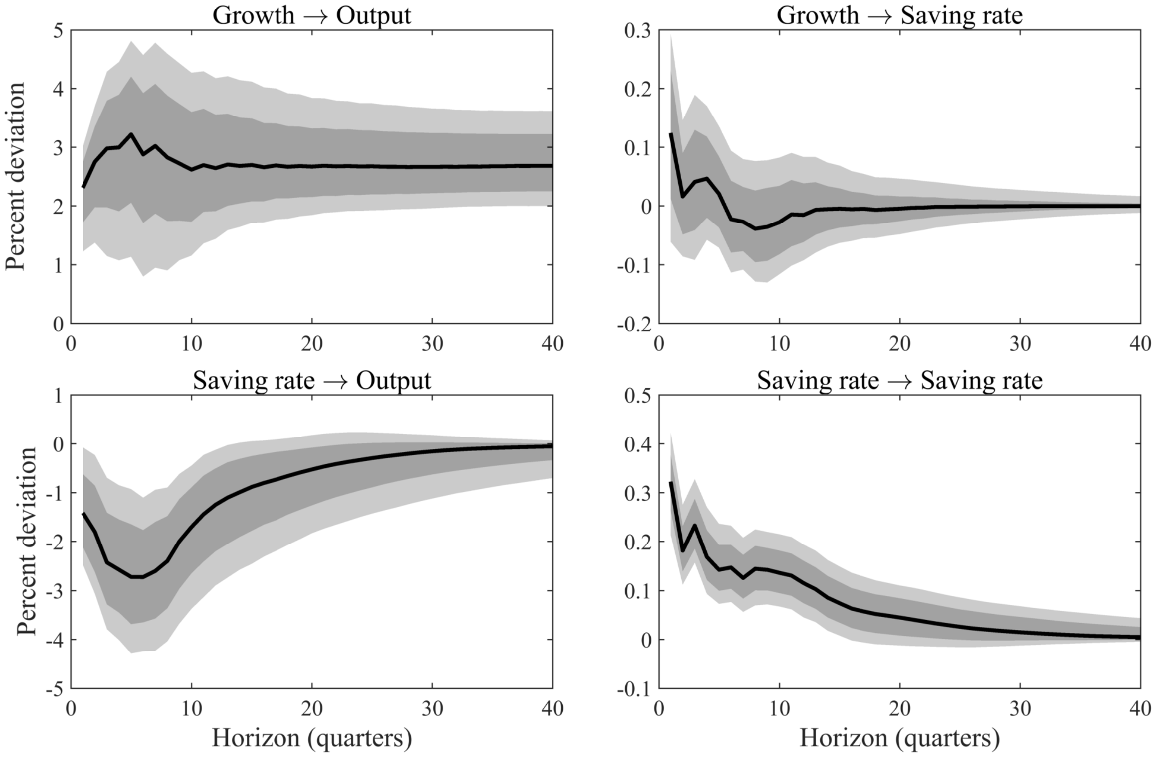

Fig. 3.3 Impulse response functions using exclusion restriction 1 that a temporary shock to the saving rate has no long-run effect on output.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The grey transparent area shows the corresponding 68% and 90% confidence bounds.

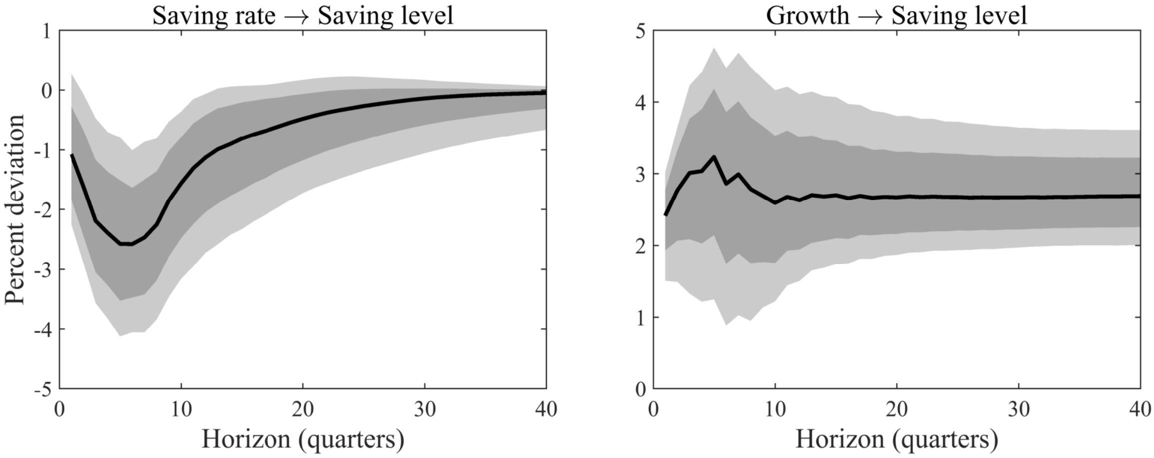

Fig. 3.4 Impulse response function of the level of private savings according to the exclusion restriction in Figure 1.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The grey transparent area shows the corresponding 68% and 90% confidence bounds.

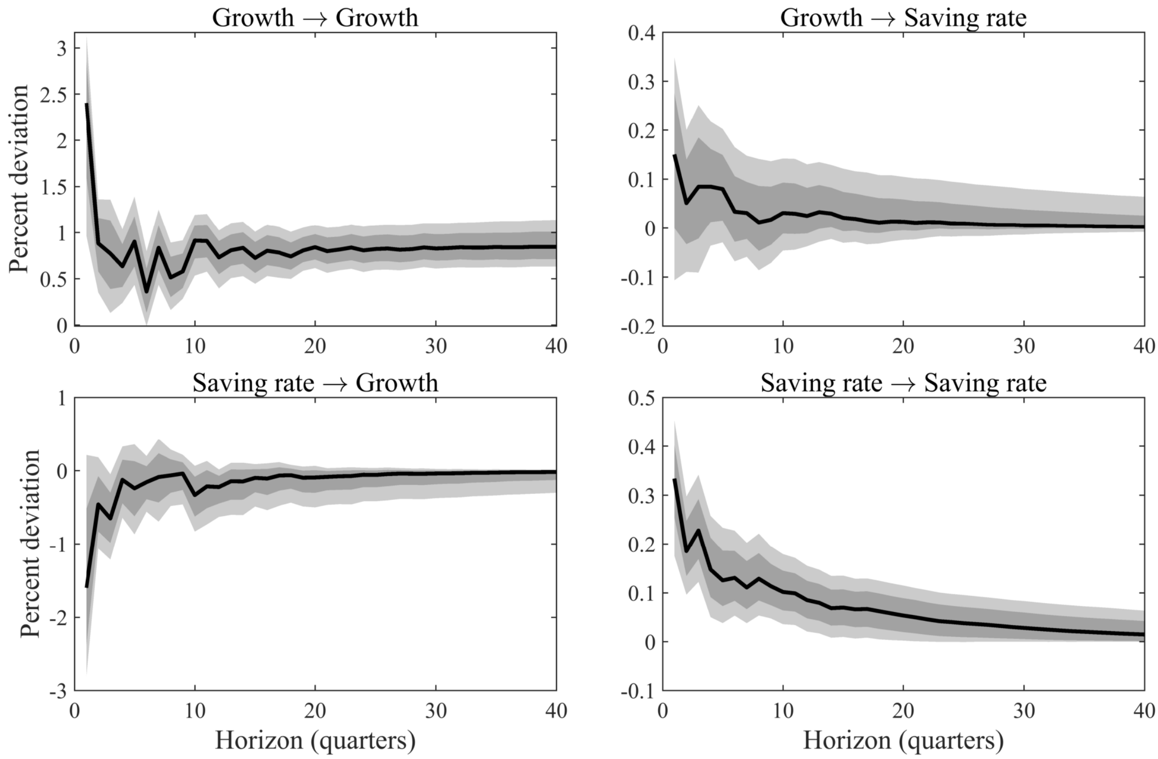

Fig. 3.5 Impulse response functions using exclusion restriction 2 that a temporary shock to the saving rate has no long-run effect on the growth of output.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The grey transparent area shows the corresponding 68% and 90% confidence bounds.

Imposing the second exclusion restriction that a temporary saving shock does not have any long-term effect on growth leads to a similar picture as depicted in Figure 3.5. A positive saving shock is followed by a drop in output growth. By definition of the second exclusion restriction, the initial drop reverts back such that growth in real GDP approaches its pre-shock level in the long-run. Furthermore, the identified temporary positive saving shock leads to a permanent decrease in the level of savings which is in line with the paradox-of-thrift argument above. This is due to the fact that while the saving rate returns to its initial level, real GDP will be permanently lower, leading to less total personal savings in the long run.

The third long-run restriction results in a permanently lower level of real GDP growth following a permanent positive saving shock as shown in Figure C.8. However, this identification restriction hinges on real GDP being an I(3) process and the saving rate being of order I(1). Under this assumption, the level of savings is decreasing over time as real GDP is growing slower than its counter-factual in the absence of a saving shock.

3.4.2 Uncertainty and saving shocks

So far, we have identified that shocks to the saving rate with no long-run impact on output or growth, have a contractionary effect on output in the short-term. However, we might be missing a link which is well established in the literature. Uncertainty could act as a latent mediator, driving both, the saving rate up due to precautionary motives and contracting output at the same time. Hence, the question arises, whether we have truly identified the effect of saving rate shocks or rather uncertainty shocks.

In order to tackle this question, we combine long-run and short-run exclusion restrictions in a three-variate SVAR with output, saving rate and proxies for uncertainty. As before, using exclusion restriction 1, we assume that a saving rate shock has no long-term impact on output. An uncertainty shock is defined as having no-long term impact on output but can have an immediate impact on both, output and the saving rate. In order for the system to be just identified, we further assume that saving rate shocks have no contemporaneous impact on uncertainty. The short-run exclusion restriction are in line with the identification strategy proposed by Basu & Bundick (2017). Following the notation by Rubio-Ramirez et al. (2010), equation 3.29 illustrates the identification assumptions imposed on the contemporaneous and long-run impact matrix Q and Θ, respectively.

$$f(Q,\mathbf{F}) = \begin{bmatrix} Q \\ \Theta \end{bmatrix} = \begin{array}{c} \\ \Delta log Y \\ U \\ S \\ \\ \Delta log Y \\ U \\ S \end{array} \begin{array}{c} \begin{array}{ccc} \epsilon^y & \epsilon^u & \epsilon^s \end{array} \\ \left[\begin{array}{ccc} x & x & x \\ x & x & 0 \\ x & x & x \\ \hdashline x & 0 & 0 \\ x & x & x \\ x & x & x \end{array}\right] \end{array} \tag{3.29}$$

As a measure for uncertainty we use the historical newspaper-based economic uncertainty index by Baker et al. (2016) as this spans almost our entire data sample. As a robustness check we also consider the Equity Market Volatility Index (EMV) by Baker et al. (2019) starting in 1985 and the Exchange Volatility Index (VXO) starting in 1986.

Figure 3.6 plots the estimated responses of each shock on output and the saving rate along with 68% and 90% confidence intervals. On impact, a saving rate shock which is now orthogonal to the identified uncertainty shock, still has a contractionary effect on output with a confidence interval of one standard deviation. The effect of uncertainty on output is hardly significant at the 68% level. In fact, when comparing the responses to a saving rate shock and an orthogonal uncertainty shock, the latter looks relatively similar just lower in magnitude.

To test whether uncertainty shocks are primarily propagated through changes in the saving rate, we mute the response in the saving rate. Specifically, we counter the uncertainty shock with saving rate shocks and thereby offsetting the response in the saving rate. As plotted in the last column of figure 3.6, when muting the saving rate channel, the effect of an uncertainty shock on output becomes insignificant. Hence, it seems that while saving shocks have a contractionary effect independent of uncertainty shocks, uncertainty shocks are mainly transmitted via changes in the saving rate due to precautionary motives.

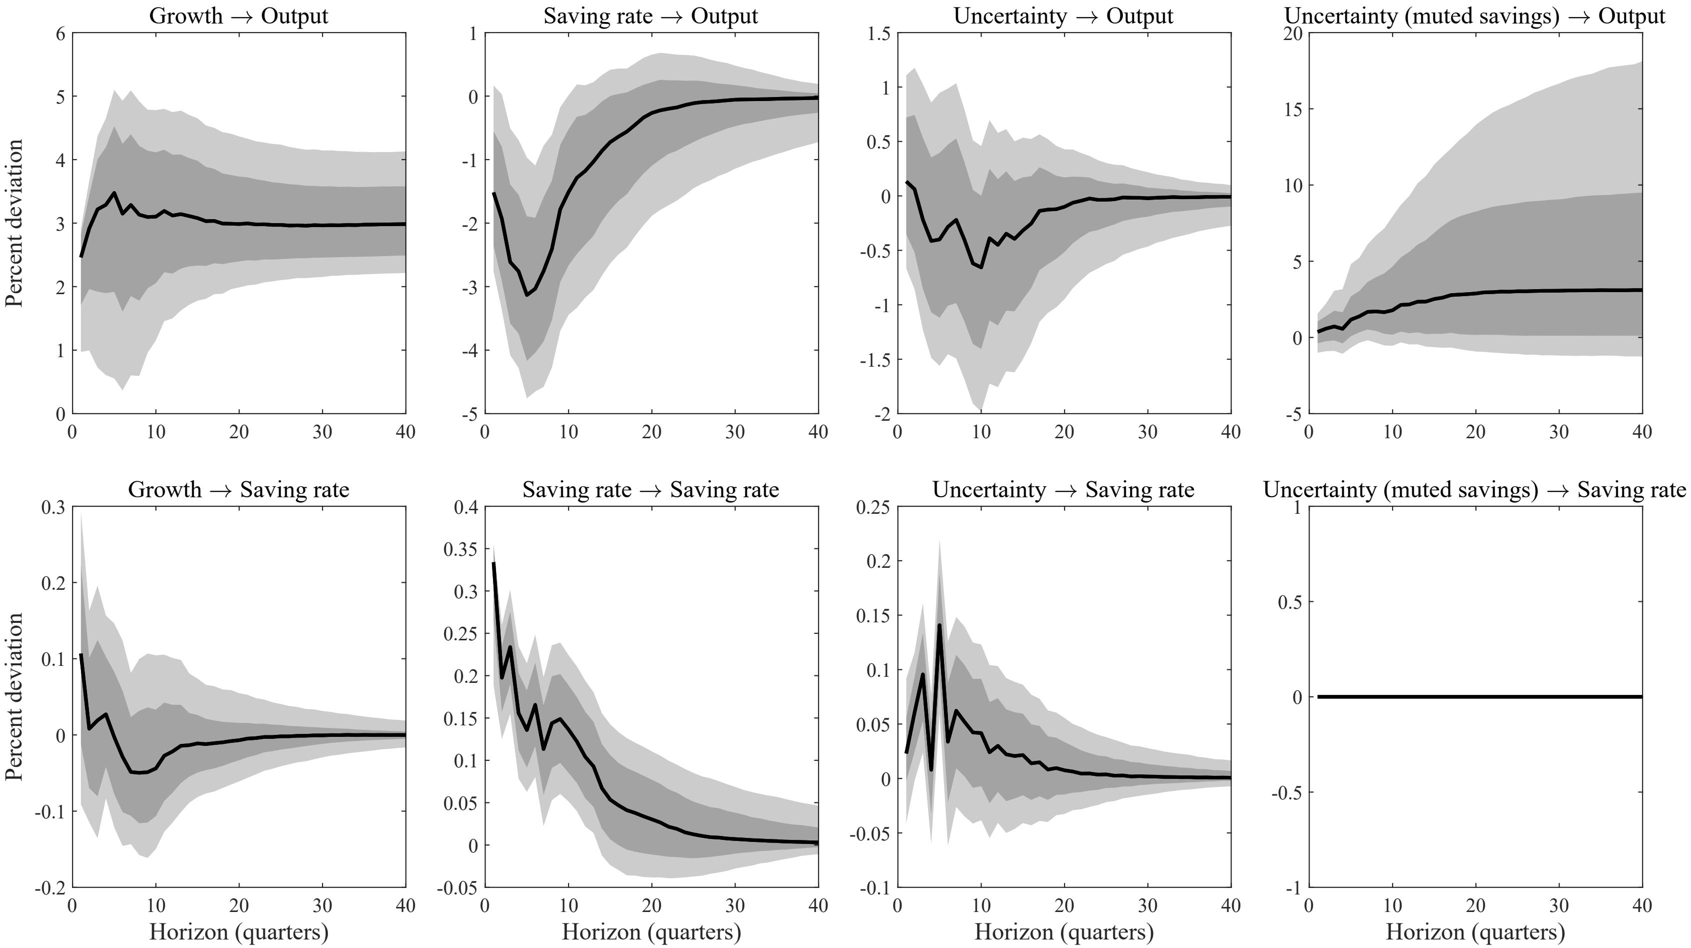

Fig. 3.6 Impulse response functions using long- and short-run exclusion restrictions for uncertainty and saving shocks.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The dark grey transparent area shows the corresponding 68% and 90% confidence bounds.

3.4.3 Robustness checks

Long run identification. Under the baseline specification, the finding that a positive saving shock can have a contractionary effect on economic activity appeared already robust for three different long run identification restrictions and two definitions of the saving rate, the aggregate saving rate and the personal saving rate in terms of disposable household income. The rest of this section examines the robustness of this finding for alternative choices of lags, trends and time periods. Table 3.2 lists the mean cumulative percentage impulse responses in real GDP after a one percent positive shock in the aggregate saving rate for alternative specifications of the respective structural VAR models[12]. The specification details are defined at the top of the table. The baseline specification uses eight lags, a linear time trend and the sample from 1960 to end of 2019, excluding the Covid crisis, and is included in the fourth column of table 3.2. Bootstrapped standard errors are reported in parentheses beneath the mean estimates.

Across all specifications we observe a negative cumulative effect in real GDP growth after eight quarters following a positive MPS shock. By definition of the exclusion restrictions, the cumulative response in real GDP growth is likely to be lower for the second and third long-run restriction as they do not impose any return of the output level in the long-run. For most of the cases, the negative response is significant for a confidence interval of at least one standard deviation. While reducing the lag length to four yields a weaker response, higher lag lengths in general reduce the estimation bias in VARs as shown by Plagborg-Møller & Wolf (2021). Using twelve lags, results are very close to the baseline estimates confirming that we capture enough flexibility by including eight lags. When choosing a quadratic trend instead of a linear trends, the response is only slightly weaker but still robust. When including the period of the Covid crisis until the end of 2020, results using exclusion restriction 1 are broadly in line, while exclusion restriction 2 predicts a much stronger response. Also when excluding the Financial Crisis or periods before 1990, results are broadly robust.

[12] Robustness checks for the saving rate definition in terms of disposable household income are in table C.1 in the appendix

Table 3.2 Cumulative percentage response in real GDP growth for alternative estimation specifications

		Lags			Trend		Time		
Identification	At horizon h	4	8	12	linear	quadratic	1960Q1-2020Q4	1960Q1-2006Q4	1990Q1-2019Q4
Exclusion restriction 1	1	-0.93	-1.50	-1.52	-1.50	-1.24	-0.97	-0.61	-1.23
		(0.70)	(0.73)	(0.69)	(0.73)	(0.74)	(1.06)	(1.00)	(0.44)
	2	-1.15	-1.89	-1.93	-1.89	-1.58	-2.27	-1.00	-1.45
		(0.83)	(0.85)	(0.80)	(0.85)	(0.85)	(0.81)	(1.13)	(0.56)
	4	-1.82	-2.61	-2.68	-2.61	-2.26	-2.60	-1.55	-2.00
		(0.91)	(0.92)	(0.88)	(0.92)	(0.90)	(0.96)	(1.21)	(0.62)
	8	-1.36	-2.38	-2.59	-2.38	-2.05	-2.18	-1.65	-1.63
		(0.79)	(0.91)	(0.86)	(0.91)	(0.85)	(1.08)	(1.07)	(0.68)
Exclusion restriction 2	1	-0.86	-1.39	-1.28	-1.39	-1.29	-0.89	-1.68	-0.94
		(0.78)	(0.79)	(0.76)	(0.79)	(0.75)	(0.84)	(0.91)	(0.51)
	2	-1.13	-1.80	-1.67	-1.80	-1.65	-2.60	-2.29	-1.12
		(1.09)	(1.06)	(0.99)	(1.06)	(1.01)	(0.80)	(1.68)	(0.71)
	4	-1.97	-2.62	-2.41	-2.62	-2.39	-4.12	-3.16	-1.69
		(1.63)	(1.48)	(1.24)	(1.48)	(1.39)	(1.19)	(1.68)	(0.95)
	8	-3.16	-3.06	-2.41	-3.06	-2.67	-6.325	-4.02	-1.53
		(2.62)	(2.12)	(1.86)	(2.12)	(1.98)	(1.82)	(2.56)	(1.38)
Exclusion restriction 3	1	-0.72	-0.67	-1.07	-0.67	-0.73	-2.58	-0.91	-0.25
		(0.34)	(0.42)	(0.48)	(0.42)	(0.41)	(1.26)	(0.53)	(0.41)
	2	-1.08	-0.88	-1.46	-0.88	-0.95	-1.96	-1.36	-0.16
		(0.55)	(0.64)	(0.68)	(0.64)	(0.62)	(1.46)	(0.77)	(0.64)
	4	-2.11	-1.64	-2.36	-1.64	-1.74	-2.34	-2.23	-0.65
		(0.95)	(1.05)	(1.11)	(1.05)	(1.01)	(2.69)	(1.27)	(1.01)
	8	-3.85	-2.60	-3.25	-2.60	-2.75	-4.80	-3.96	-0.63
		(1.77)	(1.85)	(1.92)	(1.85)	(1.77)	(5.44)	(2.34)	(1.67)

Note: Standard errors are parentheses; $'$ $p < 0.32$, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.3 Cumulative percentage response in real GDP growth for different uncertainty measures

		Uncertainty measure		
Response	Horizon	Historical Uncertainty Index	EVM Index	VXO Index
Saving rate → GDP	1	-1.09	-1.21	-1.06
		(0.72)	(0.82)	(0.74)
	2	-1.23	-1.42	-1.20
		(0.88)	(1.06)	(0.91)
	4	-1.93	-2.09	-1.89
		(1.00)	(1.24)	(1.05)
	8	-1.78	-1.90	-1.75
		(1.19)	(1.39)	(1.24)
Uncertainty → GDP (muted savings)	1	-0.19	-0.11	-0.16
		(0.70)	(0.96)	(0.71)
	2	-0.28	0.32	-0.25
		(0.94)	(1.40)	(0.94)
	4	0.62	0.65	0.62
		(1.69)	(1.78)	(1.66)
	8	2.78	1.30	2.74
		(3.10)	(1.92)	(3.16)

Note: Standard errors are parentheses; $'$ $p < 0.32$, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Long and short run identification. Combining long and short run exclusion restrictions, we find that saving shocks lead to a decline in output orthogonal to uncertainty shocks. This result appears robust for several different proxies for uncertainty as documented in table 3.3. Corresponding impulse response functions are reported in the appendix, see figures C.12 - C.15. Saving shocks are followed by a negative response in output across all uncertainty measures. In fact, even the magnitude of the response at different horizons is pretty robust for all proxies considered, although their data availability varies substantially.

Uncertainty shocks on the other hand have no significant effect on real GDP when shutting down any response in the saving rate suggesting that precautionary savings are the main transmission channel of uncertainty shocks.

SVAR-IV approach. As another robustness check, we use monetary policy surprises as an external instrument to identify the effect of structural saving shocks. We use the shock series by Gertler & Karadi (2015), who, using high-frequency financial data, obtain an external instrument for monetary policy. In estimating the SVAR-IV, we follow Plagborg-Møller & Wolf (2021) and order the external instrument first in an otherwise recursive

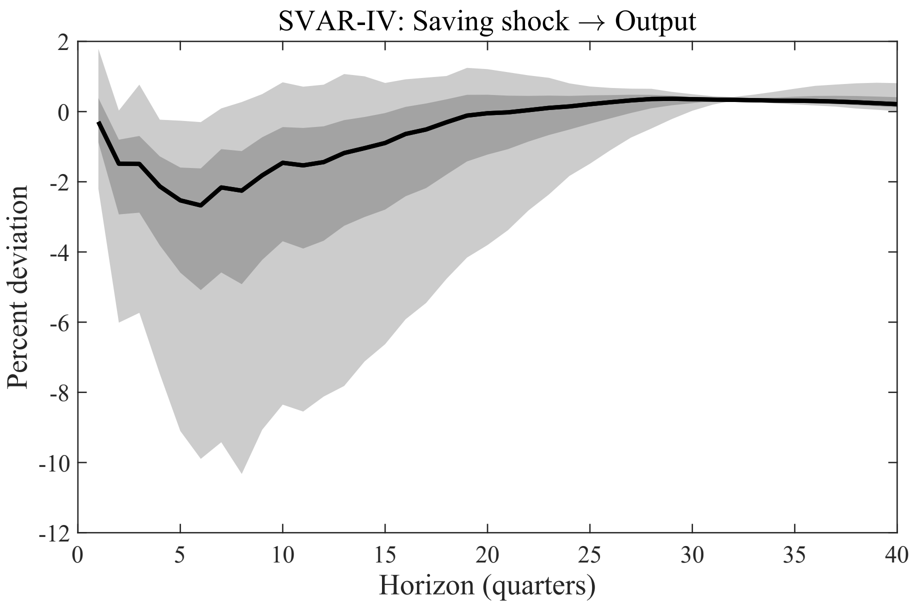

Fig. 3.7 Impulse response function of real GDP after a saving rate shock instrumented by high-frequency monetary policy surprises by Gertler & Karadi (2015)

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The dark grey transparent areas denote the corresponding 68% and 90% confidence bands.

VAR[13]. Identification is then based on the validity and relevance of the instrument. The instrument is valid if it is uncorrelated to the structural shocks to real GDP, i.e. $\mathbb{E}[z_t \epsilon_t^y] = 0$. Given the high-frequency identification approach used by Gertler & Karadi (2015) there is consensus that the instrument is exogenous to shocks in output and hence can be used to identify e.g. effects of monetary policy. We argue that the relevance assumption holds as changes in consumption and thereby savings are one of the main transmission channels of monetary policy (see Holm et al. (2021) among others). Furthermore, if it was rather unrelated and therefore at best a weak instrument we would hardly get any significant results. Figure 3.7 depicts the estimated impulse response in real GDP after a shock in the saving rate induced by a monetary policy surprise[14]. As before, we find evidence that saving shocks have a significant and contractionary effect on economic activity.

[13]Also referred to as 'internal instrument' recursive VAR identification and similar to the VARX approach by Paul (2020)

[14]The plot is virtually the same for both saving rate definitions, see figure C.16 in the appendix using the saving rate in terms of disposable household income.

3.5 Conclusion

This paper uses the long-run properties arising from MPS shocks in both exogenous- and endogenous growth models with sticky prices in order to identify their causal effect on output. We find that time series data from the United States is strongly supportive of the notion that MPS shocks indeed have a causal, and contractionary, effect on economic activity, lending support to the most common approach of studying the financial crisis. Furthermore, we find evidence of a paradox-of-thrift, namely that a rise in the saving rate can lead to a decrease in total savings as output contracts. This corroborates the narrative of the financial crisis being ultimately the response to a negative shock to the demands for goods all across the board and underlines the claim by Christiano et al. (2017) to reconsider the theory of the paradox-of-thrift in modern macroeconomics. Last but not least, we provide evidence that saving shocks can have a negative impact orthogonal to uncertainty shocks. In fact, uncertainty shocks appear to be transmitted almost entirely via changes in the saving rate. Our findings are robust to a battery of various robustness checks.

References

Aghion, P. & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica (1986-1998)*, 60(2), 323.

Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *The quarterly journal of economics*, 131(4), 1593–1636.

Baker, S. R., Bloom, N., Davis, S. J., & Kost, K. J. (2019). *Policy news and stock market volatility*. Technical report, National Bureau of Economic Research.

Barro, R. J. & Martin, S.-i. (2004). *Economic Growth*. MIT press.

Basu, S. & Bundick, B. (2017). Uncertainty shocks in a model of effective demand. *Econometrica*, 85(3), 937–958.

Blanchard, O. J. & Quah, D. (1989). The dynamic effects of aggregate demand and supply disturbances. *The American Economic Review*, 79(4), 655–673.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., & Terry, S. J. (2018). Really uncertain business cycles. *Econometrica*, 86(3), 1031–1065.

Caballero, R. J. (1991). Earnings uncertainty and aggregate wealth accumulation. *The American Economic Review*, (pp. 859–871).

Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly journal of economics*, 112(1), 1–55.

Carroll, C. D., Hall, R. E., & Zeldes, S. P. (1992). The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings papers on economic activity*, 1992(2), 61–156.

Christiano, L., Eichenbaum, M., & Rebelo, S. (2011). When is the government spending multiplier large? *Journal of Political Economy*, 119(1), 78–121.

Christiano, L. J. et al. (2017). *The great recession: A macroeconomic earthquake*. Technical report, Federal Reserve Bank of Minneapolis.

Christiano, L. J., Motto, R., & Rostagno, M. (2014). Risk shocks. *American Economic Review*, 104(1), 27–65.

Deaton, A. (1989). *Saving and liquidity constraints*. Technical report, National Bureau of Economic Research.

Eggertsson, G. B. et al. (2003). Zero bound on interest rates and optimal monetary policy. *Brookings papers on economic activity*, 2003(1), 139–233.

Engen, E. M. & Gruber, J. (2001). Unemployment insurance and precautionary saving. *Journal of monetary Economics*, 47(3), 545–579.

Gali, J. (1999). Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations? *American economic review*, 89(1), 249–271.

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

Gourinchas, P.-O. & Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1), 47–89.

Gourio, F. (2012). Disaster risk and business cycles. *American Economic Review*, 102(6), 2734–66.

Grossman, G. M. & Helpman, E. (1991). Quality ladders in the theory of growth. *The Review of economic studies*, 58(1), 43–61.

Guerrieri, V. & Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics*, 132(3), 1427–1467.

Holm, M. B., Paul, P., & Tischbirek, A. (2021). The transmission of monetary policy under the microscope.

Keen, B. & Wang, Y. (2007). What is a realistic value for price adjustment costs in new keynesian models? *Applied Economics Letters*, 14(11), 789–793.

Kilian, L. & Lütkepohl, H. (2017). *Structural vector autoregressive analysis*. Cambridge University Press.

Leland, H. E. (1978). Saving and uncertainty: The precautionary demand for saving. In *Uncertainty in Economics* (pp. 127–139). Elsevier.

Lugilde, A., Bande, R., & Riveiro, D. (2019). Precautionary saving: a review of the empirical literature. *Journal of Economic Surveys*, 33(2), 481–515.

Paul, P. (2020). The time-varying effect of monetary policy on asset prices. *Review of Economics and Statistics*, 102(4), 690–704.

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

Rotemberg, J. J. (1982). Sticky prices in the united states. *Journal of Political Economy*, 90(6), 1187–1211.

Rubio-Ramirez, J. F., Waggoner, D. F., & Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *The Review of Economic Studies*, 77(2), 665–696.

Skinner, J. (1988). Risky income, life cycle consumption, and precautionary savings. *Journal of Monetary Economics*, 22(2), 237–255.

Zeldes, S. P. (1989). Optimal consumption with stochastic income: Deviations from certainty equivalence. *The Quarterly Journal of Economics*, 104(2), 275–298.

Appendix C

Appendix for Chapter 3

C.1 Model derivations

C.1.1 Aggregation and market clearing

Intermediate Production: The final good sector's demand for intermediate production can be found by aggregating equation (8) over all intermediate goods

$$Y^I(t) = \int_{i=0}^{1} \left(\frac{p_i(t)}{P(t)}\right)^{-\sigma} \left(\frac{\sigma-1}{\sigma}A(t)\right)^{\sigma} L(t) q^{(\sigma-1)k_i} di \tag{C.1}$$

$$= \left(\frac{\sigma-1}{\sigma}A(t)\right)^{\sigma} L(t)\psi^{-\sigma} Q(t). \tag{C.2}$$

Final output: Analogously final output is given by

$$Y(t) = A(t)L(t)^{1/\sigma} \int_{i=0}^{1} \left(q^{k_i(t)} y_i(t)\right)^{\frac{\sigma-1}{\sigma}} di \tag{C.3}$$

$$= A(t)L(t)^{1/\sigma} \int_{i=0}^{1} \left(q^{k_i(t)} \left(\frac{p_i(t)}{P(t)}\right)^{-\sigma} \left(\frac{\sigma-1}{\sigma}A(t)\right)^{\sigma} L(t) q^{(\sigma-1)k_i}\right)^{\frac{\sigma-1}{\sigma}} di \tag{C.4}$$

$$= \left(\frac{\sigma-1}{\sigma}A(t)\right)^{\sigma} L(t)\psi^{1-\sigma} Q(t). \tag{C.5}$$

Remark: Final output is larger than intermediate output as $\psi > 1$, hence $\psi^{1-\sigma} > \psi^{-\sigma}$ for $\sigma > 0$.

Research sector's demand for final good: We get this by rearranging the equation for research success and integrating over all research sectors:

$$Z(t) = \int_0^1 \frac{\mu}{\phi(k_i(t))} di = \int_0^1 \frac{\mu}{\frac{1}{\lambda} q^{(1-\sigma)(k_i+1)}} di = \lambda \mu q^{(\sigma-1)} Q(t) \tag{C.6}$$

Goods market clearing:

$$Y(t) = C(t) + Y^I(t) + Z(t) + \frac{\theta}{2} (\pi(t) - 1)^2 Y^I(t) \tag{C.7}$$

Labour market clearing:

$$L^d(t) = L^s(t) \tag{C.8}$$

C.1.2 Balanced growth path of the endogenous growth model

In case of the endogenous growth model all growth is coming through growth in the quality index of products $Q(t)$ and we impose a constant total factor productivity A. Recall the definition of the quality index

$$Q(t) = \int_{i=0}^1 \left(q^{k_i}\right)^{\sigma-1} di. \tag{C.9}$$

In each sector i, the term $\left(q^{k_i}\right)^{\sigma-1}$ does not change if no innovation occurs but changes to $\left(q^{k_i+1}\right)^{\sigma-1}$ in the case of successful research. The probability of success is μ as defined above which is independent of time and rung on the quality ladder. Furthermore, we assume the probability of success to be independent across sectors. Hence, we can write the expected change in the quality index as follows

$$E[\Delta Q(t)] = \int_0^1 \mu\left((q^{k_i+1})^{\sigma-1} - (q^{k_i})^{\sigma-1}\right) di \tag{C.10}$$

$$= \mu(q^{\sigma-1} - 1) \int_0^1 \left(q^{k_i}\right)^{\sigma-1} di \tag{C.11}$$

$$= \mu(q^{\sigma-1} - 1) Q(t). \tag{C.12}$$

Hence, we find growth in the quality index as

$$g_Q = \frac{\Delta Q(t)}{Q(t)} = \mu(q^{\sigma-1} - 1), \tag{C.13}$$

where we use the law of large numbers that the expected change is the same as the actual change for the continuum of intermediate firms. Note that for $\sigma = 1$, the growth rate goes to zero as the production of the final good is independent of the intermediate good's quality. Analogously, if $q = 1$ there is no growth through quality improvements as the quality ladder collapses to one rung.

The balanced growth paths for the final good, the intermediate good and the investment in research in case of endogenous growth can then be written as

$$\bar{Y}(t) = \frac{Y(t)}{Q(t)} = \left(\frac{\sigma-1}{\sigma}A\right)^{\sigma} L(t)\psi^{1-\sigma}, \tag{C.14}$$

$$\bar{Y}^I(t) = \frac{Y(t)^I}{Q(t)} = \left(\frac{\sigma-1}{\sigma}A\right)^{\sigma} L(t)\psi^{-\sigma}, \tag{C.15}$$

$$\bar{Z}(t) = \frac{Z(t)}{Q(t)} = \lambda\mu q^{(\sigma-1)}, \tag{C.16}$$

and from the goods market clearing condition we get that consumption grows at the same rate g_q and must equal the difference between total final good production $Y(t)$ and the sum of the demands for final goods by the intermediate and research sector as well as the Rotemberg adjustment cost.

$$\frac{C(t)}{Q(t)} = \frac{Y(t)}{Q(t)} - \frac{Y(t)^I}{Q(t)} - \frac{Z(t)}{Q(t)} - \frac{\theta}{2}(\pi(t)-1)^2 \frac{Y(t)^I}{Q(t)}$$

$$\bar{C}(t) = \bar{Y}(t) - \bar{Y}^I(t) - \bar{Z}(t) - \frac{\theta}{2}(\pi(t)-1)^2 \bar{Y}^I(t) \tag{C.17}$$

C.2 Balanced growth path of the exogenous growth model

An exogenous version of the growth model is nested when setting $Z(t) = 0$ and shutting down thereby shutting down the R&D channel [make investment productive yet not growth enhancing, like in Solow model]. Under this assumption the growth in the quality index $Q(t)$ is zero as the economy stays on the same quality rung. The only source of growth is then total factor productivity which is imposed exogenously. The growth rate in the final good and intermediate good is as follows

$$g_Y = g_{Y^I} = \sigma g_A. \tag{C.18}$$

Hence, we can define the balanced growth paths for the final good and the intermediate inputs as

$$\bar{Y}_{ex}(t) = \frac{Y(t)}{A(t)^{\sigma}} = \left(\frac{\sigma-1}{\sigma}\right)^{\sigma} L(t)\psi^{1-\sigma} Q, \tag{C.19}$$

$$\bar{Y}^I_{ex}(t) = \frac{Y^I(t)}{A(t)^{\sigma}} = \left(\frac{\sigma-1}{\sigma}\right)^{\sigma} L(t)\psi^{-\sigma} Q, \tag{C.20}$$

where Q is fixed exogenously. When shutting down the investment channel, the goods market clearing condition simplifies to

$$\bar{C}_{ex}(t) = \bar{Y}_{ex}(t) - \bar{Y}^I_{ex}(t) - \frac{\theta}{2}(\pi(t)-1)^2\, \bar{Y}^I_{ex}(t). \tag{C.21}$$

C.3 Theoretical results

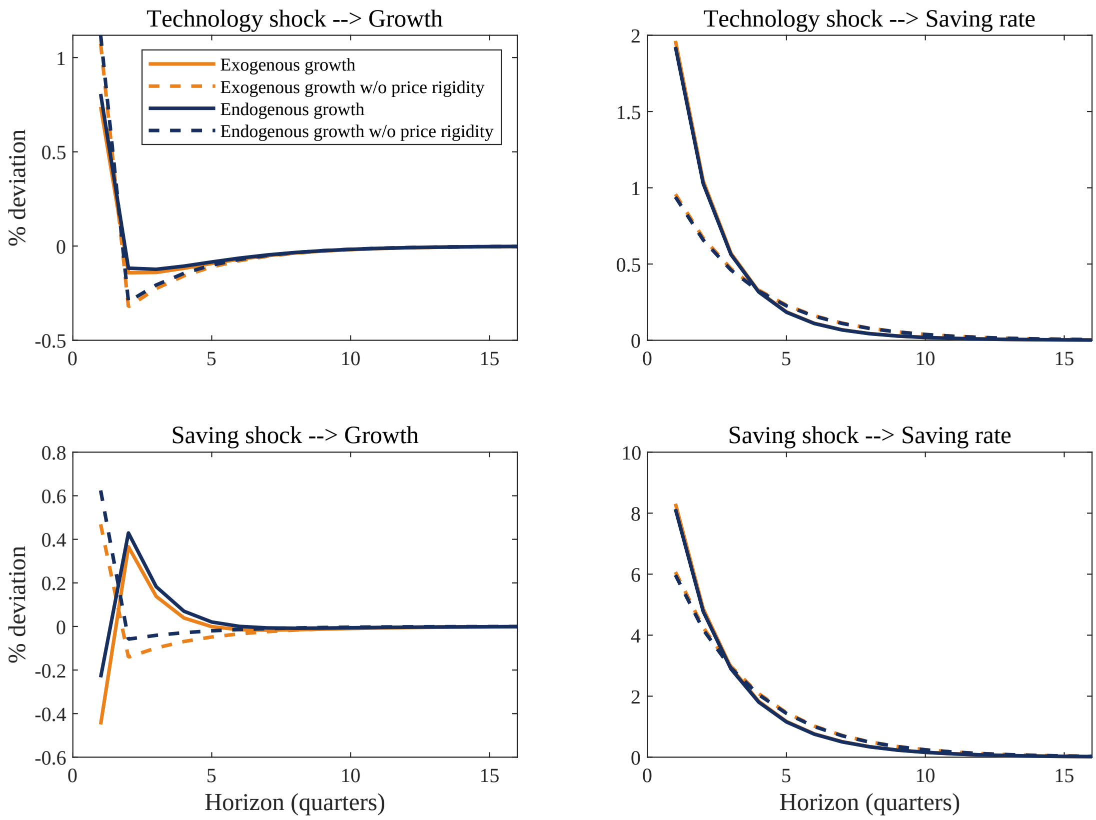

Fig. C.1 Impulse response functions for the exogenous and endogenous growth model presented in case of a temporary shock to the discount factor.

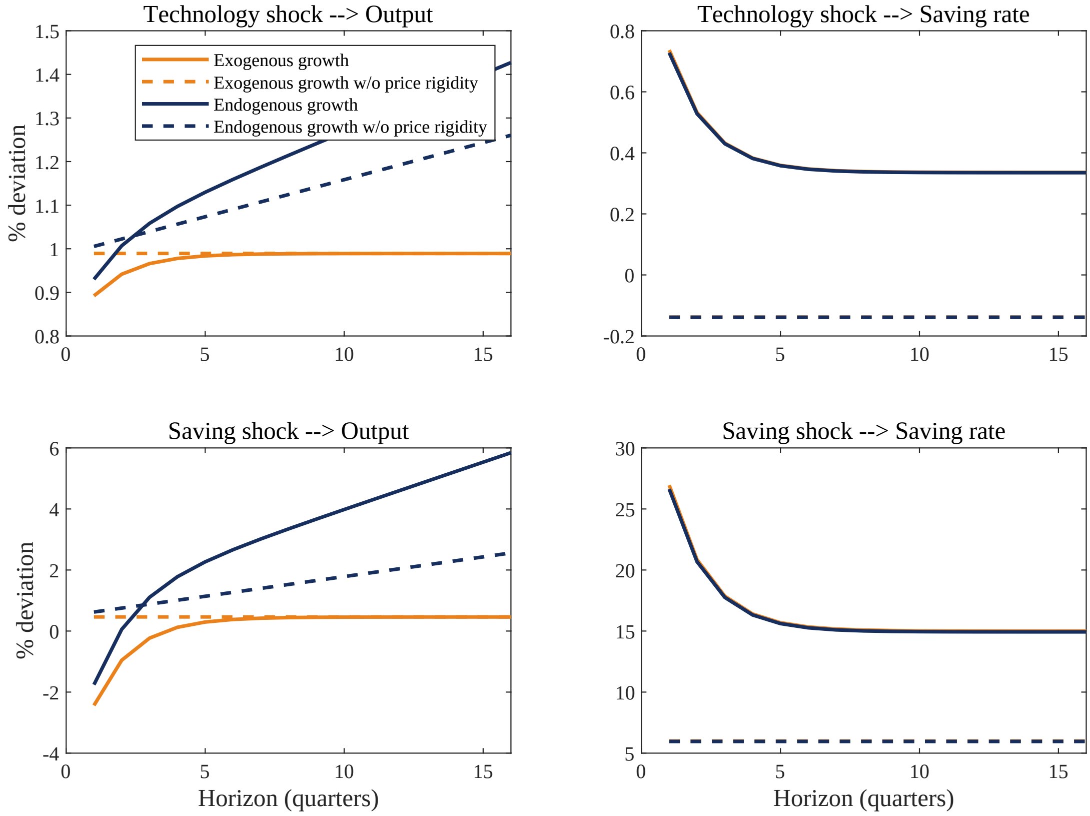

Fig. C.2 Impulse response functions for the exogenous and endogenous growth model presented in case of a permanent shock to the discount factor.

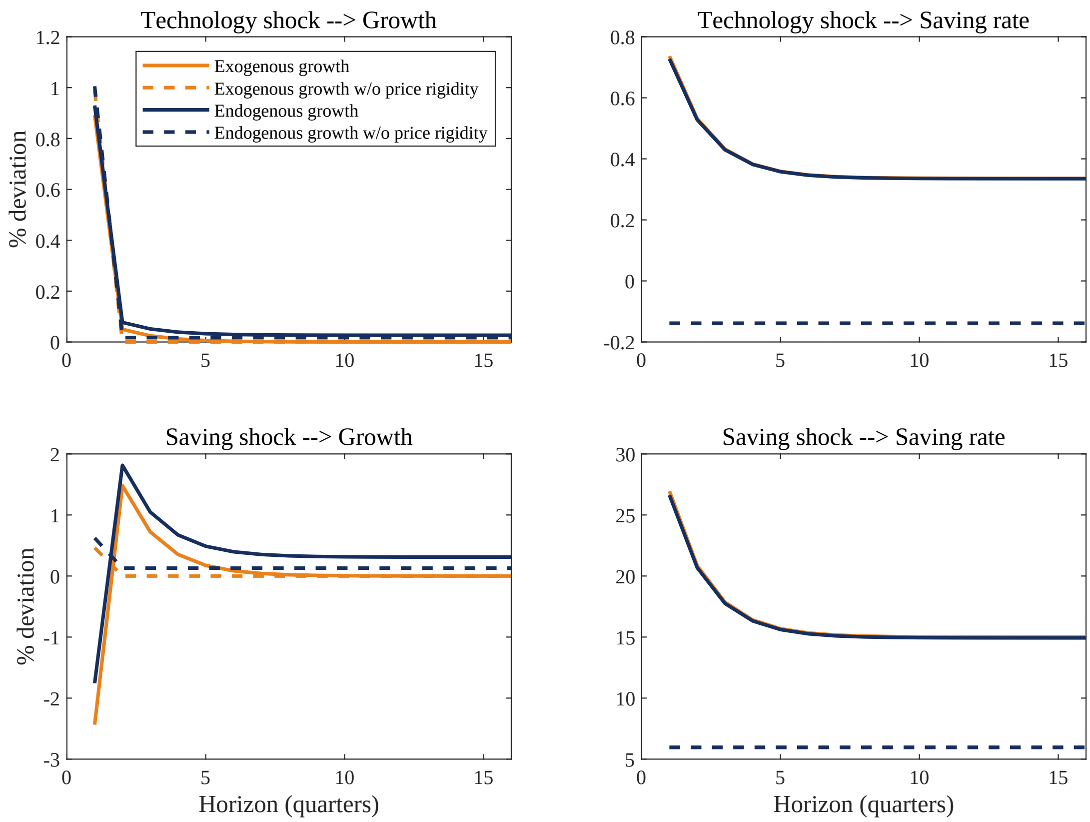

Fig. C.3 Impulse response functions for the exogenous and endogenous growth model presented in case of a permanent shock to the discount factor.

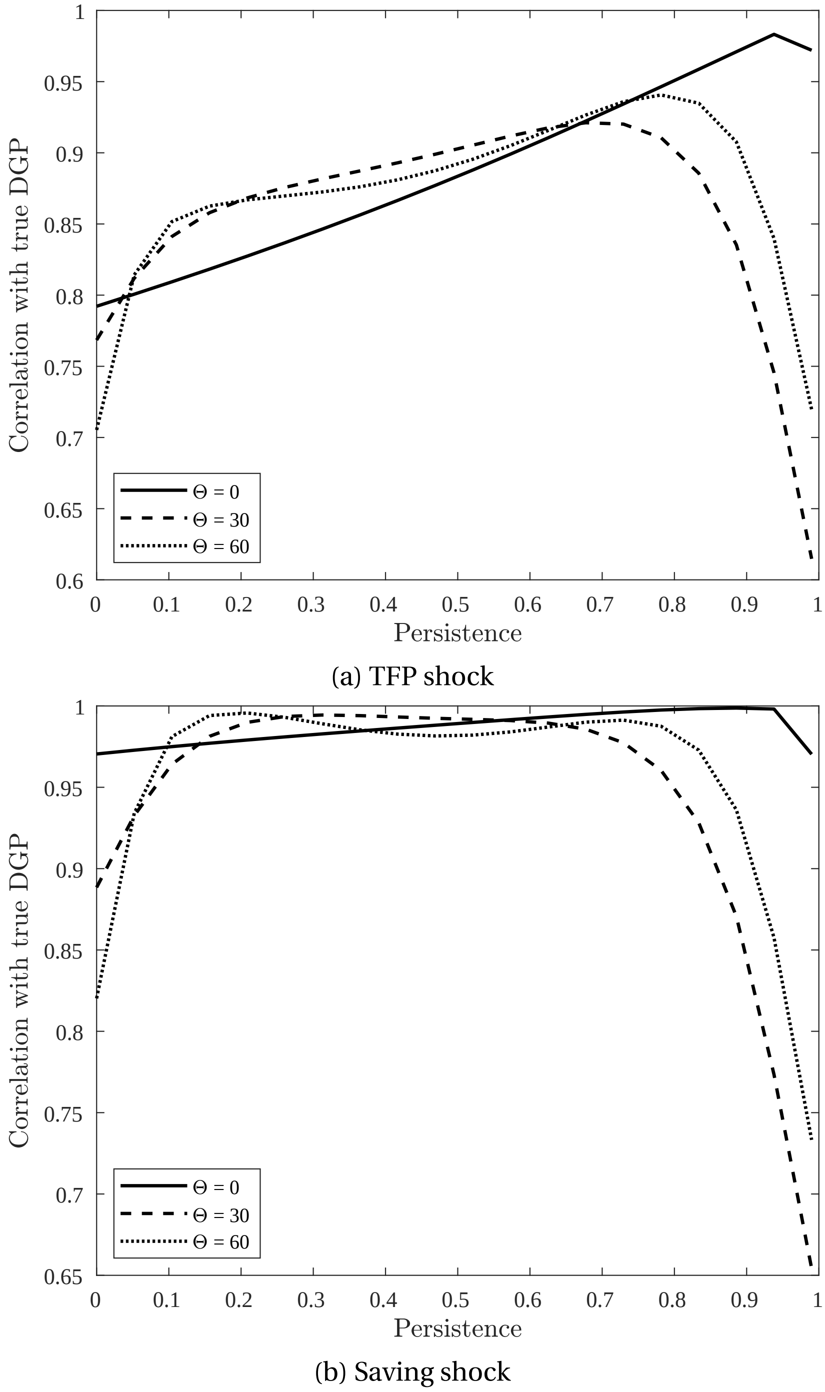

Fig. C.4 Identification of structural shocks in exogenous growth model

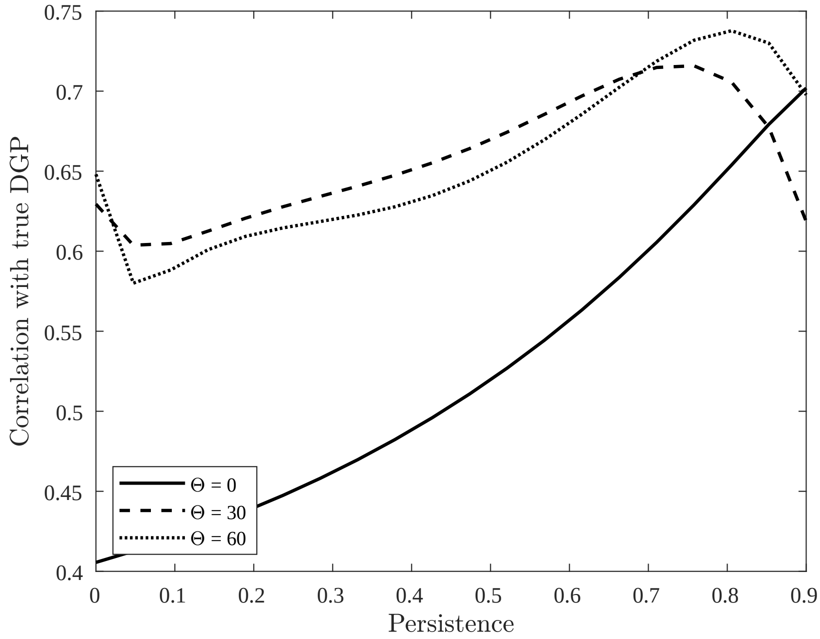

(a) TFP shock

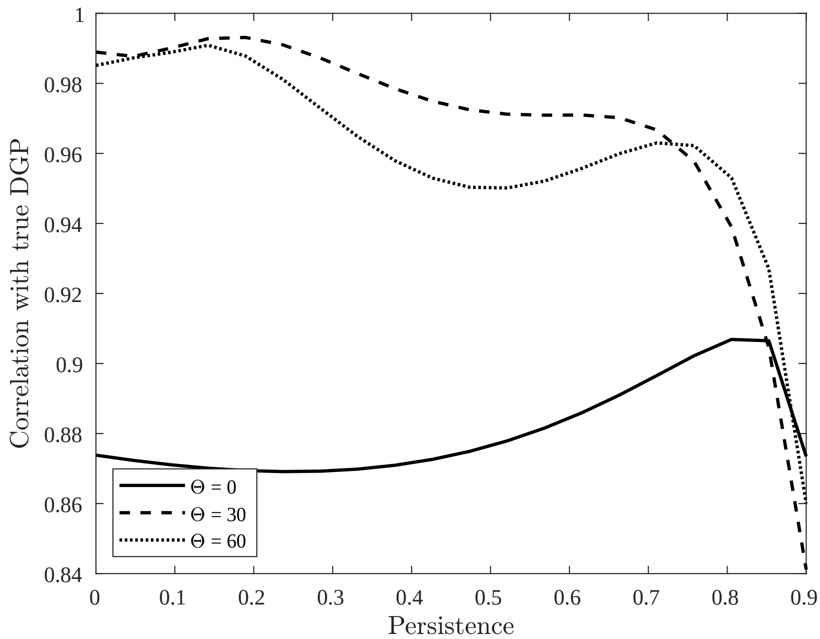

(b) Saving shock

Fig. C.5 Identification of structural shocks in endogenous growth model

C.4 Empirical results

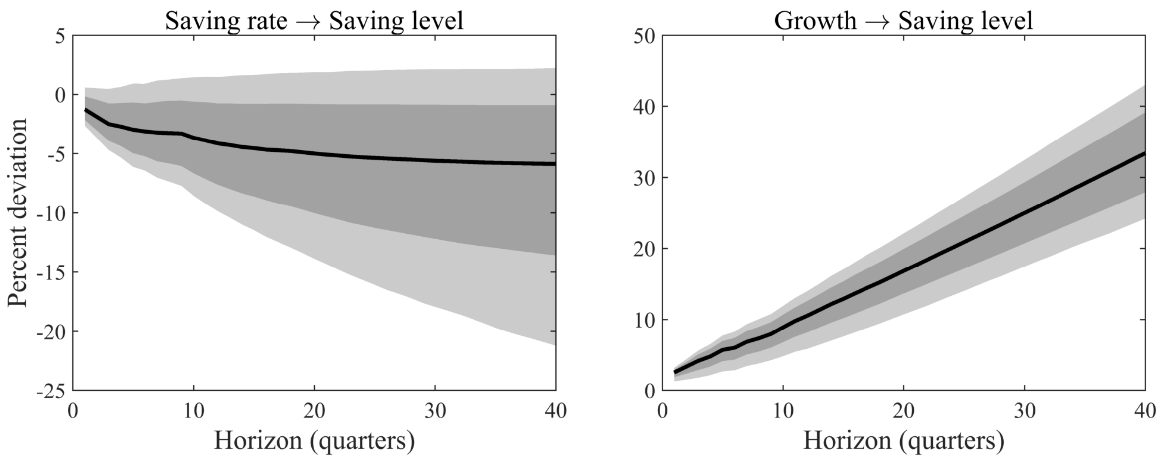

Fig. C.6 Impulse response function of the level of private savings according to the exclusion restriction 2 in Figure 4.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The grey transparent area shows the corresponding 68% confidence bounds.

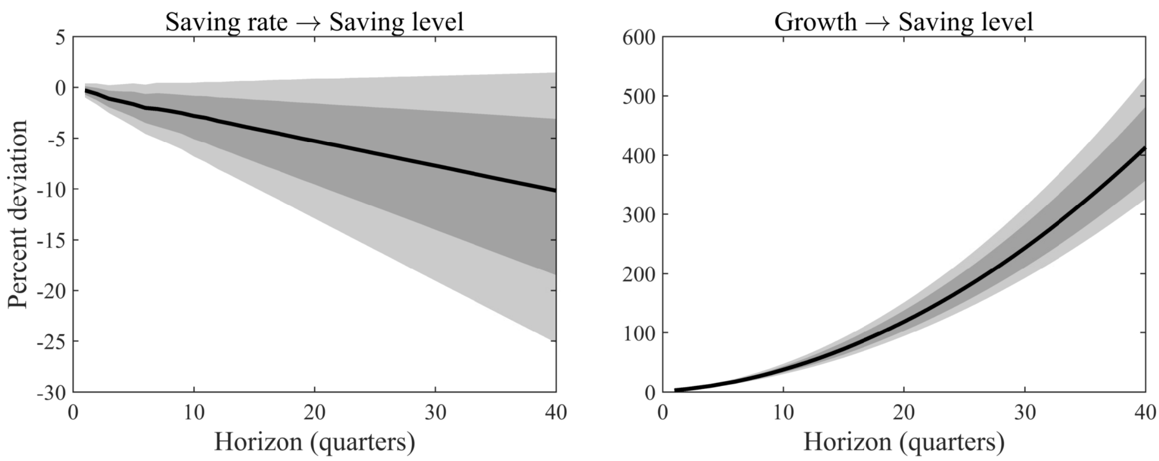

Fig. C.7 Impulse response function of the level of private savings according to the exclusion restriction 3 in Figure 5.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The grey transparent area shows the corresponding 68% confidence bounds.

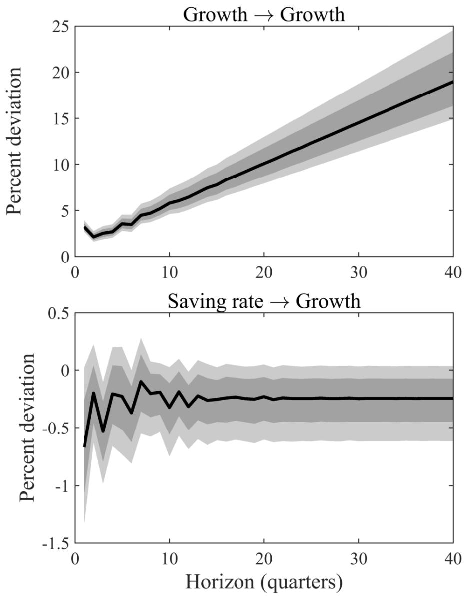

Fig. C.8 Impulse response functions using exclusion restriction 3 that a permanent shock to the saving rate has no long-run effect on the change in growth of output.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The grey transparent area shows the corresponding 68% confidence bounds.

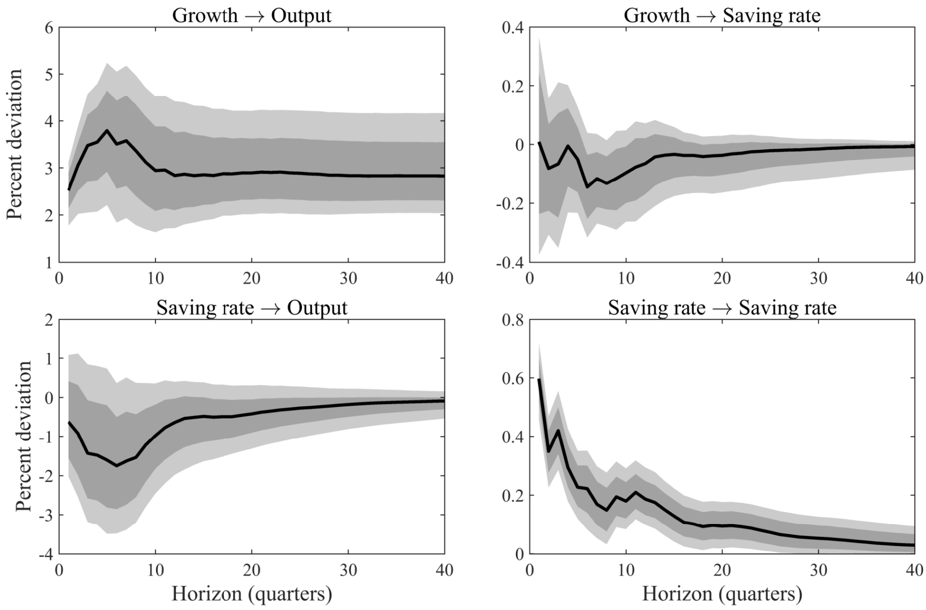

Fig. C.9 Impulse response functions for a temporary shock to the savings rate using exclusion restriction 1. The saving rate is defined as personal savings share of disposable income.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The grey transparent area shows the corresponding 68% confidence bounds.

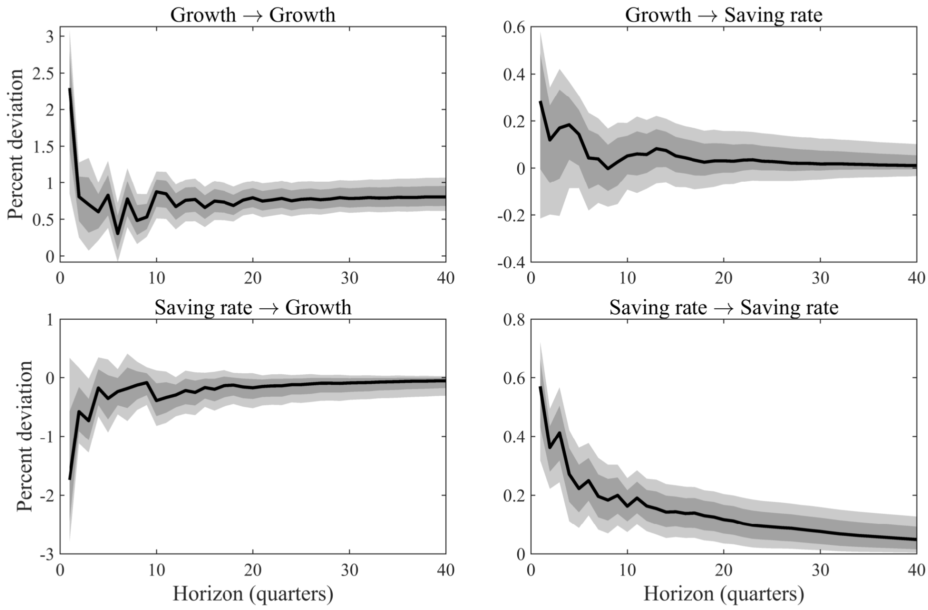

Fig. C.10 Impulse response functions for a temporary shock to the savings rate using exclusion restriction 2. The saving rate is defined as personal savings share of disposable income.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The grey transparent area shows the corresponding 68% confidence bounds.

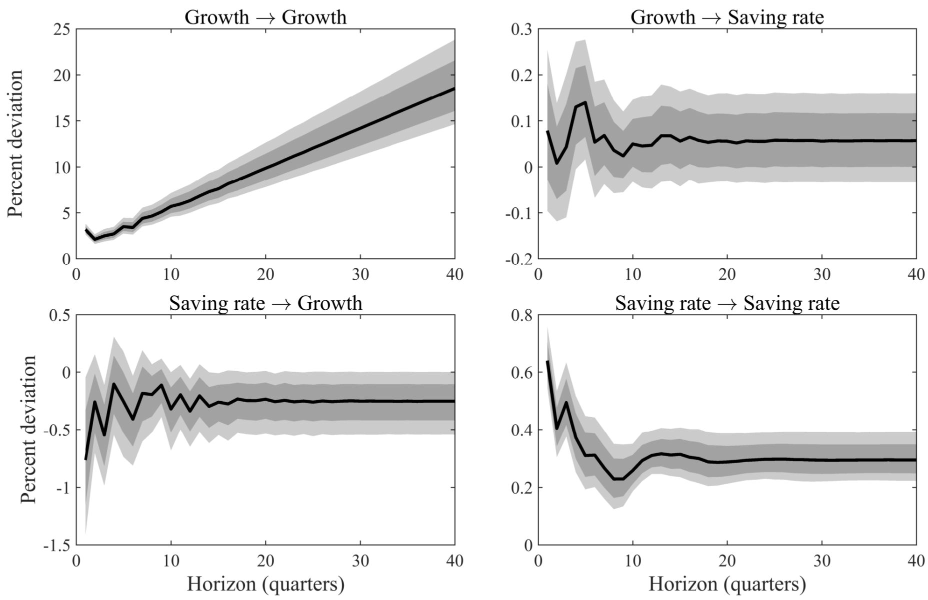

Fig. C.11 Impulse response functions for a permanent shock to the savings rate using exclusion restriction 3. The saving rate is defined as personal savings share of disposable income.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The grey transparent area shows the corresponding 68% confidence bounds.

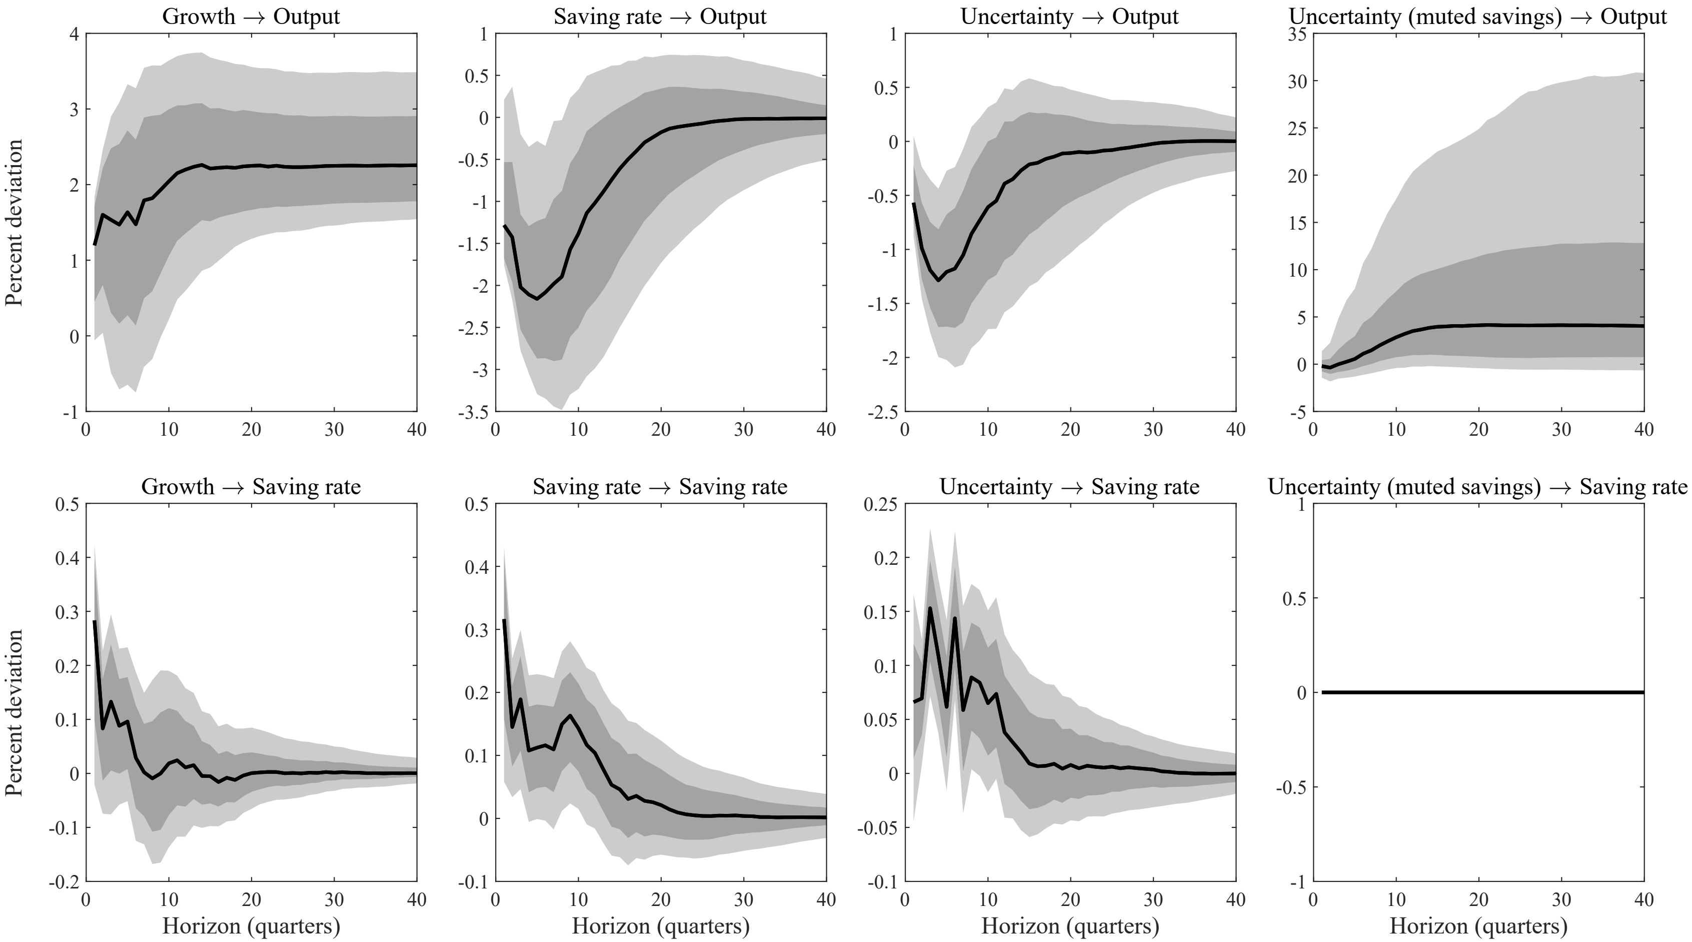

Fig. C.12 Impulse response functions using long- and short-run exclusion restrictions for uncertainty and saving shocks.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The dark grey transparent area shows the corresponding 68% confidence bounds.

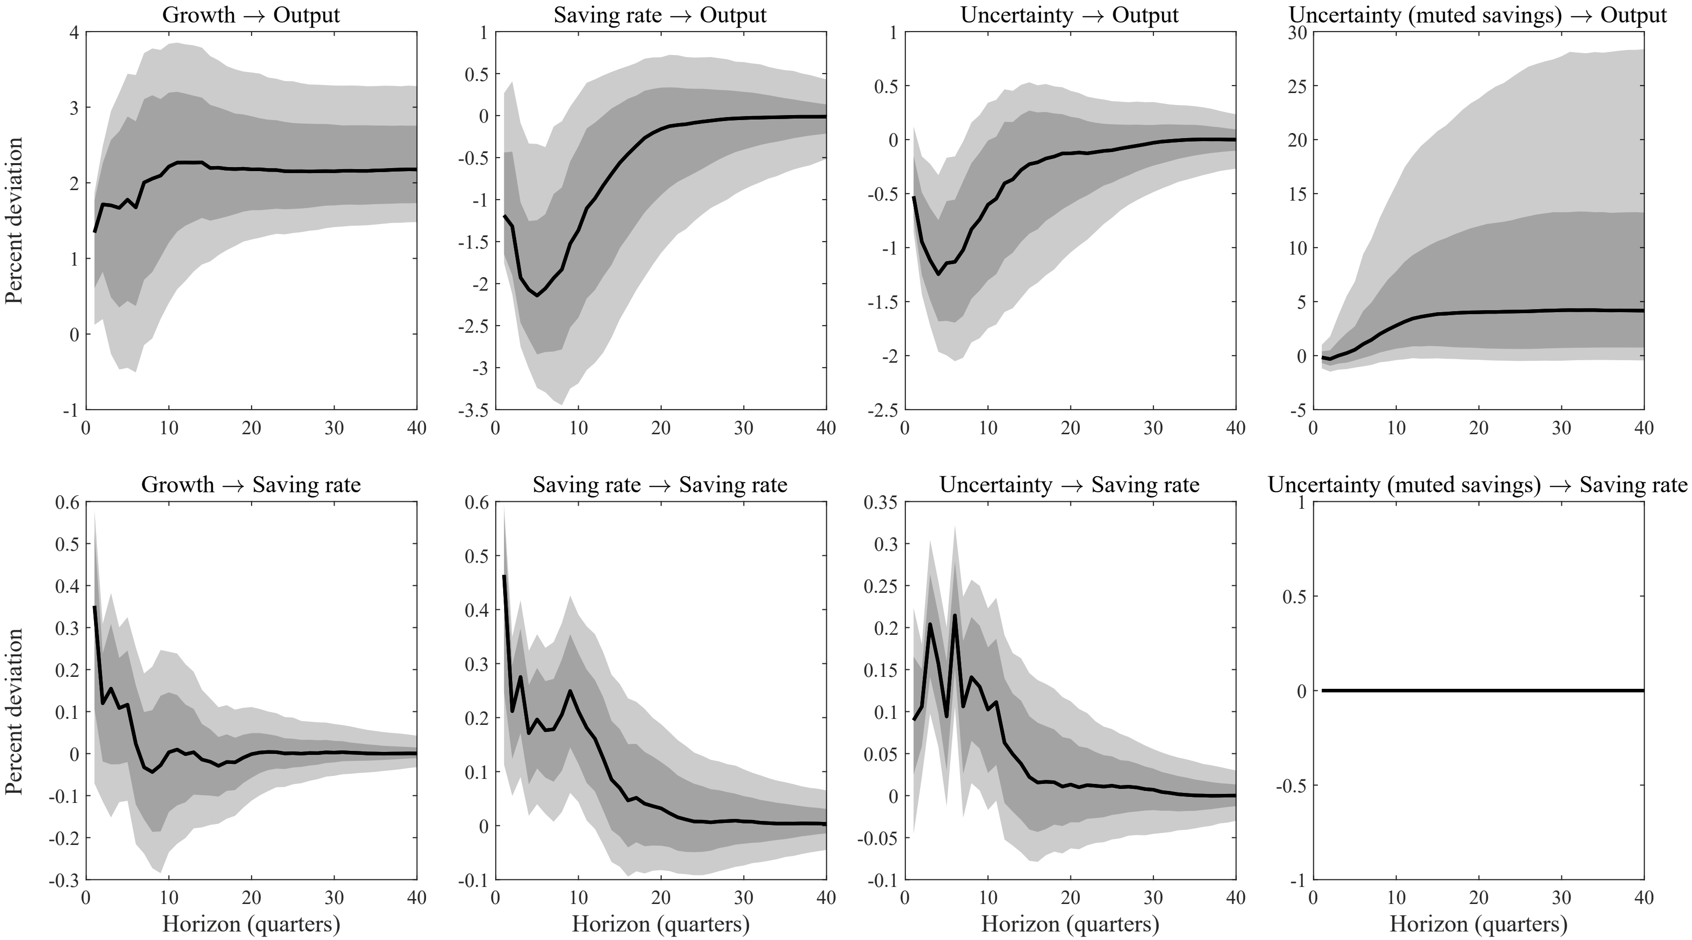

Fig. C.13 Impulse response functions using long- and short-run exclusion restrictions for uncertainty and saving shocks.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The dark grey transparent area shows the corresponding 68% confidence bounds.

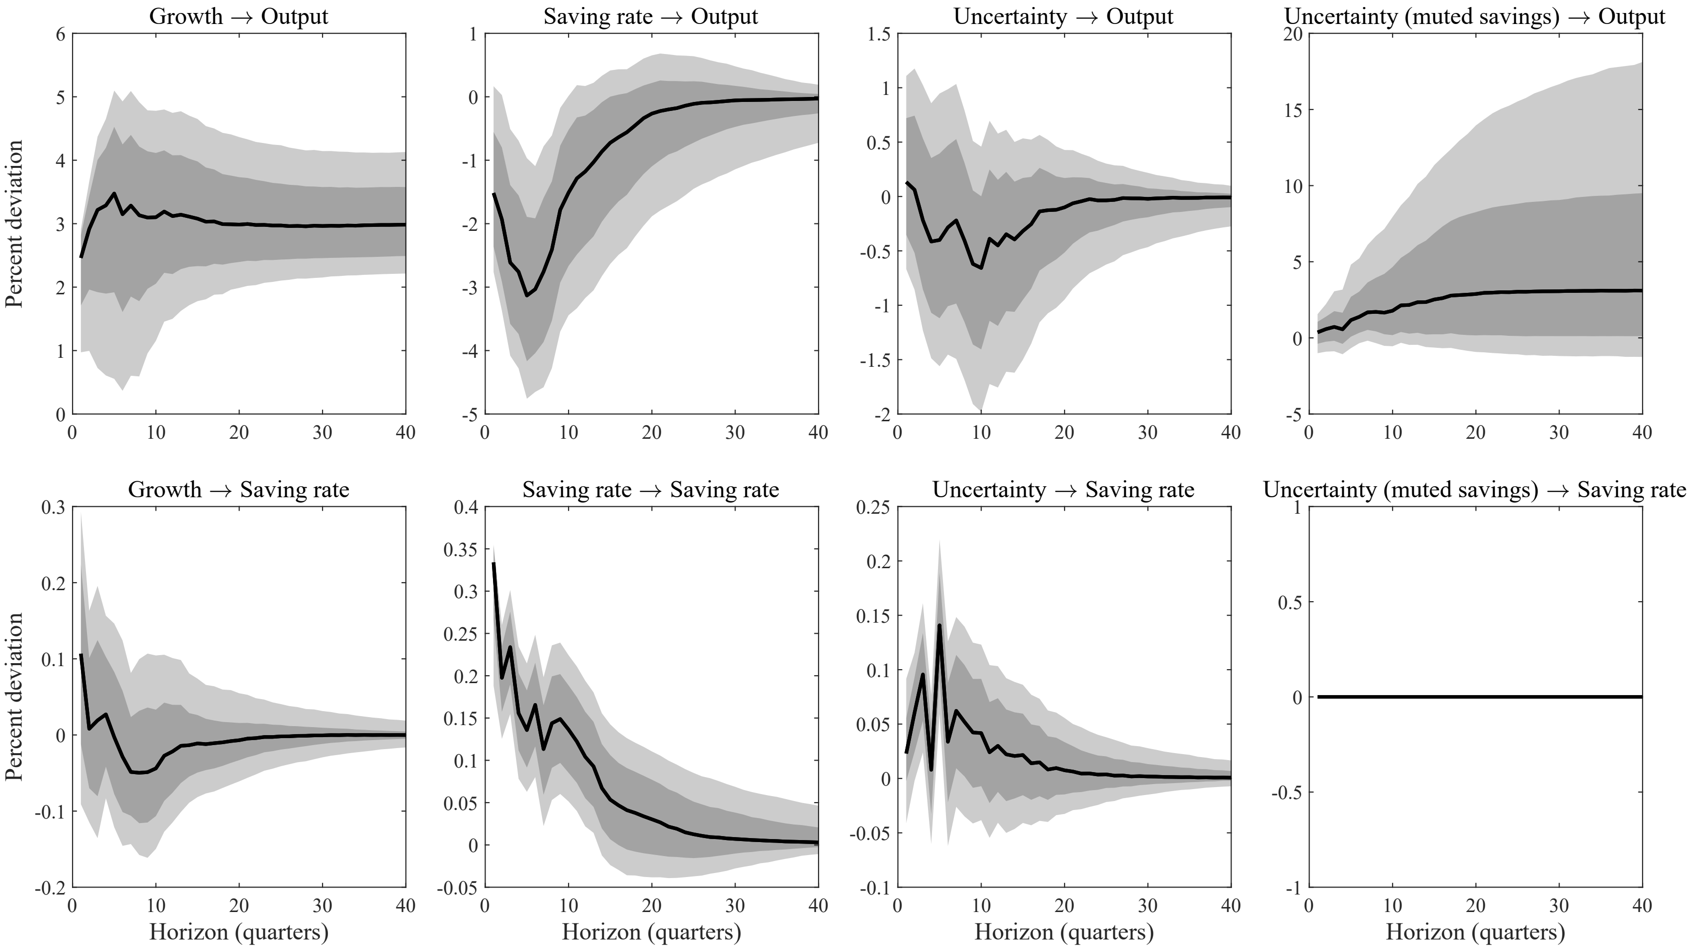

Fig. C.14 Impulse response functions using long- and short-run exclusion restrictions for uncertainty and saving shocks.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The dark grey transparent area shows the corresponding 68% confidence bounds.

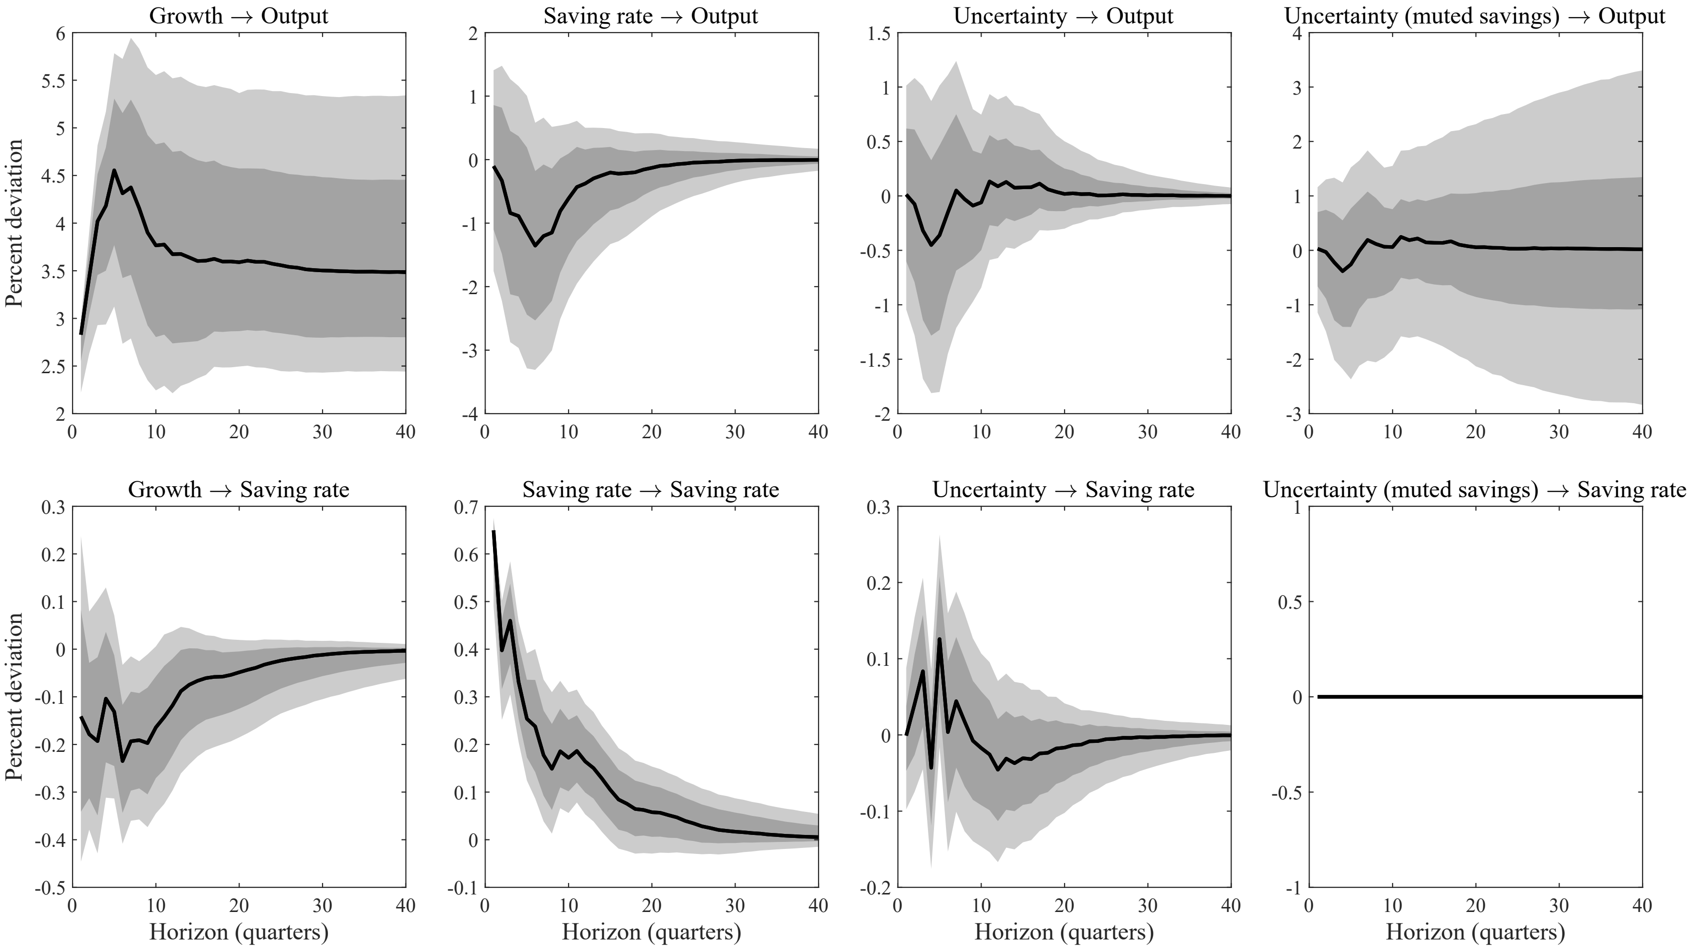

Fig. C.15 Impulse response functions using long- and short-run exclusion restrictions for uncertainty and saving shocks.

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The dark grey transparent area shows the corresponding 68% confidence bounds.

Table C.1 Cumulative percentage response in real GDP growth for alternative estimation specifications

		Lags			Trend			Time	
Identification	At horizon h	4	8	12	linear	quadratic	1960Q1-2020Q4	1960Q1-2006Q4	1990Q1-2019Q4
Exclusion restriction 1	1	-0.17	-0.59	-0.39	-1.50	-0.35	-1.48	0.51	-1.16
		(0.81)	(0.94)	(0.99)	(0.73)	(0.93)	(2.42)	(0.83)	(0.46)
	2	-0.27	-0.87	-0.64	-1.89	-0.58	-1.34	0.30	-1.37
		(0.96)	(1.10)	(1.17)	(0.85)	(1.09)	(2.21)	(0.96)	(0.56)
	4	-0.73	-1.42	-1.11	-2.61	-1.09	-1.32	-0.02	-1.93
		(1.06)	(1.21)	(1.34)	(0.92)	(1.19)	(2.29)	(1.02)	(0.62)
	8	-0.58	-1.53	-1.28	-2.38	-1.24	-1.25	-0.44	-1.57
		(0.79)	(1.08)	(1.15)	(0.91)	(1.00)	(2.43)	(0.83)	(0.68)
Exclusion restriction 2	1	-0.67	-1.56	-1.27	-1.39	-1.47	-2.55	-1.56	-0.94
		(0.88)	(0.87)	(0.93)	(0.79)	(0.86)	(1.70)	(0.88)	(0.50)
	2	-0.88	-2.08	-1.68	-1.80	-1.97	-3.19	-2.17	-1.12
		(1.23)	(1.17)	(1.19)	(1.06)	(1.14)	(1.46)	(1.14)	(0.68)
	4	-1.44	-3.00	-2.31	-2.62	-2.82	-4.16	-3.00	-1.69
		(1.84)	(1.63)	(1.63)	(1.48)	(1.59)	(1.82)	(1.59)	(0.91)
	8	-2.33	-3.84	-2.38	-3.06	-3.54	-6.04	-3.89	-1.45
		(2.94)	(2.36)	(2.23)	(2.12)	(2.28)	(2.73)	(2.28)	(1.35)
Exclusion restriction 3	1	-0.74	-0.75	-1.27	-0.67	-0.77	-2.67	-0.81	-0.22
		(0.38)	(0.41)	(0.49)	(0.42)	(0.42)	(1.14)	(0.50)	(0.41)
	2	-1.18	-0.88	-1.77	-0.88	-1.04	-1.89	-1.20	-0.17
		(0.60)	(0.62)	(0.70)	(0.64)	(0.63)	(1.16)	(0.72)	(0.64)
	4	-2.17	-1.64	-2.74	-1.64	-1.71	-2.25	-1.79	-0.70
		(1.00)	(0.98)	(1.13)	(1.05)	(0.99)	(1.83)	(1.14)	(1.01)
	8	-3.96	-2.60	-3.83	-2.60	-2.82	-3.90	-3.16	-0.73
		(1.84)	(1.64)	(1.94)	(1.85)	(1.64)	(3.01)	(1.92)	(1.69)

Note: Standard errors are parentheses; $'$ $p < 0.32$, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

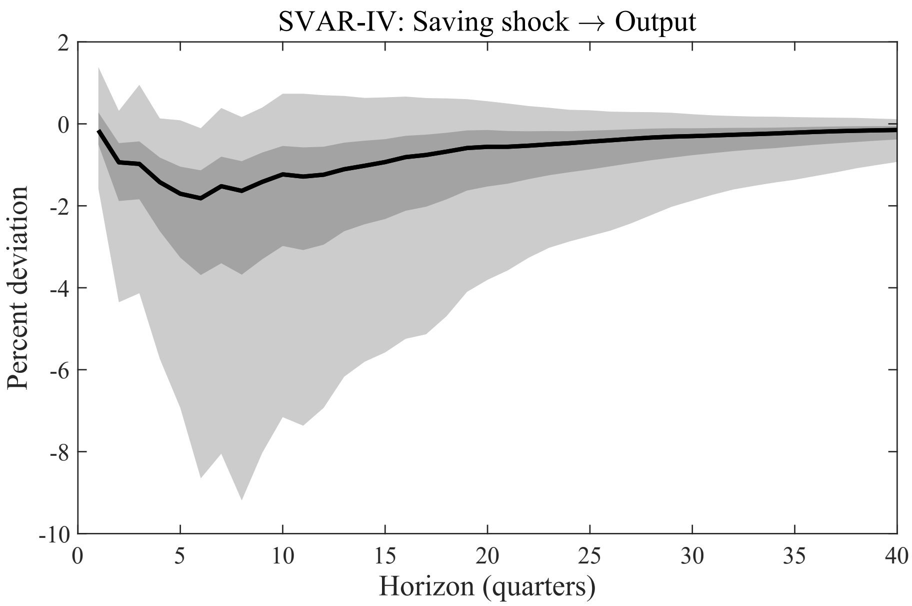

Fig. C.16 Impulse response function of real GDP after a saving rate shock instrumented by high-frequency monetary policy surprises by Gertler & Karadi (2015)

Notes. The black solid line shows the median impulse response function for 10,000 bootstrap replications. The dark grey transparent areas denote the corresponding 68% and 90% confidence bands.