Short-Term Tax Cuts, Long-Term Stimulus*

James Cloyne Joseba Martinez Haroon Mumtaz Paolo Surico

July 2022

Abstract

We study the persistent effects of temporary changes in U.S. federal corporate and personal income tax rates using a narrative identification approach. A corporate income tax cut leads to a sustained increase in GDP and productivity, with peak effects between five and eight years. R&D spending and capital investment display hump-shaped responses while hours worked and employment are much less affected. In contrast, personal income tax cuts trigger a short-lived boost to GDP, productivity and hours worked but have no long-term effects. We develop and estimate an endogenous growth model with variable factor utilization and show that these features generate a pro-cyclical response of productivity which is key to account for our empirical findings.

JEL classification: E23, E62, H24, H25, H31, H32, O32
Key words: corporate taxes, narrative identification, fiscal stimulus, productivity, R&D.

*We are grateful to Tim Besley, Òscar Jordà, Elias Papaioannou and Maarten de Ridder for very useful comments and insightful suggestions. James Cloyne (University of California Davis, NBER and CEPR) jcloyne@ucdavis.edu; Joseba Martinez (London Business School and CEPR) jmartinez@london.edu; Haroon Mumtaz (Queen Mary, University of London) h.mumtaz@qmul.ac.uk; Paolo Surico (London Business School and CEPR) psurico@london.edu.

1 Introduction

Can temporary tax cuts stimulate the economy over the longer-term? Over the last twenty years, extensive research efforts have focused on the short-run impact of stabilization policies, but evidence about the long-run effects of tax changes remains sparse. This is despite recurrent academic and policy debates about the potential longer-term impacts of tax reforms. Our paper provides empirical evidence and a theoretical mechanism through which temporary changes in corporate and personal taxes may have persistent effects.

We use Bayesian Local Projections and post-WWII U.S. data on output, taxes, productivity and R&D spending to estimate the dynamic effects of income tax changes.[1] Federal tax changes are identified based on the narrative approach of Romer and Romer (2010), which excludes all tax changes that were motivated by fluctuations in current or prospective economic conditions. We use the decomposition of these data into different tax types by Mertens and Ravn (2013) and focus on personal and corporate income tax changes separately.

Our paper has three main empirical findings. First, changes in tax rates tend to be temporary, reverting to historical averages between two and four years after the shock. Second, the effects of personal income taxes on output are large and significant during the first six quarters but return to zero within two years. In contrast, a corporate income tax cut generates a smaller stimulus on impact but has larger expansionary effects over the medium- and long-term, with peak effects occurring around eight years after the shock. Third, while the dynamic effects of personal income taxes are short-lived on productivity and negligible on R&D expenditure, changes in corporate income taxes are associated with a sustained, but temporary, rise in R&D spending and a persistent increase in aggregate productivity.

We interpret our empirical findings through the lens of a structural model with variable factor utilization (for both capital and labor), R&D spending and technological adoption. We estimate the structural model via Bayesian methods by minimizing the distance between the impulse responses of the structural model and those estimated using Local Projections (LPs). We choose prior distributions that are conventionally used in the empirical macro literature and that imply: (i) productivity is virtually a-cyclical and (ii) long-run effects are unlikely. Yet, the posterior distributions of the structural parameters imply impulse responses that replicate both the short-run

[1] As discussed in Section 2, the Bayesian approach provides us with an efficient method to characterise the joint and marginal distribution of the impulse responses to both tax shocks in the short-run and in the long-run. It also offers tools to navigate the small sample bias-variance trade-off highlighted by Jordà et al. (2020) and Li et al. (2021).

effects of personal income tax changes and the long-run effects of corporate income tax changes on output and productivity.

To uncover the mechanism behind our findings, we use the estimated structural model to perform three counterfactual simulations, which in turn: (i) switch R&D spending off, (ii) make factor utilization (both capital and labor) inelastic, (iii) switch R&D spending off *and* make factor utilization (both capital and labor) inelastic. The first exercise shows that R&D is key for the model's ability to generate a persistent response of productivity and output following a temporary change in corporate income taxes. The second exercise shows that variable factor utilization is crucial to produce the short-run response of productivity to temporary tax cuts. The third counterfactual shows that without R&D and variable factor utilization, an otherwise standard New-Keynesian model would not deliver pro-cyclical productivity, thereby missing the estimated output responses to either shock.

Related literature. Our paper relates to several strands of empirical and theoretical work. In terms of stabilization policy, an influential literature exemplified by Romer and Romer (2010), Barro and Redlick (2011), Mertens and Ravn (2013), Cloyne (2013) and Caldara and Kamps (2012), among many others, estimate the short-run dynamic effects of tax changes on output. This literature has focused on solving the identification problem that tax changes affect the economy but fiscal policy may also react to economic conditions.

This strand of research does not typically look at the long-run, or examine the impact on productivity or R&D expenditure. There have been a few recent studies, however, considering the effect of tax reforms on innovation in different contexts. Akcigit et al. (2022) exploit cross-sectional variation in income taxes across individual inventors and U.S. states to estimate large and positive effects of tax cuts on innovation and capital investment. Cram and Olbert (2022) measure the effects of the 2021 global corporate tax reform on stock prices across companies with different shares of foreign earnings and intangible assets and on sovereign debt risk of countries with different shares of multinational companies' tax bases. Baley et al. (2022) study the monetary policy response to permanent corporate tax reforms.

Growing research efforts, surveyed by Cerra et al. (2022) and including Comin and Gertler (2006), Benigno and Fornaro (2017), Anzoategui et al. (2019), de Ridder (2019), Beaudry et al. (2020), Jordà et al. (2020), Queraltó (2022), Fornaro and Wolf (2020), Furlanetto et al. (2021) and Antolin-Diaz and Surico (2022), among others, examine the long-run effects of non-technology

shocks working via strategic complementarities and financial frictions, monetary policy and government spending. A distinctive feature of our analysis is the focus on the long-run effects of corporate and personal income taxes on productivity, aggregate R&D spending and GDP, a key part of the policy debate around fiscal policy.

Our analysis is structured as follows. In Section 2, we present the identification strategy and the empirical framework. In Section 3, we present the evidence based on LPs and post-WWII U.S. data. In Section 4, we develop a structural model with variable factor utilization and endogenous growth, which we estimate in Section 5 by minimizing the distance between the model impulse responses and those based on the LP estimates of Section 3. In Section 6, we perform a number of counterfactual analyses that highlight the main channels that drive our empirical findings. Section 7 revisits the role of investment adjustment costs in the propagation of macroeconomic shocks. Section 8 concludes. The Appendix contains further details and robustness exercises.

2 Empirical Framework

In this section, we describe the narrative approach and dataset we use to isolate exogenous variation in taxes. We then present the empirical model that we use to identify the longer-run effects of corporate and personal income tax changes, and provide details on the data and estimation.

2.1 Identification and approach

Our goal is to examine the longer-term aggregate effects of different tax policy reforms. We face at least three empirical challenges. First, we need information on when tax policy was changed, including how different types of taxes were adjusted as these may have different long-run properties. Second, tax policy is often endogenous because policy levers tend to be adjusted in response to changes in current or prospective economic conditions: tax policy affects macroeconomic outcomes, but economic conditions also influence tax policy decisions. Third, we need an empirical approach well-suited to studying longer-term impacts.

We address the first two challenges using the identified corporate and personal taxes changes from Mertens and Ravn (2013). These data build on the original dataset of Romer and Romer (2010), who identify tax changes for the United States from 1950 to 2006. To isolate changes in tax policy that are plausibly "exogenous", Romer and Romer (2010) examine the motivations given by policymakers for all major pieces of Federal tax legislation over this period. Tax changes that were

not implemented for reasons related to changes in current or prospective future economic conditions are regarded as "exogenous".

A quantitative measure of each exogenous reform is constructed using historical revenue projections for the impact of the policy change, as announced at the time of the intervention. These are scaled by nominal GDP and, thus, approximate changes in the average tax rate (all else equal). Mertens and Ravn (2013) refine this series by excluding potentially anticipated reforms, defined as tax changes implemented more than 90 days after announcement. Key for our purpose, Mertens and Ravn (2013) sub-divide the Romer and Romer (2010) shocks into corporate and personal tax reforms. This so-called "narrative" approach of looking for quasi-natural experiments from historical episodes has a long tradition in macroeconomic research, as exemplified by Barro and Redlick (2011), Cloyne (2013), Mertens and Ravn (2012, 2014), Guajardo et al. (2014), Hayo and Uhl (2014), Cloyne and Surico (2017), Gunter et al. (2018), Nguyen et al. (2018), Hussain and Liu (2018), Cloyne et al. (2021).[2]

The literature studying the effects of tax changes using narrative methods tends to find large macroeconomic effects, but typically these papers focus on the shorter term effects over 2 to 5 years. A sizable part of the macroeconomic policy debate, however, has focused on the potential longer-term effects of tax reforms, yet there remains little direct evidence on this issue. In fact, policy recommendations often have to rely on inferring long-run results from the short-run estimates in a number of the papers referenced above. The datasets in Romer and Romer (2010) and Mertens and Ravn (2013), however, span a period of nearly 60 years and contain numerous reforms to personal and corporate taxes. It should also be feasible to examine the longer-term impacts of tax reforms using these data.

As for the empirical model, we need an econometric approach that allows us to draw inference about longer-term effects. Recent work by Jordà et al. (2020) for monetary policy has shown that the longer-term effects of policy interventions tend to be incorrectly captured when impulse response functions (IRFs) are estimated using a traditional Vector Autoregression (VAR) approach with short lag lengths (as is common in the empirical macro literature that focuses on relatively short time-series samples after WWII). This is because impulse responses are constructed as a projection from a fixed model using all the lags in the VAR. In finite samples, the lag structure has to be truncated

[2]The narrative approach arguably dates back to, at least, Friedman and Schwartz (1963) who examine episodes of unusual monetary policy in the United States. In a modern setting the approach has been popularized by Romer and Romer (1989) and Romer and Romer (2004). On the government spending side, a number of papers have employed a narrative approach to examine the impact of defence spending, e.g. Ramey and Shapiro (1998), Ramey (2011), Crafts and Mills (2013), Ramey and Zubairy (2018) and Barro and Redlick (2011).

and the VAR impulse response function, particularly at longer horizons, will be sensitive to the number of lags included (Li et al., 2021). Jordà et al. (2020) recommend estimation of impulse response functions using local projections (LPs), following Jordà (2005). This is a direct estimate of the impulse response function and does not use coefficient estimates on all the lagged controls to construct the IRF. As a result, this approach is less sensitive to the choice of lag structure and lag truncation issues that afflict VAR methods in finite samples. For estimation we use Bayesian methods, which provide an efficient way to compute and characterize joint and marginal posterior distributions.

One contribution of Mertens and Ravn (2013) is to introduce a methodology for treating the narratively identified tax changes derived from historical documents as potentially noisy "proxies" (or instruments) for the genuinely exogenous variation in tax policy (the "shock"). The Mertens and Ravn (2013) technology, however, is based on a vector autoregression framework. Accordingly, our econometric specification starts as close as possible to Mertens and Ravn (2013) but, following Jordà et al. (2020), we conduct estimation via local projections, given our focus on the longer-term impacts of tax changes.

We begin from a structure close to Mertens and Ravn (2013) where the joint dynamics of a vector of observables Z can be described by a reduced form including all the lags of the variables in Z. This is the conventional starting point for a vector autoregression approach. To construct the impulse response function, however, we follow Jordà (2005) and estimate the following sequence of local projections:

$$Z_{t+h} = c^{(h)} + B_1^{(h)} Z_{t-1} + \sum_{j=1}^{P} b_j^{(h)} Z_{t-1-j} + u_{t+h}, \qquad u_{t+h} \sim N(0, \Omega_h) \tag{1}$$

where Z_t denotes the M variables of interest described below, h is the impulse response horizon and u_{t+h} denote residuals that contain a combination of forecast errors and may be serially correlated and heteroscedastic. Here will follow the Jordà (2005) approach to estimating IRFs via local projections rather than a local projection-instrumental variables (LP-IV) approach (see, for example, Jordà and Taylor, 2015, Ramey, 2016). There are two reasons for this. First, the formulation in Jordà (2005) allows us to remain as close as possible to the set up in Mertens and Ravn (2013) while still conducting estimation via local projections. Indeed, the shorter-term effects we will estimate below are very close to the short-run IRFs estimated by Mertens and Ravn (2013), which provides a useful benchmark. Second, the approach in Mertens and Ravn (2013) considers two types of tax changes

using two instruments that are correlated. The two instruments identify a convolution of the tax shocks but we do not know the true causal relationship between the personal and the corporate income tax changes in the data. Mertens and Ravn (2013) consider different causal orderings when simulating their results from their proxy-VAR. We therefore implement the same approach here.[3]

The identification issue centers around the fact that the reduced form residuals are an unknown combination of all the underlying structural shocks, ε_t, including the exogenous variation in tax policy. The goal is to identify the contemporaneous impact of a structural shock to taxes on the vector of reduced form residuals u_t. The mapping from the reduced from residuals in period t to the structural shocks can be written as:

$$u_t = A_0 \varepsilon_t \tag{2}$$

Given knowledge of the relevant elements of A_0, Jordà (2005) shows that the impulse response at horizon h can be calculated as $B_1^{(h-1)} A_0$. On the other hand, Mertens and Ravn (2013) point out that the relevant elements of A_0 can be identified by treating narratively identified tax changes as proxies for the true structural variation in taxes. This is akin to using the narrative shocks as instruments for observed tax policy changes. The identification restriction is that the narrative shocks are uncorrelated with other structural shocks that may influence the economy, at least conditional on the lags of Z.[4]

Identification of A_0 directly follows Mertens and Ravn (2013). In particular, the identification problem can be written in terms of the reduced form residuals:

$$\begin{aligned} u_{T,t} &= \eta u_{X,t} + S_1 \varepsilon_{T,t} \\ u_{X,t} &= \zeta u_{T,t} + S_2 \varepsilon_{X,t} \end{aligned}$$

where $u_{T,t}$ and $\varepsilon_{T,t}$ are vectors containing the two reduced form and structural tax shocks, while $u_{X,t}$ and $\varepsilon_{X,t}$ are the remaining residuals and innovations for the other variables of interest (collected

[3] An alternative LP-IV setup would be: $\Delta^h Z_{t+h} = \alpha^h + \beta^h \Delta T_t + \Gamma^h X_{t-1} + u_{t+h}$ where Z are the same outcome variables of interest above, ΔT_t is the observed and potentially endogenous variation in tax policy (containing two tax variables) and X is a vector of controls, potentially including lagged values of Z. $\Delta^h Z_{t+h} = Z_{t+h} - Z_{t-1}$. ΔT_t would then be instrumented using the narrative "proxies" from Mertens and Ravn (2013). Because corporate and personal tax changes are correlated, we would need to be careful comparing the coefficient estimates with those in Mertens and Ravn (2013) (who explicitly consider the relationship between the two taxes when simulating the IRFs). More generally, Stock and Watson (2018) and Plagborg-Møller and Wolf (2021) discuss the equivalence of LP-IV and proxy VAR methods. For transparency and completeness, we also implement a LP-IV approach in the robustness section.

[4] Stock and Watson (2018) call this lag-lead exogeneity. This is a form of weak exogeneity where the narrative shocks are identified as orthogonal to current and future economic shocks but can, in principle, reflect past events.

together in a vector X).

In a nutshell, the approach amounts to using the narratively identified proxies as instruments for $u_{T,t}$ in the second equation above, constructing an estimate of $S_2\varepsilon_{X,t}$ and then using this as an instrument for $u_{X,t}$ in the first equation. This strategy identifies the contemporaneous impacts of the tax variables ζ and provides an associated estimate of the matrix η, although these latter terms have no structural interpretation because we are not identifying the effects of the non-policy shocks. Mertens and Ravn (2013) show that the structural parameters η, S_1, ζ and S_2 underlie the elements of A_0. The first two columns of the A_0 matrix, which refer to the tax shocks, are given by:

$$\beta_1 = \begin{pmatrix} I + \eta(I - \zeta\eta)^{-1}\zeta \\ (I - \zeta\eta)^{-1}\zeta \end{pmatrix} S_1$$

As mentioned above, because we are identifying two shocks — to corporate and personal income taxes — these policy instruments may well be correlated (as they are sometimes changed together in the same piece of legislation). We can estimate the effect of one policy holding the other constant, but we do not know how the two policies causally respond to each other in the data. As a result, in order to produce impulse response functions we need to take a stand on the precise policy experiment being conducted. Mathematically, this can be seen above. In order to pin down β_1 for construction of the impulse response function, a decomposition of S_1S_1' is required. As in Mertens and Ravn (2013), we use a Cholesky decomposition and order the tax rate being perturbed last for this decomposition. This restricts the direct contemporaneous effect of this shock on the remaining tax rate to be zero while still allowing for indirect effects via u_{2t}. Like Mertens and Ravn (2013), we show in the robustness section that our results are not sensitive to the ordering assumptions.

2.2 Data and estimation

In our benchmark specification, we use the same data as Mertens and Ravn (2013). The control variables on the right hand side of the sequence of local projections (1) include four lags of the following eight variables: (i) $APITR_t$, (ii) $ACITR_t$, (iii) $\ln\left(B_t^{PI}\right)$, (iv) $\ln\left(B_t^{CI}\right)$, (v) $\ln\left(G_t\right)$, (vi) $\ln\left(GDP_t\right)$, (vii) $\ln(DEBT_t)$, $(viii)$ PC_t.[5] The average personal and corporate tax rates are denoted by $APITR_t$ and $ACITR_t$, respectively, while $\ln\left(B_t^{PI}\right)$ and $\ln\left(B_t^{CI}\right)$ are the corresponding tax bases. Finally, $\ln\left(G_t\right)$ denotes government spending, $\ln(DEBT_t)$ stands for federal debt and GDP

[5] Montiel Olea and Plagborg-Møller (2021) demonstrate that lag-augmented local projections are particularly well-suited to draw robust inference about impulse responses at long horizons. Furthermore, they show that lag augmentation obviates the need to correct standard errors for serial correlation in the regression residuals.

is represented by $\ln(GDP_t)$. All variables, expect $APITR_t$ and $ACITR_t$, are expressed in real per-capita terms. The sample runs from 1950Q1 to 2006Q4 and the data are obtained from the replication files of Mertens and Ravn (2013). An initial estimation of the structural tax shocks using the variables (i) to (vii) above for $h = 0$ reveals that the estimated personal tax rate shock can be predicted by the lags of a principal component (denoted PC_t) obtained from a large quarterly dataset of Macro and Financial variables of the US economy.[6] Following Forni and Gambetti (2014), we add this principal component as eighth control variable on our LPs to ameliorate the effects of information insufficiency. Note that, as in Mertens and Ravn (2013), any additional variables of interest (that we will consider below) are added one by one to the benchmark model. These are personal consumption expenditures, non-residential investment, Research and Development (R&D) expenditure and productivity. The appendix provides a detailed description of the variables and data sources.

As noted above, we estimate the local projections via Bayesian methods. The posterior distributions are derived as in Miranda-Agrippino and Ricco (2015), who present an MCMC algorithm to approximate the posterior while taking into account the non-spherical nature of the forecast errors u_{t+h}. In eliciting priors for the LP parameters, we depart from Miranda-Agrippino and Ricco (2015) and use a flat prior on the LP coefficients. This implies that the results for impulse responses we present below are still largely data-driven.

3 Empirical results

In this section, we present the main results on the short-run and long-run effects of corporate and personal income tax changes using the local projections approach and data described in the previous section. We first focus on transitory versus persistent effects and show clear heterogeneity in the response of output and productivity to each type of tax shock. Then, we focus on the dynamic effects on investment, R&D expenditure and consumption to shed light on the potential mechanism. The final parts of the section discuss the forecast error variance decomposition and a set of robustness exercises, which are covered in more detail in the Appendix.

[6] The large dataset is obtained from Mumtaz and Theodoridis (2020). In order to implement the "structuralness" test of Forni and Gambetti (2014), we use up to 4 lags of the first 5 principal components obtained from this dataset.

3.1 Transitory versus persistent dynamics

Using the approach outlined above, we now present our baseline set of empirical estimates about the longer-term effects of cuts to corporate and personal income taxes. We first focus on the response of the average tax rate, GDP and productivity, with the latter being a key and novel focus of our analysis. We will then extend our empirical evidence to examine investment, R&D and consumption to shed light on the most likely mechanism driving the GDP responses. Each further variable is added to the benchmark data vector Z one at the time to avoid a sharp increase in the number of parameters to be estimated.

In Figure 1, we present our first set of main results. The figure contains two columns. On the left, we show the IRFs to a reduction in the average corporate tax rate. On the right, we show the results for a reduction in the average personal tax rate. The impact effect is normalized so that both shocks reduce their respective average tax rate by 1 percentage point in the first period. The solid red lines are the posterior medians and the shaded bands refer to 68% and 90% (Bayesian) credible intervals. Impulse response functions are computed using posterior draws of the coefficients A_0 and B_1. The solid blue lines come from the estimated structural model that will be presented, solved and estimated in Section 5.

The first row in Figure 1 reveals that, following a shock to corporate and personal income taxes, the average tax rates decline temporarily. The change in the average corporate tax rate (first column) loses significance after 8 quarters and goes back to zero after around 20 quarters. The changes in the average personal income tax rate are somewhat less persistent, losing significance after 6 quarters and touching zero after around 16 quarters. Despite the different method (i.e. local projections vs. VAR), these results largely replicate the findings in Figures 2 and 3 of Mertens and Ravn (2013), where the results are plotted for the first 20 quarters. In short, the estimated tax cuts are rather transitory.

The second row in Figure 1 shows the impulse response functions for the percentage response of real GDP. As expected, the IRFs for the first 20 quarters are very comparable to the main figures in Mertens and Ravn (2013). What is new is our estimate of the longer-term effects beyond quarter 20. Looking at the first column it is clear that, despite the transitory nature of the corporate tax reduction, there are very persistent effects on real GDP, whose short-run increase of 0.5% persists throughout the ten year period shown in the figure. In other words, the corporate income tax cut has disappeared after 5 years, but the effect on the level of economic activity is still sizable and significant after 8 years. The second column, however, reveals that the average personal tax rate

Figure 1: Response of the Tax Rate, GDP and Productivity to Corporate and Personal Tax Changes

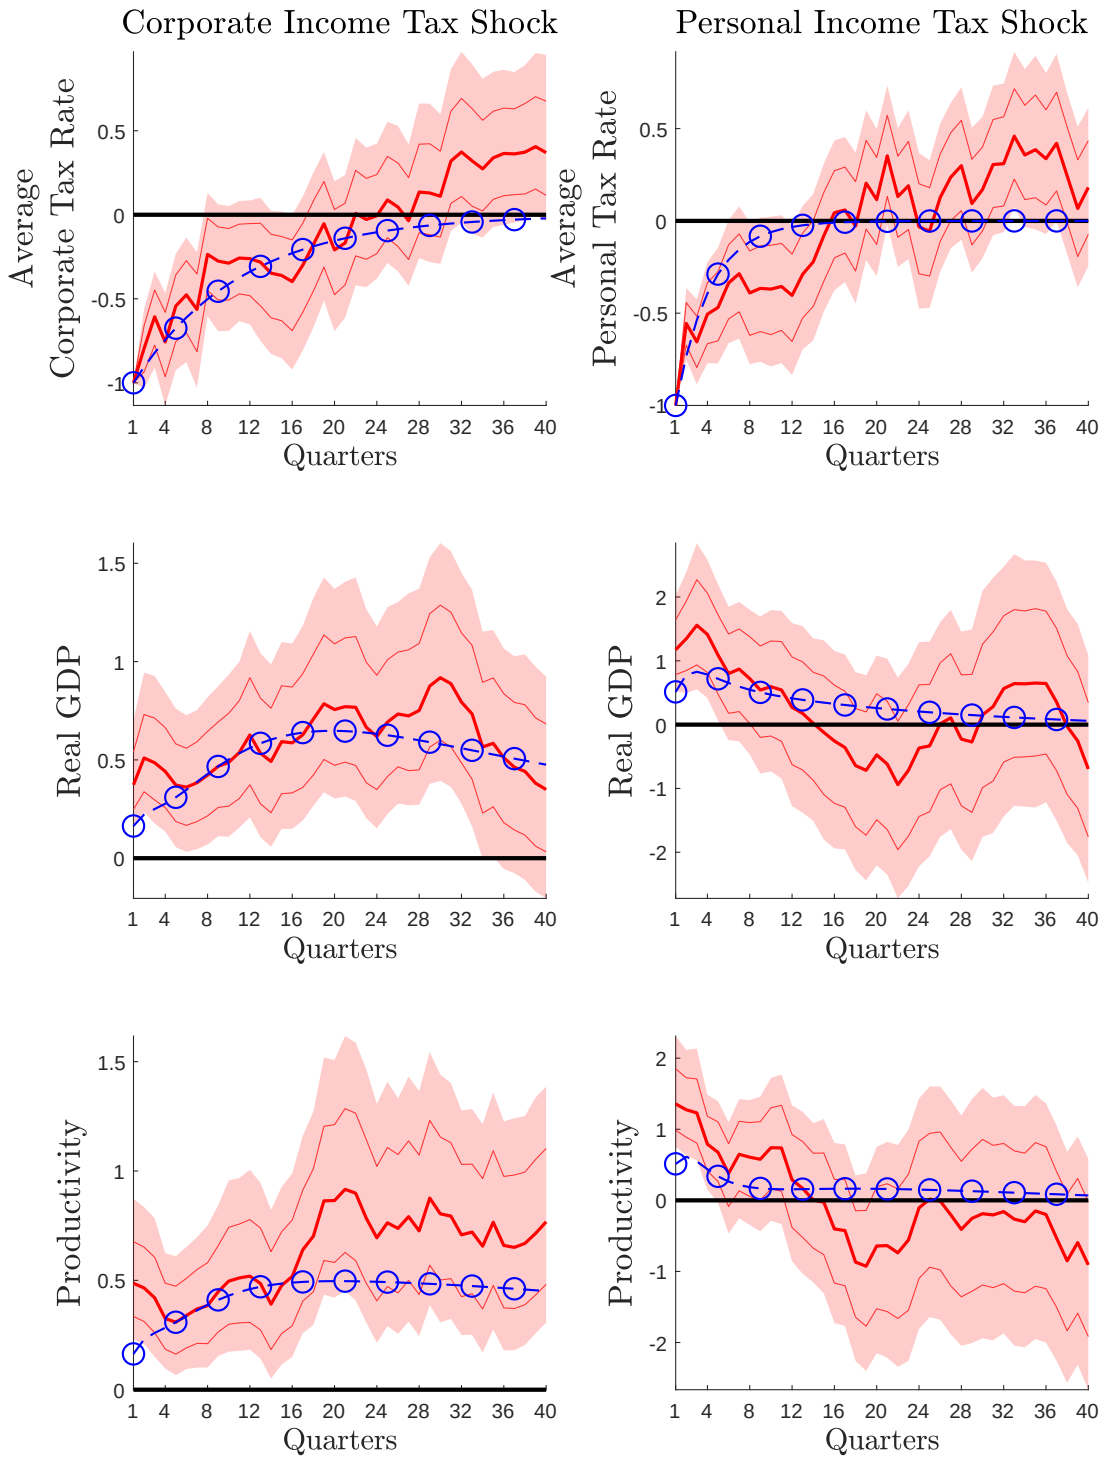

This figure shows the responses of the average tax rates, real GDP, and labor productivity to a 1% cut in the average rate of corporate income taxes (left column) and the average rate of personal income taxes (right column). Red shadow bands represent central posterior 68th and 90th credible sets. Blue lines with circles represent the impulse responses of the model in Section 4 evaluated at the posterior median of estimated model parameters. These model-produced estimates will be discussed later in the text.

cut does not produce such long-lasting dynamics. The underlying personal income tax cut is only slightly more transitory than the corporate tax cut but its effects on GDP are far less persistent and appear to die out already after two to three years after the shock hits.

A similar picture emerges for productivity, as shown in the third row of Figure 1. Both tax rate cuts boost productivity on impact, with the size of the initial response to a personal income tax cut being much larger than for a cut to corporate taxes. On the other hand, the effects of corporate tax cuts grow over time and remain significant even after 10 years. In sharp contrast, the response of productivity to a change in personal income tax rates is not statistically different from zero already after two years.

The clear difference in the short-run and long-run properties of the two taxes can be illustrated further by looking at the joint posterior distribution of transitory and persistent responses of the variable of interest in Figure 2. The top row refers to GDP while the bottom row represents productivity. We use 'transitory' (or shorter-term) to mean the effect estimated within 2 years after the shock, while 'persistent' (or longer-term) represents the dynamic effects estimated beyond the two year horizon.[7] The horizontal axis shows the response to a corporate tax cut, while the vertical axis shows the *associated* responses to a personal income tax cut.

For the sake of exposition, each panel also reports the 45° degree line, which is the locus of points where the effects of personal and corporate income taxes on the variable of interest are numerically identical. Posterior draws below (above) the 45° degree line indicate a larger impact of corporate (personal) taxes. The share of draws *below* the 45° degree line, which we denote as δ in Figure 2, can therefore be seen as a measure of the probability that corporate tax changes have larger effects. In Appendix Figure A.2, we show that our flat priors for the LP parameters imply very disperse joint prior distributions for the short-run and long-run effects of the two shocks on GDP and productivity. These joint prior distributions are centered at (0,0). As a result, there is an even chance that the (persistent and transitory) effects of one type of tax will dominate the other.

The left column indicates that for only about 6% of posterior draws do corporate income tax changes have larger short-term effects on GDP (top row) and productivity (bottom row) than personal income tax changes. In sharp contrast, the right column of Figure 2 reveals that corporate income tax cuts have significantly larger long-term effects than personal income tax cuts in about 93% of posterior draws for GDP and 98% for productivity. We conclude that the evidence of het-

[7] We obtain similar results using as a cutoff any other horizon within the first 4 years, which is the period by which the effects of personal (corporate) tax changes on output and productivity return to zero (start growing).

Figure 2: Joint Posterior Distribution of Shorter-term versus Longer-term Effects on GDP

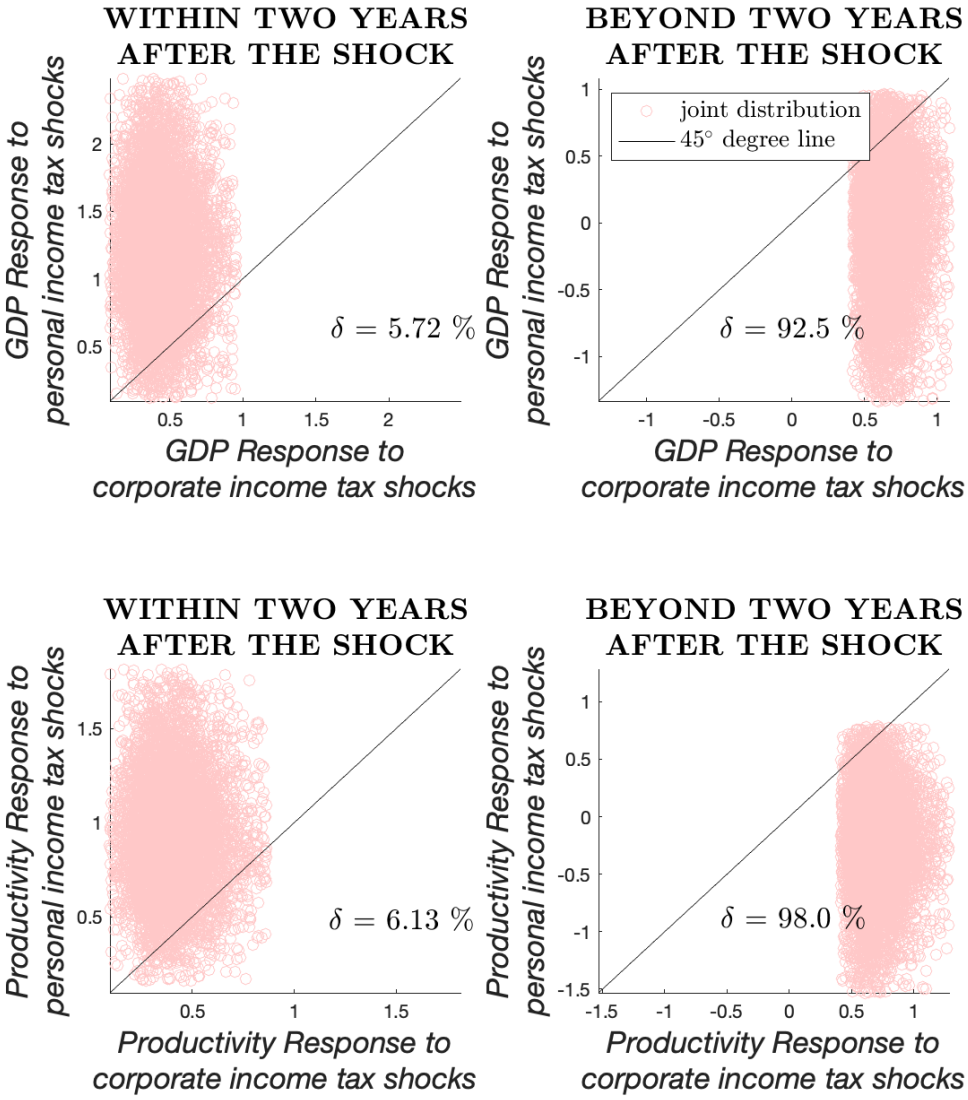

This figure shows the joint posterior distribution of the GDP responses (top row) and productivity responses (bottom row) to shocks to the personal income average tax rate (on the vertical axis) and the corporate income average tax rate (on the horizontal axis) in the shorter-term (left column) and longer-term (right column). The shorter-term (longer-term) column refers to the cumulated average responses of GDP and productivity to each shock over the quarters within (beyond) the first two years after the shock. The red dots show draws from the posterior distributions of estimated parameters based on the local projections. The black slope represents the 45° degree line, which is the locus of points along which the two shocks have GDP effects of exactly the same magnitude. Points above (below) the 45° degree line indicates the mass of the joint posterior distribution for which the effects of a personal income tax change are larger (smaller) than the effects of corporate income tax changes. Each δ statistic is the probability that the response to a corporate income tax change is larger than the response to a personal income tax change for each variable at the two different horizons.

erogeneous responses across both forecast horizons and type of income tax in Figure 1 is significant at conventional levels.

In Appendix A3, we report the responses of Total Factor Productivity (TFP), total hours and employment. Our theoretical model will not, however, feature an extensive margin so the employment IRF will not be used in the structural estimation in Section 5. The three main takeaways from this additional analysis are that: (i) the response of total hours and employment to a corporate income tax cut is typically modest and insignificant; (ii) in contrast, changes in personal income taxes have a significant but short-lived impact on total hours but muted effects on employment; (iii) the effects on TFP are very similar to those based on labour productivity.

In summary, in the short-run, the effects of personal income tax changes on output, productivity and hours worked are significantly larger than the effects of corporate income tax changes. Over longer horizons, however, the responses of output and productivity to corporate income tax cuts are large and highly significant. In contrast, the long-run effects of a personal income tax cut are indistinguishable from zero, both in economic and statistical terms. In the next section, we will extend our empirical analysis to investment, R&D expenditure and consumption in an effort to shed light on the possible mechanism behind the heterogeneity documented in Figure 1 and Figure 2.

3.2 On the mechanism

The impulse responses in Figures 1 and 2 replicate the results in Mertens and Ravn (2013) over the first 20 quarters, which is the horizon at which most previous literature stops. On the other hand, we have shown that there are significant longer-term effects of corporate tax cuts that persist beyond the typical IRF horizons presented in earlier work. These persistent dynamics are not, however, evident for the effects of personal tax cuts. In the rest of the paper we investigate what may explain these findings.

In this sub-section, we look at a number of additional variables that could offer insights on the transmission mechanism, especially at longer horizons. These are R&D expenditure, investment and consumption expenditure. The endogenous growth literature argues that R&D spending has the potential to generate persistent effects on both output and productivity. On the other hand, studies in the Real Business Cycle tradition emphasize the role of capital expenditure as an important propagation mechanism. Finally, given such longer-term output responses, we would also expect to see persistent effects on household expenditure for corporate tax changes, with short-lived effects from personal tax changes.

The findings are reported in Figure 3. The first row shows the impulse responses of R&D expenditure to a corporate tax cut (left column) and to a personal tax cut (right column). The second and third rows show the dynamic effects on investment and consumption respectively. Red lines represent medians and 68% credible sets of the impulse response posterior distributions. Shaded areas refer to 90% central intervals. As discussed in Section 2, each variable is added one at the time to our baseline dataset to avoid a sharp increase in our already richly parameterized local projections.

The evidence in the first row of Figure 3 suggests that the effects of corporate tax cuts (first column) on R&D are initially negligible but become significant at about one year after the shock. The increase is persistent and reaches a peak of 1.4% at quarter 18 before returning to zero after nine years. The effect also loses significance after six years. The response of investment to corporate tax changes (second row) is equally strong but its significance seems more short-lived. Finally, the consumption profile (third row) is similar to the pattern of the impulse responses of output and productivity in Figure 1. The significant and sustained rise in R&D seems a plausible candidate for explaining the persistent increase in productivity reported in Figure 1. In the next sections, we will explore this conjecture formally by developing and estimating a structural model with endogenous growth via R&D.

The estimated effects of a personal income tax cut (second column of Figure 3), paint a different picture. The response of R&D is never statistically different from zero while the change in investment is larger but far more transitory than for corporate tax changes. The effects on R&D and capital expenditure suggests that the sharp and short-lived increase in productivity after a personal income tax cut in Figure 1 does not come from firms' innovation activities. Later, we show that this is consistent with a short-run labor utilization story. Finally, the response of consumption in the bottom row largely inherits the shape of the GDP profile, as was the case for corporate taxes. This is consistent with the notion that corporate taxes raise labor income persistently, while personal taxes affect incomes only temporarily.

In summary, the evidence in this section is consistent with a transmission mechanism in which R&D responds to a corporate tax shock (but not to a personal tax shock) and this triggers an endogenous response of productivity, which in turn drives a persistent effect on GDP. In Appendix Figure 5, we provide further support for this interpretation by looking at sectoral real gross output from the U.S. Bureau of Economic Analysis's Industrial Accounts. We classify sectors in two groups based on their R&D intensity and estimate the heterogeneous effects of corporate and personal tax

Figure 3: The response of expenditure by component: R&D, Investment and Consumption

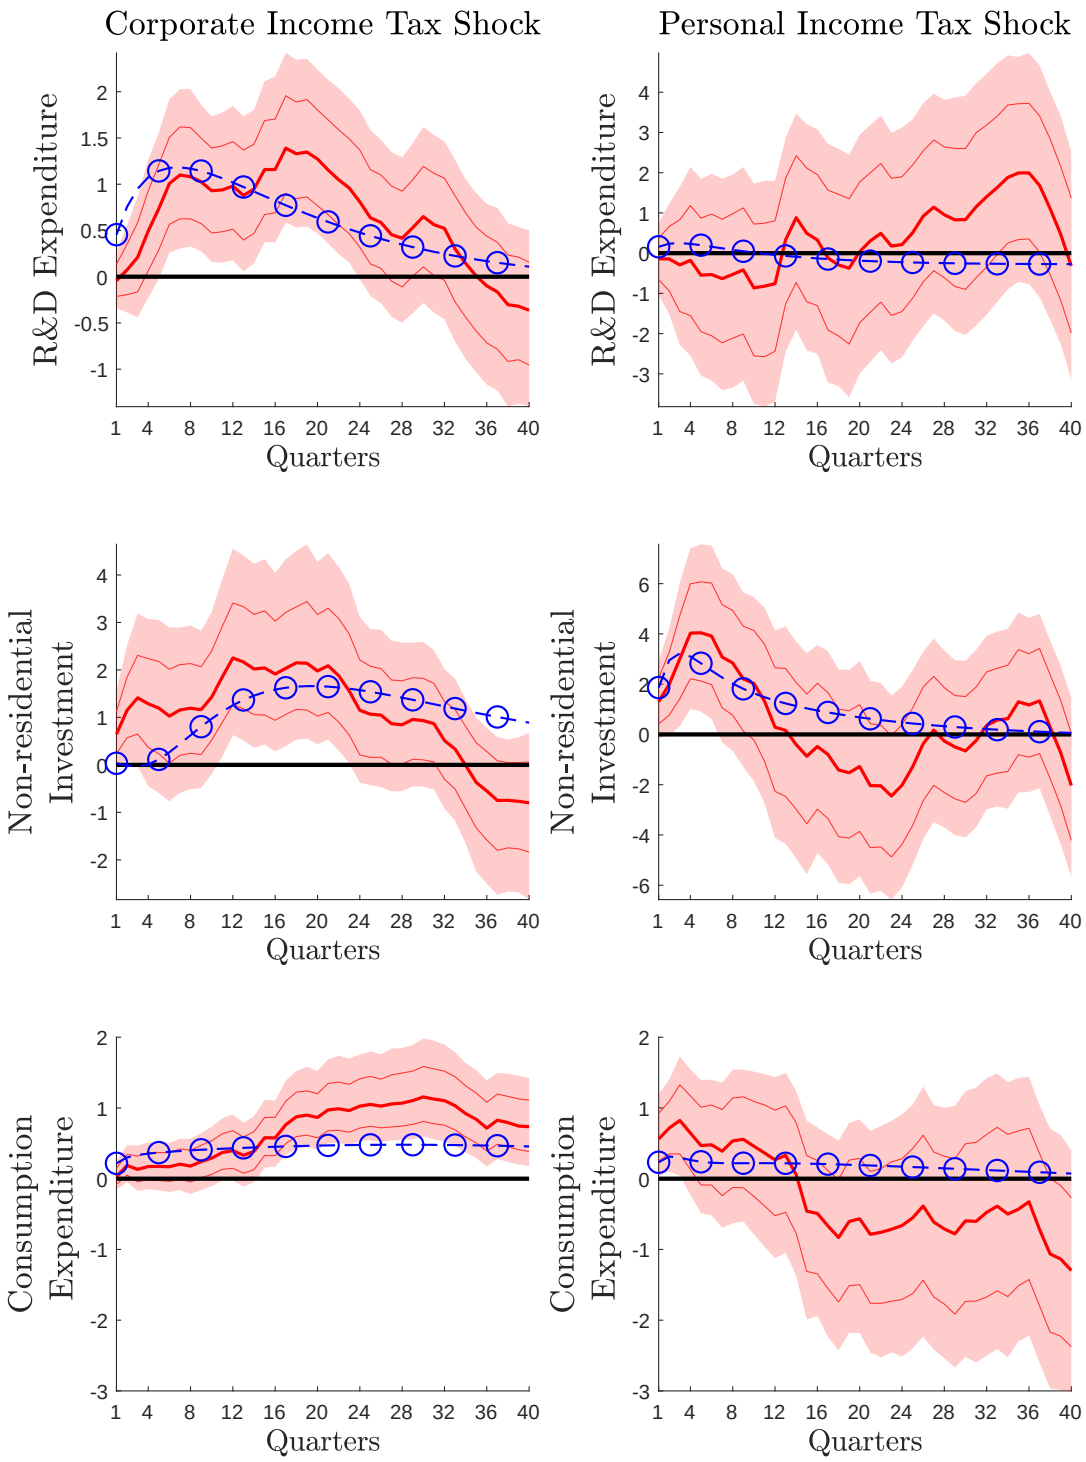

This figure shows the responses of R&D expenditure, non-residential investment and personal consumption expenditures to a 1% cut in the average rate of corporate income taxes (left column) and the average rate of personal income taxes (right column). Red shadow bands represent central posterior 68th and 90th credible sets. Blue lines with circles represent the impulse responses of the model in Section 4 evaluated at the posterior median of estimated model parameters. These model-produced estimates will be discussed later in the text.

cuts. The estimates reveal that the output response to corporate tax changes is significantly larger in sectors with high R&D intensity. In contrast, there is no statistical difference in the output responses of the two groups of sectors to personal tax changes.

3.3 Forecast Error Variance Decompositions

In this section, we use the LP estimates to assess the contribution of each shock to the variance of the endogenous variables at different forecast horizons. The results of this exercise are summarized in Appendix Figure A.4, which reports the median estimates and 90% central credible sets of the forecast error variance decomposition for the corporate income tax shock (in red) and the personal income tax shock (in blue).[8]

Two main results emerge. First, at the shorter horizon of one year, the contribution of both shocks is similar, accounting for around 20% of the variance of GDP and investment, as well as 15% to 20% of the variation in productivity and R&D spending. But as the forecast period increases, and especially at longer horizons, the contribution of the corporate income tax shock becomes dominant, peaking around year 8 and accounting for around 30% of the variance of GDP and consumption, and 20% to 25% for productivity, capital and R&D expenditure. In contrast, the contribution of personal income tax changes to longer-run fluctuations tend to be lower than 10%.[9]

3.4 Robustness

In this section, we briefly describe a variety of sensitivity analyses that confirm the robustness of our results. The full set of results are reported in the Appendix. We consider sensitivity to: (i) varying the lag length for the controls in Z, (ii) estimation via local projection instrumental variables, (iii) using the optimal prior strategy described in Giannone et al. (2015), (iv) including the defence news shock from Ramey (2011) as a further control, (v) changing the causal ordering of the two taxes as in Mertens and Ravn (2013) and (vi) estimation using Smooth Local Projections a la Barnichon and Brownlees (2019).

[8] By estimating the Mertens and Ravn (2013) VAR-type structure using local projections, we side-step practical issues associated with computing forecast error variance decompositions using local projection IV methods (see Plagborg-Møller and Wolf (Forthcoming)).

[9] These findings also echo results in earlier studies that focused more on short-term impacts. Mertens and Ravn (2012) find that Romer and Romer (2010) tax shocks explain around 20% of the in-sample variance of output at business cycle frequencies, which lines up well with our short-term results in Appendix Figure A.4. Cloyne (2013) finds that narrative-identified tax shocks in the U.K. account for around 20% of the variation in output in-sample and at the ten year horizon. McGrattan (1994) finds that labor taxes account for around 25% of the in-sample variance of output and capital taxes around 5%, again at business cycle frequencies and using a completely different VAR-based identification approach.

The results of all these sensitivity analyses are summarized in Appendix Figures A.5, A.6 and A.7. Figure A.5 shows the response of real GDP to corporate and personal tax cuts. The solid red line coincides with the benchmark specification above. The light red bands report the original 90% credible set. Overlaid are the results from each of the robustness checks noted above. The point estimates from all these exercises have the same shape as the baseline result and all lines lie well within the original confidence intervals. In Appendix Figures A.6 and A.7, we repeat these exercises for the variables related to the endogenous TFP mechanism discussed above. In particular, Figure A.6 shows that the response of R&D is not sensitive to all these different specifications. Finally, Figure A.7 tells a similar story about the robustness of the productivity response.

4 A Structural Model

The core of our model builds upon the framework developed by Anzoategui et al. (2019), which combines an otherwise standard sticky price model in the New-Keynesian tradition of Christiano et al. (2005) and Smets and Wouters (2007) with the endogenous productivity features introduced by Comin and Gertler (2006). We extend this framework along several dimensions, which give our model a chance to match the empirical IRFs estimated via local projections in Section 3. Most importantly, we introduce adjustment costs to productivity-enhancing expenditures, as well as variable labor utilization via an unobserved labor effort margin in the spirit of Galí and van Rens (2020) on the household side.

4.1 Production Sector and Endogenous Productivity

There exists a continuum of measure A_t of monopolistically competitive intermediate goods firms that each make a differentiated product: intermediate goods firm i produces output $Y_{i,t}$. The endogenous state variable A_t is the stock of intermediate goods adopted in production (equivalently, the stock of adopted technologies). The final goods composite is the following CES aggregate of individual intermediate goods:

$$Y_t = \left(\int_0^{A_t} (Y_{i,t})^{\frac{1}{\theta}} \, di \right)^{\theta} \tag{3}$$

with $\theta > 1$. Let $K_{i,t}$ be the stock of capital firm i employs, U_t capital utilization (described below), and $L_{i,t}$ the stock of labor employed. Then firm i produces output $Y_{i,t}$ according to the following

Cobb-Douglas technology:

$$Y_{i,t} = (U_t K_{i,t})^{\alpha} (L_{i,t})^{1-\alpha} . \tag{4}$$

Given a symmetric equilibrium for intermediate goods, we express the aggregate production function as:

$$Y_t = A_t^{\theta-1} \cdot (U_t K_t)^{\alpha} (L_t)^{1-\alpha}. \tag{5}$$

Endogenous total factor productivity growth is the result of expansion in the variety of adopted intermediate goods, measured by A_t. We next describe how R&D and adoption drive the dynamics of A_t.

4.2 R&D and Technological Adoption

The processes for creating and adopting new technologies are based on Comin and Gertler (2006) and Anzoategui et al. (2019). Let Z_t denote the total stock of discovered technologies. As above, A_t is the stock of adopted technologies, so $Z_t - A_t$ is the *un*adopted technology stock. R&D expenditures increase Z_t while adoption expenditures increase A_t. We distinguish between innovation and adoption to i) allow for realistic technology adoption lags; and ii) enable to model to jointly match the empirical dynamics of labor productivity and R&D expenditure.

R&D There are a continuum measure 1 of innovators that spend R&D-specific goods to create new intermediate goods. We describe below the technology for producing R&D goods. Let $X_{j,z,t}$ be R&D good expenditure by innovator j; the number of new technologies created by a unit of R&D expenditure, φ_t, is given by:

$$\varphi_t = Z_t^{\zeta+1} X_{z,t}^{\rho_z - 1}, \tag{6}$$

where $X_{z,t}$ is the aggregate amount of R&D expenditure and Z_t is the stock of technology, both of which an individual innovator takes as given. Following Romer (1990), the presence of Z_t reflects public learning-by-doing in the R&D process; as in Jones (1995), the degree of returns is parameterized by ζ.[10] We assume $\rho_z < 1$, which implies that more R&D expenditure in the aggregate reduces the efficiency of R&D at the individual level.

Let J_t be the value of an unadopted technology, $\Lambda_{t,t+1}$ the household's stochastic discount factor and $P_{z,t}$ the price of R&D goods. We can then express innovator j's decision problem as choosing

[10] The existence of a balanced growth path requires $\zeta = -\rho_z \frac{\theta-1}{1-\alpha}$, which we impose when estimating and simulating the model. See the online appendix for details on the balanced growth path of the model.

$X_{j,z,t}$ to solve:

$$\max_{X_{j,z,t}} E_t \left\{ \beta \Lambda_{t,t+1} J_{t+1} \varphi_t X_{j,z,t} \right\} - P_{z,t} X_{j,z,t}. \tag{7}$$

The optimality condition for R&D is then given by

$$E_t \left\{ \beta \Lambda_{t,t+1} J_{t+1} \varphi_t \right\} - P_{z,t} = 0$$

which implies

$$E_t \left\{ \beta \Lambda_{t,t+1} J_{t+1} Z_t^{1+\zeta} X_{z,t}^{\rho_z - 1} \right\} = P_{z,t} \tag{8}$$

The left side of Equation (8) is the discounted marginal benefit from an additional unit of expenditure, while the right side is the marginal cost. Finally, we allow for obsolescence of technologies. Let ϕ be the survival rate for any given technology. Then, we can express the evolution of the stock of technologies as:

$$Z_{t+1} = \varphi_t X_{z,t} + \phi Z_t \tag{9}$$

where the term $\varphi_t X_{z,t}$ represents the creation of new technologies. Combining equations (9) and (6) yields the following expression for the growth of new technologies:

$$\frac{Z_{t+1}}{Z_t} = Z_t^{\zeta} X_{z,t}^{\rho_z} + \phi. \tag{10}$$

Adoption We next describe how unadopted technologies become adopted, and therefore enter productive use. There is a competitive group of "adopters", indexed by k, who convert unadopted technologies into adopted ones. They buy the rights to the technology from the innovator, at the competitive price J_t, which is the value of an unadopted technology. They then convert the technology into use by employing adoption goods as an input (we describe the production technology for adoption goods below). This process takes time on average, and the conversion rate may vary endogenously. In particular, the rate of adoption depends positively on the level of resources devoted to adoption: an adopter succeeds in making a product usable in any period t with probability λ_t, which is an increasing and concave function of expenditure, $X_{k,a,t}$, according to the following expression

$$\lambda_t = \lambda \left(\frac{Z_t X_{k,a,t}}{\Psi_t} \right) \tag{11}$$

where $\lambda' > 0$, $\lambda'' < 0$.

To ensure the existence of a balanced growth path, we augment $X_{k,a,t}$ by a spillover effect from

the total stock of technologies Z_t (implying that the adoption process becomes more efficient as the technological state of the economy improves) and Ψ_t, which is a scaling factor that grows at the same rate as GDP on the balanced growth path. Once in usable form, the adopter sells the rights to the technology to a monopolistically competitive intermediate goods producer that makes the new product using the production function in Equation (4). Let $\Pi_{i,t}$ be the profits that an intermediate goods firm makes from producing a good under monopolistically competitive pricing. The price of the adopted technology, V_t, is the present discounted value of after-tax profits from producing the good, which is given by:

$$V_t = \left(1 - \tau_t^{CI}\right) \Pi_{i,t} + \phi E_t \left\{\beta \Lambda_{t,t+1} V_{t+1}\right\} \tag{12}$$

where τ_t^{CI} is the tax rate on corporate income (profits). An adopter's problem is choosing $X_{k,a,t}$ to maximize the value J_t of an unadopted technology, namely:

$$J_t = \max_{X_{k,a,t}} E_t \left\{-P_{a,t} X_{k,a,t} + \phi \beta \Lambda_{t,t+1} \left[\lambda_t V_{t+1} + (1 - \lambda_t) J_{t+1}\right]\right\} \tag{13}$$

where λ_t is as in Equation (11) and $P_{a,t}$ is the price of adoption goods. The first term in the Bellman equation reflects total adoption expenditures, while the second term stands for the discounted benefit: the probability weighted sum of the values of adopted and unadopted technologies. The first order condition for $X_{k,a,t}$ is

$$Z_t \lambda' \cdot \phi E_t \left\{\Lambda_{t,t+1} \left[V_{t+1} - J_{t+1}\right]\right\} = P_{a,t} \tag{14}$$

The term on the left is the marginal gain from adoption expenditures: the increase in the adoption probability, λ_t, times the discounted difference between the value of an adopted versus an unadopted technology. The right side is the marginal cost. Since λ_t does not depend on adopter-specific characteristics, we can sum across adopters to obtain the following relation for the evolution of adopted technologies:

$$A_{t+1} = \lambda_t \phi \left[Z_t - A_t\right] + \phi A_t \tag{15}$$

where $Z_t - A_t$ measures the stock of unadopted technologies. Note that $Z_t - A_t$ is also the measure of adopters at time t, which implies that the aggregate expenditure on adoptions goods, $X_{a,t}$, is given by $X_{a,t} = (Z_t - A_t) X_{a,k,t}$.

4.3 Households and the corporate sector

The representative household consumes, supplies labor, and receives dividends from the corporate sector (described below). There is habit formation in consumption. The model differs from the standard setup in the specification of labor supply. Households supply labor competitively, but choose employment N_{t+1} one period in advance, and face an adjustment cost $\frac{\psi_n}{2}\left(\frac{N_{t+1}}{N_t}-1\right)^2\Psi_t$ when changing employment. Following the realization of uncertainty in period t, the household chooses effort, e_t, and we assume that the effective labor supply is given by $L_t = N_t e_t$. The household's maximization problem and budget constraint are:

$$\max_{C_t, N_{t+1}, e_t} E_t \sum_{\tau=0}^{\infty} \beta^{\tau} \left\{ \log\left(C_{t+\tau} - bC_{t+\tau-1}\right) - \gamma_0 \frac{1+e_{t+\tau}^{1+\gamma}}{1+\gamma} N_{t+\tau} \right\}, \tag{16}$$

and

$$C_t = \left(1-\tau_t^{PI}\right) w_t L_t + D_t - \frac{\psi_n}{2}\left(\frac{N_{t+1}}{N_t}-1\right)^2 \Psi_t + T_t, \tag{17}$$

where C_t is consumption, D_t are dividends from the corporate sector, w_t is the real wage, and T_t are government transfers. [11] The symbol Ψ_t denotes a scaling factor that grows at the same rate as aggregate output, required to ensure that labor adjustment costs do not vanish along the balanced growth path.

The household's investment decisions are managed on their behalf by a representative investment fund that owns the physical capital stock, rents capital to intermediate goods firms, and chooses the rate of capital utilization, U_t, with associated cost $a(U_t)K_t$, where $a(U)$ is increasing and convex. The objective is to maximize lifetime dividends to households, discounted using the household's discount factor, $\Lambda_{t,t+1}$.

The investment fund owns all firms in the economy. They earn a return on capital services, receive all profits from the intermediate goods firms and collect and pay any corporate income taxes due to the government. Individual firms and innovators make the specific production, R&D and technological adoption decisions, as described earlier.

The dividend to the household in period t is given by overall corporate sector income minus corporate taxes due:

$$D_t = CI_t - \tau_t^{CI} TB_t^{CI}, \tag{18}$$

[11] Changes in dividend taxes form only a small part of the personal income tax measure in the Mertens and Ravn (2013) dataset. As a result, we abstract from explicitly modelling dividend taxes.

where CI_t is the income of the corporate sector overall:

$$CI_t \equiv \Pi_t + r_t^k U_t K_t - P_{I,t} I_t - a(U_t) K_t, \tag{19}$$

τ_t^{CI} is the tax rate on corporate income and the corporate income tax base, TB_t^{CI}, is defined as

$$TB_t^{CI} \equiv \Pi_t + r_t^k U_t K_t - \delta K_t. \tag{20}$$

Π_t are the profits from the intermediate goods firms, which are already net of R&D and adoption expenditures. We also assume that capital depreciation (but not utilization) is tax-deductible. These deductions for R&D and depreciation expenses are consistent with the US tax code during the sample period studied in the empirical section. I_t and $P_{I,t}$ are the quantity and price of investment respectively, K_t is the stock of capital (with rental rate r_t^k are), such that effective capital at time t is $U_t K_t$. Note that equation 12 implies that intermediate goods firms understand they will be required to pay taxes on their profits. The investment fund is therefore simply choosing I_t, K_t and U_t each period.

The law of motion for phyiscal capital follows the process:

$$K_{t+1} = (1 - \delta) K_t + I_t. \tag{21}$$

Taking the household and investment fund together, first order conditions are given by

1. Euler Equation for capital

$$P_{I,t} = \mathbb{E}_t \left\{ \beta \Lambda_{t,t+1} \left[\left(1 - \tau_{t+1}^{CI}\right) r_{t+1}^k U_{t+1} + \delta \tau_{t+1}^{CI} + (1 - \delta) P_{I,t+1} - a\left(U_{t+1}\right) \right] \right\}, \tag{22}$$

   where $\Lambda_{t,t+1} \equiv \frac{u_{c,t+1}}{u_{c,t}}$ is the household SDF between periods t and $t+1$.

2. Capital Utilization

$$\left(1 - \tau_t^{CI}\right) r_t^k = a'\left(U_t\right) \tag{23}$$

3. Employment

$$
\begin{aligned}
&-\psi_n u_{c,t}\Psi_t \frac{1}{N_t}\left(\frac{N_{t+1}}{N_t}-1\right)+ \\
&\beta\mathbb{E}_t\left(\gamma_0\frac{1+e_{t+1}^{1+\gamma}}{1+\gamma}+u_{c,t+1}\left(\left(1-\tau_{t+1}^{PI}\right)w_{t+1}e_{t+1}+\phi_n\Psi_{t+1}\left(\frac{n_{t+2}}{N_{t+1}^2}\right)\left(\frac{n_{t+2}}{N_{t+1}}-1\right)\right)\right)=0 \quad (24)
\end{aligned}
$$

4. Effort

$$
-\gamma_0 e_t^{\gamma}+u_{c,t}\left(\left(1-\tau_t^{PI}\right)w_t\right)=0 \quad (25)
$$

4.4 Fiscal Policy

The government's budget constraint is given by:

$$
\bar{G}\left(1+g_y\right)^t-T_t=\tau_t^{PI}w_tL_t+\tau_t^{CI}CI_t, \quad (26)
$$

where the tax rates τ_t^{CI} and τ_t^{PI} follow AR(1) processes (in logs):

$$
\log\left(\tau_t^x\right)=\left(1-\rho_{\tau x}\right)\bar{\tau}^x+\rho_{\tau x}\log\left(\tau_{t-1}^x\right)+\varepsilon_t^{\tau x}, \quad (27)
$$

for $x\in\{CI,PI\}$, with $\rho_{\tau x}\in(0,1)$, and $\varepsilon_t^{\tau x}\sim N(0,1)$ is i.i.d..

4.5 Rest of the model

We now summarize the main assumptions and derivations behind the rest of the structural model, which are standard.

Factor demands. Intermediate goods firm i chooses capital services $U_tK_{i,t}$, and labor $L_{i,t}$ to minimize costs, given the rental rate r_t^k, the real wage w_t and the desired markup ς. Expressed in aggregate terms, the first order conditions from firms' cost minimization problem are given by:

$$
\alpha\frac{MC_tY_t}{U_tK_t}=r_t^k, \quad (28)
$$

$$
\left(1-\alpha\right)\frac{MC_tY_t}{L_t}=w_t, \quad (29)
$$

where MC_t is the real marginal cost of production. We allow the actual markup ς to be smaller than the optimal unconstrained markup θ due to the threat of entry by imitators, as is common in

the literature (e.g. Aghion and Howitt (1998)).

Investment good producers. There are three types of investment goods in the economy: investment goods used to produce capital, R&D goods and adoption goods. Competitive producers use final output to produce these goods which they sell to households (investment), R&D firms (R&D goods) or adopters (adoption goods). Following Christiano et al. (2005), we assume flow adjustment costs of investment for the three types of goods. The adjustment cost functions are:

$$f_I\left(\frac{I_t}{(1+g_y)\,I_{t-1}}\right),\ f_z\left(\frac{X_{z,t}}{(1+g_y)\,X_{z,t-1}}\right),\ \text{and}\ f_a\left(\frac{X_{a,t}}{(1+g_y)\,X_{a,t-1}}\right),$$

where each function is increasing and concave, with $f_x(1) = f_x^{'}(1) = 0$ and $f_x^{''}(1) > 0$; and I_t, $X_{z,t}$ and $X_{a,t}$ are new capital, R&D and adoption goods produced in period t; g_y is the steady state growth of output. The first order conditions for each of the types of good, $x \in \{I, X_z, X_a\}$, are:

$$P_{x,t} = 1 + f_x\left(\frac{x_t}{(1+g_y)\,x_{t-1}}\right) + \frac{x_t}{(1+g_y)\,x_{t-1}} f_x'\left(\frac{x_t}{(1+g_y)\,x_{t-1}}\right) - E_t\left[\Lambda_{t,t+1}\,(1+g_y)\left(\frac{x_{t+1}}{(1+g_y)\,x_t}\right)^2 f_x'\left(\frac{x_{t+1}}{(1+g_y)\,x_t}\right)\right] \quad (30)$$

Price Setting. Nominal prices are set on a staggered basis following the Calvo adjustment rule. Denoting by ξ_p the probability that a firm cannot adjust its price, by $\hat{\pi}_t$ the inflation rate and by $\widehat{mc}_t$ the marginal cost in log-deviation from steady state, the Phillips curve reads:

$$\hat{\pi}_t = \kappa_p \widehat{mc}_t + \beta \mathbb{E}_t[\hat{\pi}_{t+1}], \quad (31)$$

with slope $\kappa_p = \frac{(1-\xi_p\beta)(1-\xi_p)}{\xi_p}$.

Monetary Policy The nominal interest rate $R_{n,t+1}$ is set according to a Taylor rule:

$$R_{n,t+1} = \left(\left(\frac{\pi_t}{\bar{\pi}}\right)^{\phi_\pi} \frac{L_t}{\bar{L}}\right)^{\phi_y} R_n\Bigg)^{1-\rho^R} (R_{n,t})^{\rho^R}, \quad (32)$$

where R_n is the steady state nominal rate, $\bar{\pi}$ the target rate of inflation, L_t total effective labor supply and $\bar{L}$ steady state labor supply; ϕ_π and ϕ_y are the feedback coefficients on the inflation and capacity utilization gaps, respectively. Following Anzoategui et al. (2019), we use the labor supply

gap instead of the output gap in the Taylor rule.

Resource Constraint. Finally, the aggregate resource constraint is given by:

$$Y_t = C_t + \sum_{x=\{I,X_z,X_a\}} \left[1 + f_x\left(\frac{x_t}{(1+g_y)\,x_{t-1}}\right)\right] x_t + a\left(U_t\right) K_t + \frac{\psi_n}{2}\left(\frac{N_{t+1}}{N_t} - 1\right)^2 \Psi_t + \bar{G}. \quad (33)$$

Given all these ingredients, we are in the position now to solve the model and compute its impulse responses to both personal and corporate income tax shocks. These will be used in the next sections to estimate the model structural parameters and perform a counterfactual analysis intended to highlight the possible drivers of the LP-based empirical evidence in Section 3.

5 Structural estimation

In this section, we show that the theoretical model outlined above can rationalize all our empirical findings. To do so, we estimate the structural model using a limited-information Bayesian approach and show that it is able to account jointly for the longer-term effects of corporate income tax changes and the shorter-term effects of personal income tax changes reported in Section 3 on the basis of LPs. In the next section, we perform a series of counterfactual exercises that will shed light on the mechanism behind both sets of results, across forecast horizons and across types of taxes.

5.1 Approach

We estimate the structural model using the limited-information Bayesian approach described in Christiano et al. (2010). We refer to the vector of structural parameters in the theoretical model as Υ and to the associated impulse responses as $\Phi\left(\Upsilon\right)$. The structural parameters are estimated by minimizing the distance between the theoretical model impulse responses, $\Phi\left(\Upsilon\right)$, and the median of the empirical LP impulse response posterior distributions from Section 3, denoted by $\hat{\Phi}$, to both tax shocks.

The limited-information approach fulfils our desire to focus on the response of the economy to corporate and personal tax cuts, and to isolate the theoretical mechanism(s) that are most likely to drive the empirical findings of Section 3. It is therefore important that the estimated parameters maximize the likelihood that the structural model generates the data conditional on both income tax shocks. We will then be able to conduct a series of counterfactual experiments where we artificially

change the value of one structural parameter at a time to evaluate the importance of different channels for explaining the empirical evidence from our LPs in Section 3.

To implement this approach, we first set up the quasi-likelihood function as follows:

$$F(\hat{\Phi}|\Upsilon) = \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} |V|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\left(\hat{\Phi} - \Phi\left(\Upsilon\right)\right)' V^{-1}\left(\hat{\Phi} - \Phi\left(\Upsilon\right)\right)\right)$$

where N denotes the number of elements in $\hat{\Phi}$ and V is a weighting matrix. In our application V is a diagonal matrix with the posterior variance of $\hat{\Phi}$ on the main diagonal. Denoting by $p\left(\Upsilon\right)$ the prior distributions, the quasi-posterior distribution is defined as:

$$F\left(\Upsilon|\hat{\Phi}\right) \propto F(\hat{\Phi}|\Upsilon)p\left(\Upsilon\right)$$

We use a random walk Metropolis Hastings algorithm to approximate the posterior distribution. The total number of iterations is set to 1,100,000 and we save every 50th draw after a burn-in of 100,000.[12] As is common in limited information estimation approaches, we partition the set of structural parameters into a set that will be calibrated and a set that will be estimated. In Table 1, we list the key objects that we calibrate. The discount factor, the depreciation rate for capital and the capital share are set at common values: 0.99, 0.02 and 0.35 respectively. The markup is calibrated to target the steady state share of profits in GDP. The government spending share and the steady state tax rates are calibrated to sample averages in the data. The coefficients of the Taylor Rule are those estimated in Anzoategui et al. (2019). We calibrate the steady technology adoption rate $\bar{\lambda}$ to 0.05 (quarterly), implying a diffusion lag of 5 years, in line with the evidence in Comin and Hobijn (2010). Following Wen (2004), the employment adjustment cost parameter is set to $\psi_N = 0.25$.

Estimates of the remaining structural parameters are shown in Table 2. Our goal is to estimate the parameters that govern the strength of the main transmission mechanisms in the theoretical model. The data are then being used to inform us about which of these channels are most likely important for explaining our empirical results. The prior distributions and the posterior estimates are discussed below.

For estimation, the vectors $\hat{\Phi}$ and $\Phi(\Upsilon)$ contain the main empirical and theoretical impulse

[12] The starting values of the parameters are obtained by maximising the log posterior using the covariance matrix adaption algorithm (CMA-ES). Then, an initial run of the Metropolis algorithm is used to approximate $\text{var}\left(\Upsilon\right)$. A scaled version of $\text{var}\left(\Upsilon\right)$ is used to calibrate the variance of proposal distribution for the main run of the Metropolis algorithm. We choose the scaling so that the acceptance rate is about 20%.

Table 1: Calibrated Parameters

Parameter	Description	Value	Source
Preference & Households			
g_y	100*SS GDP growth rate	0.45	
β	Discount factor	0.99	
ψ_N	Employment adjustment	0.25	Wen (2004)
Technology			
GY	Government spending/GDP	0.16	
α	Capital share	0.35	
δ	Capital depreciation.	0.02	
ς	Markup	1.09	Profits/GDP=8%
$\bar{\lambda}$	SS technology adoption rate	0.05	Anzoategui et al. (2019)
Taxes			
$\bar{\tau}^{CI}$	SS Corp. Tax	0.19	Sample average
$\bar{\tau}^{PI}$	SS Lab. Tax	0.3	Sample average
Monetary Policy			
ρ_r	Smoothing	0.83	Anzoategui et al. (2019)
ϕ_y	Output	0.39	Anzoategui et al. (2019)
ϕ_π	Inflation	1.64	Anzoategui et al. (2019)

response functions for the two tax rates and the two sets of responses for R&D, investment, consumption, GDP, hours worked and labor productivity. Note that, by simultaneously targeting the effects of both corporate and personal taxes, we are attempting to hit a number of key moments simultaneously over the shorter- and longer-term.

5.2 Prior predictive analysis

Before presenting and evaluating the estimation results, it is useful to examine some of the predictions that the model is able or unable to generate about the effects of corporate and personal income tax changes. In this section, we therefore discuss the assumed prior distributions for the structural parameters Υ and the implied range of outcomes $\Phi(\Upsilon)$ that, a priori, the key variables are more likely to cover. As we will see below, our priors are centered on an economy with no long-run effects and a-cyclical productivity.

In Table 2, we list the parameters to be estimated and the prior distributions chosen. The table also reports moments from the estimated posterior distribution, which will be discussed in the next section. Prior distributions are chosen to be diffuse but centered on values typically found

Table 2: Estimated Parameters

Parameter	Description	Prior			Posterior	
		Distr	Mean	Std. Dev.	Median	90% int.
Preference & HHs						
h	Consumption habit	beta	0.5	0.2	0.43	[0.15, 0.75]
γ	Inverse effort elasticity	gamma	1	0.5	0.28	[0.11, 0.67]
Frictions & Production						
f_a''	Adoption adjustment	normal	4	1.5	4.36	[1.86, 6.81]
f_z''	R&D adjustment	normal	4	1.5	6.67	[4.75, 8.71]
f_I''	Investment adjustment	normal	4	1.5	0.29	[0.05, 1.31]
ψ_u	Capital utilization adjustment	beta	0.6	0.15	0.49	[0.33, 0.69]
ξ_p	Calvo prices	beta	0.5	0.2	0.21	[0.07, 0.37]
Endogenous Technology						
θ-1	Dixit-Stiglitz parameter	gamma	0.15	0.1	0.42	[0.32, 0.53]
ρ_λ	Adoption elasticity	beta	0.5	0.2	0.59	[0.54, 0.64]
ρ_Z	R&D elasticity	beta	0.5	0.2	0.15	[0.09, 0.23]
$1-\phi$	Knowledge depreciation	beta	0.05	0.05	0.11	[0.06, 0.19]
Shocks						
$\rho_{\tau CI}$	Corporate taxes AR	beta	0.7	0.2	0.91	[0.89, 0.92]
$\rho_{\tau PI}$	Labour taxes AR	beta	0.7	0.2	0.73	[0.68, 0.78]

in the literature. The prior means for the more standard parameters, such as habit formation, the Calvo probability that governs price stickiness and investment adjustment costs are consistent with common estimates and priors used in earlier empirical contributions, such as Smets and Wouters (2007). The priors for the tax processes assume that the tax rates are adjusted smoothly over time and follow Leeper et al. (2010).

There are a number of parameters that are specific to our R&D, adoption and utilization mechanisms. Empirical estimates of the elasticity of patenting to R&D expenditures, analogous to ρ_Z in the model, range widely (Danguy et al., 2013) but are generally found to be lower than 1. Accordingly, we use a beta prior centered at 0.5. We use the same prior for the adoption elasticity ρ_λ. The prior mean for the Dixit-Stiglitz parameter θ implies an elasticity of substitution across goods of 7.6, consistent with the estimates in Broda and Weinstein (2006).[13] We center the parameter for knowledge depreciation at 0.05, at the lower end of empirical estimates (Li and Hall, 2020). To avoid tilting the balance in favor of any particular adjustment cost mechanism, we use the same prior for R&D and adoption adjustment costs as for physical capital investment adjustment costs (which is also the prior on R&D adjustment costs used by Bianchi et al., 2019). We are not aware of existing estimates of the (inverse) elasticity of effort, γ. Accordingly we choose a relatively uninformative prior centered around what we consider to be a conservative value (as we will show in the prior predictive analysis discussed below).

Prior predictive analysis involves drawing a candidate Υ_i from the marginal prior distributions of the parameters. For each candidate Υ_i, the associated set of impulse response functions, $\Phi(\Upsilon_i)$, are computed. This process is repeated 100,000 times, thereby generating a distribution of impulse responses.[14] Prior predictive analysis allows us to elicit a number of useful insights. First, we can see the range of different possible outcomes that the model is likely to generate given our prior distributions. Second, we can see what our priors imply about the shorter and longer-term effects of tax changes.

In Appendix Figures A.10 and A.11, we report the distributions of the model impulse responses implied by our prior distributions. The solid (shaded) red lines report the median and central 68% (90%) prior credible sets of the IRF *prior* distribution. The blue line with circles refers to the impulse responses of the model evaluated at the estimated *posterior* median of the parameters. The main takeaway from this exercise is that our prior distributions give far more weight to an

[13] Anzoategui et al. (2019) calibrate this parameter to 1.35; our prior is relatively conservative given that a higher θ implies a larger role for the endogenous productivity mechanism.

[14] For more details on prior predictive analysis, we refer interested readers to Leeper et al. (2017).

economy in which the effects of both personal and corporate income taxes are quite short-lived and productivity is virtually a-cyclical. As we will show in the next section (and can already be seen from Appendix Figure A.10 and Appendix Figure A.11), the posterior distributions paint a quite different picture.

In Appendix Figure A.12, we provide a different way of visualizing our prior assumptions. This exercise is similar to one in Figure 2, but uses the structural model and the prior and posterior distributions from Table 2. As before, the figure compares the short-term effect of each shock relative to their long-term effect for the same outcome variable. GDP is shown in the top row and productivity in the bottom row. In the first column, a mass of draws above the 45 degree line implies that personal income tax changes have a larger shorter-term effect than corporate income tax changes. In the second column, a mass below the 45 degree line suggests that corporate tax cuts have a larger longer-term effect. The cloud of light grey (blue) dots refers to draws from the joint *prior* (posterior) distribution of the impulse responses implied by the (estimated) structural model. The main message from Appendix Figure A.12 is that the grey dots are evenly spread either side of the 45 degree line. The two tax shocks are therefore put on equal footing a priori: they both have a similar chance of generating larger effects on output and productivity at shorter- and longer-horizons.

In summary, our prior distributions for the structural parameters are centered around a relatively standard calibration, according to which: (i) the economy does not feature any long-run effects and (ii) both types of tax cuts produce a largely a-cyclical productivity response.

5.3 Estimation results

In this sub-section, we present the estimates of the model using the approach discussed above. The final two columns in Table 2 report the median and the central 90% credible set for the posterior distributions of the key parameters of interest. The impulse responses from the estimated model (evaluated at the median values of the parameters in Table 2) are shown in Figures 1 and 3 as blue lines with circles.

Starting with fiscal policy in the last two rows of Table 2, the processes for the tax rates evolve smoothly over time, with the changes in the corporate income tax rate being less short-lived than those for the personal tax rate. Still, as shown in Figure 1, both tax rates return to zero over the forecast horizon, with their estimated tax profiles lining up closely with their LP counterparts from Section 3.

The estimates of the parameters on R&D and technological adoption are reported in the third block of Table 2 and are largely consistent with the available evidence. The posterior median of θ is close to the calibrated value in Anzoategui et al. (2019). The depreciation of ideas is at the higher end of existing estimates but in line with both the rates used by the U.S. Bureau of Economic Statistics in calculating GDP (Li and Hall, 2020) and the estimates in Ma and Samaniego (2020). Adjustment costs for R&D spending are higher than for adoption. The inverse elasticity of effort is close to the value of 0.3 that Galí and van Rens (2020) calibrate to match second moments of the U.S. labor market.

The estimation also puts some weight on investment adjustment costs, habit persistence and price stickiness in the top two parts of Table 2. Interestingly, however, by incorporating an endogenous growth mechanism (via R&D and adoption), our estimates seem to downplay these more 'traditional' ways of generating endogenous persistence and amplification relative to earlier work. In particular, traditional investment adjustment costs are estimated to be much smaller than the value reported elsewhere in the literature (Christiano et al., 2005, Smets and Wouters, 2007, Justiniano et al., 2010). Unlike conventional medium-scale models, however, our framework features a range of additional sources of endogenous persistence. The estimation therefore appears to favor much larger adjustment costs on R&D and technological adoption than on physical capital investment, consistent with the evidence from aggregate data in Bianchi et al. (2019) and from firm-level data in Bernstein and Nadiri (1989) and Bond et al. (2005).

In Section 7, we estimate a restricted version of our structural model in which we switch off all the endogenous growth mechanisms. The estimates of physical capital investment adjustment costs in this restricted specification become much larger and in line with those reported by Christiano et al. (2005), Smets and Wouters (2007) and Justiniano et al. (2010), among others. We interpret this finding as suggestive evidence that the omission of R&D spending and technological adoption might distort inference on the role of physical capital investment adjustment costs in the propagation of macroeconomic shocks.

5.4 Discussion

The results of the previous section make clear that the posterior estimates in Table 2 allow the model to closely line up with the LP evidence in Section 3. This can be seen in Figure 1 and Figure 3 by comparing the estimated model IRFs in the dashed blue lines with circles to the LP estimates in red. It is worth emphasizing again that the two taxes have very different dynamic effects and we

Figure 4: Additive decomposition of model log productivity

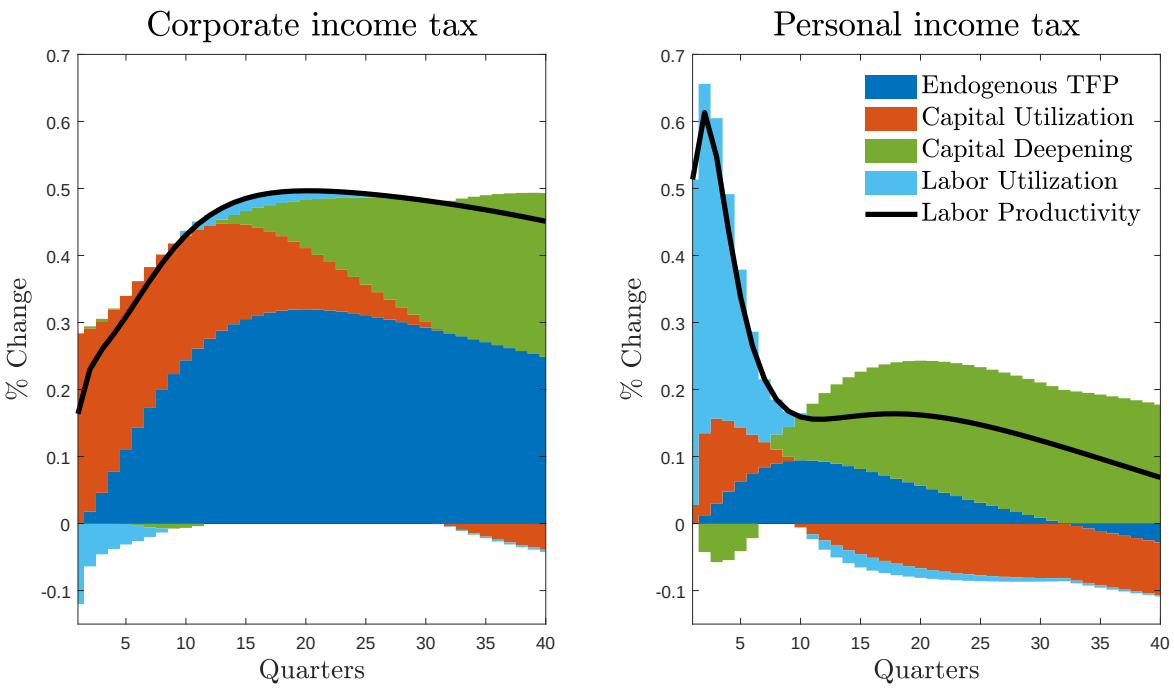

The left (right) panel of this figure shows an additive decomposition of the estimated model response of log labor productivity to a corporate (personal) income tax shock. The solid line is the total response of labor productivity and the shaded areas show the contribution of each factor.

are trying to match the two sets of responses *jointly*. What are the theoretical mechanisms behind these effects?

We answer this question through the lens of our structural model. We begin with a discussion of the effects of a temporary cut in corporate income taxes. The key part of the production process occurs via intermediate goods firms (since these goods are ultimately combined to generate the homogeneous final consumption good, see equation 3). These firms have adopted specific technologies and create output using capital and labor services. A corporate tax cut raises the ex-post profits of these firms.

On the one hand, the model features the usual corporate tax cut channels: a lower corporate tax rate raises the return on capital and encourages production. The tax cut also raises the marginal benefit of capital utilization, which helps generate sizable short-run amplification. A corporate tax cut therefore stimulates investment, capital utilization, employment (because of the complementarity with capital in production) and output. But (as shown in Figure 5, discussed below) these effects tend to be rather transitory and it is challenging to generate significant endogenous persistence that looks like our evidence in Section 3.

In our context, however, the corporate tax cut also encourages firms' R&D expenditure and technological adoption because the tax cut raises after-tax profits, which directly increases the value of an adopted technology (as made clear by equation 12). In turn, this raises the value of an unadopted technology (equation 13). As a result, R&D spending increases, generating new innovations, and encouraging further adoption. Productivity in the economy rises over time (Figure 1) and this persistent increase in productivity leads to a large and persistent rise in GDP (also shown in Figure 1).

It should be noted that these endogenous growth effects get amplified through other more standard channels, such as variable capital utilization. Non-residential investment also increases in Figure 3. Turning to the dynamics of household expenditure, the additional accumulation of capital may, in the short-run, lead to lower consumption, although this also depends on other factors such as the response of labor supply, habit formation and the endogenous growth effects. In our estimated model, as well as in the LP evidence, consumption grows steadily over the impulse response period (see Figure 3).

The mechanism behind a cut in personal income taxes is very different, as these mostly affect households but have no *direct* influence on firms. The reduction in personal income taxes encourages households to supply more labor but, due to employment adjustment costs, the increase in hours worked is gradual. Labor utilization (via increased effort), however, responds contemporaneously. As a result, households can still reap the benefits of higher after-tax wages, even in the face of employment adjustment costs. This implies that labor productivity increases on impact in response to a cut in personal income taxes. This channel, however, does not directly affect firms' incentives to innovate, and therefore tends to persist only as long as the distortionary personal income tax rate is below its steady state level.

Figure 4 illustrates the effects of each of the tax shocks on the economy through the lens of the model using an additive decomposition of the response of (log) labor productivity. From the aggregate production function (Equation 5), the change in log labor productivity is given by $\Delta \log \frac{Y}{N} = \Delta \log A^{\theta-1} + \Delta \log U^{\alpha} + \Delta \log \left(\frac{K}{N}\right)^{\alpha} + \Delta \log e^{1-\alpha}$ (respectively the contributions of adoption, capital utilization, capital deepening and labor utilization). The panel on the left shows that more than half of the level of labor productivity to the corporate income tax shock at 40 quarters is accounted for by the endogenous TFP term, with capital deepening accounting for the reminder of the long-run effect. The figure understates the contribution of the endogenous TFP channel, since capital deepening is driven largely by the complementarity between endogenous TFP and physical

capital, as we show in the counterfactual exercises in the next section. The panel on the right shows that variable labor utilization is the most important driver of the response of labor productivity to the personal income tax shock, and the fact that this channel is transitory in nature explains the overall lack of persistence in the response.

6 Counterfactual analysis

In this section, we perform a series of counterfactual exercises that highlight the central importance of pro-cyclical productivity for explaining our results at different forecast horizons. More specifically, we will show that R&D and technological adoption make the response of productivity to corporate tax changes pro-cyclical in the long-run, while variable labor utilization makes the response of productivity to personal tax changes pro-cyclical in the short-run. Whenever productivity is either counter-cyclical or a-cyclical —which is the case when we switch off the endogenous growth block of the model and variable factor utilization— the structural model becomes much closer to a standard New-Keynesian framework and misses both the longer-term effects of corporate tax changes and the short-term effects of personal tax changes on output.

6.1 No endogenous growth

In the first counterfactual experiment, we switch off the endogenous productivity part of the model. More specifically, we solve a version of the model in Section 4 with no R&D or adoption, and compute the impulse responses of the structural model using the posterior medians reported in Table 2 for all other parameters. In Figure 5, we show the results of this counterfactual exercise. As in all previous charts, the left column refers to corporate income tax changes while the right column refers to personal income tax changes. For the sake of exposition, we only report the counterfactual impulse responses of the key variables: GDP (top row), productivity (middle row) and R&D spending (bottom row). The blue lines with circles show the impulse responses from the estimated baseline model and replicate the blue lines with circles in Figures 1 and 3. The black lines with crosses correspond to the counterfactual exercise with no endogenous productivity. The green lines with diamonds show the experiment with no variable factor utilization, which is discussed in the next sub-section.

The black lines with crosses in the third row of Figure 5, show that in the counterfactual model with no endogenous growth, the response of R&D expenditure is, by construction, always zero. In

Figure 5: Counterfactuals: Response of GDP, Productivity and R&D

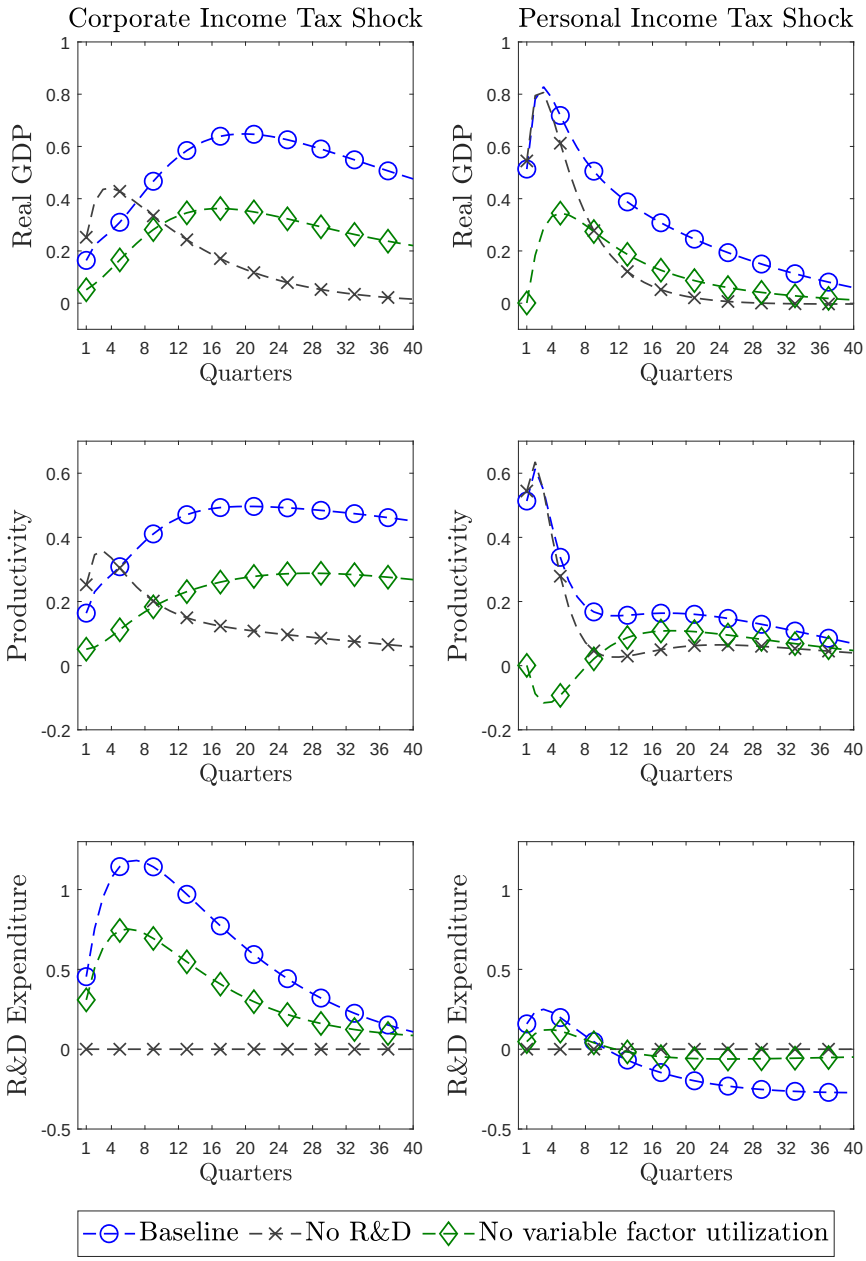

This figure shows the results of two counterfactual exercises. The first set of simulations shuts down the endogenous growth part of the model ("no R&D") while the second experiment turns-off the utilization margin for both labour and capital ("no variable factor utilization"). The baseline model results are reported as blue lines with circles, which replicate the results from Figure 1 and Figure 3. The black lines with crosses show the impulse responses from the estimated model after shutting down the endogenous growth mechanism. The green line (diamonds) shows the impulse responses from the estimated model after shutting down variable factor utilization.

this experiment, the model is similar to other standard New Keynesian models with distortionary taxation (except for variable labour utilization). Figure 5 also makes clear that this version of the model is unable to match the significant long-run response of either output (top row) or productivity (middle row) to a corporate income tax cut. Once R&D expenditure and technological adoption are switched off, investment (and to a lesser extent habit formation) becomes the main propagation mechanism that generates persistence in the model; but this is insufficient to prevent the effects of corporate tax changes from returning to zero by the end of the forecast period. The absolute magnitudes are also much smaller over most horizons, especially in the long-run. The response of productivity is still mildly pro-cyclical in the short-run but this is now entirely due to variable capital utilization. Moreover, this effect is not enough to generate the size and persistence of the dynamic effects of corporate tax cuts on output and productivity that we estimate with LPs in Figures 1 and 3.

The results in the left column of Figure 5 stand in sharp contrast to the findings in the right column, which show the effects of a cut in personal income taxes. Here, the IRFs generated by the model with no endogenous growth (black lines with crosses) are much closer to estimated IRFs of the baseline specification (blue lines with circles), especially in the short-term. A main reason for this is that the endogenous growth mechanism has little impact at shorter horizons but this is precisely where most of the personal tax effects materialize in the data. While R&D and technological adoption appear crucial to obtain longer-term effects of corporate tax changes on output and productivity, they play only a modest role in shaping the shorter-term effects of personal income tax changes.

In summary, the main takeaway from this counterfactual experiment is that the key to generating longer-run effects from corporate tax cuts is the endogenous increase in productivity that occurs as a result of higher R&D expenditure and adoption. Because it takes time for new technologies to be adopted, these channels raise productivity steadily, which in turn boosts economic activity persistently. In contrast, a personal income tax cut has no such direct effect on firms' incentives. As a result, there is no sustained increase in R&D, productivity or output.

6.2 No variable factor utilization

In this counterfactual experiment, we turn off the variable utilization margin for both labor and capital but maintain the endogenous productivity mechanism.[15] The green lines with diamonds

[15] In Appendix A10, we switch off variable utilization for each factor separately.

in the left column of Figure 5 reveal that making factor utilization inelastic (in a model with endogenous growth) has a limited impact on the effects of corporate income tax changes on GDP, productivity and R&D. The shape and *dynamics* of the green lines with diamonds are very similar to those reported as blue lines with circles from the unrestricted baseline model. On the other hand, variable factor utilization helps the estimated model to match the *level* of the IRFs. In other words, variable factor utilization (and as we show in Appendix A10, variable *capital* utilization in particular) amplifies the magnitude of the effects of corporate income tax changes but it plays only a modest role, if any, in explaining their persistence.

In sharp contrast to the 'No R&D' counterfactual, however, the right column of Figure 5 shows that, without variable factor utilization, the dynamic effects of personal income tax cuts in the counterfactual model look very different to the effects from the full estimated model (blue lines with circles), especially at shorter horizons. In particular, the response of productivity to a personal income tax cut (middle row) becomes now counter-cyclical, which is at odds with the estimates in Section 3. In Appendix A10, we show that this is mostly driven by variable *labor* utilization. As in the previous counterfactual, although through a very different mechanism and at a very different horizon, pro-cyclical productivity turns out to be crucial to account for the dynamic effects of tax changes on output.

Finally, we note that while a long literature (starting with Oi (1962)) has studied the role of labor utilization in generating a pro-cyclical response of labor productivity at shorter horizons, other theoretical mechanisms could also generate this effect. For instance, Qiu and Ríos-Rull (2022) show that if households' search efforts in the product market vary with the business-cycle, then productivity becomes pro-cyclical in an otherwise standard medium-scale New Keynesian framework.

6.3 No endogenous growth *and* no variable factor utilization

In the previous sub-section, we have shown the importance of variable labor utilization (for personal tax cuts) and variable *capital* utilization (for corporate tax cuts) for amplifying the effects on output and productivity. On the other hand, the persistence of the effects of corporate tax cuts seems mostly driven by R&D spending and technological adoption. In this section, we bring these two exercises together and switch off the endogenous growth mechanism *and* variable utilization of capital and labor. The results are reported in Figure 6 as black lines with crosses.

This exercise reveals that a standard New-Keynesian model has a hard time generating pro-cyclical productivity. For a corporate income tax cut in the left column of Figure 6, the model

Figure 6: Counterfactuals: Response of GDP, Productivity and R&D

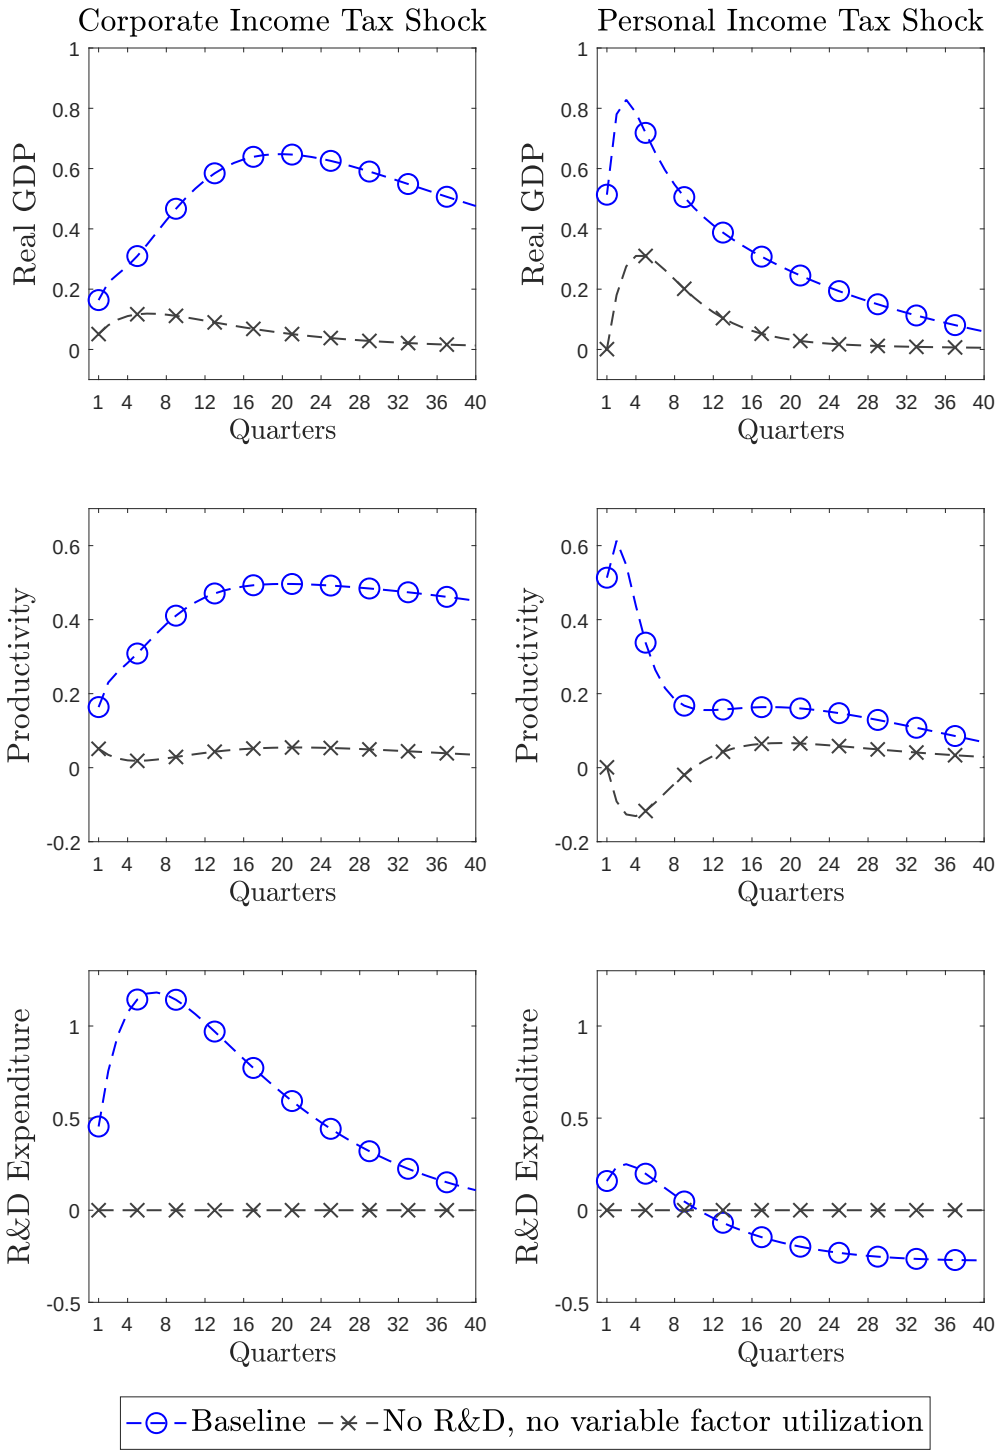

This figure shows the results of a counterfactual exercise that turns off the endogenous growth part of the model, the labor effort margin and variable capital utilization. The baseline model results are blue (circles), these repeat the results from Figure 1 and Figure 3. The black line (crosses) shows the impulse responses from the estimated model after shutting down endogenous productivity, the labor effort margin and variable capital utilization.

without endogenous growth or variable factor utilization generates an a-cyclical response of productivity and, therefore, it misses entirely the magnitude and persistence of the effects on GDP. As for personal income tax changes in the right column, the restricted model generates a counter-cyclical (and counterfactual) response of productivity, which in turn more than halves the effects on output in the first row.

In summary, the estimated structural model shows that transitory corporate income tax cuts can generate persistent dynamics and that these align well with our LP estimates from the data. In contrast, transitory personal income tax cuts generate sizable short-run effects, but appear to have no persistent long-term effects. In our theoretical model — as in the data — the key is to generate pro-cyclical productivity at the relevant horizon. The counterfactual analysis reveals that variable factor utilization combined with an endogenous growth mechanism are central to jointly explaining all our results on the heterogeneous effects of corporate and personal taxes across forecast horizons.

7 Another look at the adjustment costs on capital investment

A distinctive feature of the structural model estimates in Table 2 is that the adjustment costs on physical capital investment are different from those estimated in medium-scale DSGE models such as Smets and Wouters (2007) and Justiniano et al. (2010). Standard estimated models tend to find an important role for physical capital accumulation in explaining business cycle dynamics. Earlier studies, however, typically do not feature an endogenous growth mechanism and investment in intangible capital.

One possible explanation for this discrepancy is that we are targeting the IRFs for only two shocks, as opposed to the full-information approach adopted in other papers. Physical capital adjustment costs may simply be unimportant for the transmission of these shocks. An alternative interpretation, however, is that standard estimated medium-scale DSGE models may be relying 'excessively' on the dynamics of physical capital accumulation when in fact the persistence in the data may be better described by an (omitted) endogenous growth mechanism.

To examine this hypothesis, we now use the same IRFs-matching partial-information approach as in Section 5 but we estimate an alternative version of the model that omits the endogenous growth mechanisms. In so doing, we are making the model as close as possible to conventional New Keynesian models such as, for instance, Smets and Wouters (2007) and Justiniano et al. (2010). The full set of estimates are reported in Appendix Table A.2, which is the no-endogenous-growth

Figure 7: Prior and Posterior Distributions of the Investment Adjustment Cost Parameter

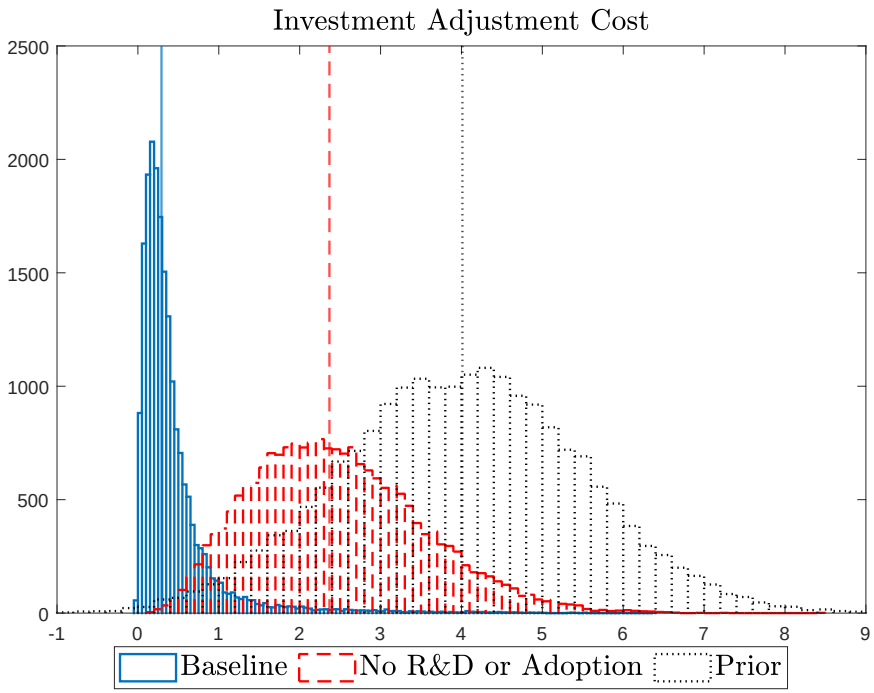

This figure compares the prior distribution (black dotted edges) and posterior distribution of the baseline model (blue solid edges) and the model without the endogenous productivity mechanism (red dashed edges). Vertical lines display the medians of each distribution.

model counterpart of Table 2. In Figure 7, we focus on $f_I^{''}$, the parameter governing adjustment costs on physical capital investment. The histogram with black dotted edges represents the prior distribution whereas the blue/solid (red/dashed) contours depict the posterior distribution of the estimate of the baseline structural model with (without) endogenous growth. Vertical lines represent the median estimates.

The main message from this exercise is that without R&D or adoption, the estimate of the investment adjustment cost parameter becomes far more in line with estimates from more standard medium-scale New Keynesian models. More specifically, the estimation of the restricted model with no endogenous growth assigns a much larger probability to investment adjustment costs $f_I^{''}$ being around 2.5, with the upper bound of the 90% credible set (in red) being 4.35. This is in sharp contrast to the estimates of our structural model with R&D, where the point estimate for the investment adjustment costs, $f_I^{''}$, becomes only 0.29 (with a much tighter band spanning the interval $[0.05, 1.31]$).

In summary, when we shut down the endogenous growth part of the structural model, the

partial-information estimates of the capital investment adjustment costs from this restricted specification are well in line with the full-information estimates from earlier studies that feature no R&D expenditure. This suggests that the omission of an endogenous growth mechanism in conventional medium-scale DSGE models may induce a bias in estimated adjustment costs on investment.

8 Conclusions

Do transitory changes in corporate and personal income taxes have persistent effects on output? And what are the channels? We answer the first question using local projections and narrative-identified tax shocks on post-WWII U.S. data. We answer the second question by running counterfactual simulations from an estimated structural model with endogenous growth, variable factor utilization and distortionary taxes.

Our main findings are that corporate income tax changes generate persistent effects on R&D expenditure, productivity and output whereas personal income tax changes trigger large but short-lived responses of capital expenditure, productivity and output. We show that matching the pro-cyclical response of productivity in the short-run and in the long-run is crucial for the ability of the estimated model to account for the dynamic effects of the two tax shocks on economic activity. Variable labor utilization appears important for replicating the short-term response of productivity and output to a personal income tax change, while R&D expenditure and technological adoption are key to account for the long-term effects of corporate income tax changes.

References

Aghion, P. and P. Howitt (1998): *Endogenous Growth Theory*, Cambridge, MA: MIT Press.

Akcigit, U., J. R. Grigsby, T. Nicholas, and S. Stantcheva (2022): "Taxation and Innovation in the 20th Century," *Quarterly Journal of Economics*, forthcoming.

Antolin-Diaz, J. and P. Surico (2022): "The long-run effects of government spending," Tech. rep., mimeographed, London Business School.

Anzoategui, D., D. Comin, M. Gertler, and J. Martinez (2019): "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence," *American Economic Journal: Macroeconomics*, 11, 67–110.

Baley, I., A. Blanco, and E. Markiewitz (2022): "Monetary Policy and Corporate Tax Reforms," Tech. rep., UPF and University of Michigan.

Barnichon, R. and C. Brownlees (2019): "Impulse Response Estimation by Smooth Local Projections," *The Review of Economics and Statistics*, 101, 522–530.

Barro, R. J. and C. J. Redlick (2011): "Macroeconomic effects from government purchases and taxes," *The Quarterly Journal of Economics*, 126, 51–102.

Beaudry, P., D. Galizia, and F. Portier (2020): "Putting the Cycle Back into Business Cycle Analysis," *American Economic Review*, 110, 1–47.

Benigno, G. and L. Fornaro (2017): "Stagnation Traps," Mimeo, London School of Economics.

Bernstein, J. I. and M. I. Nadiri (1989): *Rates of Return on Physical and R&D Capital and Structure of the Production Process: Cross Section and Time Series Evidence*, Dordrecht: Springer Netherlands, 169–187.

Bianchi, F., H. Kung, and G. Morales (2019): "Growth, slowdowns, and recoveries," *Journal of Monetary Economics*, 101, 47–63.

Bond, S., D. Harhoff, and J. V. Reenen (2005): "Investment, R&D and Financial Constraints in Britain and Germany," *Annales d'Économie et de Statistique*, 433–460.

Broda, C. and D. E. Weinstein (2006): "Globalization and the Gains From Variety," *The Quarterly Journal of Economics*, 121, 541–585.

Caldara, D. and C. Kamps (2012): "The analytics of SVARs: A unified framework to measure fiscal multipliers," *FEDS Working Paper 2012-20, Board of Governors of the Federal Reserve System.*

Cerra, V., A. Fatás, and S. Saxena (2022): "Hysteresis and Business Cycles," *Journal of Economic Literature*, forthcoming.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

Christiano, L. J., M. Trabandt, and K. Walentin (2010): "Chapter 7 - DSGE Models for Monetary Policy Analysis," Elsevier, vol. 3 of *Handbook of Monetary Economics*, 285–367.

Cloyne, J. and P. Surico (2017): "Household Debt and the Dynamic Effects of Income Tax Changes," *Review of Economic Studies*, 84(1).

Cloyne, J. S. (2013): "Discretionary tax changes and the macroeconomy: new narrative evidence from the United Kingdom," *American Economic Review*, 103(4), 1507–28.

Cloyne, J. S., N. Dimsdale, and N. Postel-Vinay (2021): "Taxes and Growth: New Narrative Evidence from Interwar Britain," .

Comin, D. and M. Gertler (2006): "Medium-Term Business Cycles," *American Economic Review*, 96, 523–551.

Comin, D. and B. Hobijn (2010): "An Exploration of Technology Diffusion," *American Economic Review*, 100, 2031–59.

Crafts, N. and T. C. Mills (2013): "Rearmament to the rescue? New estimates of the impact of "Keynesian" policies in 1930s' Britain," *The Journal of Economic History*, 73, 1077–1104.

Cram, R. G. and M. Olbert (2022): "Measuring the Effects of the Global Tax Reform - Evidence from High-frequency Data," Tech. rep., mimeographed, London Business School.

Danguy, J., G. de Rassenfosse, and B. van Pottelsberghe de la Potterie (2013): "On the origins of the worldwide surge in patenting: an industry perspective on the R amp;D–patent relationship," *Industrial and Corporate Change*, 23, 535–572.

de Ridder, M. (2019): "Intangible Investments and the Persistent Effect of Financial Crises on Output," .

Fornaro, L. and M. Wolf (2020): "The Scars of Supply Shocks," CEPR Discussion Papers 15423, CEPR Discussion Papers.

Forni, M. and L. Gambetti (2014): "Sufficient information in structural VARs," *Journal of Monetary Economics*, 66, 124–136.

Furlanetto, F., A. Lepetit, Ørjan Robstad, J. Rubio-Ramírez, and P. Ulvedal (2021): "Estimating Hysteresis Effects," Working Papers 2021-11, FEDEA.

Galí, J. and T. van Rens (2020): "The Vanishing Procyclicality of Labour Productivity," *The Economic Journal*, 131, 302–326.

Giannone, D., M. Lenza, and G. E. Primiceri (2015): "Prior Selection for Vector Autoregressions," *The Review of Economics and Statistics*, 97, 436–451.

Guajardo, J., D. Leigh, and A. Pescatori (2014): "Expansionary austerity? International evidence," *Journal of the European Economic Association*, 12, 949–968.

Gunter, S., D. Riera-Crichton, C. Vegh, and G. Vuletin (2018): "Non-linear effects of tax changes on output: The role of the initial level of taxation," *World Bank working paper 8668*.

Hayo, B. and M. Uhl (2014): "The macroeconomic effects of legislated tax changes in Germany," *Oxford Economic Papers*, 66.

Hussain, S. M. and L. Liu (2018): "Comparing the effects of discretionary tax changes between the US and the UK," *The B.E. Journal of Macroeconomics*, 18(1).

Jones, C. I. (1995): "R& D-Based Models of Economic Growth," *Journal of Political Economy*, 103, 759–784.

Jordà, Ò. (2005): "Estimation and inference of impulse responses by local projections," *American economic review*, 95, 161–182.

Jordà, Ò., S. R. Singh, and A. M. Taylor (2020): "The Long-Run Effects of Monetary Policy," Working Paper 26666, National Bureau of Economic Research.

Jordà, Ò. and A. M. Taylor (2015): "The time for austerity: estimating the average treatment effect of fiscal policy," *The Economic Journal*, 126, 219–255.

Jordà, , S. R. Singh, and A. M. Taylor (2020): "The Long-Run Effects of Monetary Policy," Working Paper 26666, National Bureau of Economic Research.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010): "Investment shocks and business cycles," *Journal of Monetary Economics*, 57, 132–145.

Leeper, E. M., M. Plante, and N. Traum (2010): "Dynamics of fiscal financing in the United States," *Journal of Econometrics*, 156, 304–321.

Leeper, E. M., N. Traum, and T. B. Walker (2017): "Clearing Up the Fiscal Multiplier Morass," *American Economic Review*, 107, 2409–54.

Li, D., M. Plagborg-Møller, and C. K. Wolf (2021): "Local Projections vs. VARs: Lessons From Thousands of DGPs," Mimeographed, Princeton University and Chicago University.

Li, W. C. Y. and B. H. Hall (2020): "Depreciation of Business R&D Capital," *Review of Income and Wealth*, 66, 161–180.

Ma, X. and R. Samaniego (2020): "Macroeconomic shocks and productivity: Evidence from an estimated ideas production function," Tech. rep.

McGrattan, E. R. (1994): "The macroeconomic effects of distortionary taxation," *Journal of Monetary Economics*, 33, 573–601.

Mertens, K. and M. O. Ravn (2012): "Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated U.S. Tax Policy Shocks," *American Economic Journal: Economic Policy*, 4(2), 145–181.

——— (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103, 1212–47.

——— (2014): "A reconciliation of SVAR and narrative estimates of tax multipliers," *Journal of Monetary Economics*.

Miranda-Agrippino, S. and G. Ricco (2015): "The Transmission of Monetary Policy Shocks," Discussion Papers 1711, Centre for Macroeconomics (CFM).

Montiel Olea, J. L. and M. Plagborg-Møller (2021): "Local Projection Inference is Simpler and More Robust Than You Think," *Econometrica*, 89, 1789–1823, arXivCode and data (GitHub)Video of online seminar (YouTube).

Mumtaz, H. and K. Theodoridis (2020): "Fiscal policy shocks and stock prices in the United States," *European Economic Review*, 129.

Nguyen, A. D., L. Onnis, and R. Rossi (2018): "The Macroeconomic Effects of Income and Consumption Tax Changes," *Unpublished manuscript*.

Oi, W. Y. (1962): "Labor as a Quasi-Fixed Factor," *Journal of Political Economy*, 70, 538–555.

Plagborg-Møller, M. and C. K. Wolf (2021): "Instrumental Variable Identification of Dynamic Variance Decompositions," Matlab code suite and data (GitHub).

——— (Forthcoming): "Instrumental Variable Identification of Dynamic Variance Decompositions," *Journal of Political Economy*.

Qiu, Z. and J.-V. Ríos-Rull (2022): "Procyclical Productivity in New Keynesian Models," Working Paper 29769, National Bureau of Economic Research.

Queraltó, A. (2022): "Monetary Policy in a Model of Growth," .

Ramey, V. (2016): "Macroeconomic Shocks and Their Propagation," *Handbook of Macroeconomics*, 2.

Ramey, V. and M. Shapiro (1998): "Costly capital reallocation and the effects of government spending," *Carnegie Rochester Series on Public Policy*, 48, 145–194.

Ramey, V. and S. Zubairy (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data," *Journal of Political Economy*.

Ramey, V. A. (2011): "Identifying Government Spending Shocks: It's All in the Timing," *Quarterly Journal of Economics*, 126(1), 1–50.

Romer, C. D. and D. H. Romer (1989): "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," *NBER Macroeconomic Annual 1989*.

——— (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94(4), 1055–84.

——— (2010): "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, 100, 763–801.

Romer, P. M. (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98, pp. S71–S102.

Smets, F. and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606.

Stock, J. H. and M. W. Watson (2018): "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *The Economic Journal*, 128, 917–948.

Wen, Y. (2004): "What Does It Take to Explain Procyclical Productivity?" *The B.E. Journal of Macroeconomics*, 4.

Appendix

A1 Data Appendix

A1.1 Macroeconomic data

Table A.1: Macroeconomic variables definitions

Variable	Description	Source
Consumption	Real personal consumption expenditure per-capita	FRED divided by population
Investment	Real Non-residential investment per-capita	MR
Productivity	Output per hour (Non-Farm business sector)	FRED
R&D spending	Investment in Research and Development	FRED divided by IPP deflator and population
Employment	Total economy employment per-capita	MR
Population	Total Population over age 16	MR

The main macroeconomic variables are taken directly from Mertens and Ravn (2013): (1) $APITR_t$, (2) $ACITR_t$, (3) $\ln\left(B_t^{PI}\right)$, (4) $\ln\left(B_t^{CI}\right)$, (5) $\ln\left(G_t\right)$, (6) $\ln\left(GDP_t\right)$, (7) $\ln(DEBT_t)$. The personal and corporate tax rates are denoted by $APITR_t$ and $ACITR_t$, respectively while $\ln\left(B_t^{PI}\right)$ and $\ln\left(B_t^{CI}\right)$ are the corresponding tax bases in real per-capita terms. $\ln\left(G_t\right)$ denotes real per-capita government spending, while $\ln(DEBT_t)$ is real per-capita federal debt. Real per-capita GDP is denoted by $\ln\left(GDP_t\right)$. For a detailed description of these series and data sources, see the appendix of Mertens and Ravn (2013). The table below provides a list of the additional macroeconomic data used in our analysis. MR denotes the replication files of Mertens and Ravn (2013) available at https://www.aeaweb.org/articles?id=10.1257/aer.103.4.1212.

A1.2 Sectoral Data

Gross output by industry is obtained from the Bureau of economic analysis (BEA). The annual data from 1947-1997 is available at the following link. We deflate Gross output by its deflator. This historical data is combined with the more recent quarterly real Gross output data to produce an annual time series for 87 sectors from 1950-2006. Real gross output is divided by population.

Data on R&D intensity is obtained from the Business Enterprise Research and Development Survey of the National Science Foundation for the period 1999 to 2007. R&D intensity is defined as funds for industrial R&D as a percent of net sales of companies. The R&D intensity data from this survey can be matched to 28 industries in the Gross output data set. These 28 industries are used in the sectoral analysis presented in the paper. Figure A.1 displays the average R&D intensity obtained over the period 1999-2007 in these 28 industries and shows the split between high and low R&D intensity groups.

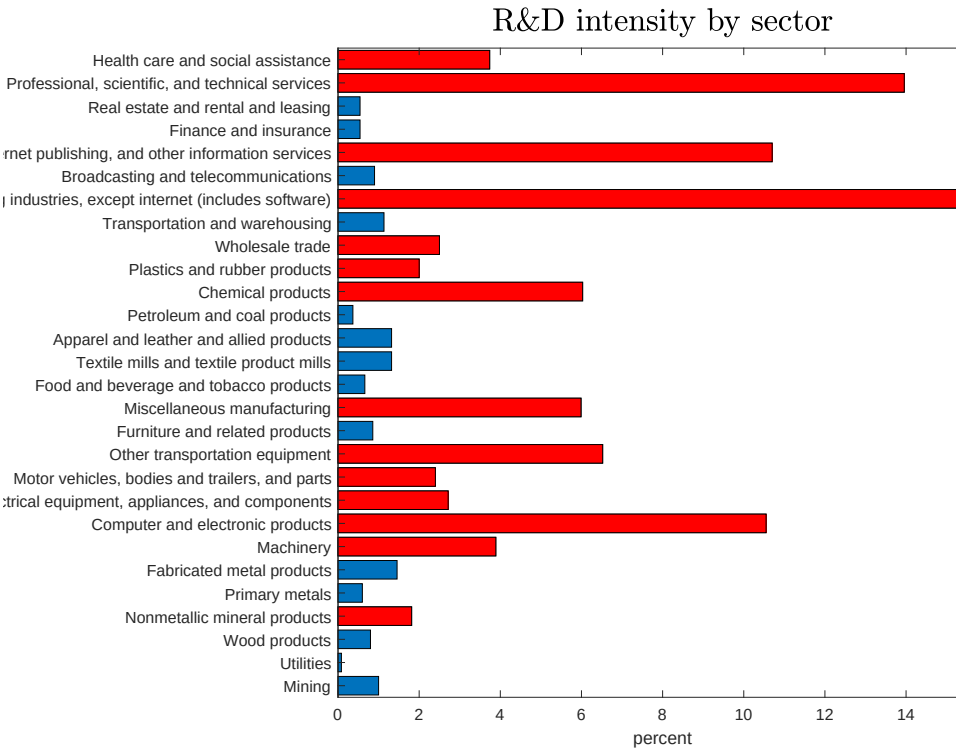

Figure A.1: Average R&D intensity 1999-2007. Red bars denote industries with higher than median R&D intensity.

A2 Shorter-Term vs. Longer-Term Effects: Prior Distributions

Figure A.2: Joint Prior and Posterior Distributions of Shorter-term versus Longer-term Effects

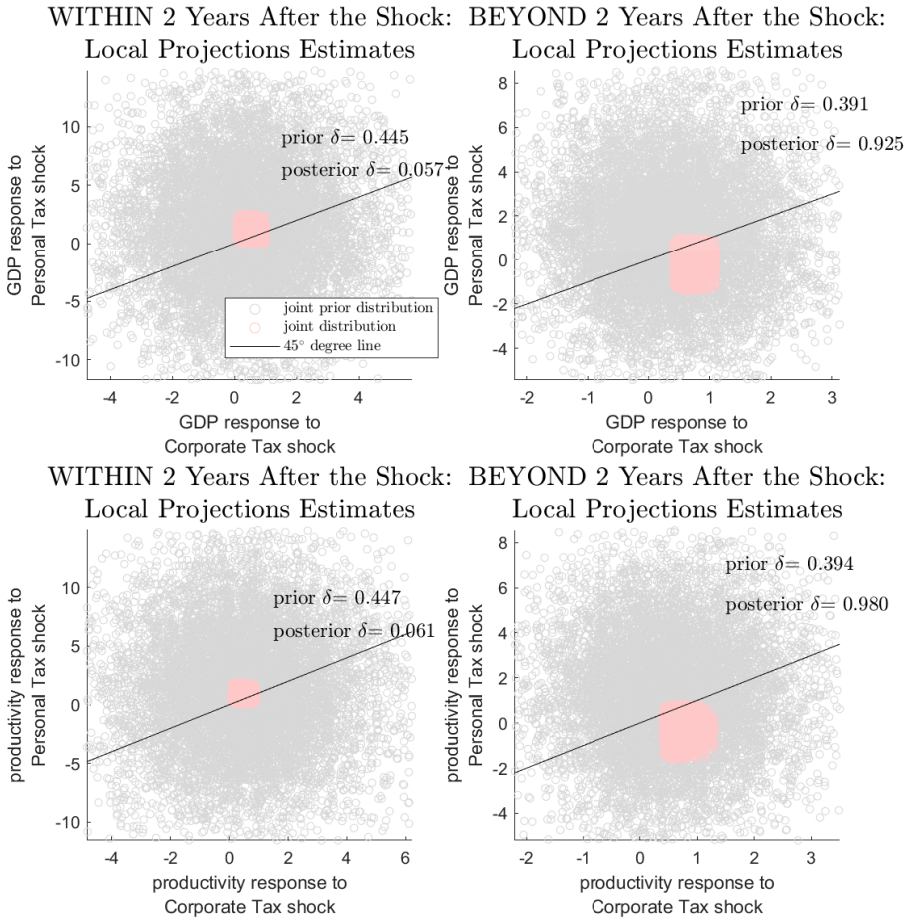

This figure shows the joint prior (reported in grey) and posterior (reported in light red) distributions of the GDP responses (top row) and productivity response (bottom row) to shocks to the personal income average tax rate (on the vertical axis) and the corporate income average tax rate (on the horizontal axis) in the shorter-term (left column) and longer-term (right column) based on the local projections in equation 1. The shorter-term (longer-term) column refers to the cumulated average responses of GDP and productivity to each shock over the quarters within (beyond) the first two years after the shock. The red dots show draws from the posterior distributions of the impulse responses. The grey dots refer to draws from the prior distribution. The black slope represents the 45° degree line, which is the locus of points along which the two shocks have GDP effects of exactly the same magnitude. Points above (below) the 45° degree line indicates the mass of the joint posterior distribution for which the effects of personal income tax shocks are larger (smaller) than the GDP effects of corporate income tax shocks. The δ statistic shows denotes the probability that the effect of a corporate income tax shock is larger than the effect from a personal income tax shock.

A3 The Responses of Hours, Employment and TFP

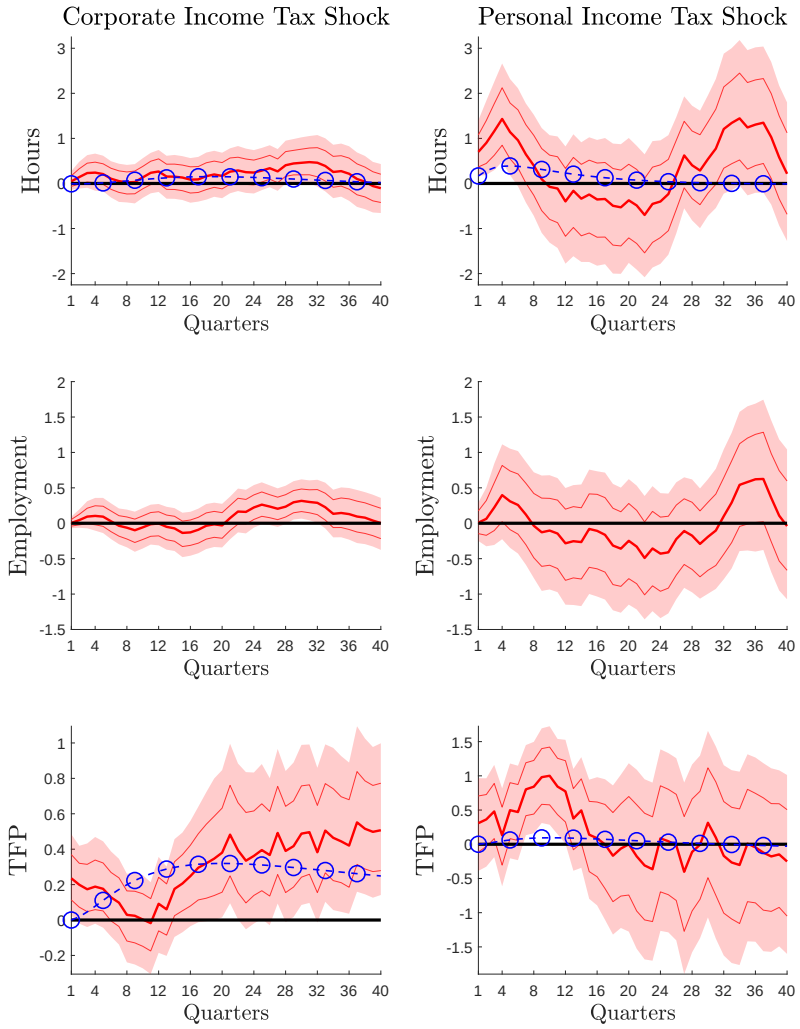

This figure shows the responses of total hours, employment and total factor productivity to a 1% cut in the average rate of corporate income taxes (left column) and the average rate of personal income taxes (right column). Red shadow bands represent central posterior 68th and 90th credible sets. Blue lines with circles represent the impulse responses of the model in Section 4 evaluated at the posterior median of estimated model parameters. Since the model used in Section 4 does not have a well-defined extensive model of employment we do not plot the model response of employment.

A4 Sectoral Evidence

We investigate the response of gross output (GO) output to the tax shocks in the high and low R&D groups. The former group is defined as the industries that have a R&D intensity larger than the median, while industries in the low group have intensity lower than the median. We construct aggregate GO in these two groups and use our benchmark LP to estimate the response of these series to the tax shocks. As the number of observations is limited, the model is kept parsimonious, with one lag of the tax rates and annual GDP as the control variables. The estimated impulse responses are shown in Figure 5. The top panel shows the response of GO in all sectors. As in the benchmark case, corporate tax shocks have their largest effect in the medium to the long-run. In contrast, personal tax shocks lead to an increase in output in the first 2 years. However, the medium and long-run impact of this shock is not statistically different from zero.

The bottom two panels of Figure 5 show the response of output in high and low R&D sectors. Consider the bottom left panel. There is clear evidence that the response of output to corporate tax shocks is larger in the high R&D group at long horizons. This heterogeneity is entirely absent when the response to personal tax shocks is considered.

Figure A.3: Sectoral Evidence: Average Effects and Results by R&D Share

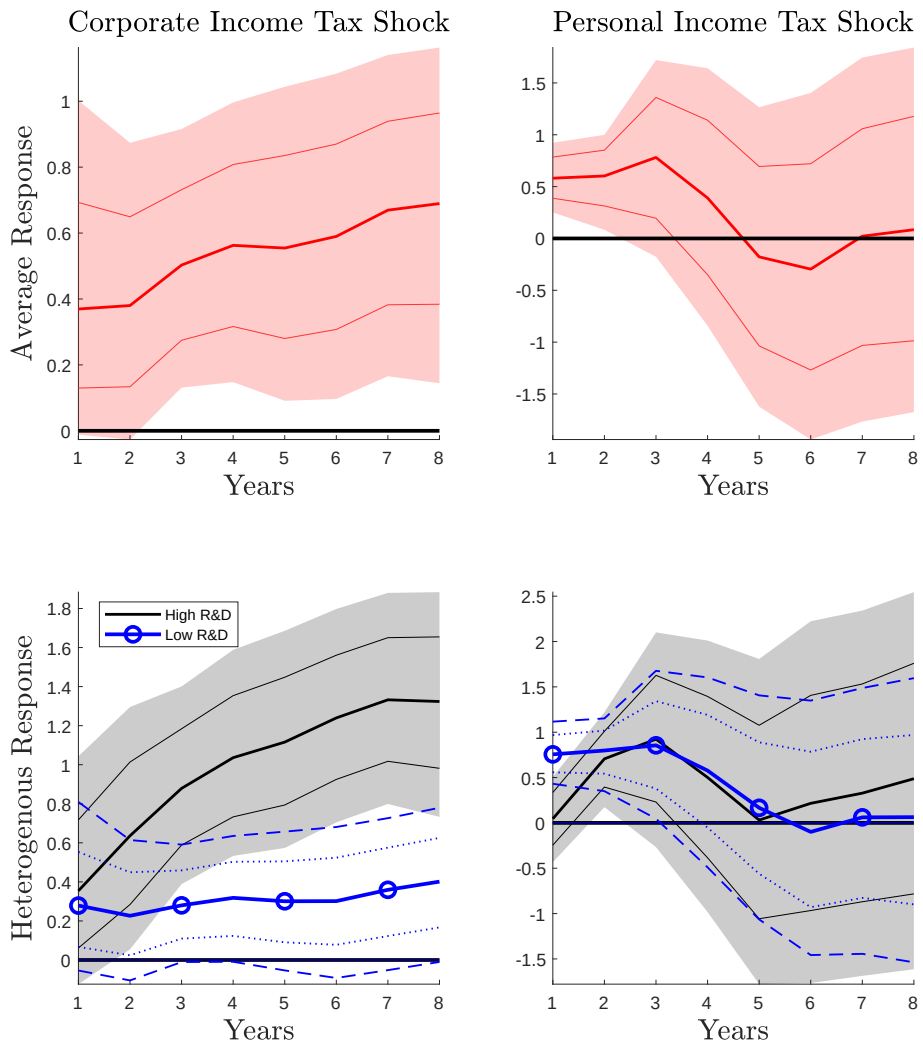

This figure shows the response of output using sectoral data from the U.S. BEA. The first row show the average effect. The bottom row further split sectors into high R&D intensive and low R&D intensive. See text for more details.

A5 Forecast Error Variance Decomposition

In this section we report the contribution of the personal and corporate tax shocks to the forecast error variance (FEV) of key variables, using the estimated local projection-based impulse responses (see Jordà, 2005). Figure A.4 presents the estimated decomposition. The contribution of both tax shocks to the FEV of GDP is about 20 percent at short horizons. However, at medium and long horizons, the contribution of corporate tax shocks is at least twice as large as that of personal tax shocks. A similar pattern holds for productivity, investment, R&D and consumption.

Figure A.4: Forecast Error Variance decomposition

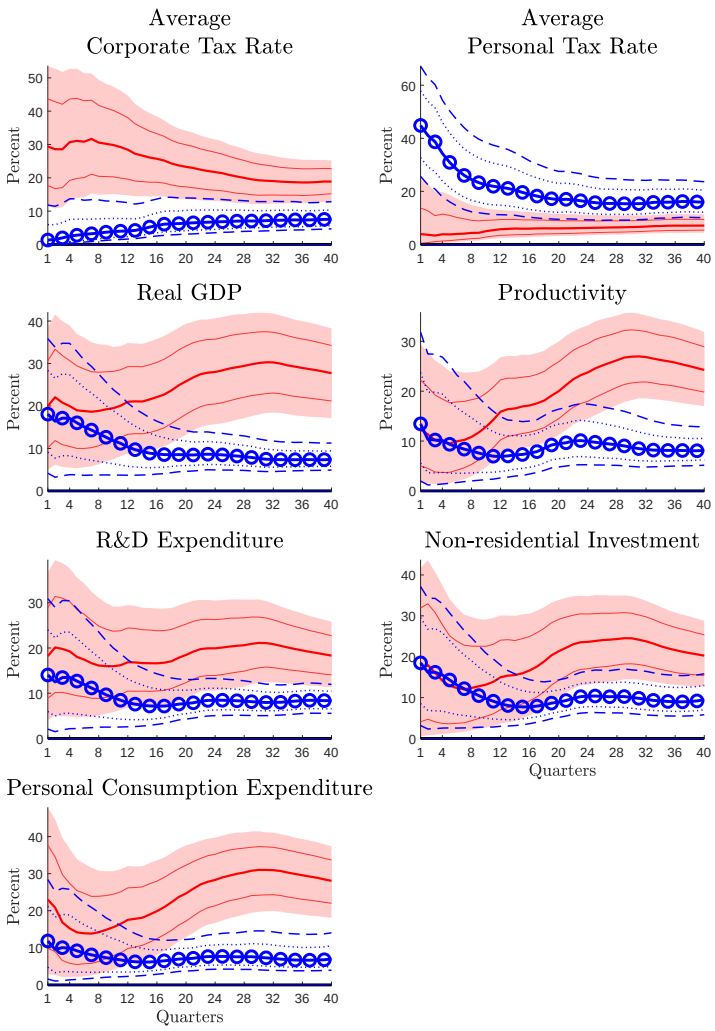

This figure shows the contribution of corporate and personal tax changes to the variance of each variable in the figure. The effects of corporate tax changes are shown in the red lines (posterior median and 68 percent band) and the shaded area (90 percent band). The line with circles show the contribution of the personal tax shock, with the posterior 68 percent (90 percent) bands shown by the dotted (dashed) lines.

A6 Robustness

Figure A.5: Response of real GDP: Different Specifications

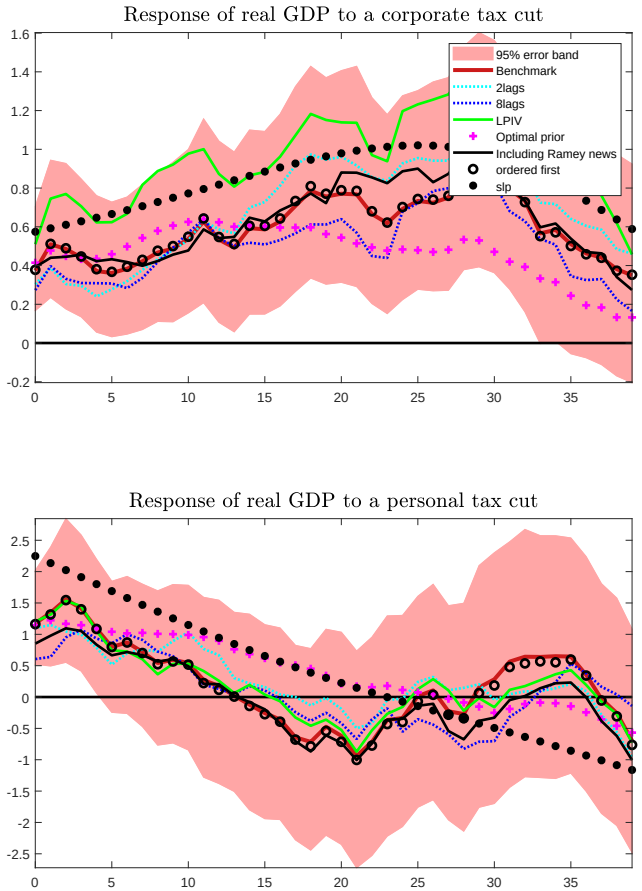

This figures shows the 90% bands for the baseline empirical real GDP result in pink, together with the point estimates from various alternative specifications. In particular, this figure shows the effects of: (i) changing number of lags used as control variables, (ii) adjusting the prior, (iii) conducting pure local projections IV rather than our Bayesian local projection setup (iv) including the Ramey (2011) defence news shock as a control (v) using a Smooth Local Projections approach (vi) changing the ordering of the tax shocks. See text for more discussion.

Figure A.6: Response of R&D: Different Specifications

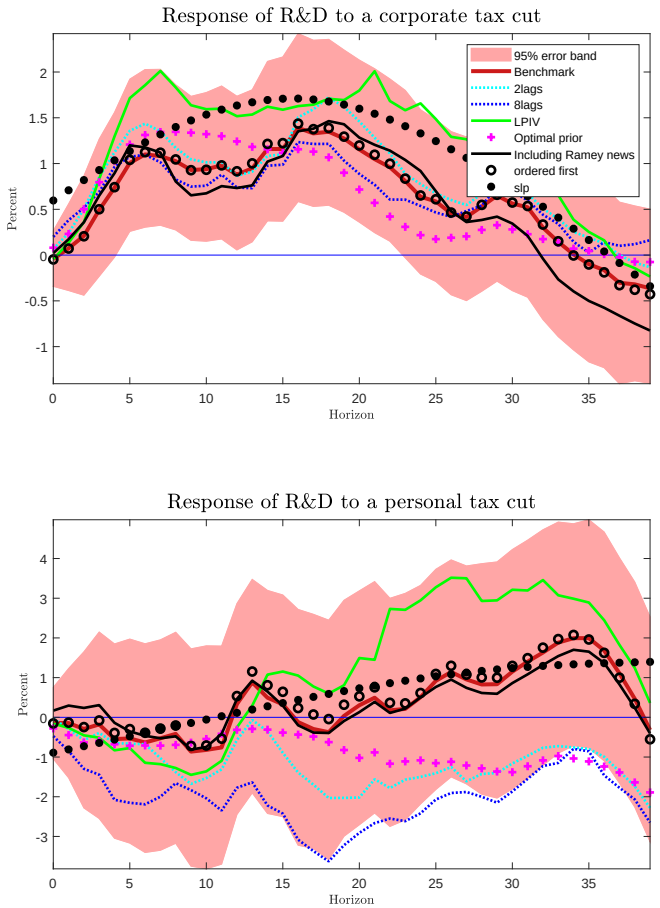

This figures shows the 90% bands for the baseline empirical R&D result in pink, together with the point estimates from various alternative specifications. In particular, this figure shows the effects of: (i) changing number of lags used as control variables, (ii) adjusting the prior, (iii) conducting pure local projections IV rather than our Bayesian local projection setup (iv) including the Ramey (2011) defence news shock as a control (v) using a Smooth Local Projections approach (vi) changing the ordering of the tax shocks. See text for more discussion.

Figure A.7: Response of Productivity: Different Specifications

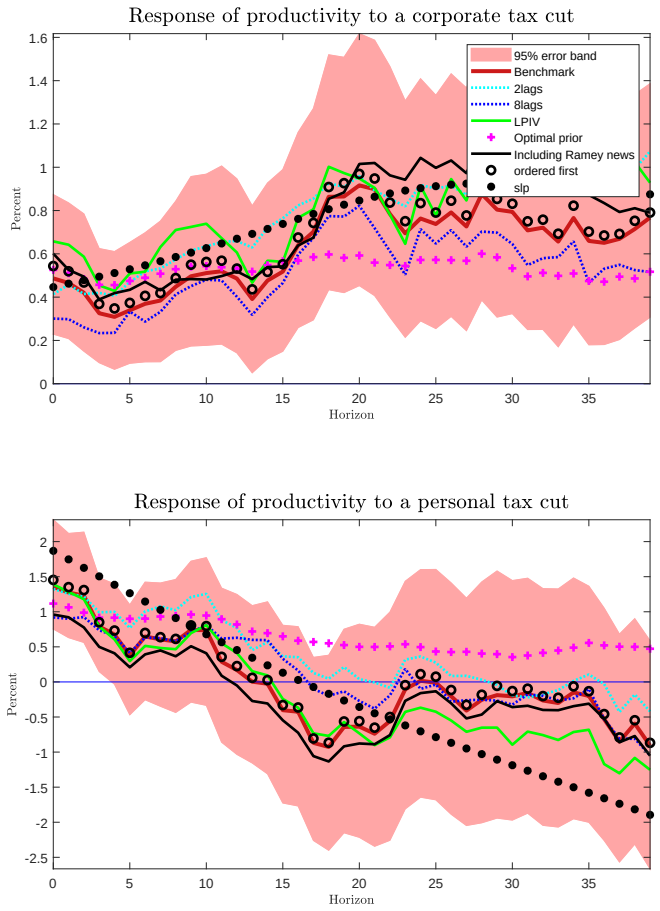

This figures shows the 90% bands for the baseline empirical productivity result in pink, together with the point estimates from various alternative specifications. In particular, this figure shows the effects of: (i) changing number of lags used as control variables, (ii) adjusting the prior, (iii) conducting pure local projections IV rather than our Bayesian local projection setup (iv) including the Ramey (2011) defence news shock as a control (v) using a Smooth Local Projections approach (vi) changing the ordering of the tax shocks. See text for more discussion.

A7 Monte-Carlo evidence on Local Projections estimates of impulse response functions at medium and long-run horizons

In this section we investigate the ability of LPs and VARs to estimate impulse response functions at medium and long-run horizons. Our Monte-Carlo analysis complements that of Jordà et al. (2020) as we consider the performance of multi-variate models.

A7.1 Data Generating Process and models

The data generating process is designed to mimic the broad features of the impulse responses of key variables to corporate tax shocks. The estimated response of variables such as GDP, consumption and productivity to corporate shocks is characterised by small increases at short horizons with larger positive changes arriving after about 20 periods. We replicate this shape by generating data from a bi-variate VAR(20)

$$Y_t = B_1 Y_{t-1} + B_2 Y_{t-2} + ... + B_{20} Y_{t-20} + A_0 E_t, E_t \sim N(0,1) \tag{34}$$

We assume that $B_1 = \begin{pmatrix} 0.7 & 0 \\ 0 & 0.75 \end{pmatrix}$ and $B_{20} = \begin{pmatrix} 0.1 & 0.1 \\ 0.1 & 0 \end{pmatrix}$ while $B_2 = B_3 = ... = B_{19} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$. The contemporaneous impact matrix is fixed at $A_0 = \begin{pmatrix} 1 & 0 \\ 0.05 & 1 \end{pmatrix}$ We generate $T_1 = T + T_0$ observations from this model where $T_0 = 50$ and $T = 230$. The first T_0 observations are discarded to account for initial values. We estimate two models using this artificial data: (1) A VAR(4) and (2) A LP that includes 4 lags of the two variables as controls. The models are used to estimate the response to the first shock. Note that we do not attempt to estimate A_0 which is kept fixed at the true value for both models.

A7.2 Results

Figure A.8 displays the main results. Consider first the true impulse response of Variable 2. The features of this function are similar to those reported in our empirical analysis for variables such as GDP, consumption and productivity. That is, this response is characterised by the feature that the main effect arrives in the medium run rather than immediately. The VAR(4) model captures the short-run impact well. However, it completely misses the increase in the variables at horizon 20. In contrast, the LP that includes the same number of lags captures both the initial increase in the variables and the subsequent rise at horizon 20. Figure A.9 shows the effect of increasing the lag length. Even with 10 lags, the VAR response of the second variable is far from the truth at long horizons. When the lag length is increased to 20, the performance of the VAR improves substantially. In the case of the LP, increasing the lag length does not materially affect the response after horizon 20. However, there is some evidence that longer lags reduce the discrepancy between the LP response and truth between horizons 10 and 20. In short, this simple stylised simulation demonstrates that VARs with a small number of lags are likely to be unreliable in estimating responses where the bulk of the movement occurs at long horizons. The LP appears to be more

robust to lag truncation.

Figure A.8: Monte-Carlo results

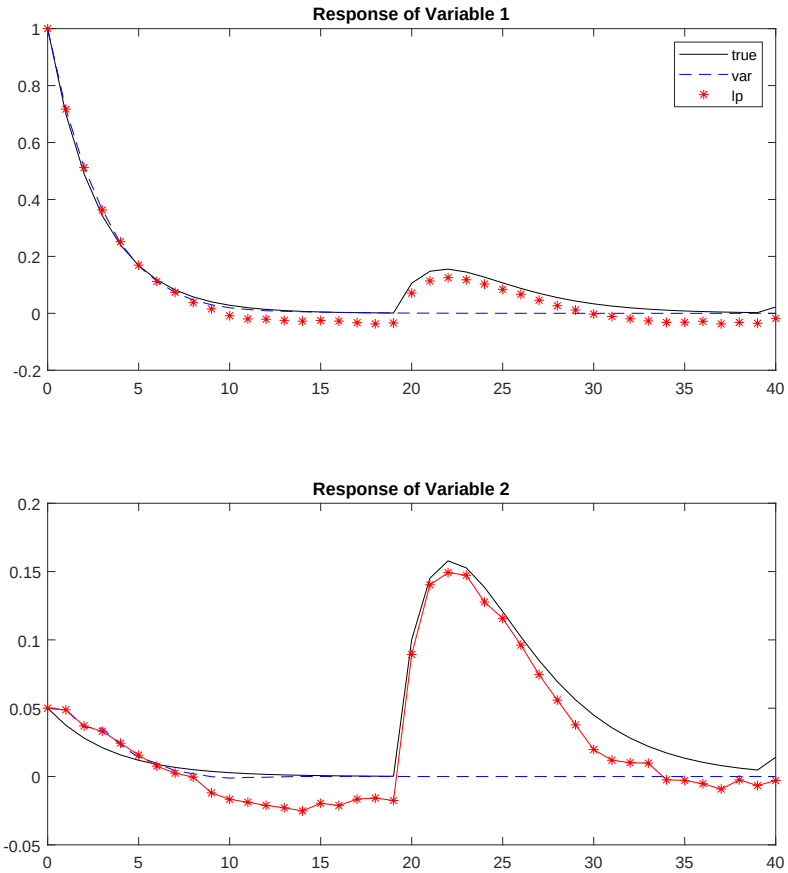

Monte-Carlo estimates of impulse responses of the two variables in Y to the first shock.

Figure A.9: Monte-Carlo results using different lag lengths

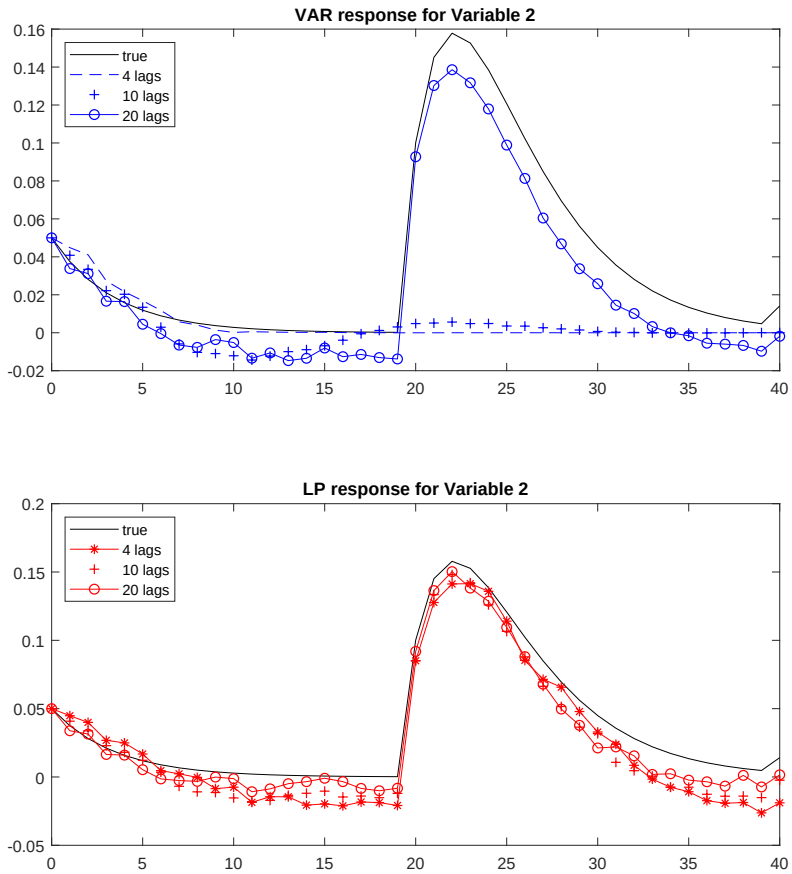

Monte-Carlo estimates of impulse responses of the second variable in Y to the first shock. The experiment is repeated for different lag lengths

A8 Prior Predictive Analysis

Figure A.10: Response of Taxes, Real GDP and Productivity

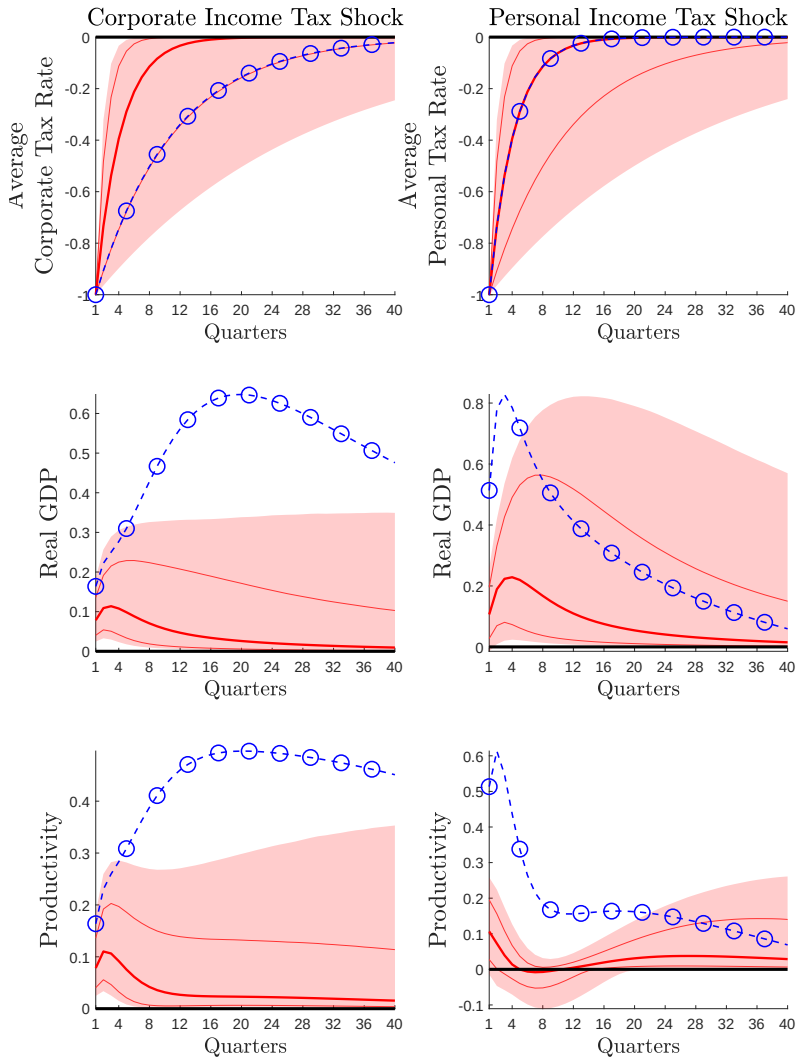

This figure shows the response of the average tax rates, real GDP and productivity to a 1% cut in the average tax rate of corporate income taxes (left column) and the average tax rate of personal income taxes (right column). Red shadow bands and solid lines represent the 90th and 68th percentiles of the prior distribution of impulse response functions. Blue lines with circles represent the impulse responses of the model in Section 4 evaluated at the posterior median of estimated model parameters.

Figure A.11: Response of Labor R&D, Investment and Consumption

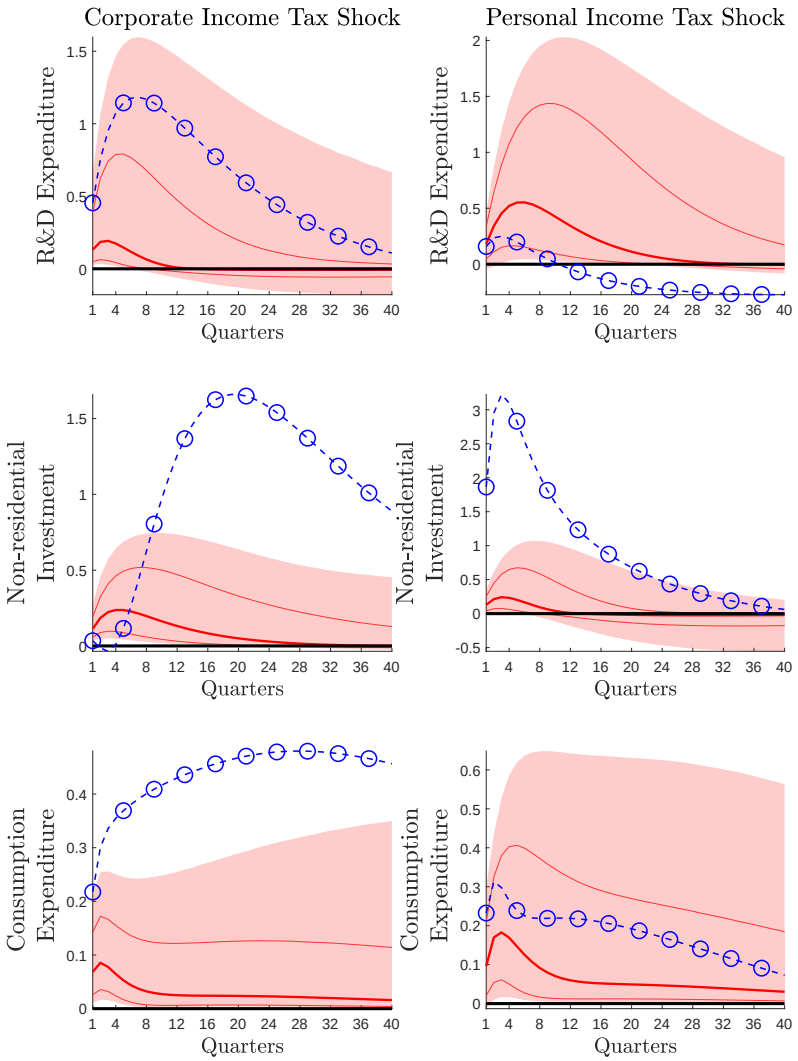

This figure shows the responses of consumption, investment and R&D to a 1% cut in the average tax rate of corporate income taxes (left column) and the average tax rate of personal income taxes (right column). Red shadow bands and solid lines represent the 90th and 68th percentiles of the prior distribution of impulse response functions. Blue lines with circles represent the impulse responses of the model in Section 4 evaluated at the posterior median of estimated model parameters.

A9 Shorter-Term vs. Longer-Term Effects: Model Estimates

Figure A.12: Joint Posterior Distribution of Shorter-term versus Longer-term Effects on GDP

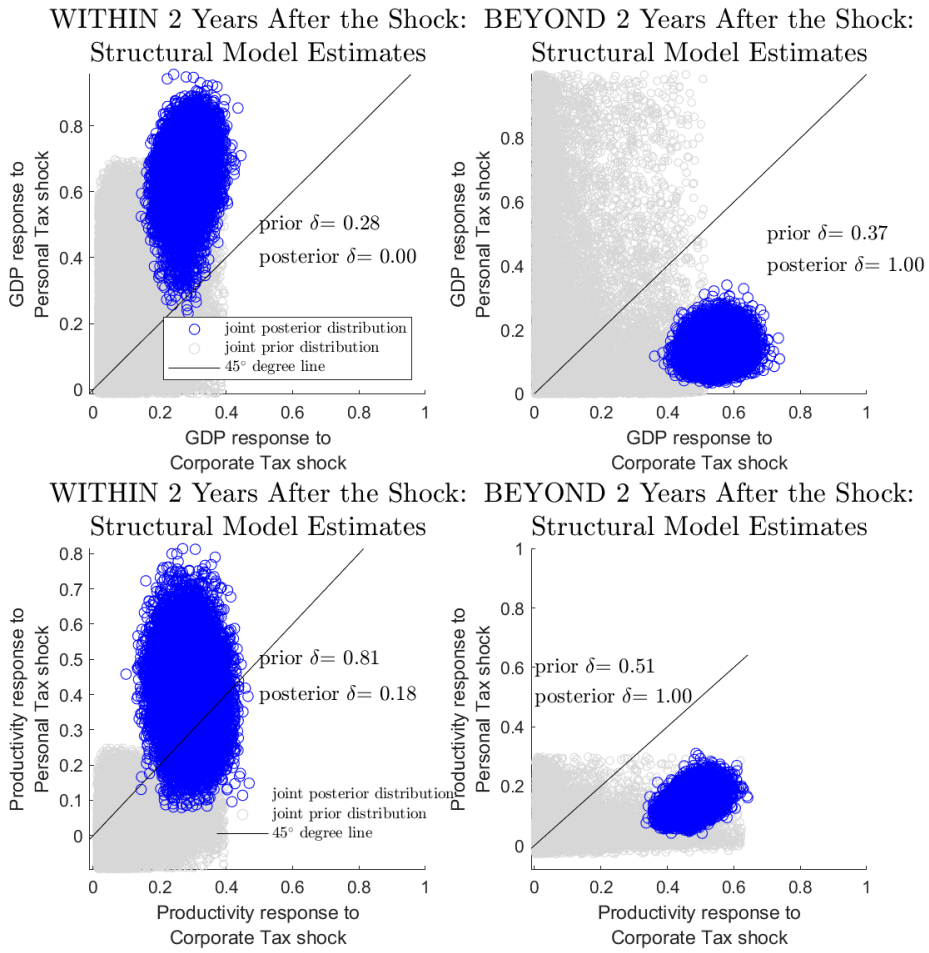

This figure shows the joint posterior distribution of the GDP responses (top row) and productivity response (bottom row) to shocks to the personal income average tax rate (on the vertical axis) and the corporate income average tax rate (on the horizontal axis) in the shorter-term (left column) and longer-term (right column) based on the prior and posterior distributions of the structural model in Table 2. The shorter-term (longer-term) column refers to the cumulated average responses of GDP and productivity to each shock over the quarters within (beyond) the first two years after the shock. The blue dots in the top (bottom) row identify draws from the posterior distributions of estimated parameters based on the structural model. The grey dots refer to draws from the prior distribution. The black slope represents the 45° degree line, which is the locus of points along which the two shocks have GDP effects of exactly the same magnitude. Points above (below) the 45° degree line indicates the mass of the joint posterior distribution for which the GDP effects (in the top row) or the productivity effects (in the bottom row) of personal income tax shocks are larger (smaller) than the GDP effects of corporate income tax shocks. The statistics δ in the top (bottom) row denotes the probability that the GDP (productivity) response to corporate income tax shocks is larger than the GDP (productivity) response to personal income tax shocks.

A10 IRFs with no variable utilization of either labor or capital

Figure A.13: Counterfactuals: Response of GDP, Productivity and R&D

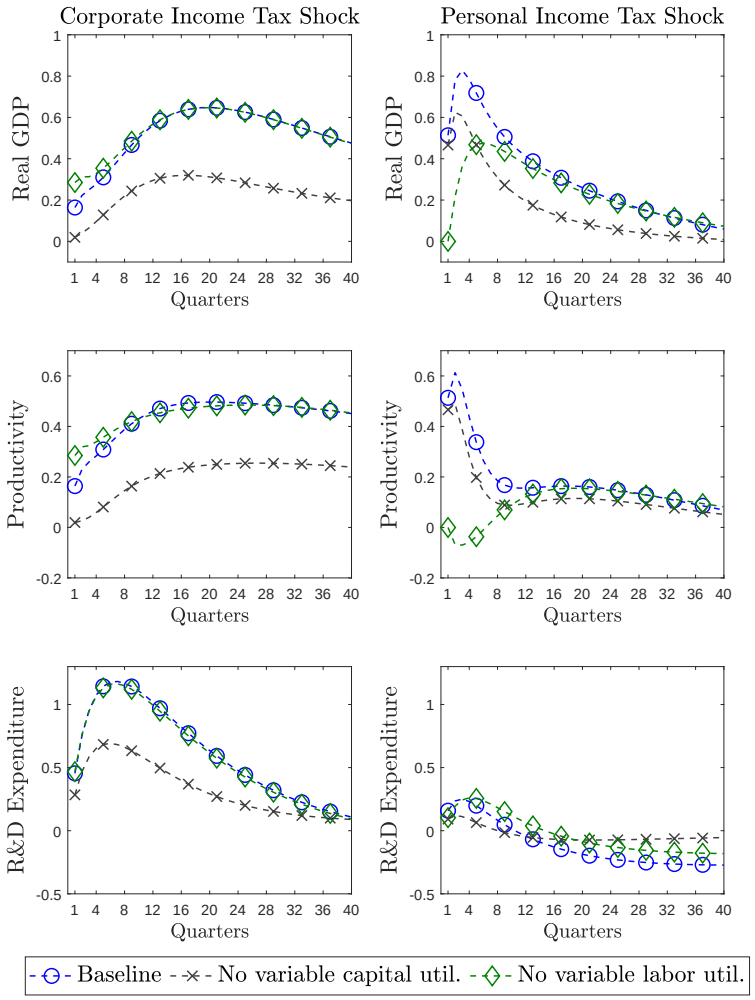

This figure shows the results of two further counterfactual exercises. The first (second) exercise in black lines with crosses (green lines with diamonds) shuts down variable labor (labor) utilization. The baseline model results from Figure 1 and Figure 3 are reported as blue lines with circles.

A11 Estimates of the structural model with no endogenous growth

This section reports the prior and posterior distributions of the parameters of the structural model in the restricted specification with neither technological adoption nor R&D expenditure. The main difference relative to Table 2 is that the investment adjustment cost parameter is significantly higher than the estimates based on the model with endogenous growth. Furthermore, and in sharp contrast to Table 2, the estimate of this parameter in Table A.2 is in line with the available estimates in the business cycle literature on DSGE model (see for instance Smets and Wouters, 2007, Justiniano et al., 2010), which typically assume an exogenous growth path.

Table A.2: Estimated Parameters - No technological adoption or R&D spending

Parameter	Description	Prior			Posterior	
		Distr	Mean	Std. Dev.	Mean	90% int.
Preference & HHs						
h	Consumption habit	beta	0.5	0.2	0.41	[0.15, 0.67]
γ	Inverse effort elasticity	gamma	1	0.5	0.58	[0.16, 1.27]
Frictions & Production						
f_I''	Investment adjustment	normal	4	1.5	2.47	[0.93, 4.35]
ψ_u	Capital utilization adjustment	beta	0.6	0.15	0.35	[0.23, 0.49]
ξ_p	Calvo prices	beta	0.5	0.2	0.21	[0.07, 0.38]
Shocks						
$\rho_{\tau CI}$	Corporate taxes AR	beta	0.7	0.2	0.94	[0.92, 0.95]
$\rho_{\tau PI}$	Labour taxes AR	beta	0.7	0.2	0.76	[0.72, 0.8]