The Dynamic Effects of Income Tax Changes in a World of Ideas*

James Cloyne Joseba Martinez Haroon Mumtaz Paolo Surico

December 2023

Abstract

Using a narrative identification of US tax changes over the post-WWII period, we show that corporate income tax cuts foster R&D spending and innovation, leading to a persistent increase in aggregate productivity and output. In contrast, changes in the average personal income tax rate have mostly short-term effects. An estimated endogenous productivity model highlights the role of "applied research" —over and above formal R&D— as a main force behind these results, and suggests a social rate of return to investment in innovation between 20% and 75%.

JEL classification: E23, E62, O32, O34, O38
Key words: corporate taxes, narrative identification, TFP, R&D, technological adoption.

*We are grateful to Martin Beraja, Tim Besley, Yuriy Gorodnichenko, Chad Jones, Òscar Jordà, Karel Mertens, Elias Papaioannou, Valerie Ramey and Maarten de Ridder for very insightful comments and suggestions. We thank Aditya Polisetty for excellent research assistance. Paolo Surico gratefully acknowledges financial support from ESRC (Grant ref. ES/Y002490/1). 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

The last two decades have witnessed a dismal record of productivity growth in advanced economies. In response, many governments have pursued a wide array of policy interventions, ranging from public spending on education, research and development subsidies, income tax cuts, infrastructure investment, and industrial policies. However, little is known about the effects of these policies on aggregate productivity. This is more surprising given the large empirical literature on fiscal policy spurred by the crisis of 2007-2009, which has mostly focused on the ability of public spending and taxation to stimulate GDP in the short run.

Our empirical analysis uncovers a novel channel through which fiscal policy can foster innovation and productivity over the medium term. A temporary cut in corporate income taxes triggers a sustained but transitory expansion in capital and R&D expenditure. This triggers a persistent increase in patenting and productivity, which in turn leads to large and significant effects on GDP and consumption at long horizons. In contrast, temporary changes in the average rate of personal income taxes mostly work through labour supply, and the bulk of their impact on productivity and output occurs in the short run.

Our results are based on Local Projections (LP) and post-WWII US data. Federal tax changes are identified using the narrative approach of Romer and Romer (2010), which excludes all tax changes motivated by current or prospective economic conditions. We adopt the decomposition of these data by Mertens and Ravn (2013) and focus on changes in personal and corporate income tax separately. To interpret the evidence from the LPs, we develop and estimate a structural business cycle model with endogenous productivity that features (i) tax amortisation allowances on the value of ideas and (ii) two margins of innovation: 'basic' and 'applied' research.

Tax amortisation allowances on intellectual property are a salient feature of the tax codes of many advanced economies, including the United States. Basic research refers to the effort to generate ideas that break new ground in science and business. This can be thought of as the type of formal R&D expenditure featured in endogenous growth models in the tradition of Romer (1990), Jones (1995) and Aghion and Howitt (1998). Applied research captures the equally important activity of turning new ideas and technologies into new products and processes. These can be thought of as 'technological adoption' in the framework of Comin and Gertler (2006) or 'applied research' in the model proposed by Jones (2022).

The estimated structural model reveals that tax amortisation allowances for intellectual property assets are crucial for corporate income tax changes to affect the market price of ideas, and thereby the incentives to innovate. Tax cuts increase the market price of ideas, which stimulates expenditure in basic and applied research, fostering innovation and productivity over long horizons. As a result, GDP and consumption increase persistently despite the corporate tax cuts being temporary. We provide independent evidence on this mechanism by showing that the empirical responses to corporate tax cuts of both the stock of patents and the share prices of the most patent-rich US firms line up closely with the predictions of our structural model. Counterfactual analyses highlight that applied research is quantitatively important for the magnitude and persistence of the effects of temporary corporate tax changes on TFP and GDP that we estimate with LPs, but also the elasticity of patents to permanent tax changes that Akcigit et al. (2022) estimate using historical US data. Finally, the estimates of the structural model imply social returns to innovation in the range of 20% to 75%.

Related literature. Our analysis is related to several strands of work. An influential empirical literature exemplified by Romer and Romer (2010), Barro and Redlick (2011), Mertens and Ravn (2013), Cloyne (2013) and Caldara and Kamps (2012), estimate the short-term response of GDP to tax shocks. However, these important contributions do not look at the impact on productivity and R&D spending, nor at the responses of macro variables at long horizons, both of which are a primary focus of our analysis.

Some recent studies have looked at the link between tax changes and innovation. Jones (2022) studies optimal taxation for top earners in a model where innovation cannot be perfectly targeted by a research subsidy. Akcigit et al. (2022) estimate large and positive effects of permanent tax cuts on patenting activities across US states and inventors. Cram and Olbert (2022) measure the impact of the 2021 global corporate tax reform on the stock prices of companies with different shares of intangible assets. Baley et al. (2022) look at the monetary policy response to corporate tax reforms. We complement these studies by documenting the persistent effects of temporary corporate tax cuts on innovation, TFP, and GDP.

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), Furlanetto et al. (2021), Antolin-Diaz and Surico (2022), Fieldhouse and Mertens (2023), examine the long-term effects of non-technology shocks working via hysteresis, financial frictions, monetary policy and government spending. A distinctive feature of our analysis is the focus on the medium-term effects of corporate and personal income taxes.

Structure of the paper. Section 2 presents the identification strategy and the empirical framework based on local projections. Section 3 summarises the evidence obtained using local projections and narrative identification using post-WWII US data. In Section 4, we develop an endogenous productivity model of the business cycle with two margins of innovation: basic and applied research. In Section 5, we estimate the structural model by minimising the distance between the model impulse responses and those based on the local projection estimates of Section 3. In Section 6, we provide further evidence on the specific mechanism highlighted by the structural model and find that the interaction of the market price of ideas and the tax amortisation period on intangible assets is a main channel through which endogenous productivity shapes the persistent response of GDP to corporate taxes. In Section 7, we show that applied research plays a significant role in accounting not only for the size and persistence of the estimated effects of temporary corporate tax cuts, but also for the long-term elasticities of patents to permanent tax changes. Section 8 concludes. The Appendix contains further results and robustness analyses.

2 Empirical Framework

In this section, we describe the narrative approach to identify exogenous variation in personal and corporate income taxes. We then present the empirical models to estimate their dynamic effects and the data we use. Finally, we provide details of the estimation procedure.

2.1 Identification

Our goal is to examine the effects of different tax policy reforms on productivity and innovation. We face at least three empirical challenges. First, we need information on when and how different types of tax were changed. Second, tax policy is often endogenous because policy levers tend to be adjusted in response to changes in current or prospective economic conditions. Third, given the focus on productivity and innovation, we need econometric methods that are well-suited to elicit potentially longer-term effects.

We address the first two challenges using the identified corporate and personal taxes changes from Mertens and Ravn (2013). These data are based on the original data set of Romer and Romer (2010), which identified 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 considered "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 the announcement. Key for our purpose, Mertens and Ravn (2013) subdivide the Romer and Romer (2010) shocks into corporate and personal tax reforms. This so-called "narrative" approach of looking for quasinatural 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. (2021), Hussain and Liu (2018), Cloyne et al. (2021).[1]

The literature on the effects of tax changes using narrative methods finds large effects on GDP, but typically these papers focus only on the shorter-term effects over 2 to 5 years and do not look at all at the response of productivity and innovation. A sizable part of the macroeconomic policy debate, however, has focused on the potential longer-term effects of tax reforms. Despite this, there is little direct evidence on whether fiscal policy can boost productivity, and policy recommendations often have to rely on inferring longer-term results from the short-run estimates in some of the papers referenced above.

The identification issue centres around the fact that the reduced form residuals are an unknown combination of all 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-form residuals in period t to the structural shocks can be written as:

$$u_t = A_0 \varepsilon_t \tag{1}$$

As shown by Mertens and Ravn (2013), 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 narrative shocks as instruments for observed tax policy changes. The identification restriction is that narrative shocks are uncorrelated with other structural shocks that may influence

[1] 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 popularised 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 (Ramey and Shapiro, 1998, Ramey, 2011, Crafts and Mills, 2013, Ramey and Zubairy, 2018, Barro and Redlick, 2011) and nondefence spending (Fieldhouse and Mertens, 2023).

the economy, at least conditional on the lags of Z.[2] Identification of A_0 follows directly 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 together in a vector X).

In summary, the approach amounts to using narratively identified proxies as instruments for $u_{T,t}$ in the second equation above, building 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 tax variables ζ and provides an associated estimate of the matrix η, although these latter terms do not have structural interpretation. 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, to produce impulse response functions, we need to take a stand on the precise policy experiment that is being conducted. Mathematically, in order to pin down β_1 for the construction of the impulse response functions, a decomposition of $S_1 S_1'$ is required. As in Mertens and Ravn (2013), we use a Cholesky factorisation and order last the tax rate being perturbed in 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}. In the robustness section, we show that our results are not sensitive to the ordering assumptions.

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

2.2 Econometric method

As for the empirical model, we need an econometric approach that allows us to draw inferences 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 on tax policies, which 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 and the VAR impulse response function at longer horizons will be sensitive to the number of lags included (as shown by 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 to lag truncation issues that afflict VAR methods in finite samples. For estimation, we use Bayesian methods, which provide an efficient way to compute and characterise 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, 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 that includes 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 a sequence of local projections:

$$Z_{t+h} = c^{(h)} + B_1^{(h)} Z_{t-1} + \sum_{j=1}^{P} b_j^{(h)} Z_{t-j} + u_{t+h}, \qquad var(u_{t+h}) = \Omega_h \tag{2}$$

where Z_t denotes the M variables of interest described below, h is the impulse response horizon and u_{t+h} denote residuals. As discussed in Section 2.4, we allow for the possibility that the distribution of u_{t+h} is non-Gaussian.

Given the knowledge of the relevant elements of A_0, Jordà (2005) shows that the impulse response at horizon h can be calculated as $B_1^{(h-1)} A_0$. This has two main advantages for our purposes. First, the formulation in Jordà (2005) allows us to remain as close as possible to the setup in Mertens and Ravn (2013) while still conducting estimation via local projections. Indeed, the shorter-term effects we 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 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 implement the same approach here. This is our baseline model.[3]

However, several empirical studies do not estimate the contemporaneous impact matrix separately from the reduced form dynamics. Instead, the outcome variable is regressed directly on the instruments and control variables. In our setting, such a LP can be written as:

$$Z_{i,t+h} = c^{(h)} + \beta_{ct}^{(h)} \epsilon_{ct,t} + \beta_{pt}^{(h)} \epsilon_{pt,t} + b^{(h)} Z_{t-1} + u_{t+h}, \qquad u_{t+h} \sim N(0, \sigma_h) \tag{3}$$

where $\epsilon_{ct,t}$ ($\epsilon_{pt,t}$) denotes the narrative measure of corporate (personal) tax shocks of Mertens and Ravn (2013). We refer to Equation (3) as 'Direct' model because it treats the narrative measures as the structural shocks and the estimates of $\beta_{ct}^{(h)}$ ($\beta_{pt}^{(h)}$) provide the response to the corporate (personal) tax shock under the assumption that the contemporaneous impact on the personal (corporate) tax shock is zero.

One concern with this 'Direct' model is the fact that it does not take into account the possibility of measurement error in the narrative tax proxies. This can be dealt with by using an instrumental variable approach (LPIV) as in Jordà and Taylor (2015):

$$Z_{i,t+h} = c^{(h)} + \beta_{i}^{(h)} \tau_{j,t} + \theta^{(h)} \epsilon_{k,t} + b^{(h)} Z_{t-1} + u_{t+h}, \qquad u_{t+h} \sim N(0, \sigma_h) \tag{4}$$

where $\tau_{j,t}$ for $j = ct, pt$ denotes the tax rate, which is instrumented by the narrative measure $\epsilon_{j,t}$. The regression also includes the narrative measure for the other tax rate $\epsilon_{k,t}, k \neq j$ as a contemporaneous control. In Section 3.4, we show that our results are robust to these alternative

[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 in 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.

estimation strategies.

2.3 Data

In our benchmark specification, we use the same data as in Mertens and Ravn (2013). The control variables in the sequence of local projections (2) 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.[4] 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 is represented by $\ln\left(GDP_t\right)$. All variables, except $APITR_t$ 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 data set of macro and financial variables for the US economy.[5] Following Forni and Gambetti (2014), we add this principal component as eighth control variable in our LPs to ameliorate concerns about 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 capital utilisation-adjusted Total Factor Productivity (TFP), hours worked, Research and Development (R&D) expenditure, non-residential investment, personal consumption expenditures and real wages. In Appendix A, we provide a detailed description of the variables and data sources.

2.4 Estimation

We estimate the local projections in Equations (2) to (4) via Bayesian methods. The Bayesian approach offers three main advantages in our setting. First, the error bands incorporate uncertainty regarding the A_0 matrix. Second, the Markov chain Monte-Carlo approach allows us to easily compute joint posterior distributions that can be used to assess statistical differences across shocks and horizons. Third, in Section 5, we use the IRFs produced by LPs to estimate the structural parameters of an endogenous growth model via IRF matching, for which Bayesian methods are routinely used.

[4] 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.

[5] The large data set is obtained from Mumtaz and Theodoridis (2020). 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 data set.

The local projections in Equation (2) can be written compactly as:

$$Z_{t+h} = \beta^h X_t + u_{t+h}, \qquad var(u_{t+h}) = \Omega_h \tag{5}$$

where $X_t = (1, Z_{t-1}, .., Z_{t-p})$ collects all the regressors and $\beta^h = \left(c^h, B_1^h, b_1^h, .., b_p^h\right)$ is the coefficient matrix. When the horizon is $h = 0$, the model reduces to a Bayesian VAR. Given a Normal prior for β_0 and an inverse Wishart *prior* for Ω_0, the conditional posterior distributions of these parameters are known in closed form and the posterior distribution can be approximated via Gibbs sampling. We use the draws of these parameters to construct the posterior for the contemporaneous impact matrix A_0.

For longer horizons, the estimation of the model is more complex. As discussed in Jordà (2005), the residuals u_{t+h} are nonspherical when $h > 0$. We deal with this issue in two ways. In the benchmark specification, we allow elements of u_{t+h} to have a nonnormal distribution. Following Chiu et al. (2017), we define $u_{t+h} = A^{-1} e_{t+h}$ where A^{-1} is a lower triangular matrix. The vector $e_{t+h} = (e_{1,t+h}, .., e_{M,t+h})$ denotes the orthogonalised residuals that follow Student's t-distributions with degrees of freedom ν_j and variances σ_j^2 for $j = 1, .., M$. As discussed in Geweke (1993) and Koop (2003), this assumption is equivalent to allowing for heteroscedasticity of an unknown form. In the frequentist case, Montiel Olea and Plagborg-Møller (2021) show that heteroscedasticity robust confidence intervals for LPs that control for lags of the regression variables deliver satisfactory coverage rates. In Appendix C, we report a simple Monte-Carlo experiment showing that: (i) the results in Montiel Olea and Plagborg-Møller (2021) extend to the Bayesian LPs with Student's t-disturbances, and (ii) the estimated error bands display reasonably good coverage rates even at long-horizons.[6]

Furthermore, we attempt to account for autocorrelation in u_{t+h} by modelling it directly. In a recent study, Lusompa (2021) show that the u_{t+h} follows an $MA(h)$ process. Therefore, we consider the following extended model:

$$Z_{t+h} = \beta^h X_t + u_{t+h} \tag{6}$$

$$u_{t+h} = \epsilon_{t+h} + \theta_1 \epsilon_{t+h-1} + .. + \theta_q \epsilon_{t+h-q}, \qquad \epsilon_{t+h} \sim N(0, \Omega_h) \tag{7}$$

where we allow q to grow with the horizon. As ϵ_t is unobserved, the estimation of this model is computationally intensive. In Appendix G, we show that the IRFs estimated using (6) and (7) corroborate our main findings.[7]

[6] We provide details of the estimation algorithms in Appendix D.

[7] Our results do not depend on the Bayesian approach. As we show in Section 3.4, frequentist LPs estimated via OLS or IV with confidence intervals based on HAC standard errors produce very similar results.

Finally, in the benchmark specification, the prior for β_h is centred on a mean that implies that each variable in Z_{t+h} follows an AR(1) process. The prior variance follows the Minnesota prior, with tightness set to a large number. As discussed in Appendix D, we use a non-informative prior for the free elements of A and σ_j^2.[8]

As for the 'Direct' model in (3) that is used in one of the sensitivity analyses of Section 3.4, we present frequentist estimates based on OLS. For the LPIV in (4), we adopt instead the ridge estimator of Barnichon and Brownlees (2019) with smoothing parameter set via cross-validation.[9] In either case, we construct asymptotic confidence intervals using Newey and West (1987) HAC standard errors with the number of lags set so as to match the length of the IRF horizon.

3 Empirical results

In this section, we present the main results on productivity and other aggregate outcomes using local projections and the data described in the previous section. We first focus on GDP, TFP and hours, and then move onto R&D expenditure, investment and consumption to shed light on the transmission mechanism. The final parts of the section discuss the forecast error variance decomposition and a set of robustness exercises that are reported in Appendices F and G.

3.1 Main findings

Using the approach outlined in Section 2, we report here the baseline estimates of the effects of corporate and personal income tax cuts. We begin with the responses of the average tax rate, GDP, hours and productivity, with the latter being a key and novel focus of our analysis. We then extend our empirical evidence to examine R&D expenditure, investment and consumption to shed light on the most likely mechanism driving the responses of productivity and output. Each additional variable is added to the benchmark data vector Z one at a time to avoid a sharp increase in the number of estimated parameters.

In Figure 1, we present our first set of main results. In the left column, we show the IRFs to a reduction in the average corporate tax rate. In the right column, we report the results for a

[8]Following Bańbura et al. (2010), we set the prior mean for β_h by running AR(1) regressions for each endogenous variable. The diagonal elements of the prior variance matrix corresponding to own lags p are defined as $\frac{\mu_1^2}{p^2}$ and as $\frac{s_i}{s_j}\frac{\mu_1^2}{p^2}$ for coefficients on lags of other variables. The variances $\frac{s_i}{s_j}$ account for the differences in scale between variables and are obtained as residual variance from the preliminary AR(1) regressions. We set the tightness parameter μ_1 to 10 which implies a loose prior belief.

[9]Plagborg-Møller and Wolf (2021) show that smooth local projections imply a reduction in the variance while leading to only a small increase in the bias of LPs. We present the unsmoothed 2SLS estimate in Appendix G. Our main findings of a significant response of GDP and TFP at longer horizons are unaffected by these modifications.

reduction in the average personal tax rate. The impact effect is normalised 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. Solid blue lines come from the estimated structural model that will be presented, solved, and estimated in Section 5.

The first row of Figure 1 reveals that, following a shock to corporate and personal income taxes, the average tax rates temporarily decline. The change in the average corporate tax rate (first column) loses significance after about 8 quarters and goes back to zero after around 20 quarters. Changes in the average personal income tax rate are somewhat less persistent, losing significance after 6 quarters and reaching zero after around 16 quarters. Despite the different method (ie, local projections versus 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 of Figure 1 shows the percentage response of real GDP. For the first 20 quarters, this is very comparable to the main findings in Mertens and Ravn (2013). What is new are our estimates 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 effects on the level of economic activity is still sizable and significant after 8 years. In contrast, the second column reveals that the average personal tax rate 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, which is an entirely novel focus of our analysis and is reported in the third row of Figure 1. Both income tax rate cuts boost total factor productivity on impact, with the size of the initial response to a personal income tax cut being larger than for a corporate income tax change. On the other hand, the effects of corporate tax cuts grow with the forecast horizon and remain significant even beyond business-cycle frequencies. 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. Finally, in the last row of Figure 1, hours worked do not respond to changes in corporate income tax at all, but witness a short-lived boost following a cut in personal

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

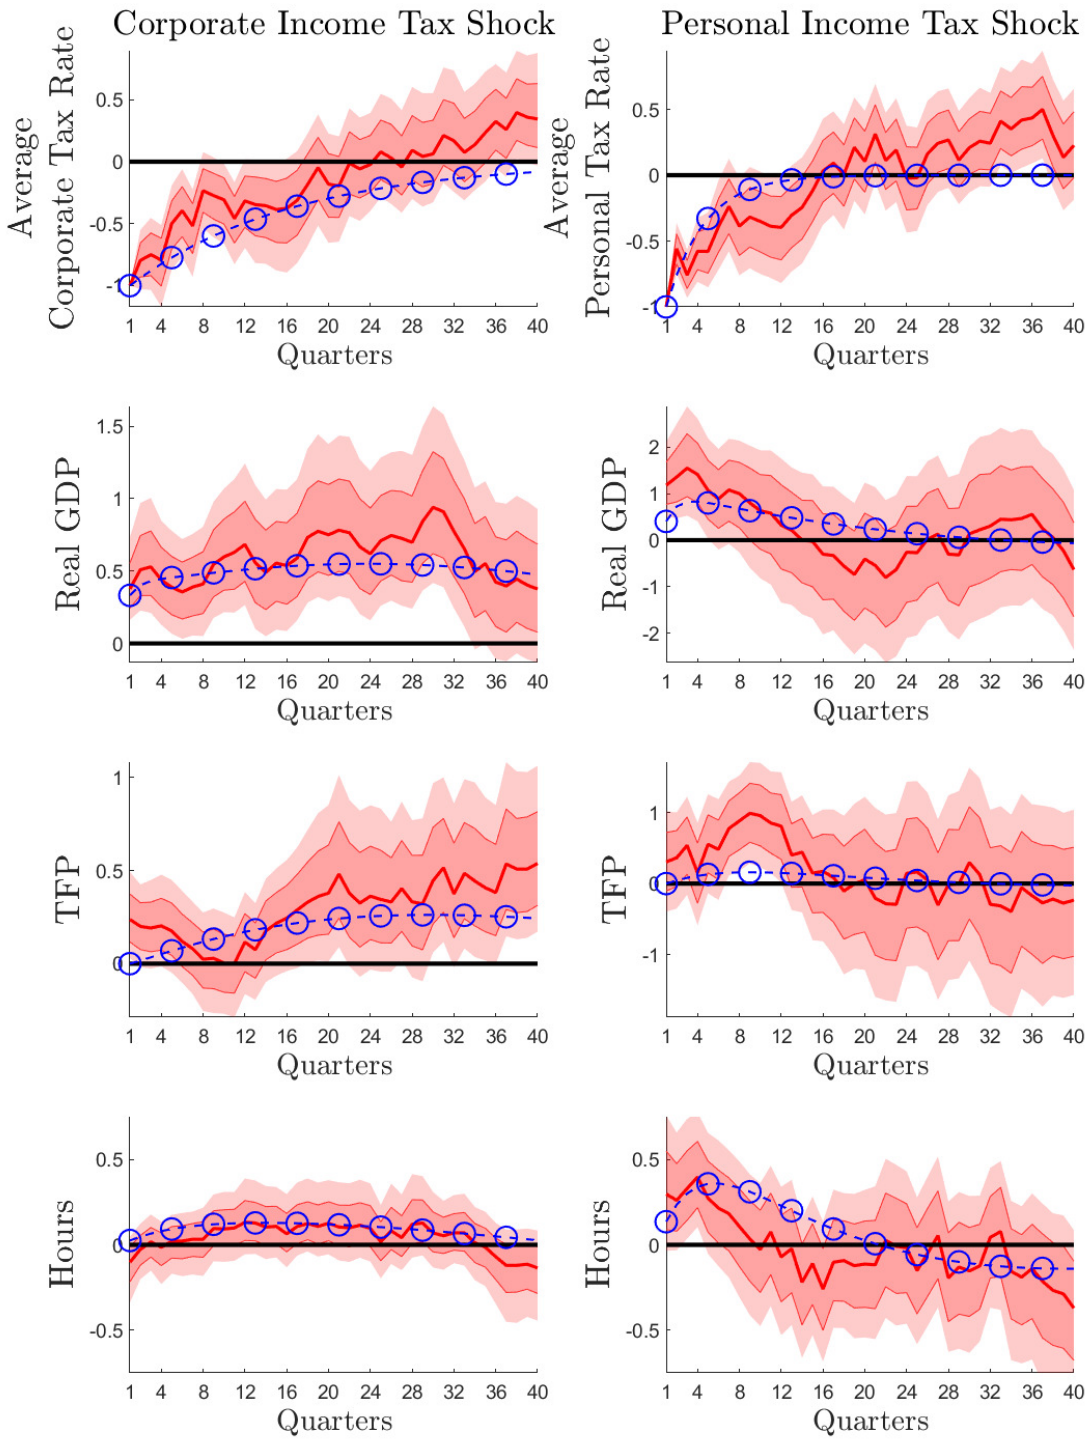

Notes: this figure shows the responses of the average tax rates, real GDP, TFP and hours to a 1% cut in the average rate of corporate income taxes (left column) and the average rate of personal income taxes (right column). Red shadow bands represent central posterior 68^{th} and 90^{th} credible sets. 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.

income tax rates.

In summary, only corporate income tax changes have large and significant effects on productivity and output at long horizons. In contrast, tax cuts on personal income have significantly larger effects on TFP, hours worked, and GDP in the short-run than corporate income tax changes. In the next section, we extend our empirical analysis to R&D expenditure, investment and consumption in an effort to shed light on the possible mechanism behind the heterogeneity documented in Figure 1.

3.2 On the mechanism

The focus on productivity and on horizons beyond 20 quarters allowed us to uncover a novel empirical finding in the previous section: changes in corporate income tax rates have very persistent effects on productivity, whereas personal income tax changes do not. In this section, we look at a number of additional variables that could offer insights into the transmission mechanism, especially at longer horizons. These are R&D expenditure, corporate investment and household expenditure. The endogenous growth literature argues that R&D spending has the potential to generate persistent effects on both productivity and output, while studies in the Real Business Cycle tradition emphasise the role of physical capital accumulation as an important propagation mechanism. Finally, given such longer-term output responses, we would also expect to see persistent effects on consumption for corporate tax changes, as opposed to short-lived effects after a personal tax change.

The findings are reported in Figure 2. 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 2 suggests that the effects of corporate tax cuts (first column) on R&D are initially negligible but become significant 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 about six years. The response of investment to corporate tax changes (second row) is equally strong, but its significance seems shorter-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 the persistent increase in productivity reported in Figure 1. In the next sections,

Figure 2: Responses of Expenditure Components: R&D, Investment and Consumption

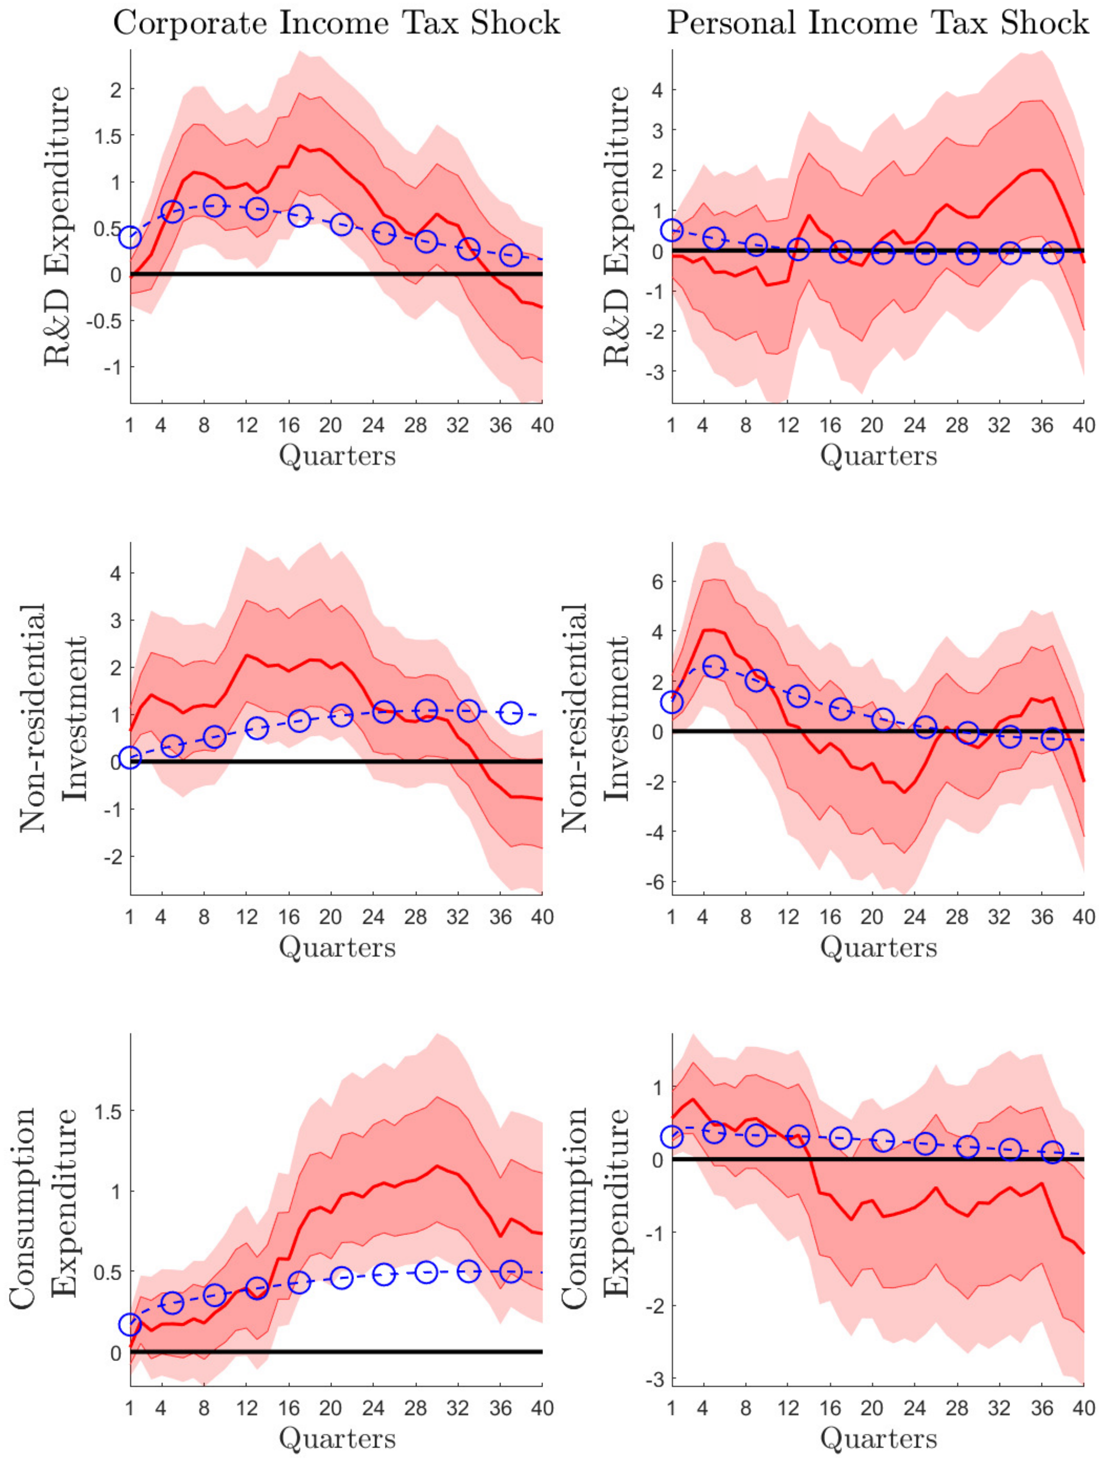

Notes: responses of R&D expenditure, non-residential investment and personal consumption expenditures to a 1% cut in the average rate of corporate income taxes (left column) and the average rate of personal income taxes (right column). Red shadow bands represent central posterior 68^{th} and 90^{th} 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.

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) paint a different picture. The response of R&D is never statistically different from zero while the change in investment is larger over the first two years but then dies out much earlier 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 innovation activities. The estimates of the endogenous growth model will reveal that this finding is consistent with a short-run labour utilization story. Finally, the response of consumption in the bottom row of Figure 2 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 labour income persistently, whereas personal taxes affect income only temporarily.

In Appendix B, we report the responses of labour productivity, real wages, and employment. Our theoretical model does not feature an extensive margin, and thus the IRF for employment will not be used in the structural estimation of Section 5. The two main takeaways from this additional analysis are that: (i) the responses of labour productivity and wages in Figure B.1 largely resemble the response of TFP in Figure 1 to either shock; (ii) the effects of both income tax shocks on employment are mostly insignificant.

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 E.1, we provide further support for this interpretation by looking at sectoral real gross output from the US Bureau of Economic Analysis' Industrial Accounts. We classify sectors into two groups based on their R&D intensity and estimate the heterogeneous effects of corporate and personal tax 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 income tax changes.

3.3 Forecast error variance decomposition

In this section, we use the LP estimates to assess the contribution of each shock to the variance of endogenous variables at different forecast horizons. The results of this exercise are summarised in Appendix Figure F.1, 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).[10]

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, investment and R&D expenditure. In contrast, the contribution of personal income tax changes to longer-term fluctuations tends to be lower than 10%.[11]

3.4 Robustness

In this section, we briefly describe a variety of sensitivity analyses that confirm the robustness of our results. Full details are reported in Appendix G. In Figure G.1, we use Bayesian LPs where the residuals are modelled as an MA process. The estimated responses of GDP, TFP and R&D to corporate tax cuts are positive and persistent. In contrast, the effects of personal tax shocks are shorter-lived. In Appendix Figure G.2, we show that the benchmark model is robust to a number of other changes in the model specification. We consider sensitivity to: (i) varying the lag length for the controls in Z, (ii) using the optimal prior strategy described in Giannone et al. (2015), (iii) including the defence news shock from Ramey (2011) as a further control, and (iv) changing the causal ordering of the two taxes as in Mertens and Ravn (2013). The solid red line and the shaded areas in Figure G.2 replicate the median estimate and 90% credible set of Figure 1. The results of each of the robustness checks mentioned above are overlayed. The main takeaway from Figure G.2, Appendix G and this section is that our main finding of significant effects of corporate income tax changes on output and productivity over the medium term is a very robust feature of post-WWII US data.

In Figures G.3, we present frequentist estimates of the responses of GDP, TFP and R&D to the two tax shocks using either the Direct model of equation (3) or the LPIV of equation (4). These two specifications employ the narrative proxies of Mertens and Ravn (2013) as exogenous regressors and instruments, respectively. While the "direct" regression of the instruments on the outcome variables

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

[11] These findings also echo results in earlier studies that focused more on short-term impact. Mertens and Ravn (2012) find that Romer and Romer (2010) tax shocks explain around 20% of the output fluctuations at business cycle frequencies, consistent with the short-term results in Appendix Figure F.1. Cloyne (2013) finds that narrative-identified tax shocks in the U.K. account for around 20% of output variation at the ten-year horizon. McGrattan (1994) finds that labour taxes account for around 25% of the in-sample variance of output and capital taxes around 5% at business cycle frequencies, using a completely different VAR-based identification approach.

produces responses (in solid grey) that are erratic, their pattern broadly matches those obtained via the smooth LPIV (in dotted red). The effects of corporate tax cuts on GDP are evident after about four years and continue up to 40 quarters ahead. The response of TFP is persistent, grows with the forecast horizon, and becomes significant over the longer term. In contrast, the effects of corporate tax changes on R&D spending occurs during the first 5 years after the shock. Regarding personal tax changes, we find little evidence of significant medium-term effects on output, productivity or R&D expenditure.

It should be noted that because the narrative proxies of Mertens and Ravn (2013) are contemporaneously correlated, the identification of the shocks in the LPIV differs from the scheme used by Mertens and Ravn (2013) and therefore the responses from these regressions are not directly comparable to our benchmark results. An alternative strategy is to employ the mutually orthogonal structural shocks from the Mertens and Ravn (2013) VAR as instruments. We estimate the following regressions:

$$Z_{i,t+h} = c^{(h)} + \beta_i^{(h)} \tau_{j,t} + b^{(h)} Z_{t-1} + u_{t+h}, \qquad u_{t+h} \sim N(0, \sigma_h) \tag{8}$$

where $\tau_{j,t}$ for $j = ct, pt$ denotes the tax rate that is instrumented by the corresponding shock from the Mertens and Ravn (2013) VAR. The impulse responses of output, productivity and R&D expenditure from these LPIV models are shown in Appendix Figure G.4 and are very similar to the results from the benchmark Bayesian LP model. Given the similarity of the impulse responses from this LPIV and the benchmark model, we use this specification to test the instrument strength at each horizon. In Appendix Figure G.5, we report the robust test statistic proposed by Lewis and Mertens (2022) and the associated 5% critical value.[12] For the corporate income tax shock, the test statistic lies above the critical value at all horizons considered in the LP. For the personal income tax shock, the test fails to reject the null for a few quarters around the 7 year mark, but this instrument appears strong at all remaining (shorter and longer) horizons.

4 A structural model with endogenous productivity

In the previous section, we have documented three main findings. Corporate income tax changes have significant effects on productivity and output in the medium term; the response of TFP to a corporate tax shock is much more persistent than the response of R&D; personal income tax changes have significant effects on productivity and output in the short term only. In this section,

[12]In our case, with one endogenous regressor, this is equivalent to the statistic of Montiel-Olea and Pflueger (2013).

we develop a theoretical framework that accounts for these results. In the next sections, we estimate this structural model by matching the empirical IRFs of Section 3 and then run counterfactual analyses on the estimated model to highlight the transmission mechanism of the tax shocks.

4.1 Overview

The model has three main ingredients: (i) endogenous productivity, to have a mechanism that could amplify the effects of tax shocks on productivity over the medium-term; (ii) two margins of innovation, so that R&D expenditure and productivity could potentially exhibit different dynamics; (iii) variable labour utilisation, to have a channel through which taxes affect productivity in the short term.

For the first ingredient, we combine elements of endogenous growth theory and business cycle analysis, following Anzoategui et al. (2019).[13] For the second feature, we introduce innovation as a two-stage process consisting of 'basic' and 'applied' research. Following Jones (2022), we refer to 'basic research' as activities that uncover fundamental truths about the world in the form of new ideas and technologies. These are the kind of R&D featured by endogenous growth models in the tradition of Romer (1990) and Aghion and Howitt (1998). But innovation is not just about new ideas or new technologies; effort and expenditure are also required to turn those ideas and technologies into new products and processes. We refer to this second type of innovation activity as 'applied research', as in Jones (2022), or 'adoption' as in Comin and Gertler (2006). In Section 6.1, we show that modelling innovation as a two-stage process not only captures realistic adoption lags but also generates a complementarity between these two margins that plays an important role in accounting for the magnitude and persistence of the TFP and GDP responses to a corporate tax shock.

For the third ingredient, we adopt the unobserved labour effort margin on the household side proposed by Galí and van Rens (2020). Finally, we model corporate and personal income taxes and amortisation tax deductions consistently with the US tax code over the sample period covered by our empirical analysis. It is worth noting that by bringing in model ingredients from business cycle analysis, such as nominal rigidities, we are able to consider New Keynesian channels of fiscal policy transmission working through aggregate demand. However, in the empirical analysis, this feature will not prove to be very important in explaining long-term responses.

[13] Growth in our model is semi-endogenous rather than fully endogenous as in Anzoategui et al. (2019). In our context, this is a more "conservative" approach because it does not build in permanent level effects from transitory changes. Furthermore, a semi-endogenous model is consistent with the observation that the US GDP growth has been relatively stable even as the average corporation tax rate has trended consistently lower in the postwar era.

4.2 Endogenous productivity: basic and applied research

In the economy, there exists a continuum of measure A_t of monopolistically competitive intermediate goods firms. Each of them manufactures a differentiated product using capital and labour with a standard production function. In Appendix I, we show that aggregate output is given by:

$$Y_t = A_t^{\theta-1}(U_t K_{g,t})^{\alpha}(L_{g,t})^{1-\alpha}. \tag{9}$$

In this section, we describe how R&D and adoption drive the dynamics of A_t. Let Z_t be the total stock of known technologies. A_t is the stock of adopted technologies, so $(Z_t - A_t)$ is the *un*adopted technology stock. Basic research expenditure –or R&D for short– increases Z_t while applied research expenditure –or adoption for short – increases A_t.

Basic Research. There is a continuum measure 1 of innovators that hire R&D-specific labour and capital to discover new ideas. Let $X_{z,j,t} = L_{z,j,t}^{\gamma} K_{z,j,t}^{1-\gamma}$ be R&D expenditure by innovator j, where $L_{z,j,t}$ and $K_{z,j,t}$ are labour and capital hired by innovator j, and γ is the labour share in innovation expenditure. The number of new technologies created by a unit of R&D expenditure (equivalently, total factor productivity in R&D), φ_t, is given by:

$$\varphi_t = Z_t^{1+\zeta} X_{z,t}^{\rho_z - 1}, \tag{10}$$

where $X_{z,t}$ is aggregate R&D spending 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.[14] We assume $\rho_z < 1$, which implies that higher aggregate R&D spending reduces the efficiency of R&D at the individual level.

Let $P_{z,t}$ denote the market price of an unadopted technology. As explained below, the relationship between the market price of an idea and the present value of ownership is determined by the tax treatment of intellectual property. Denoting $r_{z,t}$ and $w_{z,t}$ the rental rates of R&D capital and labour, respectively, we can express innovator j's decision problem as choosing $L_{j,z,t}$ and $K_{j,z,t}$ to

[14]The existence of a balanced growth path requires $\zeta = -\rho_z \left(\frac{\theta-1}{1-\alpha}\right)\left(\frac{g_y}{g_y - g_n} - \gamma\right)$, where g_y and g_n are the growth rates of GDP and the population, and the other parameters are described in text. In estimating the model, we use average GDP and population growth rates over our sample period and estimate or calibrate the remaining parameters. See Tables 1 and 2 for the estimated value of ζ and other parameters.

maximimise period t after-tax profit:

$$\max_{L_{z,j,t},K_{z,j,t}} (1-\tau_{c,t}) \left(P_{z,t}\varphi_t L_{z,j,t}^{\gamma} K_{z,j,t}^{1-\gamma} - w_{z,t} L_{z,j,t} - r_{z,t} K_{z,j,t} \right), \tag{11}$$

where the first term inside the brackets is innovator j's period t revenue, given by the product of the market price of technology ($P_{z,t}$) and the number of technologies produced ($\varphi X_{z,j,t}$). Innovator j pays corporate income tax $\tau_{c,t}$ on profits, given by revenues minus the costs of hiring workers and R&D-specific capital.

The optimality conditions for R&D (aggregated over the unit measure of innovators) equate the marginal cost and product of each factor: $w_{z,t} = \gamma P_{z,t} \frac{X_{z,t}}{L_{z,t}}$ and $r_{z,t} = (1-\gamma) P_{z,t} \frac{X_{z,t}}{K_{z,t}}$. In aggregate, $\varphi X_{z,t}$ new technologies are discovered in period t. Denoting by ϕ the one-period survival rate for any given technology, we can express the evolution of the stock of technologies as:

$$Z_{t+1} = \varphi_t X_{z,t} + \phi Z_t \tag{12}$$

Combining equations (12) and (10) yields the following expression for the growth of new technologies:

$$\frac{Z_{t+1}}{Z_t} = Z_t^{\zeta} X_{z,t}^{\rho_z} + \phi. \tag{13}$$

Applied Research. We next describe how unadopted technologies become adopted, and therefore enter productive use. There is a competitive group of "adopters", indexed by j, who convert unadopted technologies into adopted ones. They buy the rights to the technology from the innovator at the competitive price $P_{z,t}$ and convert the technology into use by employing adoption-specific labour and capital as inputs. 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, $X_{a,j,t} = L_{a,j,t}^{\gamma} K_{a,j,t}^{1-\gamma}$, according to the following expression

$$\lambda_t = \lambda \left(\frac{Z_t}{N_t^{\gamma} \Psi_t^{1-\gamma}} X_{a,j,t} \right) \tag{14}$$

where $\lambda' > 0$, $\lambda'' < 0$, $L_{a,j,t}$ and $K_{a,j,t}$ are labour and capital hired by innovator j, and γ is the labour share in innovation expenditure.

To ensure the existence of a balanced growth path, we augment $X_{a,j,t}$ by a spillover effect coming 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 $N_t^{\gamma} \Psi_t^{1-\gamma}$, where Ψ_t is a scaling factor that grows at the same rate of GDP on the balanced growth path and N_t is population. Once in usable form, the adopter sells the rights to the technology at price $P_{a,t}$, determined in a competitive market, to a monopolistically competitive intermediate goods producer that makes the new product using a Cobb-Douglas production function (described in Equation (46)). Letting $\Pi_{i,t}$ be the profits that an intermediate goods firm makes from producing a good under monopolistically competitive pricing, the present value of after-tax monopolistic profits is given by:

$$V_t = (1 - \tau_{c,t}) \, \Pi_{i,t} + \beta \phi \mathbb{E}_t \left[\Lambda_{t,t+1} V_{t+1} \right], \tag{15}$$

where $\tau_{c,t}$ is the tax rate on corporate income. An adopter's problem is choosing inputs to maximize the value J_t of an unadopted technology, namely:

$$J_t = \max_{L_{a,j,t}, K_{a,j,t}} \mathbb{E}_t \left[(1 - \tau_{c,t}) \left(\lambda_t P_{a,t} - w_{a,t} L_{a,j,t} - r_{a,t} K_{a,j,t} \right) + \phi \beta \left(1 - \lambda_t \right) \Lambda_{t,t+1} J_{t+1} \right], \tag{16}$$

where λ_t is as in Equation (14), $P_{a,t}$ is the market price of an adopted technology, and $w_{a,t}$ and $r_{a,t}$ are the rental rates of adoption-specific labour and capital, respectively. The first term in the Bellman equation reflects expected after-tax profits (expected revenues $\lambda_t P_{a,t}$ minus the costs of hiring adoption-specific labour and capital), while the second term stands for the discounted expected continuation value: $(1 - \lambda_t)$ times the discounted continuation value. The first-order conditions for labour and capital are:

$$(1 - \tau_{c,t}) \, w_{a,t} = \frac{\partial \lambda_t}{\partial L_{a,j,t}} \beta \phi \mathbb{E}_t \left[(1 - \tau_{c,t}) \, P_{a,t} - \Lambda_{t,t+1} J_{t+1} \right] \tag{17}$$

and

$$(1 - \tau_{c,t}) \, r_{a,t} = \frac{\partial \lambda_t}{\partial K_{a,j,t}} \beta \phi \mathbb{E}_t \left[(1 - \tau_{c,t}) \, P_{a,t} - \Lambda_{t,t+1} J_{t+1} \right]. \tag{18}$$

The terms on the right are the marginal benefits of 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 left 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 aggregate evolution of adopted technologies:

$$A_{t+1} = \lambda_t \phi \left[Z_t - A_t \right] + \phi A_t \tag{19}$$

where $(Z_t - A_t)$ measures the stock of unadopted technologies.

4.3 Corporate taxes and the price of ideas

The price of (un)adopted technologies (which we collectively refer to as ideas) is determined in competitive markets and, as in the model of Hall and Jorgenson (1967), given by the sum of the present value of after-tax service flows plus the tax deductions associated with ownership of ideas. Consistent with the US tax code for the period we study, we assume that the value of intellectual property assets is amortised over time, resulting in future tax deductions.[15] Following Auerbach (1989), we model amortisation as a geometric process: in every period, an owner of an intellectual property can deduct a fraction $\hat{\delta}_{IP}$ of the book value (in this case, the purchase price) of the asset from taxable profits. The remaining portion (1-$\hat{\delta}_{IP}$) is carried into the next period.

With this assumption, the present value of profits, inclusive of the purchase price $P_{a,t}$, for an entrant monopolist that buys a newly adopted technology at time t and starts production at time $t+1$ is given by:

$$\Pi_t^M = -P_{a,t} + \mathbb{E}_t \left[\beta\phi\Lambda_{t,t+1}V_{t+1} + \sum_{s=0}^{\infty} \beta^s \Lambda_{t,t+s} \hat{\delta}_{IP}^{s+1} \left(1-\hat{\delta}_{IP}\right)^s \tau_{c,t+s} P_{a,t} \right] \tag{20}$$

The first term on the right-hand side is negative because the entrant monopolist is purchasing the technology from an adopter. The second term captures the present value of monopolistic profits starting in period $t+1$ (per Equation (15)). The third term is the present value of amortisation allowances. Potential monopolists compete to buy adopted technologies and therefore, in equilibrium, lifetime profits are zero. Rearranging terms and exploiting the zero-profit condition, we can express the price of an adopted technology as:

$$P_{a,t}\left(1-d_{IP,t}\right) = \phi\beta\mathbb{E}_t\Lambda_{t,t+1}V_{t+1}, \tag{21}$$

where

$$d_{IP,t} = \sum_{s=0}^{\infty} \beta^s \Lambda_{t,t+s} \hat{\delta}_{IP}^{s+1} \left(1-\hat{\delta}_{IP}\right)^s \tau_{c,t+s} \tag{22}$$

is the present value of amortisation allowances.

As adopters compete to buy unadopted technologies and the purchase price of unadopted technologies is amortised in the same way, the analogous derivation yields the market price of an

[15] The tax amortisation periods of intangible assets in the US are defined in IRC section 704(c) AND IRC section 197 of Chapter 3 of the Audit Techniques Guide published by the Internal Revenue Service. It establishes a mandatory 15-year recovery period for assets such as goodwill, trademarks, franchises, licenses granted by governmental agencies, and customer-based intangibles. Other assets, such as patents and copyrights, are amortisable under IRC section 197 if they are purchased as part of a trade or business.

unadopted technology:

$$P_{z,t}\left(1-d_{IP,t}\right)=\phi\beta\mathbb{E}_t\Lambda_{t,t+1}J_{t+1}. \tag{23}$$

According to equations (15), (16), (21) and (23), changes in current and expected future corporate tax rates generate variation in the present value of both after-tax service flows and the amortisation allowances associated with the purchase of ideas. This leads to price fluctuations in $P_{z,t}$ and $P_{a,t}$, which in turn directly affect incentives to discover new technologies and adopt existing ones. As we show in Section 6, the tax depreciation rate (which corresponds to the time span over which amortisation is allowed in the tax code) is crucial for the ability of the structural model to generate fluctuations in the market price of ideas in response to corporate tax changes, and thereby account for the responses of output and productivity to the corporate income tax cut estimated in Section 3.

4.4 Rest of the model

The rest of the model is relatively standard and is described in Appendix I. Several features are common to many models: quadratic adjustment costs on capital (used in R&D, adoption and goods production); sticky prices à la Calvo, a Taylor rule for monetary policy; habits in consumption. The government budget constraint is balanced in every period, with lump sum taxes adjusting to balance out any difference between exogenous government consumption and the revenues raised by corporate and personal income taxation. Variable labour utilisation is modelled as an effort choice, following Galí and van Rens (2020). The household chooses hours one period in advance and faces a quadratic adjustment cost (increasing in the change in hours) in doing so. After observing the period wage, the household chooses the effort per hour, and the effective labour supply is given by hours times the effort. In the following section, we compute model impulse responses and use them to estimate the model structural parameters and perform a counterfactual analysis intended to elucidate the possible drivers of the empirical evidence in Section 3.

5 Structural estimation

In this section, we show that the theory outlined above can rationalise all our empirical findings. To do so, we estimate the model of Section 4 using a limited-information Bayesian approach and show that it accounts jointly for the responses of TFP, R&D and GDP to corporate and personal tax changes reported in Section 3. Finally, we use the estimates of our structural model to revisit a long-standing question in the endogenous growth literature: what are the social returns to innovation. In the next section, we will shed light on the mechanism behind our results by decomposing the output and productivity responses into the contributions of the various channels at play in our model.

5.1 Econometric framework

We estimate the structural model in Section 4 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 $\Phi(\Upsilon)$. The structural parameters are estimated by minimizing the distance between the theoretical model impulse responses, $\Phi(\Upsilon)$, 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 responses of the economy to corporate and personal tax cuts jointly, 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 not only conditional to both income tax shocks, but also across short and long horizons. We will then be able to conduct, in the next section, a series of counterfactual experiments where we artificially change the value of one set of structural parameters at a time to evaluate the importance of different channels for explaining the empirical evidence from LPs in Section 3. To implement this approach, we first set up the quasi-likelihood function as follows:

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

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 $p(\Upsilon)$ the prior distributions, the quasi-posterior distribution is defined as:

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

We use a random walk Metropolis-Hastings algorithm to approximate the posterior distribution. The number of iterations is set to 1,100,000 and we save every 50th draw after a burn-in of 100,000.[16] The vectors $\hat{\Phi}$ —which is based on the LPs of Section 3— and the vector $\Phi(\Upsilon)$ —which is based on

[16]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 $\text{var}(\Upsilon)$. A scaled version of $\text{var}(\Upsilon)$ 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%.

the theoretical model of Section 4— contain the IRFs (to both shocks) of the following variables: R&D, investment, consumption, GDP, hours worked and TFP. It is worth emphasizing that, by simultaneously targeting the effects of both corporate and personal taxes, we are attempting to hit a number of key moments jointly, across shocks *and* across forecast horizons.

5.2 Calibrated parameters and prior distributions

We partition the structural parameters into a calibrated set (Table 1) and an estimated set (Table 2). The discount factor, capital depreciation and the capital share are set at 0.99, 0.02 and 0.35 respectively. The markup is calibrated to target the steady-state share of profits in GDP. The coefficients of the Taylor interest rate rule for monetary policy are borrowed from Anzoategui et al. (2019). Following Wen (2004), the employment adjustment cost for the three types of labour is set to $\psi = 0.35$ (whereas the elasticities of labour effort are estimated). The government spending share and the steady state tax rates are set to their sample averages. To calibrate the tax depreciation rate for capital ($\hat{\delta}_K$), we average the estimated present value of depreciation deductions employed by Hall and Jorgenson (1967) and House and Shapiro (2008), since those two sets of estimates bookend the time period covered by our data. We calibrate the tax depreciation for intellectual property assets ($\hat{\delta}_{IP}$) to match the 15-year amortisation period allowed by the US tax code.[17] Turning to the technological parameters, we calibrate the steady technology adoption rate $\bar{\lambda}$ to 0.05 (quarterly), implying a diffusion lag of five years, in line with the evidence in Comin and Hobijn (2010); the rate of technological obsolescence, $(1-\phi)$, is 0.08 based on the estimates in Li and Hall (2020); and the labour share of production in R&D and adoption activities, γ, is set to 0.9, consistent with R&D expenditure data from the National Science Foundation.

In Table 2, we collect the parameters that will be estimated and their prior distributions. The table also reports moments from the posterior distribution, which will be discussed in the next section. Prior distributions are chosen to be diffuse but centred on values typically found 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 studies, 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 utilisation mecha-

[17] For robustness, we have also tried a version of the model in which an R&D subsidy is calibrated to the rate estimated by the OECD. We find that the inclusion of a static subsidy of empirically plausible magnitude has a very small effect on the conditional model dynamics, and thus we do not include it in the baseline model for parsimony.

Table 1: Calibrated Parameters

Parameter	Description	Value	Source
Preference & Households			
β	Discount factor	0.99	
ψ_N	Employment adjustment	0.35	Wen (2004)
Technology			
g_y	100*SS GDP growth rate	0.91	Sample average
g_n	100*SS population growth rate	0.35	Sample average
GY	Government spending/GDP	0.16	Sample average
α	Capital share	0.35	
δ	Capital depreciation	0.02	
ς	Markup	1.087	Profits/GDP=8%
$\bar{\lambda}$	SS technology adoption rate	0.05	Anzoategui et al. (2019)
$1-\phi$	Technology obsolescence	0.08	Li and Hall (2020)
γ	Labor share in R&D expenditure	0.9	NSF data
Taxes			
$\bar{\tau}_c$	SS Corp. Tax	0.19	Sample average
$\bar{\tau}_p$	SS Lab. Tax	0.3	Sample average
$\hat{\delta}_K$	Tax depreciation (capital)	0.0165	Hall and Jorgenson (1967) House and Shapiro (2008)
$\hat{\delta}_{IP}$	Tax depreciation (IP)	0.0285	US tax code (15y amortization period)
Monetary Policy			
ρ_r	Smoothing	0.83	Anzoategui et al. (2019)
ϕ_y	Output	0.385	Anzoategui et al. (2019)
ϕ_π	Inflation	1.638	Anzoategui et al. (2019)

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	0.36	[0.12, 0.65]
χ_g	Inverse effort elasticity (goods)	gamma	1	0.5	0.25	[0.08, 0.92]
χ_a	Inverse effort elasticity (adoption)	gamma	1	0.5	0.76	[0.27, 1.68]
χ_z	Inverse effort elasticity (R&D)	gamma	1	0.5	1.76	[1.03, 2.86]
Frictions & Production						
f_a''	Adoption adjustment	normal	4	1.5	3.14	[0.6, 5.86]
f_z''	R&D adjustment	normal	4	1.5	3.49	[0.7, 6.12]
f_I''	Investment adjustment	normal	4	1.5	0.81	[0.11, 4.79]
ν''	Capital utilization adjustment	beta	0.6	0.15	0.46	[0.35, 0.59]
ξ_p	Calvo prices	beta	0.5	0.2	0.26	[0.09, 0.5]
Endogenous Technology						
θ-1	Dixit-Stiglitz parameter	gamma	0.15	0.1	0.68	[0.46, 0.96]
ρ_λ	Adoption elasticity	beta	0.5	0.2	0.72	[0.57, 0.84]
ρ_Z	R&D elasticity	beta	0.5	0.2	0.27	[0.17, 0.4]
ζ	R&D returns to scale	-	-	-	-0.2	[-0.3, -0.13]
Shocks						
$\rho_{\tau,c}$	Corporate taxes AR	beta	0.85	0.07	0.94	[0.93, 0.95]
$\rho_{\tau,p}$	Labour taxes AR	beta	0.85	0.07	0.76	[0.71, 0.8]

nisms. Estimates of the elasticity of patenting to R&D expenditures, analogous to ρ_Z in the model, vary widely in the literature (Danguy et al., 2013) but are generally below 1. Accordingly, we use a beta prior centred on 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).[18] To avoid tilting the balance in favour of any particular adjustment cost mechanism, we use the same prior capital investment adjustment costs in each of the sectors. We are not aware of existing estimates of the (inverse) elasticity of effort, χ. Consequently, we choose a relatively uninformative prior centred at 1.

Finally, in Appendix H, we report the impulse response functions of output, productivity and R&D implied by our prior distributions. The goal is to check whether any of the prior choices made in this section may build in a tendency for our posterior estimates to spuriously detect significant effects at long horizons. As shown in Appendix Figure H.1, our prior distributions for the structural parameters are centred around values that imply: (i) income tax changes have no long-term effects on the economy; (ii) productivity does not move much after either tax shock.

5.3 Posterior distributions

In this section, we discuss the posterior distributions of the structural parameters of the model in Section 4 estimated by minimising the difference between the IRFs of the theoretical model, $\Phi(\Upsilon)$, and the IRFs of the baseline LPs of Section 3, $\hat{\Phi}$, to both tax rates. The posterior median and central 90% credible set of the key parameters of interest are reported in the last two columns of Table 2. The model impulse responses (evaluated at the posterior medians of Table 2) are shown in Figures 1 and 2 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 closely aligned with their LP counterparts in 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. All of these parameters are included in the calculation of the social returns to R&D, shown below, which provides a useful way to relate the implications of our estimates to the existing literature. The inverse effort elasticity is close to the value of 0.3 that Galí and van Rens (2020) calibrate to match second moments of US labour market fluctuations.

[18] 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.

The estimation places a modest weight on investment adjustment costs, habit persistence and price stickiness (top of Table 2). Interestingly, by incorporating an endogenous growth mechanism, our estimates seem to downplay significantly these more 'traditional' ways of generating persistence and amplification. In particular, adjustment costs on investment in physical capital are estimated to be much lower than the values reported by Christiano et al. (2005), Smets and Wouters (2007), Justiniano et al. (2010). Unlike conventional medium-scale business cycle models, however, our framework features a range of additional sources of endogenous persistence via research spending and innovation. More specifically, the estimation 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), Bond et al. (2005) and Chiavari and Goraya (2023). Finally, we also 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 the earlier literature cited above. We interpret this finding as suggestive evidence that the omission of R&D spending and technological adoption in the business cycle models routinely used for policy analyses might distort inference on the importance of physical capital investment and its adjustment costs for business cycle fluctuations.

5.4 The social returns to R&D

An instructive way to summarise the estimates of our structural model is to revisit a fundamental question in growth theory: what are the social returns to investment in innovation? To this end, we follow the variational approach of Jones and Williams (1998), modified to account for the two margins of innovation featured in our model: R&D and adoption. The social returns to innovation are calculated as the return in additional units of consumption relative to the balanced growth path of reallocating one unit of output from consumption to R&D today, and consuming the proceeds in the future. In our model, the future proceeds from an increase in R&D today are the sum of the two components in the Jones and Williams (1998) calculation, plus a novel dimension due to the adoption margin: (i) the additional output generated, (ii) the future reduction in R&D such that the subsequent stock of unadopted ideas is unchanged, and (iii) the future reduction in adoption expenditure such that the subsequent stock of adopted ideas is unchanged. In Appendix J, we show that the social returns, $\tilde{r}_{RD}$, are given by:

$$\tilde{r}_{RD} = \frac{\tilde{d}_t}{\tilde{P}_{Z,t}} + g_{\tilde{P}_Z}, \tag{24}$$

where $\tilde{d}_t$ is the "social dividend" (which consists of additional output, the reduction in adoption expenditure, and the productivity change associated with the increase in R&D), $\tilde{P}_{Z,t}$ is the cost to society of a new unadopted idea in units of consumption, and $g_{\tilde{P}_Z}$ is the rate of change in the social cost of producing new ideas.

Endowed with Equation (24), we use the posterior distributions in Table 2 to calculate the social returns to R&D implied by our structural model. We estimate that the social returns to investment in innovation, $\tilde{r}_{RD}$, range from 20.8% to 74.5% (95% confidence level), with a posterior median of 35.9%.[19] Excluding the consumption gains to adoption from the social dividend lowers this interval to [14.9%,40%] with a median value of 22%, highlighting the importance of the complementarity between R&D and adoption in determining the social returns to innovation.

6 Inspecting the transmission of corporate tax changes

A main novel empirical finding from the previous sections is that temporary corporate tax changes have persistent effects on aggregate productivity and output at horizons beyond business-cycle frequencies. In this section, we use the structural model to perform a set of counterfactual simulations that highlight the role of endogenous productivity, and the two main forces behind it: (i) innovation, and (ii) the tax amortisation benefits on intellectual property. We present independent evidence that the responses of both patents and the share prices of the most ideas-rich US firms, estimated with local projections, are close to the *untargeted* IRFs implied by our structural model. Finally, we show that the impact response of the market price of ideas (and therefore the medium-term effects of corporate taxes on TFP and GDP) would be far smaller and less persistent in a counterfactual world with a shorter tax amortisation period on intangible capital investment.

6.1 Endogenous productivity

The goal of this section is to elicit the role that endogenous productivity plays in accounting for the medium-term response of output to corporate tax changes. To this end, we proceed in two steps. First, we decompose the (log) GDP response into the contributions of TFP, capital and capital utilization, labour and labour utilization, using the final goods production function (Equation 9):

$$\Delta \log Y = (\theta - 1)\, \Delta \log A + \alpha \left(\Delta \log U + \Delta \log K \right) + (1 - \alpha) \left(\Delta \log e_g + \Delta \log H \right). \tag{25}$$

[19] Our estimates, which are based on US tax changes on corporate and personal income over time, are remarkably similar to those obtained by Bloom et al. (2013) exploiting variation in R&D tax credits across time and US states.

Second, we switch off the adoption margin, and then we turn off endogenous productivity altogether, so as to isolate the contribution of each of the two margins of innovation to the output response.

Figure 3: GDP Decomposition and Counterfactual Analyses

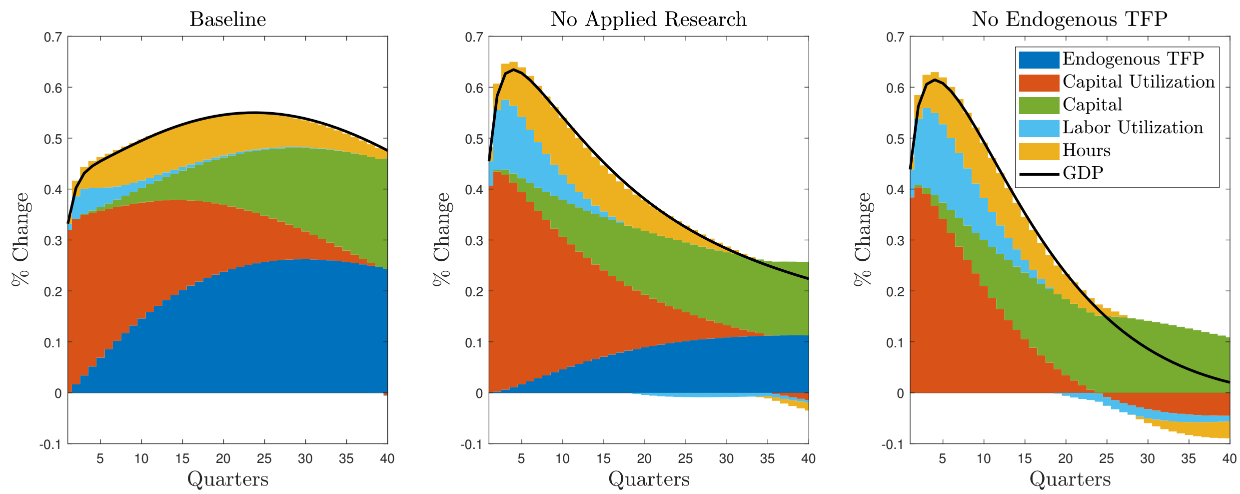

Notes: this figure plots the model impulse responses of aggregate output and its components (see Equation 9) to a corporate income tax shock. From left to right, these are the responses of the baseline model, a model in which the diffusion rate of new ideas is constant ("No Applied Research") and a model with no innovation. To construct the counterfactual plots, we re-estimate the restricted models following the procedure described in Section 5. Parameter estimates for the restricted models are in Appendix L.

The left panel of Figure 3 plots the decomposition (25) implied by our estimated model. The black line is the total response of GDP (also shown in Figure 1), and the shaded areas represent the contribution of each variable. TFP accounts for the largest share of the medium-term effects, with the rest explained by capital accumulation. A cut in corporate taxes boosts after-tax profits, which increases the market price of ideas and therefore the incentive to discover new ideas and adopt existing ones. Moreover, the rise in adoption efforts pushes up the adoption probability and reduces the expected "time to market" of innovation, further raising the incentives to innovate. In general equilibrium, more innovation also fosters capital accumulation by boosting the marginal product of capital.

To elicit the role of the adoption margin, in the middle panel of Figure 3 we show a similar decomposition, but for a reestimated version of the model in which the rate at which technology is adopted remains constant at the steady-state level. The medium-horizon response of GDP in the restricted model without adoption is half that of the baseline model, suggesting that the additional amplification provided by applied research is an important mechanism to explain the dynamic effects of corporate tax changes.

Finally, we can get a clearer sense of the total contribution of innovation and endogenous TFP in accounting for the response of the economy to a corporate tax cut by looking at the right panel of Figure 3. This plots a similar decomposition to the other two panels, except that it is based on a model without the endogenous productivity channel (akin to a RBC model with factor utilisation). To give this most restricted model the best chance to match the data, we re-estimate the parameters of this specification following the procedure described in Section 5. Without an endogenous productivity channel, the model is unable to reproduce the medium-term persistence of the GDP response reported in our empirical estimates. We conclude that endogenous productivity accounts for the vast majority of the medium-term response of GDP.

6.2 Innovation: price and quantities

At the heart of our model is a theory of investment in ideas. R&D and applied research firms allocate labour and capital as a function of the prices of inputs – wages and the rental rate of capital – and of the prices of outputs, the market prices of adopted and unadopted technologies. The first-order effect of a corporate income tax cut is a jump in the market price of ideas, which in turn stimulates investment in innovation (in the form of both higher R&D and applied research expenditure). Increased investments lead to larger stocks of unadopted and adopted technologies and, thereby, cause persistently higher productivity and GDP. We provide further empirical evidence for this transmission mechanism in Figure 4. In Panel A (left column), we plot both the empirical and the model-implied responses for measures of quantities and prices of innovation, as well as the overall response of asset prices. Empirical IRFs based on LPs are shown as red lines and shaded areas whereas the model IRFs are displayed as blue lines with circles. It is important to emphasise that we do not target these impulse responses in estimation; instead, we use them as a validation for the empirical merits of the model mechanism. In Panel B, we provide further evidence on the response of investment in intangible assets for which, however, there is not a direct model counterpart. Descriptions of the variables are detailed in the Data Appendix A.

Patents. In Panel A top row of Figure 4, we record the empirical response of the aggregate stock of patents (red shaded areas) and the model response of the stock of ideas Z_t (blue line with circles). In the data and in the model, the stock of knowledge remains above its steady-state level ten years after the shock; furthermore, the model does a good job in reproducing the magnitude of the response of the patent stock, despite the fact that this IRF has not been targeted by our IRF-matching structural estimation method. A long-standing literature (exemplified by Griliches (1990)) has argued that patents are a useful measure of technological progress. The model ability to match the magnitude and evolution of the response of both patents and real variables (TFP, GDP, etc.)

represents further evidence that not only patents contain relevant information about technological progress, but also that our estimated model can replicate the joint dynamics of innovation effort, innovation output, and the long-term productivity gains of innovation. We reiterate this point in Section 7, where we find that the model response of the stock of ideas to permanent tax cuts is close to existing empirical estimates.

Trademark assignments. In the top row of Panel B of Figure 4, we report a measure of transactions in the market for ideas: the count of transactions in the secondary market for trademarks (assignments). These transactions are registered with the USPTO when this particular type of idea - a trademark - changes owner for any reason. Consistent with the mechanism in the model, where a corporate tax cut leads to a prolonged period of increased innovation activity, trademark assignments show a positive LP response, which is significant up to 30 quarters after the shock.

The market price of ideas. The middle row of Figure 4 Panel A displays the difference in the stock price responses of value-weighted stock portfolios that are computed by ranking companies according to the value of their patents. To form these portfolios, we estimate firm-level patent stock values by applying the perpetual inventory method onto the patent-level data put together by Kogan et al. (2017). We compare this "market price of ideas" to the model response of the difference between: (i) the price of a portfolio consisting of adopted and unadopted ideas; and (ii) the price of capital. The market price of ideas jumps on impact and smoothly returns to zero over the forecast horizon, consistent with the key mechanism in the model. In the bottom row of Panel A, we plot the response of the stock market (Dow Jones index) in red and the response of a value-weighted portfolio of all assets in the model economy (i.e. ideas and capital) in blue. In sharp contrast to the effects on the market price of ideas, both in the data and in the model, the stock market responds gradually to a change in corporate taxes, as ideas and capital accumulate over time.[20]

Intangible assets. The last two rows of Panel B refer to additional measures of intangibles: (i) investment in Intellectual Property Products (IPP) other than R&D (as measured in the national accounts), and (ii) the measure of the stock of organizational capital proposed by Peters and Taylor (2017), which is calculated by capitalising SG&A expenditures. In line with the effects of corporate tax changes on other measures of investment in innovation, the response of IPP investment displays

[20]It should be noted that while the estimated structural model is able to replicate well the *dynamics* of the stock market responses to corporate tax cuts over the forecast horizon, it does not match the *level* of the stock price responses, possibly because –by design– we have not included features that could potentially account for the stock price volatility observed in the data. Accordingly, in Figure 4, the LP and the model IRFs for the stock prices are plotted on different scales, respectively on the left and right vertical axes of each stock price chart.

Figure 4: Evidence on the Mechanism

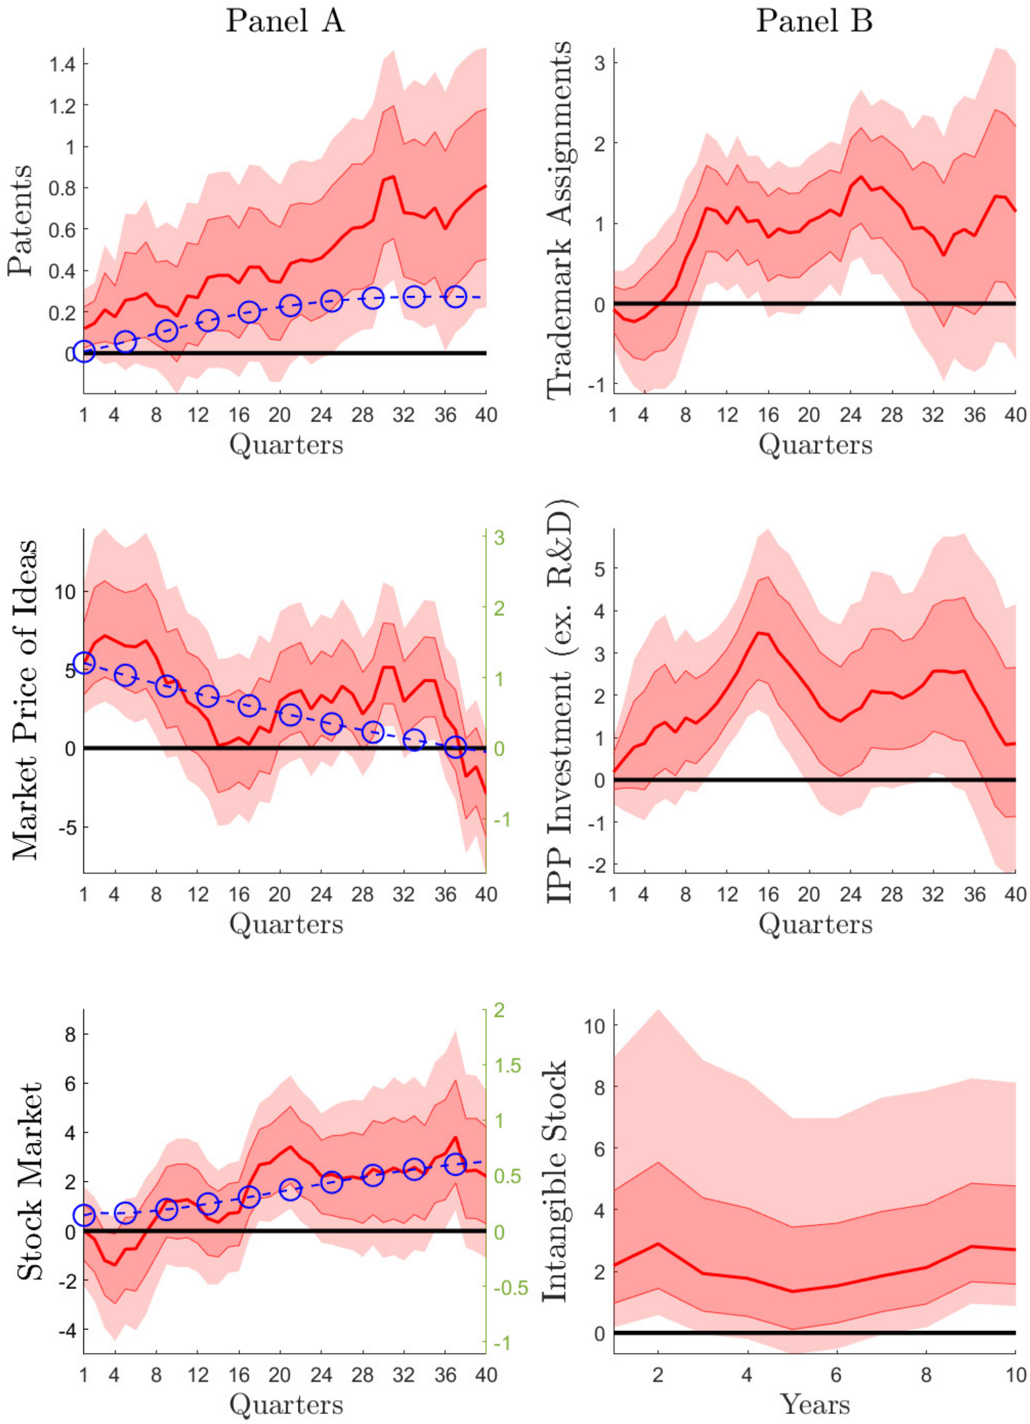

Notes: Panel A (left column) shows the responses of (from top to bottom) the stock of patents, the ideas stock price premium and the stock market to a 1% cut in the average rate of corporate income taxes. Red shadow bands represent central posterior 68^{th} and 90^{th} credible sets. Blue lines with circles represent the impulse responses of the model in Section 4 evaluated at the posterior median of estimated model parameters. The blue lines with circles in the second and third rows of Panel A are plotted on the right-hand axis. Panel B (right column) shows the responses of trademark transactions, intellectual property products investment (excluding R&D) and the stock of intangible assets. These variables are described in Section 6.2, and data sources are described in the Data Appendix A.

persistent dynamics that extend beyond the horizon of the shock itself. The stock of intangibles, like patents, is significantly higher during the 10 years after the corporate income tax cut.

6.3 Tax allowances on the amortisation of intellectual property

In Section 5, we set the tax amortisation period for ideas to 15 years, equivalent to $\hat{\delta}_{IP} = 0.0285$, consistent with Section 197 of the IRS Code which allows for straight-line amortisation of intellectual property assets over a 15-year period. In Appendix L, we further show that: (i) tax amortisation benefits on intangible assets are a salient feature of the tax codes of many countries around the world; (ii) the legal tax amortisation periods for patents, technology and trademarks vary widely across countries; (iii) advanced economies tend to set a much longer tax amortisation period (on average around 15 years) than emerging markets (over about 5 years).

In this section, we explore the role played by the tax treatment of ideas. To do so, we calculate tax depreciation rates that equate the present value of tax deductions to straight-line amortisation over varying amortisation periods. We use our estimated model to compute the responses of: (a) GDP at long horizons and (b) the market price of ideas on impact, to a corporate tax cut as a function of the amortisation period, ranging from 0 (i.e. fully amortised within the year) to 40 years (which we take as a proxy for codes in which purchases of intellectual property cannot be amortised at all).

The results from this exercise are summarized in Figure 5. A main finding from Panel A is that the medium-term response of output to a temporary corporate tax cut increases monotonically with the tax amortisation period. For short amortisation periods, corporate tax cuts would only have a modest, even negative, impact on GDP at long horizons. In contrast, if intangible assets were not deductible at all, at the far right of the chart, the medium-term effects on output would be maximized. In between these two extremes, the steeper increases occur between one and ten years. After that, the tax amortisation benefits curve flattens and values of 12, 15, 20 or the number of years in the vertical lines implied by the calculations in Hall and Jorgenson (1967) and House and Shapiro (2008) for tangible capital would all produce similar medium-term effects on GDP. Panel B paints a similar picture: the *impact* response of the market price of ideas grows monotonically with the tax amortisation horizon.

The findings of this section highlight the central role that the tax treatment of ideas plays in determining the output response to corporate income tax changes. In Equations (21) and (23), the present value of tax deductions, d_{IP}, appears as a wedge that decreases in the amortisation length. This equals exactly τ_c for instant amortisation and equals zero in the case of no amortisation at

Figure 5: GDP and price of ideas response as function of the tax amortisation period on intangibles

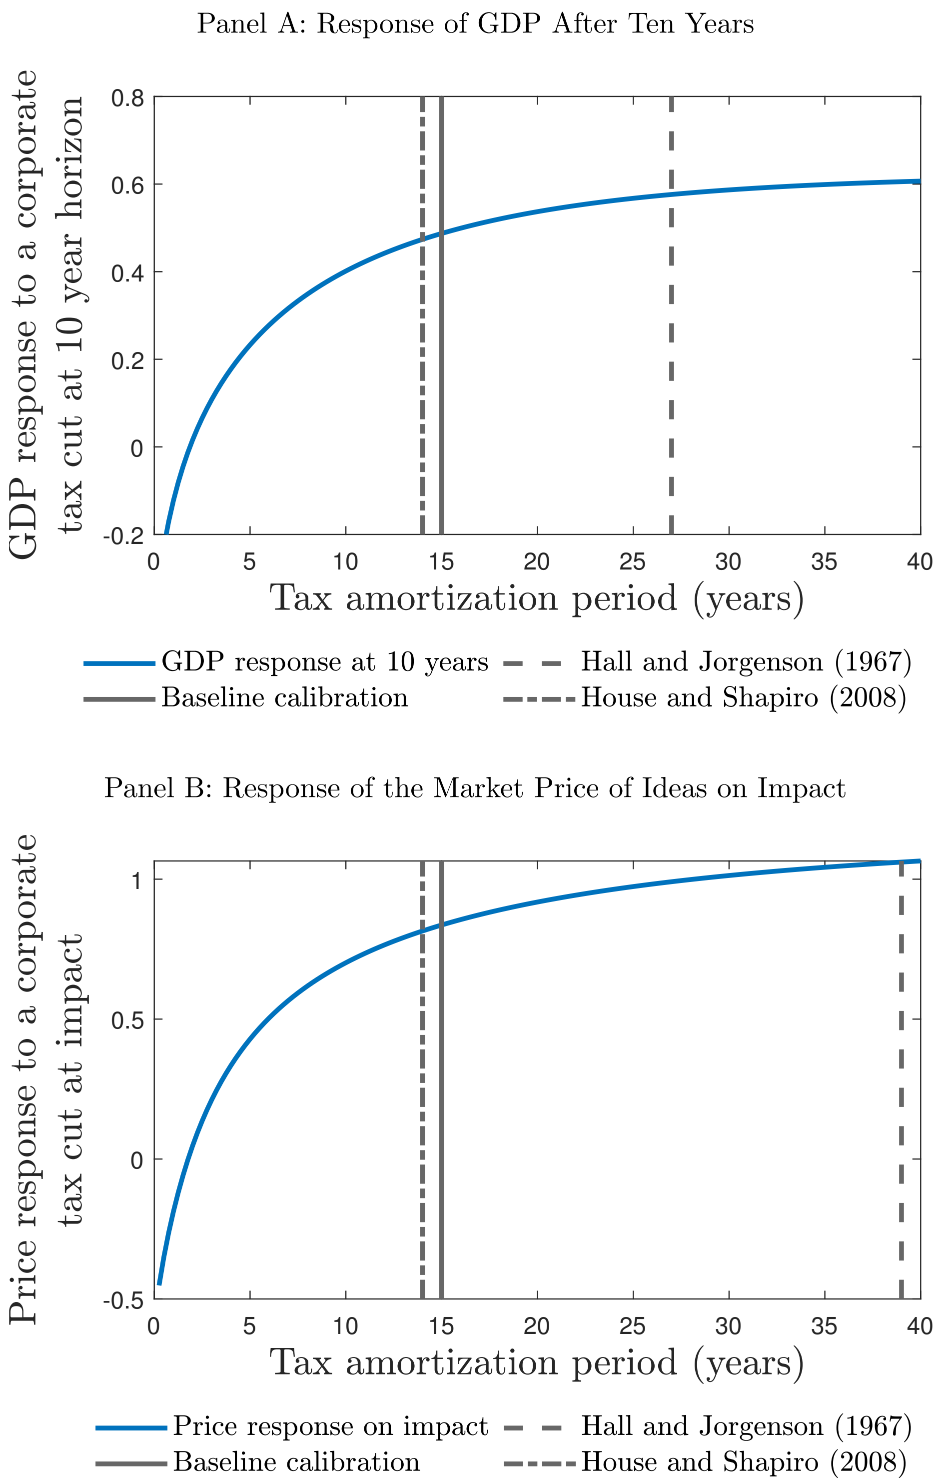

Notes: this figure shows the responses of output at the 10 year horizon (Panel A) and of the market price of ideas on impact (Panel B) to a 1% cut in the average rate of corporate income taxes as a function of the tax allowance amortisation period on intangible capital investment, implied by the estimates of the structural model presented in Section 4. Vertical lines represent the value of the tax amortisation period on intangibles used in Hall and Jorgenson (1967) and House and Shapiro (2008) for tangible assets.

all (which is equivalent to setting $\hat{\delta}_{IP}$ to 1 or 0 in equation 22, respectively). Because the right-hand sides of Equations (21) and (23) are the value functions of *after tax* profits, when $d_{IP} = \tau_c$ the tax rates cancel on both sides in steady state, and the market price of ideas is equal to the present value of pretax profits. Extending this logic to a dynamic setting, the closer is the value of the wedge to the tax rate itself (and hence the shorter the amortisation window), the smaller is the response of the market price of ideas to changes in corporate taxes. This reduces additional incentives to engage in R&D and adoption, and, therefore, dampens the response of GDP after forty quarters. In other words, shortening the tax amortisation period reduces the distortionary effects of the corporate income taxes on the market price of ideas, increasing steady-state innovation and GDP but reducing their sensitivity to the changes in the corporate tax rate.

7 Long-run elasticities

In the previous section, we have shown that applied research plays an important role to account for the persistent response of GDP to a *temporary* corporate tax change. In this section, we explore this margin of innovation further, looking at its implications for the effects of *permanent* tax changes. For this purpose, we use the estimates of our model in Table 2 to compute the elasticities of the stock of knowledge (Z) and GDP to a 1% permanent change in the marginal rates on corporate and personal taxes, and compare them to the predictions of: (i) an estimated restricted model without adoption, and (ii) the estimates available in the empirical literature. The goal is to assess the ability of estimated endogenous growth models with and without applied research to quantitatively account for other salient features of the data.

The findings of our analysis are reported in Table 3. We start by looking at the elasticity of patents to changes in the marginal rates on corporate and personal taxes, in Panel A, as these statistics can be readily compared to the estimates available from empirical studies. Then, we look at the predictions on the output elasticity, in Panel B, for which, to the best of our knowledge, there is no evidence. The first row simulates a 1% permanent change in the marginal tax rate on corporate income whereas the second (third) row refers to a 1% permanent change in the marginal tax rate on income at the top (bottom) of the income distribution. The first two columns refer to the elasticities implied by the estimates of the unrestricted model and the estimates of the restricted model (i.e. with no applied research), respectively, where scientists (workers) exemplify top (non-top) earners. The third column reports the elasticities of patents to tax rates estimated by Akcigit et al. (2022), who also consider permanent changes in the marginal tax rate at the 90^{th} and 50^{th} percentile of the income distribution.

Table 3: Elasticities of Patenting and Output to Permanent Tax Rate Changes

	Baseline	No Applied Research	Empirical Literature
Panel A – Long-run elasticity of innovation to:			
Corporate Income Tax	1.5	10.12	**1.98***** [1.50,2.46]
Top Personal Income Tax	1.14	11.62	**1.45***** [1.22,1.68]
Bottom Personal Income Tax	-0.14	-1.28	1.67 [-0.69,4.03]
Panel B – Long-run elasticity of real GDP to:			
Corporate Income Tax	1.84	1.19	
Top Personal Income Tax	1.25	0.85	
Bottom Personal Income Tax	0.61	0.7	

Note: Panel A compares the effects of permanent tax shocks on the stock of unadopted ideas (Z) in two versions of our structural model to the effects on patents reported in Akcigit et al. (2022), Table 3, panel A (corporate – corp. MTR – and top personal income tax – MTR90) and Table C8 (bottom personal income tax – MTR50). The top personal income tax in the model is paid by innovation workers, and the bottom personal income tax is paid by goods production sector workers. Panel B displays the effects of permanent tax cuts on output in the model economy. The elasticities in the first (second) column are based on the estimates of the unrestricted model in Table 2 (the estimates of the restricted version of the model with no adoption). The third column reports the estimates in Akcigit et al. (2022). Consistent with this latter study, all elasticities are computed with respect to the 'keep' rate of (1-tax rate).

Several interesting results emerge from Table 3. First, for the marginal tax rates on corporate income and top personal incomes (the first two rows), our baseline structural model generates elasticites of the stock of knowledge that are very close to the patent elasticities estimated by Akcigit et al. (2022). This is consistent with the finding in Figure 4 that the model generates responses of the knowledge stock of magnitudes similar to the empirical responses observed from patents. Second, the unrestricted model (with adoption) of Section 5 also predicts that patents should not move much following a tax rate change at the bottom of the income distribution, which is consistent with the insignificant coefficient estimated by Akcigit et al. (2022). Third, in sharp contrast, the second column of Table 3 shows that the estimated endogenous productivity model with no applied research counterfactually predicts that the effects of a permanent change in the marginal tax rate on either corporate income or top personal income would be much larger than the estimates in Akcigit et al. (2022).[21]

Finally, in Panel B of Table 3, we report the output elasticities implied by the estimates of the baseline model and the restricted model without applied research, respectively. Although, to our knowledge, we do not have any existing empirical estimates of output elasticities to compare with,

[21] The intuition is that in the model with no adoption, the elasticties of R&D spending need to be much larger to match the LP impulse responses of Section 3. In other words, an endogenous TFP model with no adoption can account for our IRFs evidence only at the cost of empirically implausible elasticities of patents to either tax shock.

it is interesting to note that the baseline model produces a much larger response of output to a permanent change in either the corporate income or the top earnings tax rate. This suggests that, in the restricted model with no adoption, the output elasticity to patents is much smaller than in the estimated baseline model with adoption, despite the finding in Panel A that the restricted model has a much higher elasticity of patents to a tax change. In other words, the omission of an applied research margin of innovation implies that in the restricted model patents are very responsive to a permanent tax change, but their contribution to the GDP response is much diminished relative to an estimated endogenous growth model that features also an applied research margin.[22]

8 Conclusions

On 16^{th} November 2017, the House of Representatives passed the *Macroeconomic Analysis of the "Tax Cuts and Jobs Act"*, prepared by the staff of the nonpartisan Joint Committee on Taxation (2017) of the US Congress. The word 'productivity' is barely mentioned —a single time in the appendix— and the only (neoclassical) mechanism through which tax cuts may have any medium-term effect is described as: '*increased investment in business capital is projected to arise from the reduction in the tax rates on corporate [...]. This investment results in a gradual accumulation of capital stock, which is forecast to reach its peak sometime in the second half of the budget period*' of ten years.[23]

This paper uncovers a novel channel through which fiscal policy can boost aggregate productivity and innovation at horizons beyond the business cycle. A cut in corporate income taxes stimulates investment in R&D. This encourages both the *creation* of new ideas and technologies —which we refer to as 'basic research'— and the *adoption* of existing ideas and technologies —which we refer to as 'applied research'. We show that 'applied research' is crucial for the ability of an estimated endogenous growth model to account for the magnitude and persistence of the effects of temporary corporate tax changes on TFP and GDP that we document on post-WWII US data. In contrast, changes in the average tax rate on personal income have a limited impact on R&D expenditure and innovation over our sample, and therefore their effects on productivity and output are stronger at shorter horizons.

Our estimated endogenous productivity model suggests that two ingredients are key to the

[22] As in the canonical model of Hall and Jorgenson (1967), if all margins of corporate expenditure (on both physical capital and intellectual properties) were fully tax deductible (i.e., $\hat{\delta}_{IP} = \hat{\delta}_K = 1$), permanent cuts in corporate income taxes would only affect the government budget constraint but have no real effect on output at either short or long horizons. See Abel (2007) for an intuitive explanation of this neutrality result in a general equilibrium setting

[23] As detailed in the Joint Committee on Taxation (2017) report, the three general equilibrium models used by the staff of the Joint Committee to simulate the effects of tax changes on GDP feature no innovation or endogenous TFP.

transmission of tax shocks: the market price of ideas and the tax amortisation period of intellectual property assets. A cut in corporate income taxes leads to a jump in the price of ideas, which in turn provides firms with an incentive to innovate more, by investing in 'basic' and 'applied' research. We provide direct empirical evidence on this mechanism. On impact, the share prices of more patent-rich firms increase more than those of firms with few patents. At business cycle frequencies, R&D spending, non-R&D IPP investment and trademark assignments all increase following a cut in corporate taxes. This leads to a significant and sustained rise in patenting and aggregate productivity at long horizons, which eventually translates into higher output and consumption, in a world of ideas.

References

Abel, A. B. (2007): "Optimal Capital Income Taxation," Working Paper 13354, NBER.

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.

Auerbach, A. J. (1989): "Tax Reform and Adjustment Costs: The Impact on Investment and Market Value," *International Economic Review*, 30, 939–962.

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.

Bańbura, M., D. Giannone, and L. Reichlin (2010): "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, 25, 71–92.

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.

Bloom, N., M. Schankerman, and J. van Reenen (2013): *Econometrica*, 81, 1347–1393.

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*.

Carriero, A., J. Chan, T. E. Clark, and M. Marcellino (2022): "Corrigendum to "Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors" [J. Econometrics 212 (1) (2019) 137–154]," *Journal of Econometrics*, 227, 506–512.

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

Chan, J. C. C. (2020): "Large Bayesian VARs: A Flexible Kronecker Error Covariance Structure," *Journal of Business & Economic Statistics*, 38, 68–79.

Chiavari, A. and S. Goraya (2023): "The Rise of Intangible Capital and the Macroeconomic Implications," Research working paper, University of Oxford and Stockholm School of Economics.

Chiu, C.-W. J., H. Mumtaz, and G. Pintér (2017): "Forecasting with VAR models: Fat tails and stochastic volatility," *International Journal of Forecasting*, 33, 1124–1143.

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," Tech. rep., UC, Davis and London School of Economics.

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., 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&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," Tech. rep., mimeographed, London School of Economics.

Fieldhouse, A. J. and K. Mertens (2023): "The Returns to Government R&D: Evidence from U.S. Appropriations Shocks," Working Papers 2305, Federal Reserve Bank of Dallas.

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.

Geweke, J. (1993): "Bayesian treatment of the independent student-t linear model," *Journal of Applied Econometrics*, 8, S19–S40.

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

Griliches, Z. (1990): "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 28, 1661–1707.

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*.

Hall, R. E. and D. W. Jorgenson (1967): "Tax Policy and Investment Behavior," *The American Economic Review*, 57, 391–414.

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

House, C. L. and M. D. Shapiro (2008): "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation," *American Economic Review*, 98, 737–68.

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).

Joint Committee on Taxation (2017): "Macroeconomic Analysis of the Conference Agreement for H.R. 1, the "Tax Cuts and Jobs Act" (JCX-69-17)," Technical report, Joint Committee on Taxation of the U.S. Congress.

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

——— (2022): "Taxing Top Incomes in a World of Ideas," *Journal of Political Economy*, 130, 2227–2274.

Jones, C. I. and J. C. Williams (1998): "Measuring the Social Return to R&D," *The Quarterly Journal of Economics*, 113, 1119–1135.

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, NBER.

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, NBER.

JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2010): "Investment shocks and business cycles," *Journal of Monetary Economics*, 57, 132–145.

KOGAN, L., D. PAPANIKOLAOU, A. SERU, AND N. STOFFMAN (2017): "Technological Innovation, Resource Allocation, and Growth*," *The Quarterly Journal of Economics*, 132, 665–712.

KOOP, G. (2003): *Bayesian Econometrics*, Wiley.

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.

LEWIS, D. J. AND K. MERTENS (2022): "A Robust Test for Weak Instruments with Multiple Endogenous Regressors," Working Papers 2208, Federal Reserve Bank of Dallas.

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.

LUSOMPA, A. (2021): "Local Projections, Autocorrelation, and Efficiency," Research Working Paper RWP 21-01, Federal Reserve Bank of Kansas City.

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

MERTENS, K. AND M. O. RAVN (2011): "Understanding the aggregate effects of anticipated and unanticipated tax policy shocks," *Review of Economic Dynamics*, 14, 27–54, special issue: Sources of Business Cycles.

——— (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*.

MONTIEL-OLEA, J. L. AND C. PFLUEGER (2013): "A Robust Test for Weak Instruments," *Journal of Business and Economic Statistics*, 31, 358–369.

MONTIEL OLEA, J. L. AND M. PLAGBORG-MØLLER (2021): "Local Projection Inference is Simpler and More Robust Than You Think," *Econometrica*, 89, 1789–1823.

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

Newey, W. K. and K. D. West (1987): "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Nguyen, A. D. M., L. Onnis, and R. Rossi (2021): "The Macroeconomic Effects of Income and Consumption Tax Changes," *American Economic Journal: Economic Policy*, 13, 439–66.

Peters, R. H. and L. A. Taylor (2017): "Intangible capital and the investment-q relation," *Journal of Financial Economics*, 123, 251–272.

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*.

Plagborg-Møller, M. and C. K. Wolf (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89, 955–980.

Queraltó, A. (2022): "Monetary Policy in a Model of Growth," Tech. rep., Board of Governors, Federal Reserve.

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.

Winberry, T. (2021): "Lumpy Investment, Business Cycles, and Stimulus Policy," *American Economic Review*, 111, 364–96.

Online Appendix

"The Dynamic Effects of Income Tax Changes in a World of Ideas"

by James Cloyne (UC Davis), Joseba Martinez (LBS), Haroon Mumtaz (QMUL) and Paolo Surico (LBS)

A Data Appendix

A.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 above 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.

A.2 Sectoral Data

Gross output by industry is obtained from the Bureau of Economic Analysis (BEA) and is provided at annual frequency from 1947 to 1997 (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 below.

A.3 Data and definitions for Figure 4

Stock of Patents. This is calculated using the perpetual inventory method by adding the count of (eventually granted) USPTO patent filings and using 8% depreciation (Li and Hall (2020)). The model plot is the IRF of Z_t. Sources: Google Patents Public Data, Li and Hall (2020).

Market Price of Ideas We compute firm-level patent stock values using patent values from the extended Kogan et al. (2017) database and the perpetual inventory method with 8% depreciation (Li and Hall (2020)). We then: (i) sort firms by their patent stock value, (ii) form two portfolios consisting of the top and bottom deciles of the patent stock value distribution, and (iii) calculate the capitalization-weighted average price of these portfolios. The empirical IRF is that of the log difference of the prices of the top minus bottom decile portfolios. The model response is the log difference between the value-weighted prices of a portfolio that holds all adopted and unadopted ideas minus a portfolio that holds all capital. Sources: Center for Research in Security Prices (CRSP), Kogan et al. (2017), Li and Hall (2020).

Stock Market. Dow Jones Industrial Average data from WRDS. The model response is the aggregate value of assets (ideas and the capital stock). Source: WRDS.

Trade in Ideas. Count of trademark transactions from USPTO Trademark Transactions Database. Source: USPTO.

IPP Investment (excluding R&D). Intellectual property products investment (excluding R&D) from the national accounts. Source: BEA.

Intangible Stock. Organizational capital stock from Peters and Taylor (2017). Source: Peters and Taylor (2017).

B The Labour Market Response

Figure B.1: Response of labour Productivity, Wages and Employment to Income Tax Changes

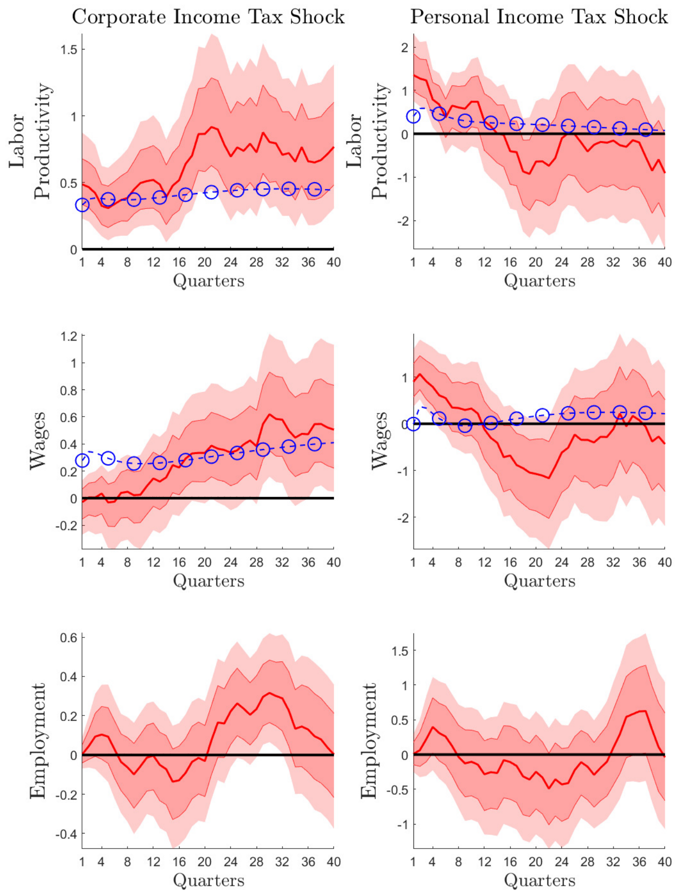

Notes: responses of labour productivity, wages and employment to a 1% cut in the average rate of corporate income taxes (left column) and of personal income taxes (right column). Red shadow bands represent central posterior 68^{th} and 90^{th} credible sets. Blue lines with circles represent the impulse responses of the model in Section 4 evaluated at the posterior median of estimated model parameters. Because the model does not have an extensive employment margin, no model response is plotted for employment.

C 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.

C.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)

$$Y_t = B_1 Y_{t-1} + B_2 Y_{t-2} + ... + B_{20} Y_{t-20} + A_0 E_t, E_t \sim N(0,1) \tag{26}$$

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: (i) A VAR(4) and (ii) 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.

C.2 Results

The top panel of Figure C.1 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, a distinctive feature of this response is that the main effect occurs 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. The bottom panel of Figure C.1 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.

D Estimation of the Bayesian Local Projection

D.1 Benchmark model

The model used to produce the benchmark results is defined as:

$$Z_{t+h} = \beta^h X_t + u_{t+h}, \qquad var(u_{t+h}) = \Omega_h \tag{27}$$

where $X_t = (1, Z_{t-1}, .., Z_{t-p})$ collects all the regressors and $\beta^h = \left(c^h, B_1^h, b_1^h, .., b_p^h\right)$ is the coefficient matrix. For $h = 0$, the model is a Bayesian VAR and estimation is standard (see for e.g. Bańbura et al. (2010)). When $h > 0$, we allow for non-normal disturbances. The covari-

Figure C.1: Monte-Carlo results

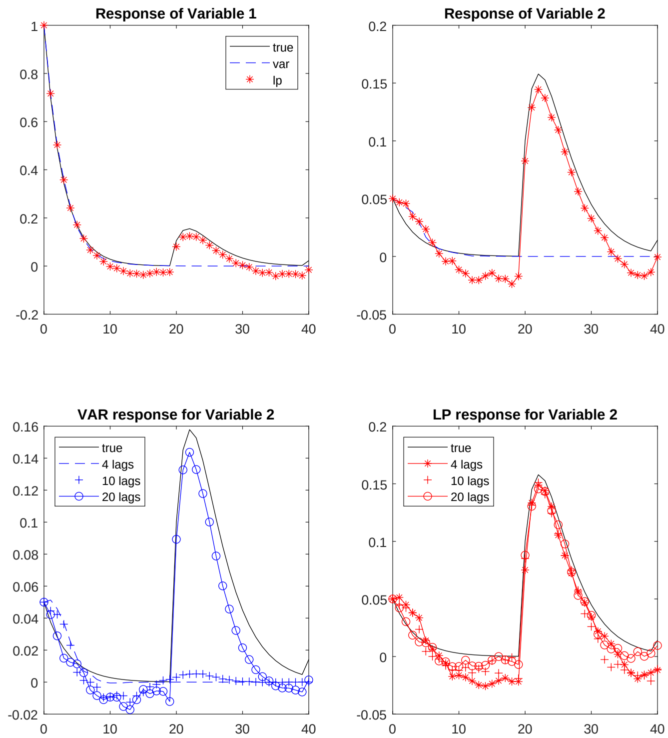

Notes: Monte-Carlo estimates of impulse responses of the two variables in Y to the first shock. In the bottom panel, the experiment is repeated for different lag lengths

ance matrix Ω_h is decomposed as $\Omega_h = A^{-1}H_tA^{-1\prime}$ where A is a lower triangular matrix while $H_t = diag\left(\frac{\sigma_1^2}{\lambda_{1t}}, \frac{\sigma_2^2}{\lambda_{2t}}, \ldots, \frac{\sigma_M^2}{\lambda_{Mt}}\right)$. Note that $\frac{1}{\lambda_{it}}$ for $i = 1, \ldots, M$ denotes the time-varying volatility of the orthogonal disturbances $e_{t+h} = Au_{t+h}$ Geweke (1993) shows that assuming a Gamma prior for λ_{it} of the form $P(\lambda_i) = \prod_{t=1}^{T} P(\lambda_{it}) = \prod_{t=1}^{T} \Gamma(1, \nu_i)$ leads to scale mixture of normal distributions for the orthogonal residuals ($\Gamma(a, b)$ denotes a Gamma distribution with mean a and degrees of freedom b). As shown in Geweke (1993), this is equivalent to assuming that each orthogonal residual e_{it} follows a Student's T-distribution with degrees of freedom equal to ν_i. This setup is used for VAR models in Chiu et al. (2017).

D.1.1 Priors

We employ the following prior distributions:

- We set a hierarchical prior for λ_{it} and ν_i (see Koop (2003)):

$$P(\lambda_{it}) = \Gamma(1, \nu_i) \tag{28}$$

$$P(\nu i) = \Gamma(\nu_0, 2) \tag{29}$$

  Note that the prior for ν is an exponential distribution, which is equivalent to a Gamma distribution with 2 degrees of freedom. We set $\nu_0 = 10$ which gives prior weight to the possibility of fat tails in the distribution of e_{it}

- The prior for σ_i^2 is inverse Gamma : $IG(T_0, D_0)$. We assume a flat prior setting the scale and degrees of freedom to 0.

- The free elements of each row of A have an independent prior of the form: $P(A_k) \sim N(a_{k,0}, s_{k,0})$ where A_k is the k_{th} row of this matrix. We set the mean of the prior to zero and the diagonal elements of $s_{k,0}$ to 1000

- We set a Minnesota type prior for the coefficients $\tilde{\beta^h} = vec(\beta^h)$: $P(\tilde{\beta^h}) \sim N(\beta_0, S_0)$. The mean β_0 implies that each variable in Z_{t+h} follows an $AR(1)$ process. The diagonal elements of the variance matrix S_0 corresponding to own lags are defined as $\frac{\mu_1^2}{p^2}$ and as $\frac{s_i}{s_j}\frac{\mu_1^2}{p^2}$ for coefficients on lags of other variables. Here p denotes the lag length while the ratio of variances $\frac{s_i}{s_j}$ accounts for differences in scale across variables. We set the tightness parameter μ_1 to 10 which implies a loose prior belief.

D.1.2 Gibbs Sampler

We use a Gibbs sampling algorithm to approximate the posterior distribution. The algorithm is based on the samplers presented in Geweke (1993), Koop (2003) and Chiu et al. (2017). In each iteration, the algorithm samples from the following conditional posterior distributions (Ξ denotes all other parameters):

- $G(\lambda_{it}|\Xi)$. Given a draw for A, the orthogonal residuals are constructed as $e_t = Au_t$. The conditional posterior distribution for $\lambda_{i,t}$ derived in Geweke (1993) applies to each column of e_t. As shown in Koop (2003) this posterior density is a gamma distribution with mean $(\nu_i + 1) / \frac{1}{\sigma_i} e_{i,t}^2 + \nu_i$ and degrees of freedom $\nu_i + 1$. Note that $e_{i,t}$ is the ith column of the matrix e_t.

- $G(\nu_i|\Xi)$. The conditional posterior distribution of ν_i is non-standard (see Koop (2003)) and given by:

$$G(\nu_i|\Xi) \propto \left(\frac{\nu_i}{2}\right)^{\frac{T\nu_i}{2}} \Gamma\left(\frac{\nu_i}{2}\right)^{-T} \exp\left(-\left(\frac{1}{\nu_0} + 0.5\sum_{t=1}^{T}\left[\ln\left(\lambda_{i,t}^{-1}\right) + \lambda_{i,t}\right]\right)\nu_i\right) \tag{30}$$

  As in Geweke (1993) we use the Random Walk Metropolis-Hastings Algorithm to draw from this conditional distribution. More specifically, for each of the M equations of the VAR, we draw $\nu_i^{new} = \nu_i^{old} + c^{1/2}\epsilon$ with $\epsilon \sim N(0,1)$. The draw is accepted with probability $\frac{G(\nu_i^{new}|\Xi)}{G(\nu_i^{old}|\Xi)}$ with c chosen to keep the acceptance rate around 40%.

- $G(A|\Xi)$: Given a draw for the coefficients β^h the model can be written as: $Au_{t+h} = e_{t+h}$ where $e_{i,t+h} \sim N(0, \frac{\sigma_i^2}{\lambda_{it}})$ for $i = 1, .., M$. This is a system of K linear regressions with known error variances. The first equation is an identity $u_{1,t+h} = e_{1,t+h}$. The second equation is: $u_{2,t+h} = -A_2 u_{1,t+h} + e_{2,t+h}$, the k$th$ equation is $u_{k,t+h} = -x_u A_k + e_{k,t+h}$ and so on, where $x_u = (u_{1,t+h}, \ldots, u_{k-1,t+h})$. By dividing both sides of the equations by the respective error standard deviation, i.e. $(\frac{\sigma_k^2}{\lambda_{kt}})^{(0.5)}$, the residual variance is normalised to 1. Given the normal prior for A_k, the conditional posterior is also normal with variance $v = \left(s_{0,k}^{-1} + \tilde{x_u}'\tilde{x_u}\right)^{-1}$ and mean $v\left(s_{0,k}^{-1}a_{0,k}+\tilde{x_u}'\tilde{u}_{k,t+h}\right)$ where $\tilde{x_u}$ and $\tilde{u}_{k,t+h}$ denote the regressors and the dependent variable after the GLS transformation described above.

- $G(\sigma_i^2|\Xi)$: The orthogonal residuals e_{t+h} can be transformed as follows: $\tilde{e_{t+h}} = e_{t+h}\lambda_{i,t}^{0.5}$. The conditional posterior for σ_i^2 is inverse Gamma with scale parameter $\tilde{e_{t+h}}'\tilde{e_{t+h}}+D_0$ and degrees of freedom $T+T_0$

- $G(\beta^h|\Xi)$ We use the algorithm of Carriero et al. (2022) to draw from this conditional posterior distribution. Carriero et al. (2022) show that the system can be re-written as:

$$AZ_{t+h} = A\beta^h X_t + e_{t+h}, \qquad e_{it,t+h} \sim N\left(0, \frac{\sigma_i^2}{\lambda_{i,t}}\right) \tag{31}$$

  Given the lower triangular structure of A, the coefficients of the jth equation can be sampled using blocks of the last $M-j+1$ equations, conditional on the remaining blocks. Carriero et al. (2022) show that these conditional posterior distributions are normal and they provide expressions for the mean and variance. This algorithm is substantially faster that drawing the coefficients of all equations in the model, jointly.

We employ 51000 iterations and drop the first 1000 as burn-in. We keep every $5th$ draws of the remainder for inference.

D.1.3 Lag augmentation and coverage

Following Montiel Olea and Plagborg-Møller (2021), we carry out a Monte-Carlo experiment to check the coverage properties of the error bands produced the Bayesian LP described above. We generate data from a 4-variable VAR(4) model. The coefficients and variance-covariance of the error terms is set equal to the OLS estimates of a VAR(4) model using data on 4 variables employed in our benchmark LP: (1) $ACITR$, (2) B^{CI}, (3) $ln(G)$ and (4) $ln(GDP)$. We generate 228 observations after discarding an initial sample of 100 observations to account for starting values. Using this artificial data, we estimate two Bayesian LPs: (1) a model with 4 lags of all 4 variables included as controls and (2) a model that is not lag-augmented and only the first lag that is required to generate the IRF is included. We employ 51000 Gibbs iterations and drop the first 1000 as burn-in. We keep every $5th$ draw of the remainder for inference. The experiment is repeated 1000 times and we compute coverage probabilities using the estimated 90 percent highest posterior density intervals. Figure D.1 Panel A shows that the benchmark model produces reasonably good coverage rates with distortions that remain below 10% even at long horizons. In contrast, when the lag augmentation is removed, the performance deteriorates substantially and coverage rates fall below 50% for all variables.

D.1.4 Convergence

To assess convergence of the Gibbs algorithm, we examine the inefficiency factors calculated using the impulse responses from the benchmark model. These estimates are below 20 for all variables and horizons (see Figure D.1 Panel B) providing support for convergence of the algorithm.

D.2 Bayesian LP with MA residuals

Our alternative specification directly models the autocorrelation in the residuals. In a recent paper Lusompa (2021) has shown that the u_{t+h} follows an $MA(h)$ process. We therefore consider the following extended model:

$$Z_{t+h} = \beta^h X_t + u_{t+h} \tag{32}$$

The residuals of each equation follow the MA process:

$$u_{t+h} = \epsilon_{t+h} + \theta_1 \epsilon_{t+h-1} + .. + \theta_q \epsilon_{t+h-q}, \qquad \epsilon_{t+h} \sim N(0, \Omega_h) \tag{33}$$

As noted in Chan (2020), this type of model can be re-written as:

$$Z_{t+h} = \beta^h X_t + \tilde{H}\epsilon_{t+h}, \qquad \epsilon_{t+h} \sim N(0, \Omega_h) \tag{34}$$

where $\tilde{H}$ is $T \times T$ banded matrix with ones on the main diagonal and the MA coefficients appearing below the main diagonal. For example, the process $u_{t+h} = \epsilon_{t+h} + \theta_1 \epsilon_{t+h-1}$ can be written as $u_{t+h} = \tilde{H}\epsilon_{t+h}$ where $\tilde{H} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ \theta_1 & 1 & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & \theta_1 & 1 \end{pmatrix}$

The model is estimated using a Gibbs sampling algorithm that is based on the methods described in Chan (2020).

D.2.1 Priors

We employ the following prior distributions:

- The prior for Ω is inverse Wishart: $IW(\Omega_0, T_0)$. We employ a flat prior and set both the scale matrix and degrees of freedom to 0.

- We set a Minnesota type prior for the coefficients $\tilde{\beta^h} = vec(\beta^h)$: $P(\tilde{\beta^h}) \sim N(\beta_0, S_0)$. The mean β_0 implies that each variable in Z_{t+h} follows an $AR(1)$ process. The diagonal elements of the variance matrix S_0 corresponding to own lags are defined as $\frac{\mu_1^2}{p^2}$ and as $\frac{\sigma_i}{\sigma_j}\frac{\mu_1^2}{p^2}$ for the coefficients on the lags of other variables. Here p denotes the lag length while the ratio of variances $\frac{\sigma_i}{\sigma_j}$ accounts for differences in scale across variables. We set the tightness parameter μ_1 to 10 which implies a loose prior belief.

- The prior for MA coefficients $\tilde{\Theta} = (\theta_1, \dots, \theta_q)$ is normal: $N(\Theta_0, V_0)$. The mean of the prior is set to 0. The variance is set using the Minnesota procedure (described above) with the coefficients on higher MA terms shrunk to 0 more quickly. We set the tightness parameter of the prior to 0.1

D.2.2 Gibbs Sampler

The Gibbs sampling algorithm for this model samples from the following conditional posterior distributions (Ξ denotes all other parameters):

- $G(\tilde{\beta^h}|\Xi)$: Given a draw for $\tilde{\Theta}$, the model can be written as

$$\tilde{Z}_{t+h} = \beta^h \tilde{X}_t + \epsilon_{t+h}, \qquad \epsilon_{t+h} \sim N(0, \Omega_h) \tag{35}$$
$$\tilde{Z}_{t+h} = \tilde{H}^{-1} Z_{t+h} \tag{36}$$
$$\tilde{X}_t = \tilde{H}^{-1} X_t \tag{37}$$

This is simply a system of linear equations with iid residuals. Let $\tilde{Z}$ and $\tilde{X}$ denote the matrices holding the transformed dependent and covariates, respectively. The conditional posterior is normally distributed with mean M and variance V:

$$V = \left(S_0^{-1} + \Omega_h^{-1} \otimes \tilde{X}'\tilde{X}\right)^{-1} \tag{38}$$
$$M = V\left(S_0^{-1}\beta_0 + \left(\Omega_h^{-1} \otimes \tilde{X}'\tilde{X}\right)\beta_{ols}\right) \tag{39}$$
$$\beta_{ols} = vec\left(\left(\tilde{X}'\tilde{X}\right)^{-1}\left(\tilde{X}'\tilde{Z}\right)\right) \tag{40}$$

- $G(\Omega_h|\Xi)$: Given a draw for β_h, the residuals ϵ_{t+h} can be easily calculated. The conditional posterior of Ω_h is inverse Wishart with scale matrix $\epsilon'_{t+h}\epsilon_{t+h} + \Omega_0$ and degrees of freedom $T + T_0$.

- $G(\tilde{\Theta}|\Xi)$: The model can be written in state-space form:

$$Z_{t+h} = \beta^h X_t + \begin{pmatrix} I_m & I_m \times \theta_1 & \ldots & I_m \times \theta_q \end{pmatrix} \begin{pmatrix} \epsilon_t \\ \epsilon_{t-1} \\ \vdots \\ \epsilon_{t-q} \end{pmatrix} \tag{41}$$

$$\begin{pmatrix} \epsilon_t \\ \epsilon_{t-1} \\ \vdots \\ \epsilon_{t-q} \end{pmatrix} = \begin{pmatrix} 0 & 0 & \ldots & 0 \\ 1 & 0 & \ldots & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \ldots & \ldots & \ldots \end{pmatrix} \begin{pmatrix} \epsilon_{t-1} \\ \epsilon_{t-2} \\ \vdots \\ \epsilon_{t-q-1} \end{pmatrix} + \begin{pmatrix} \epsilon_t \\ \\ \vdots \\ 0 \end{pmatrix} \tag{42}$$

$$var\left(\begin{pmatrix} \epsilon_t \\ 0 \\ \vdots \\ 0 \end{pmatrix}\right) = \begin{pmatrix} \Omega & 0 & \ldots & 0 \\ 0 & 0 & \ldots & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \ldots & \ldots & \ldots \end{pmatrix} \tag{43}$$

We use a random walk Metropolis-Hastings step to draw $\tilde{\Theta}$. We generate a candidate draw using $\tilde{\Theta}_{new} = \tilde{\Theta}_{old} + e, e \sim N(0, \tau)$. The draw is accepted with probability $\alpha = \frac{F\left(Z_{t+h}|\tilde{\Theta}_{new},\Xi\right) \times P\left(\tilde{\Theta}_{new}\right)}{F\left(Z_{t+h}|\tilde{\Theta}_{old},\Xi\right) \times P\left(\tilde{\Theta}_{old}\right)}$ where the likelihood function $F\left(Z_{t+h}|\tilde{\Theta}, \Xi\right)$ is calculated using the Kalman filter and the Normal prior $P\left(\tilde{\Theta}\right)$ is evaluated directly. We adjust the variance τ to ensure an acceptance rate between 20 and 40%.

We employ 51000 Gibbs iterations and drop the first 1000 as burn-in. We keep every $5th$ draw of the remainder for inference.

Figure D.1: Coverage probabilities and inefficiency factors

(a) Coverage Probabilities (b) Inefficiency Factors

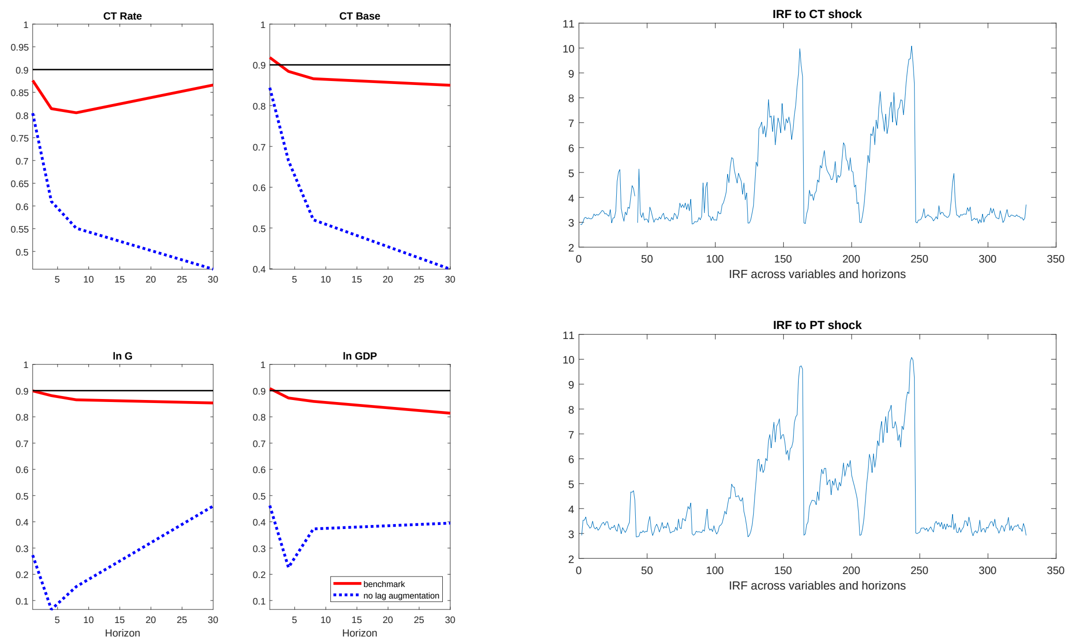

Notes: this figure shows the coverage probabilities for the Bayesian LP with and without lag augmentation (Panel A) and inefficiency factors calculated using the MCMC draws of the impulse responses from the benchmark model (Panel B).

E 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 E.1. 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 E.1 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 E.1: Sectoral Evidence: Average Effects and Results by R&D Share

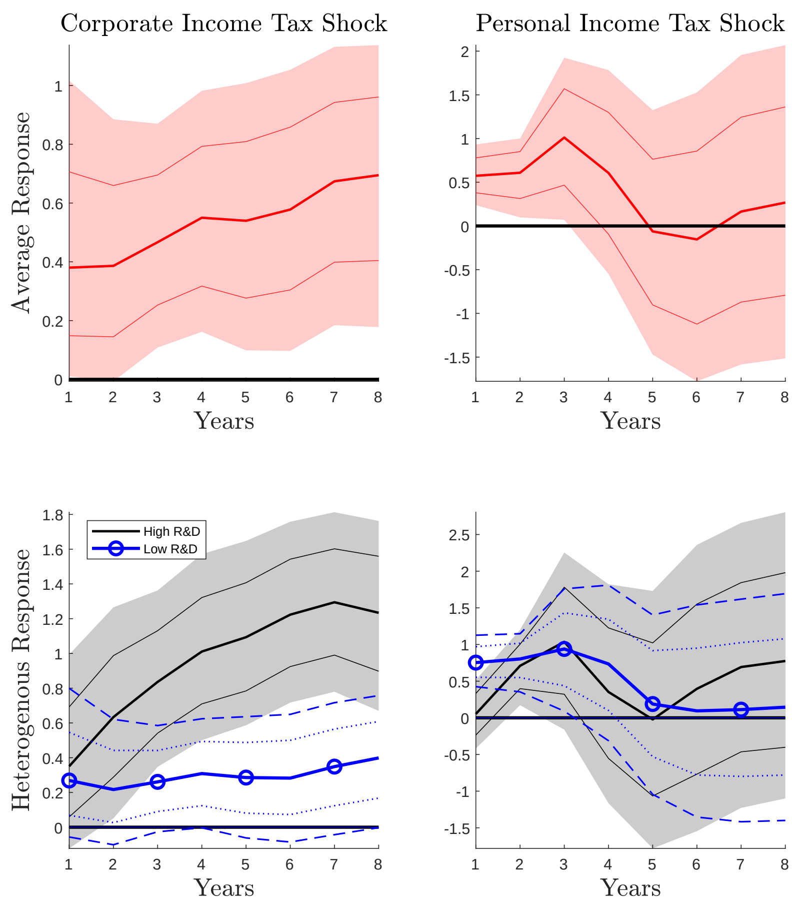

Notes: output response using sectoral data from the US BEA. The first row show the average effect. The bottom row further split sectors into high R&D intensive and low R&D intensive.

F 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 F.1 presents the estimated decomposition. The contribution of both tax shocks to the FEV of GDP is about 20% 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.

G Robustness

Bayesian LP with MA residuals. In Figure G.1, we use the Bayesian LPs described in section D.2. While the response of GDP, TFP and R&D to corporate shocks is more volatile than the benchmark, the results confirm that this shock has long-lasting effects on output and productivity. In contrast, the estimated impact of personal tax shocks is short-lived. In Figure G.2, we show that our main findings of very persistent effects of corporate tax changes on GDP and TFP are

Figure F.1: Forecast Error Variance decomposition

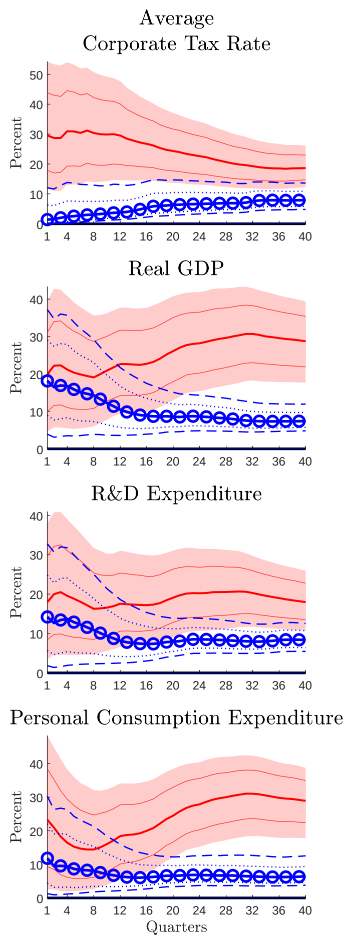

Notes: contribution of corporate and personal tax changes to the variance of each variable in the figure. The contribution of corporate tax changes are shown in the red lines (posterior median and 68 percent band) and the shaded area (90% band). The line with circles shows the contribution of the personal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

robust also to varying the number of lags, using optimal priors, adding the measure of government spending shocks proposed by Ramey (2011) and changing the ordering of the tax shocks.

Frequentist estimates of the Direct model and LPIV. In this part, we present two cases:

- **using narrative measures as instruments.** Figure G.3 presents impulse responses estimated using the frequentist approach discussed in Section 2.2. The figure shows estimates obtained using OLS and the smoothed version of LPIV (Barnichon and Brownlees (2019)). As noted in the text, these regressions using the narrative measures of Mertens and Ravn (2013) as regressors/instruments.

- **using VAR shocks as instruments.** In this exercise, we use the structural tax shocks estimated by the VAR of Mertens and Ravn (2013) as instruments. One advantage of this approach is that the VAR shocks are orthogonal by construction and each of them can be used to instrument the two tax rates separately. We proceed in the following steps:

Figure G.1: IRFs using Bayesian LP with MA residuals

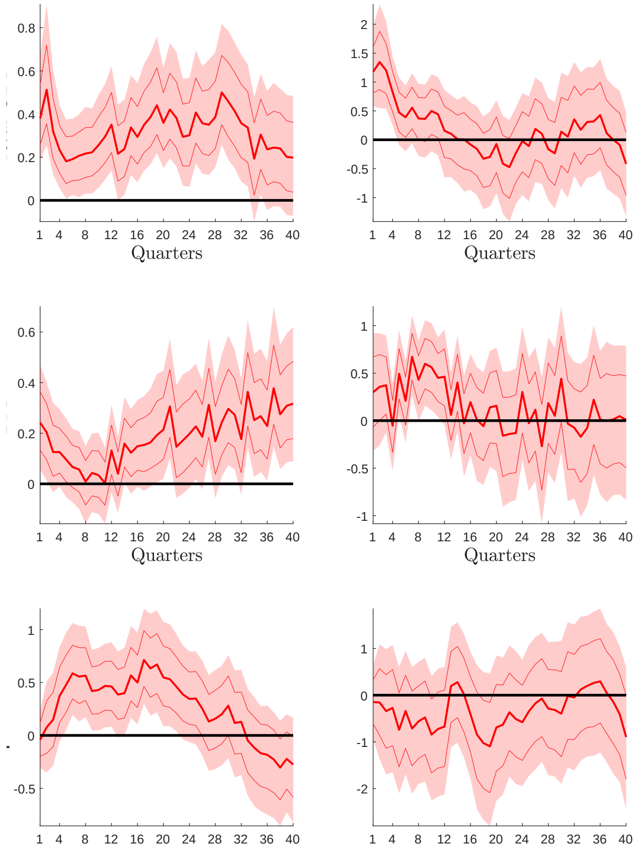

Notes: this figure shows impulse responses estimated using the Bayesian LP with residuals modelled as an MA process. The thin lines and shaded areas are the 68% and 90% error bands.

1. estimate the VAR of Mertens and Ravn (2013) and obtain the estimates of structural corporate and personal tax shocks (z_{ct} and z_{pt}, respectively). Note that these are orthogonal by construction.
2. we then estimate the following regression:

$$Z_{i,t+h} = c^{(h)} + B_1^{(h)} x_t + \sum_{j=1}^{L} b_j^{(h)} Z_{t-j} + u_{t+h}, \qquad u_{t+h} \sim N(0, \sigma_h) \tag{44}$$

   where x_t is the endogenous variable (i.e. either the corporate or the personal tax rate) which is instrumented by the appropriate shock obtained in step 1. The matrix Z denotes the 8 variables considered in the benchmark specification and L is set equal to 1.

The IRFs are given by $B_1^{(h)}$; error bands use HAC standard errors. Figure G.4 reveals that the LPIV estimates broadly support the benchmark results. We reach similar conclusions when we employ the smooth local projections (SLP) of Barnichon and Brownlees (2019).

- **Instrument Strength.** We use the robust test proposed by Lewis and Mertens (2022) to assess the strength of the VAR shocks as instruments in the LPs at each horizon. The test is carried out separately for the two orthogonal instruments using the LP in equation 44. The

Figure G.2: Response of real GDP, TFP and R&D: Different Specifications

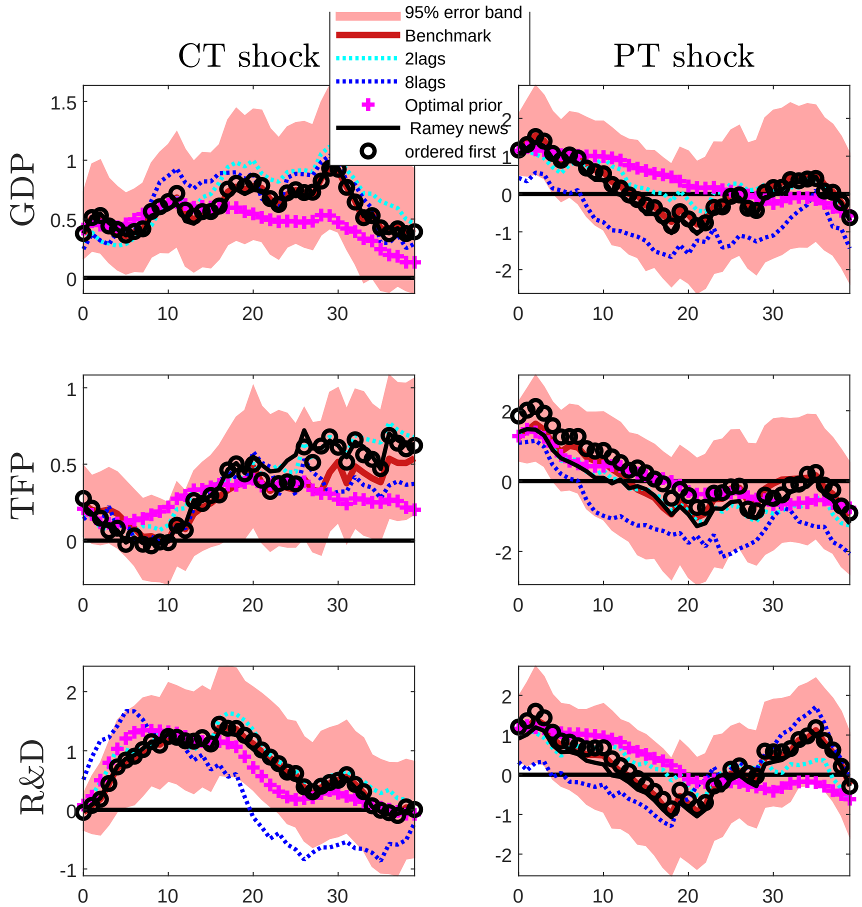

Notes: 90% bands for the baseline empirical real GDP result are shown in pink, together with the point estimates from various alternative specifications. These include: (i) changing number of lags used as control variables, (ii) adjusting the prior, (iii) including the Ramey (2011) defence news shock as a control (iv) changing the ordering of the tax shocks. See text for more discussion.

test is implemented using the Newey and West (1987) covariance matrix with the lag set to the impulse response horizon. Figure G.5 shows that the test statistic lies above the critical value for the CT shock at all horizons considered in the LP. For the PT shock, the test fails to reject the null for a few quarters but this instrument appears strong at the remaining horizons.

H Prior Predictive Analysis

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.[1] 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 Figure H.1, we report the distributions of the model impulse responses

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

Figure G.3: Responses of GDP, TFP and R&D expenditure to Corporate and Personal Tax Changes using alternative local projection models

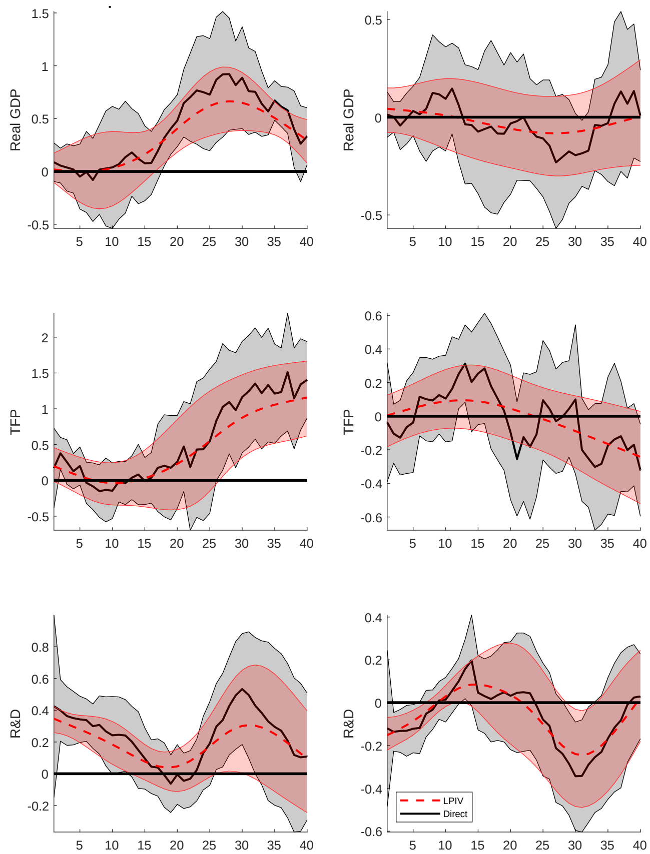

Notes: responses of the average tax rates, real GDP, and TFP to a 1% cut in the average rate of corporate income taxes (left column) and the average rate of personal income taxes (right column). Red and grey shadow bands represent 90 percent confidence intervals based on Newey and West (1987) standard errors.

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 economy in which the effects of both personal and corporate income taxes are quite short-lived and productivity is virtually a-cyclical. As shown in Section 5, however, the posterior distributions paint a quite different picture.

I Model Appendix

I.1 Production Sector and Endogenous Productivity

There exists a continuum of measure A_t of monopolistically competitive intermediate goods firms. Each of them manufactures a differentiated product: intermediate goods firm i produces output $Y_{i,t}$. The endogenous state variable A_t is the mass of intermediate goods adopted in production (equivalently, the stock of adopted technologies). As detailed in the text, A_t grows as a result of expenditures on applied research, which we call adoption. The final goods composite is the following

Figure G.4: LPIV estimates using Mertens and Ravn (2013) VAR shocks as instruments

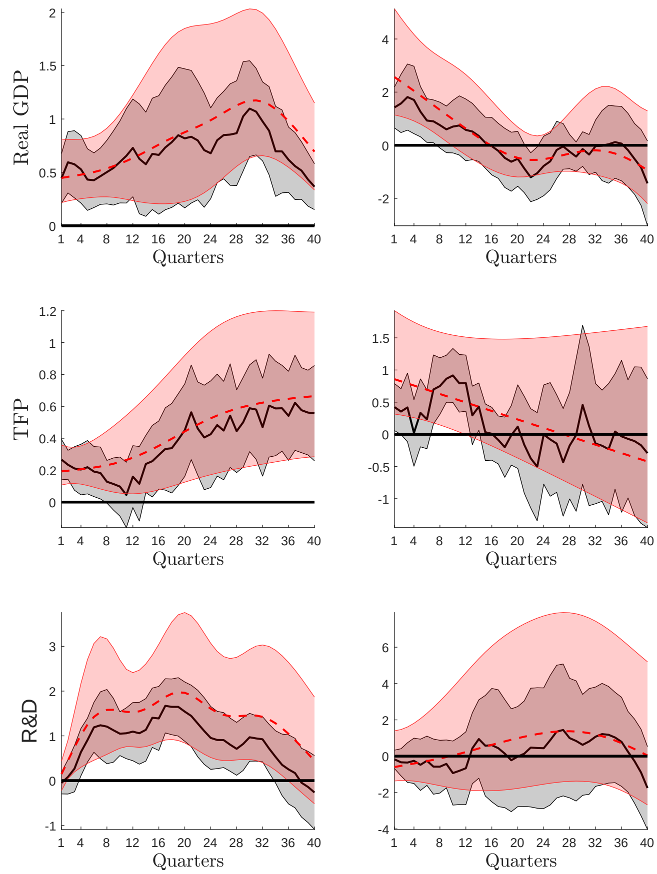

Notes: Impulse responses using IV estimates of local projections.The black lines show TSLS estimates, while the dotted red lines are smoothed local projections. The shaded areas are the the 90% error bands, respectively. These are constructed using the Newey and West (1987) HAC estimator for the variance of the coefficients. The bandwidth parameter is set to the horizon of the impulse response

CES aggregate of individual intermediate goods, with $\theta > 1$:

$$Y_t = \left(\int_0^{A_t} (Y_{i,t})^{\frac{1}{\theta}} \, di \right)^{\theta} \tag{45}$$

Let $K_{g,i,t}$ be the stock of capital that firm i uses, U_t denotes capital utilization (described below), and $L_{g,i,t}$ represents the stock of labour employed. Firm i produces output $Y_{i,t}$ according to the following Cobb-Douglas technology:

$$Y_{i,t} = (U_t K_{i,t})^{\alpha} (L_{i,t})^{1-\alpha} . \tag{46}$$

Given a symmetric equilibrium for intermediate goods, the aggregate production function is:

$$Y_t = A_t^{\theta-1} \cdot (U_t K_{g,t})^{\alpha} (L_{g,t})^{1-\alpha}. \tag{47}$$

$L_{g,t}$ and $K_{g,t}$ are aggregate capital and labour employed in the goods production sector.

Figure G.5: Test for instrument strength. Mertens and Ravn (2013) VAR shocks as instruments

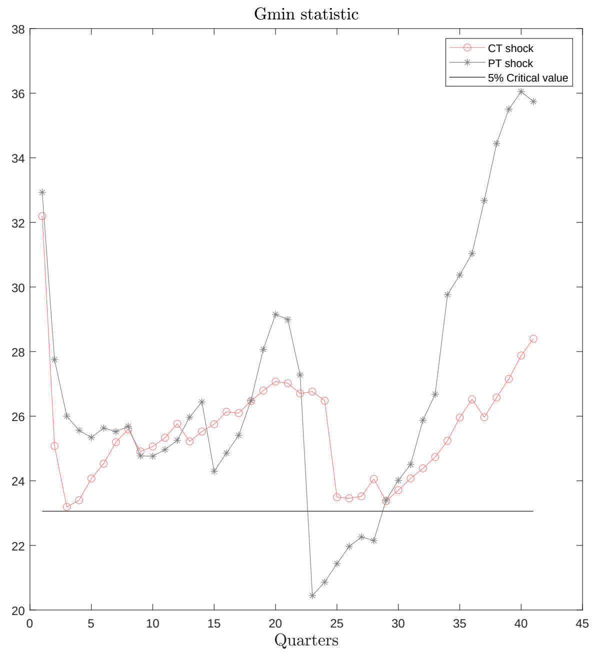

Notes: the figure shows the gmin statistic of Lewis and Mertens (2022) at each horizon. The solid line depicts the 5 percent critical value. The test is constructed using the Newey and West (1987) HAC estimator. The bandwidth parameter is set to the horizon of the impulse response.

I.2 Households and the Corporate Sector

The representative household consumes, supplies labour, saves 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 labour supply. There are three types of labour: goods production (g), R&D (z) and adoption labour (a). Households supply the three types of labour competitively but choose hours $H_{j,t+1}$ one period in advance, and face a quadratic adjustment cost when changing hours. Following the realization of uncertainty in period t, the household chooses effort, $e_{j,t}$, and we assume that the effective labour supply is given by $L_{j,t} = H_{j,t}e_{j,t}$. The household's maximization problem and budget constraint are:

$$\max_{C_t,S_{t+1},H_{j,t+1},e_{j,t}} \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \left\{ \log\left(\frac{C_t}{N_t} - b\frac{C_{t-1}}{N_{t-1}}\right) - \sum_{j\in\{g,a,z\}} \frac{1+\bar{e}_j e_{j,t}^{1+\chi_j}}{1+\chi_j} \frac{H_{j,t}}{N_t} \right\}, \tag{48}$$

and

$$C_t + P_{S,t}S_{t+1} = \sum_{j\in\{g,a,z\}} \left[(1-\tau_{p,t})\, w_{j,t}e_{j,t}H_{j,t} - \frac{\psi_j}{2}\left(\frac{H_{j,t+1}}{(1+g_n)\, H_{j,t}} - 1\right)^2 \Psi_t \right] + T_t + S_t\left(P_{S,t} + D_t\right), \tag{49}$$

where C_t is consumption, S_t are shares in the corporate sector (which trade at price $P_{S,t}$), D_t are dividends from the corporate sector, $w_{j,t}$ are real wages, and T_t are government transfers.[2] The

[2] Changes in dividend taxes are a small part of the personal income tax measure in the Mertens and Ravn (2013) data set. As a result, we abstract from explicitly modelling dividend taxes.

Figure H.1: Prior and Posterior Distributions of the response of the main variables in the model

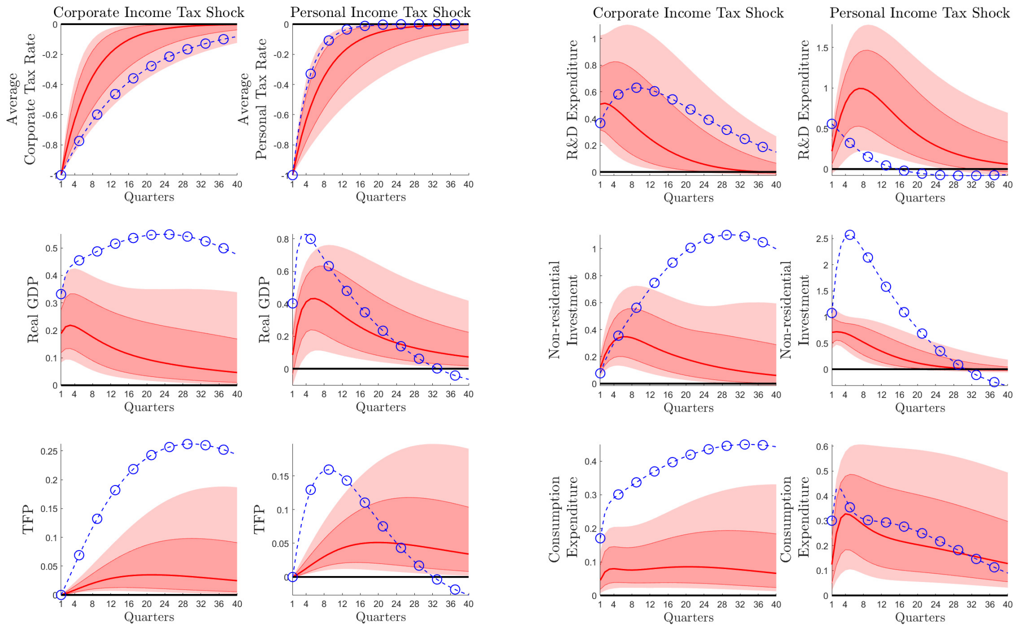

Notes: this figure shows the response of the average tax rates, real GDP, productivity, consumption, investment and R&D to a 1% cut in the average tax rate of corporate income taxes (left column) and the average tax rate of personal income taxes (right column). Red shadow bands and solid lines represent the 90^{th} and 68^{th} 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.

symbol Ψ_t denotes a scaling factor that grows at the same rate as aggregate output, required to ensure that labour 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 invests in the physical capital stock (with associated quadratic adjustment cost), rents capital to intermediate goods firms, finances innovation costs, and chooses the rate of capital utilization in the goods sector, U_t, with associated cost $\nu(U_t)K_{g,t}$, where $\nu(U)$ 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. Individual firms and innovators make the specific production, R&D and technological adoption decisions, as described earlier.

Dividends in period t are given by overall corporate sector income minus corporate taxes due:

$$D_t = CI_t - \tau_{c,t} CI_t^{TAX}, \tag{50}$$

where CI_t is net corporate income, which is GDP net of wages, investment and utilization:

$$CI_t = Y_t - \sum_{j \in \{g,a,z\}} \left[w_{j,t} L_{j,t} + I_{j,t} \left(1 + f_j \left(\frac{I_{j,t}}{(1+g_y) I_{j,t-1}} \right) \right) \right] - \nu(U_t) K_{g,t} \tag{51}$$

$\tau_{c,t}$ is the corporate income tax rate and CI^{TAX} is corporate income minus deductions for depreciation and amortisation. As with intellectual property assets (described above), we follow Auerbach (1989), Mertens and Ravn (2011) and Winberry (2021) in modelling depreciation allowances for the capital stock as a geometric process: in every period, a fraction $\hat{\delta}$ of investment can be deducted from taxable profits, with the remaining portion 1-$\hat{\delta}$ carried into the next period. Details of the derivation of amortisation allowances and taxable corporate income are in Appendix I.3.

Factor demands. Intermediate goods firm i chooses capital services $U_t K_{i,t}$, and labour $L_{i,t}$ 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:

$$\alpha \frac{MC_t Y_t}{U_t K_{g,t}} = r_{g,t}, \tag{52}$$

$$(1-\alpha) \frac{MC_t Y_t}{L_{g,t}} = w_{g,t}, \tag{53}$$

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, Anzoategui et al., 2019).

Investment good producers. There are three types of capital goods in the economy, used in the goods-producing, R&D and adoption sectors. Competitive producers use final output to produce these goods which they sell to the investment fund, which in turn rents capital to firms. Following Christiano et al. (2005), we assume flow adjustment costs of investment for the three types of capital goods. The adjustment cost functions (for $j \in \{g, z, a\}$)are $f_j\left(\frac{I_{j,t}}{(1+g_y)I_{j,t-1}}\right)$, with each function increasing and concave, with $f_x(1) = f_x^{'}(1) = 0$ and $f_x^{''}(1) > 0$; and $I_{j,t}$ is new capital of type i produced in period t. The first-order conditions are:

$$\begin{aligned} Q_{j,t} = 1 + f_j\left(\frac{I_{j,t}}{(1+g_y)\, I_{j,t-1}}\right) + \frac{I_{j,t}}{(1+g_y)\, I_{j,t-1}} f_j^{\prime}\left(\frac{I_{j,t}}{(1+g_y)\, I_{j,t-1}}\right) \\ -\beta \mathbb{E}_t \Lambda_{t,t+1}\left(\frac{1-\tau_{c,t+1}}{1-\tau_{c,t}}\right)(1+g_y)\left(\frac{I_{j,t+1}}{(1+g_y)\, I_{j,t}}\right)^2 f_j^{\prime}\left(\frac{I_{j,t+1}}{(1+g_y)\, I_{j,t}}\right), \end{aligned} \tag{54}$$

where $Q_{j,t}$ is the price of type j capital.

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 $\hat{\pi}_t = \kappa_p \widehat{mc}_t + \beta \mathbb{E}_t[\hat{\pi}_{t+1}]$ with slope $\kappa_p = \frac{(1-\xi_p\beta)(1-\xi_p)}{\xi_p}$.

Fiscal Policy. The government's budget constraint is given by:

$$\bar{G}(1+g_y)^t - T_t = \tau_{p,t}\left(\sum_{j\in\{g,a,z\}} w_{j,t} L_{j,t}\right) + \tau_{c,t} CI_t^{TAX}, \tag{55}$$

For simplicity, the government finances consumption using personal and corporate income taxes; lump sum taxes adjust to balance the budget every period. The process of tax rates $\tau_{c,t}$ and $\tau_{p,t}$

$$\log\left(\tau_t^x\right) = (1-\rho_{\tau x})\, \bar{\tau}^x + \rho_{\tau x} \log\left(\tau_{t-1}^x\right) + \varepsilon_t^{\tau x}, \tag{56}$$

follows an AR(1) process in logs for $x \in \{c, p\}$, with $\rho_{\tau x} \in (0,1)$, and $\varepsilon_t^{\tau x} \sim N(0,1)$ is i.i.d..

Monetary Policy. The nominal interest rate $R_{n,t+1}$ is set according to a Taylor rule $R_{n,t+1} = \left(\left(\frac{\pi_t}{\bar{\pi}}\right)^{\phi_\pi}\left(\frac{L_t}{\bar{L}}\right)^{\phi_y} R_n\right)^{1-\rho^R} (R_{n,t})^{\rho^R}$ where R_n is the steady state nominal rate, $\bar{\pi}$ the target rate of inflation, L_t total effective labour supply and $\bar{L}$ steady-state labour supply; ϕ_π and ϕ_y are the feedback coefficients on, respectively, the inflation gap and the capacity utilization gap, measured as in Anzoategui et al. (2019).

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

$$Y_t = C_t + \sum_{j\in\{g,a,z\}} \left[\left(1 + f_j\left(\frac{I_{j,t}}{(1+g_y) I_{j,t-1}}\right)\right) I_{j,t} + \frac{\psi_j}{2}\left(\frac{H_{j,t+1}}{(1+g_n) H_{j,t}} - 1\right)^2 \Psi_t\right] + \nu(U_t) K_t + \bar{G}(1+g_y)^t . \tag{57}$$

I.3 Derivation of Taxable Corporate Income

Taxable corporate income is corporate income minus amortisation and depreciation allowances for capital and intellectual property assets. To derive this, we start by defining corporate income:

$$\begin{aligned} CI_t = {} & \underbrace{Y_t - w_{g,t}L_{g,t} - r_{g,t}K_{g,t} - P_{a,t}\Delta A_t}_{\text{Goods-producing firms}} + \underbrace{\sum_{j\in\{g,a,z\}} (r_{j,t}K_{j,t} - Q_{j,t}I_{j,t}) - \nu(U_t) K_t}_{\text{Investment firm}} \\ & + \underbrace{\sum_{j\in\{g,a,z\}} \left(Q_{j,t}I_{j,t} - I_{j,t}\left(1 + f_j\left(\frac{I_{j,t}}{(1+g_y) I_{j,t-1}}\right)\right)\right)}_{\text{Capital-producing firms}} + \underbrace{P_{z,t}\Delta Z_t - w_{z,t}L_{z,j,t} - r_{z,t}K_{z,j,t}}_{\text{R\&D firms}} \\ & + \underbrace{P_{a,t}\Delta A_t - w_{a,t}L_{a,t} - r_{a,t}K_{a,t} - P_{z,t}\Delta Z_t}_{\text{Adoption firms}}, \end{aligned} \tag{58}$$

where $\Delta A_t \equiv A_{t+1} - \phi A_t$ and $\Delta Z_t \equiv Z_{t+1} - \phi Z_t$ are the measures of newly adopted and discovered technologies, respectively, such that the terms in red are the aggregate entry costs in the goods-producing and adoption sectors (which are equal to the aggregate revenues of the adoption and R&D sectors). Netting out terms, corporate income is given by:

$$CI_t = Y_t - \sum_{j\in\{g,a,z\}} \left[w_{j,t}L_{j,t} - I_{j,t}\left(1 + f_j\left(\frac{I_{j,t}}{(1+g_y) I_{j,t-1}}\right)\right)\right] - \nu(U_t) K_{g,t}, \tag{59}$$

which is real output minus wages and the cost of investment in each of the goods-producing, adoption and R&D sectors, and utilization cost in the goods-producing sector. Consistent with the US tax code, in the model firms deduct depreciation and amortisation from taxable profits to arrive at taxable income. We model these allowances as a geometric process in which a fraction $\hat{\delta}$ of the value of investments can be deducted from profits each period. Denoting amortisation allowances by Ξ, the laws of motion for aggregate allowances in capital and intellectual property products are given respectively by:

$$\Xi_{IP,t+1} = \left(1 - \hat{\delta}_{IP}\right)\left(\Xi_{IP,t} + P_{Z,t}\Delta Z_t + P_{A,t}\Delta A_t\right) \tag{60}$$

$$\Xi_{K,t+1} = \left(1 - \hat{\delta}_{K}\right)\left(\Xi_{K,t} + \sum_{j\in\{g,a,z\}} Q_{j,t} I_{j,t}\right) \tag{61}$$

Depreciation allowances at $t+1$ are $1-\hat{\delta}_{\bullet}$ times depreciation allowances at t plus the value of new investments in the three types of capital and the two types of intellectual property products. Using this notation, taxable corporate income is:

$$CI_t^{TAX} = CI_t + \nu\left(U_t\right) K_{g,t} - \hat{\delta}_K\left(\Xi_{K,t} + \sum_{j\in\{g,a,z\}} Q_{j,t} I_{j,t}\right) - \hat{\delta}_{IP}\left(\Xi_{IP,t} + P_{Z,t}\Delta Z_t + P_{A,t}\Delta A_t\right) \tag{62}$$

To arrive at taxable corporate income, we add back a non-deductible expense (capital utilization) and subtract the depreciation allowances that reduce the corporate sector tax liabilities.

J Social Returns to R&D

Following Jones and Williams (1998), the production function for new unadopted ideas is given by a function G of research efforts and the stock of unadopted ideas:

$$Z_{t+1} - \phi Z_t = G\left(X_{z,t}, Z_t\right) = Z_t^{1+\zeta} X_{z,t}^{\rho_z}$$

The increase in ideas associated with a marginal change in research effort is

$$\nabla Z_{t+1} = \left(\frac{\partial G}{\partial X_z}\right)_t,$$

where ∇ denotes the change relative to the balanced growth path. Note that $X_{z,t}$ is in units of the R&D good, which is produced using R&D labour and capital. Denoting by $P_{Xz,t}$ the price of this composite good, 1 unit of consumption yields $P_{Xz,t}^{-1}$ units of the R&D good. Since we are computing the return in terms of consumption, the relative prices of R&D and adoption will be used in the calculation.

To determine how much consumption is gained in time $t+1$ from the reduction in R&D that returns Z to its balanced growth path, note that $Z_{t+2} - \phi Z_{t+1} = G\left(X_{z,t+1}, Z_{t+1}\right)$ and that the deviation of Z from its balanced growth path is given by:

$$\nabla Z_{t+2} = \nabla Z_{t+1} + \left(\frac{\partial G}{\partial X_z}\right)_{t+1} \nabla X_{z,t+1} + \left(\frac{\partial G}{\partial Z}\right)_{t+1} \nabla Z_{t+1}$$

where the terms are, respectively: the deviation in Z occasioned by the increase in research effort; the reduction in Z from a cut in research effort; and the change in research efficiency as a result of additional ideas. The gain in consumption from returning Z to its balanced growth path is found by setting $\nabla Z_{t+2} = 0$:

$$\nabla X_{z,t+1} = -\frac{\left(\frac{\partial G}{\partial X_z}\right)_t}{\left(\frac{\partial G}{\partial X_z}\right)_{t+1}} \left(\left(\frac{\partial G}{\partial Z}\right)_{t+1} + 1\right).$$

Following the same logic, the change in adopted ideas at $t+1$, which determines the change in output, is given by $A_{t+2} - \phi A_{t+1} = \phi \lambda_t \left(Z_{t+1} - A_{t+1}\right)$.

Note that, because there is a one period delay between when ideas are discovered and when adopters can start working to adopt them, the initial change in R&D affects the stock of adopted ideas, and therefore output, at time $t+2$. Defining ∇A_{t+2} as the deviation in adopted ideas from the balanced growth path,

$$\nabla A_{t+2} = \nabla Z_{t+1} \frac{\partial A_{t+2}}{\partial Z_{t+1}} = \nabla Z_{t+1} \left(\phi \left(\left(Z_{t+1} - A_{t+1}\right) \left(\frac{\partial \lambda}{\partial Z}\right)_{t+1} + \lambda\right)\right).$$

The change in ideas has two components: (i) an increase in Z_t increases adoption efficiency, so any idea is more likely to be adopted; (ii) $\lambda \nabla Z_{t+1}$ extra ideas are adopted. At $t+2$, the contribution to output of these additional ideas is given by $\nabla Y_{t+2} = \left(\frac{\partial Y}{\partial A}\right)_{t+2} \nabla A_{t+2}$ Furthermore, at $t+2$, the deviation in the stock of adopted ideas is given by:

$$\nabla A_{t+3} = \nabla A_{t+2} + \left(\frac{\partial \lambda}{\partial X_a}\right)_{t+2} \nabla X_{a,t+2},$$

and as with R&D, we compute the reduction in adoption expenditure that returns the economy to the balanced growth path by solving for the value of $\nabla X_{a,t+2}$ such that $\nabla A_{t+3} = 0$:

$$\nabla X_{a,t+2} = -\frac{\nabla A_{t+2}}{\left(\frac{\partial \lambda}{\partial X_a}\right)_{t+2}}$$

Grouping all terms, the social return to R&D is given by

$$1 + \tilde{r}_{RD} = \beta^2 \left(\frac{\partial Y}{\partial A}\right)_{t+2} \frac{\nabla A_{t+2}}{P_{Xz,t}} + \beta^2 \frac{P_{Xa,t+2}}{P_{Xz,t}} \frac{\nabla A_{t+2}}{\left(\frac{\partial \lambda}{\partial X_a}\right)_{t+2}} + \beta \frac{P_{Xz,t+1}}{P_{Xz,t}} \frac{\left(\frac{\partial G}{\partial X_z}\right)_t}{\left(\frac{\partial G}{\partial X_z}\right)_{t+1}} \left(\left(\frac{\partial G}{\partial Z}\right)_{t+1} + 1\right),$$

where β and β^2 terms appear because the gains happen at different times, and each of the terms is scaled by the relative price of the R&D goods at t and $t+1$ or the adoption good at $t+2$, which converts all terms into units of consumption in the given time period, relative to price of R&D goods at time t. Defining the social cost of a new idea in units of consumption as $\tilde{P}_{Z,t} = \left(\frac{\nabla Z_{t+1}}{P_{Xz,t}}\right)^{-1}$, and $g_{\tilde{P}Z}$ as the growth rate of the social cost, and $\tilde{d}_t = \frac{\beta}{\tilde{P}_{Z,t}} \left(\beta \frac{\partial A_{t+2}}{\partial Z_{t+1}} \left(\left(\frac{\partial Y}{\partial A}\right)_{t+2} + \frac{P_{Xa,t+2}}{\left(\frac{\partial \lambda}{\partial X_a}\right)_{t+2}}\right) + \tilde{P}_{Z,t+1} \left(\frac{\partial G}{\partial Z}\right)_{t+1}\right)$, we obtain the expression in the main text,

$$\tilde{r}_{RD} = \frac{\tilde{d}_t}{\tilde{P}_{Z,t}} + g_{\tilde{P}Z}. \tag{63}$$

The "social dividend" of R&D has three components: the increase in output, the decrease in adoption expenditures, and the change in the efficiency of R&D. The social return on R&D is a function of (i) model parameters, namely g_n (the population growth rate) and β (the discount factor), plus the parameters of the endogenous growth block of the model (θ, ρ_z, ρ_λ, ϕ, ζ, $\bar{\lambda}$; see Tables 1 and 2 for definitions); and (ii) the expenditure shares of R&D and adoption, which are in turn also functions of parameters, including the tax parameters ($\hat{\tau}_c$, $\hat{\delta}_{IP}$, and to a lesser extent $\hat{\tau}_p$ and $\hat{\delta}_K$).

K Decomposing the Effects of Personal Income Taxes

Figure K.1 is the personal income tax counterpart of Figure K.1, and so also decompose the GDP response and its TFP contribution into channels of transmission. The output decomposition as well as the counterfactual analyses with no applied research and no endogenous growth makes it clear that variable labour utilisation is the most important driver of the GDP and TFP responses to the personal income tax shock at short horizons. The mechanism behind the effects of personal income taxes is very different from the one behind corporate taxes, as the former works primarily on the household side. The reduction in personal income taxes encourages workers (applied researchers) to supply more labour but, due to employment adjustment costs, the rise in hours worked (endogenous TFP) is gradual. On the other hand, labour utilisation responds contemporaneously via increased efforts, implying that labour productivity increases on impact. This channel, however, has no *direct* influence on firms' incentives to innovate, and therefore lasts as long as the distortionary personal income tax rate is below its steady state level.

Figure K.1: GDP Decomposition and Counterfactual Analysis

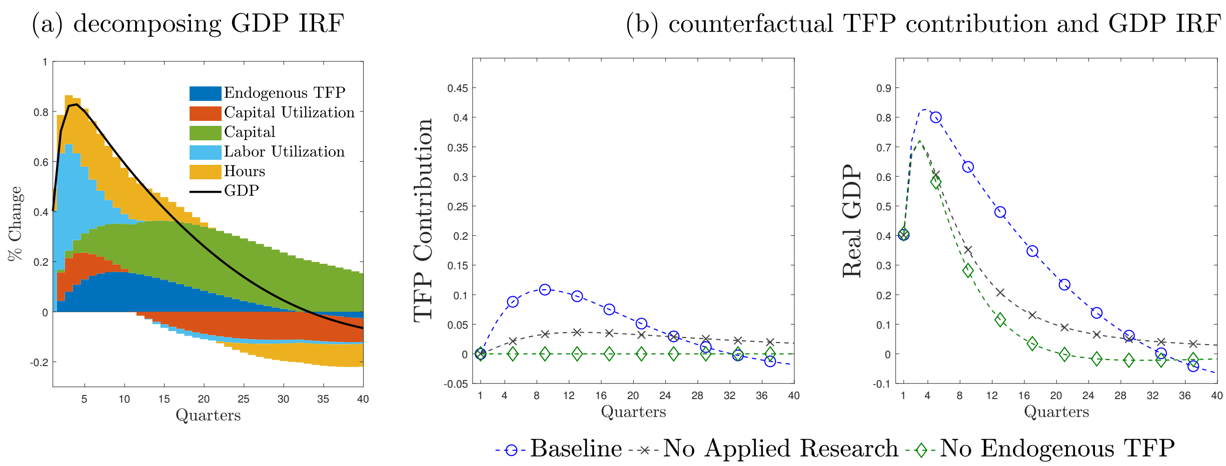

L Estimates of the model with no endogenous growth

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 L.1 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 L.1: 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]

M Tax amortisation of Intangible Assets around the World

In Table M.1, we report the legal tax amortisation periods (in years) for the main types of intangible assets in selected countries as of 2016. The tax treatment of intangible assets varies widely across jurisdictions. Interestingly, advanced economies typically have longer tax amortisation periods on intellectual property products than developing countries.

Table M.1: An International Perspective on Tax amortisation of Intangible Assets

Country	**Patents**	**Technology**	**Trademark**
Australia	20	5	no TAB
Austria	RUL	RUL	15
Canada	20	20	20
China	RUL (≥10)	RUL (≥10)	RUL (≥10)
France	RUL (≥5)	RUL (≤5)	no TAB
Germany	15	RUL (≤3)	25
Greece	RUL (≤5)	5	20
Hungary	RUL	5	10
India	10	RUL (≤3)	10
Indonesia	10	10	10
Ireland	20	15	no TAB
Italy	RUL (≥3)	5	RUL (≤5)
Japan	RUL	RUL	20
Malaysia	20	20	no TAB
Mexico	20	20	10
Netherlands	20	5	no TAB
New Zealand	15	5	no TAB
Norway	20	RUL	10
Poland	RUL	5	5
Portugal	RUL (≤5)	RUL (≤5)	no TAB
Romania	20	5	15
Russia	15	15	10
Singapore	5	5	no TAB
Slovakia	15	4	no TAB
Slovenia	RUL	5	no TAB
South Africa	20	5	no TAB
Spain	RUL (≥5)	RUL (≤5)	RUL (≤5)
Sweden	10	10	10
Switzerland	20	5	no TAB
Taiwan	20	15	no TAB
Thailand	20	5	20
Turkey	20	10	no TAB
UK	25	20	20
USA	15	15	15
Vietnam	10	10	10

Notes: RUL: for Remaining Useful Lifetime; TAB: tax amortisation benefit. The table reports the legal tax amortisation life time in years of the main types of intangible assets across selected countries in 2016. Source: http://www.taxamortisation.com/tax-amortisation-benefit.html