Innovation and the Productivity Growth Slowdown*

Patrick Moran† — University of Oxford

Albert Queralto‡ — Federal Reserve Board

April 16, 2017

Abstract

We use vector-autoregressive (VAR) methods to investigate the effects of R&D on total factor productivity (TFP) in the U.S. and in a sample of advanced economies, and find that movements in R&D have a significant and delayed effect on TFP. We then augment a New Keynesian DSGE model to allow for endogenous TFP growth via innovation and technology adoption, and discipline its key parameters using the VAR evidence. We use the model to explore the drivers of TFP growth in recent times, including the role of the Great Recession, and to draw some implications for monetary policy.

Keywords: Endogenous Technology; Business Cycles; Monetary Policy.

JEL classification: E32; F41; F44; G15.

---

*The views expressed in this paper are solely the responsibility of the authors, and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

†Department of Economics, Manor Road Building, Oxford OX1 3UQ, UK. Email: patrick.donnellymoran@economics.ox.ac.uk

‡Federal Reserve Board, 20th St and Constitution Ave NW, Washington, DC, 20551, USA. Email: albert.queralto@frb.gov.

1 Introduction

Growth in total factor productivity (TFP) has slowed dramatically in recent years: as seen in Figure 1, TFP growth in the U.S. has rarely, if ever, been as low for as long as in the post-2007 period. Weak productivity growth has been widespread across advanced economies (Figure 2). This development has caused concern for policymakers, and at the same time has sparked an intense debate on its possible causes, with the role of innovation and business dynamism the subject of increasing attention.[1] Relatedly, a growing literature in macroeconomics analyzes the link between innovation and productivity dynamics, and its implication for aggregate fluctuations, within the context of modern quantitative frameworks, following the lead of Comin and Gertler (2006).[2]

Figure 1: U.S. TFP Growth

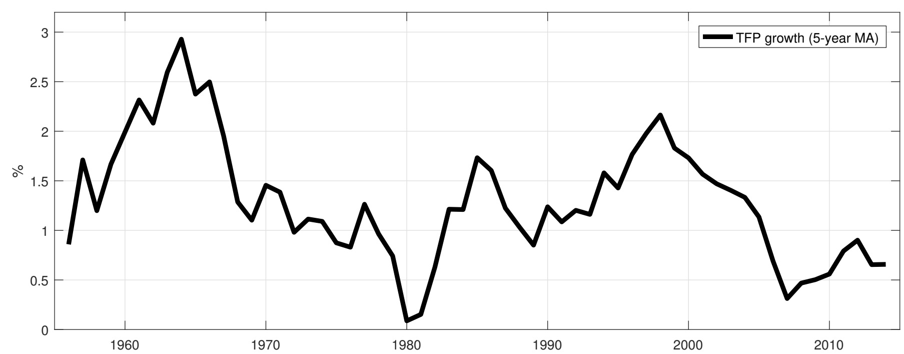

Note: 5-year moving average (two-sided) of U.S. TFP growth.

Motivated by these observations, in this paper we investigate the role of innovation in driving productivity growth, both empirically and theoretically. Fist, we use vector autoregression (VAR) methods to investigate systematically the hypothesis that movements in innovation drive medium-run TFP developments. Figure 3 provides motivation for our empirical investigation: observe that medium-term fluctuations in business-sector R&D expenditure tend to precede fluctuations in TFP, suggesting a causal link between the two variables. Second, we develop a macroeconomic model featuring endogenous technology innovation and adoption, as in Comin and Gertler (2006). Here our goal is twofold: first, we explore the extent to

[1] For example, Yellen (2016) emphasizes that "[...] understanding whether, and by how much, productivity growth will pick up is a crucial part of the economic outlook" and notes that "there is some evidence that the deep recession had a long-lasting effect in depressing investment, research and development spending, and the start-up of new firms, and that these factors have, in turn, lowered productivity growth."

[2] See, for example, Anzoategui, Comin, Gertler and Martinez (2016), Benigno and Fornaro (2016), Bianchi and Kung (2014), Guerron-Quintana and Jinnai (2014), Kung and Schmid (2015), or Queralto (2013).

which the model can account for the VAR evidence on the dynamic effects of R&D on TFP. Second, we use the model to explore the drivers of the slowdown in TFP growth, and to draw implications for monetary policy.

Figure 2: TFP Growth across Advanced Economies

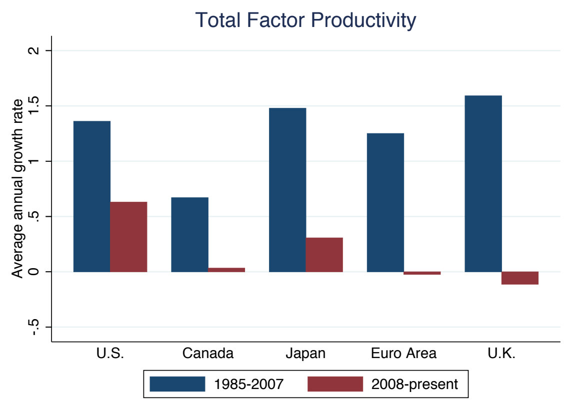

The VAR analysis suggests a significant, though delayed, effect of movements in business-sector R&D on TFP. This result holds for the U.S. as well as for a panel of advanced economies. We also find that in more R&D-intensive countries—with higher ratios of private R&D spending to GDP—the effects of R&D on TFP tend to be stronger. The impact of R&D on TFP is quantitatively large: in our preferred specification, a rise in R&D by 4 percent in the first year induces an increase in TFP of almost half a percentage point, with the peak effect occurring after about seven years. There is also some evidence of "spillovers" (Coe and Helpman (1995)) from U.S. R&D to TFP in other advanced economies, although the effects appear to materialize only at very long horizons. Interestingly, we also find that stock prices tend to immediately jump in response to the R&D shocks we identify—a result reminiscent of the findings in the "news shocks" literature (e.g. Beaudry and Portier (2006)) that high stock prices tend to be associated with future TFP increases.

Our approach relies on using the dynamic effects of R&D on TFP from the VAR to estimate some of the key model parameters, by minimizing the distance between model and empirical impulse responses to an R&D shock. An advantage of this approach is that it allows identifying key parameters by using more direct evidence on the dynamic link between R&D and TFP, rather than relying on external sources as frequently done in the literature. In particular, our approach allows identifiying a key elasticity in the model, governing the impact of R&D expenditure on the creation of new technologies. Overall, we find that the model does a reasonably good job of capturing the VAR evidence.

Figure 3: U.S. Business-Sector R&D and TFP, Medium-Term Cycle

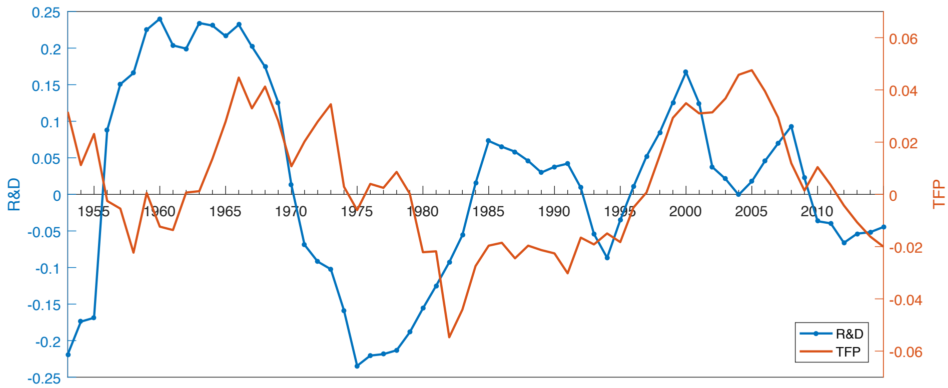

Note: Both series have been detrended using a band-pass filter that isolates frequencies between 2 and 50 years.

We next use the model to explore the drivers of the productivity slowdown of recent times. An advantage here is that our model arguably captures well the causal link between R&D and TFP, given that its parameters have been disciplined by evidence on precisely that link. Our findings suggest that the endogenous growth channel accounts for a substantial part of the productivity slowdown—about forty percent, on average, between 2001 and 2014. We also consider the question of how much of the productivity slowdown is due to the Great Recession, relative to factors that predate it—see, for example, Fernald (2014) and also the discussion in Anzoategui et al. (2016). We find that the sharp decline in R&D during the crisis likely contributed significantly to the subsequent low TFP growth, particularly after 2010. Finally, we also explore to what extent monetary policy can stimulate TFP growth going forward. The model suggests that it can, although the effect is likely to be transitory and modest in size.

The rest of the paper is organized as follows. Section 2 describes our empirical analysis. Section 3 describes the model. Section 4 describes the estimation of the model's key parameters. Section 5 presents the historical analysis of productivity. Section 6 concludes.

2 Evidence

In this section, we explore the hypothesis that business-sector innovation drives medium-run developments in productivity. Our basic approach consists in identifying shocks in private R&D expenditure, and then tracing out their dynamic effect on TFP.

We perform the analysis within different settings. We first explore a small-scale VAR for the U.S., consisting of R&D expenditure, TFP, and real GDP (section 2.1). We then analyze the same system in a panel of advanced economies (section 2.2). We next explore whether there are "spillovers" from U.S. innovation to foreign countries' TFP (section 2.3). Finally, we turn to a larger-scale VAR for the U.S., which includes a set of standard macroeconomic indicators in addition to the aforementioned variables (section 2.4). Appendix A contains details on the data.

2.1 U.S.

We begin with a small-scale empirical model for the U.S. Our reduced-form empirical specification is a first-order VAR:

$$\begin{bmatrix} y_t^{us} \\ tfp_t^{us} \\ rd_t^{us} \end{bmatrix} = c^{us} + B^{us} \begin{bmatrix} y_{t-1}^{us} \\ tfp_{t-1}^{us} \\ rd_{t-1}^{us} \end{bmatrix} + u_t^{us} \tag{1}$$

Here y_t^{us}, tfp_t^{us}, and rd_t^{us} represent, respectively, real GDP, TFP, and real business-sector R&D expenditure. All variables are in logs. The frequency is annual, and observations start in 1953. We estimate the above system by least squares. The coefficients to be estimated include a vector of constants, c^{us}, a matrix of autoregressive coefficients, B^{us}, and the variance-covariance matrix of the reduced-form residuals u_t. We include all three variables in levels, given the likely presence of cointegrating relationships among them.[3]

To identify structural shocks to R&D, we rely on a lower-triangular Choleski factorization of the variance-covariance matrix of the reduced-form residuals. Given the variable ordering in (1), this identification scheme imposes the restriction that TFP does not respond contemporaneously to structural innovations in R&D. We believe this assumption is natural: it captures the idea that it takes time for R&D expenditure (an input into the innovation process) to result in new technologies that become implemented and used in production. Macroeconomic models featuring technological innovation and adoption, like Comin and Gertler (2006) and variants of it (including the one we develop below), generally satisfy this restriction. We also believe it is important to allow both TFP and R&D to respond to shocks to GDP, which is accomplished by placing GDP first in the VAR. This allows us to control for business-cycle effects which might induce comovement between R&D and TFP if, for example, the short-run behavior of the latter partly reflects mismeasurement. However, our main results on the effects of R&D shocks on TFP do not change significantly when we instead place GDP third,

[3]The same approach is followed, for example, by Christiano et al. (2005).

thus allowing R&D to impact GDP contemporaneously.

Figure 4: Identified R&D Shock in the U.S.

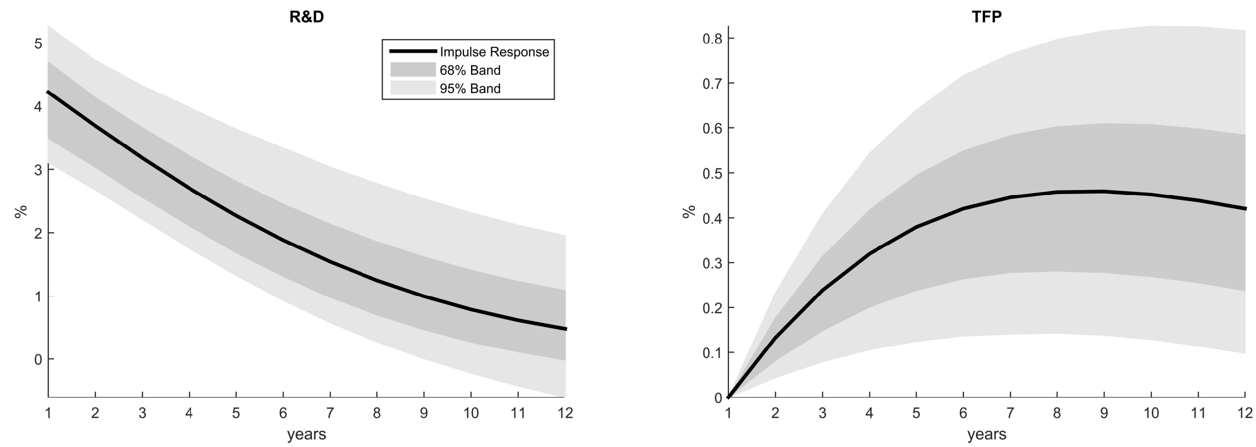

Note: Response to a one-standard-deviation identified shock to R&D expenditure obtained from (1). The black line represents the dynamic response, and the dark-grey and light-grey shaded areas respectively depict 68% and 95% confidence intervals obtained by bootstrapping with 10,000 repetitions.

Figure 4 shows the dynamic effects of a one-standard-deviation identified shock to R&D expenditure: R&D rises by about 4 percent on impact, and then gradually declines. The shock impacts TFP significantly, albeit with a delay: at its peak—which occurs after about seven years—the response of the level of TFP reaches nearly 0.5%, with half of the full effect materializing about three years after the initial shock. Further, the TFP increase induced by the shock is highly persistent, and its level stays high long after R&D has returned to baseline.

The natural interpretation of these results is that a rise in R&D for reasons unrelated to current TFP (or to the state of the economy, as captured by real GDP) accelerates the development of technological innovations which, after some time, become implemented in production and eventually improve firms' productivity. There is, however, the possibility of reverse causality: if firms can foresee the future rise in TFP, they could respond by increasing R&D expenditure, perhaps because they believe that the new technologies resulting from that expenditure will now be more profitable. To the extent that such "news" effects are not fully controlled for by real GDP, they would imply that the causal interpretation offered above might be incorrect: instead of R&D causing TFP, we could just be capturing the response of R&D to an exogenous future rise in TFP, currently anticipated by firms when making their R&D decisions.

One way to test for this possibility is to repeat the analysis reported above, but using aggregate investment in place of R&D.[4] The idea is that the anticipation of high future R&D would lead to a rise in overall investment, and not just in R&D—so that if a rise in investment similarly "leads" to a future increase in TFP, we might really be capturing the effect of news about high future TFP. Accordingly, we next examine a VAR exactly analogous to (1), but using real aggregate investment in place of R&D expenditure. As seen in Figure 5, a rise in investment is not followed by an increase in TFP—the latter actually *falls* a bit following the increase in investment, although the decline is not statistically significant (at the 95% level). Thus, we conclude that the results shown in Figure 4 likely reflect a causal effect from R&D to subsequent TFP developments, rather than an effect of anticipated future TFP on current R&D.

Figure 5: Identified Shock to Investment in the U.S.

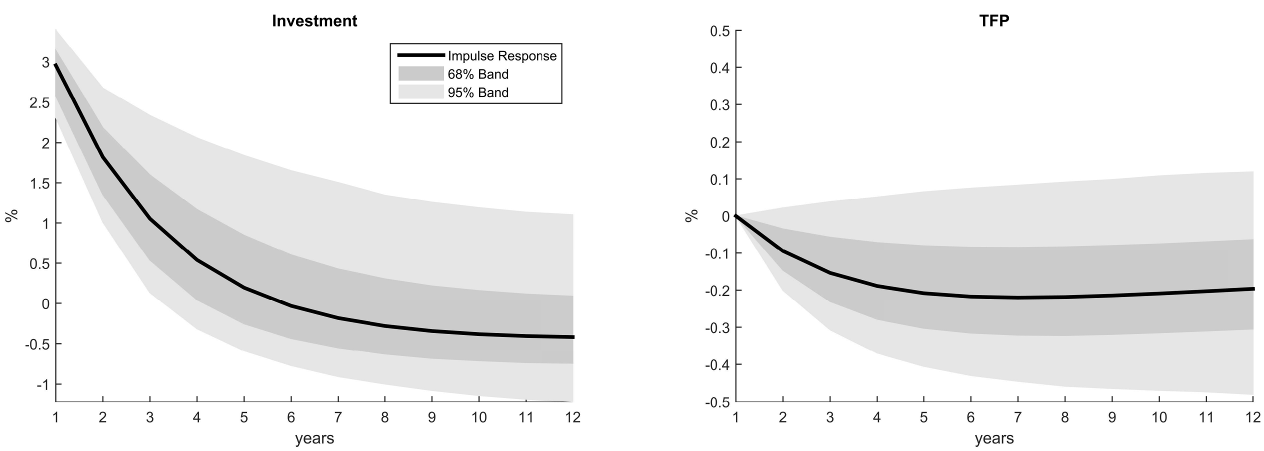

Note: Response to an identified shock to investment, obtained by estimating a system analogous to (1) with investment in place of R&D. The black line represents the dynamic response, and the dark-grey and light-grey shaded areas respectively depict 68% and 95% confidence intervals obtained by bootstrapping with 10,000 repetitions.

2.2 Panel of advanced economies

We next explore whether the effects of R&D on TFP identified in the U.S. hold more generally in a sample of advanced economies (AEs henceforth). The data consists of a panel of 21 AEs (not including the U.S.) in the post-1980 period (see Appendix A for details on the data). Data on business-sector R&D expenditure is from the OECD. We select the sample of countries based on the availability of business-sector R&D data. We specify the following empirical model, analogous to (1):

[4]We thank Andrew Atkeson for suggesting this check to us.

$$\begin{bmatrix} y_{i,t} \\ tfp_{i,t} \\ rd_{i,t} \end{bmatrix} = c_i + B \begin{bmatrix} y_{i,t-1} \\ tfp_{i,t-1} \\ rd_{i,t-1} \end{bmatrix} + u_{i,t} \tag{2}$$

We estimate the system above by least squares. The system contains a vector of country fixed effects, c_i, thus allowing estimation of the country-specific intercept term for each country in the sample. The model, however, imposes the matrix B as well as the variance-covariance matrix of the residuals $u_{i,t}$ to be common across countries. This so-called least-square dummy variable (LSDV) or fixed-effects estimator is commonly used in panel VAR settings with relatively long time series of macroeconomic data—for example, Uribe and Yue (2006), Akinci (2013) or Cerra and Saxena (2008).[5] We follow the same Choleski approach as above to identify R&D shocks.

Figure 6: Identified R&D Shock in a Panel of AEs

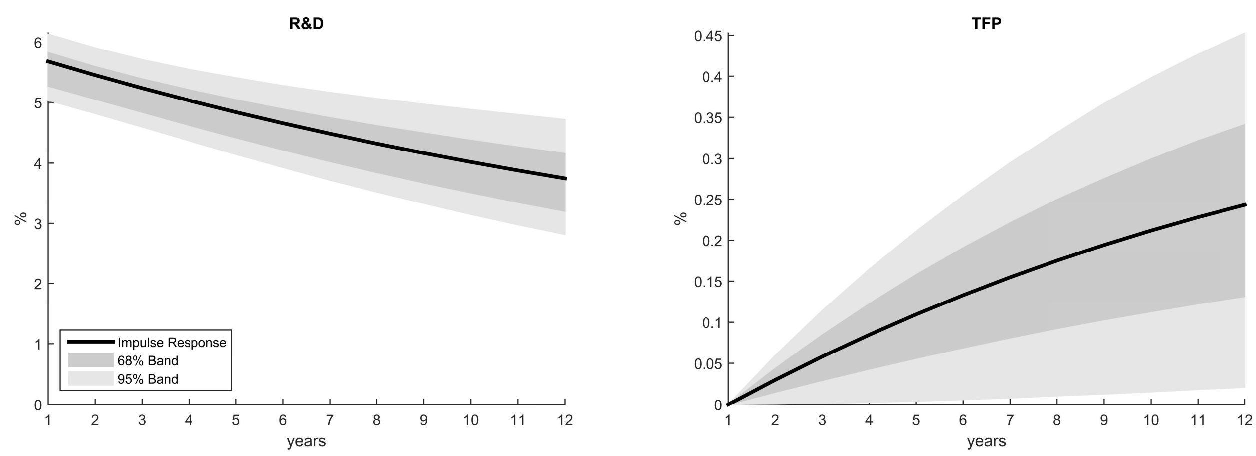

Note: Response to an identified shock to R&D expenditure obtained from estimating 2 on the full sample of AEs. The black line represents the dynamic response, and the dark-grey and light-grey shaded areas respectively depict 68% and 95% confidence intervals obtained by bootstrapping with 10,000 repetitions.

Figure 6 shows the impulse responses to an R&D shock for our full panel of AEs. As in the case of the U.S., a rise in R&D induces a gradual, persistent rise in TFP. The TFP increase is statistically significant at the 95% level. There are, however, some notable differences with the U.S. First and foremost, the impact on TFP of a rise in R&D of a given size appears to be notably weaker in the foreign economies: R&D rises about 5.75% on impact in the AEs

[5] As shown by Nickell (1981), the LSDV estimator is biased due to correlation between the country fixed effects and the lagged dependent variables. This bias, however, is likely to be small in settings like the one above where the time-series dimension is relatively large.

(more than in the U.S.), but the overall effect on the level of TFP is about 0.25%, half that in the U.S. Further, the peak effect is reached much later: TFP continues to rise by year 12, and starts to settle shortly after that (not shown). By contrast, in the U.S. the TFP response levels out after about seven years. Finally, note that R&D itself also rises much more persistently in the AEs: by year 12 it is still 3.75% above baseline, while in the U.S. it has returned to baseline by that time. This strengthens our conclusion that R&D is less powerful in affecting TFP in the AEs, and may also help explain why the TFP rise is more gradual in the AEs than in the U.S., where TFP levels out much sooner.

We have found that there is significant heterogeneity across the countries in our sample on the effects on TFP of an R&D shock. In particular, the effects on TFP tend to be stronger in countries with higher ratios of private R&D to GDP. To illustrate this point, Figure 7 repeats the same analysis as above, but this time estimating (2) on the top-5 research economies in our sample as measured by their average R&D-GDP ratios, which turn out to be Germany, Japan, South Korea, Sweden and Switzerland. We focus precisely on the top 5 for illustration and because they are widely recognized as highly innovative countries, but conclusions hold more generally when we look at reasonable variations in the set of countries, so long as they include countries that are high in the ranking by R&D-to-GDP. As seen in the Figure, TFP now rises much more than in the full sample of AEs: the peak effect is about 0.9%, much larger than for the full sample, and in fact stronger than for the U.S. in terms of peak TFP response per size of initial rise in R&D. That said, the R&D movement continues to be much more persistent in this sample of foreign economies than in the U.S.

2.3 Spillovers from U.S. R&D to foreign TFP

A natural question to ask when analyzing TFP developments across countries is whether there are cross-country R&D spillovers, i.e. if R&D expenditure in one given country may benefit productivity in other countries. Coe and Helpman (1995) and Eaton and Kortum (1996), for example, find evidence in support of such spillovers. The VAR methodology employed above can complement the existing studies—which typically focus on longer-run relationships—as it allows a richer study of the dynamic interaction between R&D and productivity. To this end, we next specify a VAR which allows for spillovers from U.S. variables to "local" (i.e., foreign-economy) variables. We restrict attention to spillovers from the technological leader—namely, the U.S. Accordingly, we estimate the following model:

Figure 7: Identified R&D Shock in a Panel of AEs, Top 5 Countries by R&D-GDP Ratio

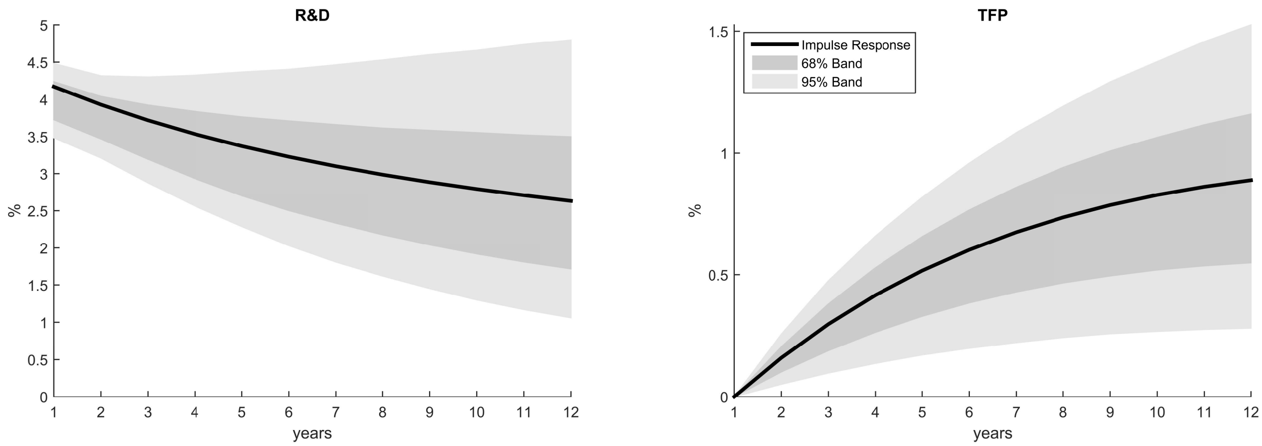

Note: Response to an identified shock to R&D expenditure obtained from estimating (2) on the top-5 economies by business-sector R&D expenditure to GDP (Germany, Japan, South Korea, Sweden and Switzerland). The black line represents the dynamic response, and the dark-grey and light-grey shaded areas respectively depict 68% and 95% confidence intervals obtained by bootstrapping with 10,000 repetitions.

$$\begin{bmatrix} y_t^{us} \\ tfp_t^{us} \\ rd_t^{us} \\ y_{i,t} \\ tfp_{i,t} \\ rd_{i,t} \end{bmatrix} = \begin{bmatrix} c^{us} \\ \tilde{c}^i \end{bmatrix} + \begin{bmatrix} B^{us} & \mathbf{0} \\ D & \tilde{B} \end{bmatrix} \begin{bmatrix} y_{t-1}^{us} \\ tfp_{t-1}^{us} \\ rd_{t-1}^{us} \\ y_{i,t-1} \\ tfp_{i,t-1} \\ rd_{i,t-1} \end{bmatrix} + \begin{bmatrix} u_t^{us} \\ \tilde{u}_{i,t} \end{bmatrix} \tag{3}$$

Above, subindex i denotes the country (other than the U.S.). We include a set of constants for the U.S., contained in c^{us}, as well as country-specific fixed effects ($\tilde{c}^i$). We suppose that the local variables cannot impact U.S. variables, neither contemporaneously nor with a lag, and accordingly set the upper-right quadrant of the autoregressive matrix above to 0. Thus, the process for the U.S. variables is unaffected by local variables. Since we have longer time series for the U.S., we estimate B^{us} separately in a first step. The matrix D captures the impact of lagged U.S. variables on local variables. We continue to identify U.S. R&D shocks by performing a Choleski decomposition of the variance-covariance matrix of $\left[u_t^{us\prime}\, \tilde{u}_{i,t}^{\prime}\right]$. Our focus is on whether local TFP responds to U.S. R&D shocks.

Although we do not find much evidence of spillovers when estimating (3) on the full sample, we do find some evidence for the panel of high research intensity countries studied earlier, as shown in Figure 8: local TFP eventually rises, by almost 0.6% at its peak.[6] The effect, however, takes a long time to materialize: local TFP does not move for the first few years, and then rises only very gradually. For comparison, the Figure also includes the response of TFP to own-R&D shocks (the dashed blue line). Note that in that case, the response of TFP is much faster (as well as larger in magnitude). Note also that the effects are estimated with significant uncertainty, as indicated by the width of the confidence bands.

Figure 8: Spillovers from U.S. R&D to Foreign TFP

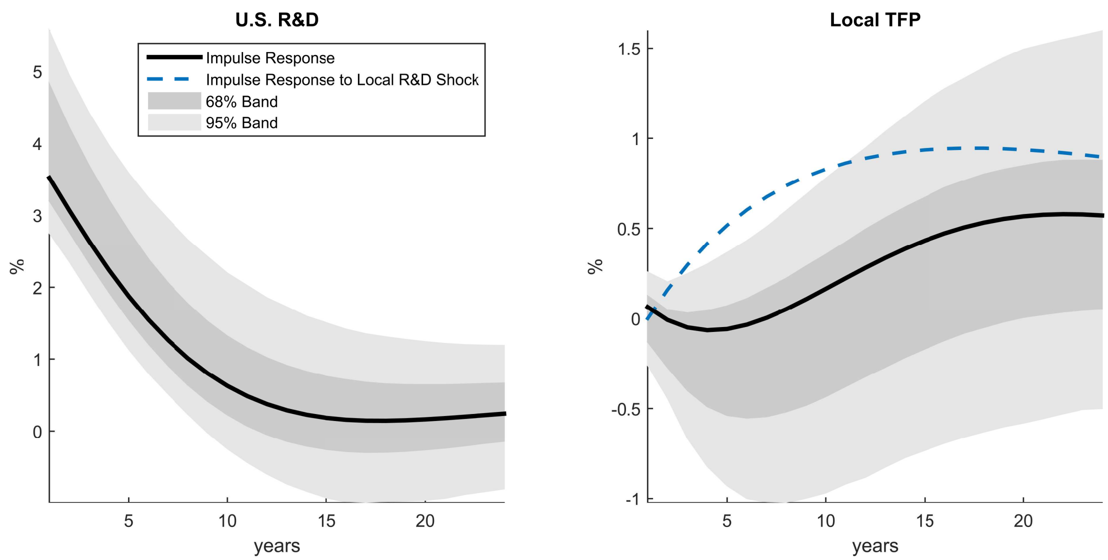

Note: Responses to an identified shock to U.S. R&D expenditure obtained from estimating (3) on the U.S. and the top-5 economies by business-sector R&D expenditure to GDP (Germany, Japan, South Korea, Sweden and Switzerland). The black line represents the dynamic response, and the dark-grey and light-grey shaded areas respectively depict 68% and 95% confidence intervals obtained by bootstrapping with 10,000 repetitions. For comparison, we include the TFP response to own-R&D shocks (see Figure 7), shown by the dashed blue line.

2.4 A larger-scale U.S. VAR

We have focused so far on a three-variable VAR—a minimal setting allowing the study of the effects of R&D on productivity. For robustness, we next estimate a higher-dimensional VAR for the U.S., including a standard set of macroeconomic variables. The estimated model provides interesting information on the effects of R&D shocks on a larger number of variables than was the case for the simple three-variable VAR. We estimate the following model:

[6] Note that the time path of U.S. R&D does not exactly match that in Figure 4. The reason is that the time period used in estimation is now different, since our AE data has a shorter time series dimension.

$$
\begin{bmatrix} y_t^{us} \\ tfp_t^{us} \\ rd_t^{us} \\ inv_t^{us} \\ c_t^{us} \\ \pi_t^{us} \\ r_t^{us} \\ sp_t^{us} \end{bmatrix} = \tilde{c}^{us} + \tilde{B}^{us} \begin{bmatrix} y_{t-1}^{us} \\ tfp_{t-1}^{us} \\ rd_{t-1}^{us} \\ inv_{t-1}^{us} \\ c_{t-1}^{us} \\ \pi_{t-1}^{us} \\ r_{t-1}^{us} \\ sp_{t-1}^{us} \end{bmatrix} + \tilde{u}_t^{us} \tag{4}
$$

In addition to the three variables considered until now, the model above includes aggregate investment (excluding R&D), consumption, inflation, the monetary policy rate, and a real stock price index (respectively, $inv_t^{us}, c_t^{us}, \pi_t^{us}, r_t^{us}, sp_t^{us}$). We measure inflation with the GDP deflator. The stock price index sp_t is the S&P500 index deflated by the GDP deflator. We measure the stance of monetary policy, r_t, with the Wu and Xia (2016) shadow rate. We include inv_t, c_t, and sp_t in logs, π_t in percent annual change, and r_t in annual percentage points. We continue to identify R&D shocks following the Choleski approach, which now allows all the variables below rd_t to respond contemporaneously to R&D shocks.

Figure 9 displays the responses of the variables in (4) to an R&D shock. Note first that the pattern of responses of R&D and TFP is largely unchanged relative to the three-variable VAR: the jump in R&D continues to induce a gradual, persistent rise in TFP. Consumption and investment initially display a muted response (investment actually declines somewhat in the initial years), but eventually rise as the boom in TFP and GDP strengthens. Inflation declines significantly as TFP rises, possibly reflecting the cost-saving benefits of higher productivity, and the policy rate declines somewhat.

Interestingly, stock prices (shown in the bottom-right panel) immediately jump in response to the R&D shock by more than 3 percent, and remain persistently high. This is the case even though the rest of the macroeconomic variables (including TFP) take several years to reach their peak response. This dynamic pattern is reminiscent of the findings by Beaudry and Portier (2006) and related literature on "news shocks." The latter authors show how, in a bivariate setting with TFP and stock prices, two distinct identification schemes (one isolating shocks to stock prices orthogonal to current TFP; the other identifying shocks that drive long-run movements in TFP) isolate almost collinear disturbances, inducing nearly-exact dynamics—with stock prices jumping on impact and TFP rising gradually. That dynamic pattern of stock prices and TFP resembles the one in Figure 9, which we obtain through a

Figure 9: Identified R&D Shock in the U.S., Larger-Scale VAR

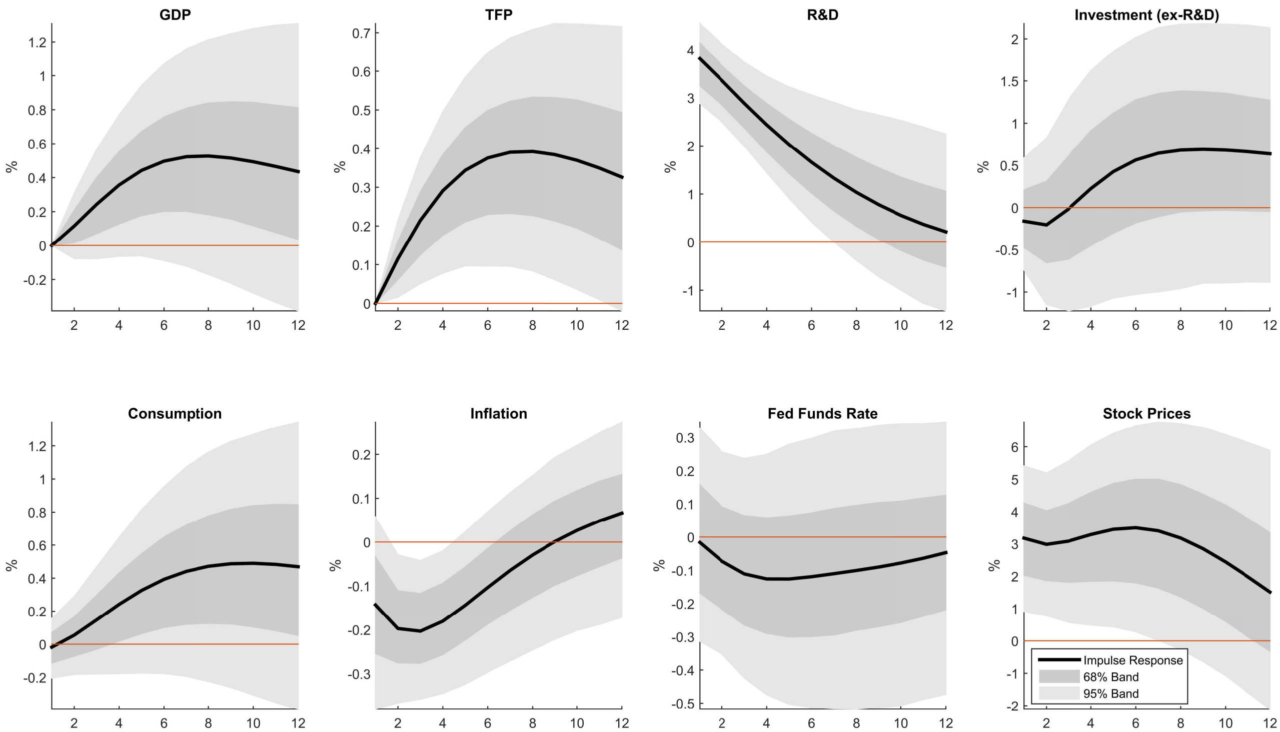

Note: Responses to an identified R&D shock in the larger-scale U.S. VAR (equation (4)). The black line represents the dynamic response, and the dark-grey and light-grey shaded areas respectively depict 68% and 95% confidence intervals obtained by bootstrapping with 10,000 repetitions.

completely different identification strategy—one that relies on shocks to a third variable not examined by Beaudry and Portier (2006), namely R&D expenditure. There are, however, some differences in the the dynamic pattern identified by Beaudry and Portier (2006) and ours. Most notably, in our case the full rise of TFP takes several years to materialize, while in Beaudry and Portier (2006) TFP peaks after about a year and a half (see their Figure 1). Still, our results provide some support for the notion that the findings emphasized by the news shocks literature may in part be due to technology innovation and diffusion effects (as highlighted by Beaudry and Portier (2006)), at least at the lower frequencies.

As a final exercise in this section, we reestimate (4) decomposing the aggregate investment series into two components: investment in information-processing equipment and software on the one hand, and the rest of investment categories (still excluding R&D) on the other. The goal is to examine whether the rise in R&D is accompanied by firms' investments in the implementation of technology, as proxied by the investment categories just mentioned. Figure 10 shows the responses with investment decomposed into the two categories (the response of the

Figure 10: Identified R&D Shock in the U.S., Investment Decomposition

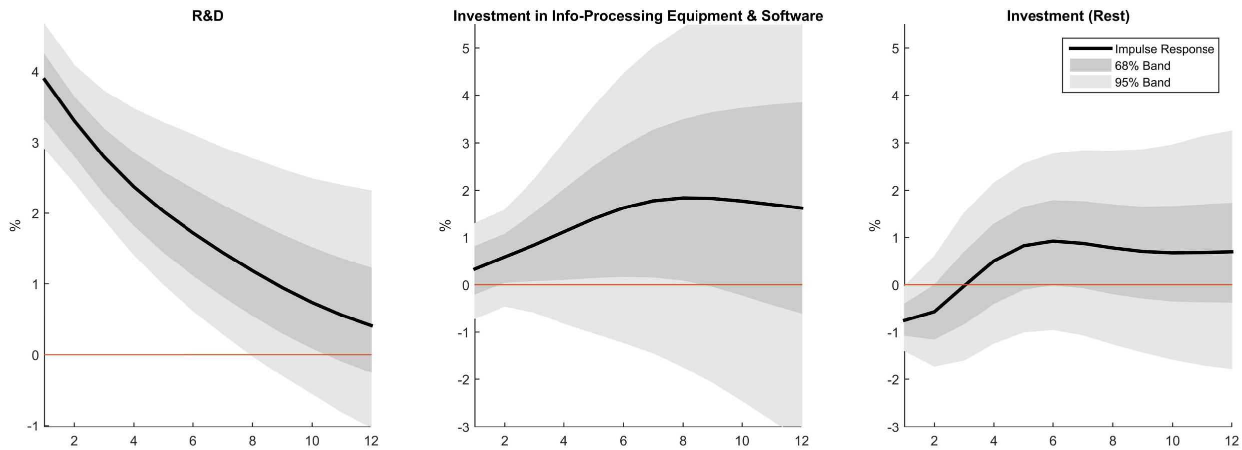

Note: Effects of an identified R&D shock in the larger-scale U.S. VAR on investment in information-processing equipment and software (middle panel) and on the rest of investment categories (right panel). The black line represents the dynamic response, and the dark-grey and light-grey shaded areas respectively depict 68% and 95% confidence intervals obtained by bootstrapping with 10,000 repetitions.

remaining variable is essentially unaffected relative to Figure 9). Note that the response of the information-processing and software investment category is positive throughout (even if featuring considerable uncertainty), and always above the response of the rest of categories, which now falls initially (by more, in percent, than overall investment does in Figure 9). Thus, the evidence supports the notion that the rise in R&D is accompanied by a rise in investments that might be seen as more complementary to R&D, even if investment in other categories initially falls.

3 Model

Our theoretical framework is a standard New Keynesian model augmented to include endogenous technology innovation and adoption, as in Comin and Gertler (2006) or Anzoategui et al. (2016). The formulation of the evolution of technology closely follows Comin and Gertler (2006). The model has six sets of agents: intermediate goods producers, innovators, adopters, households, capital producers, and retailers. Of these, the first three correspond to the endogenous technology mechanism. Capital producers use final output as input for the production of investment goods. Retailers are the source of nominal rigidity in the model. We next describe each set of agents in turn.

3.1 Intermediate Goods Producers

In period t, there exists a continuum of measure A_t of currently available varieties of intermediates, each produced by a monopolistically competitive intermediate goods producer. Wholesale output, Y_t^W, is a CES aggregate of individual intermediate goods:

$$Y_t^W = \left[\int_0^{A_t} Y_t^M(s)^{\frac{\vartheta-1}{\vartheta}} ds \right]^{\frac{\vartheta}{\vartheta-1}} \tag{5}$$

Wholesale output is used to produce final output by retailers, described below, who are subject to nominal rigidities. In (5), $Y_t^M(s)$ is output by intermediates producer s. Each intermediates producer sets (nominal) price $P_t(s)$. The price level of wholesale output associated with (5) is given by $P_t^W = \left[\int_0^{A_t} P_t(s)^{1-\vartheta} ds \right]^{\frac{1}{1-\vartheta}}$. Each intermediate goods firm s uses capital $K_t(s)$ and labor $L_t(s)$ to produce their variety, using a Cobb-Douglas production function:

$$Y_t(s) = \Psi_t K_t(s)^\alpha L_t(s)^{1-\alpha} \tag{6}$$

Here, Ψ_t is an exogenous TFP shock, which is assumed to follow an AR(1) in logs: $\log(\Psi_t) = \rho_\Psi \log(\Psi_{t-1}) + \epsilon_t^\Psi$.

Solving the intermediates goods firm's problem yields standard first order conditions for pricing, labor, and capital. Let W_t be the real wage and $\mathcal{Z}_t$ be the real rental rate of capital. Factor prices are equalized to their respective marginal products:

$$W_t = \frac{\vartheta - 1}{\vartheta} \frac{1}{\mathcal{M}_t} (1-\alpha) \frac{Y_t^W}{L_t} \tag{7}$$

$$\mathcal{Z}_t = \frac{\vartheta - 1}{\vartheta} \frac{1}{\mathcal{M}_t} \alpha \frac{Y_t^W}{K_t} \tag{8}$$

Real per-period profits by intermediates producers, denoted Π_t, are equal across firms and can be shown to be given by

$$\Pi_t = \frac{1}{\vartheta} \frac{1}{\mathcal{M}_t} \frac{Y_t^W}{A_t} \tag{9}$$

In the three equations above, $\mathcal{M}_t$ is the ratio of the final output price level, P_t, over the wholesale price level: $\mathcal{M}_t = \frac{P_t}{P_t^W}$. The determination of P_t is described below, in subsection 3.6 characterizing retailers.

Combining (5) with the first-order conditions for intermediates producers and with equilibrium in factor markets can be shown to yield the following expression for aggregate wholesale output Y_t^W:

$$Y_t^W = A_t^{\frac{1}{\vartheta-1}} \Psi_t K_t^{\alpha} L_t^{1-\alpha} \tag{10}$$

Here, K_t and L_t denote aggregate capital and labor: $K_t \equiv \int_0^{A_t} K_t(s)ds$, $L_t \equiv \int_0^{A_t} L_t(s)ds$. Equation (10) makes clear that measured total factor productivity (TFP) is driven by the measure of varieties of intermediates, A_t, as well as by the exogenous TFP shock Ψ_t. The evolution of the former is described in the next two subsections, characterizing technology innovators and adopters.

3.2 Innovators

Our modeling of innovators follows Comin and Gertler (2006). Competitive innovators spend resources in R&D to develop new intermediate goods. They then sell the rights to new goods to an adopter, who converts the idea for the new product into an employable input, as described in the next subsection.

Specifically, each innovator i has access to the following production function for new innovations:

$$V_{i,t} = \zeta Z_t \frac{1}{K_t^{\eta} N_t^{1-\eta}} N_{i,t} \tag{11}$$

Here $V_{i,t}$ denotes new products developed by innovator i and $N_{i,t}$ denotes R&D expenditure by innovator i (in units of final output). Aggregate R&D is $N_t \equiv \int_i N_{i,t} di$. As in Romer (1990), there is a positive spillover from the aggregate stock of innovations, Z_t, to individual R&D productivity. At the same time, the term $\frac{1}{K_t^{\eta} N_t^{1-\eta}}$ introduces a congestion externality from aggregate R&D: everything else equal, higher aggregate R&D reduces innovators' efficiency of developing new products. Under this formulation, in equilibrium the R&D elasticity of aggregate new technology creation is given by parameter η, satisfying $0 < \eta < 1$. This parameter is one of the key objects that we aim to identify using the evidence described in the preceding section. Also as in Comin and Gertler (2006), the congestion effect depends positively on the aggregate capital stock K_t, capturing the notion that as the economy becomes more sophisticated (as measured by the amount of capital) the efficiency of R&D declines. This term helps ensure that the growth rate of new intermediate products is stationary. The parameter ζ is a scaling factor, which helps the model match the growth rate of TFP in the balanced growth path.

Let J_t be the value of a new "unadopted" innovation. We describe how J_t is determined in the following subsection. Innovations developed at t become available starting at $t + 1$.

Accordingly, letting $\varphi_t \equiv \frac{\zeta Z_t}{K_t^\eta N_t^{1-\eta}}$ and with $\Lambda_{t,t+1}$ denoting the household's stochastic discount factor between t and $t+1$, innovator i's problem is

$$\max_{N_{i,t}} \ \mathbb{E}_t \left(\Lambda_{t,t+1} J_{t+1}\right) \varphi_t N_{i,t} - \left(1 + \Delta_t^n\right) N_{i,t}$$

Given that all innovators make the same choices, we now drop the i subindex. The first-order condition for the problem above is given by

$$\mathbb{E}_t \left(\Lambda_{t,t+1} J_{t+1}\right) \varphi_t = 1 + \Delta_t^n$$

Innovators' problem above includes an exogenous R&D tax, or "wedge," given by the variable Δ_t^n. We assume the wedge follows a first-order autoregressive process: $\Delta_t^n = \rho_N \Delta_{t-1}^n + \epsilon_t^n$. The wedge effectively introduces a gap between the marginal benefit and the marginal cost of innovation. Below, we use variation in the wedge Δ_t^n to initiate movements in R&D. One possible interpretation for the wedge is that it reflects frictions in financial intermediation, constraining credit for innovators.[7] Wedges of this type affecting various agents have been used, for instance, to characterize the recent U.S. Great Recession (see, e.g., Christiano et al. (2015)). More generally, here we think of the wedge as a reduced-form way of inducing movements in R&D, be it due to financing constraints or to other (unmodeled) sources of variation of the desirability of R&D investments.

The aggregate stock of adopted technologies, Z_t, evolves according to the following:

$$Z_{t+1} = \phi Z_t + V_t \tag{12}$$

The parameter ϕ, satisfying $0 < \phi < 1$, captures technological obsolescence. $V_t \equiv \int_i V_{i,t} di$ is the aggregate amount of new innovations introduced in period t.

3.3 Adopters

There is a competitive set of "adopters" that convert available technologies into use. Each adopter succeeds in making a product usable in any given period with probability λ_t (determined below). If the adopter is not successful in period t, he may try again in $t+1$. This success rate depends positively in the amount of adoption expenditures by the adopter. Given that the success rate will be the same across products, this formulation facilitates aggregation. Accordingly, the total number of technologies in use A_t evolves according to the

[7] See Queralto (2013) for an explicit model of this channel.

following law of motion:

$$A_{t+1} = \lambda_t \phi \left(Z_t - A_t \right) + \phi A_t \tag{13}$$

As a way to introduce adjustment costs in adoption activity, we suppose that adopters' input is a specialized good (e.g. equipment) that is produced using final output by equipment producers, described below.[8] The latter agents face adjustment costs tha are analogous to those faced by capital goods producers. We denote the price of the equipment good used by adopters by Q_t^m.

Let $M_{i,t}$ be the amount of equipment used by any given adopter. The probability of a successful adoption, λ_t, depends on $M_{i,t}$ and is given by the following:

$$\lambda_t(M_{i,t}) = \kappa_\lambda \left(\frac{V_t}{Z_t} \right)^\nu M_{i,t}^{\rho_\lambda} \tag{14}$$

with $\kappa_\lambda > 0$, $0 < \nu < 1$, and $0 < \rho_\lambda < 1$. The probability of a successful adoption is increasing and concave in adoption effort $M_{i,t}$. In addition, it includes a "spillover" term from aggregate innovation V_t (relative to the total stock of innovations). The idea here is that aggregate innovation may have a benign effect on the likelihood of adoption of existing innovations, for example because adopters learn from recently introduced innovations.[9] In addition to having some plausibility, this spillover term helps prevent a fall in adoption rates in response to a shock to the innovation wedge, as we illustrate below.[10]

An adopter i buys the rights to an unadopted technology from innovators, at competitive price J_t. The adopter then uses resources $M_{i,t}$ which lead to the technology becoming usable for production with probability $\lambda_t(M_{i,t})$. If the adopter is successful, he sells the adopted technology to goods producers obtaining for it the price H_t, given by

$$H_t = \Pi_t + \phi \mathbb{E}_t \left(\Lambda_{t+1} H_{t+1} \right)$$

where Π_t is the monopoly profit from operating the technology, given by (9).

[8] Adopters' adjustment costs help avoid excessive volatility in adoption activity, e.g. in response to monetary shocks.

[9] Griffith et al. (2004), for instance, emphasize that an important role of R&D is to facilitate the adoption of existing innovations.

[10] In previous versions of the paper we introduced an "adoption wedge" proportional to the innovation wedge, which accomplished a similar objective.

The problem of an adopter is

$$J_t = \max_{M_{i,t}} -Q_t^m M_{i,t} + \phi \mathbb{E}_t \Lambda_{t+1} \left\{ \lambda_t(M_{i,t}) H_{t+1} + [1 - \lambda_t(M_{i,t})] J_{t+1} \right\}$$

Adopters' first-order condition is given by the following:

$$\rho_\lambda \phi \left(\frac{V_t}{Z_t} \right)^\eta \mathbb{E}_t \Lambda_{t+1} \left(H_{t+1} - J_{t+1} \right) = Q_t^m M_{i,t}^{1-\rho_\lambda}$$

Since $M_{i,t}$ is the same across adopters, we now drop the i subscript. Adoption effort M_t is increasing in the expected discounted value of the difference $H_t - J_t$, i.e. in the difference in value between an adopted and an unadopted technology.

In period t there is a measure $Z_t - A_t$ of technologies which adopters are attempting to adopt, with each adopter using M_t equipment goods. Accordingly, the aggregate amount of goods used by adopters is given by $(Z_t - A_t)M_t$.

3.4 Households

The representative household chooses (real) consumption C_t, labor supply L_t, holdings of nominal riskless bonds B_{t+1}, and holdings of physical capital K_{t+1} to maximize

$$\mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \left[\log(C_t - hC_{t-1}) - \frac{\chi}{1+\epsilon} L_t^{1+\epsilon} \right]$$

subject to a sequence of budget constraints

$$C_t + \frac{B_{t+1}}{P_t} + Q_t K_{t+1} \leq W_t L_t + R_t \frac{B_t}{P_t} + \left[\mathcal{Z}_t + (1-\delta) Q_t \right] K_t + \tilde{\Pi}_t$$

Here W_t is the real wage, Q_t is the price of capital, $\mathcal{Z}_t$ is the capital rental rate and $\tilde{\Pi}_t$ is total profits distributed to the household (from both output and capital producers). The parameter h, satisfying $0 < h < 1$, governs the presence of consumption habits.

Following Christiano et al. (2015), Smets and Wouters (2007), and others, we modify the household's optimality conditions to include an exogenous "consumption wedge" Δ_t^b, which works to distort the household's Euler equation for riskless bonds. We assume that (the log of) Δ_t^b follows a first-order autoregressive process: $\log(\Delta_t^b) = \rho_b \log(\Delta_{t-1}^b) + \epsilon_t^b$. Christiano et al. (2015), for instance, use a wedge of this type (in combination with other shocks) to model the disturbances triggering the U.S. Great Recession. In a similar vein, Anzoategui et al. (2016) introduce a time-varying preference for the riskless bond in households' utility function, which works to modify consumers' Euler equation in a similar way. More generally, this type of shock has been used in the literature to capture sources of aggregate demand variation.

Letting inflation be $\pi_t \equiv P_t/P_{t-1}$, the household's optimality conditions for riskless bond holdings, physical capital, and labor supply are then given by the following:

$$1 = \mathbb{E}_t \left(\Lambda_{t,t+1} \frac{R_t}{\pi_{t+1}} \right) \Delta_t^b \tag{15}$$

$$1 = \mathbb{E}_t \left[\Lambda_{t,t+1} \frac{\mathcal{Z}_{t+1} + (1-\delta) Q_{t+1}}{Q_t} \right] \tag{16}$$

$$\chi L_t^{\epsilon} = U_{C,t} W_t \tag{17}$$

where the household's stochastic discount factor and marginal utility of consumption are respectively given by

$$\Lambda_{t,t+1} = \frac{U_{C,t+1}}{U_{C,t}} \tag{18}$$

$$U_{C,t} = \frac{1}{C_t - hC_{t-1}} - \beta h \mathbb{E}_t \left[\frac{1}{C_{t+1} - hC_t} \right] \tag{19}$$

3.5 Capital and Equipment Producers

Capital producers make new capital goods using final output as input, and are subject to adjustment costs. They sell new capital to househods at price Q_t. The objective of the representative capital producer is to choose a state-contingent sequence $\{I_t\}$ to maximize the expected discounted value of profits:

$$\mathbb{E}_t \sum_{s=0}^{\infty} \Lambda_{t,t+s} \left\{ Q_{t+s} I_{t+s} - \left[1 + f\left(\frac{I_{t+s}}{I_{t+s-1}} \right) \right] I_{t+s} \right\} \tag{20}$$

where the function f is convex and satisfies $f(\overline{g}) = f'(\overline{g})$ and $\psi_N \equiv f(\overline{g}) > 0$, with $\overline{g}$ denoting the growth rate of investment along the balanced growth path (which coincides with the growth rate of technology, output, and other aggregates).

From profit maximization, we obtain that the price of capital goods is equal to the marginal cost of investment goods production:

$$Q_t = 1 + f\left(\frac{I_t}{I_{t-1}} \right) + \frac{I_t}{I_{t-1}} f'\left(\frac{I_t}{I_{t-1}} \right) - E_t \Lambda_{t+1} \left(\frac{I_{t+1}}{I_t} \right)^2 f'\left(\frac{I_{t+1}}{I_t} \right) \tag{21}$$

The aggregate stock of physical capital then follows the law of motion below:

$$K_{t+1} = (1-\delta)K_t + I_t \tag{22}$$

Equipment producers face a problem analogous to that of capital producers, with an identical adjustment cost function. Letting equipment goods produced by the (representative) equipment producer be I_t^m, the latter's objective is the same as in (20) replacing Q_{t+s} by Q_{t+s}^m and I_{t+s} bt I_{t+s}^m. In the aggregate, the market for equipment goods must clear, so we must have $I_t^m = (Z_t - A_t)M_t$.

3.6 Retailers

There is a continuum of mass unity of retailers, who produce final output using wholesale output as input. Each producer simply purchases wholesale output, costlessly differentiates it, and sells it to final output users. Retailers are subject to nominal rigidities: each retailer can only reset its price with probability $1-\theta$, and must keep its price fixed with the complementary probability. Firms not resetting their price index partially to previous-period inflation (with elasticity ι_p), and partially to steady state inflation (with the complementary elasticity). Final output Y_t is a CES composite of retailers' output:

$$Y_t = \left[\int_0^1 Y_t^R(k)^{\frac{\omega_t - 1}{\omega_t}} dk\right]^{\frac{\omega_t}{\omega_t - 1}} \tag{23}$$

where $Y_t^R(k)$ is output by retailer $k \in [0,1]$. To allow for a source of variation in firms' desired markups, we assume that the elasticity of substitution ω_t is time-varying, and follows a first-order autoregressive process: $\log(\omega_t) = \log(\omega_{t-1}) + \epsilon_t^\omega$.

Let the price set by retailer k be $P_t(k)$. Then cost minimization by users of final output yields the following demand function for each retailer k:

$$Y_t^R(k) = \left[\frac{P_t(k)}{P_t}\right]^{-\omega_t} Y_t \tag{24}$$

where the final output price level, P_t, is

$$P_t = \left[\int_0^1 P_t(k)^{1-\omega_t} dk\right]^{\frac{1}{1-\omega_t}} \tag{25}$$

Nominal marginal cost for retailers is P_t^W. Let the indexation term be $I_{t,t+\tau} \equiv \prod_{k=1}^{\tau} \pi_{t+k-1}^{\iota_p} \pi^{1-\iota_p}$ for $\tau \geq 1$, where π denotes steady-state inflation. Given the pricing friction, the problem of a retailer is the following:

$$\max_{P_t^*} \quad \mathbb{E}_t \sum_{i=0}^{\infty} \theta^i \Lambda_{t,t+i} \left(\frac{P_t^* I_{t,t+i}}{P_{t+i}} - \frac{P_{t+i}^W}{P_{t+i}} \right) Y_{t,t+i}^R \tag{26}$$

subject to

$$Y_{t,t+i}^R = \left[\frac{P_t^* I_{t,t+i}}{P_{t+i}} \right]^{-\omega_t} Y_{t+i} \tag{27}$$

This problem leads to the usual first-order condition for the optimal reset price P_t^*:

$$P_t^* = \frac{\omega_t}{\omega_t - 1} \frac{\mathbb{E}_t \sum_{i=0}^{\infty} \theta^i \Lambda_{t,t+i} \frac{1}{\mathcal{M}_{t+i}} I_{t,t+i}^{\omega_t} P_{t+i}^{\omega_t} Y_{t+i}}{\mathbb{E}_t \sum_{i=0}^{\infty} \theta^i \Lambda_{t,t+i} I_{t,t+i}^{1-\omega_t} P_{t+i}^{\omega_t - 1} Y_{t+i}}$$

where as mentioned earlier, the variable $\mathcal{M}_t$ is given by the ratio of price levels $\frac{P_t}{P_t^W}$ (i.e., the inverse of retailers' real marginal cost).

From the law of large numbers, the evolution of the price level is

$$P_t = \left[\theta \left(\pi_{t-1}^{\iota_p} \pi^{1-\iota_p} P_{t-1} \right)^{1-\omega_t} + (1-\theta) P_t^{*1-\omega_t} \right]^{\frac{1}{1-\omega_t}} \tag{28}$$

3.7 Central Bank, Resource Constraint, and Stock Prices

We suppose that monetary policy is characterized by a simple Taylor rule with interest-rate smoothing, where the systematic component of policy responds to inflation and to the output gap (approximated by the inverse of the price markup). Accordingly, the policy rule is

$$R_t = R_{t-1}^{\gamma_r} \left[\left(\frac{\pi_t}{\pi} \right)^{\gamma_\pi} \left(\frac{Y_t}{Y_t^{pot}} \right)^{\gamma_y} \overline{R} \right]^{1-\gamma_r} r_t^m \tag{29}$$

where Y_t^{pot} denotes potential output (defined as the level of output that would result with perfectly flexible prices and no markup shocks), and the steady-state interest rate $\overline{R}$ is given by $\overline{g}\pi/\beta$. The rule includes a monetary policy shock, given by r_t^m, which follows the stochastic process $\log(r_t^m) = \rho_m \log(r_{t-1}^m) + \epsilon_t^m$.

Equilibrium in wholesale output requires $Y_t^W = \int_0^1 Y_t^R(k)dk$. Combining this condition with (27), the relation between wholesale and final output is

$$Y_t = \frac{Y_t^W}{\mathcal{D}_t} \tag{30}$$

where $\mathcal{D}_t$ is given by a measure of price dispersion across retailers: $\mathcal{D}_t \equiv \int_0^1 \left[\frac{P_t(k)}{P_t}\right]^{-\omega_t} dk$. It can be shown that $\mathcal{D}_t \geq 1$, and that $\mathcal{D}_t \approx 1$ to a first order (i.e., output losses due to price dispersion are second order).

The aggregate resource constraint is given by

$$Y_t = C_t + \left[1 + f\left(\frac{I_t}{I_{t-1}}\right)\right] I_t + \left[1 + f\left(\frac{I_t^m}{I_{t-1}^m}\right)\right] I_t^m + N_t \tag{31}$$

Final output is used for consumption, investment, adoption and innovation.

Finally, stock prices $\mathcal{S}_t$, defined as the present discounted value of dividends of the entire firm sector, are given by

$$\mathcal{S}_t = Q_t K_t + (H_t - \Pi_t) A_t + (J_t + M_t Q_t^m)(Z_t - A_t) \tag{32}$$

Stock prices in the model reflect not only the value of physical capital $Q_t K_t$ but also the value of currently adopted and unadopted technologies (the second and third summands in the right-hand-side of (32)), as originally highlighted by Comin et al. (2009) and later used by Kung and Schmid (2015), among others.

This completes the description of the model.

4 Model Estimation

We now proceed to estimate some of the model's key parameters, by using the empirical responses to an identified R&D shock documented in section 2.1. Because our focus is on the link between R&D and productivity, we focus on the impulse responses from the small-scale VAR when estimating the model, although we later check how well the model fares against the larger-scale VAR responses.

We partition the parameters into two sets: the first set contains mostly standard preference and technology parameters which we calibrate following the literature. The second set contains parameters that we estimate by minimizing the distance between empirical and model-simulated impulse responses. The key parameter within this set is the elasticity of innovation to R&D, η. We next discuss each parameter set in turn.

4.1 Calibrated Parameters

Since our data is annual, we calibrate the model at an annual frequency. The calibrated parameter values are shown in table 1. Our calibration for common preference and technology

Table 1: Calibrated Parameters

Symbol	Value	Description
β	0.9978	Discount factor
α	0.33	Capital Share
δ	0.1	Capital depreciation
ϵ^{-1}	2	Frisch labor supply elasticity
h	0.50	Habit
ϑ	2.4925	Intermediates producers' elasticity of substitution
ϕ	0.92	Obsolescence of technologies
ρ_λ	0.95	Adoption elasticity
ν	1/3	Innovation spillover to adoption
$\overline{L}$	1	Steady-state labor
$\overline{g}^{\frac{1}{\vartheta-1}}$	1.0120	Steady-state TFP growth (gross)
$\overline{\lambda}$	0.20	Steady-state adoption probability
$\overline{\omega}$	4.167	Retailers' average elasticity of substitution
θ	0.65	Probability of keeping prices fixed
ι_p	0.20	Degree of indexation to pat inflation
π	1.02	Steady-state inflation (gross)
γ_r	0.32	Smoothing parameter of the Taylor rule
γ_π	1.5	Inflation coefficient of the Taylor rule
γ_y	0.5	Output gap coefficient of the Taylor rule
ρ_Ψ	0.9	Exogenous TFP shock persistence
ρ_b	0.65	Consumption wedge persistence
ρ_ω	0.33	Markup shock persistence
ρ_Ψ	0.10	Monetary shock persistence

parameters is relatively standard. We set the discount factor, β, to 0.9978, to deliver a balanced-growth-path real interest rate of 2 percent annually. The capital share α is set to 0.33, and the capital depreciation rate is $\delta = 0.1$. We calibrate ϵ to 0.5, resulting in a Frisch elasticity of labor supply of 2. We set the habit parameter h to 0.50, somewhat below typical estimates, to account for the fact that these estimates typically result from quarterly data while our model is annual. The parameter governing the elasticity of final output with respect to intermediates, ϑ, is chosen so that the technological level A_t takes the purely labor-augmenting form, which amounts to imposing the restriction $(1-\alpha)(\vartheta-1) = 1$.[11] This restriction implies that there exists a balanced growth path along which output is proportional to TFP, and therefore profits per period Π_t are stationary (see equation (9)). Given the choice for α, the resulting value for the intermediate goods markup is $\vartheta/(\vartheta-1) = 1.67$, close to the value of 1.6 chosen by Comin and Gertler (2006). We set the technology obsolescence rate, $1-\phi$, to 10 percent annually, similar to Anzoategui et al. (2016), who rely on estimates of technological obsolescence from Caballero and Jaffe (1993). We follow Comin and Gertler (2006) and set the elasticity of the adoption probability to adoption expenditure, ρ_λ, to 0.95. This value helps deliver a realistic ratio of R&D to GDP in steady state, and is also consistent with measures of the cyclicality of technology diffusion, as Anzoategui et al. (2016) show.

To set the parameters χ (labor disutility), ζ (productivity of R&D), and κ_λ (constant in the adoption rate), we target properties of the model's balanced growth path. In particular, we normalize the level of labor $\overline{L}$ to unity, and target a TFP growth rate of 1.20% and an adoption rate of 0.20. The target growth rate to the average annual growth rate of TFP for the U.S. The adoption rate target follows Comin and Gertler (2006) and Anzoategui et al. (2016), who rely on evidence on average technology adoption lags. The average adoption lag in the model is given by $\frac{1}{\overline{\lambda}}$; the chosen value for $\overline{\lambda}$ thus implies an average adoption lag of five years. Given these targets, we then back out the parameters χ, ζ and κ_λ. In our estimation procedure below, we always keep the targets fixed as we search over the estimated parameters, thus ensuring that our estimates are always consistent with our targeted values for the balanced growth path.

We also need to assign a value to the spillover parameter ν, governing the impact of aggregate innovation on the adoption rate. Ideally we would like to calibrate or estimate this parameter based on evidence on the impact of R&D on adoption rates or adoption expenditures. Comprehensive measures of the latter, however, are not available. We thus follow an alternative strategy to calibrate ν, which relies on imposing that the response of adoption rates to R&D is not very large. In particular, we pick ν so that the response of λ_t

[11] Kung and Schmid (2015) make a similar parameter restriction.

Table 2: Estimated Parameters

Symbol	Value	Description
η	0.35	Elasticity of technology creation to R&D
ρ_n	0.78	Persistence coefficient of Δ_t^n
σ_n	0.04	Size of impulse to Δ_t^n

to Δ_t^n is zero on average over the first ten years, resulting in a value of 1/3.[12]

We set the elasticity of substitution across retailers, ω, to 4.167, following Primiceri et al. (2006). Our choice of the price rigidity parameter $\theta = 0.65$ reflects the equivalent in annual terms to the quarterly estimate in Anzoategui et al. (2016) of about 0.9. We set the degree of indexation to past inflation, ι_p, to 0.20, following estimates from Primiceri et al. (2006) and Smets and Wouters (2007). The steady-state inflation rate is set to 2 percent per year. We set the Taylor rule coefficients on inflation and output, γ_π and γ_y, to 1.5 and 0.5 respectively, both standard values. We also use a standard value for the interest rate smoothing parameter, γ_r, which we set to 0.32, corresponding to 0.75 at the quarterly frequency. Finally, we assign conventional values to the shock persistence parameters. We assume that exogenous TFP is a high-persistence process, and accordingly set $\rho_\Psi = 0.9$ (corresponding to a quarterly persistence of $0.9^{\frac{1}{4}} = 0.974$), similar to the estimate by Anzoategui et al. (2016)). We also take our estimate of the consumption wedge persistence from Anzoategui et al. (2016). Finally, we set the markup and monetary policy disturbances to low-persistence processes, in line with estimates from Primiceri et al. (2006) and Anzoategui et al. (2016).

4.2 Estimated Parameters

We estimate the following three parameters: the elasticity of new innovations with respect to R&D, η; the first-order autoregressive coefficient of the innovation wedge, ρ_N; and the size of the impulse to the innovation wedge, σ_n. Let the subset of estimated model parameters be $\varepsilon \equiv (\eta, \rho_n, \sigma_n)$. Let also $\Psi(\varepsilon)$ denote the mapping from ε to the model's impulse responses to the initiating shock to Δ_t^n, and let $\hat{\Psi}$ be the empirical impulse responses from the panel VAR in section 2.1. We use the first 20 years of each response. We estimate ε by solving

$$\min_{\varepsilon} \left[\hat{\Psi} - \Psi(\varepsilon)\right]' V^{-1} \left[\hat{\Psi} - \Psi(\varepsilon)\right] \tag{33}$$

[12] As we discuss in detail below, when $\nu = 0$ the model predicts a drop in adoption rates following a shock to Δ_t^n, so that a positive ν is required to maintain adoption rates stable following the innovation shock.

Here, V denotes a diagonal matrix with the variances of the estimated impulse responses along the main diagonal. The weighting matrix V gives relative more weight to more precise estimates in the optimization problem above.

Table 2 contains the resulting parameter estimates. Our estimate of the elasticity of aggregate new technology production with respect to aggregate R&D expenditure, η, is 0.35, a value lower than used by Comin and Gertler (2006) but in the vicinity of the value estimated by Anzoategui et al. (2016). The estimate of the "spillover" parameter α_n is 0.108, indicating that the data favors an increase in technology adoption to occur alongside the rise in R&D. The estimated size of the adoption wedge is about one-tenth of the size of innovators' wedge. The first-order autoregressive coefficient ρ_n is estimated to be 0.81, in line with the considerable persistence of R&D in the data, and the size of the impulse to Δ_n is 5.6 percent.

Figure 11: Impulse response to R&D shock, model v. data

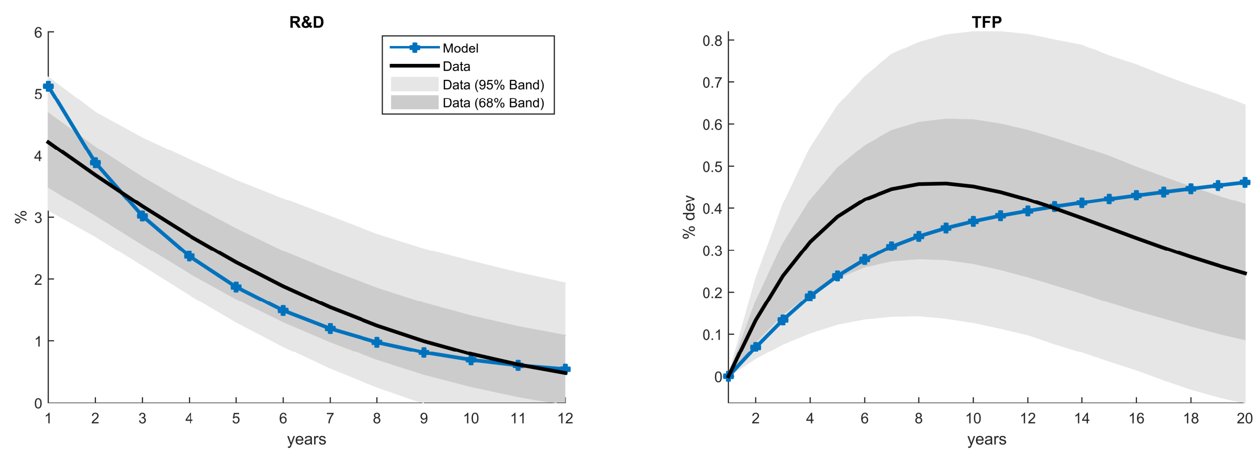

Note: As in Figure 4, the black solid lines show the empirical responses to an R&D shock in the U.S. VAR from section 2.1, and the shaded areas indicate confidence intervals. The solid blue lines with circles show the model's responses to a shock to the R&D wedge at the estimated parameter values.

Figure 11 plots the empirical impulse responses from section 2.1 along with the model-generated responses, computed using the estimated parameter values in Table 2. The model tracks the empirical movements in R&D and TFP reasonably well. As seen in the Figure, in both the model and the data, an increase in R&D of about 4 percent initially leads the level of TFP to rise about 0.4 in the medium-run.

We next document the model's transmission from R&D to TFP, and illustrate the role of the parameter η in shaping the model's responses. Figure 12 shows the impulse responses of several variables pertaining to the innovation and adoption sectors at our estimated values (blue solid line), along with the responses resulting from setting $\eta = 0.175$ (i.e. fifty percent below its estimate), shown by the green dash-dotted line. Throughout we continue to maintain the steady-state targets for $\overline{L}$, $\overline{g}$ and $\overline{\lambda}$. As seen in the Figure, the increase in R&D spurs the creation of new innovations V_t, which add to the stock of existing technologies, Z_t. As these innovations become adopted, the stock of technologies in use (A_t) rises, which accounts for the rise in TFP.

Figure 12: R&D shock transmission, sensitivity to η

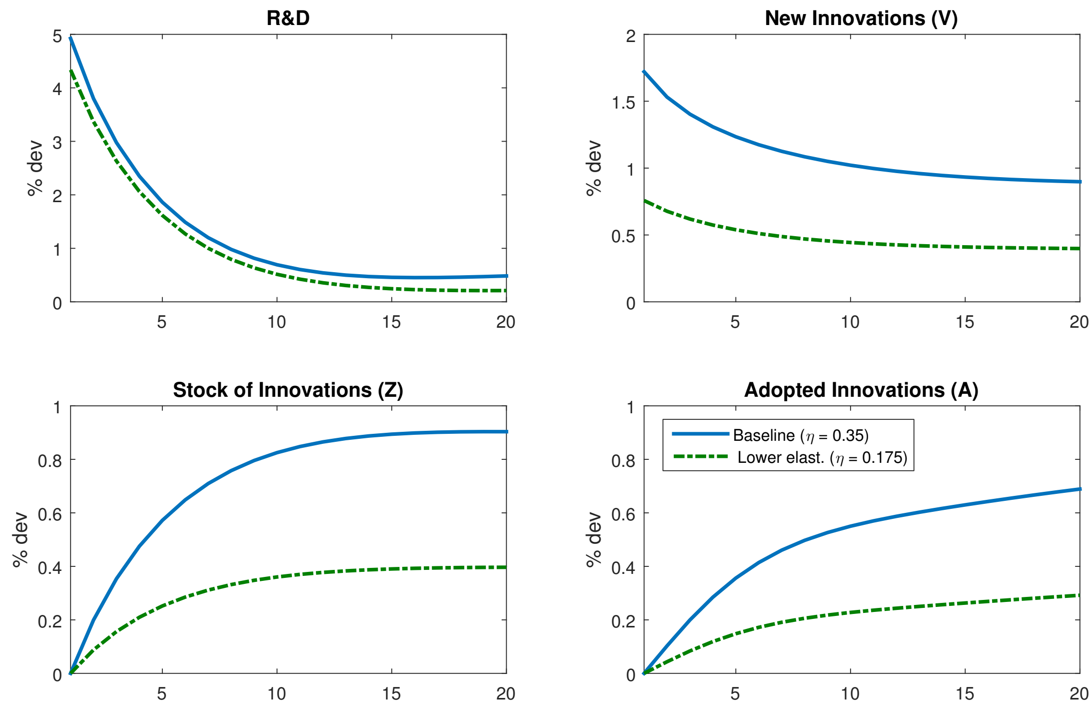

Note: The solid blue line represents the impulse responses at the estimated parameter values, and the green dash-dotted line shows the responses when lowering η by fifty percent relative to its estimated value.

Note from Figure 12 that with a lower elasticity of innovation to R&D, the increase in V_t from a given rise in R&D becomes substantially smaller. As a consequence, the total stock of technologies Z_t rises by much less. In addition, technology adoption also weakens, through the innovation spillover term. As a result, the rise in A_t is much smaller. In this way, the data helps identify the magnitude of the parameter η.

We next discuss the role of the spillover parameter ν. To this end, we first reestimate the parameter vector ε, this time imposing $\nu = 0$ (i.e. no spillover from aggregate innovation to the adoption rate). The resulting parameter estimates are $\eta = 0.73$, $\rho_n = 0.91$ and $\sigma_n = 0.04$. Figure 13 shows the model's behavior in this case (red dash-dotted line), compared to our baseline case with $\nu = 1/3$. As we discussed earlier, the baseline case has the adoption rate remain close to its steady-state value (of about twenty percent per year). By contrast, absent the spillover, the adoption rate falls significantly. This results from a type of substitution effect: given the decrease in the innovation wedge—which works to make investments in R&D more desirable—agents optimally direct resources toward that activity, and away from other activities (including technology adoption). The lower adoption rate then makes it very hard for the model to match the data: note from the bottom-left panel that the increase in A_t with $\nu = 0$ is very gradual, and thus clearly at odds with the data. This is the case even though the elasticity of innovation to R&D is more than twice as large in this case (0.73 compared to 0.35).

Figure 13: R&D shock, effect of adoption spillover

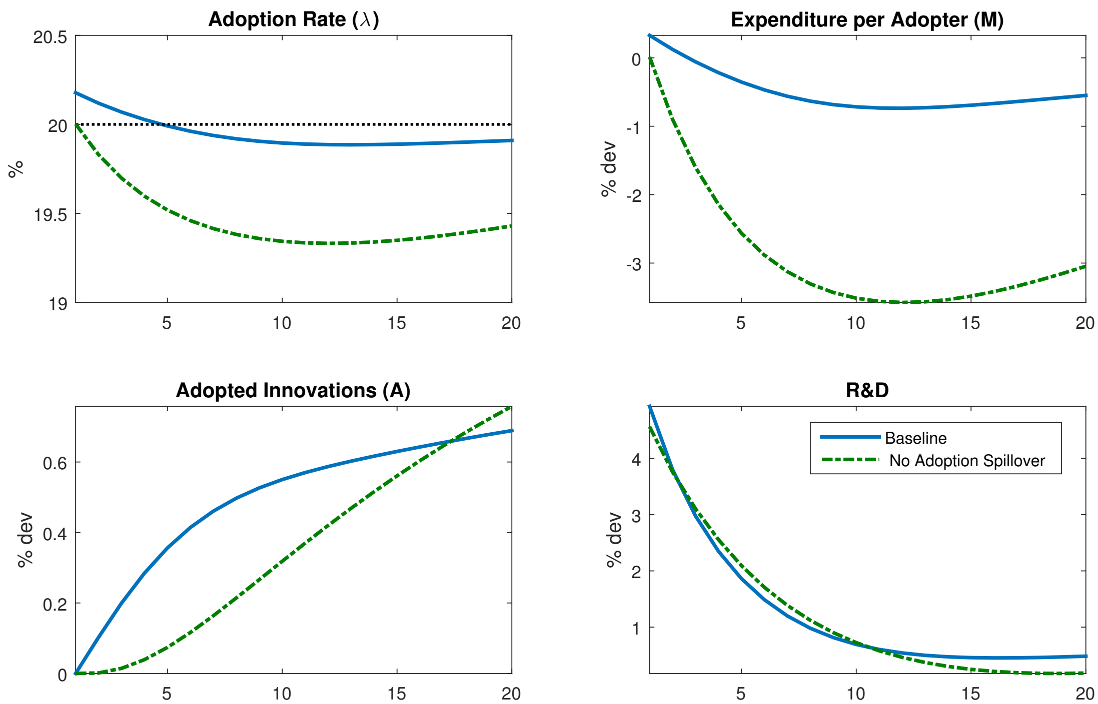

Note: The solid blue line represents the impulse responses at the estimated parameter values, and the green dash-dotted line shows the responses in a (reestimated) model that imposes $\nu = 0$ (no adoption spillover). The reestimated parameter values in the no-spillover model are $\eta = 0.73, \rho_n = 0.91, \sigma_n = 0.04$.

We conclude this section by assessing the model's fit *vis-à-vis* the larger-scale VAR estimated in section 2.4. Figure 14 plots the empirical responses for the larger set of variables, along with the model counterparts. The model matches macroeconomic aggregates reasonably well, while it has more trouble matching inflation and the fed funds rate, and particularly stock prices. In the model, the rise in productivity does eventually put downward pressure on inflation. This effect, however, is offset initially by the boost in aggregate demand that results from higher desired investment in R&D and technology adoption. This effect is also

Figure 14: Larger-Scale VAR, Model v. Data

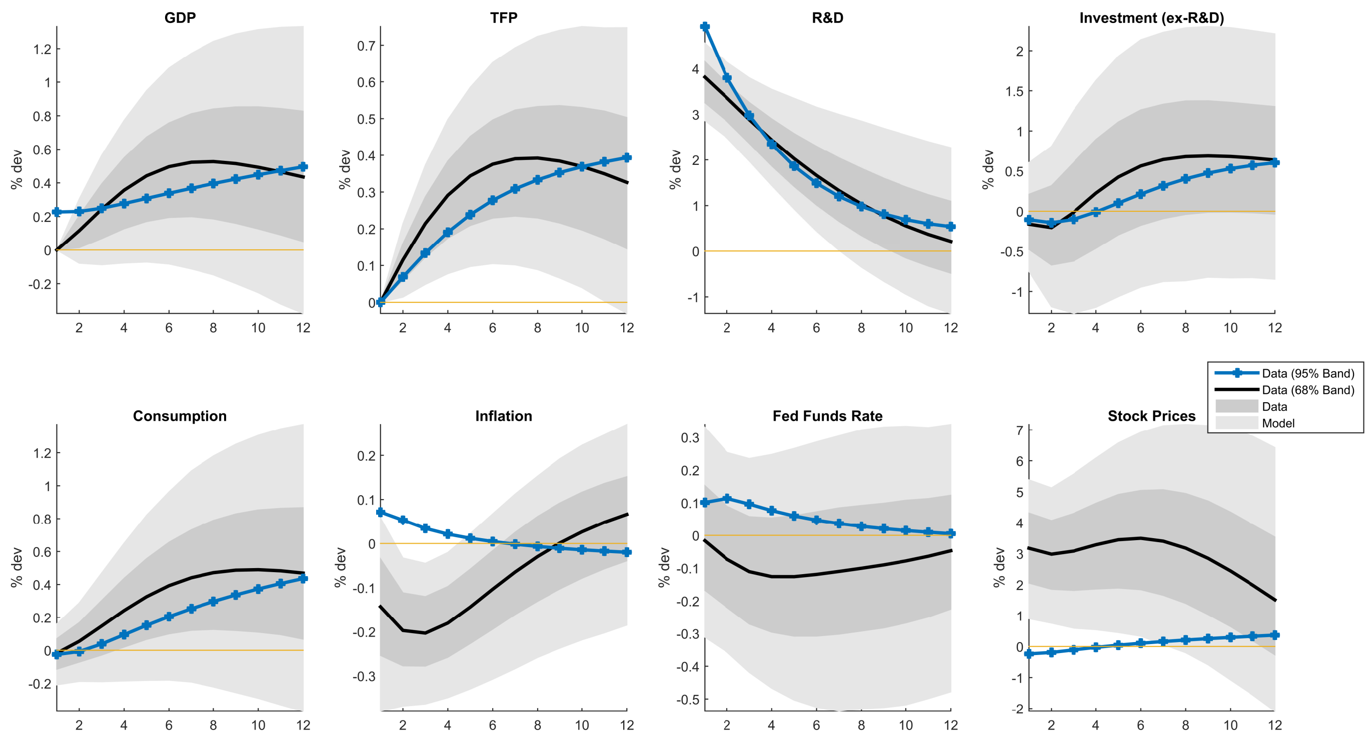

Note: Responses to an identified R&D shock in the larger-scale U.S. VAR as in Figure 9 (black solid lines and shaded areas), along with model impulse responses (blue circled lines).

partially responsible—together with a rise in labor supply due to the wealth effect (note that consumption initially falls slightly)—for the short-run increase in output.[13] The interest rate increase, on the other hand, is mainly due to a rise in the natural rate, resulting from higher expected consumption growth. The resulting more-heavy discounting of future profits, together with a rise in the profit flow Π_t that is fairly transitory, accounts for the small reaction of stock prices.

5 Historical Analysis

Given the model and the estimated parameter values, we now turn to a historical analysis of U.S. productivity growth. Our focus is on exploring the relevance of technology innovation and adoption for the evolution of observed TFP growth in recent times. We are interested,

[13]One possible way to mitigate this effect might be to assume that, instead of using final goods as inputs as in Comin and Gertler (2006), innovators and adopters use specialized labor, as assumed by Anzoategui et al. (2016).

in particular, in examining to what extent the post-Great Recession decline in R&D was responsible for the low rates of TFP growth. We also examine whether there is scope for monetary policy in providing a boost to productivity growth.

To this end, we use the model to obtain "smoothed" sequences of the exogenous innovations $\epsilon_t^n, \epsilon_t^\Psi, \epsilon_t^m, \epsilon_t^b$ and ϵ_t^ω (respectively, disturbances to the innovation wedge, exogenous TFP, monetary, consumption wedge, and price markup) using data on R&D growth, TFP growth, the shadow Fed Funds rate, output, and inflation. Armed with these, we next construct a series of experiments and counterfactuals, aimed at addressing the questions posed above.

5.1 Drivers of the productivity slowdown

We start with a decomposition of measured TFP in the data into its endogenous and exogenous components. From (10), measured TFP is

$$TFP_t = A_t^{\frac{1}{\vartheta-1}} \Psi_t \tag{34}$$

Thus, a natural question is how much of observed TFP growth is accounted for by the endogenous component $A_t^{\frac{1}{\vartheta-1}}$, and how much is picked up by the exogenous TFP shock Ψ_t. Figure 15 plots the corresponding decomposition (obtained using the log-differenced version of (34)). Because we are interested in illustrating what drives the slowdown in recent decades, we show each component relative to its own value in 1998 (when TFP growth peaked). The first observation from Figure 15 is that the endogenous component (shown by the green dash-dotted line) is quite slow moving. Accordingly, the higher-frequency movements in TFP growth in the data are picked up disproportionately by exogenous TFP shocks, as made clear by comparing the red and blue lines in Figure 15.

That said, endogenous TFP is identified to play a significant role in the post-2000s slowdown in TFP growth. After an upward trend in the second half of the 1990s (featuring a total rise, between 1996 and 2001, of about one-half percentage point), this component starts a marked decline. The decline has two phases: one coinciding with the 2001 recession (during which R&D expenditure fell sharply), and the other beginning around 2006 and accelerating during the Great Recession.

The decline in the endogenous component accounts for a significant part of the TFP growth slowdown in recent years. To illustrate this point clearly, Figure 16 shows the same decomposition as in Figure 15, but this time expressing the data (as well as each component) as 5-year moving averages. This aids in the interpretation as it smooths out the more volatile year-to-year variations in the TFP growth rate. As made clear by the Figure, the decline

Figure 15: Decomposition of TFP growth

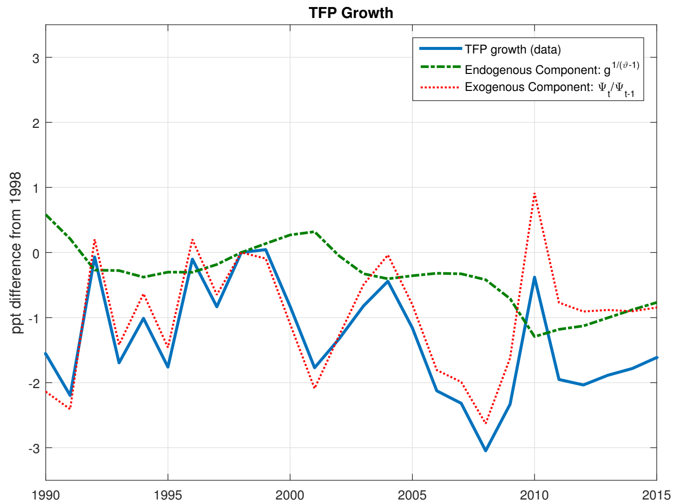

Note: The blue solid line is log-differenced TFP in the data, the green dash-dotted line is its endogenous component, and the red dotted line is its exogenous component (see equation (34)).

in the endogenous component of TFP growth since the late 1990s accounts for a large and growing portion of the fall in the TFP growth in the data—about a third percent, on average, between 2001 and 2014.

5.2 The role of the Great Recession

We next turn to the issue of how much the decline in R&D expenditure seen during and after the Great Recession contributed to the weakness in productivity growth. To address this question, we use the model to produce a counterfactual scenario in which, instead of sharply declining, R&D expenditure remains on its pre-crisis trend. In particular, as shown in the top-left panel of Figure 17, we suppose that starting in 2009, R&D continues to grow at a constant rate—equal to its average growth from 2005 through 2008. We see this as a simple, yet plausible, way of projecting the likely evolution of R&D absent the crisis.[14] Because we want to isolate the contribution of R&D expenditure, we simply engineer the counterfactual R&D path via shocks to the innovation wedge Δ_t^n, by searching for the alternative path of disturbances ϵ_t^n that accomplishes our desired path of R&D. As shown in the left panel of Figure 18, this requires a path for the disturbances ϵ_t^n that remains low since 2009, rather than rising sharply—which then has the effect of keeping R&D growth low, as shown in the left panel of Figure 18.

[14]Christiano et al. (2015) use a similar method to characterize the effects of the Great Recession on a broad set of variables.

Figure 16: Decomposition of TFP growth, 5-year moving average

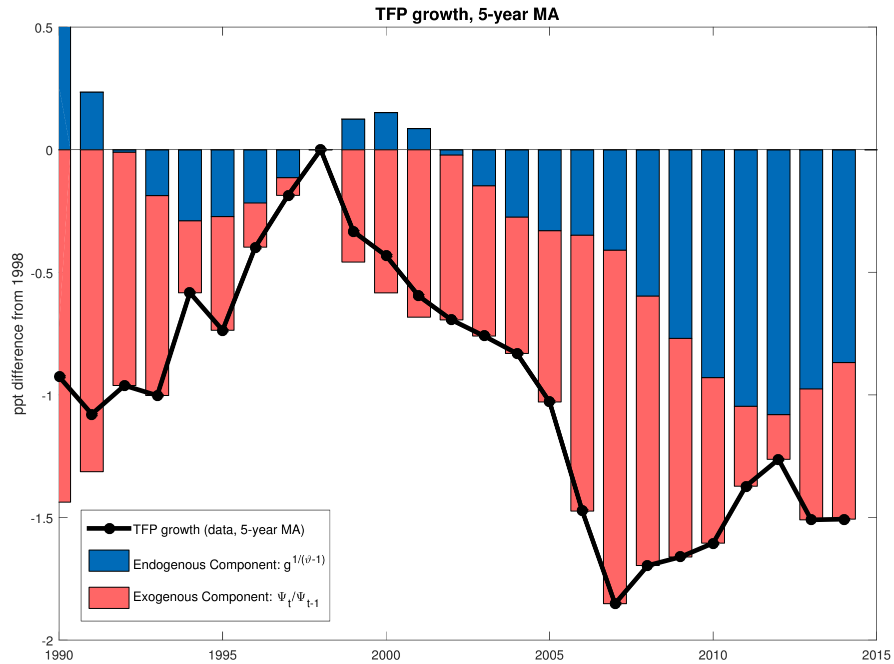

Note: The black solid line is the 5-year moving average of log-differenced TFP in the data. The blue bars indicate the contribution of the endogenous component, and the red bars show the contribution of the exogenous component.

As shown in the right and bottom panels of Figure 17 the consequences for TFP of the alternative path of R&D are substantial: the moving-average measure of TFP growth now rises gradually starting in 2009, and by the end of the sample reaches about 1.1%—nearly double the actual value, and recovering a substantial part of the decline seen since the 1998 peak. This alternative evolution has sizable implications for the level of TFP, which by 2016 is about four percent higher in the counterfactual scenario relative to its actual path. Thus, even if the TFP slowdown began prior to the Great Recession—as compellingly argued by Fernald (2014)—the analysis above suggests that the slowdown in R&D since the crisis has significantly contributed to the low TFP growth rates seen in recent years.

5.3 Effects of monetary policy

We next analyze some implications for monetary policy in this setting. We begin by analyzing the effects of a monetary policy shock. In Figure 19 we report the effects of a rise of 25 basis points in the policy rate in the baseline model with endogenous TFP. For comparison, we also include the effect of a rise in the policy rate of the same size in a model where TFP is fully exogenous. From the top row, note that the policy rate increase adversely

Figure 17: R&D and TFP, actual and counterfactual

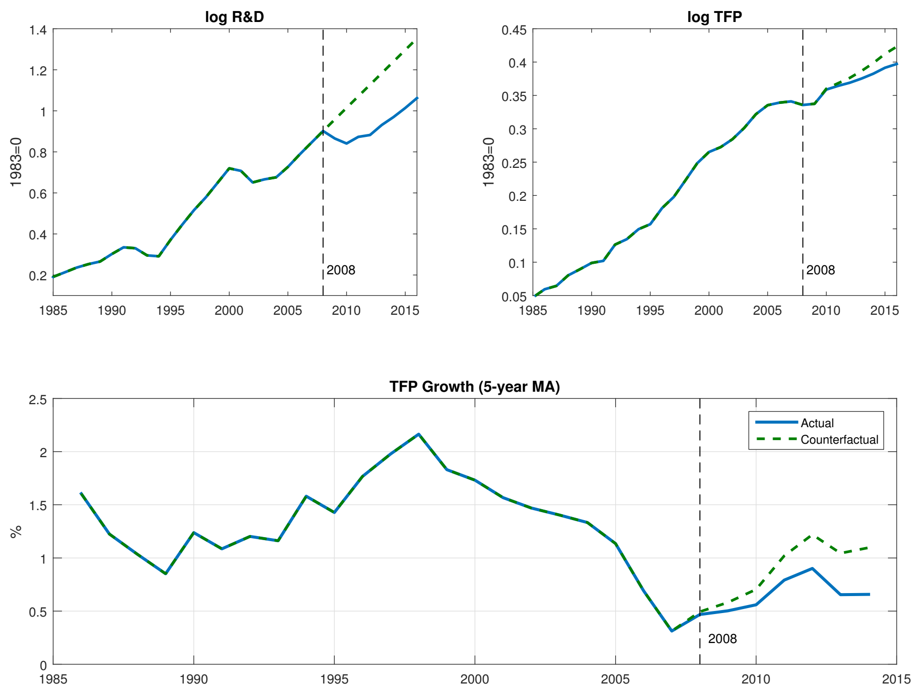

Note: The blue solid lines show the actual evolution of R&D, TFP, and TFP growth (5-year MA) in the data. The green dashed lines show a counterfactual scenario in which R&D post-2008 remains on its pre-crisis trend.

Figure 18: Innovations to Δ_t^n and R&D growth, actual and counterfactual

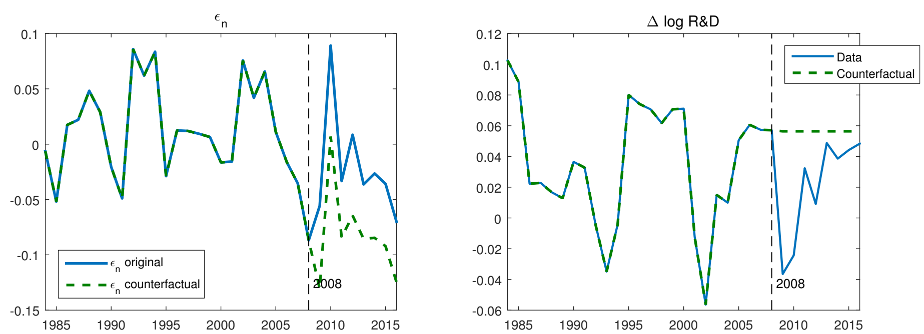

Note: In the right panel, the blue solid line shows the innovations to Δ_t^n recovered from the historical analysis, and the green dashed line represents the counterfactual innovations that are required to keep $R\&D$ on trend. The right panel shows log-differenced R&D in the data (blue solid) and in the counterfactual (green dashed).

affects both the stock of innovations and the adoption rate. R&D falls by about 0.5 percent on impact, and adoption effort M_t falls about half as much. The reason is twofold: first, the present value of future profits diminishes due to discounting. Second, the profit flow Π_t itself falls (albeit in a transitory fashion), as a consequence of the cyclical downturn engineered by the shock. Due to the slower pace of innovation and adoption, TFP falls persistently relative to the unshocked path. This then implies that the decline in macroeconomic aggregates like output, consumption and investment has a persistent component that is absent in the exogenous-growth version of the model.

Figure 19: Monetary policy shock, baseline v. model without endogenous TFP

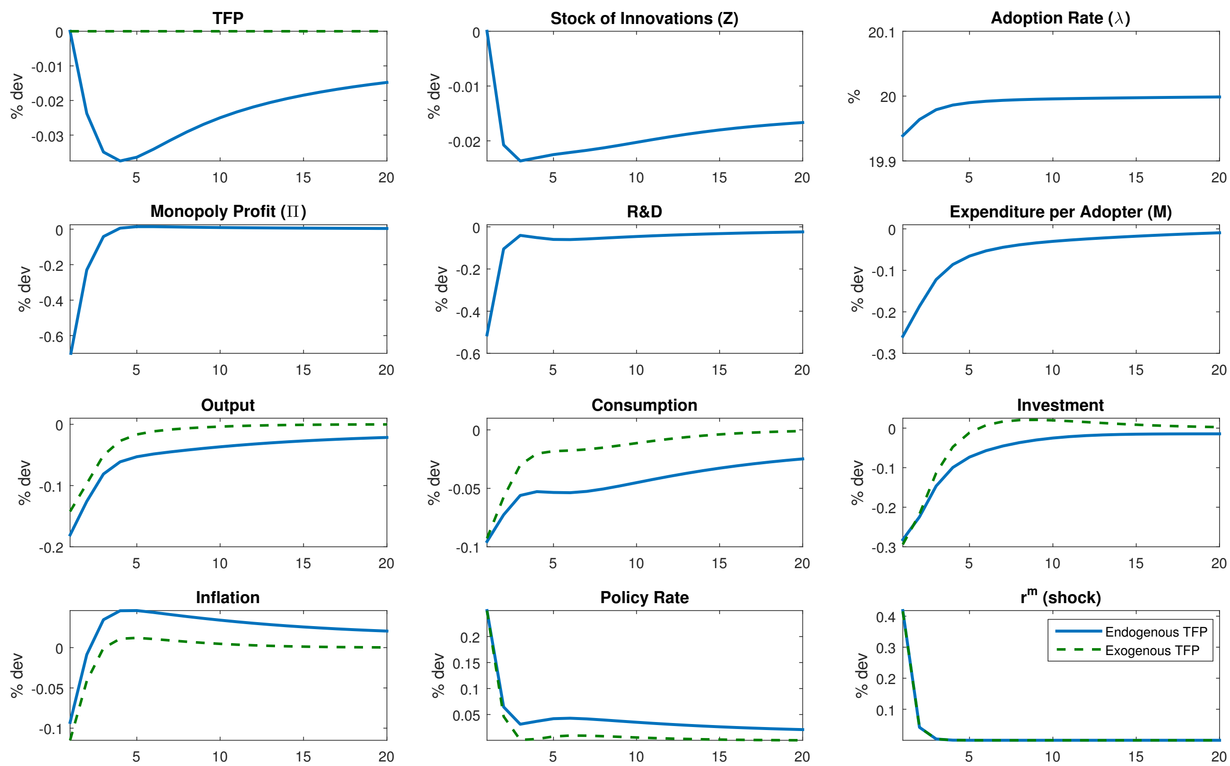

Note: We show the impulse responses to a monetary shock r_t^m in the baseline model (blue solid line) and in a model without the endogenous TFP growth mechanism. We size the impulse so that it induces a rise in the monetary policy rate of 25 basis points.

Given the effects of monetary policy on TFP just documented, how much can future monetary policy boost TFP growth within this setting? We next use the model to illustrate the consequences for future TFP growth of the pace of monetary policy tightening post-2016. In particular, we consider the following experiment. Suppose a baseline scenario in which the policy rate post-2016 is expected to follow the path shown by the blue line in the left panel of Figure 20, taken from the projected appropriate policy path by the FOMC as of March 15, 2017. Suppose also that in this baseline scenario, agents expect TFP growth to remain constant and at its average pace in the period 2011-2016 (of about 0.65 percent). We then consider an alternative scenario in which monetary policy tightens more slowly, as shown by the green dash-dotted line in Figure 20). The alternative path has the policy rate below the baseline projection by about 100 basis points, on average, from 2017 through 2021. We implement this alternative path using a sequence of shocks to the monetary rule, impacting from 2017 through 2020.

The left panel of Figure 20 shows the implications for TFP growth, again measured as a 5-year moving average. The monetary stimulus leads to a temporary boost to TFP growth, which is above the baseline path between 2017 and 2023. At its peak in 2020 TFP growth reaches 0.74 percent, about 10 basis points above the baseline. The effect of the stimulus dies out thereafter, with TFP growth returning to its baseline path.

Figure 20: Policy rate and TFP growth, baseline projection and alternative

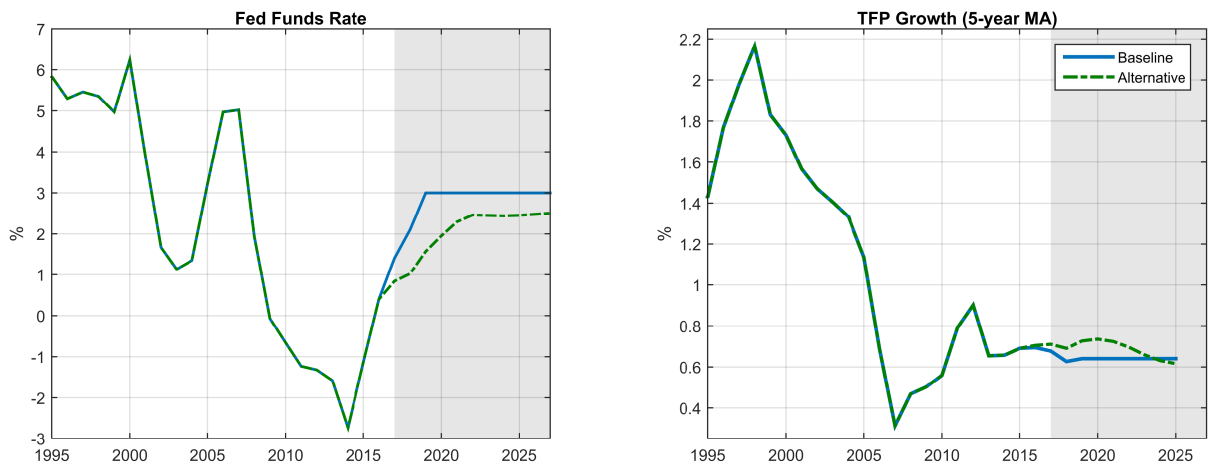

Note: The blue line shows the baseline projected Fed funds rate and TFP growth rate. The green dashed line shows an alternative scenario where the policy tightening post-2016 occurs much more slowly than in the baseline.

6 Conclusion

In this paper, we estimate the impact of R&D movements on TFP in the U.S. and in a panel of advanced economies, and we develop a model featuring endogenous TFP via technology innovation and adoption to address the evidence. We also use the model to shed light on the drivers of the productivity growth slowdown of recent times.

One notable absence from our analysis is a consideration of the implications of the zero lower bound (ZLB) on monetary policy. This is particularly relevant given that the slow productivity growth seen across advanced economies has coincided with a period in which several of them were constrained by the ZLB, suggestive of the possibility of "stagnation traps" (Benigno and Fornaro (2016)). Incorporating such traps into more quantitativelyoriented frameworks such as the one develop here is a promising area of future work. Another interesting area for future research, in light of the findings in section 2.4, is a more thorough analysis of the interaction between stock prices, R&D, and subsequent TFP developments.

Appendix

A Data

Our dataset consists of annual data for 22 advanced economies. Our panel includes the same countries as Coe et al. (2009), with just two exceptions due to data availability. We exclude Greece and Iceland due to missing R&D data. Table A1 in the Appendix contains a complete list of the countries and years included in our panel, as well as basic summary statistics.

R&D is measured as R&D expenditure performed by business enterprise, in millions of constant US dollars, converted using constant PPPs. Data comes from the OECD Research and Development Statistics. For the United States, this series is equivalent to R&D data published by the National Science Foundation. The OECD data is extracted from the NSF's Science and Engineering Indicators data on R&D, performed in the domestic United States by all companies with five or more employees, publicly or privately held.[15]

TFP comes from two separate data sources, based on availability. Our primary source is The Long-Term Productivity database published by Bergeaud et al. (2015). This data is available for 17 advanced economies including the United States. For the five remaining countries in our panel, we use TFP data from the Total Economy Database produced by The Conference Board. These series are augmented with Information and Communications Technology (ICT) and Labor Quality. These series are used for Austria, Ireland, Israel, New Zealand, and South Korea.

Gross Domestic Product is measured in millions of constant US dollars, converted using Geary Khamis PPPs. This data is from the Total Economy Database produced by The Conference Board.

Finally, our country-level stock price indexes come from MSCI Inc., formerly Morgan Stanley Capital International. We use end of period stock prices in real per capita terms,[16] following the lead of Beaudry et al. (2011). Adjustment is performed using GDP deflator and population series from the World Development Indicators published by the World Bank.

[15]See the NSF's Science and Engineering Indicators 2016, Appendix Table 4-2.

[16]The indexes are converted into per capita terms by subtracting the log population growth rate from the log growth in prices: $\log P_t/P_{t-1} - \log Pop_t/Pop_{t-1}$.

Table A1: Summary Statistics

Country	**Sample**	**Observations**	**Mean ΔTFP**
Australia	1981 - 2011	31	1.25%
Austria	1989 - 2013	25	2.14%
Belgium	1981 - 2013	33	1.86%
Canada	1981 - 2013	33	1.39%
Denmark	1981 - 2013	33	1.89%
Finland	1981 - 2013	33	2.61%
France	1981 - 2013	33	1.85%
Germany	1981 - 2013	33	2.79%
Ireland	1989 - 2012	24	1.37%
Israel	1991 - 2013	23	4.26%
Italy	1981 - 2013	33	0.86%
Japan	1981 - 2013	33	3.05%
Netherlands	1981 - 2013	33	1.57%
New Zealand	1989 - 2011	23	0.56%
Norway	1981 - 2013	33	1.46%
Portugal	1982 - 2013	32	0.47%
South Korea	1995 - 2013	19	2.92%
Spain	1981 - 2013	33	0.82%
Sweden	1981 - 2013	33	3.24%
Switzerland	1981 - 2012	32	3.00%
United Kingdom	1981 - 2013	33	1.75%
United States	1953 - 2015	63	2.48%
Full Sample		701	1.16%

References

Akinci, Ozge, "Global financial conditions, country spreads and macroeconomic fluctuations in emerging countries," *Journal of International Economics*, 2013. 2.2

Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez, "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence," 2016. 2, 1, 3, 3.4, 4.1, 4.2, 13

Beaudry, Paul and Franck Portier, "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 2006, *96* (4), 1293–1307. 1, 2.4

_ , Deokwoo Nam, and Jian Wang, "Do mood swings drive business cycles and is it rational?," Technical Report, National Bureau of Economic Research 2011. A

Benigno, Gianluca and Luca Fornaro, "Stagnation traps," 2016. 2, 6

Bergeaud, Antonin, Gilbert Cette, and Remy Lecat, "Long-run per-capita gross domestic product in France and the advanced countries: the role of productivity and labour," *Economie & Statistique*, 2015. A

Bianchi, Francesco and Howard Kung, "Growth, Slowdowns, and Recoveries," *National Bureau of Economic Research*, 2014. 2

Caballero, Ricardo J and Adam B Jaffe, "How High are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth," in "NBER Macroeconomics Annual 1993, Volume 8," MIT press, 1993, pp. 15–86. 4.1

Cerra, Valerie and Sweta Chaman Saxena, "Growth Dynamics: the Myth of Economic Recovery," *American Economic Review*, 2008, *98* (1), 439–457. 2.2

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45. 3

Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt, "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 110–67. 3.2, 3.4, 14

Coe, David T and Elhanan Helpman, "International R&D spillovers," *European Economic Review*, 1995, *39* (5), 859–887. 1, 2.3

— , — , and Alexander W Hoffmaister, "International R&D Spillovers and Institutions," *European Economic Review*, 2009, *53* (7), 723–741. A

Comin, Diego and Mark Gertler, "Medium-Term Business Cycles," *American Economic Review*, 2006, pp. 523–551. 1, 1, 2.1, 3, 3.2, 3.2, 4.1, 4.2, 13

— , — , and Ana Maria Santacreu, "Technology Innovation and Diffusion as Sources of Output and Asset Price Fluctuations," 2009. 3.7

Eaton, Jonathan and Samuel Kortum, "Trade in ideas Patenting and productivity in the OECD," *Journal of international Economics*, 1996, *40* (3), 251–278. 2.3

Fernald, John, "Productivity and Potential Output Before, During, and After the Great Recession," in "NBER Macroeconomics Annual 2014, Volume 29," University of Chicago Press, 2014. 1, 5.2

Griffith, Rachel, Stephen Redding, and John Van Reenen, "Mapping the two faces of R&D: Productivity growth in a panel of OECD industries," *Review of economics and statistics*, 2004, *86* (4), 883–895. 9

Guerron-Quintana, Pablo and Ryo Jinnai, "Liquidity, Trends, and the Great Recession," 2014. 2

Kung, Howard and Lukas Schmid, "Innovation, growth, and asset prices," *The Journal of Finance*, 2015, *70* (3), 1001–1037. 2, 3.7, 11

Nickell, Stephen, "Biases in dynamic models with fixed effects," *Econometrica: Journal of the Econometric Society*, 1981, pp. 1417–1426. 5

Primiceri, Giorgio E, Ernst Schaumburg, and Andrea Tambalotti, "Intertemporal disturbances," 2006. 4.1

Queralto, Albert, "A Quantitative Model of Slow Recoveries from Financial Crises," *International Finance Discussion Papers, Board of Governors*, 2013. 2, 7

Romer, Paul, "Endogenous Technological Change," *Journal of Political Economy*, 1990, pp. S71–S102. 3.2

Smets, Frank and Rafael Wouters, "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *The American Economic Review*, 2007, *97* (3), 586–606. 3.4, 4.1

Uribe, Martín and Vivian Z Yue, "Country Spreads and Emerging Countries: Who drives Whom?," *Journal of International Economics*, 2006, *69* (1), 6–36. 2.2

Wu, Jing Cynthia and Fan Dora Xia, "Measuring the macroeconomic impact of monetary policy at the zero lower bound," *Journal of Money, Credit and Banking*, 2016, *48* (2-3), 253–291. 2.4

Yellen, Janet, "Current Conditions and the Outlook for the U.S. Economy," 2016. The World Affairs Council of Philadelphia, 6 June 2016, Speech. 1