Business Cycles, R&D, and Hysteresis: An Empirical Investigation

Filippo Massari*
Department of Economics
Fairfield University

Hedieh Shadmani
Department of Economics
Fairfield University

May 10, 2024

Abstract

This paper investigates the permanent effect on total factor productivity (TFP) of temporary shocks. We estimate a structural vector autoregression to test the predictions of endogenous growth models over the business cycle. According to theory, the stock of technological knowledge promotes its flow as researchers "stand on the shoulders of giants." Therefore, if R&D investment is pro-cyclical — as data show and theory predicts — a recession leads to a temporary deviation of the R&D level from its trend, thus reducing new knowledge creation. The consequent technological stock loss sets the economy on a parallel but permanently lower trend. The results are in line with the main theoretical prediction. Specifically, the US economy loses approximately 1.5% in TFP following an increase in cyclical unemployment that peaks at 1 percentage point above mean. The historical variance decomposition shows a particularly strong positive effect during the boom of the late

*Dolan School of Business, Fairfield University, 1073 North Benson Road, Fairfield, CT 06824. E-mail: fmassari@fairfield.edu. Website: http://filippomassari.com/.

'60s, and particularly strong negative effects around the Volcker disinflation period and the Great Recession. Finally, we estimate the effects on R&D of an exogenous increase in TFP to discriminate between various theories. Our results are consistent with models where financial frictions or nominal rigidities drive R&D's pro-cyclicality.

JEL: E22, E32, O32, O47.

Keywords: Endogenous Growth, Hysteresis, R&D.

1 Introduction

Macroeconomics is traditionally divided into two main research areas: business cycles and long-run growth. This distinction, although useful for practical purposes, is artificial and assumes away any connection between fluctuations in economic activity and its trend. Yet, the disappointing performance of the US economy in the aftermath of the Great Recession generated a widespread belief that the two phenomena were connected.

The economics profession has not ignored the topic altogether. According to endogenous growth theory, the *growth rate* of aggregate productivity is a function of the *level* of R&D investment. This proposition becomes relevant in light of the observation that R&D investment is notoriously procyclical. Therefore, if a shock that temporarily displaces output from its trend has a similar temporary effect on the R&D level, the aggregate productivity growth rate will be temporarily affected. If this change in aggregate productivity growth is not compensated later by an equal change in the opposite direction, the productivity level is permanently changed, a phenomenon known as hysteresis.[1]

Economists credit Schumpeter (1942) for being the first scholar to discuss how fluctuations in demand could affect the rate of innovation and, thus, long-term output growth, a view that is formally and concisely illustrated in Aghion and Saint-Paul (1998). He theorized a counter-cyclicality of R&D spending due to the decreased opportunity cost of reallocating resources

[1]See Bond-Smith (2019) for a recent survey of the endogenous growth literature, and Cerra et al. (2023) for a review of the literature on hysteresis.

from production to R&D during recessionary periods. During recessions, low demand implies a low marginal profitability of production, thus increasing the attractiveness of devoting those resources to R&D. The subsequent realization that in the data R&D is instead pro-cyclical promoted a noteworthy effort to produce alternative theories. The result is a long list of influential papers that propose various explanations for the persistence of cyclical fluctuations or hysterisis through this channel (Shleifer, 1986; Fatas, 2000; Comin and Gertler, 2006; Barlevy, 2007; Francois and Lloyd-Ellis, 2009; Nuño, 2011; Aghion et al., 2012; Fatás and Summers, 2018; Bianchi et al., 2019; Mand, 2019; Queralto, 2019; Anzoategui et al., 2019).

On the empirical side, the magnitude and the shape of the response of R&D to shocks that cause the business cycle are unknown. Therefore, our main goals consist in, first, identifying the path of R&D and productivity following a shock in economic activity. Second, we estimate the size of the effect, with the purpose of understanding whether this mechanism is of quantitative relevance to the point that it merits further attention.

This paper achieves these goals by computing impulse response functions for R&D and aggregate productivity to shocks that caused the US business cycle in the post-WWII period. To do that, we employ a structural VAR to test the theoretical predictions and assess their quantitative relevance. The VAR includes three variables: R&D spending, utilization-adjusted total factor productivity (TFP), and cyclical unemployment. We impose a recursive structure in which utilization-adjusted TFP does not respond on impact to R&D and unemployment, and unemployment does not react on impact to R&D. We use the forecast error of the cyclical unemployment time series as a proxy for shocks that cause a deviation of output from its trend. Responses to an exogenous increase in TFP, instead, help us in checking the validity of different theories on R&D's cyclicality.

The results support the hysteresis hypothesis. They show a pattern of R&D and TFP following an exogenous increase in cyclical unemployment consistent with the theoretical predictions. Namely, R&D spending responds pro-cyclically. Moreover, TFP decreases gradually before stabilizing at a permanently lower trend. The magnitude of this effect is remarkable: for an increase in cyclical unemployment that peaks at 1 percentage point above its sample average, R&D growth remains considerably below trend for almost 2 years, and the permanent loss in TFP amounts to approximately 1.5%.

Interestingly, our empirical work rescues Schumpeter's original hypothesis on the counter-cyclicality of R&D. Following a positive exogenous increase in TFP, R&D declines. This result is useful for our purposes, as it allows us to rule out some explanations for the pro-cyclicality of R&D. In fact, models that achieve this pro-cyclicality through the presence of physical capital in the R&D technology or the pro-cyclicality of expected profits predict an increase in R&D following a TFP shock. Instead, theories that emphasize credit constraints as the element responsible for R&D's pro-cyclicality, such as Aghion et al. (2012), are in line with our findings. They predict counter-cyclical R&D in case of a TFP shock, and pro-cyclical R&D in the case of other shocks that affect firms' access to liquidity. Theories whose key transmission mechanism for the cyclicality of R&D relies on labor market dynamics (Aysun, 2020) are also in line with our findings: a positive TFP shock reduces employment in the short run, thus leading to a positive comovement between employment and R&D.

Empirical work on the connection between business cycle and long-run growth abounds. Cerra et al. (2023) provide a comprehensive literature review going back to the early '80s. Whereas early work was concerned with employing statistical techniques to separate trend and deviations, some more recent work is closer to our analysis as it focuses more heavily on specific transmission mechanisms. Some of it studies single events, typically financial and banking crises, see for example Cerra and Saxena (2008). On the mechanism that we consider, notable examples include Aghion et al. (2012), Ouyang (2011), and Duval et al. (2020). All these studies analyze firm level data, finding that firms that enter a financial crisis with weak balance sheet reduces their innovation effort after the shock.

In general, the idea that temporary negative shocks can cause persistent, or even permanent, output losses is already corroborated by empirical evidence. The idea that one of the transmission channels could be R&D is backed by evidence too. What remains uncertain is the size of this effect, which may vary from country to country, the shape of impulse response functions for aggregate variables, and the reason why R&D evolves pro-cyclically.

Our paper adds to this literature by providing those results for the US. It differs from studies on R&D's pro-cyclicality because we do not focus on the specific event of the Great Recession, but we take a longer-term approach by exploiting all data available since the post-WWII period. In this way, we point out periods in US history where this hysteresis effect has been relevant. It also differs from studies that focus on firm and industry level effects because through our SVAR approach we are able to estimate the magnitude of the long-run effect of shocks on productivity, and provide impulse response functions for variables that are useful for assessing and guiding theoretical work. Importantly, our analysis provides estimates of the effect of temporary shocks on aggregate variables, which are of great interest to macroeconomists and policy-makers.

Furthermore, by distinguishing between two different sources of shocks, our approach allows us to distinguish between different transmission channels. In this way, we can determine which theory on the pro-cyclicality of R&D provides predictions that align with empirical evidence.

The paper is organized as follows. Section 2 illustrates the theory, section 3 discusses the data, section 4 introduces the empirical approach, section 5 shows the results, section 6 illustrates robustness tests, and section 7 concludes.

2 From Short Run Shocks to Permanent Productivity Changes: An Illustration of the Theory

This section presents the relevant features of the theory that we test. In outlining the theory, we focus on elements pertinent to our empirical purposes, namely the production and investment side of the economy. We will only mention the theoretical ingredients responsible for driving the R&D cyclical behavior we investigate while linking to the works describing those models in detail.

2.1 The Effects of Pro-Cyclical R&D

Economists have been focusing on R&D to understand long-run growth for decades. Endogenous growth theory is an outcome of their efforts. The earlier endogenous growth models, like Romer (1990), rely on the aggregate production function:

$$Y_t = f(A_t, L_t), \tag{1}$$

where A denotes technological knowledge, and L is labor. The production function is increasing in its inputs.

The critical feature of endogenous growth theory is the R&D technology. The most straightforward specification in discrete time is the following:

$$A_{t+1} - A_t = h(A_t, R_t), \tag{2}$$

where R_t is an endogenous variable that denotes the total of resources devoted to R&D, which in this specification consists of labor from scientists and engineers. The R&D technology is assumed to be separable, increasing in its inputs, and linear in A_t. The linear dependence of technological knowledge on the existing stock of knowledge drives endogenous growth. Given these assumptions, we can re-express the equation as follows by dividing both sides by A_t:

$$g^A_{t+1} = h(R_t). \tag{3}$$

That is, the technological knowledge *growth rate* depends on the R&D *level*.[2] In the absence of population growth (or after introducing elements that sterilize the scale effect), the theory predicts a stationary R_t, thus leading to constant technological knowledge growth that depends on an endogenous variable.

Words of caution are needed when analyzing the relationship between these aggregate variables. Endogenous growth theory has made progress ever since the original Romer's contribution. Specifically, other variables stand in the way of aggregate R&D in determining the path of aggregate productivity growth. Some examples include the average firm size (Peretto and Connolly, 2007), the degree of industry concentration (Ghazi, 2019), government R&D spending (Huang et al., 2023), and entry/exit and churning within the firm size distribution (Massari, 2023). Some of these factors, such as government R&D spending, are irrelevant when studying the business cycle, because they do not display cyclical behavior. Other factors vary along the business cycle, but we nevertheless decide to exclude them for three main reasons.

First, data availability is an issue, as variables such as number of firms

[2] The careful reader will have noticed a scale effect, namely the dependence of knowledge growth on the population level, which is part of what constitutes R&D effort. We present the simplest theoretical outline that drives the relevant predictions for our empirical endeavor. Bond-Smith (2019) illustrates how the literature has modified the model to eliminate the scale effect while preserving endogenous growth in a way that is consistent with some empirical regularities. These modifications preserve the essence of the relevant mechanism for this paper.

exist only at an annual frequency from the late '70s onward. As one of our goals is to understand whether hysteresis is present and quantitatively important, we believe that the better way to proceed involves exploiting the whole sample of R&D and TFP and the highest frequency of data that is available.

Second, another goal is to understand whether R&D and productivity's comovement in response to shocks that cause the business cycle is consistent with theoretical predictions, which concern the behavior of aggregate variables. There is room for introducing additional features to these models, but we believe that it is the responsibility of theory to point out how additional elements that affect the R&D technology play a role in reinforcing, weakening, or preventing hysteresis.

Third, as we will describe later, we do not impose equation (3) in our empirical analysis. We use it as guidance when selecting variables and when imposing restrictions that would not be violated under any alternative specification of the relationship between productivity and R&D. Instead, the empirical approach allows for the positive relationship between R&D and productivity growth to be violated.

From this framework, it is easy to notice that short-run shocks can have long-run consequences if R&D effort (or resources) temporarily deviates from its trend. Indeed, this temporary deviation of R&D from its steady state can permanently change the aggregate productivity level. Figure 1 shows the long-run effect on productivity after a shock that reduces output and under the assumption that R&D effort is pro-cyclical. It is important to keep in mind the distinction between the behavior of R&D effort and R&D spending: while R&D effort reverts to its pre-shock trend, R&D spending does not. The reason why R&D spending is permanently affected is that wages paid to R&D personnel decrease after the reduction in productivity. However, the steady-state amount of resources devoted to R&D is unchanged, ensuring that the productivity growth rate reverts to its pre-shock level after the shock has fully propagated.

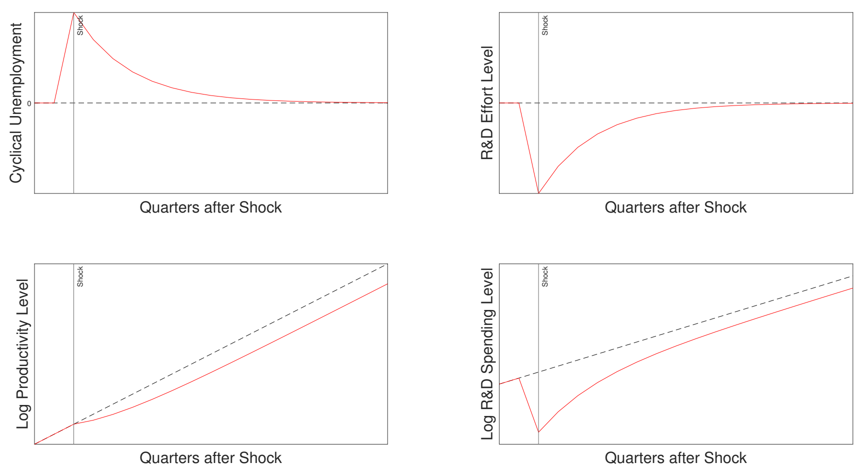

Figure 1: Theoretical Predictions Following a Negative Shock.

Note: Impulse Responses to a shock that increases cyclical unemployment rate. R&D is assumed pro-cyclical, while productivity moves according to the equations illustrated. The black line is the pre-shock trend.

2.2 Why Is R&D Pro-Cyclical?

As R&D is an endogenous variable, what does the theory predict regarding its behavior over the business cycle? Economists within this literature credit Schumpeter (1942) as the first person to discuss R&D in a business cycle context. In his view, R&D would move counter-cyclically because the opportunity cost of allocating resources to R&D increases during booms, when production is particularly profitable. As data became available, realizing that aggregate R&D spending and employment correlate positively with output led economists to reject this explanation and hunt for different mechanisms to deliver this pro-cyclicality.

The current state of the literature is one where different explanations coexist. Some of these explanations still emphasize the pro-cyclicality of profits as the driver of the cyclical behavior of R&D. However, they add some features to ensure that R&D is pro-cyclical. For example, Barlevy (2007) assumes that imitation leads to the dissipation of profits derived from R&D after some time. Therefore, firms become short-term oriented and invest in R&D when profitability increases, which happens during booms. Francois and Lloyd-Ellis (2009), picking up insights from Shleifer (1986), introduce lags between R&D, commercialization, and implementation of the innovation. As firms would find it profitable to implement the innovation during business cycle peaks, efforts in commercialization will be stronger in worse times, and R&D activities will precede them, thus tending to coincide with booms. Anzoategui et al. (2019) combine the pro-cyclicality of profits with stickiness in R&D personnel's wages. As a result, the cost of conducting R&D responds with some lag to shocks that cause the business cycle.

Other explanations rely on the idea that R&D does not consists exclusively of spending on scientists and engineers. Comin and Gertler (2006) model R&D as lab equipment, whereas Mand (2019) constructs a specification that relies on the complementarity between scientists and engineers, physical capital, and supporting staff. As capital goods are more abundant during booms, this abundance facilitates R&D.

An additional explanation relies on the inclusion of financial frictions. By assuming that firms rely on credit to finance R&D, shocks that tighten financial conditions limit their possibility to conduct this type of investment. To the extent that credit availability is pro-cyclical, and if this force is stronger than the forces that drive counter-cyclical R&D, these models predict pro-cyclical R&D. Examples of this explanation include Aghion et al. (2012), Queralto (2019), and Bianchi et al. (2019).

Finally, Aysun (2020) has recently proposed an endogenous growth model with New-Keynesian features. This theory requires labor intensive R&D and diminishing returns to labor in production. As a result, shocks that lead to higher equilibrium employment would lead to higher production employment. Due to diminishing returns to production labor, R&D employment would also rise to balance out marginal products. Instead, following a productivity shock, the model can deliver counter-cyclical R&D if labor hours or employment move in the opposite direction as productivity. Theory allows for this possibility in models with nominal rigidities or in models with particuarly strong income effects (Basu et al., 2006). Empirical studies confirm that following a positive total factor productivity shock, employment and hours decline (Galí, 1999; Basu et al., 2006; Li, 2022).

3 Data

The analysis uses US quarterly data for R&D, total factor productivity (TFP), and unemployment. The period under consideration starts in quarter 1 of 1949 and ends in quarter 4 of 2019.

The data series for R&D is the real and seasonally adjusted gross private domestic investment in intellectual property products from the National Income and Product Accounts (Bureau of Economic Analysis). We focus on private R&D because the theory predicts R&D movements in response to cyclical shocks through market effects. We tried adding public R&D spending to the analysis, concluding that it does not add anything of interest. Next, how can we be sure that the cyclicality of R&D results from changes in resources devoted to it and not to its price? We looked for a measure of workers' salaries with a bachelor's degree, which exists at a quarterly frequency starting from 2000. At an annual frequency, it dates back to 1979. As salaries of high-skilled workers do not display much cyclicality, we do not think that they are the driving force behind the cyclicality of R&D.

When considering productivity, we rely on Fernald's series on utilization-adjusted total factor productivity (Fernald, 2014), which removes variations in productivity due to changes in labor effort and capital usage. We prefer the utilization-adjusted measure because the ultimate goal of this analysis is to understand the long-run trend of productivity. Removing the utilization of other factors of production provides us with a better estimate of the product of R&D. In fact, we think of utilization-adjusted total factor productivity as the best proxy of productive ideas in circulation.

Furthermore, to better isolate changes in the stock of technological knowledge, we smooth the TFP time series by employing a 4-period moving average. We follow Comin and Gertler (2006) rationale in removing the higher frequency component in the TFP time series to isolate what they define as the medium term cycle and produce a less noisy time series. In fact, high frequency fluctuations in TFP are unlikely to reflect rapid increases and reductions in the stock of ideas. They may instead reflect changes in efficiency of input use, for example as a result of resource reallocation across firms. For this reason, an analysis with annual data may be more appropriate for our purposes, but it would reduce the number of observations and the higher frequency fluctuations in R&D. We therefore decided to focus on quarterly variables, while smoothing out quarterly fluctuations in TFP. We note that our results do not change by choosing an 8-period moving average. However, because of this smoothing, we lose a few observations and our analysis focuses on the time period from the second quarter of 1949 to the second quarter of 2019. After smoothing utilization adjusted TFP growth, its correlation with R&D growth a few quarters before increases.

Finally, we need a variable that moves reliably due to shocks that cause macroeconomic fluctuations at business cycle frequency. We choose the cyclical unemployment rate, computed as the difference between the unemployment rate provided by the Bureau of Labor Statistics and the natural rate of unemployment from the Congressional Budget Office. Cyclical unemployment is a useful measure for our purposes because it increases following any slowdown in economic activity, regardless of the reason why economic activity slows down. Therefore, we intend exogenous deviations in cyclical unemployment from its mean as a proxy for shocks that cause aggregate fluctuations. We do not distinguish between the different type of shocks as our main goal is to understand whether hysteresis thorugh the R&D channel is a meaningful research focus. We leave the identification of specific shocks that cause hysteresis for future research.

An apparently obvious choice to capture the business cycle is to use a measure of the output gap. However, the concept of output gap is tricky for someone who takes this paper's theoretical premises seriously. Specifically, what is potential TFP? This paper argues that it varies over time as fluctuations impact R&D spending, which affects TFP. Measures of potential output, such as the one produced by the Congressional Budget Office, are sensitive to assumptions on the evolution of TFP, which we consider endogenous in this paper. We note, however, that the correlation between cyclical unemployment and the output gap is 0.83.

4 Empirical design

This paper aims to understand how R&D and aggregate productivity evolve over the business cycle. The ultimate question is whether the causes of these fluctuations permanently affect the productivity level. To do that, we need to isolate shocks that drive temporary deviations of economic activity from its trend, where the trend could be endogenously affected by these shocks as well.

We rely on a structural Vector Autoregression (SVAR) to compute impulse responses to exogenous increases in cyclical unemployment and TFP. To allow for the posibility of permanent effects from transitory shocks, TFP and R&D enter the VAR in log differences. We then cumulate their impulse responses to obtain their percent deviation from their pre-shock trend. The structure we impose on the VAR consists in ruling out the simultaneous impact of some variables on others. Specifically, we impose a recursive structure such that productivity does not respond on impact to R&D and unemployment, and unemployment does not respond on impact to R&D.

The first assumption is justified by the idea that it takes time before R&D spending is translated into higher productivity. A substantial literature documents this fact, pointing out that R&D spending is one of the first steps in the innovation process — indeed, Francois and Lloyd-Ellis (2009) present this evidence as a motivation for their model that we illustrated above. A simple correlation between R&D and TFP growth supports this assumption, as the contemporaneous correlation is 0, while it peaks at a 4 quarters lag for TFP.

The second restriction consists in assuming no impact of an exogenous change in unemployment on TFP. Our reasoning is that we want to isolate as much as possible true technological change from other factors that could drive TFP dynamics. Ruling out a simultaneous causal link from unemployment to productivity can remove the cyclical shocks' effect on productivity through channels that differ from those under investigation, such as reallocating resources across firms with different productivity levels or the layoff of workers with a lower marginal productivity.

Finally, we justify the restriction on the simultaneous impact of R&D on unemployment as a way to force the unemployment rate to capture the shocks that cause the business cycle. Both R&D and unemployment react to the same shocks, and unemployment is notoriously a lagging indicator of the business cycle — although the use of quarterly data weakens this last claim. Therefore, we impose a structure that leads the exogenous deviation of cyclical unemployment from its average to pick up these shocks. Assuming that an exogenous change in R&D does not simultaneously affect unemployment, we rule out that the error term associated with the R&D time series picks up such shock. However, we check for robustness by inverting the order of these two variables — thereby assuming that a change in cyclical unemployment does not affect R&D on impact —, finding that our conclusions remain unaltered.

Formally, our empirical design relies on the following procedure illustrated in Breitung (2001). We establish the VAR model in its reduced form:

$$Z_t = B_1 Z_{t-1} + B_2 Z_{t-2} + ... + B_p Z_{t-p} + e_t, \tag{4}$$

where Z_t is a 3×1 vector of time series observations for total factor productivity, the cyclical unemployment, and research and development spending. $B_1, B_2, \ldots, B_p$ are the coefficient matrices for the lagged dependent variables.

Corresponding to this reduced form, we define a structural model as:

$$De_t \equiv S\epsilon_t. \tag{5}$$

The matrices D and S are assumed to be invertible. ϵ is a 3×1 vector of shocks. Importantly for our purposes, its covariance matrix $\mathbb{E}(\epsilon\epsilon') = \Omega$ is diagonal. This restriction removes the correlation between error terms. We are therefore imposing that exogenous changes in unemployment are uncorrelated with exogenous changes in TFP growth and R&D growth. Consequently, the unemployment shock picks up all shocks that affect the cyclical unemployment rate but are not caused by technology nor R&D. If the hysteresis hypothesis was incorrect, the shock would not cause any permanent effect.

Additionally, we impose the recursive structure illustrated at the beginning of this section on $D^{-1}S$. This structure imposes the following restrictions:

$$\begin{bmatrix} e_{\Delta TFP,t} \\ e_{U,t} \\ e_{\Delta R\&D,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ \varphi_{21} & 1 & 0 \\ \varphi_{31} & \varphi_{32} & 1 \end{bmatrix} \times \begin{bmatrix} \epsilon_{\Delta TFP,t} \\ \epsilon_{U,t} \\ \epsilon_{\Delta R\&D,t} \end{bmatrix} \tag{6}$$

We can now use $e_t \equiv D^{-1}S\epsilon_t$ to rewrite the model as:

$$DZ_t = B_1^* Z_{t-1} + B_2^* Z_{t-2} + ... + B_p^* Z_{t-p} + S\epsilon_t, \tag{7}$$

where B_t^*, for $t = 1, 2, ..., p$ are the structural coefficients that may differ from the reduced form coefficients. We use Hannan-Quinn criterion (HQC) for selecting the optimum lag length in our estimated VAR model, which is 5. Other criteria give us an optimal lag length that we consider either too high or too low, thus causing respectively a loss in power and a loss in interdependence between the variables over longer horizons. We nevertheless run the model with alternative lag lengths and discuss it in section 6.

To analyze the dynamic effects of structural shocks, we employ the following moving average representation:

$$Z_t = \Phi(L)\epsilon_t, \tag{8}$$

where L is the lag and $\Phi(L) = \Phi_0 + \Phi_1 L + \Phi_2 L^2 + ... = B(L)^{-1}D^{-1}S$. The elements of the Φ_h matrix describe the impact of the variable associated with the row on the one associated with the column after h periods. These are the impulse response functions that we present in the next section.

5 Results

This section presents the results of the paper. We illustrate these results in the form of impulse responses of the three variables to exogenous changes in unemployment and TFP. First, we focus on impulse responses following an exogenous change in cyclical unemployment, which we interpret as the result of shocks that cause economic fluctuations at business cycle frequencies. In this way, we test the hysteresis prediction generated by endogenous growth theory over the business cycle. Second, we study the effect of an exogenous change in total factor productivity. This exercise allows us to identify which specific model does not produce results that are consistent with empirical evidence. Hence, we are able to rule out some of the mechanisms for the pro-cyclicality of R&D. Finally, we concentrate on the historical variance decomposition to get a sense of the quantitative importance of the mechanism we discuss. After doing that, we revisit the discussion on the productivity growth slowdown by pointing out how the business cycle interferes with the general understanding of the post WWII aggregate productivity dynamics.

5.1 Exogenous Change in Cyclical Unemployment

The results depicted in Figure 2 show the response of R&D and utilization-adjusted TFP in levels to an exogenous increase in cyclical unemployment.

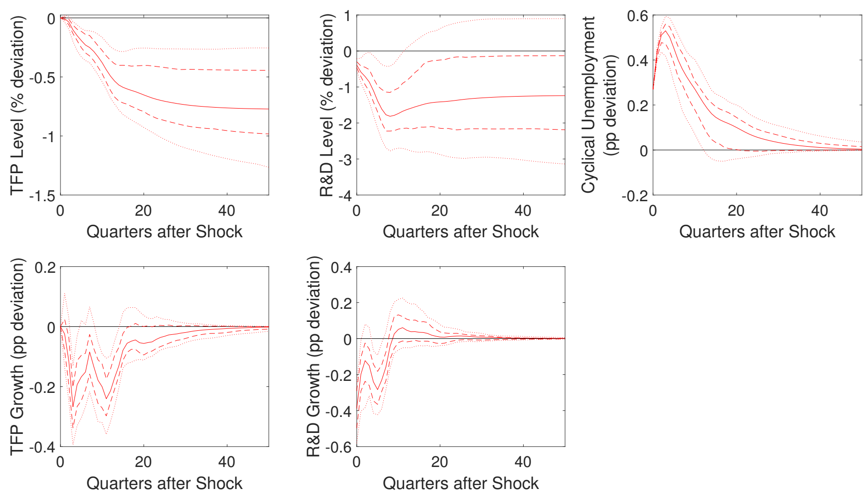

Figure 2: Unemployment Shock.

Note: Impulse Responses in growth rates to a one standard deviation shock to the cyclical unemployment rate. The dashed and dotted lines are the 68% and 95% confidence bands respectively. The black line is the pre-shock trend.

These results are in line with the theory's predictions. R&D is procyclical, and jumps immediately as the shock hits the economy. Productivity moves pro-cyclically as well but does so more slowly. The figure shows two drops in R&D growth, followed by two drops in TFP growth after about 4 quarters. Indeed, TFP's response is consistent with the idea that R&D causes it to move with some lags.

Our analysis, therefore, shows that hysteresis is present, and it is also quantitatively large. For reference, an increase in cyclical unemployment that peaks at 1 percentage points (which is two standard deviations shock to the cyclical unemployment rate) leads to a permanent loss in aggregate productivity of about 1.5% (note that the figure shows the impulse responses to one standard deviation shock).

5.2 Exogenous Change in Total Factor Productivity

The story is different when considering an exogenous increase in TFP, as shown in Figure 3. First, it is interesting to notice that the behavior of unemployment is the same as the one observed in Galí (1999) following a technology shock, despite a different identification strategy and the use of different measures.

More relevant to the scope of this paper is the behavior of R&D. R&D initially decreases following a positive exogenous increase in productivity, although the effect is small. This result is in line with Schumpeter's seminal idea that opened the field of R&D over the business cycle. In that view, R&D should move counter-cyclically as moving resources away from it and toward production in periods of higher profitability is convenient.

Over time, a few quarters after this reduction in R&D, TFP growth reverts and becomes negative. An interpretation of these results is that this

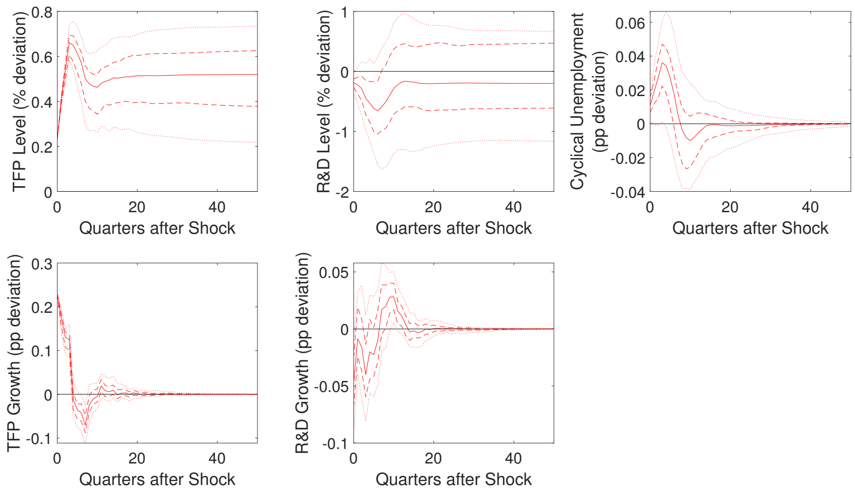

Figure 3: TFP Shock.

Note: Impulse Responses in growth rates and levels to a one standard deviation shock to utilization adjusted TFP growth. The dashed and dotted lines are the 68% and 95% confidence bands respectively. The black line is the pre-shock trend.

pro-cyclical opportunity cost of R&D produces a counter-cyclical behavior of R&D following a productivity shock. Then, the feedback mechanism from R&D to productivity produces a series of oscillations that die off after a couple of direction changes. The effect on TFP is permanent, while the long-run effect on R&D is not statistically different from 0. However, its confidence bands include an increase equal to the one on TFP.

This result is useful because it allows us to rule out a few theories on the pro-cyclicality of R&D. Theories where the pro-cyclicality emerges from the pro-cyclicality of profits would predict an increase in R&D following an increase in TFP. The theories whose predictions align with our results are those where the pro-cyclicality stems from financial frictions, or from New-Keynesian features as in Aysun (2020).

5.3 How Important Is the Business Cycle for R&D and TFP?

Turning to the forecast error variance decomposition helps to understand the main drivers of variation in the three time series. In this subsection we show the contribution of exogenous changes in unemployment to the overall variation in R&D growth and TFP growth. Based on that, we examine how the path of those variable over the sample period would have differed in the absence of shocks driving changes in unemployment.

As figure 4 shows, the variance of TFP growth is mostly explained by innovations in the TFP time series itself. The same is true for R&D: most of the deviation in R&D is explained by its own exogenous evolution. However, exogenous changes in unemployment produce deviations in TFP growth that are persistently away from its average for several quarters. In this way, this effect adds up over time producing substantial level effects.

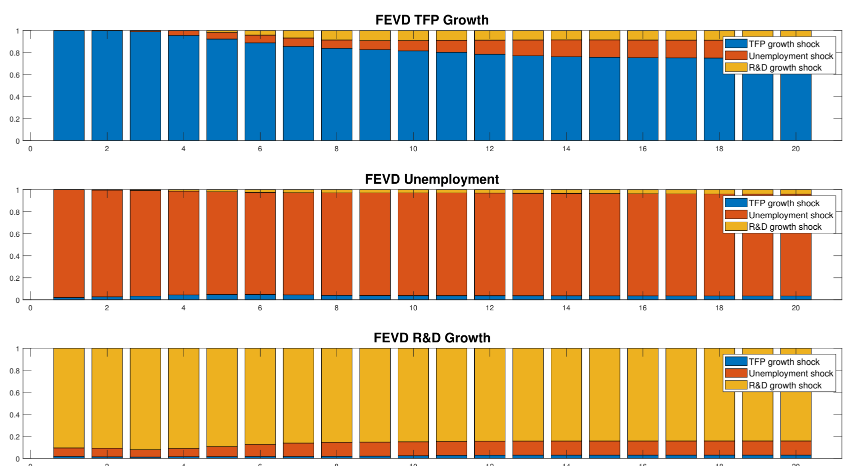

Figure 4: Forecast Error Variance Decomposition.

Before illustrating this final point further, we must point out some caveats.

We advise the reader to interpret our analysis as evidence that the mechanism that we point out produces quantitatively large effects on average. After all, our main goal is to motivate further research on this topic as opposed to pointing out precise estimates of the effect of each shock on total factor productivity in recent US history. The reader ought to keep in mind that we assume the parameters of the VAR to be constant over time. We further assume that positive and negative shocks have the same impact, just with opposite sign. We also assume that a change in their magnitude would be manifested in a proportional change in the response of all variables. Finally, we do not distinguish between different shocks as we assume that each shock that pushes unemployment away from its natural rate by a specific amount has the same effect on the other variables. We consider relaxing any of these assumption an interesting future research pursuit motivated by our quantitatively large results.

Next, we exploit the historical variance decomposition to illustrate the quantitative relevance of the mechanism we discuss in this paper over the sample period. Figure 5 illustrates the contribution of exogenous changes in unemployment to deviations of the R&D and TFP log-level from a linear trend. The graphs are constructed as the data series of R&D and TFP after removing an artificial data series where the exogenous deviations in unemployment have been shut down. Therefore, we interpret the line depicted as the deviation of R&D and TFP from their level absent the business cycle. Furthermore, Figure 6 depicts the blue lines of Figure 5, to better illustrate the connection between R&D and TFP. A noticeable feature of this graph is the large drop in R&D during recessions, followed by a gradual decline in TFP thereafter.

In those graphs, we notice three periods where the effect that we study is particularly strong: the boom in the late '60s and early '70s, the Volcker disinflation period in the early '80s, and the Great Recession. The Great Recession is of particular interest, given the existence of a lively debate on the reasons behind the productivity slowdown that occured around that period. On one side of the debate (Fernald et al., 2017), the argument is that the productivity slowdown is largely unrelated to the Great Recession as it preceded it and it consisted of a drop in the productivity growth rate to a level that was similar to the one observed a decade before. On the other side (Anzoategui et al., 2019), the idea proposed is that the productivity growth slowdown was an endogenous response to the Great Recession. Our results show that both views are partially correct: productivity growth starts declining in 2004 from reasons unrelated to the business cycle. However, this decline would not have been as severe in the absence of the Great Recession.

5.4 Revisiting the Productivity Growth Slowdown

In public discourse, the productivity growth slowdown is a hot topic. Although discussions on this topic tend to be vague on how this slowdown looks like in the data, a series of papers by Fernald Fernald (2015a,b) elucidate on the medium-term dynamics of aggregate productivity growth in the post WWII period. Specifically, he interprets these dynamics as alternations of higher productivity growth periods with lower productivity growth pe-

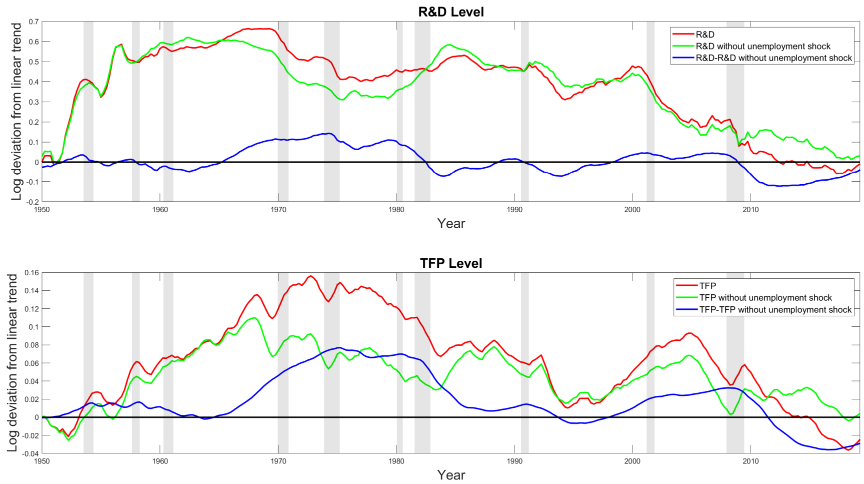

Figure 5: Log deviation from linear trend of TFP and R&D.

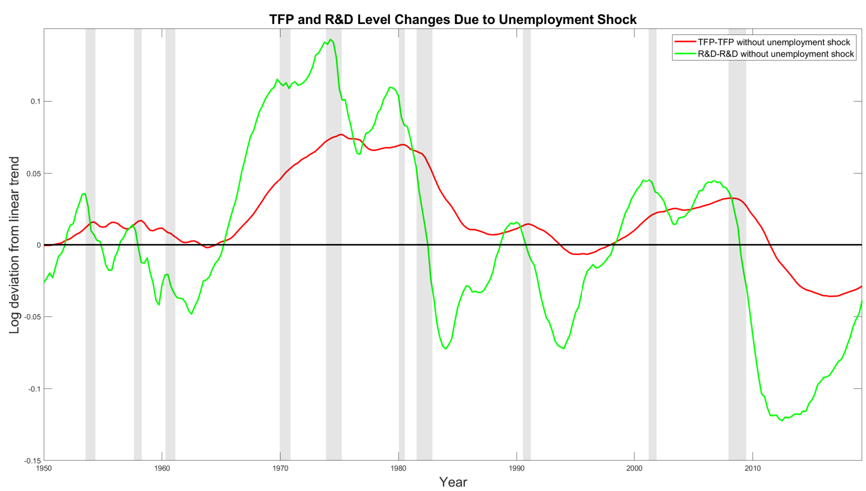

Figure 6: Log deviation from linear trend of TFP and R&D caused by exogenous changes in cyclical unemployment.

riods. Aggregate productivity growth is high between 1948Q1-1973Q1 and between 1996Q1-2004Q3. It is intead low between 1973Q2-1995Q4 and between 2004-Q4 and the latest observation available. Through our analysis, we can construct a synthetic time series of "cycle-free" aggregate productivity growth by dismissing changes in aggregate productivity growth that, according to our SVAR, result from deviations in cyclical unemployment to further inform this debate.

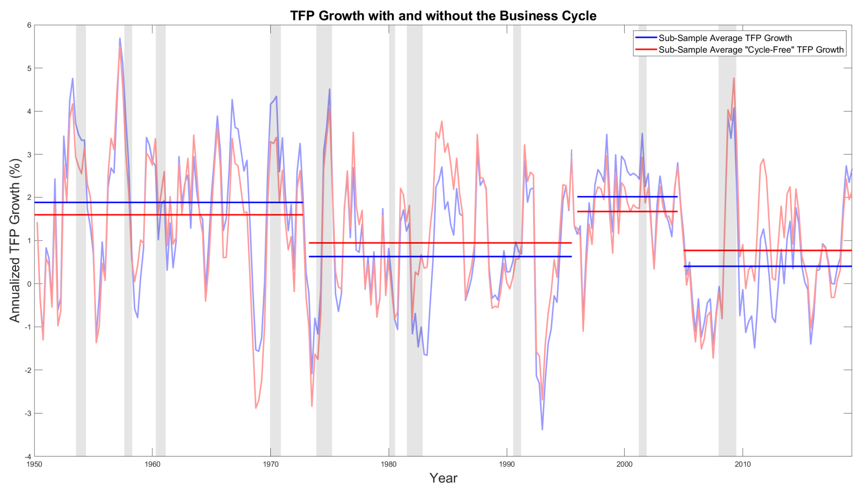

Figure 7: Sub-Sample average of productivity growth in the data and subtracting away deviations caused by the unemployment shock in the SVAR. The breaks follow Fernald.

Our results, depicted in Figure 7, show that the difference between growth rates in high versus low growth periods is noticeably shrinked once we remove the effect of the business cycle. Our procedure attributes approximately half of the difference in growth rates between high and low growth periods, corresponding to approximately 0.6 percentage points, to shocks that cause cyclical unemployment to change.

6 Robustness

The results are robust to a series of different changes in specification. In this section, we illustrate all these changes. We present in the appendix the figures for the alternative estimations that differ the most from the ones we presented in the paper. The other figures are available upon request.

First, we run the VAR adopting different lag lengths as the various criteria suggest different optimal lags. According to AIC, the optimal lag length is 13. The results are qualitatively and quantitatively similar, with the exception that the confidence bands are larger. However, the statistical significance of our results is not affected. According to SC, the optimal lag length is 2. The results, shown in the appendix, are slightly different, in that the impulse response functions are smoother, but the conclusions remain the same.

Second, when considering R&D, we would have to bear in mind a large policy change in 1982, with the introduction of an R&D tax credit. We therefore run the model twice, first using data until 1981Q4, then with data from 1982Q1. The results are almost identical.

Third, we change the time series of TFP. First, we look at utilization adjusted TFP without taking the moving average. Then, we consider TFP, without the adjustment for factor utilization. In both cases, the long-run effect is quantitatively very similar. The difference is in the short-run, as an exogenous increase in unemployment determines an increase in TFP a quarter later. This change is most likely not reflective of changes in technological knowledge.

Fourth, instead of entering TFP and R&D in log differences in the VAR, we enter them in log levels. This change rules out by assumption the presence of a permanent effect. The results, in Figure 10 and Figure 11, show that indeed TFP reverts to the mean after an exogenous change in unemployment, while the results at other time horizons are robust. The only qualitative difference is the behavior of R&D following a positive exogenous change in TFP: R&D still decreases, but only a couple of quarters after the shock, and only the 68% bands show statistical significance. On the quantitative side, the response of TFP to a change in cyclical unemployment is smaller than in our preferred estimation. Specifically, TFP declines by 0.6% for an increase in unemployment that peaks at 1 percentage point above its mean. This effect is one-third of the one estimated with the model in growth rates. Nevertheless, the effect is still quantitatively relevant, thus supporting the main conclusion of this paper.

Fifth, we change the order of the variables, assuming that an exogenous change in unemployment does not have a simultaneous impact on R&D. This change in the VAR specification ensures that more of the variation in unemployment is picked up by the innovation to the R&D time series, thus removing any potential effect that changes in investment might have on unemployment from the error term associated with the unemployment time series. The results differ in that an exogenous change in unemployment does not have an immediate impact on R&D, but everything else remains the same.

Sixth, we introduce a few variables to control for the so called "news shock". In our specification, the exogenous change in unemployment captures any shock that causes the business cycle, except for TFP shocks and exogenous variations in R&D. However, a popular stream of literature discusses the effects that expectations about future productivity may have in driving movements in aggregate variables. Of all the drivers of the business cycle, this is the only one we would want to avoid capturing because if reductions in current unemployment were due to decisions taken in response to correct predictions of an increase in future TFP, our results would misattribute future changes in TFP to shocks that cause the business cycle. On the contrary, we cannot adopt popular methods, such as Barsky and Sims (2011), to identify the news shock, because these methods rely on the assumption that the theoretical mechanism we test is not present, as they call "news shock" the combination of innovations that maximize the forecast error variance of TFP at long horizons. The compromise is to run the VAR we present in this paper while controlling for variables that have widely been adopted in the news shock literature to identify this shock. We therefore add as exogenous controls. We employ variables that incorporate information about the future following Barsky and Sims (2011), such as the S&P 500 (in log difference), consumer confidence, and CPI inflation. In doing so, we lose 10 years of data, but the results are broadly robust. The exception is the response of R&D to a TFP shock, which is now insignificant under some specifications.

Finally, we substitute the variable unemployment with the variable hours worked. We detrend hours in two different ways, first by using the HPfilter, then by considering hours per capita. The results do not change.

7 Discussion and Conclusion

This paper tests the predictions of endogenous growth theory in the presence of high-frequency disturbances. These models predict hysteresis, as shocks that temporarily move economic activity away from its trend affect the trend itself because R&D exhibits a pro-cyclical behavior. The idea behind this mechanism is that, since tehcnological knowledge is cumulative, a temporary reduction in the flow of knowledge creation permanently reduces its stock.

By focusing on post-WWII US data, we employ a recursive VAR with three variables: cyclical unemployment, aggregate R&D spending, and utilization-adjusted total factor productivity (TFP). We identify exogenous changes in cyclical unemployment and interpret them as an all-encompassing proxy for shocks that drive economic fluctuations at business cycle frequency. Our main goals are to determine whether R&D and TFP respond as the theory predicts, and to assess the quantitative relevance of these responses.

First, we confirm that R&D reacts pro-cyclically to an increase in unemployment, and TFP follows with a lag. This result aligns with the theoretical predictions and highlights the importance of these models. We also find that this effect is large: A temporary increase in cyclical unemployment that peaks at 1 percentage point above its sample average produces a permanent TFP loss of about 1.5%.

Through a historical variance decomposition, we further determine that this effect is particularly noticeable during the boom of the '60s, during the Volcker disinflation years, and during the Great Recession. Our results therefore introduce further information on the debate relative to the connection between the productivity growth slowdown and the Great Recession. While a camp, see for example Fernald et al. (2017), attributes the slowdown to causes that precede and are independent of the Great Recession, another camp, see for example Anzoategui et al. (2019), attributes the same slowdown to the Great Recession and specifically to the deteriorated financial conditions. Our analysis finds evidence of both, attributing to each a roughly equal quantitative importance.

We are conscious of the fact that our analysis suffers from some limitations, despite the various robustness checks that we account for. Specifically, the SVAR assumes constant parameters that describe the relationship between the three variables, and symmetric impulse responses. Furthermore, we do not distinguish between various sources of shocks. What we can say with confidence is that, when a shock displaces economic activity from its trend, the historical record suggests that hysteresis is present and quantitatively relevant on average. Because of this, we believe that exploring the issue further by relaxing these assumptions is a promising avenue for future research.

A second result we uncover is that, following an exogenous increase in TFP, R&D initially declines. This result is relevant as it helps in determining what drives R&D's pro-cyclical behavior. Models where R&D is pro-cyclical due to profit's pro-cyclicality are hard to reconcile with our result. In these models, a positive TFP shock increases expected short-term future profits, thus leading firms to conduct more R&D.

Instead, the result we obtain is consistent with predictions from models where financial frictions drive pro-cyclical R&D. These models include the Schumpeterian channel that gave rise to this stream of literature, namely that following a boom resources are reallocated from R&D to production as the marginal value of production increases, increasing the opportunity cost of R&D. Thus, a TFP shock would reduce R&D initially. However, credit constraints for R&D-performing firms ensure that negative shocks that drive aggregate fluctuations and tighten financial conditions dry up funds that firms would allocate to R&D. Thus, the behavior of R&D would vary depending on which shock hits the economy. Previous research emphasizes stronger pro-cyclicality in industries that are more reliant on external financing van Ophem et al. (2019). Our results reinforce that view.

Our finding is also consistent with results from an endogenous growth model with New-Keynesian features (Aysun, 2020). The key insight stressed in this model is that R&D comoves positively with aggregate employment, as opposed to output. Therefore, positive productivity shocks, which in our analysis and in previous research (Galí, 1999; Basu et al., 2006; Li, 2022) reduce employment, also reduce R&D.

However, our study does not allow us neither to rule out that a relevant part of the explanation for R&D's pro-cyclicality remains unknown, nor to estimate the extent to which financial or labor market frictions matter for hysteresis. We leave this pursuit to future research.

In conclusion, the size of the estimated effect suggests that further investigations on the connection between business cycles and long-run growth should take up a more prominent position in macroeconomists' priorities. In particular, the theoretical challenge is to provide new explanations for why R&D is pro-cyclical following non-technology shocks but responds negatively to positive TFP shocks. On the empirical side, a noteworthy exploration would concern the detection of asymmetries either in the shocks that drive R&D's behavior or in R&D's response to these shocks, which already finds some evidence in Ouyang (2011). In this way, the economy's volatility could be linked to its average growth rate.

References

Aghion, P., Askenazy, P., Berman, N., Cette, G., and Eymard, L. (2012). Credit Constraints and the Cyclicality of R&D Investment: Evidence from France. *Journal of the European Economic Association*, 10(5):1001–1024.

Aghion, P. and Saint-Paul, G. (1998). Virtues of Bad Time: Interaction Between Productivity Growth and Economic Fluctuations. *Macroeconomic Dynamics*, 2(3):322–344.

Anzoategui, D., Comin, D., Gertler, M., and Martinez, J. (2019). Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence. *American Economic Journal: Macroeconomics*, 11(3):67–110.

Aysun, U. (2020). Volatility costs of R&D. *European Economic Review*, 122:103365.

Barlevy, G. (2007). On the Cyclicality of Research and Development. *American Economic Review*, 97(4):1131–1164.

Barsky, R. B. and Sims, E. R. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58(3):273–289.

Basu, S., Fernald, J. G., and Kimball, M. S. (2006). Are Technology Improvements Contractionary? *American Economic Review*, 96(5):1418–1448.

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

Bond-Smith, S. (2019). The Decades-Long Dispute Over Scale Effects in the Theory of Economic Growth. *Journal of Economic Surveys*, 33(5):1359–1388.

Breitung, J. (2001). A convenient representation for structural vector autoregressions. *Empirical Economics*, 26(2):447–459.

Cerra, V., Fatás, A., and Saxena, S. C. (2023). Hysteresis and Business Cycles. *Journal of Economic Literature*, 61(1):181–225.

Cerra, V. and Saxena, S. C. (2008). Growth Dynamics: The Myth of Economic Recovery. *The American Economic Review*, 98(1):439–457.

Comin, D. and Gertler, M. (2006). Medium-Term Business Cycles. *American Economic Review*, 96(3):523–551.

Duval, R., Hong, G. H., and Timmer, Y. (2020). Financial Frictions and the Great Productivity Slowdown. *The Review of Financial Studies*, 33(2):475–503.

Fatas, A. (2000). Do Business Cycles Cast Long Shadows? Short-Run Persistence and Economic Growth. *Journal of Economic Growth*, 5(2):147–162.

Fatás, A. and Summers, L. H. (2018). The permanent effects of fiscal consolidations. *Journal of International Economics*, 112:238–250.

Fernald, J. (2014). A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. Technical Report 2012-19, Federal Reserve Bank of San Francisco.

Fernald, J. (2015a). The Recent Rise and Fall of Rapid Productivity Growth. *FRSB Economic Letter*.

Fernald, J., Hall, R. E., Stock, J. H., and Watson, M. W. (2017). The disappointing recovery of output after 2009. *Brookings Papers on Economic Activity*.

Fernald, J. G. (2015b). Productivity and Potential Output before, during, and after the Great Recession. *NBER Macroeconomics Annual*, 29(1):1–51.

Francois, P. and Lloyd-Ellis, H. (2009). Schumpeterian cycles with procyclical R&D. *Review of Economic Dynamics*, 12(4):567–591.

Galí, J. (1999). Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations? *American Economic Review*, 89(1):249–271.

Ghazi, S. (2019). Large Firms and Long-Run Growth: How Trends in Market Structure Affected U.S. High-R&D Industries.

Huang, C.-Y., Lai, C.-C., and Peretto, P. F. (2023). Public R&D, Private R&D and Growth: A Schumpeterian Approach. *Manuscript*.

Li, H. (2022). The Time-Varying Response of Hours Worked to a Productivity Shock. *Journal of Money, Credit and Banking*, 55(7):1907–1935.

Mand, M. (2019). On the cyclicality of R&D activities. *Journal of Macroeconomics*, 59:38–58.

Massari, F. (2023). Turbulent Growth: Business Dynamism and Aggregate Productivity.

Nuño, G. (2011). Optimal research and development and the cost of business cycles. *Journal of Economic Growth*, 16(3):257.

Ouyang, M. (2011). On the Cyclicality of R&D. *The Review of Economics and Statistics*, 93(2):542–553.

Peretto, P. F. and Connolly, M. (2007). The Manhattan Metaphor. *Journal of Economic Growth*, 12(4):329–350.

Queralto, A. (2019). A model of slow recoveries from financial crises. *Journal of Monetary Economics*.

Romer, P. M. (1990). Endogenous Technological Change. *Journal of Political Economy*, 98(5):S71–S102.

Schumpeter, J. A. (1942). *Capitalism, Socialism and Democracy*. Harper and Brothers, New York, harper colophon edition (1976) edition.

Shleifer, A. (1986). Implementation Cycles. *Journal of Political Economy*, 94(6):1163–1190.

van Ophem, H., van Giersbergen, N., van Garderen, K. J., and Bun, M. (2019). The cyclicality of R&D investment revisited. *Journal of Applied Econometrics*, 34(2):315–324.

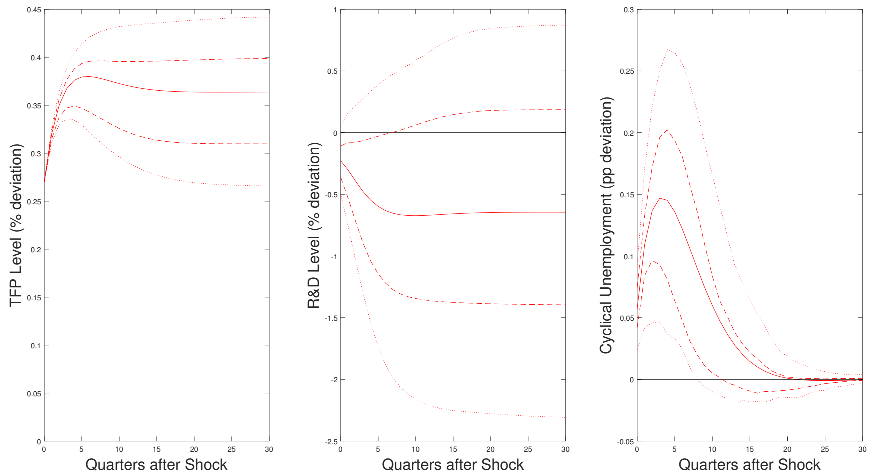

Figure 8: TFP Shock.

Note: Impulse Responses in growth rates and levels to a one standard deviation shock to utilization adjusted TFP growth with lag of 2. The dashed and dotted lines are the 68% and 95% confidence bands respectively. The black line is the pre-shock trend.

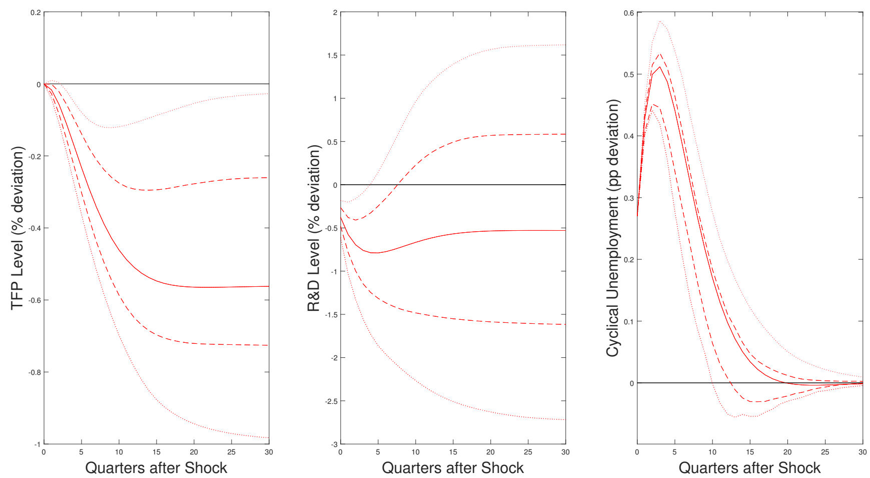

Figure 9: Unemployment Shock.

Note: Impulse Responses in growth rates and levels to a one standard deviation shock to the cyclical unemployment rate with lag of 2. The dashed and dotted lines are the 68% and 95% confidence bands respectively. The black line is the pre-shock trend.

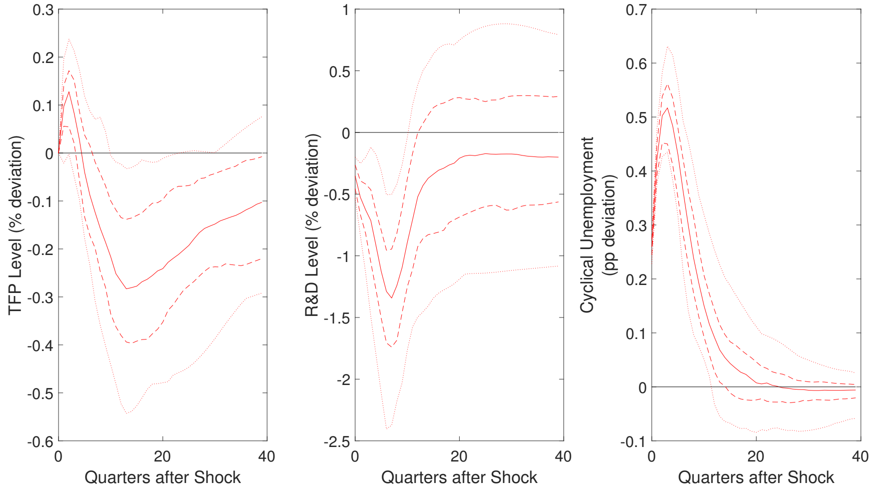

Figure 10: Unemployment Shock (Levels).

Note: Impulse Responses a one standard deviation shock to the cyclical unemployment rate. The dashed and dotted lines are the 68% and 95% confidence bands respectively. The black line is the pre-shock trend. The model is estimated with TFP and R&D entering in levels.

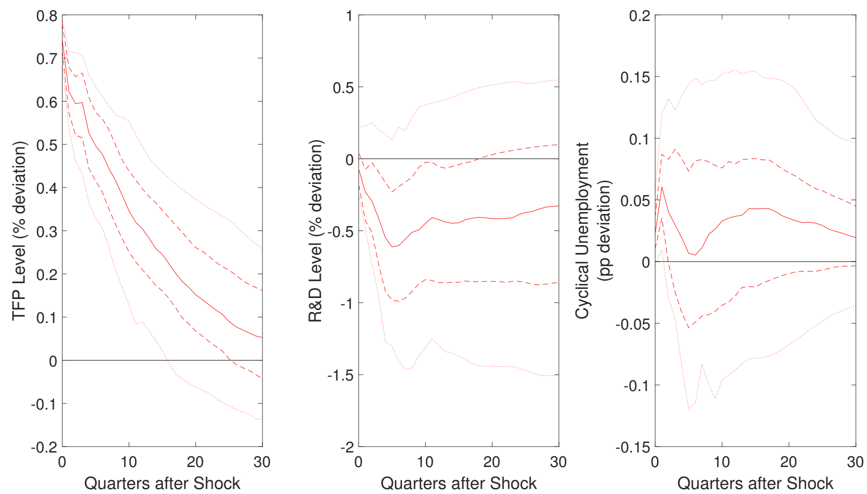

Figure 11: TFP Shock (Levels).

Note: Impulse Responses in growth rates and levels to a one standard deviation shock to utilization adjusted TFP. The dashed and dotted lines are the 68% and 95% confidence bands respectively. The black line is the pre-shock trend. The model is estimated with TFP and R&D entering in levels.