Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence*

Diego Anzoategui
Rutgers

Diego Comin
Dartmouth, NBER and CEPR

Mark Gertler
NYU and NBER

Joseba Martinez
NYU

June 2018

Abstract

We examine the hypothesis that the slowdown in productivity following the Great Recession was in significant part an endogenous response to the contraction in demand that induced the downturn. We first present some panel data evidence that technology diffusion is highly cyclical. We then develop and estimate a macroeconomic model with an endogenous TFP mechanism that allows for both costly development and adoption of new technologies. We then show that the model's implied cyclicality of technology diffusion is consistent with the panel data evidence. We next use the model to assess the sources of the productivity slowdown. We find that a significant fraction of the post-Great Recession fall in productivity was an endogenous phenomenon. The endogenous productivity mechanism also helps account for the slowdown in productivity prior to the Great Recession, though for this period shocks to the effectiveness of R&D expenditures are critical. Overall, the results are consistent with the view that demand factors have played a role in the slowdown of capacity growth since the onset of the recent crisis. More generally, they provide insight into why recoveries from financial crises may be so slow.

JEL: JEL Classification: E3, O3

Keywords: Endogenous Technology, Business Cycles.

1

1 Introduction

One of the great challenges for macroeconomists is explaining the slow recovery from major financial crises (see, e.g. Reinhart and Rogoff (2009)). Many popular explanations involve persistent demand shortfalls.[1] For example, a number of authors emphasize the deleveraging process as an important cause of a sustained decline in spending by borrowers as they saved to reduced their indebtedness. Others emphasize how constraints on macroeconomic policy likely also contributed to sluggish demand. The zero lower bound on the nominal interest rate limited the ability of monetary policy to stimulate the economy, and the political fight over the national debt ceiling effectively removed fiscal policy as a source of stimulus.

While these demand side factors have undoubtedly played a central role, it is unlikely that they alone can account for the extraordinarily sluggish movement of the economy back to the pre-crisis trend. This has led a number of authors to explore the contribution of supply-side factors. Both Hall (2014) and Reifschneider et al. (2015) have argued that the huge contraction in economic activity induced by the financial crisis in turn led to an endogenous decline in capacity growth. Hall (2014) emphasizes how the collapse in business investment during the recession brought about a non-trivial drop in the capital stock. Reifschneider et al. (2015) emphasize not only this factor but also the sustained drop in productivity. They conjecture that the drop in productivity may be the result of a decline in productivity-enhancing investments, and thus an endogenous response to the recession.

Indeed, sustained drops in productivity appear to be a feature of major financial crises. This has been the case for the U.S. and Europe in the wake of the Great Recession. The same phenomenon holds broadly for financial crises in emerging markets: in a sample of East Asian countries that experienced a financial crisis during the 1990s, Queralto (2015) finds a sustained drop in labor productivity in each case to go along with the sustained decline in output. Using panel data Huber (2018) finds that banking distress in Germany induced persistent declines in output and productivity, and he associates the latter with declines in R&D.

What accounts for reduced productivity growth following financial crises? There are two candidate hypotheses: bad luck versus an endogenous response. Fernald (2014) and Fernald et al. (2017) make a compelling case for the bad luck hypothesis. As they emphasize, the productivity slowdown began prior to the Great Recession, raising questions on whether the crisis itself can be a causal factor. Figure 1 illustrates the argument. The figure plots

---

[1]See, for example Christiano et al. (2015) for a taxonomy of possible explanations for the slow recovery.

Figure 1: Detrended Capacity Adjusted TFP and Labor Productivity
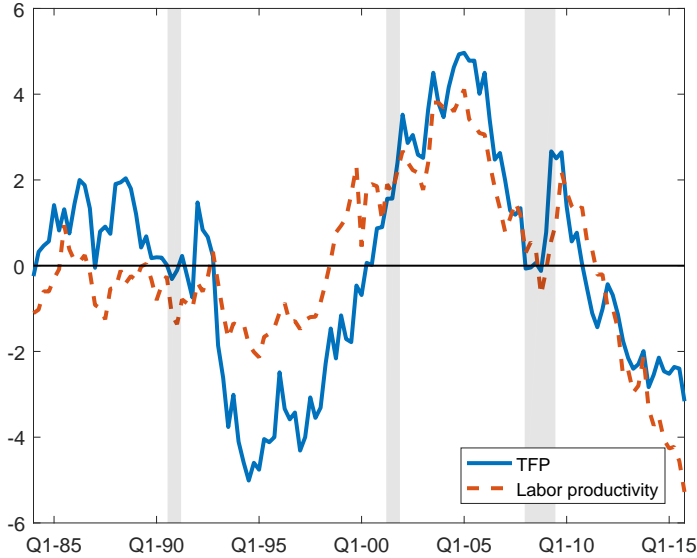
All series are log-linearly detrended. Labor productivity is GDP divided by hours worked (see Appendix A.1 for data sources). TFP is Utilization-Adjusted Total Factor Productivity (available at http://www.frbsf.org/economic-research/total-factor-productivity-tfp/; see Fernald (2012) for details).

detrended total factor productivity, specifically the utilization corrected measure described in Fernald (2012), along with labor productivity. Both measures show a sustained decline relative to trend in the years after the Great Recession, but the decline appears to begin around 2004-05.

There are several different theories of how the productivity slowdown could reflect an endogenous response to the crisis. The one on which we focus involves a reduction in productivity enhancing investments.[2] Specifically, to the extent that the crisis induced a drop in these investments, the subsequent decline in productivity could be an endogenous outcome.

We focus on two types of productivity enhancing investments: (i) the creation of new technologies through research and development (R&D) and (ii) the diffusion of new tech-

---

[2]Our hypothesis is similar to Reifschneider et al. (2015). An alternate approach stresses misallocation of productive inputs following a financial crisis. See for example Garcia-Macia (2015) who emphasizes misallocation between tangible and intangible capital. A further alternative, studied by Decker et al. (2017) is that declining business dynamism has led to declines in productivity.

nologies via adoption expenditures. In Section 2 we present evidence that each of these types of investments is highly procyclical. The cyclicality of R&D is readily apparent from aggregate data. It declined nontrivially during the Great Recession, but exhibited an even sharper decline relative to economic activity during and after the 2001-2002 recession.[3] We will argue that this decline contributed to the productivity slowdown prior to the Great Recession. Unfortunately there is no aggregate data series on technology adoption. However, using a panel of survey data on technology adoption, we estimate a highly cyclical pace of diffusion. Indeed our subsequent analysis will find much of the endogenous productivity slowdown during and after the Great Recession attributable to a drop in adoption intensity.

To assess the quantitative relevance of an endogenous technology response to the crisis on the evolution of TFP, we develop and estimate a macroeconomic model modified to allow for endogenous technology via R&D and adoption. The endogenous productivity mechanism we develop is based on Comin and Gertler (2006), which uses the approach to connect business cycles to growth. The Comin/Gertler work, in turn, is a variant of Romer (1990)'s expanding variety model of technological change, modified to include an endogenous pace of technology adoption. We include adoption to allow for a realistic period of diffusion of new technologies, and we allow for endogenous adoption intensity to capture cyclical movements in productivity that may be the product of cyclical adoption rates. We then show that the cyclical speed of diffusion generated by our model is consistent with the panel data evidence presented in Section 2. We also include an exogenous TFP shock that can capture the Fernald (2014) bad luck hypothesis.

We find that a sharp decline in endogenous adoption intensity accounts for much of the productivity decline during the Great Recession and after. On the other hand, a decline in the efficiency of R&D is mainly responsible for the pre-Great Recession slowdown. Exogenous TFP movements do not explain much of the productivity variation over this period. We conclude by providing independent evidence that supports the plausibility of our key findings. In particular, we provide a measure of R&D productivity based on the ratio of new patents relative to R&D researchers that displays the same pattern as the R&D productivity shocks we estimate. Similarly, we show that a measure of expenditures by companies in licensing new technologies developed by universities resembles the evolution of the intensity of adoption implied by our model.

To be sure, we are not suggesting that a sustained productivity slowdown should follow

---

[3]Further, the decline in R&D in 2001 was reasonably widespread across sectors. Using COMPUSTAT data, we observe that out of 29 sectors with R&D expenditures of at least $100 million in 2000, 19 experienced a decline in R&D expenditures between their peak in the period 1998-2000 and 2002. As a benchmark, computers and software experienced the 9^{th} and 10^{th} largest declines, of approximately 20% in each case.

every recession. What is key is that there are other factors, such as constraints on policy, that contribute to a slow recovery. In the case of the Great Recession, the combination of the zero lower bound on the nominal interest rate, fiscal policy constraints and the deleveraging process created an environment where declines in productivity enhancing intangible investments could have a sizable effect on productivity. This contrasts, for example, with the 1981-82 recession following which there was a rapid bounce back in output, due in large part to aggressive fiscal and monetary policy, without any tangible impact on productivity.

In addition to the literature cited above, there are several other papers related to our analysis. Schmookler (1966) and Shleifer (1986) emphasize the role of aggregate demand on the timing of innovation and technology adoption. Queralto (2015), Guerron-Quintana and Jinnai (2014) and Garcia-Macia (2015) have appealed to endogenous growth considerations to explain the persistence of financial crises. The paper most closely related to ours is Bianchi et al. (2017), who also estimate a macroeconomic model with endogenous growth and are the first to use R&D data. In addition to variation in details and focus, there are three key differences between our analysis and that in Bianchi et al. (2017). First, our model of R&D and adoption is more explicit, which imposes discipline on the lags in the diffusion process, and allows a clearer interpretation of the technology parameters. Second, and perhaps most importantly, we use the panel data evidence presented in Section 2 as an external validity check to ensure that the cyclical behavior of adoption in the model is plausible. This distinction is important because cyclical diffusion plays an important role in this literature but is not directly observable in the aggregate, necessitating some form of external validity. A final notable difference is that our solution method allows us to take account of a binding zero lower bound on monetary policy, which turns out to be an important factor propagating the endogenous decline in productivity in the wake of the Great Recession.

The rest of the paper is organized as follows. Section 2 presents evidence of the cyclical behavior of R&D and technology adoption. Section 3 presents the model. Section 4 describes the econometric implementation and presents our estimation results. In addition, we show that estimates of the cyclicality of technology diffusion from artificial data generated by the model are consistent with the estimates in 2. Finally, using a historical decomposition of productivity growth, section 5 analyzes the extent to which the endogenous growth mechanism can account for the evolution of productivity both before and after the Great Recession.

2 Evidence on R&D and the speed of technology diffusion

In this section, we present evidence on the cyclical behavior of R&D and technology adoption. The goal is twofold: first to motivate our formulation of endogenous productivity; and second to present external evidence that we use in Section 4.7 to validate the quantitative predictions of our estimated model.

Figure 2 plots expenses on R&D conducted by US corporations. The figure shows a clear procyclical pattern, consistent with the evidence found in other studies (see e.g. Comin and Gertler (2006)).[4] While there is a noticeable decline in R&D expenditures during the Great Recession, there is a substantially larger decline relative to economic activity following the 2001-2002 recession. Overall, the figure raises the possibility that the productivity slowdown prior to the Great Recession was in part a consequence of the sharp R&D contraction that preceded it.

We next turn to technology adoption. As noted earlier, an aggregate time series measuring adoption is not available. To explore the cyclicality of technology diffusion we resort to survey data on the speed of technological diffusion, of the type used in the productivity literature, for example, Griliches (1957) and Mansfield (1961). The specific data we have available is time series data of the fraction of companies that have adopted a technology, for a sample of 26 production technologies, detailed in Table 1, that diffused at various times over the period 1947-2003 in the US (5) and the UK (21). We use the data to estimate the effect of the business cycle on the speed of diffusion, after controlling for the normal diffusion process.

Specifically, we denote by m_{it} the fraction of potential adopters that have adopted a specific technology i in t. The ratio of adopters to non-adopters r_{it} is

$$r_{it} = m_{it}/(1 - m_{it}). \tag{1}$$

The speed of diffusion is then the percentage change in r_{it} :

$$Speed_{it} = \triangle \ln(r_{it}) \tag{2}$$

As shown by Mansfield (1961), if the diffusion process follows a logistic curve, the speed of diffusion (2) is equal to a constant α_i. In reality, however, the speed of diffusion is not constant; it tends to be faster in the early stages. Therefore, r_{it} declines with the age of

---

[4]There is a long literature documenting the cyclicality of R&D expenditures (see Barlevy (2007), for a summary). Barlevy (2007) presents evidence based on firm-level data on the importance of both sectoral demand as well as firms' financial conditions for the pro-cyclicality of R&D expenditures.

Figure 2: R&D Expenditures by US Corporations, 1983-2013
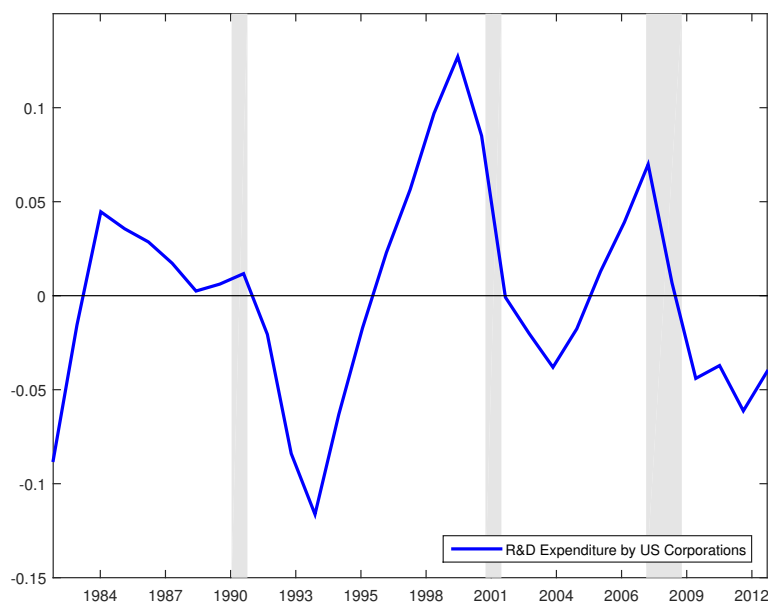
Log-linearly detrended data. Source: R&D Expenditure by US corporations (National Science Foundation). Data are deflated by the GDP deflator and divided by the civilian population older than 16 (see Appendix A.1 for data sources).

Table 1: Technologies in diffusion panel data

UK Technologies	
Special presses	Foils
Wet suction boxes	Gibberellic acid
Automatic size boxes	Accelerated drying hoods
Basic oxygen process	Vacuum degassing
Vacuum melting	Continuous casting
Electrical hygrometers	Synthetic fabrics
Tunnel kilns	Computer-controlled processing
Tufted carpets	Computer typesetting
Photo-electrically controlled cutting	Shuttleless looms
Numerical control printing presses	Numerical control turning machines
Numerical control turbines	
US Technologies	
CT scanners	Computerized numerical control machines
Automated inspection sensors	3-D CAD
Flexible manufacturing systems	

Source: UK data from Davies (1979). US data from Trajtenberg (1990) and Bartel et al. (2009)

the technology. Additionally, we want to explore whether the speed of technology diffusion varies over the cycle. To this end, we consider the following specification

$$Speed_{it} = \alpha_i + G(lag_{it}) + \beta * \widehat{y}_t + \epsilon_{it}, \tag{3}$$

where $G(.)$ is a polynomial in the years since the technology was first introduced, and $\widehat{y}_t$ is a measure of cyclical fluctuations in GDP, specifically detrended real GDP per capita (following Comin and Gertler (2006) we detrend GDP using a bandpass filter to capture high- and medium- frequency fluctuations).

Table 2 presents the estimates of equation (3). The main finding is that the estimates of the elasticity of the speed of diffusion with respect to the cycle, β, are robustly positive and significant. In particular, the point estimate is between 3.6 and 4.1 depending on the specification. These estimates of β provide a benchmark on the cyclicality of the speed of technology diffusion in the micro-data. In our analysis, we use this benchmark to externally validate the sensitivity of diffusion to the cycle in our model and in this way ensure that the productivity dynamics induced by the endogenous adoption mechanisms in the model are consistent with the micro-evidence.

Table 2: Cyclicality of the Speed of Technology Diffusion

|	I	II	III	IV
$\widehat{y}_t$	3.73	3.7	3.64	4.12
	(3.59)	(2.81)	(3.94)	(3.17)
$\widehat{y}_t$ * US		0.07		-0.74
		(0.04)		(0.53)
lag_{it}	-0.057	-0.057		
	(5.22)	(4.76)		
lag_{it}^2	0.001	0.001		
	(2.52)	(2.12)		
$ln(lag_{it})$			-0.29	-0.29
			(6.68)	(6.65)
R2 (within)	0.11	0.11	0.13	0.13
N technologies	26	26	26	26
N observations	327	327	327	327

Notes: (1) dependent variable is the speed of diffusion of 26 technologies, (2) all regressions include technology specific fixed effects. (3) t-statistics in parenthesis, (4) $\widehat{y}_t$ denotes the cycle of GDP per capita in the country and represents the high and medium frequency components of output fluctuations, (5) $\widehat{y}_t$*US is the medium term cycle of GDP per capita times a US dummy, (6) lag represents the years since the technology first started to diffuse.

The effect of years since the technology started diffusing is negative and convex (i.e. it vanishes over time). The results are robust to specifying the function G as a second order polynomial or in logarithms. Finally, we do not observe any significant differential effect of the cycle on US versus UK technologies.

To illustrate the cyclicality of the speed of technology diffusion for U.S. data, Figure 3 plots the speed of diffusion for the balanced panel of four US technologies for which we have data from 1981 to 2003. Specifically, for each of the technologies we remove the acyclical component of the diffusion rate $(\alpha_i + G(lag_{it}))$. We then average the residual $(\beta * \widehat{y}_t + \epsilon_{it})$ over the four technologies. The dashed line is a plot of this average, while the solid line is a three year centered moving average. The figure reveals a positive correlation between the speed of diffusion and the cycle. Diffusion speed was lowest in the deep 1981-82 recession; it recovered during the 80s and declined again after the 1990 recession. It increased notably during the expansion in the second half of the 90s and declined again with the 2001 recession.

Unfortunately we do not have comprehensive data on technology diffusion during the Great Recession. We do, however, have two types of limited evidence. First, Eurostat

Figure 3: Speed of Diffusion in the US
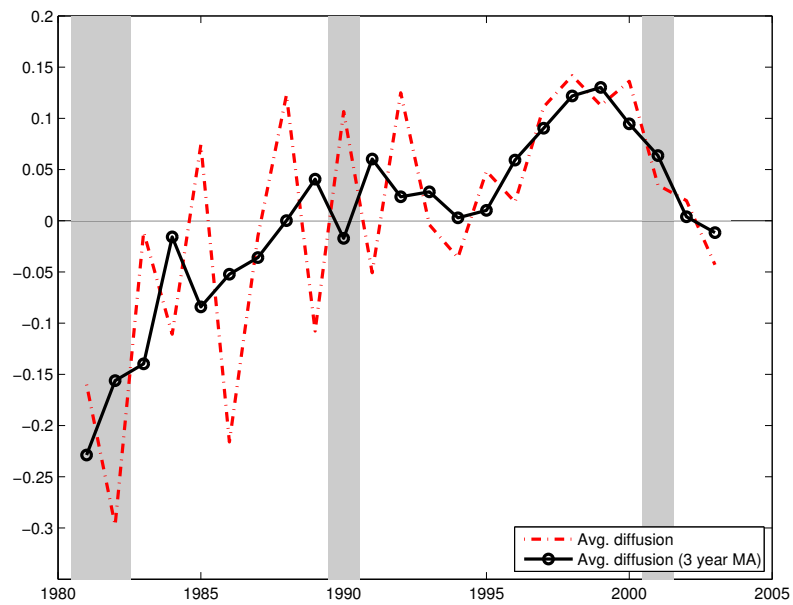
Average speed of diffusion for US technologies reported. Diffusion speed is net of age of technology effects. Specifically, for each technology we remove the acyclical component of the diffusion rate $\alpha_i + G(lag_{it})$ (see equation (3)).

provides information on the diffusion of three relevant internet-related technologies in the UK.[5] Figure 4 plots their average diffusion from 2004 until 2013 with the business cycle downturns in the UK. The figure confirms the pro-cyclicality of the speed of diffusion of these technologies. In particular, during the downturn corresponding to the Great Recession (2008-2009), the average speed of diffusion of our three technologies sharply declined by 75%. After the Great Recession, the speed of diffusion recovered but remained below trend, converging to approximately 10% below average. Beyond its cyclicality, the second observation we want to stress from the Figure is that fluctuations in the speed of diffusion are very wide, ranging from 86% above average in 2004 to 74% below the average diffusion speed in 2009.[6]

Second, for the U.S., we have an aggregate time series for fees paid for licensing new technologies. Licensing fees may be interpreted as expenses paid for adopting new technologies. We defer further discussion of this data, which shows a marked a decline during the Great Recession, to Section 5.

We make two uses of the data presented in this section in our subsequent analysis. We use the R&D series directly as an observable in our estimation procedure. And, in Section 4, we develop and estimate a model analogy to Equation (3), and use the panel regression results presented in this section as an external validity check of the cyclicality of technological diffusion generated by our estimated model.
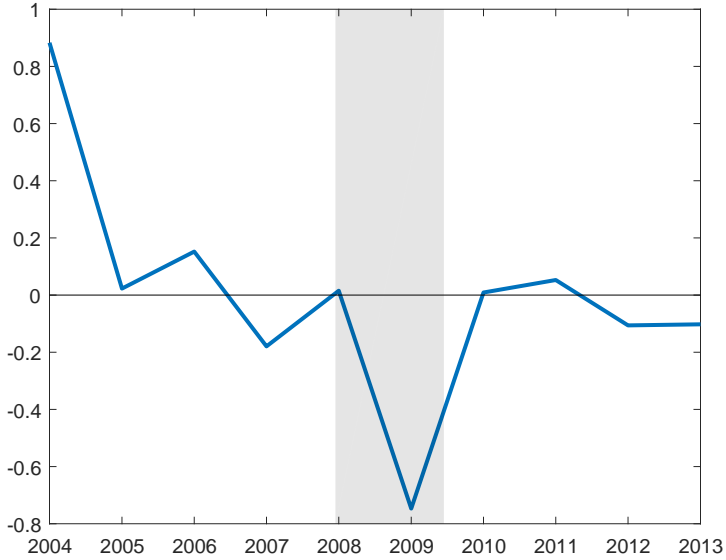
3 Model

Our starting point is a New Keynesian DSGE model similar to Christiano et al. (2005) and Smets and Wouters (2007). We include the standard features useful for capturing the data. In addition, monetary policy obeys a Taylor rule with a binding zero lower bound constraint. The key non-standard feature is that total factor productivity depends on two endogenous variables: the creation of new technologies via R&D and the speed of adoption

---

[5] The measures we consider are the fraction of firms that (i) have access to broadband internet, that (ii) actively purchase online products and services and that (iii) actively sell online products and services (actively is defined as constituting at least 1% of sales/purchases). For each of these three measures we construct the speed of technology diffusion using expression (2), and then filter the effect of the lag since the introduction of the technology using expression (3) and the estimates from column 3 of Table 2. The resulting series are demeaned so that they can be interpreted as percent deviations from the average speed of technology diffusion.

[6] Andrews et al. (2015) have recently provided complementary evidence that technology diffusion in OECD countries may have slowed during the Great Recession. In their study, they show that the gap in productivity between the most productive firms in a sector (leaders) and the rest (followers) has increased significantly during the Great Recession. They interpret the increase in the productivity gap as evidence that followers have slowed down the rate at which they incorporate frontier technologies developed by the leaders.

Figure 4: Diffusion of Technologies on Business use of Internet in UK, 2004-2013

Source: Eurostat; see footnote 5 for details of calculations. Shaded areas are UK recession dates as dated by UK ONS.

of these new technologies. Skilled labor is used as an input for the R&D and adoption processes. We do not model financial frictions explicitly; however, we allow for a shock that transmits through the economy like a financial shock, as we discuss below. We begin with the non-standard features of the model before briefly describing the standard ones.

3.1 Production Sector and Endogenous TFP: Preliminaries

In this section we describe the production sector and sketch how endogenous productivity enters the model. There are two types of firms: (i) final goods producers and (ii) intermediate goods producers. There are a continuum, measure unity, of monopolistically competitive final goods producers. Each final good firm i produces a differentiated output Y_t^i. A final good composite is then the following CES aggregate of the differentiated final goods:

$$Y_t = \left(\int_0^1 (Y_t^i)^{\frac{1}{\mu_t}} di \right)^{\mu_t} \tag{4}$$

where $\mu_t > 1$ and $\log(\mu_t)$ follows an exogenous stochastic process:

$$\log(\mu_t) = (1 - \rho_\mu)\mu + \rho_\mu \log(\mu_{t-1}) + \sigma_\mu \epsilon_t^\mu \tag{5}$$

where ϵ_t^μ is i.i.d. $N(0,1)$. Each final good firm i uses Y_{mt}^i units of intermediate goods composite as input to produce output Y_t^i, according to the following simple linear technology

$$Y_t^i = Y_{mt}^i \tag{6}$$

We assume each firm sets its nominal price P_t^i on a staggered basis, as we describe later. There exists a continuum of measure A_t of monopolistically competitive intermediate goods firms that each make a differentiated product. The endogenous predetermined variable A_t is the stock of types of intermediate goods adopted in production, i.e., the stock of adopted technologies. Intermediate goods firm j produces output Y_{mt}^j. The intermediate goods composite is the following CES aggregate of individual intermediate goods:

$$Y_{mt} = \left(\int_0^{A_t} (Y_{mt}^j)^{\frac{1}{\vartheta}} dj \right)^{\vartheta} \tag{7}$$

with $\vartheta > 1$. Let K_t^j be the stock of capital firm j employs, U_t^j be how intensely this capital is used, and L_t^j the stock of labor employed (we allow for variable utilization intensity of capital U_t^j so as not to mistakenly attribute all high frequency variation in the Solow residual

to endogenous technology). Then firm j uses capital services $U_t^j K_t^j$ and unskilled labor L_t^j as inputs to produce output Y_{mt}^j according to the following Cobb-Douglas technology:

$$Y_{mt}^j = \theta_t \left(U_t^j K_t^j \right)^\alpha (L_t^j)^{1-\alpha} \tag{8}$$

where θ_t is an aggregate productivity shock whose growth rate follows a stationary AR(1) process,

$$\log(\theta_t) = \rho_\theta \log(\theta_{t-1}) + \sigma_\theta \varepsilon_{\theta,t} \tag{9}$$

where $\varepsilon_{\theta,t}$ is i.i.d. $\sim N(0,1)$. Finally, we suppose that intermediate goods firms set prices each period. That is, intermediate goods prices are perfectly flexible, in contrast to final good prices.

Given a symmetric equilibrium for intermediate goods, it follows from equations (7) and (8) that to a first order we can express the aggregate production function for the final good composite Y_t as[7]

$$Y_t = \left[A_t^{\vartheta-1} \theta_t \right] \cdot (U_t K_t)^\alpha (L_t)^{1-\alpha} \tag{10}$$

The term in brackets is total factor productivity, which is the product of a term that reflects endogenous variation, $A_t^{\vartheta-1}$, and one that reflects exogenous variation θ_t. In sum, endogenous productivity effects enter through the expansion in the variety of adopted intermediate goods, measured by A_t. As per Equation 9 θ_t is stationary so the driving force of long term growth is the endogenous TFP mechanism. We next describe the mechanisms through which new intermediate goods are created and adopted.

3.2 R&D and Adoption

The processes for creating and adopting new technologies are based on Comin and Gertler (2006). Let Z_t denote the stock of technologies, while as before A_t is the stock of adopted technologies (intermediate goods). In turn, the difference $Z_t - A_t$ is the stock of unadopted technologies. R&D expenditures increase Z_t while adoption expenditure increase A_t. We distinguish between creation and adoption because we wish to allow for realistic lags in the adoption of new technologies. We first characterize the R&D process and then turn to adoption.

---

[7]The production function implies $Y_t = \Omega_t \cdot \overline{Y}_t$ where $\overline{Y}_t$ is average output per firm and $\Omega_t = \left(\int_0^1 (Y_t^i / \overline{Y}_t)^{\frac{1}{\mu_t}} di \right)^{\mu_t} = 1$ to a first order. Next, given the total number of final goods firms is unity, given the production function for each final goods producer (6), and given that $Y_t \approx \overline{Y}_t$, it follows that to a first order $Y_t \approx Y_{mt}$. Equation (10) then follows from (7) and (8).

3.2.1 R&D: Creation of Z_t

There are a continuum of measure unity of innovators that use skilled labor to create new intermediate goods. Let L_{srt}^p be skilled labor employed in R&D by innovator p and let φ_t be the number of new technologies available at time $t+1$ that each unit of skilled labor at t can create. We assume φ_t is given by

$$\varphi_t = \chi_t Z_t L_{srt}^{\rho_z - 1} \tag{11}$$

where χ_t is an exogenous disturbance to the R&D technology which follows an exogenous process

$$\log(\chi_t) = (1 - \rho_\chi) \log(\bar{\chi}) + \rho_\chi \log(\chi_{t-1}) + \sigma_\chi \epsilon_t^\chi \tag{12}$$

where ϵ_t^χ is i.i.d. $N(0,1)$, and L_{srt} is the aggregate amount of skilled labor working on R&D, which an individual innovator takes as given. Following Romer (1990), the presence of Z_t, which the innovator also takes as given, reflects public learning-by-doing in the R&D process. We assume $\rho_z < 1$ which implies that increased R&D in the aggregate reduces the efficiency of R&D at the individual level. This congestion externality permits us to have constant returns to scale in the creation of new technologies at the individual innovator level, which simplifies aggregation, but diminishing returns at the aggregate level. Our assumption of diminishing returns is consistent with the empirical evidence (see Griliches (1990)); further, with our specification the elasticity of creation of new technologies with respect to R&D becomes a parameter we can estimate, as we make clear shortly.[8]

Let J_t be the value of an unadopted technology, $\Lambda_{t,t+1}$ the representative household's stochastic discount factor and w_{st} the real wage for a unit of skilled labor. We can then express innovator p's decision problem as choosing L_{srt}^p to solve

$$\max_{L_{srt}^p} E_t\{\Lambda_{t,t+1} J_{t+1} \varphi_t L_{srt}^p\} - w_{st} L_{srt}^p \tag{13}$$

The optimality condition for R&D is then given by

$$E_t\{\Lambda_{t,t+1} J_{t+1} \varphi_t\} - w_{st} = 0$$

which implies

$$E_t\{\Lambda_{t,t+1} J_{t+1} \chi_t Z_t L_{srt}^{\rho_z - 1}\} = w_{st} \tag{14}$$

---

[8]An added benefit from having diminishing returns to R&D spending is that, given our parameter estimates, steady state growth is relatively insensitive to tax policies that might affect incentives for R&D. Given the weak link between tax rates and long run growth, this feature is desirable.

The left side of Equation (14) is the discounted marginal benefit from an additional unit of skilled labor, while the right side is the marginal cost. Given that profits from intermediate goods are pro-cyclical, the value of an unadopted technology, which depends on expected future profits, will be also be pro-cyclical. This consideration, in conjunction with some stickiness in the wages of skilled labor which we introduce later, will give rise to pro-cyclical movements in L_{srt}.[9]

Finally, we allow for obsolescence of technologies. Let ϕ be the survival rate for any given technology. Then, we can express the evolution of technologies as:

$$Z_{t+1} = \varphi_t L_{srt} + \phi Z_t \tag{15}$$

where the term $\varphi_t L_{srt}$ reflects the creation of new technologies. Combining equations (15) and (11) yields the following expression for the growth of new technologies:

$$\frac{Z_{t+1}}{Z_t} = \chi_t L_{srt}^{\rho_z} + \phi \tag{16}$$

where ρ_z is the elasticity of the growth rate of technologies with respect to R&D, a parameter that we estimate.

3.2.2 Adoption: From Z_t to A_t

We next describe how newly created intermediate goods are adopted, i.e. the process of converting Z_t to A_t. Here we capture the fact that technology adoption takes time on average (as documented, for example by Comin and Hobijn (2010)), but the adoption rate can vary pro-cyclically, consistent with the evidence in Comin (2009) and Section 2. In addition, we would like to characterize the diffusion process in a way that minimizes the complications from aggregation. In particular, we would like to avoid having to keep track, for every available technology, of the fraction of firms that have and have not adopted it. To do so, we suppose there is a competitive group of "adopters" who convert unadopted technologies into ones that can be used in production. They buy the rights to the technology from the innovator, at the competitive price J_t, which is the value of an unadopted technology. They then convert the technology into use by employing skilled labor as input. This process takes time on average, and the conversion rate may vary endogenously. In particular, the pace of adoption depends positively on the level of resources devoted to adoption in the

---

[9]Other approaches to motivating procyclical R&D, include introducing financial frictions (Aghion et al. (2010)), short term biases of innovators (Barlevy (2007)), or capital services in the R&D technology function (Comin and Gertler (2006)).

following simple way: an adopter succeeds in making a product usable in any given period with probability λ_t, which is an increasing and concave function of the amount of skilled labor employed, L_{sat}:

$$\lambda_t = \lambda(Z_t L_{sat}) \tag{17}$$

with $\lambda' > 0$, $\lambda'' < 0$.[10] We augment L_{sat} by a spillover effect from the total stock of technologies Z_t - think of the adoption process as becoming more efficient as the technological state of the economy improves. The practical need for this spillover is that it ensures a balanced growth path: as technologies grow, the number of new goods requiring adoption increases, but the supply of labor remains unchanged. Hence, the adoption process must become more efficient as the number of technologies expands. Unlike the specification used for R&D, there is no separate shock to the productivity of adoption activities in (17). We are forced to introduce this asymmetry because we do not have a direct observable to measure adoption labor or λ_t.

Our adoption process implies that technology diffusion takes time on average. If $\bar{\lambda}$ is the steady state value of λ_t, then the average time it takes for a new technology be adopted is $1/\bar{\lambda}$. Away from the steady state, the pace of adoption will vary with skilled labor input L_{sat}. We turn next to how L_{sat} is determined.

Once in usable form, the adopter sells the rights to the technology to a monopolistically competitive intermediate goods producer that makes the new product using the production function described by Equation (8). Let Π_{mt} be the profits that the intermediate goods firm makes from producing the good, which arise from monopolistically competitive pricing. The price of the adopted technology, V_t, is the present discounted value of profits from producing the good, given by

$$V_t = \Pi_{mt} + \phi E_t\{\Lambda_{t,t+1} V_{t+1}\} \tag{18}$$

Then we may express the adopter's maximization problem as choosing L_{sat} to maximize the value J_t of an unadopted technology, given by

$$J_t = \max_{L_{sat}} E_t\{-w_{st}L_{sat} + \phi\Lambda_{t,t+1}[\lambda_t V_{t+1} + (1 - \lambda_t)J_{t+1}\} \tag{19}$$

with λ_t as in Equation (17). The first term in the Bellman equation reflects total adoption expenditures, while the second is the discounted benefit: the probability weighted sum of

---

[10]In the estimation, we assume that

$$\lambda(\bullet) = \kappa_\lambda * (\bullet)^{\rho_\lambda}.$$

where κ_λ and $0 < \rho^\lambda < 1$ are constants.

the values of adopted and unadopted technologies. The first order condition for L_{sat} is

$$Z_t \lambda' \cdot \phi E_t\{\Lambda_{t,t+1}[V_{t+1} - J_{t+1}]\} = w_{st} \tag{20}$$

The term on the left is the marginal gain from adoption expenditures: the increase in the adoption probability λ_t times the discounted difference between the value of an adopted versus an unadopted technology. The right side is the marginal cost. The term $V_t - J_t$ is pro-cyclical, given the greater influence of near term profits on the value of adopted technologies relative to unadopted ones. Given this consideration and the stickiness in w_{st} which we alluded to earlier, L_{sat} varies pro-cyclically. The net implication is that the pace of adoption, given by λ_t, will also vary pro-cyclically.

Since λ_t does not depend on adopter-specific characteristics, we can sum across adopters to obtain the following relation for the evolution of adopted technologies

$$A_{t+1} = \lambda_t \phi[Z_t - A_t] + \phi A_t \tag{21}$$

where $Z_t - A_t$ is the stock of unadopted technologies.

3.2.3   Technology diffusion: mapping to the data

Before continuing with the model, we make a detour to map the notion of diffusion in our framework to the econometric analysis presented in Section 2. The data diffusion corresponds to the number of companies at time $t + j$ that have adopted a single technology invented at some time t. Within our model, only one firm adopts a technology but multiple technologies are invented in a given period. Accordingly, the natural notion of diffusion in the model is the share of technologies invented at time t that have been adopted at time $t + j$.

Formally, denote by Z_{t+k}^t the mass of technologies that was invented at time t that survives (i.e. is not obsolete) at time $t + k$, and A_{t+k}^t the mass of vintage t technologies that have been adopted at time $t + k$. Then, we can define the fraction of vintage t technologies adopted at time $t + k$ by

$$m_{t+k}^t \equiv \frac{A_{t+k}^t}{Z_{t+k}^t}. \tag{22}$$

Analogously to Equations 1 and 2, we define r and the speed of diffusion in the model as:

$$r_{t+k}^t \equiv \frac{m_{t+k}^t}{1 - m_{t+k}^t} \tag{23}$$

and

$$Speed_{t+k}^t \equiv \log\left(\frac{r_{t+k}^t}{r_{t+k-1}^t}\right) = \log\left(\frac{1 + \lambda_{t+k-1}/r_{t+k-1}^t}{1 - \lambda_{t+k-1}}\right) \tag{24}$$

Where the second equality comes from substituting in the law of motion for r_{t+k}^t, as detailed in Appendix A.3. Equation (24) has two relevant implications. First, the speed of technology diffusion, $Speed_{t+k}^t$, is pro-cyclical because it varies positively with the adoption probability λ_{t+k-1}, which is pro-cyclical. Second $Speed_{t+k}^t$ declines with the diffusion level of the technology as measured by r_{t+k-1}^t. This prediction is consistent with the panel data evidence from Table 2.[11]

3.3 Households

The representative household consumes and saves in the form of capital and riskless bonds which are in zero net supply. It rents capital to intermediate goods firms. As in the standard DSGE model, there is habit formation in consumption. Also as is standard in DSGE models with wage rigidity, the household is a monopolistically competitive supplier of differentiated types of labor.

The household's problem differs from the standard setup in two ways. First it supplies two types of labor: unskilled labor L_t^h which is used in the production of intermediate goods and skilled labor which is used either for R&D or adoption, L_{st}^h. Second, we suppose that the household has a preference for the safe asset, which we motivate loosely as a preference for liquidity and capture by incorporating bonds in the utility function, following Krishnamurthy and Vissing-Jorgensen (2012). Further, following Fisher (2015), we introduce a shock to liquidity demand ϱ_t.[12] As we show, the liquidity demand shock transmits through the economy like a financial shock. As the estimated model reveals, it turns out to be the main source of cyclical variation, particularly during the Great Recession period.

Let C_t be consumption, B_t holdings of the riskless bond, Π_t profits from ownership of monopolistically competitive firms, K_t capital, Q_t the price of capital, R_{kt} the rate of return, and D_t the rental rate of capital. Then the households' decision problem is given by

$$\max_{C_t, B_{t+1}, L_t^h, L_{st}^h, K_{t+1}} E_t \sum_{\tau=0}^{\infty} \beta^\tau \left\{ \log(C_{t+\tau} - bC_{t+\tau-1}) + \varrho_t B_{t+1} - \left[\frac{\upsilon(L_t^h)^{1+\varphi} + \upsilon_s(L_{st}^h)^{1+\varphi}}{1+\varphi}\right] \right\} \tag{25}$$

---

[11]To control for this non-linear "vintage" effect, which is due to the geometric nature of diffusion in the model, we introduce (see Section 4.7) a vintage-control when estimating the cyclicality of diffusion in the model, analogous to that included in the panel data regression in Section 2.

[12]As Fisher (2015) discusses, the liquidity demand shock is an explicit formulation of the risk shock in Smets and Wouters (2007). Del Negro et al. (2017) incorporates a similar shock.

subject to

$$C_t = w_t^h L_t^h + w_{st}^h L_{st}^h + \Pi_t + R_{kt} Q_{t-1} K_t - Q_t K_{t+1} + R_t B_t - B_{t+1} \tag{26}$$

with $R_{kt} \equiv (D_t + Q_t)/Q_{t-1}$. Let $\Lambda_{t,t+1} \equiv \beta u'(C_{t+1})/u'(C_t)$ be the household's stochastic discount factor and $\zeta_t \equiv \varrho_t/u'(C_t)$ be the liquidity demand shock in consumption units. The shock itself is given by the following exogenous process:

$$\zeta_t = (1 - \rho_\zeta)\,\bar{\zeta} + \rho_\zeta \zeta_{t-1} + \sigma_\zeta \epsilon_t^\zeta, \tag{27}$$

where ϵ_t^ζ is i.i.d. $N(0,1)$. Then we can express the first order necessary conditions for capital and the riskless bond as, respectively:

$$1 = E_t\{\Lambda_{t,t+1} R_{kt+1}\} \tag{28}$$

$$1 = E_t\{\Lambda_{t,t+1} R_{t+1}\} + \zeta_t \tag{29}$$

As equation (29) indicates, the liquidity demand shock distorts the first order condition for the riskless bond. A rise in ζ_t acts like an increase in risk: given the riskless rate R_{t+1} the increase in ζ_t induces a precautionary saving effect, as households reduce current consumption in order to satisfy the first order condition (which requires a drop in $\Lambda_{t,t+1}$). It also leads to a drop in investment demand, as the decline in $\Lambda_{t,t+1}$ raises the required return on capital, as equation (28) implies. The decline in the discount factor also induces a drop in adoption and R&D. Overall, the shock to ζ_t generates positive co-movement between consumption and investment similar to that arising from a monetary shock. To see this, combine equations (28) and (29) to obtain

$$E_t\{\Lambda_{t,t+1}(R_{kt+1} - R_{t+1})\} = \zeta_t \tag{30}$$

To a first order an increase in ζ_t has an effect on both R_{kt+1} and $\Lambda_{t,t+1}$ that is qualitatively similar to that arising from an increase in R_{t+1}. In addition, note that an increase in ζ_t raises the spread $R_{kt+1} - R_{t+1}$. In this respect it transmits through the economy like a financial shock. Indeed, we show later that our identified liquidity demand shock is highly correlated with credit spreads. We defer to the next section a description of the household's wage-setting and labor supply behavior.

3.4 Standard DSGE model features

Since the remaining features of the model are standard, we present only the equilibrium conditions (a full derivation of the model is available as an online appendix).

3.4.1 Intermediate goods firms: factor demands

Intermediate goods firm j chooses capital K_t^j, utilization U_t^j, and labor L_t^j to minimize costs given the relative price of the intermediate goods composite p_{mt}, the price of capital Q_t, the rental rate D_t, the real wage w_t and the desired markup ς. Following Greenwood et al. (1988), we endogenize the capital utilization decision by assuming that the depreciation rate $\delta(U_t^j)$ is an increasing and convex function of capital utilization U_t^j.

The first order conditions from the firm's cost minimization problem for K_t^j, U_t^j and L_t^j are then given by

$$\alpha \frac{p_{mt} Y_{mt}^j}{K_t^j} = \varsigma[D_t + \delta(U_t^j)Q_t] \tag{31}$$

$$\alpha \frac{p_{mt} Y_{mt}^j}{U_t^j} = \varsigma \delta'(U_t^j)Q_t K_t^j \tag{32}$$

$$(1-\alpha) \frac{p_{mt} Y_{mt}^j}{L_t^j} = \varsigma w_t \tag{33}$$

We allow ς to be smaller than the optimal unconstrained markup ϑ due to the threat of entry by imitators as is common in the literature (e.g., Aghion and Howitt (1998)).

3.4.2 Capital producers: investment

Competitive capital producers use final output to make new capital goods, which they sell to households, who in turn rent the capital to firms. Let I_t be new capital produced and p_{kt} the relative price of converting a unit of final output into new capital (the replacement price of capital), and γ_y the steady state growth in I_t. Following Christiano et al. (2005), we assume flow adjustment costs of investment: the adjustment cost function $f(I_t/(1+\gamma_y)I_{t-1})$ is increasing and concave, with $f(1) = f'(1) = 0$ and $f''(1) > 0$. The first order condition for I_t relates the ratio of the market value of capital to the replacement price (i.e. "Tobin's

Q") to investment, as follows:

$$\frac{Q_t}{p_{kt}} = 1 + f\left(\frac{I_t}{(1+\gamma_y)I_{t-1}}\right) + \frac{I_t}{(1+\gamma_y)I_{t-1}}f'\left(\frac{I_t}{(1+\gamma_y)I_{t-1}}\right)$$
$$- E_t\Lambda_{t,t+1}\left(\frac{I_{t+1}}{(1+\gamma_y)I_t}\right)^2 f'\left(\frac{I_{t+1}}{(1+\gamma_y)I_t}\right) \quad (34)$$

We assume that $\log(p_{kt})$ follows an AR(1) process with parameters ρ_{pk} and σ_{pk}. Finally, the law of motion for capital is

$$K_{t+1} = I_t + (1 - \delta(U_t))K_t \quad (35)$$

3.4.3 Price and Wage Setting

Following Smets and Wouters (2007), we assume that both nominal prices and wages are set on a staggered basis, following the "Calvo" adjustment rule. Let ξ_p be the probability a firm cannot adjust its price and let ξ_w be the probability a firm cannot adjust its wage. Conversely, let ι_p be the degree of indexing prices to past inflation and let ι_w the analogue for wages. The only difference from the standard model is that households in our economy supply two types of labor, skilled and unskilled. We assume that each type of labor has the same frequency of wage adjustment.

Denoting by π_t the inflation rate and by mc_t the marginal cost of final goods producers in log-deviation from steady state, the price Phillips curve is:

$$\pi_t = \kappa mc_t + \frac{\iota_p}{1+\iota_p\beta}\pi_{t-1} + \frac{\beta}{1+\iota_p\beta}\pi_{t-1}\mathbb{E}_t\left[\pi_{t+1}\right] + \varepsilon\mu_t \quad (36)$$

with $\kappa \equiv \frac{(1-\xi_p\beta)(1-\xi_p)}{\xi_p(1+\iota_p\beta)}$ and $\varepsilon\mu_t$ is a shock to the final goods markup that follows an AR(1) process with parameters ρ_μ and σ_μ.

The unskilled wage Phillips curve is:

$$(1+\kappa_w)\tilde{w}_t = \frac{1}{1+\beta}\left(\tilde{w}_{t-1} + \iota_w\pi_{t-1} - (1+\tilde{\beta}\iota_w)\pi_t\right) + \frac{\beta}{1+\beta}\mathbb{E}_t\left[\tilde{w}_{t+1} + \pi_{t+1}\right]$$
$$+ \kappa_w(\widetilde{u}_{c,t} - \varphi\tilde{l}_t) + \varepsilon\mu_{w,t} \quad (37)$$

with $\kappa_w \equiv \frac{(1-\xi_w\beta)(1-\xi_w)}{(\varphi(1-1/\mu_w)^{-1}+1)\xi_w(1+\beta)}$ is the steady state wage markup. $\widetilde{u}_c$, $\tilde{w}$, $\tilde{l}$, are the marginal utility of consumption and unskilled wage and hours in log deviation from steady state, and $\varepsilon\mu_{w,t}$ is a shock to the wage markup that follows an AR(1) process with parameters $\rho_{\mu w}$

and $\sigma_{\mu w}$. The skilled wage Phillips curve is identical, replacing unskilled wage and hours for skilled equivalents.[13]

3.4.4 Monetary policy

The nominal interest rate R_{nt+1} is set according to the following Taylor rule:

$$R_{nt+1} = r_t^m \left(\left(\frac{\pi_t}{\pi^0} \right)^{\phi_\pi} \left(\frac{L_t}{L^{ss}} \right)^{\phi_y} R_n \right)^{1-\rho^R} (R_{nt})^{\rho^R} \tag{38}$$

where R_n is the steady state nominal rate, π^0 the target rate of inflation, L_t total employment and L^{ss} steady state employment; ϕ_π and ϕ_y are the feedback coefficients on the inflation gap and capacity utilization gap respectively and $\log(r_t^m)$ follows an AR(1) process with parameters ρ^{mp} and σ^{mp}. We use the employment gap to measure capacity utilization as opposed to an output gap for two reasons. First, Berger et al. (2015) show that measures of employment are the strongest predictors of changes in the Fed Funds rate. Second, along these lines, the estimates of the Taylor rule with the employment gap appear to deliver a more reasonable response of the nominal rate to real activity within this model than does one with an output gap.[14] In addition, we impose the zero lower bound constraint on the net nominal interest rate, which implies that the gross nominal rate cannot fall below unity.

$$R_{nt+1} \geq 1 \tag{39}$$

3.5 Resource constraints and equilibrium

The resource constraint is given by

$$Y_t = C_t + p_{kt} \left[1 + f \left(\frac{I_t}{(1+\gamma_y)I_{t-1}} \right) \right] I_t + G_t \tag{40}$$

where government consumption G_t is financed by lump sum taxes and follows (in logs) an AR(1) process:

$$\log(G_t/(1+\gamma_y)^t) = (1-\rho_g)\bar{g} + \rho_g \log(G_{t-1}/(1+\gamma_y)^{t-1}) + \epsilon_t^g, \tag{41}$$

---

[13]In estimating the model we introduce wage markup shocks to the wage setting problem of unskilled labor only, so the markup for skilled labor is constant at its steady state level.

[14]Part of the problem may be that the behavior of the flexible price equilibrium output is quite complex in the model, particularly given the endogenous growth sector. As a robustness check on our specification of the Taylor rule, we estimate a version of the model in which we adjust the employment gap for demographic effects on the size of the labor force; our estimation results are robust to this change.

The market for skilled labor must clear:

$$L_{st} = [Z_t - A_t] L_{sat} + L_{srt} \qquad (42)$$

Finally, the market for risk-free bonds must clear, which implies that in equilibrium, risk-free bonds are in zero net supply $B_t = 0$.

This completes the description of the model.

4 Estimation

4.1 Identification of the Endogenous Productivity Mechanism

What is different about our model from the standard DSGE framework is the presence of the endogenous productivity mechanism. To identify this mechanism we use data on R&D expenditures and the restrictions of the model. Ideally we would like to also use data on expenditures on technological adoption, but we are not aware of a data series that could serve this purpose. In our estimation procedure we therefore treat the stock of adopted goods A_t as a latent variable and use the Kalman filter/smoother to identify it's temporal evolution. To do so we estimate the model using the "standard" macro series used to estimate DSGE models (see below), augmented by a measure of US R&D expenditures.

In the model the driving force of long term growth (that is to say, growth on the balanced growth path) is the endogenous productivity mechanism driven by the adoption of new technologies, A_t. In the estimation procedure, the steady state growth rate of A_t is identified by estimating the trend growth rate in real output. For this reason we do not detrend the variables we use in our estimation procedure, since the information contained in the trend identifies the long run behavior of A_t. From Equations 17 and 21, the trend growth rate in A_t together with the calibrated adoption lag ($\bar{\lambda}$) and obsolescence (ϕ) parameters pin down the steady state level of skilled labor used in the adoption of new technologies. Further, since in the balanced growth path A_t and Z_t grow at the same rate, Equation 16 pins down the steady state level of skilled labor used in R&D. Finally, since R&D expenditure in the model is equal to the product of the skilled labor wage and labor hours used in R&D, the steady state skilled labor wage is pinned down by targetting a steady state ratio of R&D expenditure to GDP of 1.8%, which is the calibration target for the elasticity parameter ρ_λ (the details of the model's steady state are in the Online Appendix).

How do we separate endogenous from exogenous variation in productivity? The endogenous component varies with the stock of adopted technologies A_t, which from Equation 21

in turn depends on the stock of technologies Z_t and labor used in adoption. Our model decomposes the utilization adjusted Solow residual into an exogenous, stationary component (the pure TFP shock θ_t) and an endogenous component, A_t. Variations in R&D expenditure do not mechanically translate into TFP for two reasons. The first is that technologies need to be adopted, and this occurs with an uncertain lag. The second is the R&D productivity shock, which is identified from the first order condition for R&D expenditures, Equation 14.

4.2   Estimation Methodology

We estimate our model using Bayesian methods (see for example An and Schorfheide (2007), Smets and Wouters (2007) and Justiniano et al. (2010)). We estimate using quarterly data from 1984:I to 2008:III on eight US series: the growth rates of real output, consumption, investment, and real wages, the log-level of hours worked, inflation (as measured by the growth rate of the GDP deflator), the nominal risk-free interest rate and the growth rates of expenditures on R&D by US corporations. Unlike the other series, R&D expenditures are annual. We deal with the mixed frequency of the data in estimation using a version of the Kalman filter adapted for this purpose. Appendix A.1 describes the data in detail.

We do not use data beyond 2008:III in the estimation of the structural parameters because the zero lower bound (ZLB) on the nominal interest was binding after that period, rendering estimation using a log-linear approximation of our baseline model problematic. However we do use the data from 2008:III to 2015:IV to identify shocks and other latent variables of our model, including the endogenous component of TFP. We do so by modifying the standard log-linear approximation of the model with the technique introduced by Guerrieri and Iacoviello (2015) to deal with occasionally binding constraints, as described in Appendix A.2.

We estimate all the standard parameters that appear in the conventional DSGE model with the exception of the markup in the final goods sector. The presence of an additional markup in the intermediate goods sector along with the elasticity of substitution between goods in this sector makes identification problematic, leading us to calibrate these parameters.

Of the four technological change parameters, we estimate the elasticity of the creation of new technologies with respect to R&D. The first order condition for R&D (14) allows us to identify ρ_z, along with the shock to R&D productivity, χ_t. We calibrate the other technological parameters using evidence from other studies and long-run restrictions. Additionally, we use the panel data evidence from section 2 as a check that the elasticity of the diffusion

25

rate with respect adoption expenditure ρ_λ that results from our model is reasonable. We next discuss the calibrated parameters and describe the prior and posterior estimates of the remaining parameters.

4.3 Calibrated parameters

As is standard we calibrate the steady state depreciation rate δ and the steady state ratio of government expenditures to output to match the data. The markups on final (μ) and intermediate goods (ς) are set to 1.1 and 1.18 respectively. We set the markup on final goods towards the lower end and the markup on intermediate goods in the middle of the range of estimates in the literature.[15] We set markups conservatively low because the R&D share of GDP is increasing in markups and decreasing in ρ_λ, so setting markups low makes our calibration of ρ_λ more conservative. The parameter ϑ is set to 1.35 to produce an elasticity of substitution of 3.85 between intermediate goods, in line with the estimates from Broda and Weinstein (2006). We calibrate the steady state liquidity demand shock $\bar{\zeta}$ to match an annual liquidity premium of 50 bps, consistent with the estimates in Del Negro et al. (2017).

The three endogenous technological change parameters we calibrate include: the steady state adoption lag $\bar{\lambda}$, the obsolescence rate $(1 - \phi)$, and the elasticity of the adoption probability λ with respect to adoption expenditures, ρ_λ. $\bar{\lambda}$ is set to produce an average adoption lag of 5 years which is consistent with the estimates in Cox and Alm (1996), Comin and Hobijn (2010) and Comin and Mestieri (2015). $1 - \phi$ is set to 2% (quarterly) which is the average of the estimates of the obsolescence rate that come from the rate of decay of patent citations (see Caballero and Jaffe (1993)) and patent renewal rates (Bosworth (1978)). Finally, ρ_λ is set to 0.925 to induce a ratio of private R&D to GDP consistent with post-1970 U.S. data (approximately 1.8% of GDP). As discussed earlier, though, since ρ_λ governs the cyclicality of technological diffusion, we check that cyclicality in the model is consistent with the panel data evidence. Table 3 presents the calibrated parameters and their values.

4.4 Parameter estimates

Table 4 presents the prior and posterior distributions for the parameters that we estimate. For the conventional parameters we use similar priors to the literature (e.g. Justiniano et al. (2010)). For the new parameter we estimate, the elasticity of R&D parameter (ρ_z) we use

---

[15]Jaimovich (2007) reports markup estimates in gross output data between 1.05 and 1.15 and in value added data from 1.2 to 1.4.

Table 3: Calibrated Parameters

Parameter	Description	Value
δ	Capital depreciation	0.02
$\frac{G}{Y}$	SS government consumption/output	0.2
μ	SS final goods mark up	1.1
ς	SS intermediate goods mark up	1.18
ϑ	Intermediate goods elasticity of substitution	1.35
$\bar{\zeta}$	SS liq. demand	50/4 bps
$1 - \phi$	Obsolescence rate	0.08/4
$\bar{\lambda}$	SS adoption lag	0.2/4
ρ_λ	Adoption elasticity	0.925

a beta prior centered around a mean of 0.6, which is at the lower end of estimates provided in Griliches (1990).

Table 4: Prior and Posterior Distributions of Estimated Parameters

Parameter	Description	Prior			Posterior	
		Distr	Mean	St. Dev.	Mean	St. Dev.
ρ^R	Taylor rule smoothing	Beta	0.70	0.15	0.833	0.0007
ϕ_π	Taylor rule inflation	Gamma	1.50	0.25	1.638	0.0540
ϕ_y	Taylor rule labor	Gamma	0.30	0.10	0.385	0.0049
φ	Inverse Frisch elast.	Gamma	2.00	0.75	2.726	0.7445
f''	Investment adj. cost	Gamma	4.00	1.00	5.630	0.6452
$\frac{\delta'(U)}{\delta}$	Capital util. elast.	Gamma	4.00	1.00	4.045	0.9615
ξ_p	Calvo prices	Beta	0.50	0.10	0.932	0.0002
ξ_w	Calvo wages	Beta	0.75	0.10	0.933	0.0008
ι_p	Price indexation	Beta	0.50	0.15	0.252	0.0115
ι_w	Wage indexation	Beta	0.50	0.15	0.386	0.0175
μ_w	SS Wage Markup	Normal	0.15	0.05	0.151	0.0025
b	Consumption habit	Beta	0.70	0.10	0.486	0.0017
ρ_z	R&D elasticity	Beta	0.60	0.15	0.376	0.0105
α	Capital share	Normal	0.30	0.05	0.200	0.0009
$\tilde{\beta}$	Discount factor	Gamma	0.25	0.10	0.512	0.0110
$100 * \gamma_y$	SS output growth	Normal	0.46	0.03	0.454	0.0005

Most of our estimates are similar to those in the literature. The price and wage rigidity parameters are higher than the estimates in, for example, Smets and Wouters (2007), likely reflecting that inflation was low and stable over our sample.[16] Our estimate of the elasticity

---

[16]In a more recent study, Del Negro et al. (2015) report estimates of wage and price rigidity parameters of similar magnitude to our results.

of new technologies with respect to R&D, ρ_z, is 0.376, which is somewhat below the Griliches (1990) estimates. The value of ρ_z is identified from the co-movement between the series on R&D expenditures and the model estimates of the value of un-adopted technologies. The discrepancy in the estimate of ρ_z may reflect the fact that (effectively) we use quarterly data while the literature uses annual data: one would expect greater diminishing returns to R&D (and hence smaller co-movement with the market value of new technologies) at higher frequencies due to frictions in adjusting skilled labor input.

Finally, with respect to the shocks, we find lower estimates of the persistence of exogenous TFP than in the literature (the prior and posterior distributions of parameters of the shock processes can be found in Table 9 in the appendix). This reflects the fact that our model produces significant endogenous persistence in TFP.[17]

4.5 Sources of Variation and the Liquidity Demand Shock

Here we establish that demand shocks - and the liquidity demand shock in particular - is not only an important source of variation in output, but an important source of variation in endogenous productivity as well. In addition, this importance is especially pronounced during recessions.

Our key finding is that the liquidity demand shock is the most important source of cyclical variation as well as the most important driver of recessions, including the Great Recession. We ascertain the relative importance of each shock by calculating a set of variance decompositions. To do so we simulate the model a large number of periods taking into account the ZLB as described in Appendix A.2. Table 5 presents the results. There are several important takeaways. As noted, liquidity demand is the main source of variation. It explains 42.7% of output growth, 54.7% of hours and 52.3% of endogenous productivity. In addition, the "demand" shocks overall are dominant. The two main demand shocks (liquidity demand and money) combined account for more than half the volatility of output and more than two thirds of the variation in hours and endogenous productivity variation. The next most important shock is the exogenous component of total factor productivity, which accounts for 18.5% of output variation, 10.8% of hours variation, and 9.5% of the variation in endogenous TFP.

Next we show that the liquidity demand shock is by far the most important shock driving recessions. Figure 5 plots the historical evolution of per capita output growth as well as

---

[17]Estimating our model without the endogenous technology mechanism over the same sample period gives an estimate for the AR(1) parameter of the exogenous TFP process of 0.9685 versus 0.953 in the model with endogenous technology.

Table 5: Variance Decomposition (%)

Variables	Liquidity Demand	Money	Govt Exp	Price of Capital	TFP	R&D	Mark up	Wage mark up
Output Growth	42.7	13.6	16.6	6.1	18.5	0.0	2.2	0.3
Consumption	45.8	14.6	16.7	0.7	19.9	0.0	1.9	0.3
Investment	16.8	4.9	2.9	64.6	7.7	0.3	1.8	0.9
Inflation	0.3	0.0	0.2	0.0	3.3	0.0	80.6	15.5
Nominal R	35.4	33.0	1.3	2.0	7.5	0.3	14.2	6.4
Hours	54.7	18.2	5.7	4.6	10.8	0.2	4.3	1.5
Endo. TFP	52.3	18.0	4.2	0.7	9.5	8.4	1.6	5.3

Variance decomposition with ZLB (10,000 simulations, HP filtered series, filter parameter = 1600).

the components that are accounted for by the liquidity demand and the exogenous TFP shock, the disturbance that is second most important in driving recessions.[18] In each of the three recessions, the liquidity demand shock accounts for most of the decline in output. In addition to moving the economy to the trough during the Great Recession, the liquidity demand shock is also responsible for its duration. In particular, the historical decomposition shows that if the only shock that had hit the economy was the liquidity demand shock the recovery of output growth after the GR would have been even slower. It is important to mention that the model does not require a sequence of large negative shocks to explain the drop in output and binding ZLB observed during the GR. Instead, one large shock (of the order of 4.9 standard deviations) is sufficient . This is due partly to the estimated persistence of the liquidity shock but importantly also to the amplification and persistence generated by the endogenous productivity mechanism, as illustrated by the impulse responses plotted in Figure 7.

Finally, we present some evidence in support of our interpretation of the liquidity demand shock as a financial shock. As noted in Section 3.3, a negative liquidity demand shock causes the spread between the return on capital and the riskless rate to widen. Figure 6 compares the spread implied by our estimated liquidity shocks to the spread between the twenty year Baa corporate bond rate and the ten year US treasury bond rate. To compute the model spread, we use the expectations hypothesis to convert the one period spreads implied by the model to the maturity of the Baa spread.

---

[18]The decomposition takes into account the ZLB (as described in Appendix A.2), which makes the model nonlinear for the period 2008:I-2015:IV. Because of this nonlinearity, the sum of the contribution of each shock does not equal the value of the smoothed variable being decomposed (output growth in this case) for the mentioned period. This "nonlinear residual" emerges because the interaction between shocks is relevant in nonlinear models. However, our results indicate that the only shock that moves the economy to the ZLB is the liquidity demand shock. We therefore assign the nonlinear residual to this shock.

Figure 5: Output Growth Decomposition
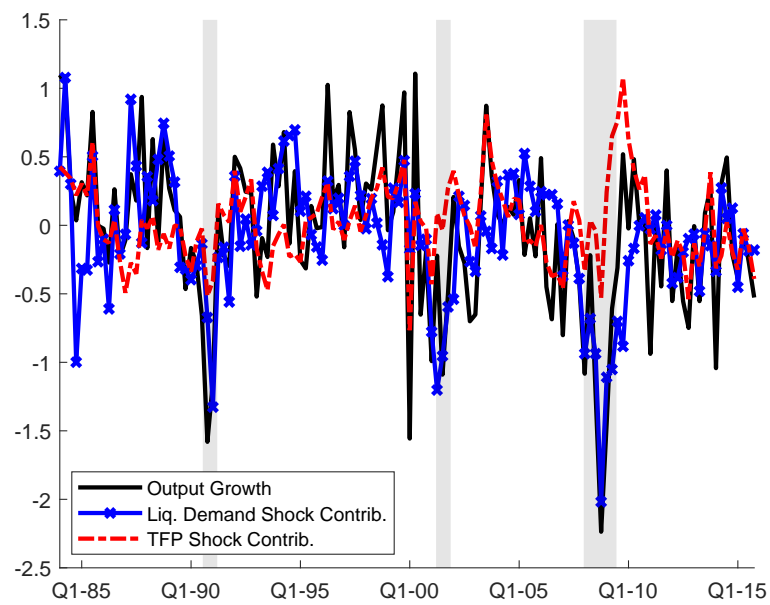
Data sources are described in Appendix A.1. Smoothed shocks from model estimated using data as described in Section 4.4 and Appendix A.1.

Figure 6: Model Spread vs Data
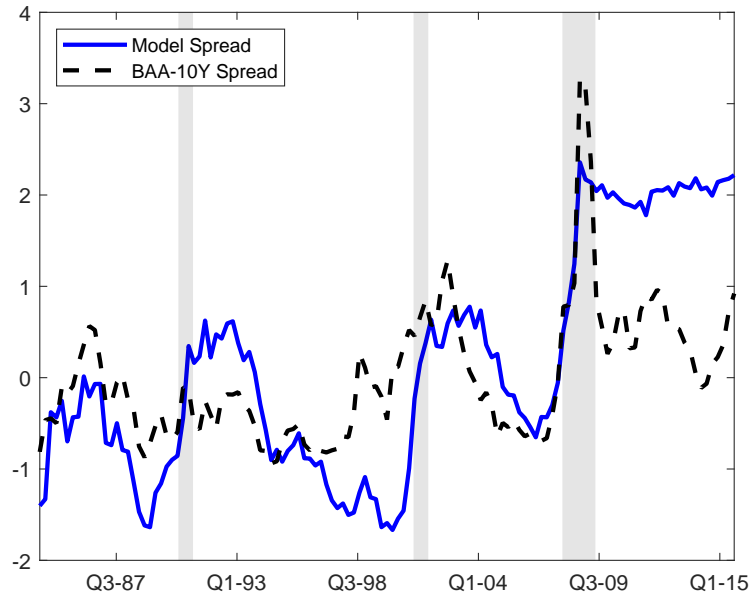
Source: Moody's Seasoned Baa Corporate Bond Yield Relative to Yield on 10-Year Treasury Constant Maturity from FRED. Both spreads series are demeaned. Model implied long run spreads are computed using the expectations hypothesis. To be consistent with the Baa spread, we use a 20-year horizon for the capital return and a 10-year one for the risk free rate.

The model estimated spread has a correlation of 0.66 with the bond spread. It displays similar countercyclical movement over recessions and expansions. Also, the magnitude of the model and data spreads are similar. The one important difference, however, is that the model spread displays more persistence following the Great Recession than the Baa spread. Del Negro et al. (2017) obtain similar results for their measure of the model spread. One possibility is that the Baa spread, which reflects credit costs for publicly traded companies, is not always representative of the wedge between borrowing and safe rates that households face. For example, while financial conditions may have normalized for large firms shortly after the Great Recession, households and small and medium-sized companies continued to face borrowing frictions. Though we do not report the results here, the two year personal loan rate for households has remained persistently high following the Great Recession.

Figure 7: Impulse Response to 1 std. dev. Liquidity Demand Shock
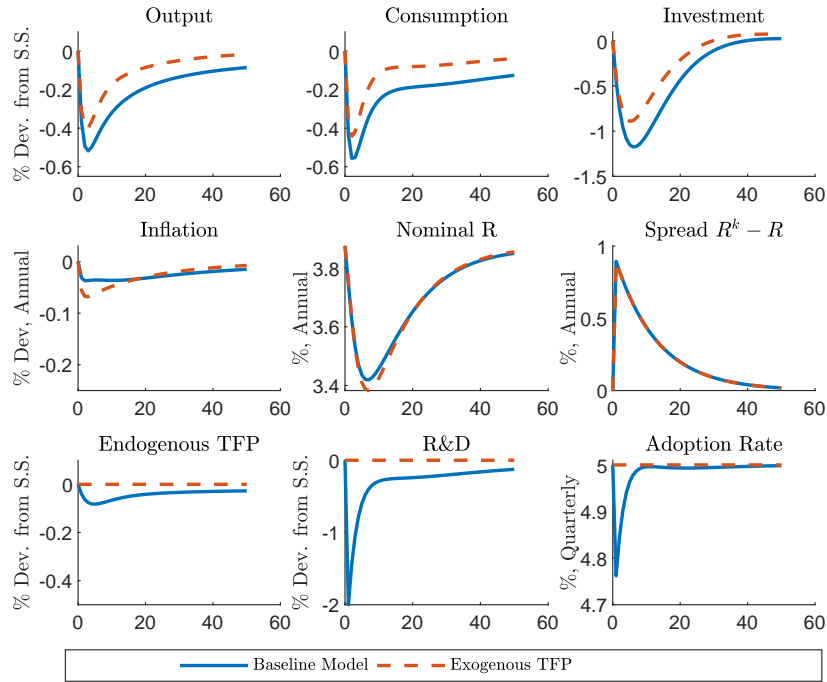
4.6 Endogenous Technology Mechanism

Before analyzing how our model can account for productivity dynamics, we do two exercises. First, in this sub-section we analyze how shocks to the economy transmit into endogenous movements in productivity. We consider a shock to liquidity demand, given the importance of this disturbance as a source of variation. In the next sub-section, we show that the cyclical movements in technology diffusion the model generates are consistent with the panel data evidence in section 2.

Figure 7 presents the responses of some key variables to a one standard deviation liquidity demand shock.[19] To isolate the effects of our endogenous productivity mechanism, we plot the responses of our model and a version where technology is purely exogenous.

An increase in the demand for the liquid asset, all else equal, induces households to reduce their consumption demand and their saving in risky assets (See equations 28-30). As a result there is upward pressure on the required return to capital, R_{kt+1}, and downward

---

[19]In the online Appendix we report the impulse response functions to the money shock and the shock to the R&D productivity.

pressure on the safe real rate R_{t+1}. The former leads to a fall in both physical investment demand as well as in the demand for productivity enhancing investments, including both R&D and adoption expenditures. The latter cushions the drop in consumption. Given nominal rigidities, the overall drop in both investment and consumption demand leads to a decline in output. The drop in productivity enhancing investments, further, induce a decline in productivity, magnifying both the overall size and persistence of the output decline relative to the version of the model where technology is exogenous.

One additional interesting result is that the endogenous productivity mechanism mutes the decline in inflation following the contractionary demand shock. As in conventional New Keynesian models, inflation declines when aggregate demand falls. However, the endogenous decline in productivity growth lessens the decline in marginal costs, which in turn dampens the decline in inflation, making it almost negligible. This feature can offer at least part of the explanation for the surprising failure of inflation to decline by any significant amount during the Great Recession.

Finally, the main part of our analysis involves analyzing productivity over a period where the ZLB is binding. Our historical decomposition, further, suggests that it is the liquidity demand shock that moves the economy into the ZLB. Accordingly it is useful to understand the implications of the ZLB for how a contractionary liquidity demand shock influences endogenous productivity. Figure 8 plots the impulse response functions with and without a binding ZLB. When the ZLB is binding, monetary policy cannot accommodate a recessionary shock. This results in higher interest rates than when the ZLB is not binding. The higher real rates amplify the drops in investment, R&D and adoption intensity. In the short term, this leads to lower aggregate demand and a larger output drop. It also leads to larger declines in the growth rate of the number of adopted technologies and to lower levels of TFP in the medium and long term.[20]

4.7 Technology Diffusion: Model vs. Data

In this section we investigate whether the cyclicality of diffusion in our model is reasonably similar to that in the micro data. This exercise is important to validate the realism of the diffusion mechanism in the model. To do so, we derive a model analogue to the estimating Equation (3) of Section 2.

---

[20]One interesting observation on how the endogenous technology mechanism interacts with the ZLB is that in contrast with standard neo-keynessian models with exogenous technology, in our model once the economy enters in the ZLB region, it naturally remains there without the need of additional contractionary shocks. This is the case because of the additional amplification and propagation generated by the endogenous contraction in TFP. This feature is also analyzed in Benigno and Fornaro (2017).

Figure 8: Liquidity Demand Shock and the ZLB
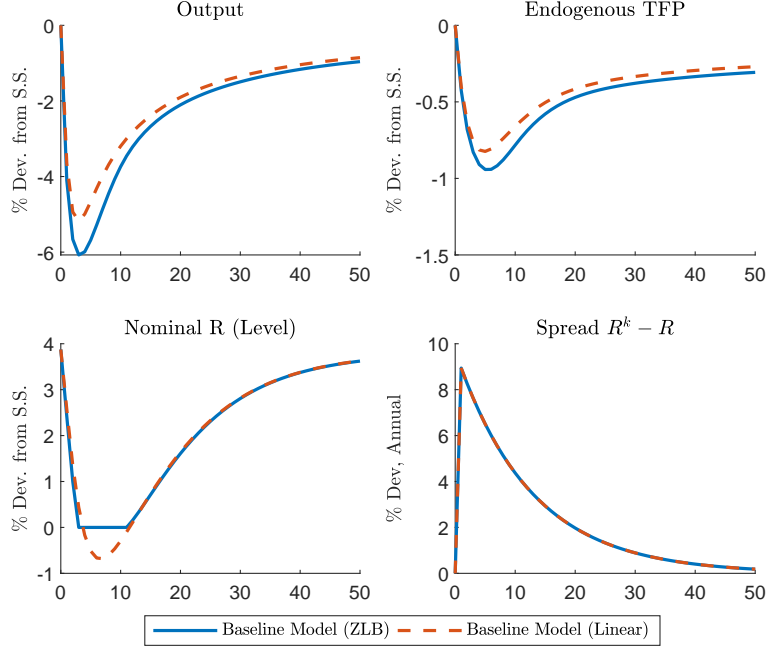
Impulse response functions to a large (10 standard deviation) positive liquidity demand shock.

Recall from Equation (24) that due to the geometric nature of diffusion in the model, the vintage of the technology affects the speed of diffusion. In particular, absent any cyclical fluctuations, r_{t+k}^t is increasing in k, the time elapsed since invention. Accordingly, we define the vintage effect as the speed that a given vintage would have in the absence of business cycle fluctuations; that is, if the adoption rate was equal to its constant steady state level, $\bar{\lambda}$. Formally, the vintage effect on speed is

$$\overline{Speed}_{t+k}^t = \log\left(\frac{1 + \bar{\lambda}/r_{t+k-1}^t}{1 - \bar{\lambda}}\right) \tag{43}$$

Using this definition, we model the regression equation for the effect of the cycle on the speed of diffusion as

$$Speed_{t+k}^t = \alpha + \overline{Speed}_{t+k}^t + \beta \hat{y}_{t+k} + \epsilon_{t+k} \tag{44}$$

where $\hat{y}_{t+k}$ denotes the same measure of the output gap as in Section 2. Note that the vintage-effect on the speed is akin to the lag control we included in regression (3) to capture

34

Table 6: Cyclicality of diffusion speed: model versus data

	Data	Model
$\hat{\beta}$	3.73	1.88
	[1.7,5.8]	[1.53,2.23]

$\hat{\beta}$ is the estimated elasticity of speed of diffusion with respect to the output gap. Numbers in brackets are 95% confidence intervals. Model confidence intervals are constructed by averaging across regression results of 100 period long sub-samples of the model simulated data.

the deterministic effect of the lag on the average speed of diffusion.

We estimate equation (44) in a synthetic sample constructed from 100,000 period simulations of our estimated model. As described in Appendix A.3, we estimate the panel by weighting each observation by the share of age k technologies in the steady state. Table 6 reports the point estimates together with the 95% confidence intervals of the estimates of β in the data simulations as well as from the panel estimates in section 2. The point estimate in the model is smaller (1.88 versus 3.73) but falls within the 95% confidence interval of the point estimate in the data.

Table 7 reports a sensitivity analysis to ρ_λ of the cyclicality of the speed of diffusion and our calibration target, the share of R&D expenditures in GDP. To obtain these results we reestimate the model for alternative values of ρ_λ and simulate the model using the resulting parameter estimates. The lower the values of ρ_λ the higher the share of R&D in GDP in steady state. This is because, for a given λ, a lower ρ_λ produces greater curvature in the value of unadopted technologies raising the rents earned from engaging in R&D. For a ρ_λ of 0.85, R&D represents 3.3% of GDP while for a ρ_λ of 0.95 the R&D share is 1.2%. The cyclicality of the speed of diffusion increases with the value of ρ_λ. For ρ_λ equal to 0.95, the elasticity of the speed of diffusion with the cycle produced by the model is 2.14 which falls within the confidence interval for the technology panel in Section 2. However, for values of ρ_λ smaller than our baseline of 0.925 the elasticities of the speed of diffusion with the cycle fall outside the confidence interval. For example, for a ρ_λ of 0.85, the elasticity is 1.19. Nonetheless, for completeness, we show in the Online Appendix that our main results on the evolution of endogenous productivity are robust to calibrating ρ_λ to 0.85.

We conclude from this analysis that the cyclical response of the speed of diffusion in our model is similar to that estimated in the panel data, falling in the lower part of the confidence interval.

Table 7: Effect of varying ρ^λ on estimated elasticity $\hat{\beta}$ and R&D expenditure share

ρ^λ	$\hat{\beta}$	R&D/GDP
0.85	1.19	3.3%
0.9	1.59	2.3%
0.925	**1.88**	**1.8%**
0.95	2.14	1.2%

Values in bold are for baseline calibration (ρ^λ=0.925). Values for alternative ρ^λ are obtained by re-estimating and simulating the model as described in Sections 4 and 4.7.

5   Analysis of Productivity Dynamics and Inflation

We now explore the model's implications for the evolution of productivity, with particular emphasis on the periods before, during and after the Great Recession. We focus on TFP but also consider labor productivity. The latter allows us to consider the impact of the demand shortfall during the Great Recession on the supply side that operates via the conventional capital accumulation channel (as emphasized by Hall (2014) and others), as well as our endogenous productivity channel.

To begin, we use equation (10) to derive the following expression that links labor productivity with TFP and capital intensity:[21],[22]

$$\frac{Y_t}{L_t} = \underbrace{\theta_t \cdot (A_t)^{\vartheta-1}}_{TFP} \cdot (U_t K_t/L_t)^\alpha . \tag{45}$$

The first two terms capture total TFP, which is the product of an exogenous component (θ_t) and an endogenous one ($(A_t)^{\vartheta-1}$). The third term measures capital intensity which includes both capital per hours worked and the capital utilization rate.

Figure 9 plots the evolution of (detrended) labor productivity together with TFP and the endogenous component of TFP. Labor productivity corresponds exactly to the data. The other two series are identified from the model. It is worth noting, though, that the evolution of TFP and labor productivity are qualitatively similar.[23]

---

[21]This expression holds to a first order approximation.

[22]We focus on labor productivity for two reasons. First, our measure of capital includes residential investment. Therefore, there is a discrepancy between our measure of TFP and that from standard sources (e.g., BLS). Second, labor productivity also captures the effect of variation in capital per hour. This is another channel by which fluctuations in demand can affect the potential supply in the economy.

[23]One can obtain the capital intensity component of labor productivity from the figure by taking the difference between labor productivity and TFP.

**Figure 9:** Endogenous TFP, TFP and Labor Productivity
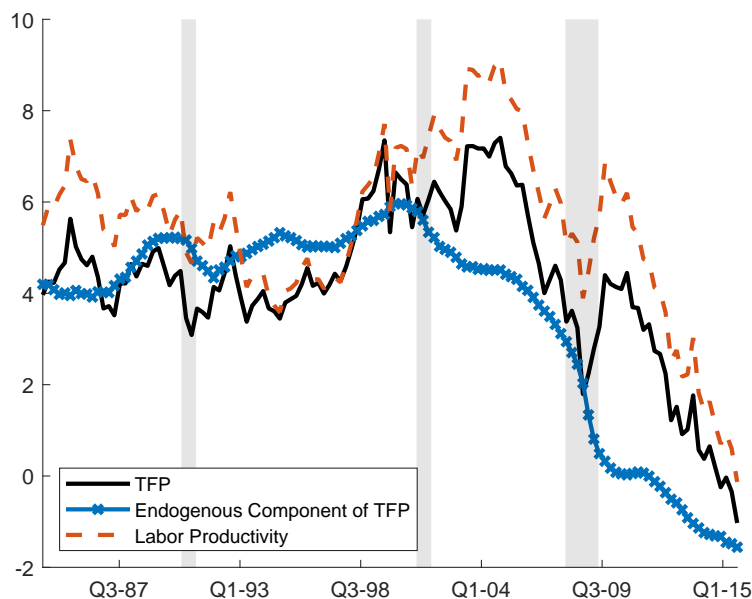
Labor productivity is GDP divided by hours worked (see Appendix A.1 for data sources). Smoothed shocks from model estimated using data as described in Section 4.4 and Appendix A.1.

Except for the middle to late 1990s, the endogenous component of TFP accounts for much of the cyclical variation in TFP. The model attributes the rise in TFP during the late 90s mainly to its exogenous component; the labor productivity surge in this period is explained by both exogenous innovations to TFP and capital deepening. After 2000, however, the endogenous component plays a predominant role in the evolution of TFP. Importantly, the endogenous component explains most of the decline in TFP between 2005 and 2008, as well as the decline during and after the Great Recession. In particular, between the starting point of the recent productivity slowdown, 2005, and the end of our sample, 2015, total TFP declined by approximately 8 percentage points (relative to trend). The endogenous component accounts for around 6 percentage points of decline. This factor also accounts for most of the drop in labor productivity, which declined 8.5 percentage points over the same period. A drop in capital intensity after 2009 mainly accounts extra the drop in labor productivity relative to TFP (consistent with Hall (2014)).

While endogenous TFP declines steadily after 2005, the main sources of the drop varies

37

Figure 10: Endogenous TFP Decomposition
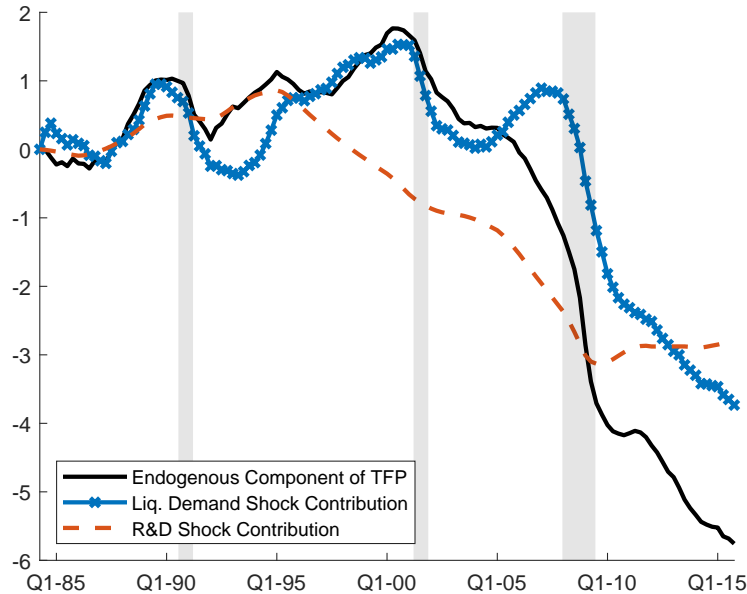
Smoothed variables from model estimated using data as described in Section 4.4 and Appendix A.1.

over time. Figure 10 presents a historical decomposition of endogenous productivity that isolates the effects of the two shocks that were the main causes of the decline: (i) shocks to the productivity of R&D and (ii) the liquidity demand shock. We note first that the liquidity demand shock accounts for nearly all of the decline in endogenous TFP after the start of the recession at the end of 2007. This result is consistent with our earlier findings that: (i) the liquidity demand shock was the main disturbance driving the recession (see Figure 5); and (ii) the liquidity demand shock has a significant impact on endogenous TFP, especially at the ZLB (see Figure 8).

In the period just prior to the Great Recession, 2005-2007, however, the liquidity demand shock is unimportant. Instead the decline in endogenous TFP is mainly the result of negative shocks to the productivity of R&D. The downward trend in R&D productivity actually begins in the mid 1990s, which is consistent with Gordon (2012)'s hypothesis of a secular decline in the contribution of technological innovations to productivity. After a brief upturn following the 2000-01 recession, shocks to R&D productivity induce a sharp downturn in TFP from 2005 until the height of the crisis.

Intuitively, the exogenous decline in R&D productivity generated fewer technologies for a given level of R&D spending, which ultimately slowed the pace of adoption of new technologies. The slow diffusion of technologies generates a lag between the decline in R&D productivity and the reduction in TFP growth. In this respect, a shock to R&D productivity is very different from a shock to exogenous TFP, which shows up immediately in measured TFP. An additional difference comes from the identification of the shocks. While exogenous TFP is identified from the Solow residual, shocks to R&D productivity are identified from the difference between observed R&D and R&D predicted by the free entry condition (14). The magnitude of the decline in R&D around the 2001 recession indicates a significant drop in R&D productivity. In the next sub-section, we present direct evidence that supports this finding.

We next explore the relative importance of the specific mechanisms that drove endogenous TFP. From equation (21), fluctuations in the stock of adopted technologies, A_t, (and hence endogenous TFP), are driven by the product of two factors: the adoption rate λ_t and the total stock of unadopted technologies, $Z_t - A_t$. Fluctuations in λ_t reflect the effect of cyclical variation in adoption intensity on endogenous productivity while fluctuations in Z_t reflect the effect of cyclical variation in R&D. To analyze the relevance of these two channels, Figure 11 plots (relative to trend) the evolution of Z_t, A_t and λ_t.[24] Note that the evolution of A_t mirrors the evolution of endogenous productivity ($A_t^{\vartheta-1}$).

We emphasize several points. First, cyclicality in λ_t is the main driver of cyclical fluctuations in endogenous productivity. That is, λ_t co-moves closely with A_t while Z_t does not. During each of the recessions, λ_t declines along with A_t, implying that the slowdown in adoption activity in turn accounts well for the the cyclical contraction in endogenous TFP. These results are consistent with our earlier findings that: (i) liquidity demand shocks are important drivers of recessions (see Figure 5) and (ii) that these shocks can induce contractions in adoption rates and endogenous productivity (See Figure 7).

Fluctuations in Z_t, however, also play a role in the evolution of endogenous productivity. Following the 2000-01 recession there is a steady decline in Z_t, consistent with the negative shocks to R&D productivity over this period that Figure 10 identifies.[25] This drop in Z_t, in turn, helps account for the pre-Recession drop in productivity that Fernald emphasizes, complementing the analysis of Figure 9. After the start of the Great Recession, though, the contraction in the adoption rate becomes the main driver of the productivity decline. The failure of the adoption rate to return to normal levels, after a brief recovery in 2010, is the

---

[24]For λ_t we plot on the right-hand axis the level of the quarterly adoption rate.

[25]The decline in endogenous productivity induced by the negative shocks to R&D productivity lags the decline in Z_t (compare figures 9 and 10) due to the lags in the adoption process.

Figure 11: Sources of Endogenous Technology
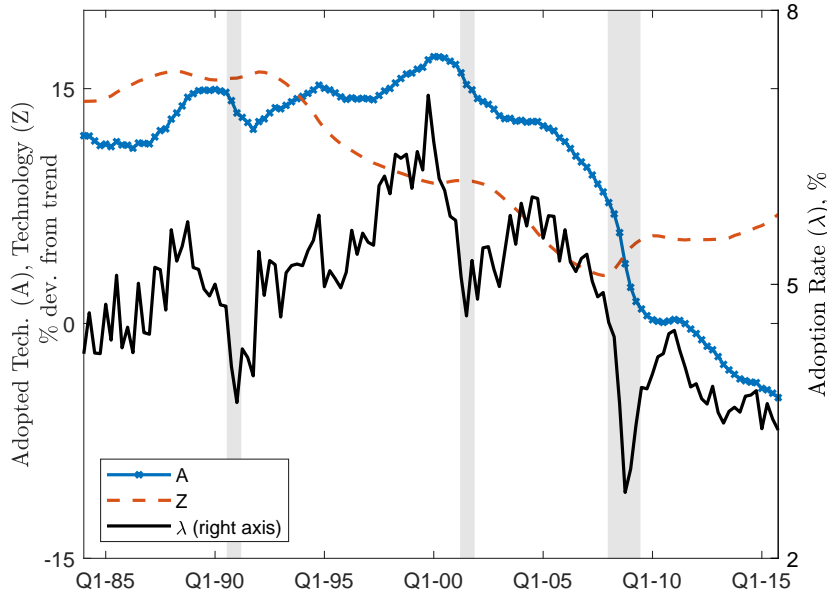
Smoothed variables from model estimated using data as described in Section 4.4 and Appendix A.1.

reason endogenous TFP continues to decline.

Interestingly, while λ_t remains low following the Great Recession, the stock of unadopted technologies, $Z_t - A$, reaches a peak over the sample. This occurs mostly because the stock of adopted technologies, A_t, declines, but also because there is a modest increase in Z_t. This finding is consistent with the evidence by Andrews et al. (2015) that suggests that innovation by leading edge firms continued after the Great Recession but adoption by followers slowed. An important implication is that the economy may not be doomed to low productivity growth for the foreseeable future. Given the high stock of unadopted technologies, to the extent enhanced economic activity pushes up the adoption rate, productivity growth should pick up. Conversely, if the economy continues to stagnate, adoption rates will remain low, keeping productivity growth low.

5.1 Evidence on the two key findings

In the previous section we presented two main findings regarding the recent productivity slowdown: first, the decline in R&D productivity after the 2001 recession contributed to the

pre-Great Recession slowdown that began in 2005. Second, the drop in adoption intensity during and after the Great Recession was mainly responsible for the rest. We conclude our analysis by providing independent historical evidence in support of these two conclusions.

First we focus on the evolution of R&D productivity. A natural way to measure R&D productivity is by the number of patent applications relative to the number of R&D researchers. Patent applications are a proxy for R&D outputs while researchers employed is a proxy for its inputs. Because the outcomes of today's R&D efforts may lead to applications at some point in the near future, we propose the following measure of the average productivity of R&D in years $t-1$ and t :

$$R\&D\_prod_{t-1,t} = \log\left(\frac{Patents_t + Patents_{t-1}}{2 * R\&D\_emp_{t-2}}\right) \tag{46}$$

where $Patents_t$ is the number of patent applications in year t, and $R\&D\_emp_{t-2}$ denotes the number R&D researchers at year $t-2$.[26]

Figure 12 plots the linearly detrended level of $R\&D\_prod_{t-1,t}$ together with the average of the estimated log-level of the model measure of R&D productivity, χ_t, between years $t-1$ and t (which we label $\bar{\chi}_{t-1,t}$).[27] The main observation is that the model and direct measures of R&D productivity evolve similarly around the three recessions in the sample. Both series drop significantly around 1991, and 2001. In particular, the magnitude of the decline after 2001 is similar in the model estimates of $\bar{\chi}_{t-1,t}$ and in the data proxy, $R\&D\_prod_{t-1,t}$. This finding supports our model's prediction that the pre-GR productivity slowdown may partly reflect the decline in R&D productivity after 2001. The measure of $R\&D\_prod_{t-1,t}$ is also consistent with our finding that R&D productivity was relatively high during and after the GR. Admittedly, the model estimate, $\bar{\chi}_{t-1,t}$, remains higher than $R\&D\_prod_{t-1,t}$ during 2009 and 2010, but overall the patterns in both series between 2008 and 2013 are similar. Based on these observations, we conclude that the independent measure of R&D productivity is consistent with the evolution of R&D productivity that we have estimated.

As we have noted, our analysis also finds that the critical driver of the productivity slowdown during and after the Great Recession is the slowdown in the intensity of adoption of new technologies in response to the liquidity demand shock. One approach to measuring adoption activity is through the investments of companies in adopting new technologies.

[26]The patent applications come from the United States Patent and Trademark Office and measure the total number of applications in the US during the calendar year. The series on the number of researchers in the US comes from the OECD. The patterns for the number of US researchers closely resembles that for R&D expenditures in Figure 2.

[27]One reason to detrend the measure of R&D_prod are the changes in the law that strengthened patent protection during the 80s inducing patent applications.

41

Figure 12: R&D efficiency in data versus model
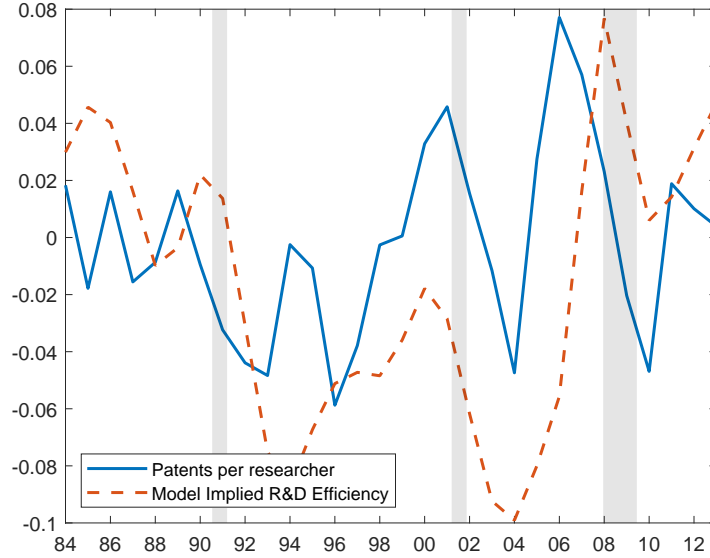
The solid line is detrended level of $R\&D\_prod_{t-1,t}$ and the dashed line is average of the estimated log-level of χ_t between years $t-1$ and t ($\bar{\chi}_{t-1,t}$). Source: Patent data from the United States Patent and Trademark Office; R&D research employment from the OECD; model implied R&D efficiency estimated using data as described in Section 4.4 and Appendix A.1.

While there does not exist a measure that covers all adoption expenditures for the US economy, the Association of University Technology Managers (AUTM) compiles the revenues by universities and research hospitals from selling licenses of technologies to companies.[28] The expenditures to license new technologies are one component of firms' overall investments in technological adoption.

Figure 13 shows the evolution of linearly detrended licensing revenues from 1995 to 2014 together with linearly detrended GDP.[29] The two series are highly correlated with a coefficient of 0.69. Both around 1999 and 2006 licensing fees start to decline in a protracted way coinciding with the cyclical declines in GDP. The decline in licensing fees continued after the Great Recession and by the end of the sample in 2014 there was no sign of a recovery in the revenues from licensing university technologies. This evidence is thus consistent with

---

[28]Approximately 180 institutions complete the survey. Their combined R&D budgets in 2011 was $60 billion of which $53 billion corresponded to universities. This sample represents a large majority of total R&D activity by higher education institutions which according to the NSF amounted to $62 billion in 2011.

[29]Both series are deflated by the GDP deflator and scaled by population older than 16 years old.

Figure 13: Revenues from Licensing Technologies, 1993-2013
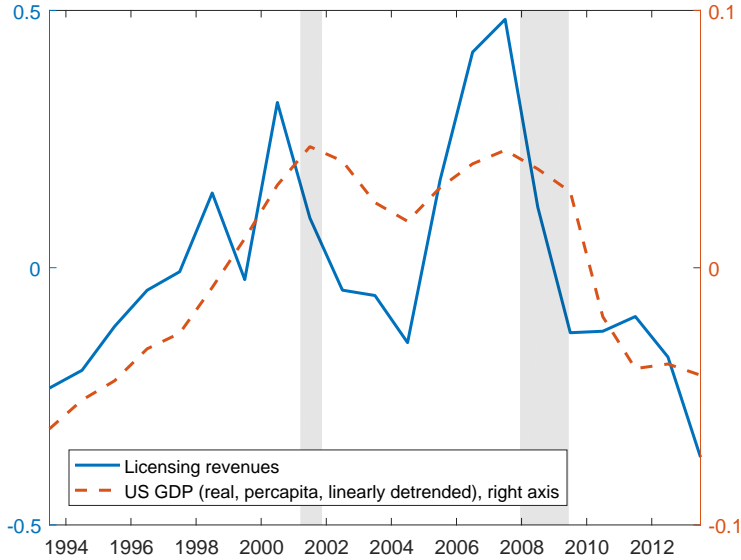
Source: Licensing revenues from Association of University Technology Managers.

our key finding that a significant and protracted decline in adoption activity underlies the slowdown in productivity growth since 2008.

5.2 Inflation

Finally, we explore the extent to which the endogenous productivity mechanism can help account for the higher than expected inflation during and after the Great Recession. To do so, we take the structural shocks identified from our baseline model with endogenous productivity and then feed them into the model with exogenous productivity. Figure 14 then reports the behavior of (GDP deflator) inflation for the baseline model versus the model with exogenous productivity. From the beginning of the recession through 2011, the inflation rate is very similar in each case. However, starting in in 2012, inflation in the baseline model moves persistently above its value in the exogenous productivity case. The differences range between thirty and eighty basis points.

Intuitively, the persistent endogenous decline in total factor productivity in the baseline model increases marginal cost, pushing up inflation relative to the exogenous productivity setting. Absent the endogenous productivity mechanism, GDP deflator inflation would have

Figure 14: Inflation Rates: Baseline vs Exogenous TFP Model
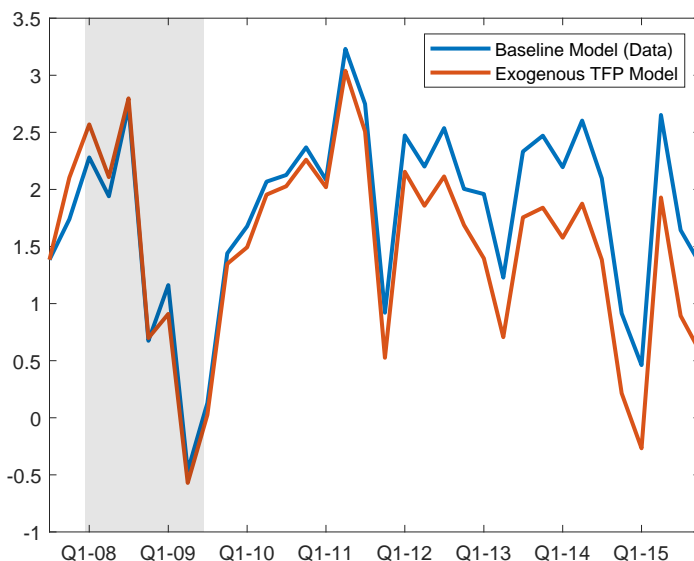
Note: The Exogenous TFP model inflation is the rate that this model predicts when the economy is hit by the shocks identified by the baseline model.

been roughly sixty basis points lower over the period 2012 to 2016.

6   Conclusions

We have estimated a monetary DSGE model with endogenous productivity via R&D and adoption. We then used the model to assess the behavior of productivity, with particular emphasis on the slowdown following the onset of the Great Recession. Our key result is that this slowdown mainly reflected an endogenous decline in the speed at which new technologies are incorporated in production. The endogenous decline in adoption, further, was a product of the recession. We also find that our endogenous productivity mechanism can help account for the productivity slowdown that preceded the Great Recession. Shocks to the productivity of the R&D process play an important role, consistent with Fernald (2014)'s view that acyclical factors were important over this period. Finally, we find a very limited role for an exogenous decline in TFP in the slowdown of productivity. Overall, the results suggest that the productivity slowdown following the start of the Great Recession

was not simply bad luck, but rather another unfortunate by-product of the downturn.

Our analysis also sheds light on two open debates. First, it provides a time series for the productivity of R&D activities that can be used to explore the hypothesis advanced by Gordon (2012) that the U.S. economy is experiencing a secular deterioration in its innovation capacity. Consistent with Gordon's hypothesis we find low levels of productivity of R&D activities between 2002 and 2007 that contributed to the decline in TFP between 2005 and 2009. However, this episode is short-lived and the estimates suggest that the slowdown in productivity reflects medium term cyclical factors rather than secular ones. We provide independent evidence on the evolution of patent applications relative to research labor that supports this interpretation of our estimates. The second relevant debate concerns the stability of inflation during the Great Recession in spite of the very significant decline in economic activity. Our model and estimates suggests that the endogenous decline in TFP has increased production costs (relative to trend) counteracting to some degree the traditional Phillips- curve effect of economic contractions on inflation.

Overall, our results emphasize the importance of the effects that demand shocks have on the supply side over the medium term. This is an important take away that can be used to explain productivity dynamics more generally.

A Appendix

A.1 Data

The data used for estimation are available from the FRED (https://research.stlouisfed.org/fred2/) and NSF (http://www.nsf.gov/statistics/) websites. Descriptions of the data and their correspondence to model observables follow (the standard macro series used are as in Del Negro et al. (2015)).[30] To estimate the model we use data from 1984:I to 2008:III. Real GDP (GDPC), the GDP deflator (GDPDEF), nominal personal consumption expenditures (PCEC), and nominal fixed private investment (FPI) data are produced by the BEA at quarterly frequency. Average weekly hours of production and nonsupervisory employees for total private industries (AWHNONAG), civilian employment 16 and over (CE16OV) and civilian noninstitutional population 16 and over (CNP16OVA) are released at monthly frequency by the Bureau of Labor Statistics (BLS) (we take quarterly averages of monthly data). Nonfarm business sector compensation (COMPNFB) is produced by the BLS every quarter. For the effective federal funds rate (DFF) we take quarterly averages of the annualized daily data (and divide by four to make the rates quarterly). Letting Δ denote the temporal difference operator, the correspondence between the standard macro data described above and our model observables is as follows:

- Output growth = 100 $\times \Delta$LN((GDPC)/CNP16OVA)

- Consumption growth = 100 $\times \Delta$LN((PCEC/GDPDEF)/CNP16OVA)

- Investment growth = 100 $\times \Delta$LN((FPI/GDPDEF)/CNP16OVA)

- Real Wage growth = 100 $\times \Delta$LN(COMPNFB/GDPDEF)

- Hours worked = 100 $\times$ LN((AWHNONAG $\times$ CE16OV/100)/CNP16OVA)

- Inflation = 100 $\times \Delta$LN(GDPDEF)

- FFR = (1/4) $\times$ FEDERAL FUNDS RATE

The R&D data used in estimating the model is produced by the NSF and measures R&D expenditure by US corporations. The data is annual, so in estimating the model and extracting model-implied latent variables (see Appendix A.2) we use a version of the Kalman filter adapted for use with mixed frequency data.

---

[30]Del Negro et al. (2015) include consumer durables in consumption as opposed to investment. Our results are robust to including them in investment. Neither approach, of course, is ideal.

A.2 Extracting Model-Implied Latent Variables during ZLB period

The piece-wise linear solution from the OccBin method developed by Guerrieri and Iacoviello (2015) can be represented in state space form as

$$S_t = C(N_t, \theta) + A(N_t, \theta)S_{t-1} + B(N_t, \theta)\epsilon_t$$
$$Y_t = H(N_t, \theta)S_t$$

Where θ is a vector of structural parameters, S_t denotes the endogenous variables at time t, Y_t are observables, and ϵ_t are normally and independently distributed shocks. N_t is a vector that identifies whether the occasionally binding constraint binds at time t and whether it is expected to do so in the future. In particular, N_t is a vector of zeros and ones indicating when the constraint is or will be binding. For example, the vector $N_t = (0, 1, 1, 1, 0, 0, 0...)$ is a situation in which the constraint does not bind at time t (denoted by the first zero in the vector), but is expected to bind in $t+1, t+2$ and $t+3$. Note that the matrices A, B and C, which in a standard linear approximation depend only on parameters are here also functions of N_t. The matrix H might also be a function of N_t because some observables might become redundant when the occasionally binding constraint binds. This is the case for the Taylor rule interest rate when the ZLB binds. OccBin provides a way of computing the sequence of endogenous variables $\{S_t\}_{t=1}^T$ and regimes $\{N_t\}_{t=1}^T$ for a given initial condition S_0 and sequence of shocks $\{\epsilon_t\}_{t=1}^T$. The vector N_t is computed by a shooting algorithm and its resulting value will depend on the initial state and shocks at time t. We refer the reader to Guerrieri and Iacoviello (2015) for a detailed description of the method. We construct the Kalman filter and smoother from the nonlinear state space representation presented above by taking advantage of the fact that a given sequence of regimes, say $\{\hat{N}_t\}_{t=1}^T$, uniquely defines a sequence of matrices $\{\hat{C}_t, \hat{A}_t, \hat{B}_t, \hat{H}_t\}_{t=1}^T$. It follows that for that specific set of regimes the state space representation becomes linear:

$$S_t = \hat{C}_t + \hat{A}_t S_{t-1} + \hat{B}_t \epsilon_t$$
$$Y_t = \hat{H}_t S_t$$

For this linear state space representation it is straightforward to compute the Kalman filter and smoother. We use this fact in our algorithm by running two blocks: (i) one in which we compute the Kalman filter and smoother for a given set of regimes $\{N_t\}_{t=1}^T$; and (ii) another where we use OccBin to compute the regimes given a sequence of shocks $\{\epsilon_t\}_{t=1}^T$. The algorithm steps are the following.

1. Guess a sequence of regimes $\{N_t^{(0)}\}_{t=1}^T$;

2. Use the guess from the previous step and define the sequence of matrices $\{C_t, A_t, B_t, H_t\}_{t=1}^T$ using OccBin;

3. With the matrices from the previous step, compute the Kalman Filter and Smoother using the observables $\{Y_t\}_{t=1}^T$, and get the Smoothed shocks $\{\hat{\epsilon}_t\}_{t=1}^T$ and initial conditions of endogenous variables;

4. Given the smoothed shocks and initial conditions from the previous step, use OccBin to compute a new set of regimes $\{N_t^{(1)}\}_{t=1}^T$;

5. If $\{N_t^{(0)}\}_{t=1}^T$ and $\{N_t^{(1)}\}_{t=1}^T$ are the same, stop. If not, update $\{N_t^{(0)}\}_{t=1}^T$ and go to step 2.

Once it converges, this algorithm yields a sequence of smoothed variables and shocks, consistent with the observables, and taking into account the occasionally binding constraint.

A.3 Comparing diffusion speed in the model to the data

We calibrate ρ^λ, the elasticity of adoption with respect to skilled labor input, by targeting a ratio of R&D expenditure to GDP consistent with the data (around 1.8%). In our baseline calibration, this results in a value of ρ^λ of 0.925. To check that this calibration does not lead to a rate of technological diffusion that is at odds with the data, we compare the sensitivity of speed of diffusion in the model to the regression analysis presented in Table 2. There are three conceptual obstacles to overcome in carrying out this comparison. The first is that the data in the regressions of Table 2 concerns the diffusion of specific technologies in the cross section of potential adopters over time. In our model instead each new technology is adopted either fully or not at all. The second is that the diffusion process in the data is approximately logistic, whereas the diffusion process in our model is geometric. Finally, in the model, unlike in the data available for analysis, technologies become obsolete over time. To address the first challenge, we define speed of diffusion in our model as relating to the speed at which technologies invented at different times are adopted. Formally, denote by Z_{t+k}^t the mass of technologies invented at time t that survives (i.e. is not obsolete) at time $t + k$, and A_{t+k}^t the mass of vintage t technologies that have been adopted at time $t + k$.

Then, we can define the fraction of vintage t technologies adopted at time $t + k$ as

$$m_{t+k}^t \equiv \frac{A_{t+k}^t}{Z_{t+k}^t}. \tag{47}$$

Analogously to Equation 1, we define

$$r_{t+k}^t \equiv \frac{m_{t+k}^t}{1 - m_{t+k}^t} \tag{48}$$

The stock of vintage t total and adopted technologies evolve as follows:

$$Z_{t+k}^t = \phi Z_{t+k-1}^t \tag{49}$$

$$A_{t+k}^t = \phi A_{t+k-1}^t + \lambda_{t+k-1} \phi \left(Z_{t+k-1}^t - A_{t+k-1}^t \right) \tag{50}$$

With initial conditions $Z_t^t = Z_t - \phi Z_{t-1}$ and $A_t^t = 0$. These laws of motion and initial conditions imply that m_{t+k}^t and r_{t+k}^t follow:

$$m_{t+k}^t = m_{t+k-1}^t + \lambda_{t+k-1} \left(1 - m_{t+k-1}^t \right)$$

and

$$r_{t+k}^t = \frac{r_{t+k-1}^t + \lambda_{t+k-1}}{1 - \lambda_{t+k-1}} \tag{51}$$

With initial conditions $m_t^t = 0$ and $r_t^t = 0$. In each period, a fraction $1 - \phi$ of technologies becomes obsolete, so the total stock of vintage t technologies decreases over time. The stock of adopted vintage t technologies increases as a fraction λ_{t+k-1} of the remaining unadopted technologies (note that all unadopted technologies, irrespective of vintage, have the same probability λ_{t+k-1} of being adopted). By analogy to Equation (2) we define the speed of diffusion at time $t + k$ for a vintage t technology as

$$Speed_{t+k}^t \equiv \Delta \log \left(r_{t+k}^t \right) = \log \left(\frac{r_{t+k}^t}{r_{t+k-1}^t} \right) \tag{52}$$

The regression analysis we conduct measures the sensitivity of the speed of technological diffusion to fluctuations in the output gap. If the fraction of adopters is a logistic function of time (see Mansfield (1961)), diffusion speed is constant absent any cyclical fluctuations. In contrast, the diffusion process in the model is geometric, which implies that speed is a declining function of the age of a technology. To remove this non-cyclical variation,

we construct a detrended model measure of speed, which we denote $\widehat{Speed}_{t+k}^{t}$, defined as follows:

$$\widehat{Speed}_{t+k}^{t} \equiv Speed_{t+k}^{t} - \overline{Speed}_{t+k}^{t} \tag{53}$$

Where

$$\overline{Speed}_{t+k}^{t} = log\left(\frac{\bar{r}_{t+k}^{t}}{r_{t+k-1}^{t}}\right) \tag{54}$$

and

$$\bar{r}_{t+k}^{t} = \frac{r_{t+k-1}^{t} + \bar{\lambda}}{1 - \bar{\lambda}} \tag{55}$$

Intuitively, $\bar{r}_{t+k}^{t}$ is the value that r_{t+k}^{t} would take if the diffusion process returned to steady state. Our detrended measure therefore captures cyclical variation in the diffusion process. Figure 15 illustrates the effect of detrending on our measure of speed.

The data used in our regression analysis is a panel of technologies. In the model however the relative masses of technologies of different vintages is not constant over time, due to obsolescence, adoption and trend growth in the stock of technologies. In calculating a population average of speed of diffusion, we take account of the effect of these three factors to make the data and model regressions comparable. To do so, we run the following regression using model-simulated data

$$\widehat{Speed}_{t+k}^{t} = \alpha + \beta^{k}\hat{y}_t + \epsilon_{k,t} \tag{56}$$

where k denotes the age of a technology, for a range of values of k. To construct a population average, we weight each β^{k} by the relative steady state fraction of technologies of age k in the population, w^k, defined as:

$$w^k = \underbrace{\left(1 - m^k\right)}_{\text{adoption}} \underbrace{\phi^{k-1}}_{\text{obsolescence}} \underbrace{(1 + g^a)^{-(k-1)}}_{\text{growth in tech}}$$

The population average elasticity of speed with respect to the output gap is then

$$\beta = \sum_{j=1}^{K} \frac{w^j}{\bar{w}} \beta^j,$$

where $\bar{w} = \frac{1}{K}\sum_{j=1}^{K} w^j$. For the estimation results reported in Section 5 we set K to a large number (50) so that the effect of adding additional vintages to the average is negligible. To obtain the point estimates and confidence intervals reported in 6 we estimate regression (56)

in 100 period subsamples of our model-simulated data and average across these subsample results.

Figure 15: *Speed* and $\widehat{Speed}$: Simulated and Steady State
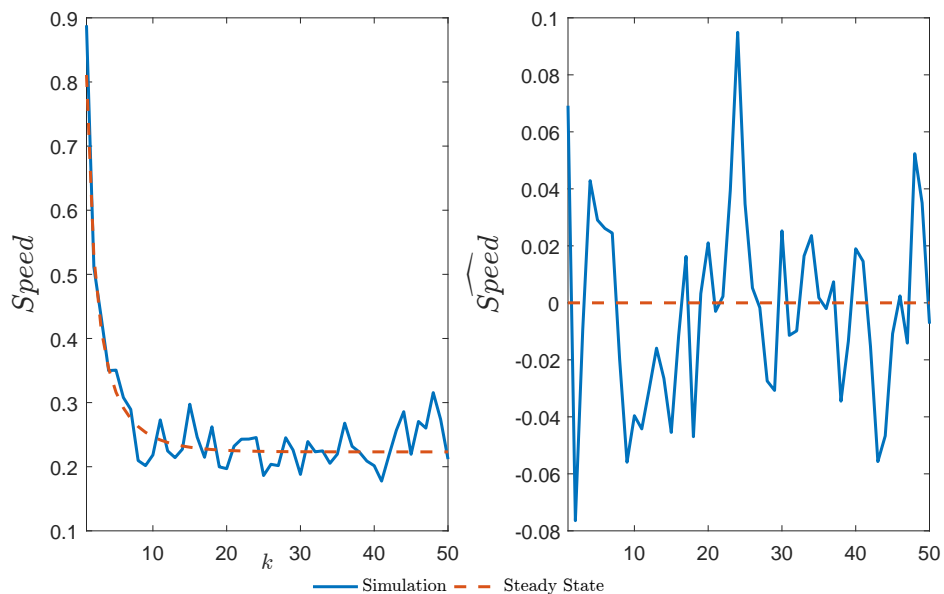
This figure plots Speed (left panel) and detrended $\widehat{Speed}$ (right panel) as a function of the age of a technology k . In both cases the dotted line is the steady state value of speed and the solid line is a typical model simulation.

A.4 Other Tables and Figures

As a check of the fit of the estimated model, Table 8 presents the theoretical standard deviations of the observable variables generated by the model and compares them with the data in our sample. Roughly speaking the model is in line with the actual volatilities of the key variables.

Table 8: Comparison of Standard Deviations

Variable	Data	Model
Output Growth	0.55	0.63
Consumption Growth	0.51	0.71
Investment Growth	1.54	1.52
Inflation	0.23	0.36
Nominal R	0.60	0.55
Hours (level)	1.82	1.53
R&D Expenditure Growth	4.00	6.83

Table 9: Prior and Posterior Distributions of Shock Processes

Parameter	Description	Prior			Posterior	
		Distr	Mean	St. Dev.	Mean	St. Dev.
σ_ζ	Liq. Demand	Inv. Gamma	0.10	2.00	0.225	0.0013
σ_χ	R&D	Inv. Gamma	0.10	2.00	2.202	0.1749
σ_g	Govt. exp.	Inv. Gamma	0.10	2.00	2.559	0.0350
σ_{mp}	Monetary	Inv. Gamma	0.10	2.00	0.097	0.0001
σ_μ	Markup	Inv. Gamma	0.10	2.00	0.093	0.0002
σ_{pk}	Investment	Inv. Gamma	0.10	2.00	1.270	0.0092
σ_θ	TFP	Inv. Gamma	0.10	2.00	0.489	0.0013
σ_{μ_w}	Wage markup	Inv. Gamma	0.10	2.00	0.284	0.0014
ρ_ζ	Liq. Demand	Beta	0.50	0.20	0.924	0.0006
ρ_χ	R&D	Beta	0.50	0.20	0.803	0.0073
ρ_g	Govt. exp.	Beta	0.50	0.20	0.968	0.0001
ρ_{mp}	Monetary	Beta	0.50	0.20	0.465	0.0068
ρ_μ	Markup	Beta	0.50	0.20	0.401	0.0186
ρ_{pk}	Investment	Beta	0.50	0.20	0.899	0.0014
ρ_θ	TFP	Beta	0.50	0.20	0.953	0.0007
ρ_{μ_w}	Wage markup	Beta	0.50	0.20	0.288	0.0101

References

AGHION, P., G.-M. ANGELETOS, A. BANERJEE, AND K. MANOVA (2010): "Volatility and growth: Credit constraints and the composition of investment," *Journal of Monetary Economics*, 57, 246 – 265.

AGHION, P. AND P. HOWITT (1998): *Endogenous Growth Theory*, Cambridge, MA: MIT Press.

AN, S. AND F. SCHORFHEIDE (2007): "Bayesian Analysis of DSGE Models," *Econometric Reviews*, 26, 113–172.

ANDREWS, D., C. CRISCUOLO, AND P. GAL (2015): "Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence from OECD Countries," .

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

BARTEL, A., C. ICHNIOWSKI, K. L. SHAW, AND R. CORREA (2009): *International Differences in the Adoption and Impact of New Information Technologies and New HR Practices: The Valve-Making Industry in the United States and United Kingdom*, University of Chicago Press, 55–78.

BENIGNO, G. AND L. FORNARO (2017): "Stagnation Traps," Mimeo, London School of Economics.

BERGER, D., I. DEW-BECKER, L. D. SCHMIDT, AND Y. TAKAHASHI (2015): "Layoff risk, the welfare cost of business cycles, and monetary policy," .

BIANCHI, F., H. KUNG, AND G. MORALES (2017): "Growth, Slowdowns, and Recoveries," *Journal of Monetary Economics*, Forthcoming.

BOSWORTH, D. L. (1978): "The Rate of Obsolescence of Technical Knowledge–A Note," *The Journal of Industrial Economics*, 26, 273–279.

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

CABALLERO, R. J. AND A. B. JAFFE (1993): "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*, National Bureau of Economic Research, Inc, NBER Chapters, 15–86.

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

CHRISTIANO, L. J., M. S. EICHENBAUM, AND M. TRABANDT (2015): "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, 7, 110–67.

COMIN, D. (2009): "On the integration of growth and business cycles," *Empirica*, 36, 165–176.

COMIN, D. AND M. GERTLER (2006): "Medium-Term Business Cycles," *American Economic Review*, 96, 523–551.

COMIN, D. AND B. HOBIJN (2010): "An Exploration of Technology Diffusion," *American Economic Review*, 100, 2031–59.

COMIN, D. AND M. MESTIERI (2015): "If Technology Has Arrived Everywhere, Why has Income Diverged?" Working paper.

COX, M. W. AND R. ALM (1996): "The economy at light speed: technology and growth in the information age and beyond," *Annual Report*, 2–17.

DAVIES, S. (1979): *The Diffusion of Process Innovations*, Cambridge University Press.

DECKER, R. A., J. HALTIWANGER, R. S. JARMIN, AND J. MIRANDA (2017): "Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown," *American Economic Review*, 107, 322–26.

DEL NEGRO, M., D. GIANNONE, M. P. GIANNONI, AND A. TAMBALOTTI (2017): "Safety, Liquidity, and the Natural Rate of Interest," Working Paper 812, Federal Reserve Bank of New York.

DEL NEGRO, M., M. P. GIANNONI, AND F. SCHORFHEIDE (2015): "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics*, 7, 168–96.

FERNALD, J. (2012): "A quarterly, utilization-adjusted series on total factor productivity," Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

FERNALD, J. G. (2014): *Productivity and Potential Output before, during, and after the Great Recession*, 1–51, in Parker and Woodford (2015).

FERNALD, J. G., R. E. HALL, J. H. STOCK, AND M. W. WATSON (2017): "The Disappointing Recovery of Output after 2009," *Brookings Papers on Economic Activity.*

FISHER, J. D. (2015): "On the Structural Interpretation of the Smets–Wouters "Risk Premium" Shock," *Journal of Money, Credit and Banking*, 47, 511–516.

GARCIA-MACIA, D. (2015): "The Financing of Ideas and the Great Deviation," .

GORDON, R. J. (2012): "Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds," NBER Working Papers 18315, National Bureau of Economic Research, Inc.

GREENWOOD, J., Z. HERCOWITZ, AND G. W. HUFFMAN (1988): "Investment, Capacity Utilization, and the Real Business Cycle," *The American Economic Review*, 78, 402–417.

GRILICHES, Z. (1957): "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica*, 25, pp. 501–522.

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

GUERRIERI, L. AND M. IACOVIELLO (2015): "OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily," *Journal of Monetary Economics*, 70, 22 – 38.

GUERRON-QUINTANA, P. AND R. JINNAI (2014): "Liquidity, trends, and the great recession," .

HALL, R. E. (2014): "Quantifying the Lasting Harm to the U.S. Economy from the Financial Crisis," in Parker and Woodford (2015), 71–128.

HUBER, K. (2018): "Disentangling the Effects of a Banking Crisis: Evidence from German Firms and Counties," *American Economic Review*, 108, 868–98.

JAIMOVICH, N. (2007): "Firm dynamics and markup variations: Implications for sunspot equilibria and endogenous economic fluctuations," *Journal of Economic Theory*, 137, 300 – 325.

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

Krishnamurthy, A. and A. Vissing-Jorgensen (2012): "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, 120, 233 – 267.

Mansfield, E. (1961): "Technical Change and the Rate of Imitation," *Econometrica*, 29, pp. 741–766.

Parker, J. A. and M. Woodford (2015): *NBER Macroeconomics Annual 2014, Volume 29*, University of Chicago Press.

Queralto, A. (2015): "A Model of Slow Recoveries from Financial Crises," .

Reifschneider, D., W. Wascher, and D. Wilcox (2015): "Aggregate Supply in the United States," *IMF Econ Review*, 63, 71–109.

Reinhart, C. M. and K. Rogoff (2009): *This time is different: Eight centuries of financial folly*, Princeton University Press.

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

Schmookler, J. (1966): *Invention and Economic Growth*, Harvard University Press.

Shleifer, A. (1986): "Implementation Cycles," *Journal of Political Economy*, 94, 1163–90.

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

Trajtenberg, M. (1990): *Economic Analysis of Product Innovation: The Case of CT Scanners*, Cambridge University Press.