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

Diego Anzoategui
NYU

Diego Comin
Dartmouth, NBER and CEPR

Mark Gertler
NYU and NBER

Joseba Martinez
NYU

February 2017

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.

*Thanks to Francesco Bianchi, Bob Hall, Howard Kung, and Chris Tonetti for helpful comments. Financial assistance from the NSF is greatly appreciated.

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.

What accounts for reduced productivity growth following financial crises? There are two candidate hypotheses: bad luck versus an endogenous response. Fernald (2014) makes a compelling case for the bad luck hypothesis. As he emphasizes, 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 detrended total factor productivity, specifically Fernald's utilization corrected measure, along with labor productivity. Both measures show a sustained decline relative to trend in the years after the Great

[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

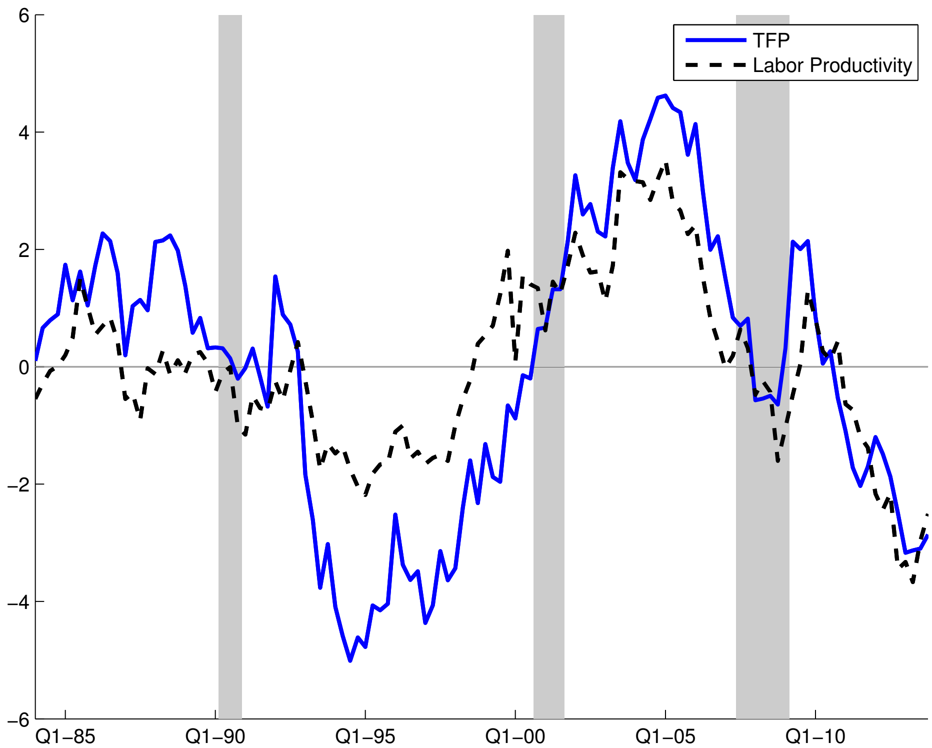

All series are log-linearly detrended. Labor productivity is GDP divided by hours worked (see Appendix A.1 for data sources). TFP is Utilization-Adjusted Total Factor Productivity (available at http://www.frbsf.org/economic-research/total-factor-productivity-tfp/; see Fernald (2012) for details).

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

[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.

technologies through research and development (R&D) and (ii) the diffusion of new technologies 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. We will argue that this latter decline contributed to the productivity slowdown prior to the Great Recession. Unfortunately there is no aggregate 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 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 framework, 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 generate by our model is consistent with the panel data evidence presented in section 2. We also allow for an exogenous TFP shock that can capture the Fernald (2014) bad luck hypotheses.

In addition to the literature cited above, there are several other papers related to our analysis. Queralto (2015), Guerron-Quintana and Jinnai (2014) and Garcia-Macia (2013) have appealed to endogenous growth considerations to explain the persistence of financial crises. The paper most closely related to ours is Bianchi and Kung (2014), who first estimated a macroeconomic model with endogenous growth using R&D data. There are key differences between our analysis and that in Bianchi and Kung (2014). First, our model of R&D and adoption is more explicit in a way that permits discipline on the lag in the diffusion process. Second, and perhaps most significant, we use the panel data evidence as an external validity check to ensure that the cyclical behavior of adoption the model generates is plausible. Finally, 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. As we mentioned, 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 present the estimates. 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. Here we find that a sharp in endogenous decline in 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, on the other, do not explain much of the productivity variation over this period.

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 to motivate our formulation of endogenous productivity.[3]

We begin with Figure 2 which 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 we noted earlier, an aggregate time series measuring adoption is not available. To explore the the cyclicality of technology diffusion we accordingly resort to survey data on the speed of technologically diffusion, of the type used in the productivity literature[5] The specific data we have available is a sample of 26 production technologies that diffused at various times over the period 1947-2003 in the US (5) and the UK (21).[6] We use the data to estimate the effect of the business cycle on the

[3]For evidence on the role of technology diffusion on productivity growth see for example, Skinner and Staiger (2015), Comin (2000), Comin and Hobijn (2010) and Comin and Mestieri (2013).

[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.

[5]See for example, Griliches (1957) and Mansfield (1961).

[6]The data on UK technologies comes from Davies (1979) and covers special presses, foils, wet suction boxes, gibberellic acid, automatic size boxes, accelerated drying hoods, basic oxygen process, vacuum degassing, vacuum melting, continuous casting, tunnel kilns, process control by computer, tufted carpets,

Figure 2: R&D Expenditures by US Corporations, 1983-2013

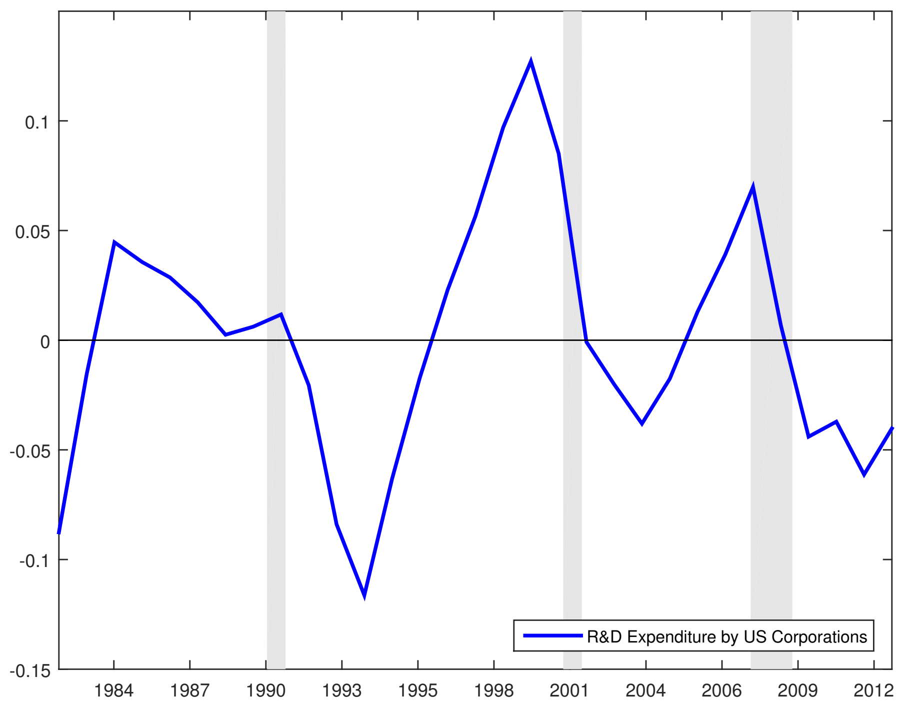

Log-linearly detrended data. Source: R&D Expenditure by US corporations (National Science Foundation). Data are deflated by the GDP deflator and divided by the civilian population older than 16 (see Appendix A.1 for data sources).

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 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 cyclical indicator, specifically detrended GDP per person over 16 years old.

Table 1 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.

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

computer typesetting, photo-electrically controlled cutting, shuttleless looms, numerical control printing presses, numerical control turning machines and numerical control turbines. The data for the five technologies in the US comes from Trajtenberg (1990), and Bartel et al. (2009) and covers the diffusion of CT scanners, Computerized numerical controlled machines, automated inspection sensors, 3-D CAD, and flexible manufacturing systems.

Table 1: 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 term 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.

Figure 3: Speed of Diffusion

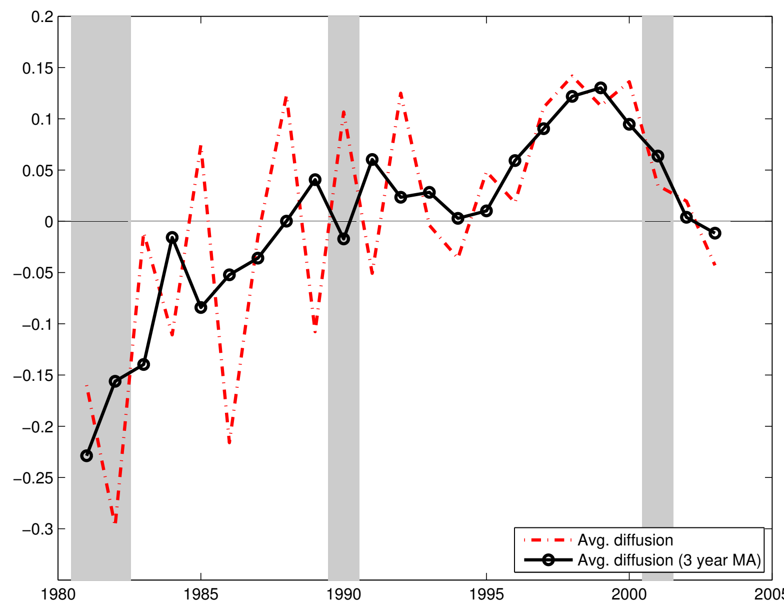

plots the speed of diffusion for the balanced panel of four US technologies for which we have data from 1981 to 2003. Specifically, for each of the technologies we remove the acyclical component of the diffusion rate $(\alpha_i + G(lag_{it}))$. We then average the residual $(\beta * \widehat{y}_t + \epsilon_{it})$ 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 U.S. data on technology diffusion during the Great Recession. We do, however, have some limited evidence for the U.K. In particular, Eurostat provides information on the diffusion of three relevant internet-related technologies in the UK.[7] Figure 4 plots their average diffusion from 2004 until 2013 with the business cycle

[7] 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 1. The

Figure 4: Diffusion Speed for 3 Internet Technologies in the UK, 2004-2013

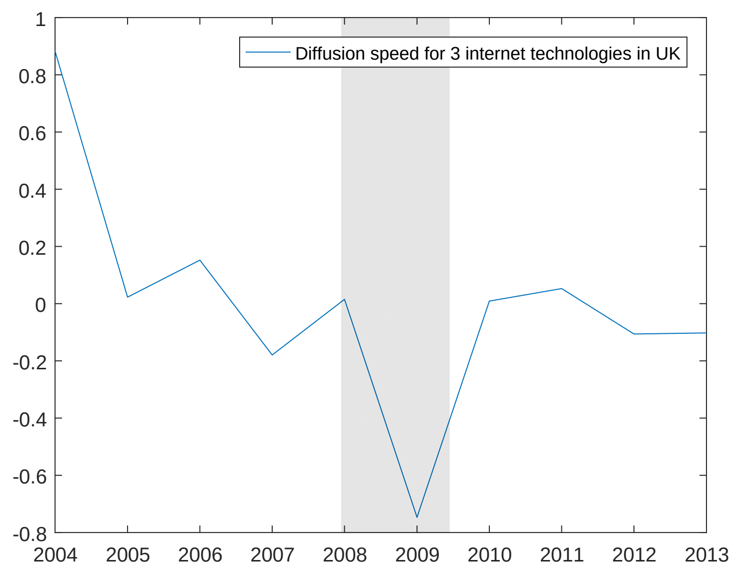

Source: Eurostat; see footnote 7 for details of calculations. Shaded areas are UK recession dates as dated by UK ONS.

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 8% 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.

An alternative way to study the evolution of the speed of diffusion of technologies is to measure the investments of companies in adopting new technologies. Assembling such a measure is a daunting task because the nature of adoption investments is very diverse including the costs of licensing technologies, training workers, redesigning production processes, purchasing the capital goods that embody them, etc. 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 expenses by companies to acquire

resulting series are demeaned so that they can be interpreted as percent deviations from the average speed of technology diffusion.

Figure 5: Revenues from Licensing Technologies, 1993-2013

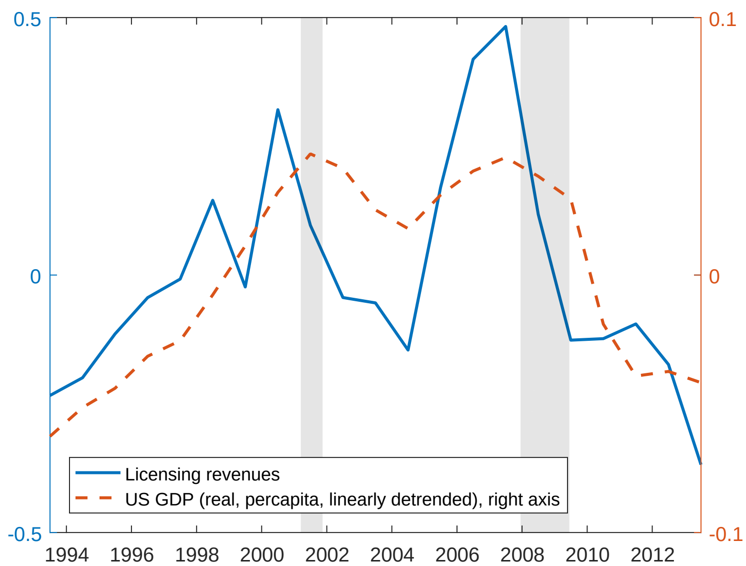

licenses of technologies developed by universities and research hospitals.[8]

Figure 5 shows the evolution of linearly detrended licensing revenues from 1995 to 2014 together with linearly detrended GDP.[9] The two series are highly correlated with a coefficient of 0.69. Both before 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 therefore consistent with the hypothesis that investments in technology adoption strongly co-move with the business cycle.[10]

[8] 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 institutions which according to the NSF amounted to $62 billion in 2011.

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

[10] 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.

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, including: habit formation in consumption, flow investment adjustment costs, variable capital utilization and "Calvo" price and wage rigidities. 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 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. In a subsequent section we present the firm optimization problems.

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 goods 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 $\mu_t > 1$ is given exogenously.

Each final good firm i uses Y_{mt}^i units of intermediate goods composite as input to produce output, according to the following simple linear technology

$$Y_t^i = Y_{mt}^i \tag{5}$$

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:

$$Y_{mt} = \left(\int_0^{A_t} (Y_{mt}^j)^{\frac{1}{\vartheta}} dj \right)^{\vartheta} \tag{6}$$

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. 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:

$$Y_{mt}^j = \theta_t \left(U_t^j K_t^j \right)^{\alpha} (L_t^j)^{1-\alpha} \tag{7}$$

where θ_t is an exogenous random disturbance. As we will make clear shortly, θ_t is the exogenous component of total factor productivity. 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.

Let $\overline{Y}_t$ be average output across final goods producers. Then the production function (4) implies the following expression for the final good composite Y_t

$$Y_t = \Omega_t \cdot \overline{Y}_t \tag{8}$$

where Ω_t is the following measure of output dispersion

$$\begin{aligned} \Omega_t &= \left(\int_0^1 (Y_t^i / \overline{Y}_t)^{\frac{1}{\mu_t}} di \right)^{\mu_t} \\ &= 1 \text{ to a 1st order} \end{aligned} \tag{9}$$

In a first order approximation, Ω_t equals unity, implying that we can express Y_t simply as $\overline{Y}_t$.

Next, given the total number of final goods firms is unity, given the production function for each final goods producer (5), and given that Y_t equals $\overline{Y}_t$, it follows that to a first order

$$Y_t = Y_{mt} \tag{10}$$

Finally, given a symmetric equilibrium for intermediate goods (recall prices are flexible in this sector) it follows from equation (6) that we can express the aggregate production function for the finally good composite Y_t as

$$Y_t = \left[A_t^{\vartheta - 1} \theta_t \right] \cdot (U_t K_t)^{\alpha} (L_t)^{1-\alpha} \tag{11}$$

where 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. Note that equation (11) holds to a first order since we impose Ω_t equals unity.

In sum, endogenous productivity effects enter through the expansion in the variety of adopted intermediate goods, measured by A_t. 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.

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 at time $t+1$ that each unit of skilled labor at t can create. We assume φ_t is given by

$$\varphi_t = \chi_t Z_t L_{srt}^{\rho_z - 1} \tag{12}$$

where χ_t is an exogenous disturbance to the R&D technology 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. We introduce this congestion externality so that we can 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.[11]

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^p_{srt} to solve

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

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

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

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}.[12]

Finally, we allow for obsolescence of technologies.[13] 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)

[11] 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.

[12] 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).

[13] We introduce obsolescence to permit the steady state share of spending on R&D to match the data.

and (12) 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, but the adoption rate can vary pro-cyclically, consistent with evidence in Comin (2009). 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.

Accordingly, we proceed as follows. We suppose there are 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 adopted 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 adoption expenditures in the 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 $\lambda' > 0$, $\lambda'' < 0$.[14] 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

[14] In the estimation, we assume that

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

where κ_λ is a constant.

for R&D, there is no separate shock to the productivity of adoption activities in (17). We are constrained to introduce this asymmetry because we do not have a direct observable to measure adoption labor or λ_t. The identified series of adoption hours, L_{sat}, can be interpreted as the effective number of adoption hours.

Our adoption process implies that technology diffusion takes time on average, consistent with the evidence. 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 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 (11). Let Π_{mt} be the profits that the intermediate goods firm makes from producing the good, which arise from monopolistically competitive pricing. The adopter sells the new technology at the competitive price V_t, which 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

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

subject to 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 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 an adopted versus 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.

Given that 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 exposition of the model, we make a detour to map the notion of diffusion in the model and in our econometric analysis of 26 specific technologies. At a high level, instead of measuring diffusion by the number of companies that use a technology, we will measure it by the share of all technologies invented at a given period that have been adopted at some future time. Formally, let Z_{t+k}^t is 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 is 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 Equation 1, we define

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

The speed of technology diffusion is

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

where the second equality comes from the law of motion of r_{t+k}^t derived in the Appendix.

Equation (24) has two relevant implications. First, the speed of technology, $Speed_{t+k}^t$, is pro-cyclical because it varies positively with λ_{t+k-1} which is pro-cyclical. Second, in addition to λ_{t+k-1} which is common across vintages, $Speed_{t+k}^t$ declines with the diffusion level of the technology as measured by r_{t+k-1}^t. This "vintage" effect, which is due to the geometric nature of diffusion in the model, motivates the introduction of a vintage-control when estimating the cyclicality of diffusion in the model similar to that included in the data analysis of section 2.

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 $\varrho_t > 0$. As we show, the liquidity demand shock transmits through the economy like a financial shock. It is mainly for this reason that we make use of it, as opposed to a shock to the discount factor.[15]

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

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

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

$$R_{kt} = \frac{D_t + Q_t}{Q_{t-1}} \tag{27}$$

$\Lambda_{t,t+1}$, the household's stochastic discount factor, is given by

$$\Lambda_{t,t+1} \equiv \beta u'(C_{t+1})/u'(C_t) \tag{28}$$

where $u'(C_t) = 1/(C_t - bC_{t-1}) - b/(C_{t+1} - bC_t)$. In addition, let ζ_t be the liquidity preference

[15] Another consideration is that the liquidity demand shock induces positive co-movement between consumption and investment, while that is not always the case for a discount factor shock.

shock in units of the consumption good:

$$\zeta_t = \varrho_t / u'(C_t) \tag{29}$$

Then we can express the first order necessary conditions for capital and the riskless bond as, respectively:

$$1 = E_t\{\Lambda_{t,t+1} R_{kt+1}\} \tag{30}$$

$$1 = E_t\{\Lambda_{t,t+1} R_{t+1}\} + \zeta_t \tag{31}$$

As equation (31) 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 (30) implies. The decline in the discount factor also induces a drop in R&D and investment.

Overall, the shock to ζ_t generates positive co-movement between consumption and investment similar to that arising from a monetary shock. To see, combine equations (30) and (31) to obtain

$$E_t\{\Lambda_{t,t+1}(R_{kt+1} - R_{t+1})\} = \zeta_t \tag{32}$$

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.

Since it is fairly conventional, we defer until later a description of the household's wage-setting and labor supply behavior.

3.4 Firms

3.4.1 Intermediate goods firms: factor demands

Let's denote by ς the markup charged by intermediate goods firms. In principle 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). Let p_{mt} be the relative price of the intermediate goods composite. Then from (6) and the production function (7), cost minimization by each intermediate goods producer yields the following standard first order conditions for capital, capital utilization, and unskilled labor:

$$\alpha \frac{p_{mt} Y_{mt}}{K_t} = \varsigma [D_t + \delta(U_t) Q_t] \tag{33}$$

$$\alpha \frac{p_{mt} Y_{mt}}{U_t} = \varsigma \delta'(U) Q_t K_t \tag{34}$$

$$(1-\alpha) \frac{p_{mt} Y_{mt}}{L_t} = \varsigma w_t \tag{35}$$

3.4.2 Final goods producers: price setting

Let P_t^i be the nominal price of final good i and P_t the nominal price level. Given the CES relation for the final good composite, equation (4), the demand curve facing each final good producer is:

$$Y_t^i = \left(\frac{P_t^i}{P_t} \right)^{-\mu_t/(\mu_t - 1)} Y_t \tag{36}$$

where the price index is given by:

$$P_t = \left(\int_0^1 (P_t^i)^{-1/(\mu_t - 1)} di \right)^{-(\mu_t - 1)}, \tag{37}$$

Following Smets and Wouters (2007), we assume Calvo pricing with flexible indexing. Let $1-\xi_p$ be the i.i.d probability that a firm is able to re-optimize its price and let $\pi_t = P_t/P_{t-1}$ be the inflation rate. Firms that are unable to re-optimize during the period adjust their price according to the following indexing rule:

$$P_t^i = P_{t-1}^i \pi_{t-1}^{\iota_p} \pi^{1-\iota_p} \tag{38}$$

where π is the steady state inflation rate and ι_p reflects the degree of indexing to lagged inflation.

For firms able to re-optimize, the optimization problem is to choose a new reset price P_t^* to maximize expected discounted profits until the next re-optimization, given by

$$E_t \sum_{\tau=0}^{\infty} \xi_p^\tau \Lambda_{t,t+\tau} \left(\frac{P_t^* \Gamma_{t,t+\tau}}{P_{t+\tau}} - p_{mt+\tau} \right) Y_{t+\tau}^i \tag{39}$$

subject to the demand function (36) and where

$$\Gamma_{t,t+\tau} \equiv \prod_{k=1}^{\tau} \pi_{t+k-1}^{\iota_p} \pi^{1-\iota_p} \tag{40}$$

The first order condition for P_t^* and the price index that relates P_t to P_t^*, P_{t-1} and π_{t-1} are then respectively:

$$0 = E_t \sum_{\tau=0}^{\infty} \xi_p^{\tau} \Lambda_{t,t+\tau} \left[\frac{P_t^* \Gamma_{t,t+\tau}}{P_{t+\tau}} - \mu_{t+\tau} p_{mt+\tau} \right] Y_{t+\tau}^i \tag{41}$$

$$P_t = \left[(1-\xi_p) \left(P_t^* \right)^{-1/(\mu_t - 1)} + \xi_p \left(\pi_{t-1}^{\iota_p} \pi^{1-\iota_p} P_{t-1} \right)^{-1/(\mu_t-1)} \right]^{-(\mu_t - 1)} \tag{42}$$

Equations (41) and (42) jointly determine inflation. In the loglinear equilibrium, current inflation is a function of current real marginal cost p_{mt}, expected future inflation, and lagged inflation.

3.4.3 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 investment expenditures into new capital (the replacement price of capital), and γ_y the steady state growth in I_t. In addition, following Christiano et al. (2005), we assume flow adjustment costs of investment. The capital producers' decision problem is to choose I_t to solve

$$\max_{I_t} E_t \sum_{\tau=0}^{\infty} \Lambda_{t,t+\tau} \left\{ Q_{t+\tau} I_{t+\tau} - p_{kt+\tau} \left[1 + f\left(\frac{I_{t+\tau}}{(1+\gamma_y) I_{t+\tau-1}} \right) \right] I_{t+\tau} \right\} \tag{43}$$

where the adjustment cost function is increasing and concave, with $f(1) = f'(1) = 0$ and $f''(1) > 0$. We assume that p_{kt} follows an exogenous stochastic process.

The first order condition for I_t the relates the ratio of the market value of capital to the replacement price (i.e. "Tobin's Q") to investment, as follows:

$$\frac{Q_t}{p_{kt}} = 1 + f\left(\frac{I_t}{(1+\gamma_y) I_{t-1}} \right) + \frac{I_t}{(1+\gamma_y) I_{t-1}} f'\left(\frac{I_t}{(1+\gamma_y) I_{t-1}} \right) - E_t \Lambda_{t,t+1} \left(\frac{I_{t+1}}{(1+\gamma_y) I_t} \right)^2 f'\left(\frac{I_{t+1}}{(1+\gamma_y) I_t} \right) \tag{44}$$

3.4.4 Employment agencies and wage adjustment

As we noted earlier, the household is a monopolistically competitive supplier of labor. Think of the household as supplying its labor to form a labor composite. Firms then hire the labor composite. The only difference from the standard DSGE model with wage rigidity, is that households now supply two types of labor, skilled and unskilled. It also sets wages for each type.

Let $X_t = \{L_t, L_{st}\}$ denote a labor composite. As is standard, we assume that X_t is the following CES aggregate of the differentiated types of labor that households provide:

$$X_t = \left[\int_0^1 X_t^{h\,\frac{1}{\mu_{wt}}} dh\right]^{\mu_{wt}}. \tag{45}$$

where $\mu_{wt} > 1$ obeys an exogenous stochastic process[16].

Let W_{xt} denote the wage of the labor composite and let W_{xt}^h be the nominal wage for labor supplied of type x by household h. Then profit maximization by competitive employment agencies yields the following demand for type x labor:

$$X_t^h = \left(\frac{W_{xt}^h}{W_{xt}}\right)^{-\mu_{wt}/(\mu_{wt}-1)} X_t, \tag{46}$$

with

$$W_{xt} = \left[\int_0^1 W_{xt}^{h\,-\frac{1}{\mu_{wt}-1}} dh\right]^{-(\mu_{wt}-1)}. \tag{47}$$

As with price setting by final goods firms, we assume that households engage in Calvo wage setting with indexation. Each period a fraction $1-\xi_w$ of households re-optimize their wage for each type. Households who are not able to re-optimize adjust the wage for each labor type according to the following indexing rule:

$$W_{xt}^h = W_{xt-1}^h \pi_{t-1}^{\iota_w} \pi^{1-\iota_w} \gamma. \tag{48}$$

where γ is the steady state growth rate of labor productivity.

The remaining fraction of households choose an optimal reset wage W_{xt}^* by maximizing

$$E_t \left\{ \sum_{\tau=0}^{\infty} \xi_w^\tau \beta^\tau \left[-\upsilon \frac{X_{t+\tau}^{h\;\,1+\varphi}}{1+\varphi} + u'(C_{t+\tau}) \frac{W_{xt}^* \Gamma_{wt,t+\tau}}{P_{t+\tau}} X_{t+\tau}^h \right] \right\} \tag{49}$$

[16] 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.

subject to the demand for type h labor and where the indexing factor $\Gamma_{xt,t+\tau}$ is given by

$$\Gamma_{wt,t+\tau} \equiv \prod_{k=1}^{\tau} \pi_{t+k-1}^{\iota_w} \pi^{1-\iota_w} \gamma \tag{50}$$

The first order condition for the re-set wage and the equation for the composite wage index as a function of the reset wage, inflation and the lagged wage are given, respectively, by

$$E_t \left\{ \sum_{\tau=0}^{\infty} \xi_w^{\tau} \Lambda_{t,\tau} \left[\frac{W_{xt}^{*} \Gamma_{wt,t+\tau}}{P_t} - \mu_{wt} \upsilon \frac{X_{t+\tau}^{h}{}^{\varphi}}{u'(C_{t+\tau})} \right] X_{t+\tau}^{h} \right\} = 0 \tag{51}$$

$$W_{xt} = \left[(1-\xi_w) \left(W_{xt}^{*} \right)^{-1/(\mu_{wt}-1)} + \xi_p \left(\gamma \pi_{t-1}^{\iota_w} \pi^{1-\iota_w} W_{xt-1} \right)^{-1/(\mu_{wt}-1)} \right]^{-(\mu_{wt}-1)} \tag{52}$$

3.4.5 Fiscal and monetary policy

We assume that government consumption G_t is financed by lump sum taxes T_t.

$$G_t = T_t \tag{53}$$

Further, the (log) deviation of G_t from the deterministic trend of the economy follows an AR(1) process. Formally,

$$\log(G_t/(1+\gamma_y)^t) = (1-\rho_g)\bar{g} + \rho_g \log(G_{t-1}/(1+\gamma_y)^{t-1}) + \epsilon_t^g, \tag{54}$$

Next, we suppose that monetary policy obeys a Taylor rule. Let R_{nt+1} denote the gross nominal interest rate, R_n the steady state nominal rate, π^0 the target rate of inflation, L_t total employment and L^{ss} steady state employment. The (nonlinear) Taylor rule for monetary policy that we consider is given by

$$R_{nt+1} = \left[\left(\frac{\pi_t}{\pi^0} \right)^{\phi_\pi} \left(\frac{L_t}{L^{ss}} \right)^{\phi_y} R_n \right]^{1-\rho} \cdot R_{nt}^{\rho} \tag{55}$$

where the relation between the nominal and real rate is given by the Fisher relation:

$$R_{nt+1} = R_{t+1} \cdot \pi_{t+1} \tag{56}$$

and where ϕ_π and ϕ_y are the feedback coefficients on the inflation gap and capacity utilization gap respectively. 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.[17]

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.

$$R_{nt+1} \geq 1 \tag{57}$$

3.5 Resource constraints and equilibrium

The resource constraint is given by

$$Y_t = C_t + p_{kt}\left[1 + f\left(\frac{I_t}{(1+\gamma_y)I_{t-1}}\right)\right] I_t + G_t \tag{58}$$

Capital evolves according to

$$K_{t+1} = I_t + (1 - \delta(U_t))K_t \tag{59}$$

The market for skilled labor must clear:

$$L_{st} = L_{sat} + L_{srt} \tag{60}$$

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

We estimate our model using Bayesian methods (see for example An and Schorfheide (2007)). As is common practice in the literature (for example Smets and Wouters (2007)

[17] 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.

and Justiniano et al. (2010)), we calibrate a subset of the parameters of the model and estimate the remainder.

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 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 next discuss the calibrated parameters and then proceed to describe the prior and posterior estimates of the remaining parameters.

4.1 Calibrated parameters

We calibrate standard real business cycle model parameters (i.e., the rates of time preference and capital depreciation, and the capital share); the steady state share of government spending in output; the trend growth rate; steady state markups for intermediate and final goods and for wages (ς, μ and μ_w respectively); the elasticity of substitution of intermediate goods, $\vartheta/(\vartheta-1)$, and three of the four endogenous technological change parameters.

Of the four endogenous technological change parameters, we calibrate the expenditure elasticity of the adoption probability, ρ_λ, the obsolescence rate $(1-\phi)$ and the steady state adoption lag $\bar{\lambda}$. The elasticity of λ with respect to adoption expenditures, ρ_λ is set to 0.925 to induce a ratio of private R&D to GDP consistent with the U.S. post-1970 experience (of approximately 1.9% of GDP). $\bar{\lambda}$ 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). Finally, the obsolescence rate $(1-\phi)$ is set to 8% 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 the patent renewal rates (Bosworth, 1978). Finally, ϑ is set to 1.35 to produce an elasticity of substitution of 3.85 in line with the estimates from Broda and Weinstein (2005).

Table 2 presents the calibrated parameters and their values.

Table 2: 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
μ_w	SS wage mark up	1
ϑ	Intermediate goods elasticity of substitution	1.35
$1-\phi$	Obsolescence rate	0.08/4
$\bar{\lambda}$	SS adoption lag	0.2/4
ρ_λ	Adoption elasticity	0.925

4.2 Parameter estimates

Table 3 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 a beta prior centered around a mean of 0.6, which is at the lower end of estimates provided in Griliches (1990).

Most of our estimates are similar to those in the literature. The price and wage rigidity parameters, though, are on the high side, likely reflecting that inflation was low and stable over our sample. Our estimate of the elasticity of new technologies with respect to R&D, ρ_z, is 0.37, 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. This reflects the fact that our model produces significant endogenous persistence in TFP.

Table 3: 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	0.829	0.0007
ϕ_π	Taylor rule inflation	Gamma	1.50	0.25	1.575	0.0518
ϕ_y	Taylor rule labor	Gamma	0.30	0.10	0.386	0.0050
φ	Inverse Frisch elast.	Gamma	2.00	0.75	2.703	0.7062
f''	Investment adj. cost	Gamma	4.00	1.00	5.639	0.6431
$\frac{\delta'(U)}{\delta}$	Capital util. elast.	Gamma	4.00	1.00	4.044	0.9789
ξ_p	Calvo prices	Beta	0.50	0.10	0.935	0.0001
ξ_w	Calvo wages	Beta	0.75	0.10	0.906	0.0013
ι_p	Price indexation	Beta	0.50	0.15	0.241	0.0104
ι_w	Wage indexation	Beta	0.50	0.15	0.382	0.0171
b	Consumption habit	Beta	0.70	0.10	0.483	0.0017
ρ_z	R&D elasticity	Beta	0.60	0.15	0.371	0.0103
α	Capital share	Normal	0.30	0.05	0.201	0.0009
$\tilde{\beta}$	Discount factor	Gamma	0.25	0.10	0.421	0.0093
$100 * \gamma_y$	SS output growth	Normal	0.46	0.03	0.453	0.0005

4.3 Shocks and Volatilities

In this section we establish the key sources of cyclical variation not only in output and hours but also in endogenous productivity.

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 4 presents the results. There are several important takeaways. First, "demand" shocks are important for both output and employment volatility and for the cyclicality of endogenous productivity. The liquidity demand shock explains 42.5 percent of output growth, 55.3 percent of hours and 50.6 percent of endogenous productivity. The liquidity demand and money shocks combined account for more than half the volatility of output and hours and more than two thirds of the variation in hours and endogenous production variation. The other important shock is the exogenous component of total factor productivity, which accounts for 19.1 percent of output variation, 10.9 percent of hours variation, and 9.9 percent of the variation in endogenous TFP.

The liquidity demand shock is by far the most important shock driving recessions. Figure 6 plots the historical evolution of per capita output growth as well as the components that are accounted for by the liquidity demand and the exogenous TFP shock, the disturbance

Table 4: Variance Decomposition (%)

Variables	Liquidity Demand	Money	Govt Exp	Price of Capital	TFP	R&D	Mark up	Wage mark up
Output Growth	42.5	13.3	16.5	6.1	19.1	0	2.1	0.4
Consumption Growth	45.7	14.4	16.6	0.7	20.5	0	1.8	0.3
Investment Growth	16.4	4.7	2.6	65.3	7.7	0.4	1.8	1.2
Inflation	0.1	0	0.1	0	3.3	0	80.8	15.7
Nominal R	35.2	33.2	1.2	2.2	7.8	0.3	13.3	6.7
Hours	55.3	18.3	5.4	4.7	10.9	0.2	3.9	1.2
Endogenous TFP	50.6	18.7	3.6	0.9	9.9	9.4	1.6	5.3

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

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. The TFP shock accounts for a comparable decline only in the 1990-91 recession. In addition to comparing to the though 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.

5 Analysis

We begin by analyzing the mechanisms through which the liquidity demand shock affects endogenous productivity in our model. We subsequently use our estimated model to provide a decomposition of the forces driving TFP before, during and after the Great Recession.

5.1 Endogenous Productivity Mechanism

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.

[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.

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

Figure 6: Output Growth Decomposition

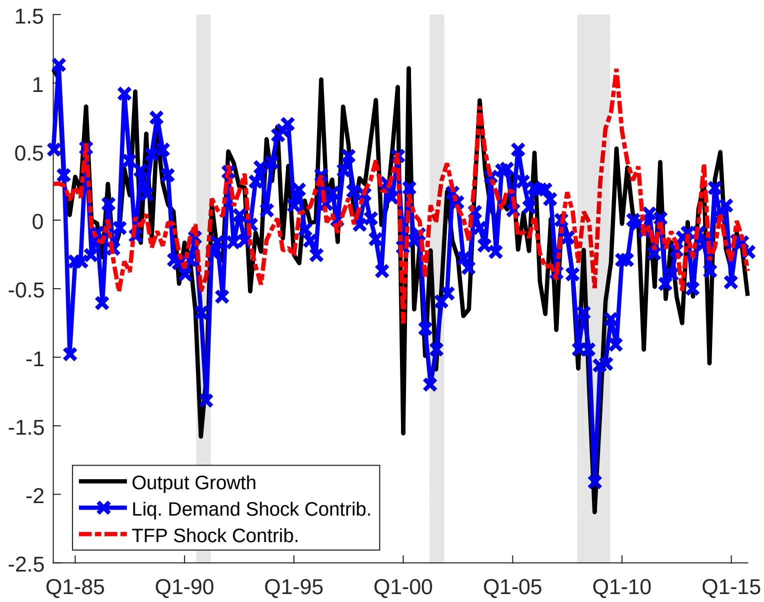

Data sources are described in Appendix A.1. Smoothed shocks from model estimated using data as described in Section 4.2 and Appendix A.1.

Figure 7: Impulse Response to 1 std. dev. Shock

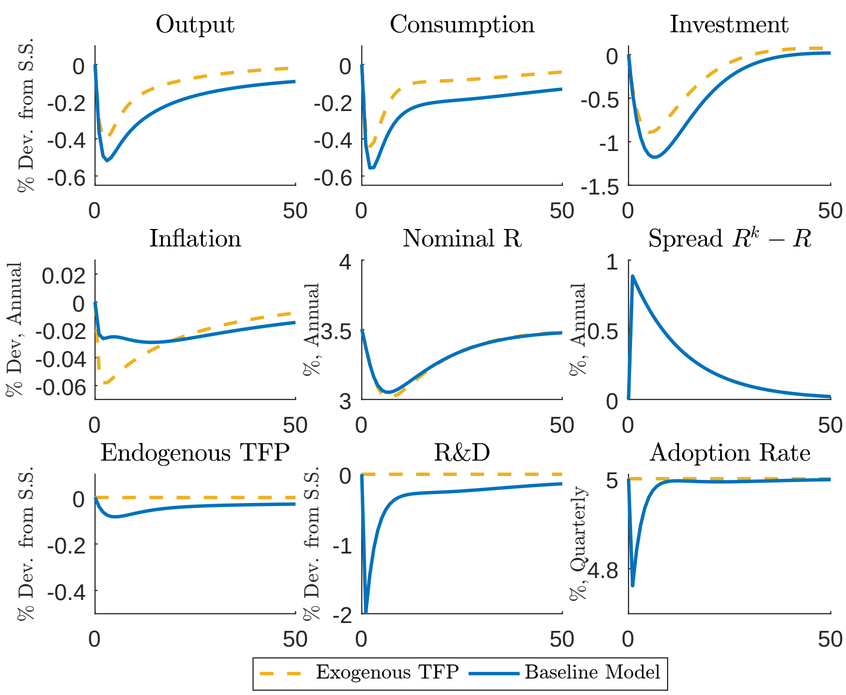

An increase in the demand for the liquid asset, everything else equal induces households to reduce their consumption demand and their saving in risky assets (See equations 30-32). As a result there is upward pressure on the required return to capital, R_{kt+1}, and downward pressure on the safe real 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 lessens 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, accordingly, can offer at least part of the explanation for the surprising failure of inflation to decline by any significant amount during the Great Recession.[20]

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 shocks influence endogenous movements in productivity. Figure 8 plots the impulse response functions with and without a binding ZLB.[21] 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.[22]

[20] Indeed, in the estimation, there is no need for positive realization of the cost push shock to match the observed evolution of inflation during and after the Great Recession.

[21] The binding ZLB is achieved with a ten-standard deviation positive shock to the liquidity preference.

[22] 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 the additional amplification and propagation generated by the endogenous contraction in TFP.

Figure 8: Liquidity Demand Shock and the ZLB

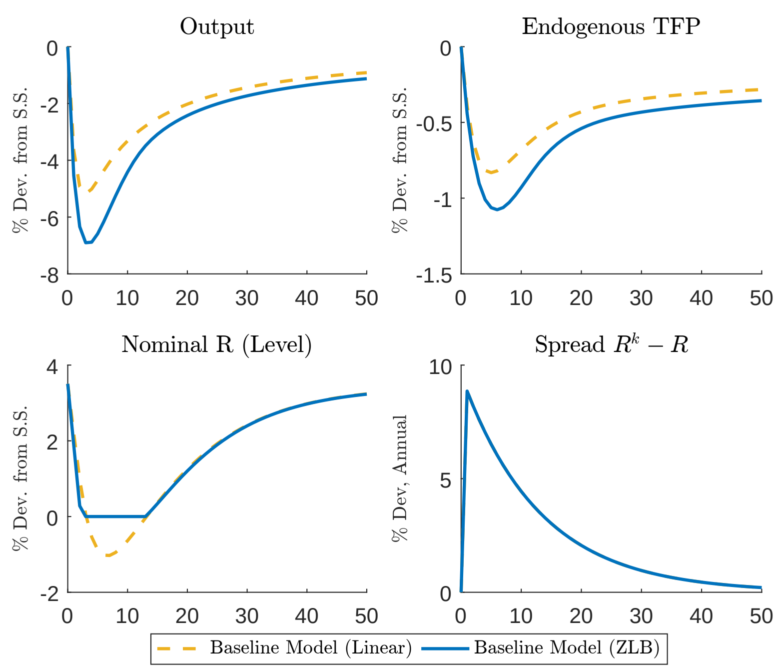

5.2 Technology Diffusion: Model vs. Data

Before turning to the historical evolution of productivity growth, 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.

Recall from Section 3 that the speed of diffusion for the technologies invented k periods ago is

$$Speed_{t+k}^{t} = \log\left(\frac{1+\lambda_{t+k-1}/r_{t+k-1}^{t}}{1-\lambda_{t+k-1}}\right), \tag{61}$$

where the geometric nature of diffusion of the model introduces a vintage-effect on the speed of technology diffusion. We can define this vintage effect as the speed that the technology of a given vintage would have in the absence of business cycle fluctuations. That is, if the adoption rate was equal to the constant steady state level, $\bar{\lambda}$. Formally, the vintage effect on speed is

$$V\_Speed_{t+k}^{t} = \log\left(\frac{1+\bar{\lambda}/r_{t+k-1}^{t}}{1-\bar{\lambda}}\right), \tag{62}$$

Accordingly, we can model the regression equation for the effect of the cycle on the speed of diffusion as

$$Speed_{t+k}^{t} = \alpha + V\_Speed_{t+k}^{t} + \beta * \hat{y}_{t+k} + \epsilon_{t+k} \tag{63}$$

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 the deterministic effect of the lag on the average speed of diffusion. We can rewrite equation (63) as where the left hand side is the speed of diffusion adjusted for the deterministic vintage effects in the model.

$$\widehat{Speed}_{t+k}^{t} \equiv Speed_{t+k}^{t} - V\_Speed_{t+k}^{t} = \alpha + \beta\hat{y}_{t+k} + \epsilon_{t+k} \tag{64}$$

We estimate equation (64) in a synthetic sample constructed from 100,000 period long run of our estimated model. Each period we construct the $\widehat{Speed}$ measure for each vintage. We estimate the panel by weighting each observation by the vintage share of unadopted technologies in the steady state.[23]

Table 5 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 (1.94 versus 3.76) but falls within the 95% confidence

[23] See Appendix for details about how we compute the weights.

interval of the point estimate in the data.

Table 5: Cyclicality of diffusion speed: model versus data

	Data	Model
β	3.73	1.94
	[1.7,5.8]	[1.59,2.29]

β is the elasticity of speed of diffusion with respect to the output gap. Numbers is brackets are 95% confidence intervals.

Table 6 reports a sensitivity analysis of the cyclicality of the speed of diffusion to the values of ρ_λ. We consider values that range from 0.85 to 0.95. The lower the values of ρ_λ the higher the share of R&D in GDP in steady state.[24] For a ρ_λ of 0.85, R&D represents 3.5% of GDP while for a ρ_λ of 0.95 the R&D share is 1.3%. 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 2.26 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.

Table 6: Effect of varying ρ^λ

Value	Diffusion Speed Coefficient	SS R&D Share of GDP (%)
0.85	1.19	3.5
0.9	1.61	2.4
0.925	1.94	1.9
0.95	2.26	1.3

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

5.3 Productivity dynamics

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

[24]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.

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 (11) to derive the following expression that links labor productivity with TFP and capital intensity:[25,26]

$$\frac{Y_t}{L_t} = \underbrace{\theta_t \cdot (A_t)^{\vartheta-1}}_{TFP} \cdot (U_t K_t / L_t)^{\alpha}. \tag{65}$$

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.[27]

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 virtually all of the decline in TFP betweem 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 7.75 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 vary over time. Figure 10 presents a historical decomposition of endogenous productivity that

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

[26] 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.

[27] 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

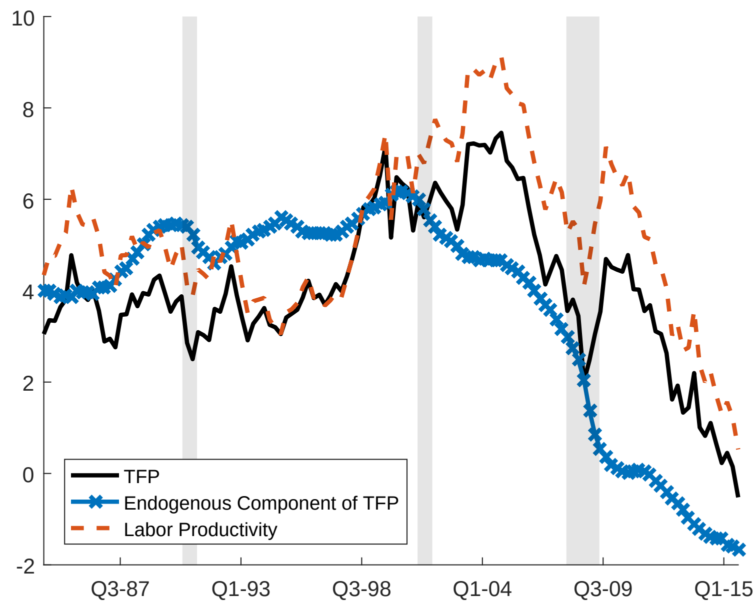

Labor productivity is GDP divided by hours worked (see Appendix A.1 for data sources). Smoothed shocks from model estimated using data as described in Section 4.2 and Appendix A.1.

Figure 10: Endogenous TFP Decomposition

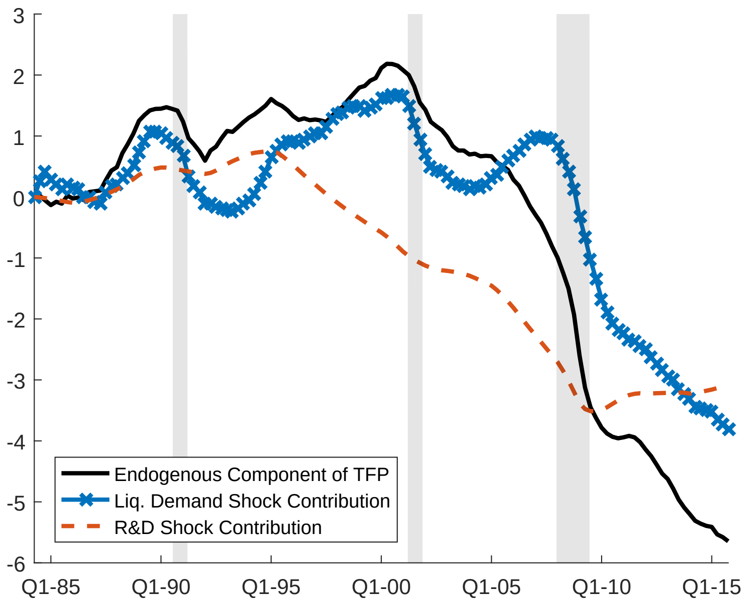

Smoothed variables from model estimated using data as described in Section 4.2 and Appendix A.1.

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 6); 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 given level of R&D spending, which ultimately slowed the pace of new technology adoption. 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 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. Below, we present direct evidence that supports this finding.

Our finding that shocks to R&D productivity mainly account for the pre-recession slowdown in TFP is consistent at a high level with Fernald's (2014) hypothesis that mainly exogenous as opposed to cyclical factors were at work. There are at least two significant differences. Fernald attributes the decline of TFP to the loss of steam of the IT revolution while in our model the decline in R&D productivity is generic. These hypotheses have differnt implications for whether the drop in R&D activity should be largely circumscribed to the computer and software sectors. Using COMPUSTAT, we observe that out of 29 sectors with R&D expenditures of at least \$100 million in 2000,[28] 19 experienced a decline in R&D expenditures between their peak in the period 1998-2000 and 2002. Computers and software were the 9^{th} and 10^{th} with largest declines, of approximately 20% each.[29] This seems to suggest that the decline in the productivity of R&D was also wide-spread as oppposed to just reflecting the exhaustion of the pipeline for computers and software.

A second important difference with Fernald (2014) is that our endogenous productivity mechanism allows cyclical shocks as well shocks to R&D productivity to drive TFP. In this regard, our accounting suggests that once the recession began, it was cyclical shocks in the form of liquidity demand shocks that largely drove the subsequent decline in endogenous TFP.

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 technigies, $Z_t - A_t$, where Z_t is the total stock of technologies. Fluctuations λ_t reflect cyclical variation in adoption on endogenous productivity while

[28] In 2009 dollars.

[29] The largests percentage reductions in R&D are in agriculture (81%), textile and apparel (58%), business services other than software (45%), and petroleum and rubber (38%). The largest increases were found in tobacco (24%), finance, insurance and real state (17%) and chemicals (14%). The average decline was 14% with a median decline of 6%.

fluctuations in Z_t reflect the effect of cyclical varition 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[30] – measured on the right-hand side axis. 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 results that: (i) liquidity demand shocks are important drives of recessions (see Figure 4) and that these shocks can induce contractions in adoption rates and endogenous productivity (See Figure 7). Furthermore, the fact that the elasticity of λ_t in the model is consistent with the data makes plausible the predicted evolution of endogenous TFP

Fluctuations in Z_t do, however, 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 9 identifies.[31] 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 reason endogenous TFP continues to decline.

Interestingly, while λ_t remains low following the Great Recession, the stock of unadoptied technologies, $Z_t - A$, reaches a peak over the sample. The latter 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 increasing aggregate demand 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.

[30]For λ_t we plot teh level of the quarterly adoption rate.

[31]The decline in endogenous productivity induced by the negative shocks to 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

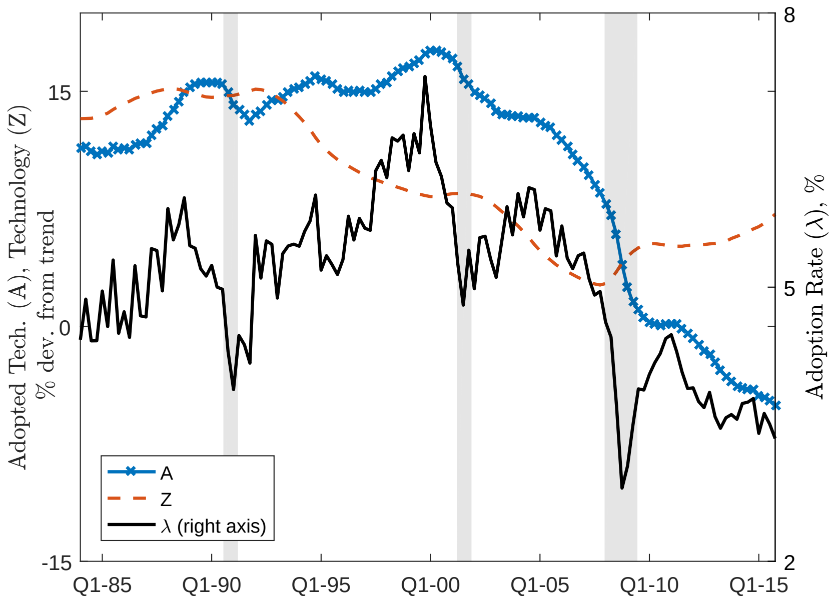

Figure 12: R&D efficiency in data versus model

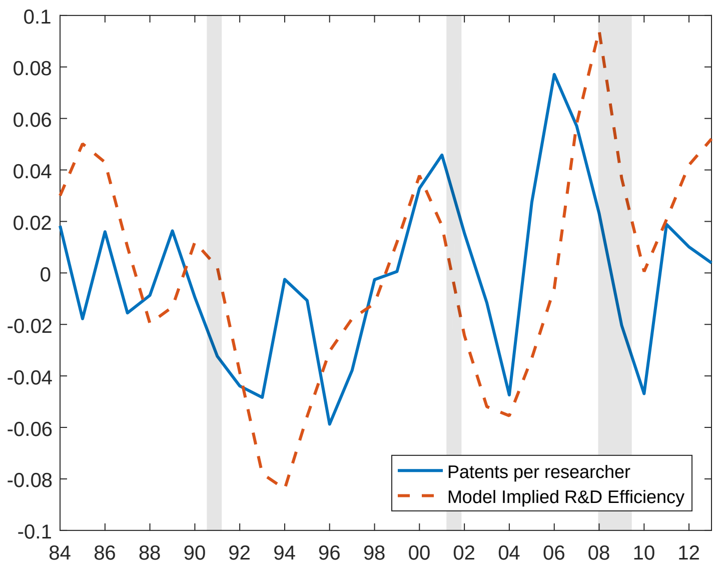

Source: ; linearly detrended level of $R\&D\_prod_{t-1,t}$ *and average of the estimated log-level of* χ_t *between years* $t-1$ *and* t $(\bar{\chi}_{t-1,t})$.

5.3.1 Evidence on the evolution of R&D productivity

We conclude our analysis by providing independent historical evidence 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 are 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 :

$$R\&D\_prod_{t-1,t} = \log\left(\frac{Patents_t + Patents_{t-1}}{2 * R\&D\_emp_{t-2}}\right) \tag{66}$$

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 $t-2$.[32]

Figure 12 plots the linearly detrended level of $R\&D\_prod_{t-1,t}$ together with the average

[32] The patent applications come from the USPTO 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.

of the estimated log-level of χ_t between years $t-1$ and t ($\bar{\chi}_{t-1,t}$).[33] There are three key observations. First, the correlation between $R\&D\_prod_{t-1,t}$ and $\bar{\chi}_{t-1,t}$ is 0.44.[34] Second, as predicted by our estimates, the independent measure of R&D productivity shows a decline after 2001. Indeed, the magnitude of teh decline is very similar in the model estimates of $\bar{\chi}_{t-1,t}$ and in the data. This finding supports our model's prediction that the pre-GR productivity slowdown may partly reflect the decline in R&D productivity after 2001. Third, the measure of $R\&D\_prod_{t-1,t}$ seems 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.[35] Based on these three observations, we conclude that the independent measure of R&D productivity supports our estimates of the evolution of R&D productivity.

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 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

[33] 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.

[34] It increases to 0.48 for the period after 1990, when, presumably, the impact of the new patenting laws on patenting activities have stabilized reducing the noise in $R\&D\_prod$.

[35] Another episode where both series coincide is the decline in R&D productivity after 1991.

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 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))[36]. 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.

[36] 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

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

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.[37]

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:

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

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

[37]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.

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 1.9%). In our baseline calibration, this results in a value of ρ^λ of 0.925. To check that this calibration does not result in 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 1.

There are three conceptual obstacles to overcome in carrying out this comparison. The first is that the data in the regressions of Table 1 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. Analogously to Equation 1, we define

$$m_{t+k}^{t} \equiv \frac{A_{t+k}^{t}}{Z_{t+k}^{t}} \tag{67}$$

and

$$r_{t+k}^{t} \equiv \frac{m_{t+k}^{t}}{1 - m_{t+k}^{t}} \tag{68}$$

where Z_{t+k}^{t} is 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} is the mass of vintage t technologies that have been adopted at time $t+k$. These two measures evolve as follows:

$$Z_{t+k}^{t} = \phi Z_{t+k-1}^{t} \tag{69}$$

$$A_{t+k}^{t} = \phi A_{t+k-1}^{t} + \lambda_{t-1}\phi \left(Z_{t+k-1}^{t} - A_{t+k-1}^{t} \right) \tag{70}$$

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:

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

and

$$r_{t+k}^{t} = \frac{r_{t+k-1}^{t} + \lambda_{t-1}}{(1 - \lambda_{t-1})} \tag{71}$$

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 of the remaining unadopted technologies is adopted in every time period (note that all unadopted technologies, irrespective of vintage, have the same probability λ_t of being adopted). We define the speed of diffusion at time $t+k$ for a vintage t technology as

$$Speed_{t+k}^{t} \equiv \Delta \log \left(r_{t+k}^{t} \right) = \log \left(\frac{r_{t+k}^{t}}{r_{t+k-1}^{t}} \right) \tag{72}$$

The regression analysis we conduct measures the sensitivity of the speed of technological diffusion to fluctuations in the output gap. If, as is commonly posited in the literature, the fraction of adopters is a logistic function of time, 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 $\widehat{Speed}^t_{t+k}$, defined as follows:

$$\widehat{Speed}^t_{t+k} \equiv \log\left(\frac{r^t_{t+k}}{\hat{r}^t_{t+k}}\right) \tag{73}$$

where

$$\hat{r}^t_{t+k} = \frac{r^t_{t+k-1} + \bar{\lambda}}{1 - \bar{\lambda}} \tag{74}$$

Intuitively, $\hat{r}^t_{t+k}$ is the value that r^t_{t+k} would take if the diffusion process returned to its steady state evolution, $\lambda_{t-1} = \bar{\lambda}$. Our detrended speed therefore measures cyclical variation in the diffusion process. The data used in our regression analysis is a panel of technologies.

Figure 13: *Speed* and $\widehat{Speed}$: Simulated and Steady State

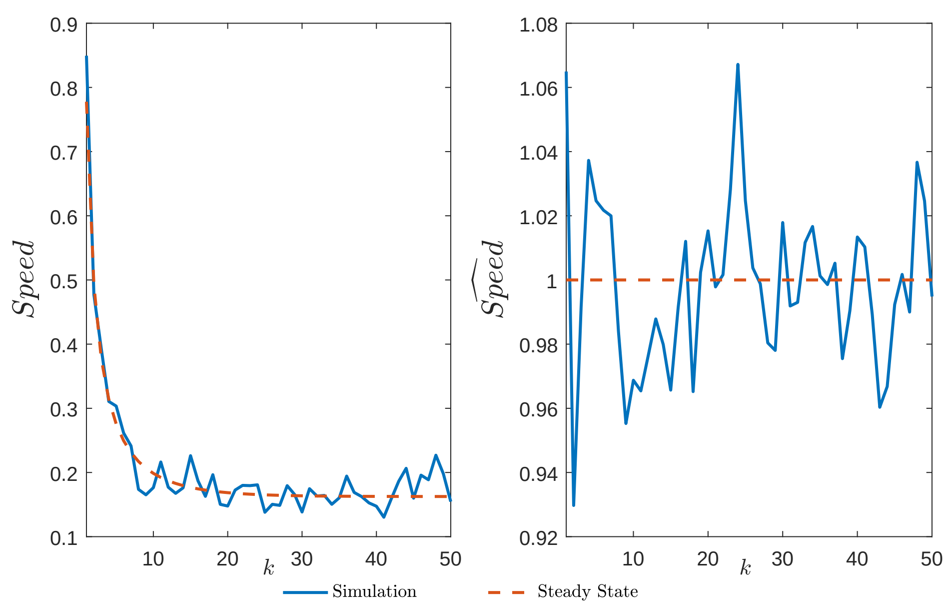

This figure plots Speed (left panel) and detrended $\widehat{Speed}$ *(right panel) as a function of the age of a technology* k*. In both cases the dotted line is the steady state value of speed and the solid line is a typical model simulation.*

In the model however the relative masses of technologies of different vintages is not constant over time, due to three factors: 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 first run the following regressions using model-simulated data

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

where k denotes the age of a technology. 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 $\bar{w} = \frac{1}{K}\sum_{j=1}^{K} w^j$.

A.4 Other Tables and Figures

As a check of the fit of the estimated model, Table 7 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 7: Comparison of Standard Deviations

Variable	Data	Model
Output Growth	0.55	0.61
Consumption Growth	0.51	0.71
Investment Growth	1.54	1.51
Inflation	0.23	0.36
Nominal R	0.60	0.55
Hours	1.82	1.53
R&D Expenditure	4.00	6.82

Table 8: Prior and Posterior Distributions of Shock Processes

Parameter	Description	Prior			Posterior	
		Distr	Mean	St. Dev.	Mean	St. Dev.
σ_{ϱ}	Liq. Demand	Inv. Gamma	0.10	2.00	0.222	0.0013
σ_{χ}	R&D	Inv. Gamma	0.10	2.00	2.212	0.1796
σ_{g}	Govt. exp.	Inv. Gamma	0.10	2.00	2.560	0.0351
σ_{mp}	Monetary	Inv. Gamma	0.10	2.00	0.096	0.0001
σ_{μ}	Markup	Inv. Gamma	0.10	2.00	0.093	0.0002
σ_{pk}	Investment	Inv. Gamma	0.10	2.00	1.271	0.0094
σ_{θ}	TFP	Inv. Gamma	0.10	2.00	0.489	0.0013
σ_{μ_w}	Wage markup	Inv. Gamma	0.10	2.00	0.287	0.0014
ρ_{ϱ}	Liq. Demand	Beta	0.50	0.20	0.926	0.0005
ρ_{χ}	R&D	Beta	0.50	0.20	0.802	0.0076
ρ_{g}	Govt. exp.	Beta	0.50	0.20	0.968	0.0001
ρ_{mp}	Monetary	Beta	0.50	0.20	0.471	0.0071
ρ_{μ}	Markup	Beta	0.50	0.20	0.405	0.0180
ρ_{pk}	Investment	Beta	0.50	0.20	0.900	0.0014
ρ_{θ}	TFP	Beta	0.50	0.20	0.951	0.0008

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.

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.

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. and H. Kung (2014): "Growth, Slowdowns, and Recoveries," Working Paper 20725, National Bureau of Economic Research.

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.

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.

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," Tech. rep.

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

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. (2013): "The Financing of Ideas and The Great Deviation," .

——— (2015): "The Financing of Ideas and the Great Deviation," .

Gilchrist, S. and E. Zakrajsek (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102, 1692–1720.

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.

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.

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.

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.