

An Exploration of the Japanese Slowdown during the 1990s

Diego Comin*

Harvard University and NBER

November 2008

Abstract

Why was the 1990s a lost decade for Japan? How is it possible that the Japanese economy stagnated for a decade if none of the shocks that arguably hit the economy seemed to have persisted for much more than three years or so? In this paper I show that the endogenous development and adoption of technologies can propagate these shocks making their effect much more persistent. When feeding the markup shocks observed in Japan during the early 1990s, the model is able to generate time series for output, TFP, employment, consumption and investment that track closely the actual data. In particular, the productivity slowdown in the model is as protracted as in the data. Reassuringly, I also find evidence that, as predicted by the model, the speed of technology diffusion slowed down in Japan during the 1990s and R&D expenditures also stopped growing.

1. Introduction

Two striking aspects of the Japanese stagnation of the 1990s are its severity and especially its persistence. Over the 1983-1991 period, TFP grew at a more than respectable rate

*I would like to thank for very helpful comments my dicussant Gialuca Violante, the editors Koichi Hamada, Anyl Kashyap and David Weinstein and the participants at the ESRI workshop at Columbia University, and, for assistance, the ESRI. All remaining errors are my own.

of 2.4 percent. It fell to an average of 0.2 percent for 1991-2000. We know from Cogley and Nason (1995) that the persistence in real business cycles models comes from the persistence of the shocks. However, none of the shocks that arguably have hit the Japanese economy seem to have persisted for a decade. This implies that a successful account of the Japanese stagnation requires a model where the propagation mechanisms generate a significant amount of persistence.

One such approach is to modify a simple real business cycle model to allow for endogenous productivity as in Comin and Gertler (2006). This approach provides a unified explanation for the co-movements of TFP, technology diffusion and R&D over the short and medium term. To capture endogenous productivity dynamics, Comin and Gertler (2006) use a variation of Romer's (1990) model of R&D expanding the variety of intermediate goods extended to allow for an endogenous rate of adoption of new technologies. Building on the evidence from Hall (1997) and others, they use (low persistence) wage markup shocks as source of fluctuations. They show that, a version of the model calibrated to match the US economy, can closely reproduce US postwar fluctuations in output, consumption, hours worked, investment, TFP, R&D and the relative price of capital at the high and medium term frequencies.

In this paper, I use the Comin and Gertler (2006) model to explore whether it can amplify and propagate the wage markup fluctuations observed in Japan over the 90s to generate the Japanese productivity slowdown. As in Comin and Gertler (2006), rather than taking a strong stand in the fundamental nature of the shocks, we take the wage markup as a reasonable reduced form measure of several shocks that can be represented as markups.

As argued in Hall (1997) and Gali, Gertler and Lopez-Salido (2007), among others, there are many shocks that are isomorphic to the wage markup shocks. Frictions in the labor markets, labor income taxes, shocks to the entry costs, changes in competition policies, even monetary policy in the presence of nominal rigidities all those can be expressed in a way whose reduced form collapses to our wage markup shocks.

Even more important is the finding in Comin, Gertler and Santacreu (2008) that once the process of endogenous technology adoption is modeled, the responses of the economy both in terms of the amplification and propagation properties to shocks of nature as diverse as credit disruptions, shocks to the price of capital, markup shocks and TFP

shocks are quite similar specially qualitatively but also in a quantitative sense. In my view, this finding shifts the focus of explanations for the Japanese lost decade from the shocks to the propagation mechanisms.

This paper is related to Hayashi and Prescott (2002) who show that a real business cycle model with TFP shocks can reproduce the Japanese fluctuations over the 1990s in output, consumption and investment. This implies that standard macro models can go a long way in explaining why Japan suffered its growth slowdown conditional on the evolution of TFP. However, technology and TFP are to a large extent endogenous to the agents decisions. Thus, the follow up question is why Japanese TFP was depressed during the 90s.

In section 2, I explore in more detail the Japanese performance during the 90s. Two interesting findings are that R&D expenditures stopped growing and that the speed of diffusion of new technologies slowed down (to the point that South Korea surpassed Japan in both the diffusion of computers and internet). This suggests that the endogenous technology mechanisms played a significant role in the propagation of the shocks that caused the Japanese slowdown in productivity.

At this point there is not a large number of alternative explanations that could explain the persistence of the Japanese productivity slowdown. One possibility is that credit frictions persisted. This, however does not seem to have been the case. Beyond the Hayashi and Prescott arguments, survey data shows that tighter lending practices by Japanese banks in the early part of the decade lasted for about a year while after the financial difficulties of 1998 they lasted for no more than that (Hoshi and Kashyap, 2008).

An alternative hypothesis is that credit was misallocated to unproductive companies (zombies) that had close connections to banks. This clearly will lead to a lower productivity of misallocated capital (Peek and Rosengren, 2005). However, the mere fact that these companies became zombies in the first place is a sign that their capital was not very productive previously and hence it is difficult that this by itself explains the large decline in productivity experienced by Japan during the 90s. Caballero, Hoshi and Kashyap (2008) argue that zombies can reduce the productivity of non-zombies because, in addition to capital, they utilize unproductively other sector-specific inputs. This argument is very interesting but, it is not obvious what the sector-specific inputs are. In any case, I regard these explanations for the persistence of the Japanese response to the initial shocks as

complementary to the mechanisms introduced in the model below.

Section 3 describes the model and the intuition of why it generates pro-cyclical and highly persistent fluctuations in productivity despite the fact that shocks are relatively transitory. Section 4 presents the impulse responses to a wage markup shock and the simulations of the Japanese economy once I hit it with the observed series of markup shocks. The model does a good job in reproducing the evolution of output, consumption, investment, TFP and hours worked in Japan during the lost decade, specially up to 1998. During the last two years of the decade, the propagation mechanisms in the model seem to run out of steam, while the Japanese economy continued to deteriorate. However, I think it is fair to say that overall the model provides a powerful propagation mechanism to the fluctuations in the markups we saw in Japan. Section 5 concludes.

2. Facts

In this section, I present a series of facts with two purposes in mind. First, to describe the main features of the data that I consider a model of the Japanese slowdown should account for. Second, to motivate the mechanisms built in the model.

Figure 1 presents the evolution of GNP per working age person¹ detrended with a linear trend given by the average post-war growth rate of the same variable in the U.S. which has been approximately 2.4% per year. In this and most of the subsequent figures, I take 1990 as the base year. This Figure illustrates clearly the magnitude of both the post-war catch up and the slow down of growth during the 1990s. between 1960 and 1990 per working age person GNP grew at an average rate of 4.6% per year. This rate declined to 0.7% per year during the lost decade.

Hayashi and Prescott (2002) show that the slowdown in TFP growth (or GNP growth) during the 1990s was not due to a constraint in investment due to a lack of capital to finance profitable investment projects. In particular, they show that there was no decline during the 1990s in investment by non-financial corporations as a share of GDP. The collapsed in the banking system, however, forced firms to find other sources of funds different from the bank loans that had helped finance 52% of the investment by non-

¹That is defined as persons with ages between 20 and 69 years.

financial corporations during the 80s. These were a higher corporate savings rate and the sale of the land and financial assets that had been accumulated during the 80s. Hayashi and Prescott (2002) also show that small firms, which rely more on bank loans, used cash and deposits as a buffer to finance their investment.

As argued by Hayashi and Prescott (2002), the slowdown in growth is driven by—largely by the slowdown in TFP growth. Figure 2 compares the evolution of TFP in Japan and the U.S. both detrended at the post-war rate of 1.02% per year. Despite the impressive performance of US TFP during the 60s, Japan closed a significant part of the gap with the U.S. Between 1970 and 1990, Japan continued to close the gap while growing faster than the US long-run average. The 1990s, however, is a different story. While US managed to approximately sustain its long-run TFP growth, Japan’s TFP grew significantly slower at 0.6% per year.

Two of the most significant sources of TFP growth in developed economies are innovation and technology adoption. Figures 3 show that Japanese R&D spending also suffered a significant slow down after 1990. Figure 4 shows that the growth rate of real wages in research activities slowdown more than in the rest of the sectors. This suggests slower demand for scientists and engineers in R&D than in other activities.

Figures 5 and 6 show the diffusion of computers and internet in Japan and compares it to the diffusion in the U.S. and in South Korea. Computers and the internet are the most significant technologies during the 1990s. Their significance resides partly in their generality which makes them important to many activities in many different sectors of the economy. Thus, they provide a good proxy for the overall technology adoption in the economy.

During the 1990s the gap in computer adoption between Japan and the U.S. increased. The slow speed of computer adoption in Japan is reflected by the fact that Korea, a significantly poorer economy, surpassed Japan in computers per capita. An almost identical picture emerges from the diffusion of the internet.

The question that these findings beg to ask is why did R&D and the speed of adoption decline in Japan. It is natural to think that this decline was driven by a decline in the profitability of the investments to improve technology. But, what drove this lower profitability?

A priori, there are many possible answers to this question. One is that aggregate

demand was lower over this period. Aggregate demand can also be driven by a variety of shocks. Gali, Gertler and Lopez-Salido (2007) (GGLS, henceforth) present a variable, that they call the total mark up, which can be thought as a reduced form of several shocks which affect aggregate demand. Specifically, they define the total markup (or markup, in short) as the wedge between the marginal product of labor and the marginal rate of substitution between consumption and leisure. To compute this variable it is necessary to make some parametric assumptions about the utility and production functions. I assume a Cobb-Douglas production function and a unit elasticity between consumption and leisure. Of course, demographic and other very slow moving variables may be affecting this measure of the markup. So, I filter these using a Band-Pass filter that keeps cycles with periods between 2 and 100 quarters. Gali, Gertler and Lopez-Salido (2007) show that the markup is highly counter-cyclical at business cycle frequencies. Comin and Gertler (2006) show that it is highly countercyclical also at the medium term (i.e. fluctuations associated to cycles with periods between 8 and 50 years). Figure 7 plots the filtered markup for Japan. In this figure we can see that the Japanese wage markup was high during the first half of the 70s and between 1988 and 2000. This is the variable we will use in section 4 to measure the disturbances in the Japanese economy.

3. The Model

I now develop a model of medium term business fluctuations. The model is annual as opposed to quarterly because I am interested in capturing fluctuations over a longer horizon than is typically studied. To this end, I abstract from a number of complicating factors that would otherwise might be useful for understanding quarterly dynamics, such as money and nominal rigidities.

The Comin and Gertler (2006) model is a two-sector version of a reasonably conventional real business cycle model, modified to allow for, among other things, endogenous productivity and endogenous countercyclical markups. Procyclical entry and exit by final goods firms induces procyclical variation in the degree of competition and, hence, countercyclical variation in (final goods) price markups, in the spirit of Rotemberg and Woodford (1995). The precise formulation we use, though, is based on Gali and Zilibotti (1995).

Final goods producers use intermediate goods as inputs, along with capital and labor. As in Romer (1990), creation of new specialized intermediate goods, stemming from R&D is the source of technological change. As I noted in the introduction, we modify the Romer framework to allow for an endogenous rate of adoption of new technologies, along with endogenous R&D. By doing so, I can allow for empirically reasonable diffusion lags but still generate endogenous medium term swings in productivity. Roughly speaking, endogenous R&D permits disturbances to have permanent effects on productivity movements, while endogenous adoption rates accelerate the transition of productivity to the new steady state. Because the model has two sectors, consumption goods versus capital goods, we are able to distinguish between embodied and disembodied technological change.

Household's are conventional, except I allow them a bit of market power in labor supply in order to motivate a wage markup. I then allow for exogenous variation in this market power. This mechanism combined with the entry/exit mechanism in the goods market permits the model to generate countercyclical movements in both price and wage markups that are consistent with the data. As I alluded to earlier, this simple structure is meant as a stand-in for a richer model of markup variation.

3.1. Main Mechanisms

Rather than laying formally the whole model,² next I just explain in detail the mechanisms that propagate the wage markup shocks generating large and persistent fluctuations in productivity and TFP. I do that by explaining the intuition behind the key equilibrium relationships in the model.

In our model, wage markup shocks, $\mu_{w,t}$, increase or reduce the wedge between the wage, W_t , and the marginal rate of substitution between labor and consumption, $(L_t)^\zeta C_t$.

$$W_t = \mu_{w,t} \cdot (L_t)^\zeta C_t \tag{1}$$

A positive wage markup, reduces, *ceteris paribus*, the number of hours that the representative worker supplies for a given wage rate. This reduction in the number of hours supplied in the market causes a decline in aggregate output. The expectations about future

²For that, see Comin and Gertler (2006).

lower output also drive a decline in aggregate consumption and investment. This reduces the profits earned by intermediate goods producers in both the capital and consumption goods sector, $\Pi_{x,t}$ for $x = c, k$. The value of a company that produces an intermediate good, $V_{x,t}$, is given by the discounted stream of earnings from selling intermediate good:

$$V_{x,t} = \Pi_{x,t} + R_{t+1}^{-1} \phi E_t \{V_{x,t+1}\} \quad (2)$$

where R_{t+1} is the discount rate between t and $t+1$, and ϕ is the probability the good has not exogenously disappeared the next period. At times when aggregate demand is lower, so are current and future profits and the value of intermediate good producers is also lower. When intermediate goods are first invented, they are not yet usable for production. For that, the inventor must succeed in making the good usable for production (i.e. adopt it). The more output he invests in adopting the good, the higher the probability he will succeed in adopting it. The value of an unadopted technology, $J_{x,t}$, can be defined by the following Bellman equation:

$$J_{x,t} = \max_{H_{x,t}} -H_{x,t} + R_{t+1}^{-1} \phi [J_{x,t+1} + \lambda(\Gamma_{x,t} H_{x,t})(V_{x,t+1} - J_{x,t+1})] \quad (3)$$

where $H_{x,t}$ is the amount of output invested in adopting the intermediate good, $\lambda(\cdot)$ (with $\lambda' > 0$, and $\lambda'' < 0$) is the probability the good becomes usable for production next period, and $\Gamma_{x,t}$ is a smooth scaling factor which ensures balanced growth. The owners of an intermediate good that has not been adopted yet, incur in the adoption costs until they succeed in making the good usable for production. At this point, they obtain a capital gain because the value of the intermediate good becomes $V_{x,t+1}$ which is greater than $J_{x,t+1}$. In recessions, the future value of adopted intermediate goods is lower, and so is the value of unadopted intermediate goods.

Next, I describe the laws of motion for the total number of intermediate goods in sector x , $Z_{x,t}$, and for the number of adopted intermediate goods, $A_{x,t}$. The law of motion for the total number of intermediate goods is given by

$$\frac{Z_{x,t+1}}{Z_{x,t}} = \chi_x \left(\frac{S_{x,t}}{\Psi_t} \right)^\rho + \phi \quad (4)$$

where $S_{x,t}$ is the amount of output spend in *R&D* activities in sector x , and where χ_x and

Ψ_t are scaling factors. Similarly, the law of motion for the number of adopted intermediate goods in sector x , $A_{x,t}$, is:

$$\frac{A_{x,t+1}}{A_{x,t}} = \lambda(\Gamma_{x,t}H_{x,t})\phi\left[\frac{Z_{x,t}}{A_{x,t}} - 1\right] + \phi \quad (5)$$

Intuitively, $Z_{x,t} - A_{x,t}$ is the stock of unadopted intermediate goods in sector x . A fraction λ_t of those will be adopted at time t . Given $V_{x,t}$ and $J_{x,t}$ and the laws of motion for $A_{x,t}$ and $Z_{x,t}$, the following two equations determine the amount of resources invested in R&D, $S_{x,t}$, and adoption, $H_{x,t}$.

$$(R_{t+1}^{-1}\phi) \cdot E_t J_{x,t+1}(Z_{x,t+1} - \phi Z_{x,t}) = S_{x,t} \quad (6)$$

$$\Gamma_{x,t}\lambda' \cdot R_{t+1}^{-1}\phi[V_{x,t+1} - J_{x,t+1}] = 1 \quad (7)$$

Equation (6) is a free entry condition that states that the discounted value of new intermediate goods created by undertaking R&D activities must equal the total cost. Equation (7) determines the optimal investment in adoption which is the level at which the marginal cost of investing output in adoption (RHS) equalizes the expected marginal benefit (LHS) which is given by the marginal increase in the probability of successfully adopting the intermediate good times the capital gain from being adopted (i.e. $V_{x,t+1} - J_{x,t+1}$).

These two equations illustrate why both R&D and adoption investments are procyclical. When the economy is hit by a positive wage markup shock causing a recession, the value of not adopted intermediate goods declines. Equation (6) implies that this will lead to a lower flow of R&D investments. The value of adopted intermediate goods declines even more than the value of unadopted intermediate goods, since the former is affected immediately by the current decline in profits while the latter is affected only in the future when the technology is adopted. This decline in $V_{x,t+1} - J_{x,t+1}$ reduces the marginal benefit from investing in adoption, and from (7) the flow of output devoted to adopt new technologies.

These temporary declines in the investments in improving the existing technology through R&D and adoption have temporary effect on the growth rate of technology, measured by $Z_{x,t}$ and $A_{x,t}$, as illustrated by (4) and (5). As a result, the effects on the

level of technology are permanent. This is the key to understanding the propagation power of the model.

I conclude this section by linking the growth rate of new technologies with TFP growth. It is possible to express the value added in sector x , $Y_{x,t}^v$, as:

$$Y_{x,t}^v = \theta_{x,t} K_{x,t}^\alpha L_{x,t}^{1-\alpha} \quad (8)$$

where $\theta_{x,t}$ is the Solow residual in sector x , and $K_{x,t}$ and $L_{x,t}$ denote respectively its capital and labor. In equilibrium, $\theta_{x,t}$ is given by

$$\theta_{x,t} = \left(1 - \frac{\gamma}{\mu_{x,t}\vartheta}\right) N_{x,t}^{\mu_{x,t}-1} \left(\frac{\gamma(A_{x,t})^{\vartheta-1}}{\mu_{x,t}}\right)^{\frac{\gamma}{1-\gamma}} U_t^\alpha. \quad (9)$$

There are a few new variables in this expression. $\mu_{x,t}$ is the gross price markup, $N_{x,t}$ is the number of final output producers in sector x , and U_t is the utilization rate. The capital utilization rate, U_t , enters $\theta_{x,t}$ because it reflects unmeasured intensity of input usage. The term $N_{x,t}^{\mu_{x,t}-1}$ reflects gains from variety due to short run fluctuations in the number of final output producers due to entry and exit. The markup $\mu_{x,t}$ also enters $\theta_{x,t}$ due to the roundabout input-output nature of production. On balance, U_t , $N_{x,t}^{\mu_{x,t}-1}$, and $\mu_{x,t}$ contribute to procyclical variation in $\theta_{x,t}$, particularly at the high frequencies.

Observe that $\theta_{x,t}$ also increases with $A_{x,t}$ (since $\vartheta > 1$). We should expect the influence of $A_{x,t}$ on $\theta_{x,t}$ to play out mainly over the medium and low frequencies, since $A_{x,t}$ is a state variable that moves smoothly in response to cyclical variations in endogenous investments in R&D and technology adoption processes.

4. Model Simulations

In this section I explore the ability of the model to generate medium term cycles. I solve the model by loglinearizing around the deterministic balanced growth path and then employing the Anderson-Moore code, which provides numerical solutions for general first order systems of difference equations. I use the same calibrated values for the model parameters as in Comin and Gertler (2006). When calibrating those parameters we make extensive use of the steady state relationships of the growth model. By using the same

values for Japan and the US, I am implicitly assuming that along the balanced growth path the variables in both countries grow at the same rate. That seems a reasonable assumption. It is however difficult to prove given the Japanese experience over the last 50 years. Between 1950 and 1985, Japan dramatically reduced its productivity gap with the US and quickly converged to approximately 80% of the US per capita income level. Since 1990, however, Japan has grown much slower than the US and the gap between the two has increased. Neither of these seem to correspond to the balanced growth path of Japan, but rather to persistent transitional dynamics which a model like ours should attempt to explain. Hence, as a first pass, I find reasonable the strategy of assuming similar parameters in the US and Japan and see how far such a model can go in explaining the Japanese fluctuations.

4.1. Some Numerical Experiments

To be clear, the exercises that follow are meant simply as a first pass at exploring whether the mechanisms I emphasize have potential for explaining the data: They are not formal statistical tests. As in Comin and Gertler (2006), I treat innovations to the wage markup, μ_t^w , stemming from exogenous fluctuations in household labor market power as the main source of disturbances to the economy. As discussed above, one should keep in mind, though, that this simple mechanism is meant as a short-cut for a richer description of countercyclical wage markup behavior.³

Ideally, we would like to evaluate the model against moments of the data that are conditional on a shock to μ_t^w . There are several difficulties with this approach, however. First, while identified VAR methods may be useful for measuring the near term effects of shocks to the economy, they usually are very imprecise at measuring the kind of medium term effects that are the focus of our analysis.⁴ Second, identification of orthogonalized movements in μ_t^w is highly problematic, particularly given that in practice there is likely a strong endogenous component to this variable.

Rather than attempt to match imprecisely estimated conditional moments, we instead

³As discussed in GGLS (2007), a model with either nominal or real wage rigidities can generate a countercyclical wage markup.

⁴In particular, long run effects of impulse response functions from identified VARs are measured much less precisely than are short run effects.

follow the lead of the traditional real business cycle literature by exploring how well a single shock presumed to be the principle driving force can explain the unconditional patterns in the data. We view this kind of exercise as a check on the plausibility of the scenario we offer, as opposed to a clear test against alternatives. While the RBC literature focused on the ability of technology shocks to account for the high frequency variation in the data, we instead consider the ability of our model driven by shocks to μ_t^w to explain the evolution of the Japanese macro variables after removing a linear trend. As a result, I will try to match not only the high frequency fluctuations in the Japanese data but also medium term fluctuations.

Impulse response functions

Before confronting the data, we first gain some intuition for the workings of the model by examining the model impulse responses to a shock to μ_t^w . In Figure 8, the solid line in each panel is the response of the respective variable within our model. For comparison purposes, I also include the responses (dotted lines) generated by a model that is very similar to the Neoclassical model since it does not have endogenous productivity or endogenous markup variation (though I allow for endogenous factor utilization). The period length is a year.

The rise in μ_t^w effectively raises the price of labor, reducing labor demand and output. Both the initial impact and the impact over time on output is larger in our full blown model than in the benchmark model. Within the full blown model, the decline in output leads to a rise in both the consumption and capital goods price markups (due to a decline in entry within each sector), enhancing the initial drop in output. Over time, output climbs back to trend, but does not make it back all the way due to the endogenous permanent decline in productivity, that we describe shortly. After a decade or so it levels off at a new steady state. The permanent decline in output is about one half of the initial decline. In the benchmark model without endogenous productivity, output simply reverts back to its initial steady state.

Even though the shock is not technological, the rise in μ_t^w generates a persistent decline in TFP and labor productivity and a persistent increase in the relative price of capital. The initial decline in measured TFP and labor productivity results mainly from variable

factor utilization (in conjunction with overhead costs) and the rise in the price markups.⁵ Over time, there is a decline in true productivity relative trend. In particular, the initial contraction in economic activity induces a drop in both R&D and the rate of technology adoption.⁶ The initial decline in R&D slows the rate of creation of new intermediate products in both the goods and capital sectors, ultimately leading to a permanent drop relative to trend in total factor productivity and labor productivity. The slowdown in the adoption rate leads to a near term decline in disembodied productivity growth. Absent this endogenous slowdown in technology adoption, the decline in productivity following the drop in R&D would take much longer, assuming realistic diffusion lags.

Similarly, variable price markup behavior accounts for the initial rise in the relative price of capital, while the increase over time is due to a decline in productivity in the capital goods sector relative to the consumption goods sector. Because investment initially falls proportionately more than output, the capital goods markup rises relative to the consumption goods markup, leading to a jump in the relative price of capital. The disproportionate drop in investment also implies that expected profits in the capital goods sector fall relative to the consumption goods sector. As a consequence, there is a relatively greater drop in both R&D and technology adoption in the capital goods sector, leading to a slowdown in the creation and adoption of specialized intermediate goods, relative the consumption goods sector. This leads to a permanent increase in the relative price of capital.

The (simulated) Japanese slowdown

The final step in the evaluation of the potential for the model to account for the Japanese Slowdown is to simulate the time series once the model economy is hit by the wage markup shocks observed in the data. As mentioned in Section 2, there were positive

⁵The formal definition of TFP that we adopt is the following: $TFP_t = \frac{Y_t A_{k,t}^{\psi\alpha(\theta-1)}}{K_t^\alpha L_t^{1-\alpha}}$, where ψ is the fraction of the variety externality in the creation of new capital not captured by the linear aggregation methodology used by the BEA to compute the capital stock. We calibrate ψ to 0.31 to match the long run growth rate of TFP implied by the model to the postwar data. In case the variety externality was completely mismeasured, annual TFP growth would be 1.72% instead of the actual 1.12%. In addition, this definition assumes that the BLS measures of TFP do not correct properly for capacity utilization.

⁶I am assuming that that R&D involves the same factor intensities as goods production. It is thus not as sensitive to wage behavior as in the Schumpeterian models of Aghion and Howitt (1992) and Aghion and Saint-Paul (1991). Barlevy (2003) shows that if productivity shocks in the final goods sector are sufficiently persistent, R&D can be procyclical even in a Schumpeterian model.

innovations to the markup until 1993 when it reached the maximum and since then started to decline (see Figure 7).

Figure 9 plots the evolution of GDP implied by the model and compares it to the actual evolution of GDP in Japan.⁷ To a first order, the model seems to predict an evolution for GDP over the 90s that tracks quite closely the data. In particular, it predicts that between 1990 and 1995 GDP declines by 11% relative to trend while in the data it declines by approximately 9.5 %, Then it recovers slightly for the next couple of years very much like the data. The biggest discrepancy between model and data seem to occur after 1998, where the model predict a mild recovery of GDP by 2% while in the data it decline by 2.5%. However, the two series are quite close and the fluctuations in GDP induced by the model are both very persistent and large. Next section, we argue that a deterioration of the credit conditions around 1997-1998 could have generated a second shock. This shock was however quite transitory and to explain the further deterioration in the economy until the end of the decade some propagation mechanism, such as endogenous adoption, would be needed.

Figure 10 reproduces the model predictions and actual data for TFP relative to a linear trend.⁸ As with output, the model predictions follow reasonably closely the actual TFP series for Japan during the 1990s. One could argue that the fit of the TFP series in Figure 10 is not surprising in the light of the finding in Hayashi and Prescott (2002) (and the RBC literature) that once we shock the economy with the observed TFP series, standard models are able to reproduce output fluctuations. However, the important difference is that in our model TFP is completely endogenous, and, as shown in Figure 7, the persistence of the macro variables (including output and TFP) is much larger than the persistence of the shocks.

As emphasized above, this endogenous propagation of the model is driven by the endogenous technology mechanisms. One of these is the development of new technologies through R&D. Figure 11 compares the evolution of R&D in the model and in the data. Qualitatively the model follows the data fairly closely. In both detrended (real) R&D

⁷Both of these series have been detrended using a linear trend of 2.4 % per year. In the data GDP has been also divided by the number of people with ages between 20 and 69 years to control for very low frequency fluctuations driven by demographic factors.

⁸To be consistent with the calibration, we detrend TFP to the post-war average in the US of 1.12% per year, as measured by the BLS.

expenditures decline from 1990 until 1995 and then they rebound. It is remarkable then that the model does not predict more persistent R&D expenditures than what we observe in the data. However, a feature of the model already pointed out in Comin and Gertler (2006) is that it generates R&D expenditures that are approximately twice as volatile as in the data. As a result, R&D declines more during the 90s in the model than it actually declined in the data.

Next we turn to hours worked. Figure 12 displays the evolution of actual and predicted hours worked per person between 20 and 69 years old. The model does a fair job in reproducing the initial decline in hours worked. In the data, the minimum is reached around 1994 where hours worked is down 4.5%. The model predicts also a minimum around the same time but the decline is a bit less than 7%. Over the next three years, hours worked increase in both model and data. In the last two or three years is where model and data differ the most. In the model hours worked increases by 2.5% relative to the level in 1998 while in the data it declines until reaching a level 6% lower than in steady state.

Figure 13 and 14 display the evolution of consumption and investment per working age population. For both of these variables, the model tracks quite closely the data with the caveat that the model does not fully capture the declines in these variables over the last couple of years of the decade. In any case, I think it is fair to say that the model does a surprisingly good job in explaining the Japanese slowdown during the 90s specially given that I have calibrated it using the US parameter values and that there are a number of factors completely missing from the analysis which must play a role. I proceed to discuss them in the concluding section.

5. Conclusions

Why was the 90s a lost decade for Japan? This paper has argued that though the shocks that hit the Japanese economy were not so persistent, the investments that Japanese companies and entrepreneurs did not undertake to improve the technology and the production methods during this decade propagated those shocks and made their effects very long-lasting. These propagation mechanisms are central to understanding the persistence of the slowdown since the standard neoclassical model generates a quick recovery once the

shock is over. Of course, these propagation mechanisms are not unique to Japan. Indeed, as shown by Comin and Gertler (2006), the exact same model we have used to explain the Japanese Lost Decade does a very good job in accounting for the macro fluctuations in the US over the postwar period.

One interesting point of departure between the simulated and the actual series took place around 1997/8 where the model predicts that GDP should be flat relative to trend while in reality the gap increased further. This further decline in GDP may reflect the deterioration of credit conditions around 1997-1998. Several paper in this volume provide micro evidence of this hypothesis and Figure 15 provides some macro evidence. In particular, it plots the gap between the marginal product of capital and the marginal rate of substitution between consumption today and tomorrow filtered using the same procedure I have used above with the markup shocks. By measuring the gap between the return to capital and the compensation required by savers, this measure captures disruptions in credit markets. The tightening of credit conditions that peaked, according to Figure 15 in 1997, corresponds with the Japanese banking crisis.

This shock was however quite transitory. As a result, its direct effect on investment and output is unlikely to be sufficiently persistent to explain the deterioration in the Japanese economy between 1997 and 2000. Nevertheless, the endogenous technology adoption mechanisms presented in this paper could easily propagate this shock and generate the decline in output Japan experienced over the last 3 years of the Lost Decade,

There are several important questions that I have not addressed and which merit further research. First, and surely most important, what is the exact nature of the markup shocks that hit the Japanese economy during the early 90s?

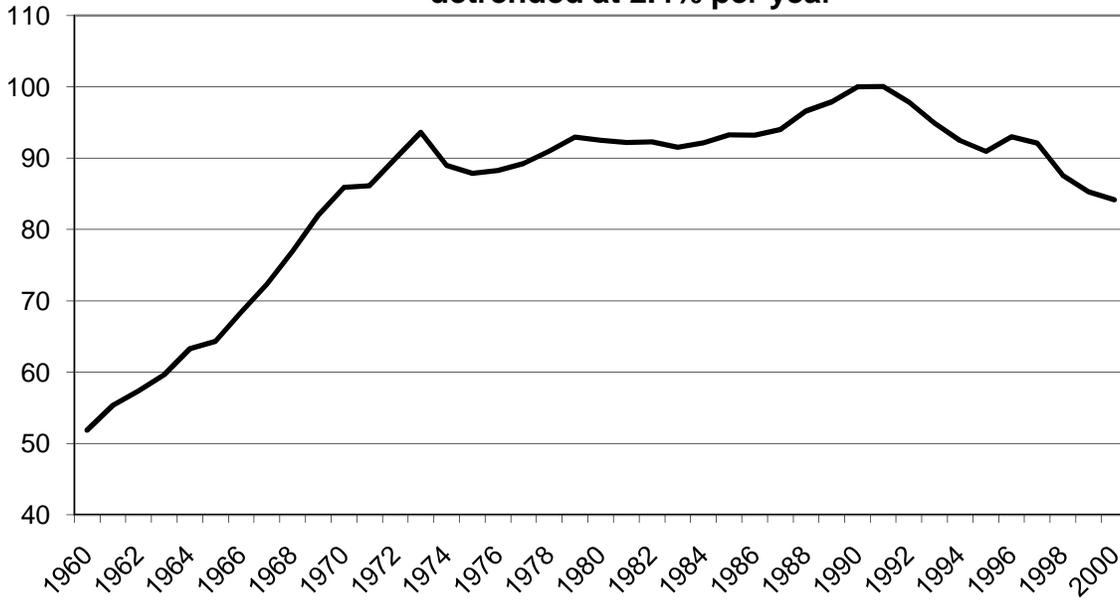
Hayashi and Prescott (2002) identify one important shock in 1987 that introduced rigidities in the Japanese labor market, the reduction of the maximum number of working hours allowed per person. However, more work is needed to explore whether there were other distortionary government interventions in labor markets afterwards. Further, financial frictions are isomorphic to shocks to the relative price of capital which, as shown by Comin, Gertler and Santacreu (2008) in a framework similar to the one used in this paper, have similar effects in macro variables to our wage markup shocks once endogenous technology adoption is introduced in the model.

References

- [1] Aghion, P. and P. Howitt (1992), “A Model of Growth through Creative Destruction”, *Econometrica* 60, 323-51.
- [2] Aghion, P. and G. Saint-Paul (1991), “Uncovering Some Causal Relationships between Productivity Growth and the Structure of Economic Fluctuations: A Tentative Survey”, NBER wp #4603.
- [3] Barlevy, G. (2003), “On the Timing of Innovation in Schumpeterian Models of Growth ” Northwestern mimeo.
- [4] Caballero, R., T. Hoshi and A. Kashyap (2008), “Zombie Lending and Depressed Restructuring in Japan” *The American Economic Review*, Vol. 98, No. 5, 1943-1977.
- [5] Cogley and Nason (1995) “Output Dynamics in Real-Business-Cycle Models,” *The American Economic Review*, Vol. 85, 3, pp. 492-511
- [6] Comin, D. and M. Gertler (2006), “Medium Term Business Cycles ” *The American Economic Review*, Vol. 96, No. 3, 523-551.
- [7] Comin, D., M. Gertler and A. M. Santancreu (2008), “Innovations in Growth Potential as Sources of Output and Asset Price Fluctuations" mimeo.
- [8] Gali, J. and F. Zilibotti, (1995), “Endogenous Growth and Poverty Traps in a Cournotian Model,” *Annales d'Economie et de Statistique*, Vol. 37/38, 197-214.
- [9] Gali, J., M. Gertler and D. Lopez-Salido (2007), “Markups, Gaps, and the Welfare Cost of Economic Fluctuations”, *Review of Economics and Statistics*, vol. 89 (February), pp. 44-59.
- [10] Hall, R. (1997) “Macroeconomic Fluctuations and the Allocation of Time,” *Journal of Labor Economics*, S223-S250.
- [11] Hayashi, F. and E. Prescott (2002), “Japan in the 90s: A Lost Decade," *Review of Economic Dynamics*.

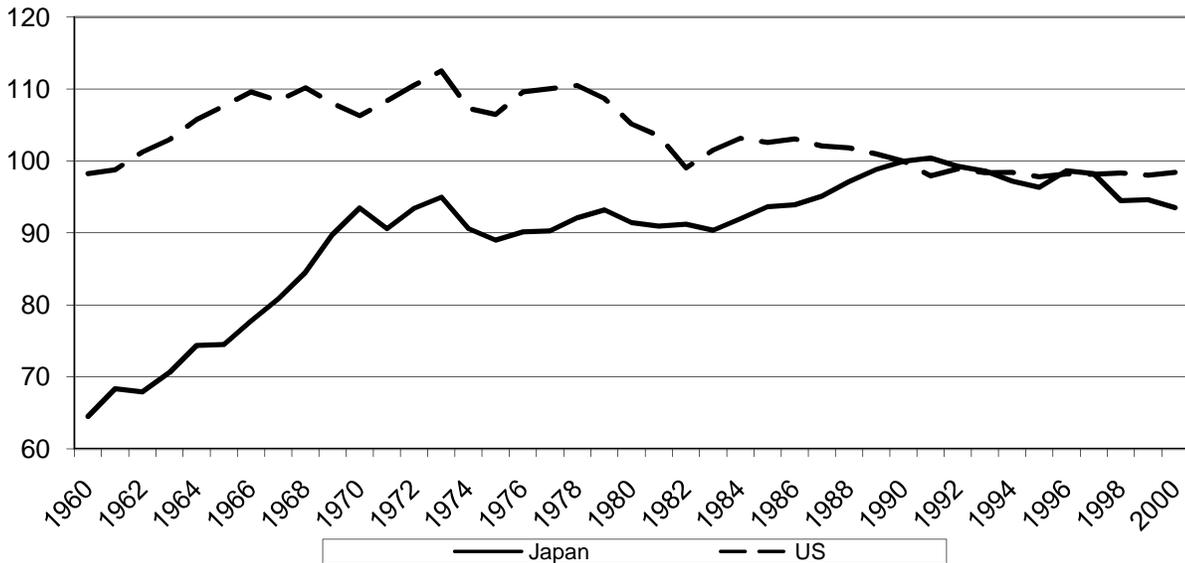
- [12] Hoshi, T. and A. Kashyap (2008), “Will the U.S. Bank Recapitalization Succeed? Lessons from Japan” NBER wp#14401.
- [13] Peek, J. and E. Rosengren (2005), “Unnatural Selection: Perverse Incentives and the Misallocation of Credit in Japan”, *The American Economic Review*, Vol. 95, 4, 1144-1166.
- [14] Romer, P. (1990), “Endogenous Technological Change” *The Journal of Political Economy*, Vol. 98, No. 5, Part 2. S71-S102.
- [15] Rotemberg, J. and M. Woodford (1995), “Dynamic General Equilibrium Models with Imperfectly Competitive Product Markets.” In *Frontiers of Business Cycle Analysis*, T. Cooley and E. Prescott, editors, Princeton University Press.

Figure 1: Real GNP per working age population in Japan detrended at 2.4% per year



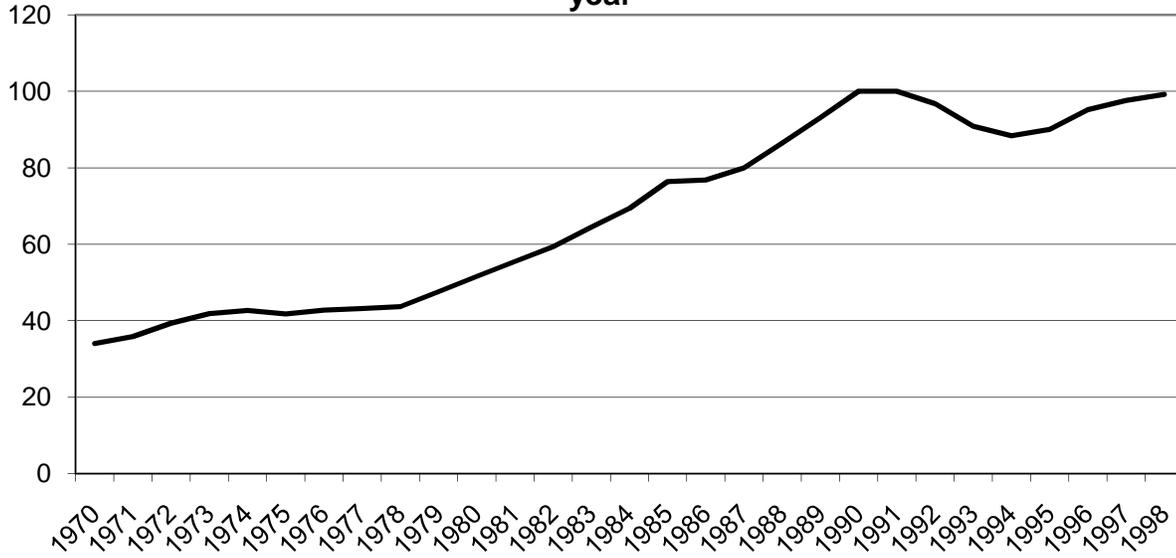
Source: Hayashi and Prescott (2002)

Figure 2: TFP detrended at 1.4% per year

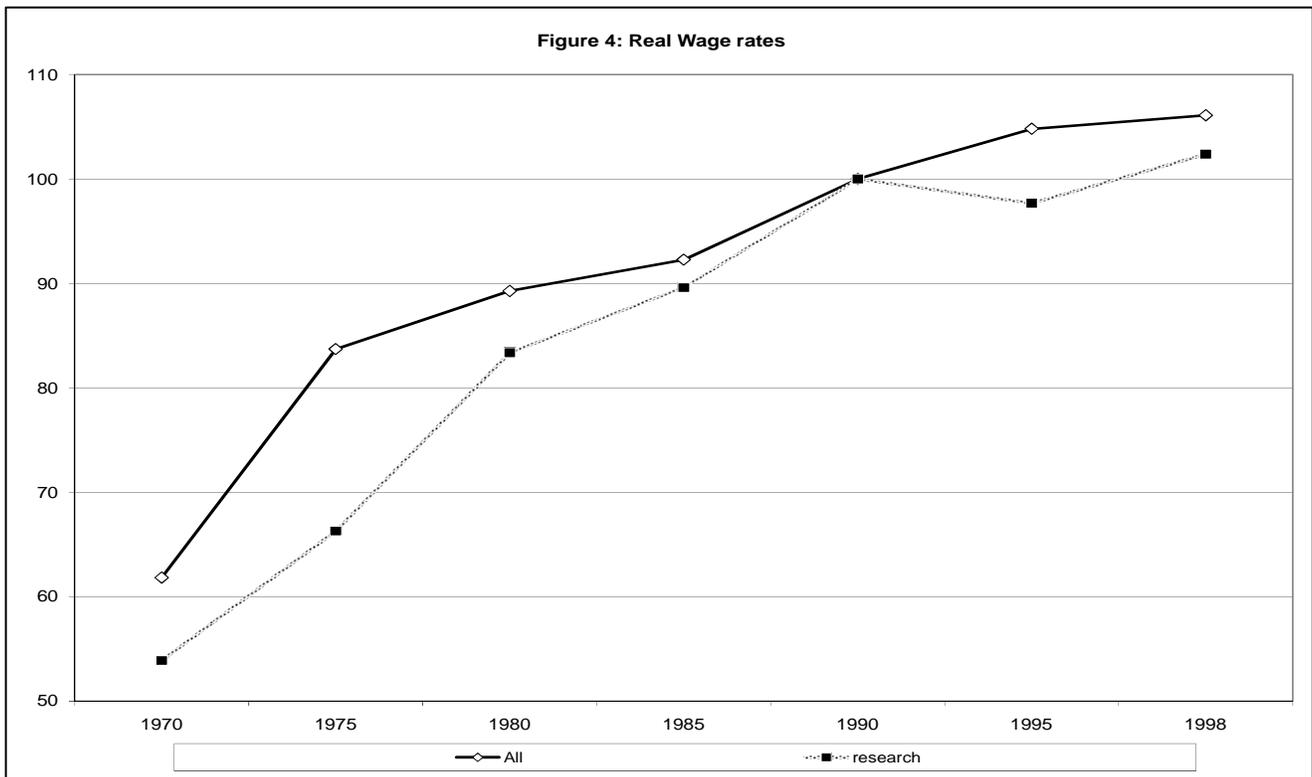


Source: Hayashi and Prescott (2002) and Bureau of economic analysis (BEA)

Figure 3: Real expenses in R&D detrended at 2.4 percent per year

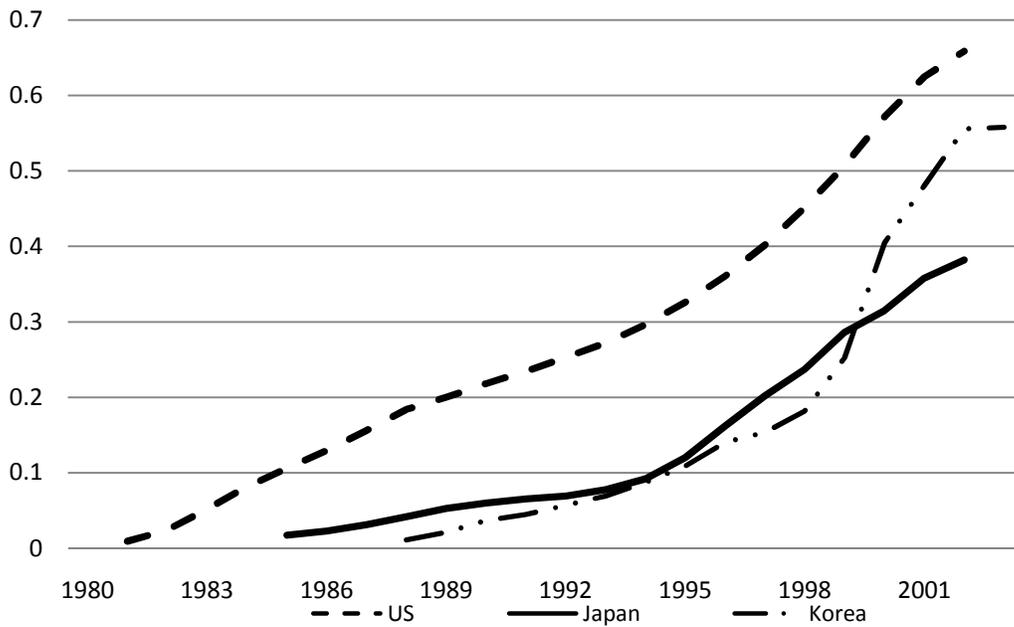


Source: ESRI



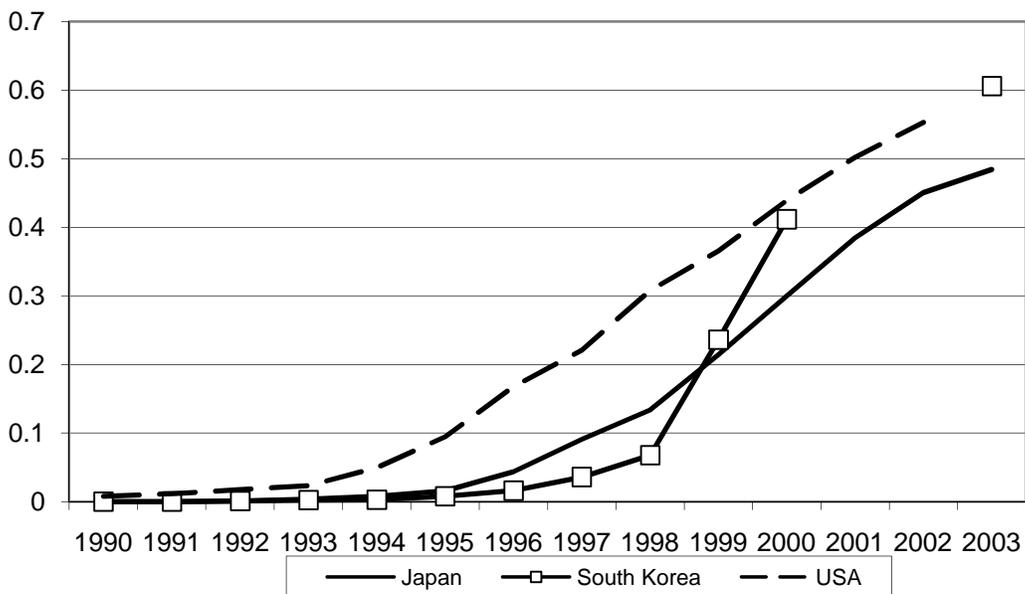
Source: ESRI

Figure 5: Computers per capita



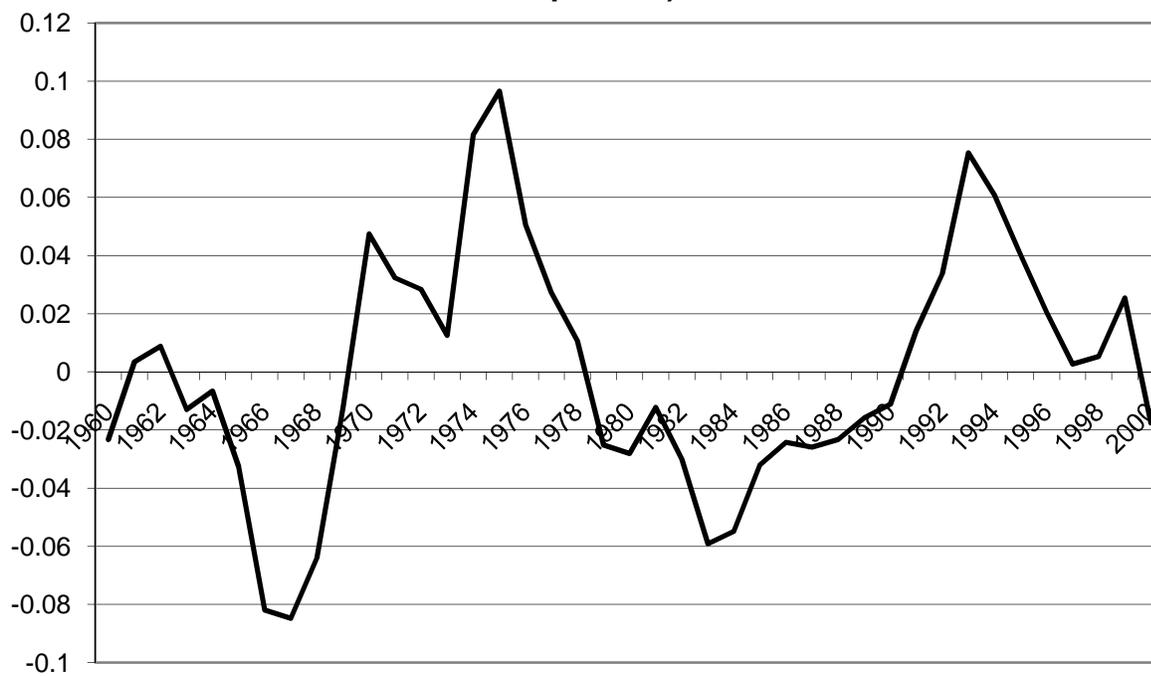
Source: Cross-country historical technology adoption (CHAT) dataset. Comin and Hobijn 2004.

Figure 6: Internet Users per capita



Source: Cross-country historical technology adoption (CHAT) dataset. Comin and Hobijn 2004.

Figure 7: Total Markup in Japan (band-pass filtered 2-100 quarters)



Source: Hayashi and Prescott (2002) and author calculations.

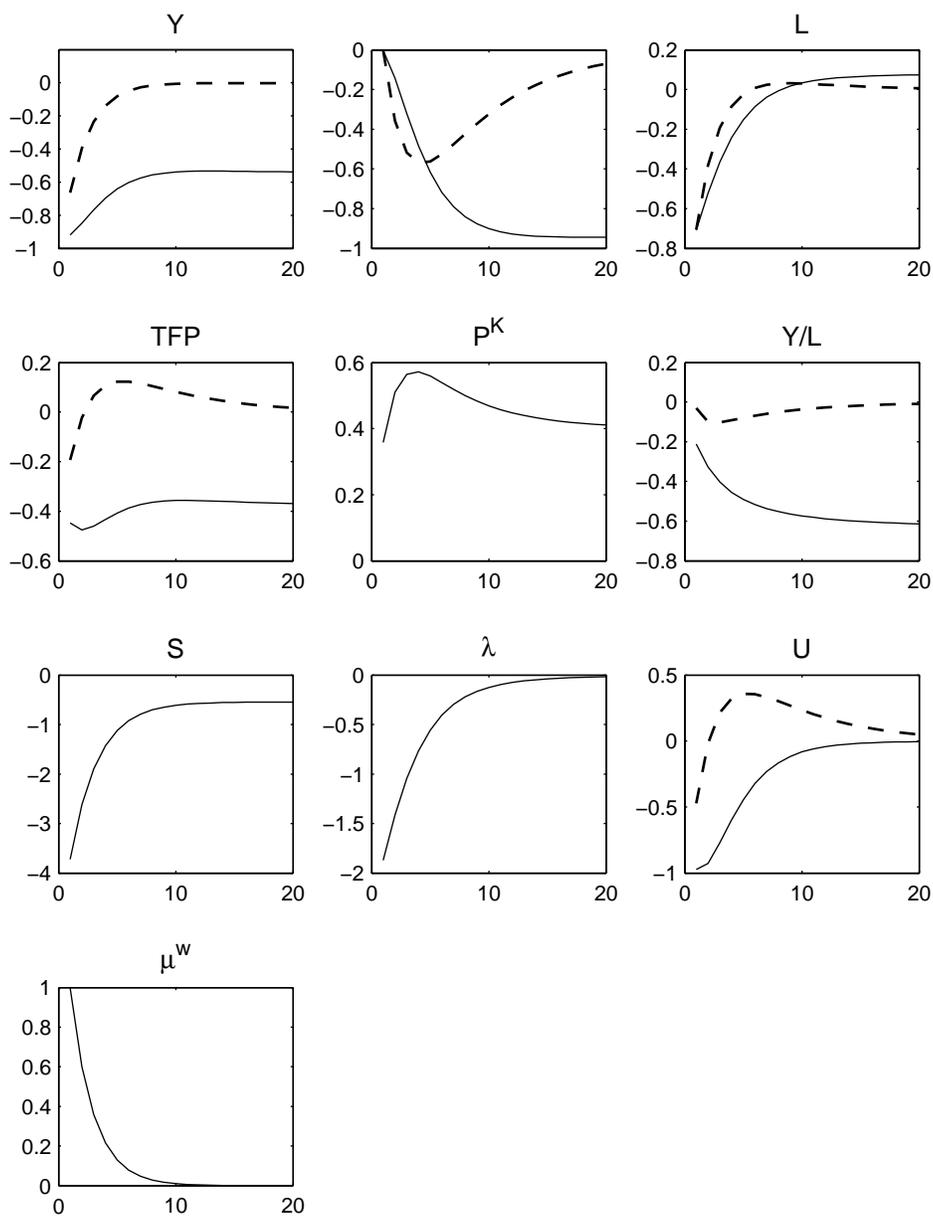


Figure 9: Output per Working Age population (1990 = 100)

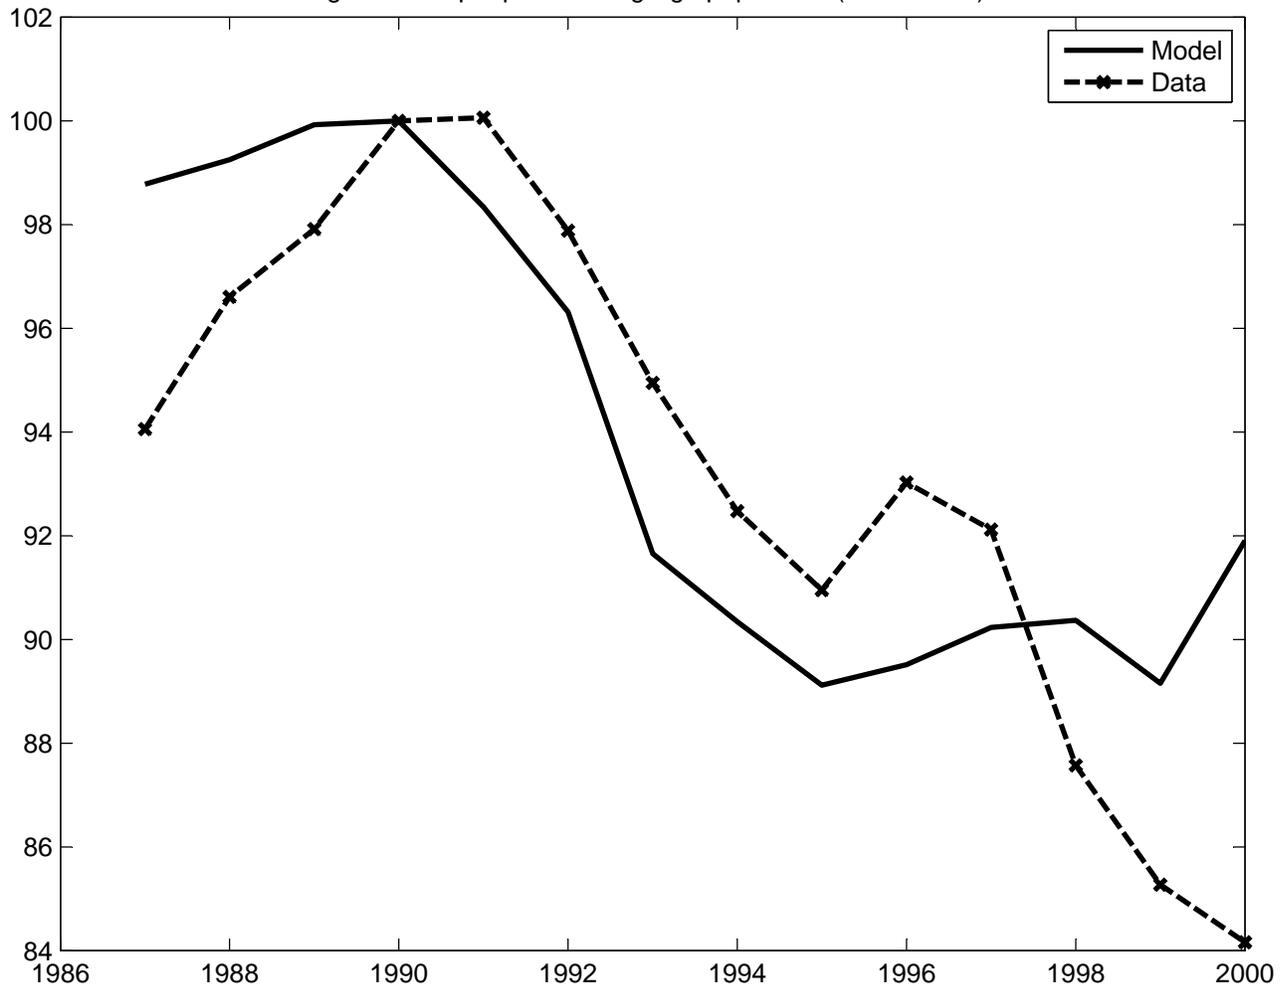


Figure 10: Total Factor Productivity (1990 = 100)

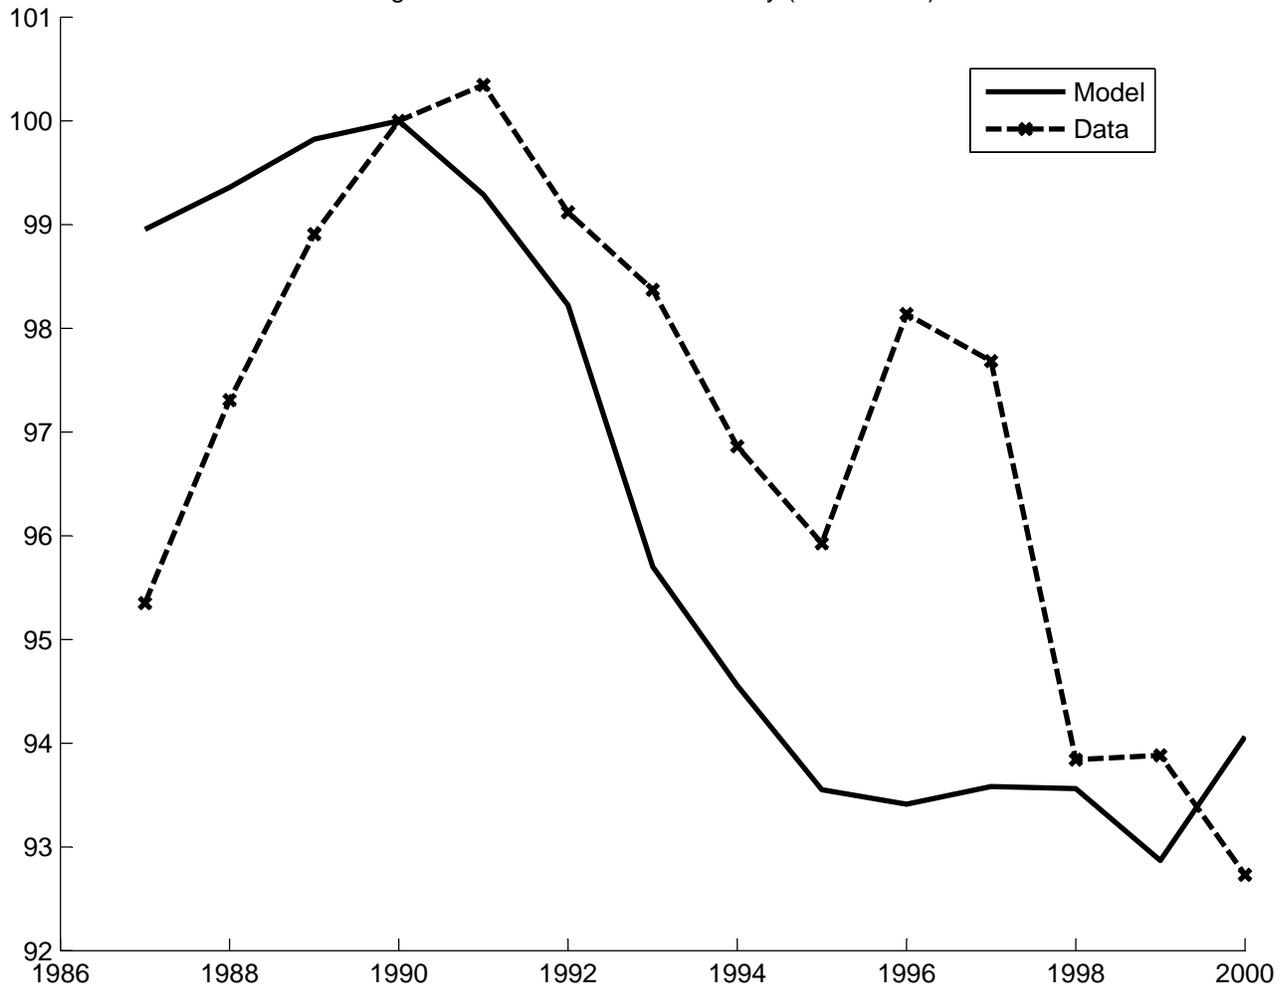


Figure 11: R&D (1990 = 100)



Figure 12: Hours Worked(1990 = 100)



Figure 13: Consumption (1990 = 100)

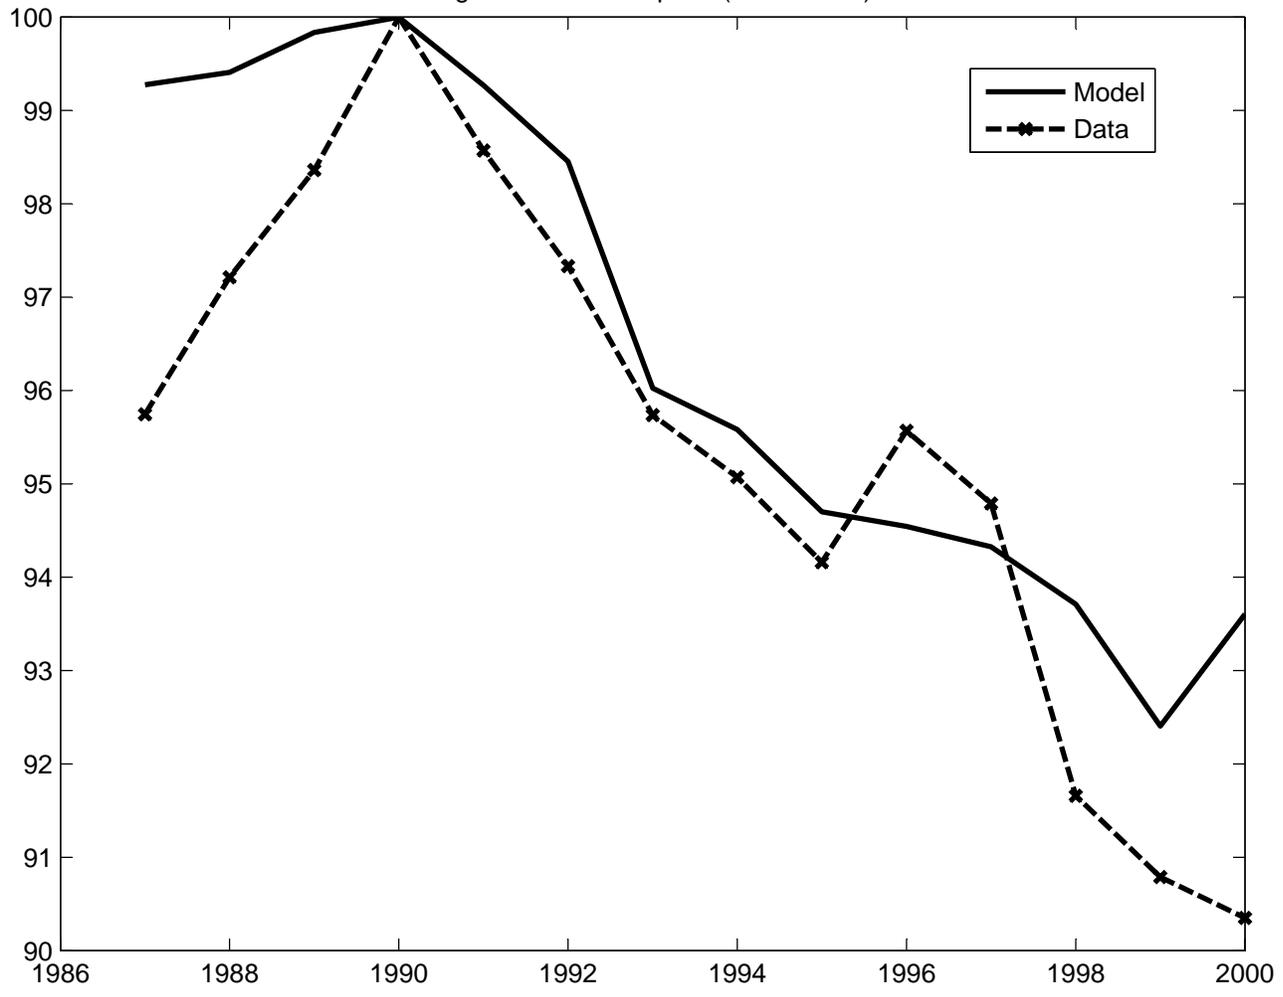
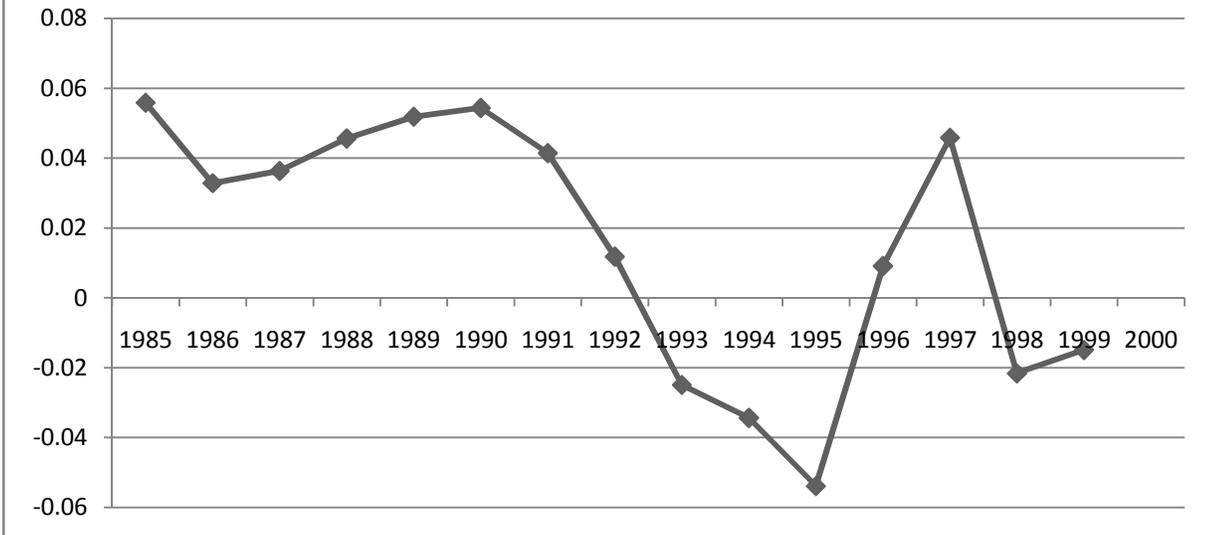


Figure 14: Investment (1990 = 100)



Figure 15: Capital markup



Note: Capital Markup is measured as the gap between the marginal product of capital and the marginal rate of substitution between consumption today and tomorrow filtered using a Band Pass filter that isolated cycles with periods smaller than 25 years.