IEF

Technological Revolutions and Debt Hangovers: Is There a Link?

by

Dan Cao

(Georgetown University)

Jean-Paul L'Huillier

(EIEF)

Technological Revolutions and Debt Hangovers: Is There a Link?

Dan Cao and Jean-Paul L'Huillier[*]

February 2013

Abstract

The Great Recession, the Great Depression, and the Japanese crisis of the 1990s were all preceded by periods of major technological innovation. This raises the question of whether such "technological revolutions" can subsequently cause great, private-debt-driven, recessions. We address this question through structural estimation. We find similar patterns in the three cases: The technological boom creates a wave of optimism in the household sector. This optimism is highly persistent and correlated with the observed accumulation of debt. The eventual deleveraging leads to a long consumption slump. Our model-produced measures of optimism line up with out-of-sample measures of consumer confidence. Overall, our findings point to a channel through which technological revolutions can lead to private-debt-driven recessions, and suggest that the Great Recession, the Great Depression, and the Japanese crisis were all caused by the preceding technological booms.

1

"Shifts in the economy are rarely forecast and often not fully recognized until they have been underway for some time."

Larry Summers, *Financial Times, March 25, 2012*

1   Introduction

A medium frequency look at the three most important private-debt recessions in developed economies reveals that they have all been preceded by periods of great technological innovation and economic transformation. Specifically, the recent Great Recession in the United States was preceded by a technological revolution, happening in the late 1990s, related to Information Technology (henceforth IT) (Atkeson and Kehoe 2007; Pastor and Veronesi 2009). Similarly, the Japanese debt crisis of the 1990s was preceded by a period of unprecedented industrial innovation in the 1980s. During this period, Japanese corporations developed and exported several electronic products that were massively consumed in many parts of the world, for instance the walkman, the VHS, the Betamax, and ATARI. We view this period as characteristic of a technological revolution concentrated in Japan. Finally, before the Great Depression, the United States witnessed the so-called 2nd Industrial Revolution, happening at the beginning of the 20th century. Two key general purpose technologies here were the combustion engine and electricity: The combustion engine made possible the mass production of cars for the American household by the Ford Motor Company, starting in 1909, and about 70% of household and corporate electrification happened approximately between 1910 and 1925.

Motivated by these facts, we set out to write a model in which technological revolutions lead to debt crises. Our goal is to understand how these facts can be rationalized in an optimizing framework. A key feature of our model is that anticipations about the future determine output, similar to the recent 'news' business cycles literature (Beaudry and Portier 2006; Christiano, Ilut, Motto, and Rostagno 2008; Jaimovich and Rebelo 2009). Importantly, these anticipations here are formed based on noisy information. Our main mechanism is the rational formation of beliefs about the future around a technological revolution.

The joint dynamics of the technological revolution and belief formation lead to a slow moving process that takes twenty to twenty-five years to be completed. Because our goal is to explain this process, we focus exclusively on medium frequencies. The process can be summarized by the following sequence of events. First, a technological

2

revolution creates a boom in aggregate productivity. This technological boom slowly creates a wave of optimism in the economy, which increases spending. Because of noisy information, spending tends to coincide with the end of the revolution. The end of the revolution implies a drop in the growth rates of productivity, together with a decrease in income. High spending combined with the decrease in income imply a wedge that creates an accumulation of debt. The increasing stock of debt leads to a debt crisis and a consumption slump due to a deleveraging process.

In our model, spending is simply characterized by consumption through a standard permanent income hypothesis combined with noisy signals about future income. In order to keep our framework as transparent as possible, we abstract away from investment and stock market valuation.

We use our model to interpret the data through structural estimation. The estimation of our model delivers three main results. First, there is a significant delay in the adjustment of anticipations about the future due a high amount of noise. This implies that spending dynamics considerably lag productivity dynamics, peaking about five years after the peak in productivity. Second, there is a large accumulation of debt that coincides with the peak in spending and the slowdown of productivity. The reason is the long delay in the adjustment of beliefs about the future, which in our model determine consumption and spending. Third, the deleveraging process after a technological revolution – necessary to return to steady state after the build-up of debt – is notoriously slow. The reason is that the end of the technological revolution produces a decline in the growth rates of productivity. Because in our economy productivity creates income for the representative agent, this agent finds itself trapped in an adverse scenario for two reasons: high debt, and low income. Therefore, the ensuing deleveraging is slow. In fact, a simulation of our estimated model suggests that it takes 15 to 20 years for consumption to get back to its steady state level. We name this situation a debt "hangover", by analogy to a post-party "recovery". According to this analogy, U.S. consumers would have "partied" between, say 1998 and 2005, and would then need to "recover", by bringing their debt/output ratio back to steady state.

2 The Model

2.1 Productivity process and Information Structure

We model an open economy similar to Aguiar and Gopinath (2007), adding a "news and noise" information structure (Blanchard, L'Huillier, and Lorenzoni 2012, henceforth BLL).[1] Specifically, productivity a_t (in logs) is the sum of two components, permanent, x_t, and transitory z_t:

$$a_t = x_t + z_t \quad .$$

Consumers do not observe these components separately. The permanent component follows the unit root process

$$\Delta x_t = \rho \Delta x_{t-1} + \varepsilon_t \quad . \tag{1}$$

The transitory component follows the stationary process

$$z_t = \rho z_{t-1} + \eta_t \quad . \tag{2}$$

The coefficient ρ is in $[0, 1)$, and ε_t and η_t are i.i.d. normal shocks with variances σ_ε^2 and σ_η^2. Similar to BLL, we assume that these variances satisfy

$$\rho \sigma_\varepsilon^2 = (1 - \rho)^2 \, \sigma_\eta^2 \quad , \tag{3}$$

which implies that the univariate process for a_t is a random walk, that is

$$E[a_{t+1} | a_t, a_{t-1}, ...] = a_t \quad .$$

This assumption is analytically convenient and, as will be seen below, also broadly in line with actual productivity data. To see why this property holds, note first that the implication is immediate when $\rho = \sigma_\eta = 0$. Consider next the case in which ρ is positive and both variances are positive. An agent who observes a productivity increase at time t can attribute it to an ε shock and forecast future productivity growth or to an η shock and forecast mean reversion. When (3) is satisfied, these two considerations exactly balance out and expected future productivity is equal to

---

[1] Boz, Daude, and Durdu (2011) use a similar framework. We simplify it further by removing capital and investment, which allows us to derive a tight connection to the New Keynesian model (see Subsection 2.3).

current productivity.[2]

Consumers have access to an additional source of information, as they observe a noisy signal about the permanent component of productivity. The signal is given by

$$s_t = x_t + \nu_t \quad , \tag{4}$$

where ν_t is i.i.d. normal with variance σ_ν^2.

We think of ε_t as the "news" shock, because it builds up gradually, has permanent effects on productivity and delivers (noisy) information about the future through the signal. We think of ν_t as the "noise" shock.

2.1.1 Slow Adjustment of Beliefs and Technological Revolutions

Here we focus on an important property of the signal extraction problem for our purposes.

First, we borrow the idea from Greenwood and Jovanovic (1999) (among others) that "technological revolutions come in waves". According to this idea, the start of a technological revolution should create an increase in the growth rate of permanent productivity – away from the old, deterministic trend – and the end of a technological revolution should create a decrease in the growth rate of permanent productivity – away from the new trend.

Turning to the implications of these movements of the permanent component of productivity for beliefs, notice that agents, trying to forecast the future path of x_t, will form beliefs about the level of x at infinity

$$x_{t+\infty} \equiv \frac{x_t - \rho x_{t-1}}{1 - \rho} \tag{5}$$

using a Kalman filter. As will become clear, in the model these beliefs will largely determine beliefs about long-run income, and therefore consumption.[3] Denoting as $x_{\tau|t} \equiv E[x_\tau]$ the conditional expectation of x_τ on information available at time t, we have

$$x_{t+\infty|t} = \frac{x_{t|t} - \rho x_{t-1|t}}{1 - \rho} \quad . \tag{6}$$

Because of noisy information, agents will be slow to adjust their beliefs $x_{t+\infty|t}$.

---

[2]See BLL for the proof.

[3]Movements in the real interest rate will also affect consumption, but they turn out to be of little importance in the calibrated model.
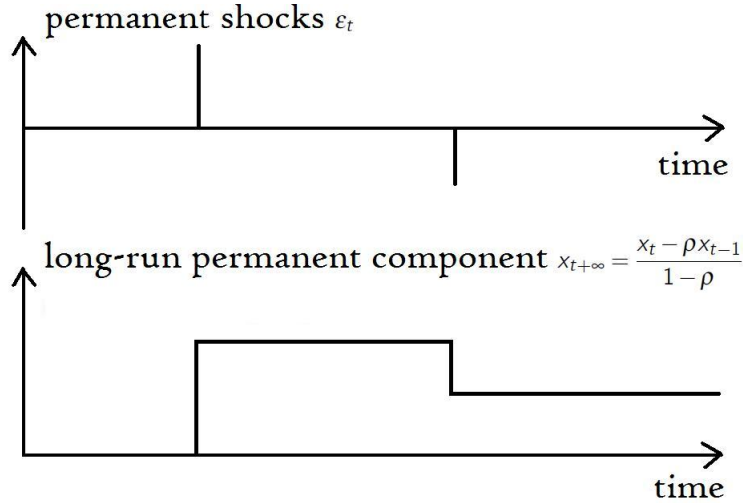
Figure 1: An Example of a Technological Revolution

In particular, they will be slow to adjust their beliefs after the decrease of Δx_t and will remain "optimistic" for a while.

For illustrative purposes, consider the example of a technological revolution given by Figure 1. The upper panel of the figure plots off-trend permanent shocks, and the lower panel plots the implied long run levels of the permanent component x_∞. Notice that the second off-trend shock appears here as negative, however, this shock is in fact positive – but smaller than the first one – once the deterministic trend is added back. The technological revolution initially increases the long run level of the permanent component, and then decreases it in off-trend terms.[4] As we will show below, our structural estimations, which allow us to estimate permanent shocks, will give support to this view of technological revolutions for the three cases we consider.

Figure 2 sketches the evolution of beliefs around this technological revolution, represented by the dotted line in the lower panel. Agents slowly learn about the increase in the long run level of the component, and about the subsequent decrease. Overall, the adjustment of beliefs lags the actual changes in the permanent component, which is the key property needed for our results. Notice that the persistence of the process for Δx_t given by ρ implies a persistence of beliefs which is compounded to the lag implied only by noisy information.[5]

---

[4]In this example, the technological revolution has a total positive effect on the off-trend level of the long run permanent component.

[5]One may ask why agents in this economy do not anticipate the second (negative) shock when learning about the first one. We think about these technological revolutions as happening rarely, for instance
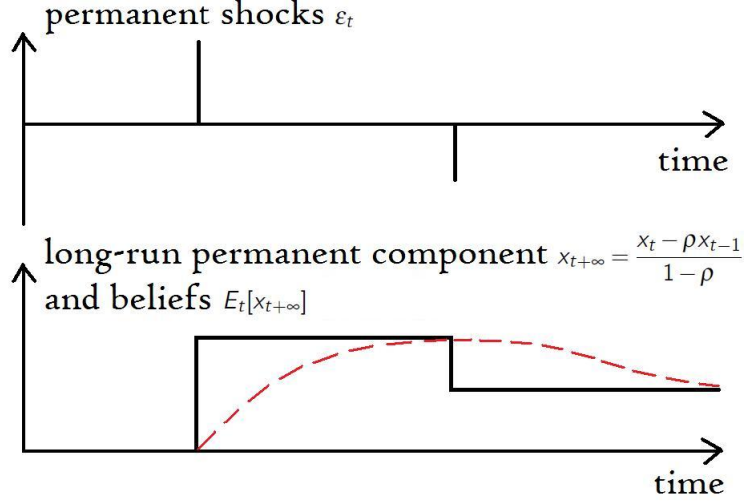
Figure 2: Beliefs Around a Technological Revolution

2.2 Consumption, Production and Net Exports

We now describe the rest of the model. A representative consumer maximizes

$$E\left[\sum_{t=0}^{\infty}\beta^t\left(\ln\left(C_t\right)-\frac{\varphi}{1+\phi}N_t^{1+\phi}\right)\right]$$

where $E[\ \cdot\ ]$ is the expectation operator conditional on information available contemporaneously. The maximization is subject to

$$C_t+B_{t-1}=Y_t+Q_tB_t \quad , \tag{7}$$

where B_t is the external debt of the country, Q_t is the price of this debt, and Y_t is the output of the country. Output is produced using only labor through the linear production function:

$$Y_t = A_tN_t. \tag{8}$$

In order to obtain the stationarity of the loglinearized model around the steady state, we assume that the interest rate that the representative consumer faces is increasing in the total amount of external debt (Schmitt-Grohe and Uribe 2003):

$$\frac{1}{Q_t} = R_t = R^* + \psi\left\{e^{\frac{B_t}{Y_t}-b}-1\right\} \quad , \tag{9}$$

---

once every century. Accordingly, in our simple specification of the evolution of technology agents are "surprised" by the sharp slowdown of aggregate productivity implied by the second shock.

normalized debt.[6]

The two first-order conditions from the optimization problem of the consumer are:

1. Intertemporal

$$\frac{1}{C_t} = \beta R_t E \left[\frac{1}{C_{t+1}} \right] \quad , \tag{10}$$

2. Intratemporal

$$A_t = \frac{\varphi N_t^\phi}{\frac{1}{C_t}} \quad . \tag{11}$$

2.2.1 Normalization and Loglinearization

Steady State. We look for a steady state in which the following variables (normalized and non-normalized) are constant: $c = C/A$, $y = Y/A$, N, $b = B/Y$, R, and Q. We assume that the steady state level of normalized debt b is determined exogenously.[7]

From the intertemporal condition (10), we have

$$\frac{1}{C} = \beta R \frac{1}{C_+} \quad ,$$

where the subscript $+$ is used to denote value one period ahead. Equivalently, then

$$\frac{1}{C/A} = \beta R \frac{A}{A_+} \frac{1}{C_+/A_+} \quad .$$

Given that $c = C/A = C_+/A_+$ in the steady state, it implies that

$$1 = \beta R \frac{A}{A_+} \quad , \tag{12}$$

or

$$R = \frac{1}{\beta} \quad .$$

---

[6]The fact that the international interest rate does not adjust to developments in the domestic economy makes this essentially a small open economy model. In Subsection 2.3 we consider a two-country open economy model where the international rate is endogenous. At the limit when $\psi \to 0$ both models are equivalent regarding the behavior of consumption, and therefore for parametrizations typically used in this literature (fairly small ψ, see Aguiar and Gopinath 2007) both models deliver similar results. See Proposition 1.

[7]Aguiar and Gopinath (2007) make the same assumption.

The resource constraint (7) gives

$$C + B = Y + \frac{1}{R} B_+ \quad,$$

or

$$\frac{C}{A} + \frac{B}{A} = \frac{Y}{A} + \frac{1}{R} \frac{B_+}{A_+} \frac{A_+}{A} \quad.$$

The intra-temporal first-order condition (11) implies

$$\frac{1}{C/A} = \varphi N^\phi \quad. \tag{13}$$

Moreover, using and (12) and (8), we can simplify the budget constraint in the steady state further to

$$\varphi N^\phi + b = N + \beta b \quad.$$

From this equation, given b, we can solve for the steady state level of employment, N. Given N, we can solve for c using (13) and then y.

Loglinearization. Define

$$c_t \equiv \log C_t - \log A_t - \log (C/A) \quad,$$

$$n_t = \log N_t - \log N \quad,$$

$$r_t \equiv \log R_t \quad,$$

and

$$b_t = \frac{B_t}{Y_t} - \frac{B}{Y} \quad,$$

$$nx_t = \frac{NX_t}{Y_t} - \frac{NX}{Y} \quad.$$

To loglinearize the intertemporal condition

$$\frac{1}{C_t} = \beta R_t E \left[\frac{1}{C_{t+1}} \right] \quad,$$

we proceed as follows:

$$\frac{1}{C_t}\frac{A_t}{A_t} = \beta R_t E\left[\frac{1}{C_{t+1}}\frac{A_{t+1}}{A_{t+1}}\right] \quad,$$

to obtain

$$c_t = -r_t + E[c_{t+1} + \Delta a_{t+1}] \quad.$$

The loglinearization of the intratemporal condition is immediate

$$c_t + \phi n_t = 0 \quad.$$

Similarly for the production function:

$$y_t = n_t \quad,$$

and the interest rate equation

$$r_t = \psi \cdot b_t \quad.$$

Approximating the resource constraint delivers

$$\left[\frac{C}{A}(c_t + 1)\right]\left[\frac{A}{Y}(-y_t + 1)\right] + \frac{NX}{Y} + nx_t = 1 \quad,$$

which leads to

$$c_t + \frac{1}{C/Y}nx_t = y_t \quad.$$

The current account surplus is

$$NX_t = B_{t-1} - Q_t B_t \quad,$$

and therefore, approximating

$$\frac{NX}{Y} + nx_t = \left[\frac{B}{Y} + b_{t-1}\right]\left[\frac{Y}{A}(y_{t-1} + 1)\right]\left[\frac{A}{Y}(-y_t + 1)\right][-\Delta a_t + 1] - \left[\frac{1}{R}(-r_t + 1)\right]\left[\frac{B}{Y} + b_t\right]$$

to obtain finally

$$nx_t = b_{t-1} - \beta b_t + \frac{1 - C/Y}{1 - \beta}\left(y_{t-1} - y_t - \Delta a_t + \beta r_t\right) \quad.$$

2.2.2   Summary of Loglinearized Model and Simulation

The loglinearized model has 6 endogenous variables: c_t, y_t, n_t, s_t, r_t, and b_t. The model is given by the equations for the shock processes and 6 equations:

$$c_t = -r_t + E[c_{t+1} + \Delta a_{t+1}] \quad , \tag{14}$$

$$c_t + \phi n_t = 0 \quad , \tag{15}$$

$$y_t = n_t \quad , \tag{16}$$

$$r_t = \psi \cdot b_t \quad , \tag{17}$$

$$c_t + \frac{1}{C/Y} nx_t = y_t \quad , \tag{18}$$

$$nx_t = b_{t-1} - \beta b_t + \frac{1 - C/Y}{1 - \beta}(y_{t-1} - y_t - \Delta a_t + \beta r_t) \quad , \tag{19}$$

Figure 3 shows a simulation of the model for given parameter values. The figure shows the responses of consumption c_t, net exports nx_t (or equivalently in this model, the current account), and external debt b_t, to a one-standard deviation increase in ε_t (the permanent technology or "news" shock). The time unit on the x-axis is four quarters (one year). The scale of productivity is relative percentage deviations from steady state plus 1, times 100. The scale of net exports is absolute percentage deviation from the steady state value of net exports-to-output, NX/Y. The scale of the debt-to-output ratio is absolute percentage deviation from the steady state value of debt-to-output, b. The parameter values are those obtained when estimating the model using data for the United States (1990–2010). The discount factor β is set at 0.99. The inverse of the Frisch elasticity of labor supply is set at 2. The value of ψ is taken from Aguiar and Gopinath (2007) (0.0010).[8] The steady state value of the consumption-to-output ratio C/Y is estimated at 0.9979. The parameter ρ is estimated at 0.98, implying slowly building permanent shocks and slowly decaying transitory shocks. The standard deviation of productivity growth, σ_u, is estimated at 0.63. These values for ρ and σ_u yield standard deviations of the two technology

---

[8]Schmitt-Grohe and Uribe (2003) use a similar value.

shocks, σ_ε and σ_η, equal to 0.01% and 0.62%, respectively. The standard deviation of the noise shock, σ_ν, is set to 10.80%, implying a fairly noisy signal.[9]
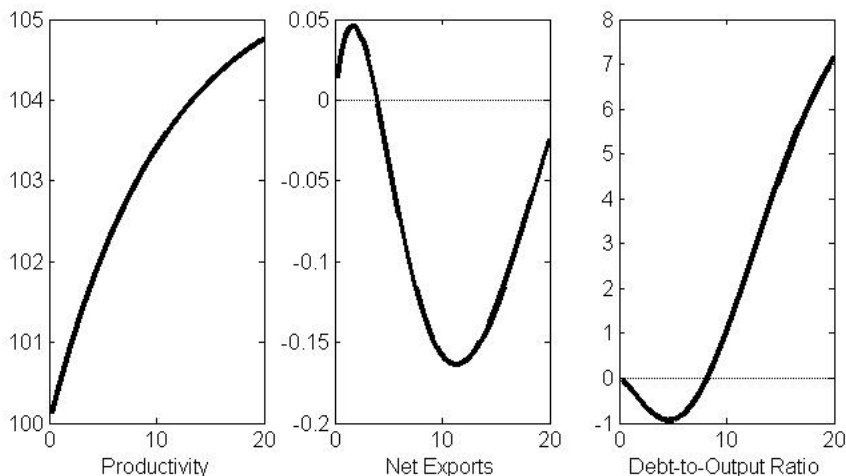

Figure 3: Impulse Response Functions After Permanent Technology Shock.

In response to a one-standard-deviation increase in ε_t, the permanent technology shock, productivity increases slightly on impact, and then gradually continues to increase until it reaches a new long-run level. This sustained increase is slow; in fact, half of the productivity increases are reached only after 7 years. Initially, net exports rise, mainly because productivity increases faster than beliefs about long run productivity. This is a reflection of the high amount of noise in this simulation. After 5 years net exports fall, because agents have received enough "news" and a standard income effect kicks in. This is translated into a sharp accumulation of external debt. In the long run, productivity reaches a new level (at 1.063) and net exports and the debt-to-output ratio go back to zero.

2.3 Some Useful Convergence Results

Here we establish that the loglinearized model presented here (a version of Aguiar and Gopinath 2007 without capital) has arbitrarily similar consumption dynamics to a closed-economy New Keynesian model with a high degree of price stickiness.[10]

---

[9]It is important to underline that these values are of a similar order of magnitude as those obtained in BLL, and also to those estimated using data for either Japan or the Great Depression. The reasons are presented in Subsections 2.3 and 3.

[10]BLL use such a model to estimate the degree of imperfect information in the U.S. economy (1948-2011), and extend the framework to a model with estimated price stickiness that incorporates other

Proposition 1 (Convergence to a Closed Economy New Keynesian Model)
If $C/Y = 1$, the consumption dynamics in the economy represented by equations (14) to (19) approaches the consumption dynamics of a closed economy New Keynesian model with highly sticky prices as ψ tends to 0 and β tends to 1.

The proof is in the appendix. The intuition behind this result is that when ϕ is small, the interest rate is not very sensitive to deviations of consumption from the steady state. Similarly, in a New Keynesian economy with high degree of price stickiness, inflation is not very sensitive to deviations of consumption from the steady state, and by implication the nominal interest. Therefore, at the limit, both models have similar consumption dynamics. Notice of course that this result is independent of the degree of noise σ_s in the model.[11]

This limiting result is useful in practice because, for parameters usually considered in the literature (Schmitt-Grohe and Uribe 2003; Aguiar and Gopinath 2007), these limiting conditions are numerically satisfied. Indeed, for a discount factor $\beta = 0.99$ and a reaction of the interest rate r_t to deviations from the steady state of the debt-to-output ratio $\phi = 0.0010$, the dynamics of our model and the dynamics of a New Keynesian model with highly sticky prices (BLL) are similar.

An immediate corollary of this main result – important to justify the use of this model for the U.S. economy and Japan – is that our model, in which the world interest rate R^* is fixed, has also arbitrarily close dynamics as a two-country open economy model. For parameters usually considered in the literature (mainly $\psi = 0.0010$), then, our model is a good approximation of a two-country model in which R^* is endogenous. We develop this theme further in the section presenting the results of structural estimation.

3   Empirical Results

In this section we report empirical results for the three episodes considered. Our goal is to look at the data through the lens of our model, and compare the lessons from each episode. For instance, the estimated model can be used, in each of these cases, to back out estimated series for the long-run permanent component of productivity, and series of corresponding contemporaneous consumer beliefs. These series can be

---

frictions frequently used in the literature (Christiano, Eichenbaum, and Evans 2005; Smets and Wouters 2007).

[11]The reason β needs to go to 1 is just an artifact of the loglinearization and has no economic content. See the proof for the details.

then used to test our hypothesis on two dimensions. First, the estimated permanent component of productivity should feature an increase, together with a subsequent decrease. Second, consumers' beliefs should lag these movements, adapting slowly on the upside, peaking *after* the permanent component, and then reverting back.

3.1   Structural Estimation

For the Great Recession, we use the open economy model presented in the previous section and fix a subset of the parameters to conservative values that we draw from the literature. These parameters are shown in Table 1. An important parameter for our purposes is ψ, which we fix to 0.0010, as Schmitt-Grohe and Uribe (2003) and Aguiar and Gopinath (2007). The remainder set is composed by the steady state debt-to-output ratio B/Y, which we estimate at 0.2143 using flow of funds data, implying a value of the consumption-to-output ratio $C/Y = 0.9979$, and the informational parameters (the persistence of the technology processes ρ, the standard deviation of productivity innovations σ_u, and the standard deviation of noise shocks σ_s). These informational parameters are estimated using maximum likelihood. From BLL we know that, in this type of model, identification of these parameters is usually obtained – in the sense of finding a unique maximum for the likelihood function – when the other parameters of the model are fixed and two time series are used, as for instance a series for productivity and a series for consumption. Therefore, we limit ourselves to this approach without recurring to Bayesian statistics.

Parameter	Description	Value
ϕ	Inv. Frisch Elasticity	2
β	Discount Factor	0.99
ψ	Sensitivity Interest Rate	0.0010

Table 1: Fixed Parameters.

In the case of Japan, we feel less comfortable with a model in which the domestic economy finances consumption through external liabilities. Moreover, for both Japan and the Great Depression, and especially for the Great Depression, we do not have a good sense of how the consumption boom was financed. Therefore, we limit ourselves to a demand determined economy without explicit borrowing and lending, in which anticipations about the future determine output (the simple New Keynesian economy of BLL). Notice however that, having established Proposition 1, this economy should behave similarly to the open economy model for the values of Table 1. We have

verified this claim in a number of ways, in particular by checking that all of the results mentioned here would be similar if instead we had estimated the open economy model for both Japan and the Great Depression.

Data and Observationally Equivalent Full Information Model. In the case of the U.S. 1990–2010, the time series used are (the logarithm of) labor productivity and per capita real net exports. We use quarterly data. In the case of Japan, the time series used are (the logarithm of) labor productivity and national accounts per capita real consumption.[12] We use quarterly data.

In the case of the U.S. 1919-1933 we do not have data for net exports, instead we use private consumption from the Gordon-Krenn data set.[13] Gordon and Krenn (2010) used the Chow-Lin method (Chow and Lin 1971) for interpolating annual national accounts series and obtain cyclical variation at quarterly frequency, thereby obtaining an estimated series for GDP components. In order to obtain a series for labor productivity, we obtained an estimate for GDP from the Gordon-Krenn data set, and we used the Kendrick data set for employment, using a linear interpolation out of this annual series. The labor productivity, net exports and consumption series are detrended prior to the estimation, by separately removing the average growth rates from each of them.

Following BLL, we exploit the existence of an observationally equivalent full information model to the model with noisy information. This immensely facilitates the implementation of our estimation by allowing us to use standard computational tools. Our loglinearized model can be represented in state space form and estimated through Maximum Likelihood. Our environment is parsimonious and therefore there is no need to recur to Bayesian methods. In fact, we know that in all estimations we hit a unique global maximum for the likelihood function. This is of course an advantage because it makes the exercise transparent, and therefore it is fairly easy to get a sense of how our model fits the data.

Results. Table 2 contains the parameter estimates. The persistence parameter ρ was estimated at 0.98 both in the case of the Great Recession and Japan, and at 0.95 in the case of the Great Depression, implying very persistent processes both for the permanent and the transitory components of productivity. The standard deviation of productivity was estimated at 0.69% in the case of the Great Recession, at 0.98%

---

[12]For more details, see Appendix A.1.

[13]In this case our sample length choice is restricted by the fact that we do not have quarterly private consumption data before the end of World War I in 1918.

in the case of Japan, and at 1.70% in the case of the Great Depression. Given the random walk assumption for productivity, the high values of ρ imply a standard deviation for permanent technology shocks that is fairly small, of 0.01%, 0.02%, and 0.09% respectively, and a fairly big standard deviation for the transitory technology shock, of 0.66%, 0.97%, and 1.66% respectively. The standard deviation of noise shocks is large, 10.80%, 12.08%, and 10.70%, implying fairly noisy signals.[14] The fact that permanent changes in productivity tend to be small compared to the size of transitory shocks, and that the amount of noise in the signal is large, suggests that learning is slow, because both a_t and s_t are fairly uninformative signals about x_t. This illustrates the major signal extraction problem that consumers face according to our estimation.

Parameter	Description	Great Rec.	Japan	Great Depr.
ρ	Persistence tech. shocks	0.98	0.98	0.95
σ_u	Std. dev. productivity	0.63	0.98	1.70
σ_v	Std. dev. permanent tech. shock	0.01	0.02	0.09
σ_z	Std. dev. transitory tech. shock	0.62	0.97	1.66
σ_s	Std. dev. noise	10.80	12.08	10.70

Table 2: Parameter Estimates.

The simplicity of our model and Proposition 1 provides a fairly clear intuition of the identification of these parameters. Indeed, the limiting model admits a closed form VAR representation, together with some equations for other moments needed to identify the parameters. For the details see BLL. In a nutshell, the persistence ρ is identified by the covariance between the productivity-to-consumption ratio and future changes in productivity. The intuition is that according to our model agents should be able to forecast these changes, and these forecasts are reflected in consumption. On average, the lower the productivity-to-consumption ratio, the larger the subsequent increases of productivity, given that high consumption should on average be associated with an increasing path of the permanent component of productivity. For a given variance of the ratio, the higher the covariance, the larger are permanent shocks with respect to transitory shocks, and the lower ρ.[15] Having pinned down ρ, the noise in the signal can be identified by the variance of changes

---

[14]Notice that the signal is about the level of the permanent component and therefore it is misleading to directly compare the size of the noise shocks to the size of permanent shocks.

[15]It is possible to show that on average this relationship is not affected by the degree of noise, given that the analytic expression takes into account the variance of this ratio.

in consumption with respect to the variance of changes in productivity. For a given ρ, the smaller the variance of the signal σ_s^2, the more consumption changes due to permanent shocks because of the income effect, and the more volatile is consumption. Since in our estimation ρ is estimated as large, the income effect is large and we need a high degree of noise to limit the reaction of consumption and match the data.

The state-space representation of the estimated model can be used in order estimate the shocks and states of the model using a Kalman smoother[16]. Among the states, the most interesting are the evolution of the state variable x_t (permanent component of productivity), and corresponding contemporaneous consumers' beliefs $x_{t|t} = E[x_t]$.

Figure 4 plots our estimated permanent technology shocks for the case of the Great Recession. Consistent with our idea of the effects of a technological revolution – spelled out in the previous section – we estimate a series of positive shocks in the early 1990s, up to 1997, and a subsequent series of (off-trend) negative shocks.[17] Figure 5 plots our estimated transitory shocks (upper panel) and our estimated noise shocks (lower panel). These series of shocks do not show any particular pattern.

---

[16]A Kalman smoother is a Bayesian filter that uses all the sample (as opposed to the Kalman filter, which uses only contemporaneous and past observables) to compute ex-post estimates of the states of the system (see Hamilton 1994, pp. 394–397).

[17]The serial correlation of our estimated permanent shocks might seem a violation of the white noise assumption about these shocks, however, it is just a reflection of the invertibility problem present in this model and discussed in detail in BLL.
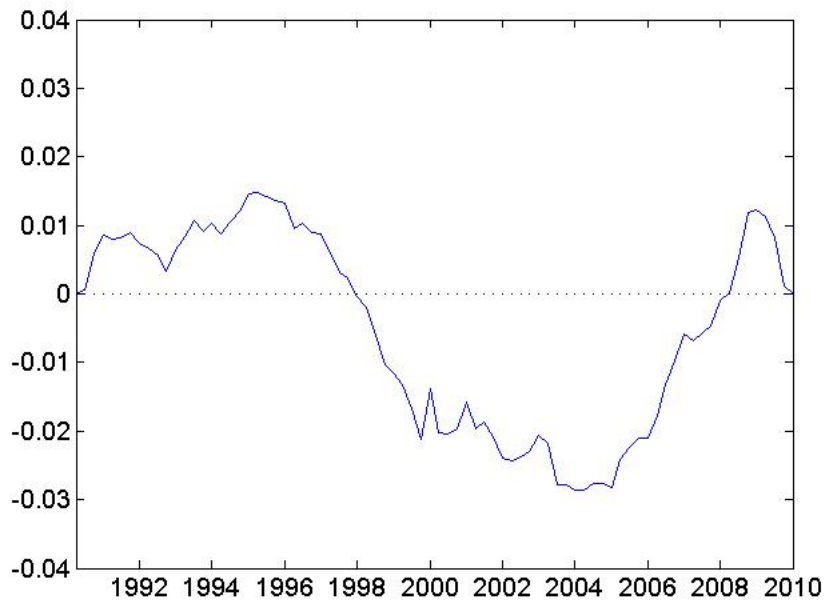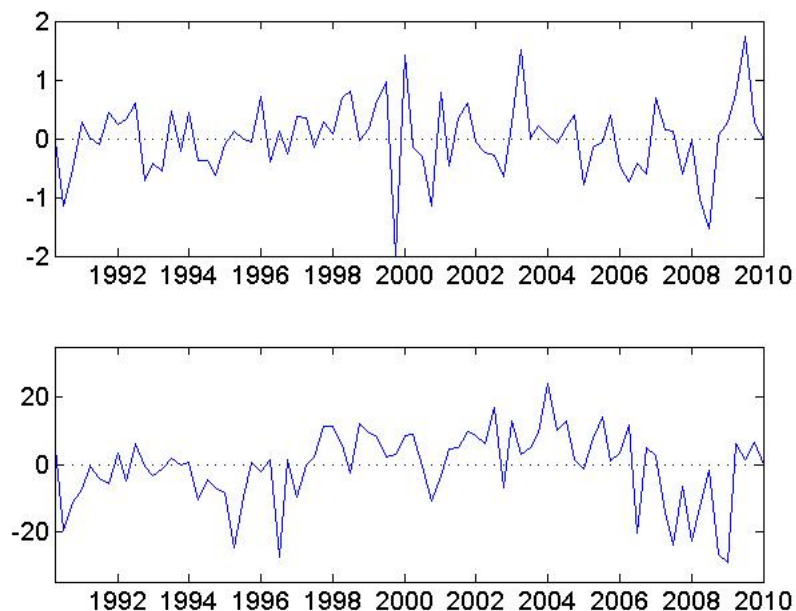
Figure 4: Smoothed permanent shocks (U.S. 1990-2010).

Figure 5: Smoothed transitory and noise shocks (upper and lower panel respectively) (U.S. 1990-2010).

Figure 6 plots a historical decomposition of series (5) and (6) (the long-run level of the permanent component as of time t and consumers' contemporaneous beliefs as of time t) when only the effect of estimated permanent shocks is taken into account. The long-run component is plotted using a dotted black line, and beliefs are plotted using a solid blue line. The (detrended) long-run component has an inverse-V shape, consistent with the idea of an acceleration of productivity (the upward slopping region), followed by a deceleration (the downward slopping region), and peaking around 1998. Beliefs follow a similar pattern but with a lag, peaking around 2002.

A possible out-of-sample check of the series of beliefs in Figure 6 can be done by comparing it to survey evidence. For instance, Consensus Forecasts publishes trend-growth expectations 6–10 years ahead for several countries. Taking the difference of the forecast for the U.S. and a weighted forecast aggregate other available countries in the data set gives the series plotted on the lower panel of Figure 7, that we reproduce from Hoffmann, Krause, and Laubach (2011) (p. 6). In the upper panel we plot our beliefs series, aligning the time axis to the Hoffmann et al. (2011) series. It is quite remarkable that such a simple model is able to produce realistic predictions regarding aggregate beliefs over the sample: A simple qualitative comparison between the two
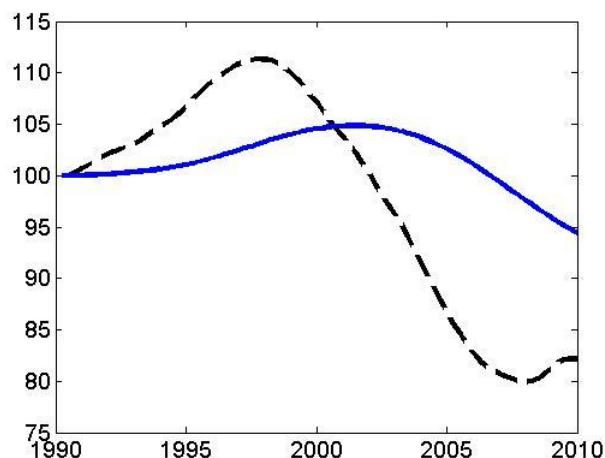
19

Figure 6: Historical decomposition of long-run component of productivity (black dotted line), and consumers' contemporaneous beliefs about it (solid blue line), permanent shocks, (U.S. 1990-2010).

series reveals that in both cases a representative agent seems to have been most optimistic around 2002–2004.

An issue that arises here concerns the sharp decline in several U.S. stock market indexes around 2001. In a model where stock prices are driven by beliefs about the future, this decline, taken at face value, constitutes a challenge to our story. However, some care is needed when interpreting this decline in the context of our model. First, depending on which index one looks at, the decline and the related financial turmoil around 2001 were much less important than in 2008. For instance, a look at the evolution of both the SP 500 and the Dow Jones Industrial Average reveals that their drop around 2001 was small compared to their drop starting in 2008. As such, the decline around 2001 could be the reflection of a high frequency event – aggravated perhaps by some orthogonal events as the September 11th terrorist attacks – that our medium frequency focus cannot possibly explain. Also, similar declines in the stock market did not happen in other episodes we consider here. For instance, in the case of the Great Depression, the stock market collapsed when the recession started and not before. Second, the crash of the dot-com bubble could have other roots than beliefs about long-run productivity. For instance, according to Pastor and Veronesi (2009), the adoption of a new technology can create time-varying shifts in discount rates that can contribute to bubble-like patterns in stock prices. Specifically, a technological revolution can generate bad news about "systematic risk", leading to
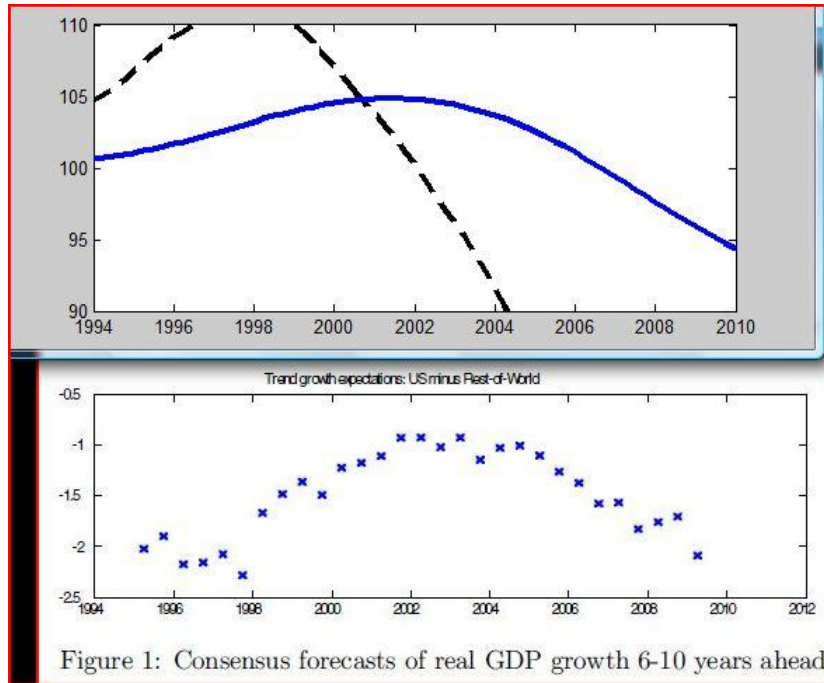
20

Figure 1: Consensus forecasts of real GDP growth 6-10 years ahead

Figure 7: Out-of-sample check: comparison of historical decomposition of beliefs (upper panel, solid blue line) and survey evidence on growth forecasts (lower panel, dotted blue line, reproduced from Hoffman et al. 2011).

an ex-post drop in stock prices. We plan to investigate this issue more deeply in future research.

The estimated shocks and states for Japan and the Great Depression have similar shapes and for brevity we do not report them. Instead, we now focus on a feature of the data that delivers estimated levels of the permanent component in the long run and beliefs as shown in Figure 6.

3.2 Dynamics of the Productivity-to-Consumption Ratio

A way to understand the results summarized by Figure 6 – the early pickup of permanent productivity, the delay in the adjustment of beliefs, and the implied lagged adjustment of consumption – is to look at the dynamics of the productivity-to-consumption ratio over the period 1990–2010, plotted in Figure 8. Productivity is the ratio of GDP divided by employment. Consumption is NIPA consumption divided by population. The series plotted is the logarithm of the ratio of productivity

21

to consumption. The vertical axis is centered around the average of the ratio over the period considered.

As the figure shows, the ratio has relatively high values at the start of this time window, with a slight increasing portion between 1990 and 1992. This is because during this period productivity is growing at a higher rate than consumption. The ratio starts declining around 1992, and this decline becomes more dramatic starting in 1997, where consumption grows at a considerably stronger rate than productivity. The ratio reaches its lowest point around 2007, after which a "correction" starts in which the ratio quickly goes back to its level from 20 years earlier. The correction is dramatic and coincides with the start of the Great Recession in 2007. Overall, the ratio appears to follow a slow moving "wave".

Importantly, these dynamics in the ratio are not due solely to the behavior of consumption. In fact, aggregate productivity growth rates are declining over the period – on average 1.87% from the first quarter of 1990 to the first quarter of 2005, dropping to 1.18% later on – a feature of the data consistent with the idea of a technological revolution happening at early stages of the time period considered.
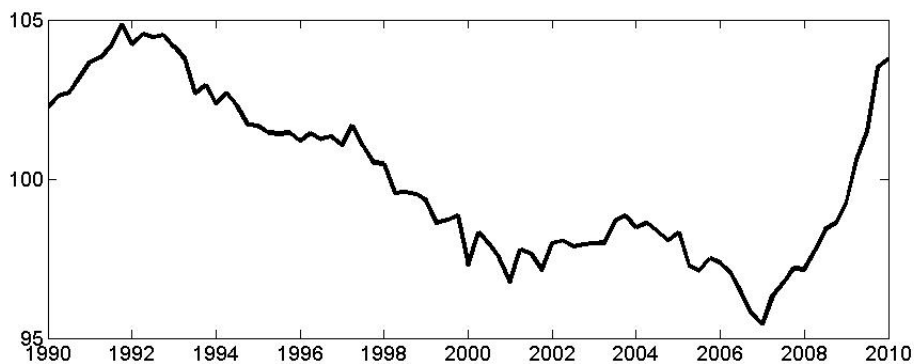
Figure 8: Productivity-to-consumption ratio, in logs (U.S., 1990-2010)

Figure 9 plots the same ratio for Japan. In this case we can see a more gradual increase of the ratio from its average over the period considered, reaching a peak in 1985. From this point on, the average growth rate of consumption is higher than the growth rate of productivity, and therefore the ratio decreases up to 1994. The lowest point of the ratio is reached in 1997, after which an upward movement brings the ratio back to its level in 1975, suggesting that similar to the previous case, the ratio followed a slow moving "wavy" path. Similar to the case of the Great Recession, Japanese growth rates of productivity are declining over the period: on average

3.22% between the first quarter of 1975 and the first quarter of 1990, and 1.06% from then on. The drop in growth rates is more important than in the previous case.
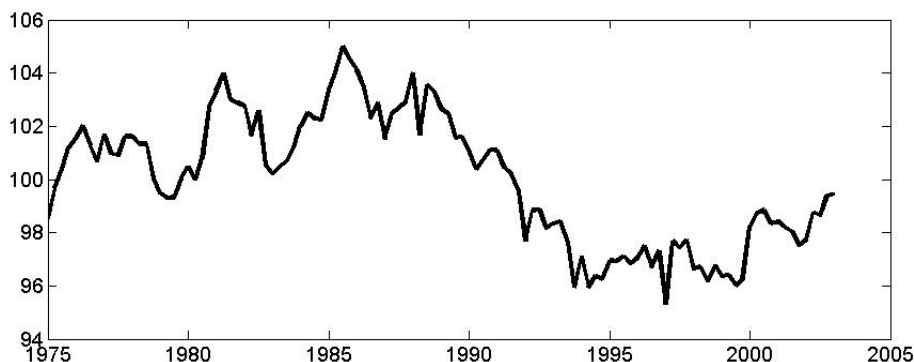

Figure 9: Productivity-to-consumption ratio, in logs (Japan, 1975–2003)

Figure 10 plots the ratio for the Great Depression. As explained previously, data limitations prevent us from looking at quarterly data before 1919. However, the ratio in this case seems to follow a similar "wavy" pattern as in the two previous figures. It starts at high values, then decreases, reaches a lowest point at the onset of the Great Depression in 1929, and then reverts back to its level of 14 years before. Also, average productivity growth rates are declining: 2.82% from the last quarter of 1919 to the first quarter of 1926, and -.91% later on.
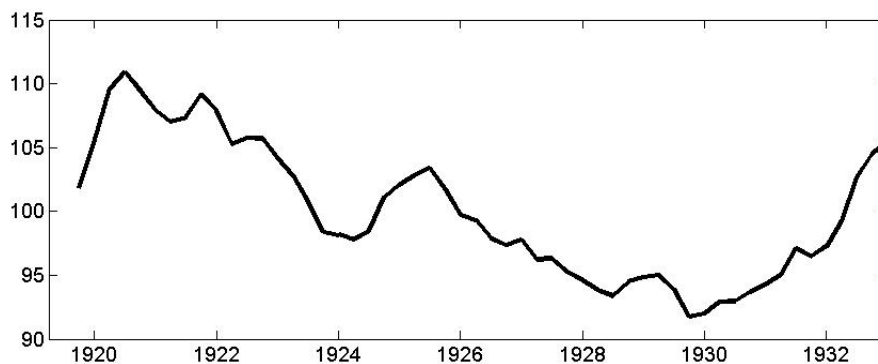

Figure 10: Productivity-to-consumption ratio, in logs (U.S., 1919–1933)

To summarize, this reduced-form analysis points to some similarities between the three episodes considered, and improves our understanding of the results obtained through structural estimation. First, in the three cases considered, the productivity-to-consumption ratio appears to follow a "wavy" shape. Second, average growth

rates of productivity over the periods considered are declining. These two sets of facts indicate that, in the three cases considered, there was a slow-moving boom and slowdown of aggregate productivity, together with similar dynamics of consumption, consumption adjusting with a significant lag with respect to productivity. We would also like to highlight that this way of looking at the data has the advantage of not requiring any particular detrending method – a sensitive issue in medium frequency analysis, and more generally, in macroeconomic time series analysis. In fact, Figures 8, 9 and 10 were constructed with raw (undetrended) data.

3.3 Implications for Debt Dynamics up to 2010, and Simulation up to 2025

Here we show the historical decomposition of the debt-to-output ratio up to 2010, and then use our estimated model to simulate the dynamics of this ratio, together with the evolution of consumption, up to 2025. Our goal is to use our model to get a sense of the speed of deleveraging and ask: when can one expect consumption to return to its steady state level after the Great Recession?

To this end, Figure 11 plots a historical decomposition of (detrended) productivity, net exports and the debt-to-output ratio using only the effect of smoothed permanent shocks. The level of productivity starts off at a normalized level of 100, and the other two are normalized at 0. In the case of productivity, the vertical axis represents log-deviations from the trend. In the case of both net exports-to-output and debt-to-output ratios, the vertical axis represents absolute percentage deviations from the trend. Productivity first increases due to the permanent positive shocks shown in Figure 6, peaks around the turn of the century, and starts declining after that, due to the effects of the negative permanent shocks that, according to our estimates, start hitting the economy in 1998. The developments of net exports (and consumption) lag those of productivity, with net exports reaching a minimum around 2003, and returning to positive values only around 2008. When net exports are negative the economy accumulates debt, with the debt-to-output ratio reaching its highest point around 2008.[18]

The dynamics of debt are determined by three elements. First, they depend on the persistence of the technology process ρ, because it governs the size of the income effect and the persistence of beliefs. The higher ρ, the larger the effect of a shock ε_t

---

[18]A close inspection of equation (19) reveals that changes of debt away from the steady state are slightly persistent, which is why the ratio starts declining a bit after net exports turn positive.
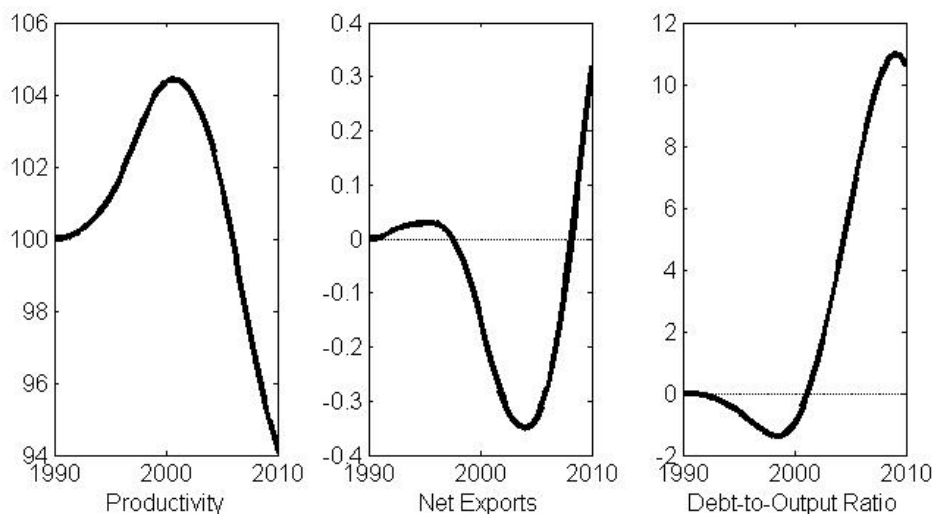
Figure 11: Historical decomposition of productivity, net exports, and debt (effect of permanent shocks only)

of given size in the long run, and the larger the income effect. The larger the income effect, the larger the accumulation of debt for a shock ε_t of given size. The higher ρ, the more persistent the process for Δx_t, and the more persistent agents' beliefs. The more persistent agents' beliefs, the more time they take to realize the end of the technological revolution. Second, the dynamics of debt depend on the relative size of the standard deviations σ_v^2, σ_z^2, and σ_s^2, because, as usual, these determine signal-to-noise ratios and thus the speed of learning. The smaller σ_v^2 with respect to the other two, the slower the learning, and the longer it takes for beliefs and consumption to adjust. Third, the dynamics are determined by the timing of the positive and negative shocks. Suppose there is only one positive and only one negative shock, of same size, and that they hit one after the other in two consecutive quarters. In this case, the effect of the shocks in the economy would be virtually nil. As shocks spread out, they can have an effect in the economy, in particular, agents can be optimistic when the negative shock hits. In the opposite extreme, if the negative shock never hits, agents are never "surprised".

Notice that despite the simplicity of the model, and the little amount of data used for the estimation (only a series of aggregate productivity and consumption), the dynamics of net exports and debt shown in Figure 11 appear to fit the evidence on global imbalances, housing markets, and household debt. Indeed, most of the large deterioration of the U.S. current account deficit took place in the late 1990s and early-mid 2000s. Also, most of the build-up of household debt started after the

25

year 2000.

A fun exercise is to use the values of state variables obtained by this historical decomposition and simulate a forecast of consumption. Figure 12 shows the results. According to this simulation, consumption would return to its steady state value only around 2025. Of course, this model is much too simple to take this quantitative prediction seriously, for instance, the effect of fiscal and monetary policy to fight against the consumption slump is not taken into account. However, the qualitative prediction of the model is clear: because the amount of debt in 2010 is large, and productivity low, the deleveraging process is notoriously slow.
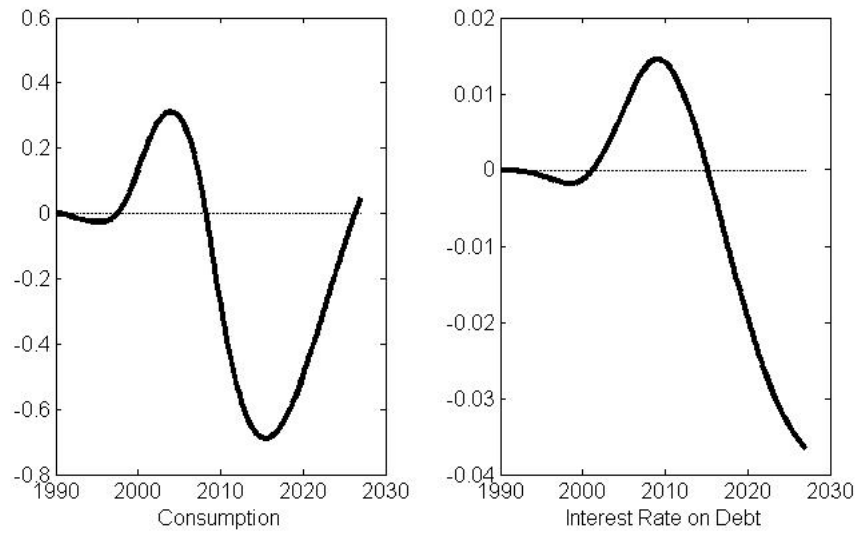
Figure 12: Simulation of the path of consumption and interest rate up to 2025 (on the vertical axis, zero is the steady state).

4 Preliminary Conclusions

We have explored the possibility that technological revolutions lead to great consumption slumps, using a model with imperfect information and anticipations about the future. The predictions of the model are intuitive, and provide a simple account of the Great Depression, the Japanese crisis of the 1990s, and the Great Recession.

In this exercise we have abstracted away from other factors that certainly affected the most salient of these episodes, i.e. the Great Recession, and also probably the other two. Among the most important ones, it is relevant to mention the role of financial deregulation, housing, and financial innovation. However, this abstraction has allowed us to pinpoint to simple medium-run dynamics common to the three episodes. These dynamics find a precise meaning in the context of our model.

References

Aguiar, M. and G. Gopinath (2007). Emerging market business cycles: The cycle is the trend. *Journal of Political Economy 115*(1), 69–102.

Atkeson, A. and P. Kehoe (2007). Modeling the transition to a new economy: Lessons from two technological revolutions. *American Economic Review 97*(1), 64–88.

Beaudry, P. and F. Portier (2006). Stock prices, news and economic fluctuations. *American Economic Review 96*(4), 1293–1307.

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni (2012). News, noise, and fluctuations: An empirical exploration. *Mimeo, IMF, Northwestern, and EIEF*.

Boz, E., C. Daude, and C. B. Durdu (2011). Emerging market business cycles: Learning about the trend. *Journal of Monetary Economics 58*(6–8), 616–631.

Chow, G. C. and A.-I. Lin (1971). Best linear unbiased estimation, distribution, and extrapolation of time series by related series. *Review of Economics and Statistics 53*(4), 372–375.

Christiano, L., C. Ilut, R. Motto, and M. Rostagno (2008). Monetary policy and stock market boom-bust cycles. *European Central Bank Working Paper No. 955*.

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), 1–45.

Gordon, R. J. and R. Krenn (2010). The end of the great depression 1939–41: Policy contributions and fiscal multipliers. *NBER Working Paper No. 16380*.

Greenwood, J. and B. Jovanovic (1999). The information-technology revolution and the stock market. *American Economic Review Papers and Proceedings 89*(2), 116–122.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton.

Hoffmann, M., M. Krause, and T. Laubach (2011). Long run expectations and "global imbalances". *Deutsche Bundesbank Discussion Paper No 01/2011*.

Jaimovich, N. and S. Rebelo (2009). Can news about the future drive the business cycle? *American Economic Review 99*(4), 1097–1118.

Pastor, L. and P. Veronesi (2009). Technological revolutions and stock prices. *American Economic Review 99*(4), 1451–1463.

Schmitt-Grohe, S. and M. Uribe (2003). Closing small economy models. *Journal of International Economics 61*, 163–185.

Smets, F. and R. Wouters (2007). Shocks and frictions in US business cycles: A bayesian DSGE approach. *American Economic Review 97*(3), 586–606.

A Appendix

A.1 Data Appendix

In the case of the Great Recession, the series for productivity is constructed by dividing GDP by labor input. GDP is measured by taking the series for Real GDP from the Bureau of Economic Analysis (available through the Federal Reserve Bank of Saint Louis online database). The labor input is measured by the employment series (Bureau of Labor Statistics online database, series IDs LNS12000000Q). The series for per capita net exports is constructed by dividing Real Net Exports by Population. The Real Net Exports are measured by the difference between Real Exports and Real Imports from the St. Louis Fed database (series EXPGSC96 and EXPGSC96 respectively). Population is from the BLS (series IDs LNS10000000Q).

In the case of Japan, all series come from the OECD website. Real GDP and Consumption are contained in the measure named VOBARSA. Employment comes from the OECD website. It is published in monthly frequency, and thus its frequency was changed to quarterly by computing the quarterly arithmetic average at every quarter. Population comes from the ALFS Summary tables in annual frequency, and thus a linear interpolation was performed to obtain quarterly frequency data.

In the case of the Great Depression, all series were obtained from Robert Gordon's website.

Our dataset is available upon request.

A.2 Proof of Proposition 1

When $C/Y = 1$, we can rewrite the equations describing the dynamics of the economy as

$$nx_t = b_{t-1} - \beta b_t \tag{20}$$

$$c_t + nx_t = y_t \tag{21}$$

$$c_t = -r_t + E[c_{t+1} + \Delta a_{t+1}] \tag{22}$$

$$c_t + \phi n_t = 0$$

$$y_t = n_t \tag{23}$$

$$r_t = \psi \cdot b_t \tag{24}$$

Substituting in nx_t from (21) y_t from (23) into (20), we obtain

$$
\begin{aligned}
b_t &= \frac{1}{\beta} b_{t-1} - \frac{1}{\beta} nx_t \\
&= \frac{1}{\beta} b_{t-1} - \frac{1}{\beta} (y_t - c_t) \\
&= \frac{1}{\beta} b_{t-1} + \frac{1}{\beta} \left(\frac{1}{\phi} + 1 \right) c_t
\end{aligned}
$$

Similarly, simplifying (22) using (24), we arrive at the two simple equations that fully characterize the dynamics of debt and consumption

$$
\begin{aligned}
c_t &= -\psi \cdot b_t + E[c_{t+1} + \Delta a_{t+1}] \\
b_t &= \frac{1}{\beta} b_{t-1} + \frac{1}{\beta} \left(\frac{1}{\phi} + 1 \right) c_t
\end{aligned}
$$

Let $\widehat{c}_t = c_t + a_t$ then

$$
\begin{aligned}
\widehat{c}_t &= -\psi \cdot b_t + E[\widehat{c}_{t+1}] \\
b_t &= \frac{1}{\beta} b_{t-1} + \frac{1}{\beta} \left(\frac{1}{\phi} + 1 \right) \widehat{c}_t - \frac{1}{\beta} \left(\frac{1}{\phi} + 1 \right) a_t \quad .
\end{aligned}
$$

This system has a solution under the form

$$
\widehat{c}_t = D_b b_{t-1} + D_k \begin{pmatrix} a_t \\ x_{t|t} \\ x_{t-1|t} \\ z_{t|t} \end{pmatrix} ,
$$

where the coefficients D_b and D_k are obtained using the method of undetermined coefficients. After some lengthy algebra:

$$
D_b = \frac{ -\left(1 - \beta - \psi \left(\frac{1}{\phi} + 1 \right)\right) - \sqrt{\left(1 - \beta - \psi \left(\frac{1}{\phi} + 1 \right)\right)^2 + 4\psi \left(\frac{1}{\phi} + 1 \right)} }{2 \left(\frac{1}{\phi} + 1 \right)}
$$

and

$$D_{k,1} = -\frac{x}{1-x}$$

$$D_{k,2} = \frac{(1-x)(1+\rho)-\rho}{(1-x)(1-\rho-x)}$$

$$D_{k,3} = -\frac{\rho}{1-\rho-x}$$

$$D_{k,4} = \frac{D_{k,1}\rho}{1-\rho-x}$$

where $x = (D_b - \psi)\frac{1}{\beta}\left(\frac{1}{\phi}+1\right)$.

We have $\lim_{\psi \longrightarrow 0} D_b = -\frac{\left(\frac{1}{\beta}-1\right)}{\frac{1}{\beta}\left(\frac{1}{\phi}+1\right)}$ so $\lim_{\substack{\psi \longrightarrow 0 \\ \beta \longrightarrow 1}} D_b = 0$. At the same time

$$\lim_{x \longrightarrow 0}\begin{pmatrix} D_{k,1} \\ D_{k,4} \end{pmatrix} = 0$$

and

$$\lim_{x \longrightarrow 0}\begin{pmatrix} D_{k,2} \\ D_{k,3} \end{pmatrix} = \frac{1}{1-\rho}\begin{pmatrix} 1 \\ -\rho \end{pmatrix} \quad .$$

In the end, the limiting dynamics of consumption are

$$\widehat{c}_t = \frac{1}{1-\rho}\left(x_{t|t} - \rho x_{t-1|t}\right) \quad ,$$

which is the limit of a New Keynesian economy in which the frequency of price adjustment tends to zero (see BLL for this proof).

∎

A.3   A Two-country Open Economy Model

The model in Section 2 can be extended to two countries. For each variable X of the home country, denote X^* the corresponding variable for the foreign country. The interest rate equation (9) is modified to:

$$R_t = R_t^* + \psi\left\{e^{\frac{B_t}{Y_t}-b} - 1\right\} \tag{25}$$

Let m and m^* denote the population sizes of the home and foreign country respectively.

An equilibrium is a set of choices $\{C_t, N_t, B_t, C_t^*, N_t^*, B_t^*\}_{t=0}^{\infty}$ and equilibrium interest rates $\{R_t, R_t^*\}_{t=0}^{\infty}$ such that

$$mB_t + m^* B_t^* = 0$$

and the interest rate spread $R_t - R_t^*$ follows (25).

We assume that the two countries have the same steady state growth rate so in steady state:

$$R = R^* = \frac{1}{\beta} \quad .$$

In the log-linearized version of this model, we replace the interest rate equations for the home and the foreign countries, equation (24), by:

$$r_t = r_t^* + \psi \cdot b_t \quad . \tag{26}$$

Moreover, we need to add the linearization for the bond market clearing conditions:

$$mb_t + m^* b_t^* = 0 \quad . \tag{27}$$