Technological Revolutions and the Three Great Slumps: A Medium-Run Analysis

Dan Cao and Jean-Paul L'Huillier*

This draft: July 2014

Abstract

The Great Recession, the Great Depression, and the Japanese slump of the 1990s were all preceded by periods of major technological innovation, which happened about 10 years before the start of the decline in economic activity. In an attempt to understand these facts, we estimate a model with noisy news about the future. We find that beliefs about long-run income adjust with an important delay to shifts in trend productivity. This delay, together with estimated shifts in the trend of productivity in the three cases, are able to tell a common and simple story for the observed dynamics of productivity and consumption on a 20 to 25 year window. Our analysis highlights the advantages of a look at this data from the point of view of the medium run.

Keywords: Aggregate productivity, permanent income, learning.

JEL codes: E21, E27, E32, N10.

---

*Georgetown University and EIEF. This is a revised version of the paper: "Technological Revolutions and Debt Hangovers: Is There a Link?" First draft: December, 2012. We benefited from useful feedback from Mark Aguiar, David Backus, Susanto Basu, Paul Beaudry, Olivier Blanchard, Emine Boz, Claire Brennecke, Robert Chirinko, Chris Faulkner-MacDonagh, Jean Flemming, Gita Gopinath, Pierre-Olivier Gourinchas, Aubhik Khan, Jennifer La'O, Francesco Lippi, Virgiliu Midrigan, Ezra Oberfield, Guillermo Ordoñez, Luigi Pistaferri, Franck Portier, Giorgio Primiceri, Vincenzo Quadrini, Chris Sims, Daniele Terlizzese, Julia Thomas, Nicholas Trachter, Eric Van Wincoop, Laura Veldkamp, Pietro Veronesi, Jing Zhang, and seminar and conference participants at the Federal Reserve Banks of Chicago and Philadelphia, Boston College, Northwestern University, Princeton University, the University of Virginia, the 5th Joint French Macro Workshop, and the European Summer Symposium in Economic Theory. Dan thanks EIEF for hospitality, and Jean-Paul thanks NYU.

"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 25th, 2012*

1 Introduction

A medium-run look at the three most important private-debt recessions in developed economies reveals that they were all 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) (Greenwood and Jovanovic 1999; Hobijn and Jovanovic 2001; Pastor and Veronesi 2009). Similarly, the Japanese slump of the 1990s was preceded by a period of unprecedented industrial innovation in the 1980s. During this period, Japanese corporations developed several key electronic products with great success.[1] We view this period as containing the elements of a technological revolution, which in the particular case was mostly concentrated in Japan. Finally, before the Great Depression, roughly between 1915 and 1925, the United States witnessed the so-called 2nd Industrial Revolution, a period of important industrial developments (David and Wright 2000).[2]

Thus, in each of these cases there seems to be roughly a 10-year gap between the technological revolution and the start of the economic slump. At face value, this suggests the existence of slow-moving, joint dynamics of technological progress and economic activity, *common* to all three episodes. In this paper, we investigate whether there is indeed evidence of these dynamics in the data, and make an effort to characterize them within a simple and intuitive framework. We take a simple permanent income model in which a representative agent learns slowly about his future income. Future income is determined by technological progress, which in turn can be gauged from the trend of productivity. However, detecting changes in the trend of productivity can be a challenging task for the agent due to imperfect information. Consumers that

[1]The two main players here were the Sony Corporation and JVC, who developed a large number of these electronic products. To name a couple of salient examples, consider the Walkman, the VHS, or the Betamax.

[2]The key general purpose technology here was the combustion engine. Among other things, this technology made possible the mass production of automobiles for the American household by the Ford Motor Company. This also brought drastic improvements in management as, for instance, the use of the moving assembly line (Bardou et al. 1982).

update their beliefs on the basis of fairly noisy signals adjust their behavior gradually. This allows us to fit the slow-moving dynamics in spending present in the data.

In this exercise, our main object of interest is the beliefs about long-run income determined by learning, or "beliefs about the long run". When we look at the data through the lens of the model, we find that in the three cases the joint dynamics of productivity and beliefs about the long run can be characterized by a slow-moving cycle. This cycle can be summarized by the following sequence of events. First, there was an initial increase in trend productivity, which was contemporaneous to the waves of innovation mentioned previously. Second, this pickup of productivity generated an increase in beliefs about the long run. This increase in beliefs increased consumption. Importantly, this (rational) "wave of optimism" came some years after the initial increase in trend productivity. The reason is a delay in the adjustment of beliefs due to noisy information. Third, there was a decline in trend productivity, probably caused by a slowdown in the pace of innovation.[3] Fourth, when consumers received enough information to learn about the drop in trend productivity, they decreased their beliefs about the long run. This "wave of pessimism" generated a persistent decline in consumption.[4]

Our model has two main ingredients. The first one is the presence of both permanent and transitory shocks to productivity (Aguiar and Gopinath 2007; Justiniano, Primiceri, and Tambalotti 2010). As Aguiar and Gopinath, we consider an open economy framework and use permanent shocks to generate shifts in trend productivity. We complement their analysis by deriving a closed-form solution to study these effects, and extend it to the study of the three episodes mentioned above. The second one is the presence of news about the future (Beaudry and Portier 2006; Jaimovich and Rebelo 2009; Barsky and Sims 2012). The novelty in our framework is the presence of noisy news (Blanchard, L'Huillier, and Lorenzoni 2013), together with agents' rational reaction to this noise in the news. However, our focus is on the effect of shifts in the trend of productivity, instead of the effect of noise shocks.

To estimate shocks to the trend of productivity, we use a tractable frame-

[3]This finding is closely related to the evidence recently collected by Fernald (2012a), who in the case of the Great Recession documents that the growth of U.S. labor productivity slowed down after 2004. The slowdown was more pronounced in IT-intensive industries. For more details, see Section 3.

[4]Using a different approach, Eggertsson (2008) emphasizes the role of expectations in the recovery *out of* the Great Depression. In particular, he emphasizes the role of a policy-driven shift in expectations that took place when Roosevelt took office at the end of 1932.

work in which beliefs about the long run drive the behavior of consumption. As econometricians, the permanent income logic together with rational expectations allow us to infer the underlying movements in trend productivity by looking at consumption. Here, we borrow the basic idea of an important body of work on household income dynamics, e.g. Blundell and Preston (1998), or Blundell, Pistaferri, and Preston (2008). Within this literature, Guvenen (2007) uses a life cycle learning model.

Our model has implications for the dynamics of debt, determined by the difference between households' spending and income. This connects our work to a growing literature analyzing the leveraging and deleveraging of U.S. households (Midrigan and Philippon 2011; Justiniano, Primiceri, and Tambalotti 2014a, 2014b). In some sense, our goal is less ambitious because we use a simpler model and solve it using a first order log-linear approximation. Therefore, we do not intend to provide a full quantitative account of the leveraging and deleveraging. At the same time, taking a longer-run perspective allows us to clarify that slow-moving movements in productivity explain part of the slow-moving dynamics of debt.

Specifically, we proceed as follows. First, we estimate our model through standard methods and then use the variance decomposition of beliefs at different horizons in order to gauge which of the shocks present in the model explain its variability on the medium run. We define the "medium run" as an horizon of about 5 years or more after the impulse of a particular shock. This decomposition indicates that most of the variability of consumption in the medium run is explained by permanent productivity shocks.[5]

Having established the importance of permanent shocks to understand the medium-run dynamics of the beliefs, we estimate these shocks using a Kalman smoother. We then feed the estimated permanent shocks into the model and shut down other shocks. We do this in order to simulate the associated beliefs about the long run, which we label "model-predicted beliefs about the long run". We then perform an out-of-sample check of these model-predicted beliefs by comparing them to survey evidence for the U.S. economy, 1994–2010. Notice that in this exercise we shut down all other shocks in the model.[6] We find that according to both the model-predicted beliefs about the long-run and the survey,

[5]Pintus and Suda (2013) also stress the importance of gradual learning to understand the recent recession in the U.S.

[6]Our empirical exercise is related to the theoretical contribution by Boz (2009), in which optimism following a "miracle" performance can lead to a downturn. Independently, Piazza (forthcoming) and Pintus and Wen (2013) model a similar interaction between development, demand, and credit.

the U.S. consumer was most optimistic about his long-run income around 2004.

In order to shed light on the properties of the data that deliver the shape of the model-predicted beliefs about the long run, we also present some reduced-form evidence by focusing on the observed dynamics of the ratio of productivity-to-consumption. We argue that – within our estimated model – this ratio is particularly informative for the estimation of permanent productivity shocks. Indeed, in the model, productivity determines income, and beliefs about the long run determine consumption. Therefore, given the variance decomposition of consumption, the joint medium-run evolution of these two variables should be determined by permanent technology shocks. Accordingly, we find that this ratio has a similar medium-run shape in the three cases.[7]

Altogether, the exercises we perform deliver three main substantive results. First, there is a significant delay in the adjustment of beliefs about the long run. The reason is the estimated amount of noise in the information consumers receive about future income, which is quite large. We quantify this delay by computing the half-life of beliefs after an impulse to the trend-growth of productivity in our estimated model. The exact measure of the delay varies from one episode to the other, but in all three cases we find a sizeable delay of at least 1 year. Second, the medium-run trend of the productivity-to-consumption ratio computed using an HP-filter has the shape of an "up-and-down wave": first increases, then decreases, and then again increases, reverting back to its value at the start of the cycle. Although the whole length of this cycle varies from case to case, it seems to be of 20 to 30 years. As argued below, in either the "no-news" or perfect foresight benchmarks, this ratio would have a different shape. Therefore, our learning model is a way of accommodating this particular feature of the data. Third, a simulation of debt dynamics in our model indicates that the leveraging and deleveraging of households lagged the up-and-down movements in productivity. The reason was the delay in the adjustment of beliefs. Moreover, this simulation indicates that households were leveraging up *precisely* when productivity was slowing down. Thus, they accumulated more debt than intended because they failed to immediately recognize the slowdown in productivity.

In the literature, little attention has been devoted to the study of medium-

[7] This historical stylized fact is akin to the well-known work by Reinhart and Rogoff (2008,2011). However, we do not seek to address the abrupt financial meltdowns emphasized there. Reinhart and Reinhart (2010) look at a number of other aggregate indicators as unemployment, housing prices, inflation and credit, using a reduced form approach. See Syverson (2013) for other interesting parallels between the 1920s and the 1990s.

term aggregate consumption dynamics, as one can infer from the large literature on empirical DSGEs which focuses on the short run. A noticeable exception is the paper by Comin and Gertler (2006). They generate medium-term dynamics using an endogenous determination of productivity through the explicit modeling of R&D. In the case of our paper, we simplify the determination of productivity by making it exogenous, and instead focus on the dynamics of learning. These turn out to generate smooth, medium-term, dynamics of consumption. Also, Section 4 establishes a further link to Comin and Gertler's work by the use of filtering techniques to look at the medium-term dynamics of time series data.

Our application provides an interpretation of consumption disasters based on shifts in beliefs about the long run. A key reference here is Nakamura, Steinsson, Barro, and Ursua (2013) who, using different techniques, separate the large *short-run* drop from the sustained *long-run* impact of disasters.[8] Our paper focuses solely on the medium-run properties of consumption, and links them to underlying movements in trend productivity using a learning channel.

The rest of the paper proceeds as follows. We first present the model (Section 2). We then discuss its estimation and present these results (Section 3). Here, we generate the model-predicted beliefs about the long run, and perform the out-of-sample check. We then turn to the properties of the productivity-to-consumption ratio in the data (Section 4). Afterwards we use our model to analyze the implication of our results of household debt (Section 5). We then conclude (Section 6). The Appendix contains several proofs and a detailed description of our data. The Supplementary Material presents a number of supplementary results.

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 2013, henceforth BLL).[9] Specifically, productivity a_t (in logs) is the sum of two

[8] For related work on rare disasters, see Barro (2006) and Barro and Ursua (2011).

[9] Boz, Daude, and Durdu (2011) use a similar framework. We simplify it further by removing labor supply and capital. Those extra ingredients do not change anything to our analysis, as we explain below (p. 11).

components, permanent, x_t, and transitory z_t:

$$a_t = x_t + z_t \quad . \tag{1}$$

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

$$\Delta x_t = \rho_x \Delta x_{t-1} + \varepsilon_t \quad . \tag{2}$$

The transitory component follows the stationary process

$$z_t = \rho_z z_{t-1} + \eta_t \quad . \tag{3}$$

The coefficients ρ_x and ρ_z are 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

$$\rho_x = \rho_z \equiv \rho \quad , \tag{4}$$

and that the variances satisfy

$$\rho \sigma_\varepsilon^2 = (1-\rho)^2 \sigma_\eta^2 \quad , \tag{5}$$

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

$$\mathbb{E}[a_{t+1} | a_t, a_{t-1}, ...] = a_t \quad . \tag{6}$$

This assumption is analytically convenient and broadly in line with 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_t shock and forecast future productivity growth or to an η_t shock and forecast mean reversion. When (4) and (5) are satisfied, these two considerations exactly balance out and expected future productivity is equal to current productivity.[10]

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

$$s_t = x_t + \nu_t \quad , \tag{7}$$

[10] See BLL for the proof.

where ν_t is i.i.d. normal with variance σ^2_ν.

We think of ε_t as the "news" shock because it builds up gradually and thus provides (noisy) advance information information about the future level of productivity (through the signal (7)). Our focus throughout the paper is on the dynamics implied by this shock. It is useful to say a word about the methodological role of the signal (7) in our exercise. It plays a key role in our identification by providing an extra source of information to consumers regarding the permanent component. Indeed, through this assumption the econometrician will be able to make inferences about trend-productivity by looking at the behavior of consumption. As mentioned in the introduction, this connects our paper to the work of Blundell and Preston (1998) and Blundell, Pistaferri, and Preston (2008). (Our identification strategy is discussed in detail below.)[11]

2.1.1 Slow Adjustment of Beliefs

Here we focus on an important property of the signal extraction problem for our purposes. Agents optimally form beliefs about the permanent component x_t using a Kalman filter.[12] Then, they form beliefs about the future path of x_t. The following definition is useful to make these ideas precise.

Definition 1 (BLR) *Given information at time t, the agent's best estimate of the productivity in the future is*

$$\lim_{\tau \longrightarrow \infty} \mathbb{E}_t \left[a_{t+\tau} \right] = \frac{\mathbb{E}_t \left[x_t - \rho x_{t-1} \right]}{1-\rho} = \frac{x_{t|t} - \rho x_{t-1|t}}{1-\rho}, \tag{8}$$

*where $x_{\tau|t}$ denotes the conditional expectation $\mathbb{E}_t[x_\tau]$ of x_τ on information available at time t. We call the estimate of long-run productivity, **beliefs about the long run (BLR)** and denote it by $x_{t+\infty|t}$.*

The first equality is proved in the Appendix and the second equality comes directly from the definition of $x_{\tau|t}$. In Proposition 1 below we show that these BLR will determine consumption.

Because of noisy information, agents will be slow to adjust their beliefs $x_{t+\infty|t}$. In particular, they will be slow to adjust their beliefs following a per-

[11] Related and important contributions on the impact of noise, or more broadly, changes in expectations are by Angeletos and La'O (2009,2013). As it will become clear, our noisy news approach is different, especially because it captures medium-term fluctuations once the model is estimated. Forni, Gambetti, Lippi, and Sala (2013b) also use the term "noisy news", but they use a different specification of the information structure. See also Forni et al. (2013a).

[12] The construction of the filter is standard, but see p. 13 for more details on this.

manent shock ε_t.

Definition 2 (Delayed adjustment of beliefs) *After a permanent shock,* $\epsilon_t = 1$, *under perfect information, BLR jumps immediately to the long-run level* $1/(1-\rho)$ *and stays at that level in the absence of future shocks. However, under imperfect information, it takes time for the BLR to reach the long-run level. We define the* ***delay*** *by the time it takes BLR to reach half of the long-run level.*

2.2 Consumption, Production and Net Exports

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

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \log C_t\right]$$

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

$$C_t + B_{t-1} = Y_t + Q_t B_t \quad , \tag{9}$$

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:

$$Y_t = A_t N \quad , \tag{10}$$

where $A_t = e^{a_t}$. We abstract from fluctuations on employment, i.e. the consumer supplies labor N inelastically.[13] The resource constraint is

$$C_t + NX_t = Y_t \quad .$$

The price of debt is sensitive to the level of outstanding debt, taking the form used by Schmitt-Grohe and Uribe (2003), and Aguiar and Gopinath (2007), among others:

$$\frac{1}{Q_t} = R_t = R^* + \psi\left\{e^{\frac{B_t}{Y_t}-b} - 1\right\} \quad , \tag{11}$$

[13]This approach is, to some extent, justified by our focus on the medium-run. However, we have used labor supply in previous versions of this model and obtained very similar results. We comment more on this feature of the model below (p. 11).

where b represents the steady state level of the debt-to-output ratio.[14]

The only first-order conditions from the optimization problem of the consumer is:

$$\frac{1}{C_t} = \beta R_t \mathbb{E}_t \left[\frac{1}{C_{t+1}} \right] \quad , \tag{12}$$

In order to log-linearize the model, we define four endogenous variables c_t, r_t, b_t, and nx_t as follows:

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

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

and

$$b_t \equiv \frac{B_t}{Y_t} - b \quad ,$$

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

In the definition of c_t, we need to use the ratio of C_t over A_t to ensure stationarity, and C/A is the steady-state value of this ratio. Similarly, NX/Y is the steady-state value of the net exports-to-output ratio. In order to examine the dynamics of consumption, we also define another variable which is the log-deviation of consumption:

$$\widehat{c}_t = c_t + a_t \quad .$$

In the Supplementary Material, we derive the log-linearization of the equilibrium. This equilibrium is given by the equations for the shock processes (1), (2), and (3), and other four equations:

$$c_t = -r_t + \mathbb{E}_t[c_{t+1} + \Delta a_{t+1}] \quad , \tag{13}$$

$$r_t = \psi \cdot b_t \quad , \tag{14}$$

$$c_t + \frac{1}{C/Y} nx_t = 0 \quad , \tag{15}$$

[14] It is straightforward to generalize our model to a two-country economy, and our main results do not change in that case. See the discussion in Appendix C.

$$nx_t = b_{t-1} - \beta b_t + \frac{1 - C/Y}{1-\beta}\left(-\Delta a_t + \beta r_t\right) \quad . \tag{16}$$

This model admits a closed-form solution. It is presented in Appendix B.

To illustrate the effect of a permanent shock on the endogenous variables of this system, we parameterize the model as follows. The period length is one quarter. The discount factor β is set at 0.99. The elasticity of the interest rate, ψ, is set to a low value, 0.0010, following previous literature (Neumeyer and Perri 2004; Schmitt-Grohe and Uribe 2003; Aguiar and Gopinath 2007). Under this common parametrization, BLR is the main driver of consumption, as established by the following proposition.

Proposition 1 *As $\beta \longrightarrow 1$, and $\psi/(1-\beta) \longrightarrow 0$, consumption is only a function of BLR. Specifically,*

$$\widehat{c}_t = \frac{1}{C/Y} x_{t+\infty|t} \quad .$$

The proof is in the Appendix. In Cao and L'Huillier (2014a) (available on our webpages) we prove a version of this theorem for a more general model that includes labor supply and capital. Therefore, for the standard parametrization in the literature, including those ingredients in our framework does not change our results. Our focus on consumption finds empirical support in the work by Mian and Sufi (2012), which provides disaggregated evidence on the role of consumption in the U.S. economic slump.

The rest of the parameters is taken from the estimation of the model for the United States (1990–2013) below. The parameter ρ is set at 0.97, implying slowly building permanent shocks and slowly decaying transitory shocks. The standard deviation of productivity growth, σ_a, is set at 0.64. These values for ρ and σ_a yield standard deviations of the two technology shocks, σ_ε and σ_η, equal to 0.02% and 0.63%, respectively. The standard deviation of the noise shock, σ_ν, is set to 7.39%, implying a fairly noisy signal.

Figure 1 shows a simulation of the model for these parameter values. The figure shows the responses of productivity a_t, net exports nx_t, and debt-to-output ratio 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 one year (four quarters). The scale of productivity is relative percentage deviations from steady state. The scale of both net exports and the debt-to-output ratio are absolute percentage deviation from the steady state value of net exports-to-output, NX/Y, and debt-to-output, b.

Figure 1: Impulse Response Functions to Permanent Technology Shock

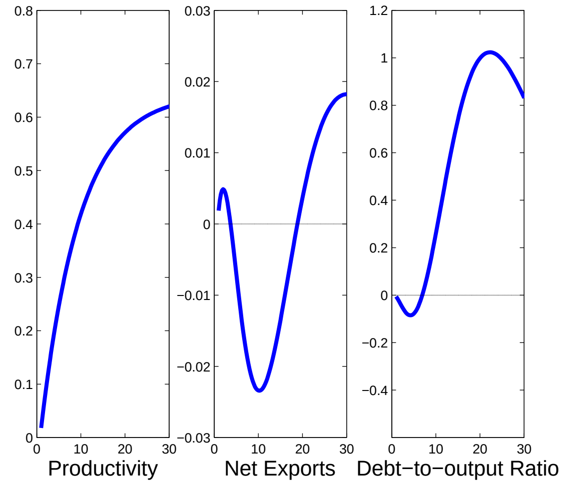

Notes: Impulse response functions of, from left to right, a_t, nx_t, and b_t. Units in the vertical axis are percentage relative deviations from steady state in the case of productivity, and absolute percentage deviations from steady state in the case of net exports and the debt-to-output ratio. The time unit on the x-axis is one year (four quarters).

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 6 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 3 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 0.63) and net exports and the debt-to-output ratio go back to zero.

3 Estimation

In this section we first explain how we estimate the model. We then show the results for the Great Recession, and we perform an out-of-sample check of this estimation by comparing the estimated model-predicted BLR to survey evidence. We then show the results for Japan and the Great Depression.

3.1 Data Sets and Estimation Procedure

Data. Our data set includes series on productivity, TFP, consumption, and net exports. We use quarterly data. The series for the Great Recession were obtained from the Bureau of Economic Analysis and the Bureau of Labor Statistics. The series for Japan were obtained from the OECD.

In the case of the Great Depression, we have data for the components of GDP from the Gordon-Krenn data set.[15] Gordon and Krenn (2010) used the Chow and Lin (1971) method 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 (1961, Appendix A, Table XXIII, 2nd column) data set for employment, using a linear interpolation out of the annual series.

The Data Appendix contains further details on the data used and the construction of the variables.

Procedure. For our baseline estimations, we fix β and ψ. The discount factor β is set at 0.99. ψ is set to low value, 0.0010, following previous literature (Neumeyer and Perri 2004; Schmitt-Grohe and Uribe 2003; Aguiar and Gopinath 2007). We estimate the remaining parameters as described below. Notice, given the random walk assumption (6) for a_t, σ_ε and σ_η are determined by ρ and σ_a. For robustness, we present below an estimation including ψ among the parameters to estimate.

Our log-linearized model can be represented in state-space form. The information structure in this model is identical to the one used in BLL, and more details are provided there on how to compute the likelihood function for a general representative-agent model with signal extraction.[16] The main idea is first to solve the consumer's Kalman filter to obtain the dynamics of consumer's expectations, and next to build the econometrician's Kalman filter, including in the list of unobservable state variables the consumer's expectations. The model can then be estimated through Maximum Likelihood (ML).

We follow Aguiar and Gopinath (2007), and thus in our baseline estimations we include the demeaned first differences of the logarithm of labor productivity

[15] In this case our sample length is restricted by the fact that there are no quarterly data on GDP components before the end of World War I in 1918.

[16] See Appendix 5.1 of BLL.

Δa_t and the demeaned ratio of net exports-to-GDP nx_t as observable variables. Using consumption instead of net exports did not change the results.

3.2 Great Recession

Here we present our baseline estimates for the Great Recession.

Table 1 contains the parameter estimates. The persistence parameter ρ is estimated at 0.97, implying very persistent processes both for the permanent and the transitory components of productivity. The standard deviation of productivity is estimated at 0.64% in the case of the Great Recession. Given the random walk assumption (6) for productivity, the high values of ρ imply a standard deviation for permanent technology shocks that is fairly small, of 0.02%, and a fairly big standard deviation for the transitory technology shock, of 0.63%. The standard deviation of noise shocks is large, 7.39%.

Although it may seem like a natural comparison, it is in fact misleading to compare the standard deviation of noise shocks (7.39%) to the standard deviation of permanent shocks (.02%). The reason is that the signal is about the permanent component x_t itself, and not about the shock ε_t (equation 7). This has two (somewhat subtle) implications for the interpretation of the estimates. First, even though the shocks are estimated to be small, they are also very persistent, and this results as in considerable volatility in x_t. Second, this permanent component is hit by a permanent shock every quarter, and previous shocks are not revealed to the agent in real time. This introduces substantial uncertainty in the learning process. A useful quantifier is provided by the one-step ahead uncertainty in the inference about x_t at time t, $\sqrt{Var_{t-1}(x_t)}$, which is given by the solution to the Kalman filter. For the estimates shown in Table 1 this computation results in $\sqrt{Var_{t-1}(x_t)} = 1.40\%$, which is of the same order of magnitude as σ_ν. Even though the signal is fairly imprecise, the agent's uncertainty is also large, and therefore the signal is not disregarded by the agent.

Notice also that permanent shocks to productivity are small compared to transitory shocks. This implies that, conditional on having observed the previous period's productivity a_{t-1}, current productivity a_t is also a fairly imprecise signal about x_t. To sum up, this discussion illustrates the major signal extraction problem that consumers face according to our estimation. Accordingly, the delay in learning is quite long, computed to 5.25 years for the parameters above.

Table 1: Parameter Estimates, Great Recession

Parameter	Description	Value	s.e.
ρ	Persistence tech. shocks	0.97	0.01
σ_a	Std. dev. productivity	0.64	0.04
σ_ε	Std. dev. permanent tech. shock (implied)	0.02	–
σ_η	Std. dev. transitory tech. shock (implied)	0.63	–
σ_ν	Std. dev. noise	7.39	2.04

Notes: ML estimates of the log-linearized state-space representation of the model. The observation equation is composed of the first differences of the logarithm of U.S. labor productivity and the ratio of net exports-to-GDP. The sample is from 1990:Q1 to 2013:Q1. Standard errors are reported to the right of the point estimate. The values for σ_ε and σ_η are implied by the random walk assumption (6) for productivity.

In this estimation there is no need to recur to Bayesian methods. In fact, we hit a unique global maximum for the likelihood function. Next we provide some intuition for the identification of these parameters.

3.3 Intuition for Parameter Identification

We now discuss the identification of the parameters. The derivation of all expressions discussed here and not presented previously can be found in Appendix F (Supplementary Material).

The intuition for identification comes from considering (4), (5), and the VAR representation of the limiting model (Proposition 1). This representation is given by the following two equations

$$\widehat{c}_t = \widehat{c}_{t-1} + u_t^c \tag{17}$$

$$a_t = \rho a_{t-1} + \frac{C}{Y}(1-\rho)\,\widehat{c}_{t-1} + u_t^a \quad , \tag{18}$$

where u_t^c and u_t^a are innovations (from the perspective of the econometrician). According to (17), consumption in the limiting model is a random walk, which simply follows from the law of iterated expectations. Equation (18) clarifies an interesting property of productivity in this model. Even though productivity a_t was restricted by (4) and (5) to have a univariate random walk representation, it is *no longer* a random walk in the bivariate representation, i.e. when conditioning its expected changes on the past value of consumption $\widehat{c}_{t-1}$. The reason is as follows. Past consumption $\widehat{c}_{t-1}$ carries extra information beyond the previous realization of productivity a_{t-1} about the permanent component x_t. This information comes from the signal s_{t-1} that consumers have received which, due to the persistence of the permanent component, helps them forecast its future path.

The parameter σ_a is identified by the standard deviation of the growth rates of productivity Δa_t. Identification of ρ comes from equation (18), which can be estimated by OLS in the following form:

$$\Delta a_t = -(1-\rho)(a_{t-1} - C/Y \cdot \widehat{c}_{t-1}) + u_t^a \quad . \tag{19}$$

The intuition provided by equation (19) is closely related to the permanent income hypothesis. Indeed, how much consumption deviates from current productivity reflects beliefs of consumers about future income, i.e. BLR, and by implication contains information about future changes in a_t.[17] The higher is consumption with respect to current productivity at $t-1$, the higher the expected productivity growth at t. The coefficient in front of the productivity-to-consumption ratio identifies ρ. If the permanent component is not very persistent (ρ is low), its expected long-run level is close to its current level, and the correlation between the ratio and productivity changes one quarter ahead is high. Instead, when the permanent component is very persistent (ρ is high), its expected long-run level is different from its current level, and the correlation between the ratio has and productivity changes one quarter ahead is low. Notice, this does not reflect a failure of consumers to forecast trend productivity, because at higher horizons the equation is

$$a_{t+j} - a_t = -(1-\rho^j)(a_{t-1} - C/Y \cdot \widehat{c}_{t-1}) + u_{t+j}^a \quad , \tag{20}$$

and thus for long horizons (high j) the coefficient in front of the ratio goes to -1. In other words, the longer the horizon, for a given variance of the productivity-to-consumption ratio, the higher the correlation between the productivity-to-consumption ratio and trend productivity. Notice also that these equations are valid for any degree of noise in the signal. In particular, the relationship between trend productivity and the productivity-to-consumption ratio (19) holds on average and takes into consideration the extra volatility of the ratio coming from the noise in the signal.[18]

[17] Similar to Campbell (1987) the consumer "saves for a rainy day", i.e., negative $\hat{c}_{t-1}$ predicts low future productivity growth Δa_t.

[18] Notice that in the model productivity follows a random walk, and therefore productivity by itself is

Having identified ρ, the sizes of the permanent and transitory shocks σ_ε and σ_η can be derived from

$$\begin{aligned} \sigma_\varepsilon &= (1-\rho)\sigma_a \\ \sigma_\eta &= \sqrt{\rho} \cdot \sigma_a \quad , \end{aligned}$$

which follow from (4), (5) and (6).

It remains to discuss the identification of the standard deviation of noise shocks, σ_ν. This is determined by the correlation between innovations to consumption u_t^c and innovations to productivity u_t^a. If the signal is not informative ($\sigma_\nu \longrightarrow \infty$), the only information available to consumers is productivity itself, and this correlation is 1. If the signal is perfectly informative ($\sigma_\nu \longrightarrow 0$), this correlation attains a lower bound.[19] The relation is monotonic and uniquely pins down σ_ν.

To sum up, first we have discussed how the permanent income hypothesis together with rational expectations help us decompose the productivity series into a permanent component and a transitory component. This is translated into parameters ρ, σ_ε, and σ_η that define the decomposition. The data favors a decomposition of productivity with a very smooth permanent component, with small and highly persistent permanent shocks that have large effects in the long run. Interestingly, these results are connected to the influential finance literature on risks for the long run (Bansal and Yaron 2004), in which a key ingredient is the presence of a small but persistent growth rate component. Second, we have discussed how our procedure pins down the accuracy of consumers' inferences. This leads to an estimate for σ_ν. Consumption innovations seem on average fairly disconnected from productivity innovations in (17), which leads to estimating a significant amount of noise.[20]

not useful to identify ρ.

[19]See BLL for the computation of this bound.

[20]It would have been possible to estimate the model using the relationships above. However, we decided to use ML instead because it does not require assuming the conditions of Proposition 1 hold exactly, and it should be able to extract more information from the data.

3.4 Variance Decomposition and Estimated Permanent Shocks

Figure 2 shows the variance decomposition of BLR in the estimated model at different horizons. At short horizons, the forecast error of BLR is mostly accounted for by both transitory noise and shocks, and the opposite holds at a medium horizon (after, say, 7 years). Given our emphasis on the medium run, we focus on the effect of permanent shocks throughout the paper.

Figure 2: Variance Decomposition of BLR at Different Horizons

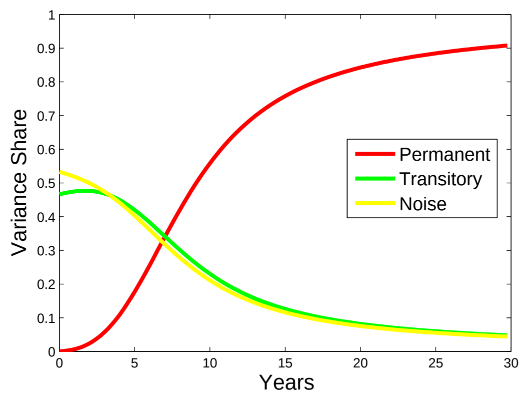

Notes: Percentage of forecast error explained by each shock.

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. Figure 3 shows our estimated permanent technology shocks for the case of the Great Recession.[21] We estimate positive shocks in the early 1990s, up to 1998, and negative shocks in the second part of the sample. Notice that the serial correlation of our estimated permanent shocks is not a violation of the i.i.d. assumption in the model, but instead purely a reflection of the information available to the econometrician. Given the small size of permanent shocks, it difficult to the econometrician to pin point with precision the quarter when each particular shock hits. This introduces an estimation error that it autocorrelated, and the smoothed shocks turn out autocorrelated as well. This has implications for the interpretation of the estimated series. Indeed, there is fairly strong evidence in the data of either a large positive shock or several positive shocks somewhere in the early 90s, although it is not possible to know exactly when. The opposite

[21]For brevity we do not show the estimated transitory and noise shocks here, see the Supplementary Material.

holds starting 1998.[22]

Figure 3: Smoothed Permanent Shocks (U.S., 1990–2013)

Notes: Shocks estimated using a Kalman smoother on the U.S. 1990–2013 sample. The data is composed by the first differences of the logarithm of labor productivity and the ratio of net exports-to-GDP. The unit on the y-axis is percentages. Shocks are scaled by their ML estimated standard deviation.

The estimated permanent shocks in our sample imply that we should have observed a productivity acceleration in the mid-90s, and a subsequent slowdown, arriving some years before the start of the Great Recession. This can be verified by evidence outside our exercise. For instance, in an impressive paper, Fernald (2012a) documents detailed evidence, at different levels of aggregation, that the growth of both labor and total-factor productivity slowed down after 2004 in most industries.[23] The slow down was most pronounced in IT-intensive industries.[24] Consistent with our structural results, we find similar evidence when looking only at annualized productivity growth rates in our sample. These are on average 1.93% from the first quarter of 1990 to the first quarter of 2004 on a yearly basis, and 1.12% from the second quarter of 2004 to the first quarter of 2013.[25]

[22]We have verified that Kalman smoothed shocks out of simulated data have a similar degree of autocorrelation.

[23]See Fernald's Figure 3 (March 24, 2014 version) for a plot of aggregate trend productivity.

[24]See Fernald's Table 1 (March 24, 2014 version).

[25]On its special report on the world economy, *The Economist* also documented a slowdown of GDP per hour worked in the U.S. that started around 2001 (Figure 12). (with data available Oct. 7th 2010)

3.5 Model-predicted Beliefs About the Long Run and Out-of-sample Check

Now we characterize the medium-run dynamics of BLR. For external validity, we compare our results to out-of-sample evidence coming from a survey.

We proceed by a standard historical decomposition as follows. We feed into the model the series of estimated permanent shocks shown in Figure 3, setting the other two shocks η_t and ν_t to zero. We then simulate the associated BLR using our model. Figure 4 shows the resulting BLR. According to this series, the U.S. was relatively most optimistic about his long-run income around 2004. Notice here the impact of the delay in learning. Even though the positive ε_t are estimated to hit the U.S. economy around 1995, it takes quite some time for the agent to become optimistic. In fact, according to our results the agent is most optimistic right when productivity started to slow down, and it also took some time to revise its beliefs downwards. In this model, this decline in beliefs produces a fall in consumption. In a model with elastic labor supply and nominal rigidities this decline would result in a fall in output (see BLL.)[26]

Figure 4: Model-predicted BLR (U.S., 1990–2013)

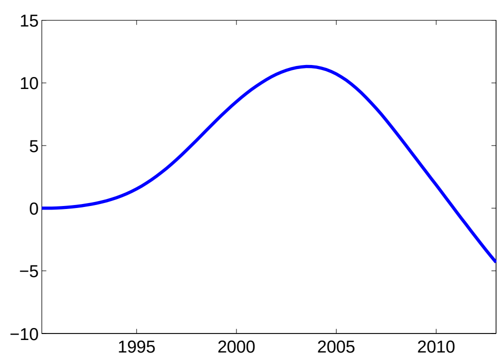

Notes: Historical effect of permanent shocks on BLR.

Consensus Forecasts publishes a survey including a question of participants' expectations of GDP growth up to 10 years ahead. The survey is done in major industrialized economies[27], and this is the longest horizon available in the survey. Figure 5 reproduces from Hoffmann, Krause, and Laubach (2011),

[26] Jaimovich and Rebelo (2008) propose other alternatives to produce comovement in open economy models.

[27] The countries included are the U.S., Japan, Germany, France, the U.K., Italy, Canada, China, Korea and Taiwan.

(p. 6) a series of GDP weighted average answers of these Forecasts of real GDP growth 6–10 years ahead (upper panel, top series, marked with a '+'), a series of U.S. answers of these Forecasts (upper panel, bottom series, marked with an 'x'), and the difference of these Forecasts between the U.S. and the Rest-of-the-World (RoW, bottom panel, unique series, marked with an 'x'). Given that on average the RoW was more optimistic about domestic growth than the U.S., the series on the bottom is negative. This is consistent with the higher average growth rate of countries like China, Korea and Taiwan. From the perspective of our exercise, we are interested in the relative evolution of trend-growth expectations in the U.S. versus the RoW, i.e. the evolution of the series in the bottom panel (unique series, marked with an 'x').

Figure 5: Survey Evidence on Long-run Growth Forecasts, U.S. versus RoW (Lower Panel)

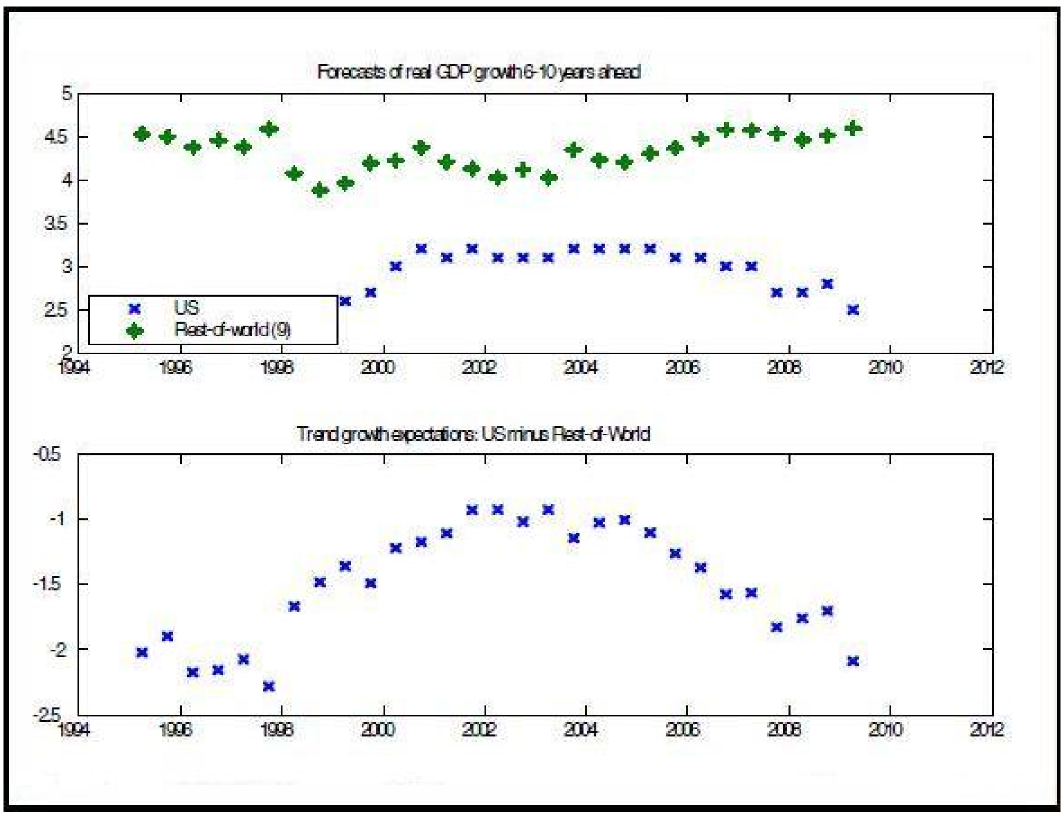

Notes: Reproduced from Hoffman et al. (2011). The upper panel plots the weighted average response in a sample of major industrialized countries countries (Japan, Germany, France, the U.K., Italy, Canada, China, Korea and Taiwan, upper panel, top series, marked with a '+'), and the average response in the U.S (upper panel, bottom series, marked with an 'x'). The lower panel plots the difference between the two (unique series, marked with an 'x').

Figure 6 compares the evolution of trend-growth expectations according to the survey (the bottom panel of Figure 5) and BLR about the long run generated by our estimated model. In the upper panel we plot our beliefs series, aligning the time axis to the Hoffmann et al. (2011) series. A qualitative comparison between the two series suggests that according to both measures the U.S. agent seems to have been relatively most optimistic between 00 and 05.[28]

Figure 6: Out-of-sample Check: Comparison of Model-predicted BLR from Figure 4 (Upper Panel) and Survey Evidence on Long-run Growth Forecasts (Lower Panel)

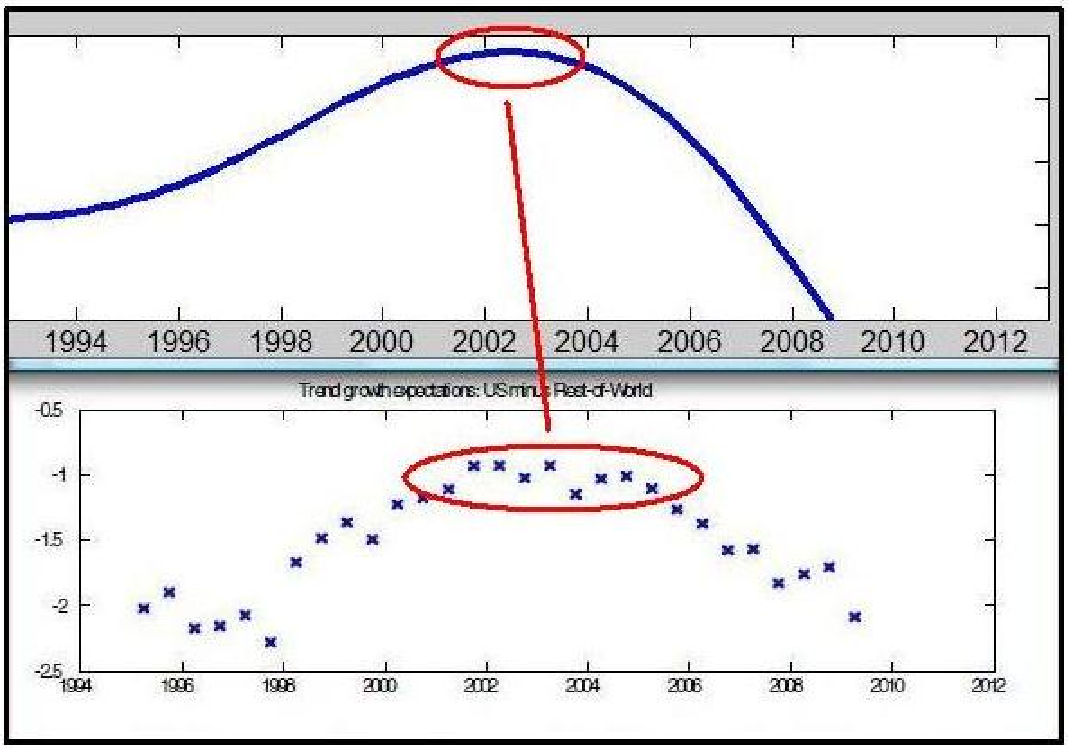

Notes: The survey data is reproduced from Hoffman et al. (2011), see Figure 5.

3.6 Robustness

To assess the robustness of the estimates presented previously and the implied delay, here we report the results of a number of alternative robustness estimations. We explore the implications of alternative sample spans, alternative set of observables, and alternative model specifications (i.e. using the limit model suggested by Proposition 1).

[28] In Subsection 3.3 of their paper, Hoffmann et al. (2011) perform a similar exercise by computing a Kalman filtered trend of productivity and comparing it to the sample. Their exercise and ours complement each other. The most important difference is our use of both productivity and net exports – following the permanent income logic mentioned previously – while they use only productivity (of course not imposing the random walk Assumption 6). Another difference is our computation of the variance decomposition for BLR and focus on permanent shocks for getting at their model-implied medium-run dynamics.

More specifically, first we report in Table 2 the **baseline estimation (1.)**, using the 1990–2013 sample, U.S. data, where observables are labor productivity and net exports. To investigate whether our estimates rely crucially on the sample span of our baseline estimation, we change the sample span to a **1985–2013 sample (2.)**, to a **1980–2013 sample (3.)**, and to the **full 1948–2013 sample (4.)**. In these three estimations we obtain similar parameter estimates (highly persistent technology processes and relatively large noise), implying large delays. In 4. the delay is particularly large (9.75 years). Numerical exercises on the mapping from parameters to the delay suggest that the reason is the even higher estimate of ρ in this case. One issue with this estimation are the low frequency changes in the trend of productivity over this long sample, which is why we feel more confident with our baseline, which does not suffer from this caveat.

Table 2: Robustness

	Estimation	ρ	σ_a	σ_ν	ψ	Delay
1.	Baseline 1990–2013 sample	0.9720 (0.0057)	0.6433 (0.0446)	7.3922 (2.0451)	0.0010	5.25
2.	1985–2013 sample	0.9737 (0.0046)	0.6183 (0.0390)	6.2959 (1.7023)	0.0010	5.50
3.	1980–2013 sample	0.9777 (0.0034)	0.7020 (0.0415)	9.0652 (2.1890)	0.0010	6.50
4.	Full 1948–2013 sample	0.9893 (0.0010)	1.0999 (0.0544)	10.6144 (3.2643)	0.0010	9.75
5.	Estimating ψ	0.9659 (0.0099)	0.5604 (0.0382)	10.0022 (2.4024)	0.0001 (0.0002)	4.75
6.	TFP	0.9739 (0.0056)	0.6918 (0.0483)	9.3421 (2.4848)	0.0010	6.00
7.	Consumption (limit model Prop. 1)	0.9736 (0.0073)	0.5254 (0.0303)	2.0509 (0.6153)	$\longrightarrow 0$	3.75

Notes: To assess the robustness of the estimates, this table presents several ML estimates changing the time-span of the sample, the set of observables, and the exact model specification. Standard errors are reported in parenthesis below the point estimate. The values of σ_ε and σ_η are computed using ρ and σ_u through the random walk assumption (6). The delay was defined in p. 9. For estimation 7., we set $C/Y = 1$. All estimates are for the U.S.

We then **estimate ψ (5.)** as well, in order to check that our baseline calibration (in which we follow the literature) is consistent with the data. We find that the data favor an even smaller value of ψ than the one adopted throughout the paper. This suggests that the assumptions of Proposition 1 are empirically valid.

We then investigate how a different set of observables changes our results. First, we estimate the model using the logarithm of **TFP (6.), using Fer-**

nald's utilization-adjusted series in first differences (Fernald 2012b). We find similar results to 1. Second, we estimate the model using the first differences of the logarithm of **consumption (7.)**, using the exact expression for consumption derived in Proposition 1. We find a smaller delay of learning, but still important, of 3.75 years.[29]

To sum up, over a range of estimation exercises we obtain similar results for our parameter estimates ρ, σ_a, and σ_ν. The estimates for ρ are quite stable at high values, suggesting a smooth and persistent permanent component. Not surprisingly, the estimates for σ_a and σ_ν change across samples due to changes in the variance of productivity and net exports (consumption) across these samples, but in all cases we obtain a large standard deviation of noise shocks σ_ν, when compared to σ_a. Finally, the delay is computed to be strictly positive[30] and sizeable in all cases.

3.7 Japan and Great Depression

Here we present our baseline results for the case of Japan (1975–2005) and the Great Depression in the U.S. (1920–1935).

Table 3 contains the parameter estimates. The persistence parameter ρ is estimated at 0.94 in the case of Japan, and at 0.86 in the case of the Great Depression. Both values imply persistent processes both for the permanent and the transitory components of productivity. The standard deviation of productivity is estimated at 1.00% in the case of Japan, and at 1.66% in the case of the Great Depression. These values are considerably larger than the ones obtained for the Great Recession. Given the random walk assumption (6) for productivity, these values imply a standard deviation for permanent technology shocks of 0.06% in the case of Japan, and of 0.24% in the case of the Great Depression, and a standard deviation for the transitory technology shock of 0.97% in the case of Japan, and of 1.53% in the case of the Great Depression. The standard deviation of noise shocks is large, 14.49% and 20.05% respectively.

The standard deviations of all shocks and, in particular, the noise shock are larger in both cases than in the Great Recession. However, the overall amount of noise in the news agents receive – quantified by the delayed learning – is actually smaller. This delay is of 2.50 years in the case of Japan, and of 1

[29] In the estimations of the limit model we set $C/Y = 1$.

[30] Perfect information is a special case of this model when $\sigma_\nu = 0$, in which the delay is equal to zero. The estimation tells us that this case fails to fit the data.

Table 3: Parameter Estimates, Japan and Great Depression

		Japan		Great Dep.	
Parameter	Description	Value	s.e.	Value	s.e.
ρ	Persistence tech. shocks	0.94	0.02	0.86	0.05
σ_a	Std. dev. productivity	1.00	0.06	1.66	0.15
σ_ε	Std. dev. permanent tech. shock (implied)	0.06	–	0.24	–
σ_η	Std. dev. transitory tech. shock (implied)	0.97	–	1.53	–
σ_ν	Std. dev. noise	14.49	3.50	20.05	8.06

Notes: ML estimates of the log-linearized state-space representation of the model. The observation equation is composed of the first differences of the logarithm of labor productivity and the ratio of net exports-to-GDP. In the case of Japan, the sample spans 1975–2005. In the case of the Great Depression, the sample spans 1920–1935 (due to data availability it does not start earlier). Standard errors are reported to the right of the point estimate. The standard deviations of the permanent and transitory shocks are implied by the random walk assumption (6) for productivity.

year in the case of the Great Depression. It is difficult to analyze the mapping between the parameter values to the delay, but from numerical exercises it seems that the main reason the delay is smaller here is that in both cases, and especially in the case of the Great Depression, the value of ρ is smaller, implying a less persistent permanent technology process and, through Assumption (6), larger and easier to detect permanent shocks. It is not quite surprising that we estimate a smaller delay in these two cases, since as the plots in Section 4 below show, consumption seems to have reacted a bit less slowly in these two cases.[31]

The variance decomposition for these cases are similar to the one shown in Figure 2 and for brevity we do not show them. Figure 7 plots the estimated permanent shocks. As for the Great Recession, we estimate positive shocks in the first part of the two samples, and negative shocks later on. In the case of Japan, the positive shocks hit roughly between 1980 and 1987. This estimated permanent shocks imply that we should have observed a productivity acceleration and deceleration. Consistently, Japanese annualized growth rates of productivity averaged to 3.22% between 1975 and 1990, and 1.14% from then on. In the case of the Great Depression, the positive shocks hit roughly between 1920 and 1922, the negative shocks roughly between 1926 and 1932, and then again positive shocks hit starting 1932, probably related to the strong economic recovery that started around 1933. In the later case our sample does not seem to start early enough (due to data availability) to appreciate the full extent of

[31]Specifically, the distance between the peak and the through of the productivity-to-consumption ratio is smaller in Japan and Great Depression (11 and 9 years) when compared to the Great Recession (15 years).

the productivity pickup related to the 2nd Industrial Revolution, because the range that mostly contains positive shocks is rather short. Looking at the dates in which some of the technological innovations were implemented – for instance the Ford Model-T was introduced in 1908 – suggest that one would like to have a reliable sample for quarterly consumption and productivity starting at least 10 years before 1920. Still, starting in 1920 captures some of the trend productivity increases of the period. Consistently, annualized productivity growth rates for the U.S. economy average 2.75% between 1920 and 1925, and drop to -.48% between 1925 and 1933. Productivity growth recovers later, between 1933 and 1935, to 4.59%.[32]

Figure 7: Smoothed Permanent Shocks (Japan 1975–2005, and U.S. 1920–1935)

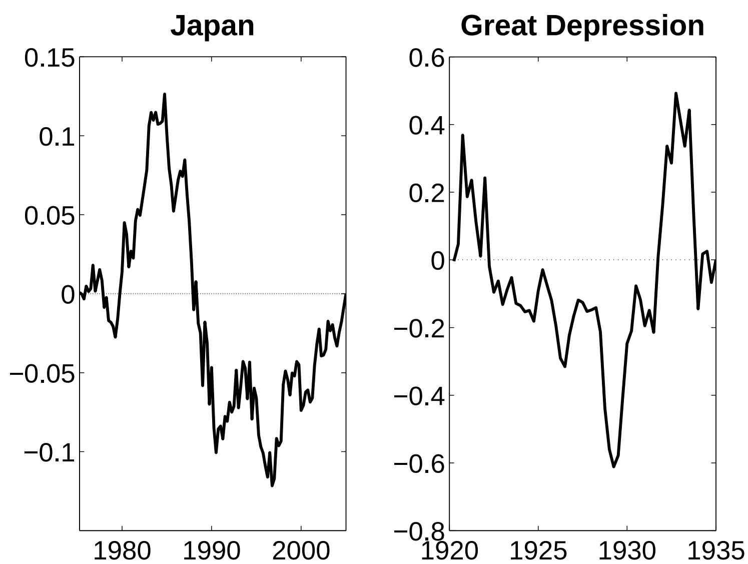

Notes: Shocks estimated using a Kalman smoother on the Japanese 1975–2005 sample, and on the U.S. 1920–1935 sample. The latter is restricted by data availability. The data is composed by the first differences of the logarithm labor productivity and the ratio of net exports-to-GDP. The time unit on the x-axis is percentages. Shocks are scaled by their ML estimated standard deviation.

Figure 8 plots the model-predicted BLR for Japan and the Great Depression. Even though the positive shocks seem to have hit the Japanese economy mostly in the mid-1980s, consumers there seem to have been most optimistic around 1990. In the case of the Great Depression, the consumer is most optimistic around 1923, which implies a shorter delay with respect to the positive permanent shocks. The reason is the smaller delay in learning.

[32]Productivity growth rates seem to have been high for a number of years during the recovery from the Great Recession, a fact noted by Field (2003), among others.

Figure 8: Model-predicted BLR (Japan 1975–2005, and U.S. 1920–1935)

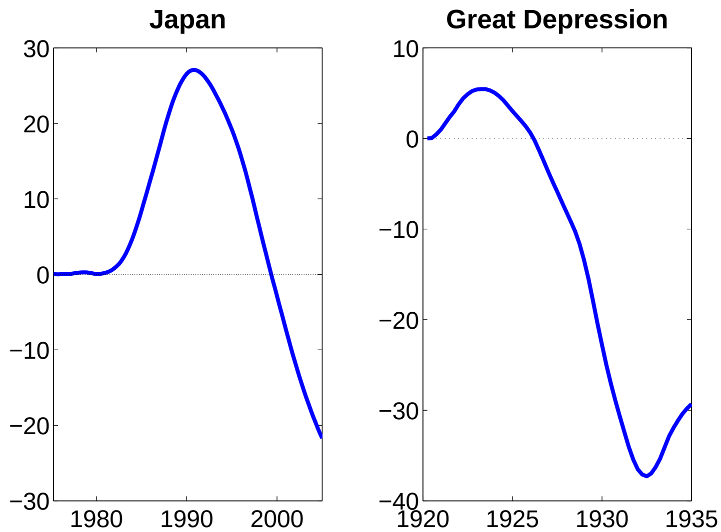

Notes: Historical effect of permanent shocks on BLR.

To sum up, even though the exact measure of the delay varies from episode to episode, we do find structural evidence of this delay also in the cases of Japan and the Great Depression. There is also evidence of an increase in the trend of productivity, followed by a decrease in this trend. To shed light at these and other facts, the next section looks directly at the ratio of productivity-to-consumption.

4 The Productivity-to-Consumption Ratio in the Data

In order to understand which feature of the data deliver the results above, here we focus on the shape of the productivity-to-consumption ratio in the three cases. As explained above, this ratio contains information about consumers' beliefs about their future income according to the permanent income hypothesis. We do *not* detrend any of these series for these plots.

Great Recession. Figure 9 plots the logarithm of the ratio of productivity-to-consumption around the Great Recession (U.S. 1990–2013). The vertical axis is centered around the average of the ratio over the period considered. The trend of this series computed using an HP-Filter ($\lambda = 800$) is also plotted. Using a band pass filter isolating the medium-run frequency of this series (between 32 and 200 quarters, following Comin and Gertler 2006) delivers a similar shape for this trend.

Figure 9: Productivity-to-consumption ratio, in logs (U.S., 1990–2013), and Trend

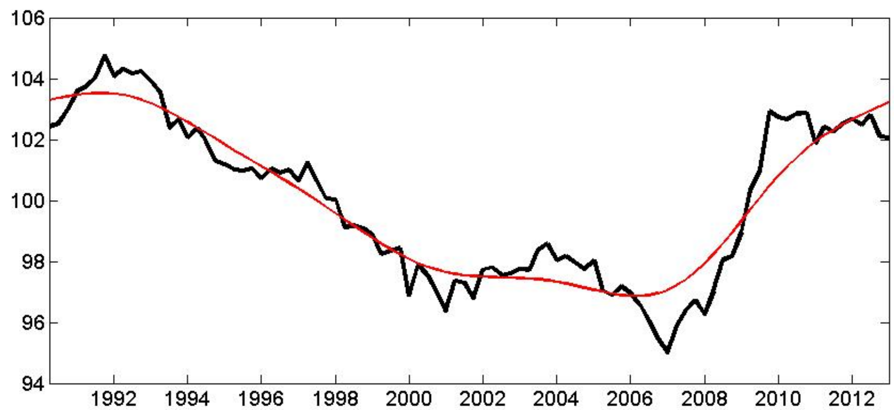

Notes: Productivity is real GDP divided by employment. Consumption is NIPA consumption divided by population. Neither series is detrended. The trend is computed with an HP-Filter ($\lambda = 800$).

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 reversal starts in which the ratio quickly goes back to its level from 20 years earlier. The reversal is quite sharp and coincides with the start of the Great Recession in 2007. Overall, the ratio appears to follow a slow-moving "up-and-down" wave.

To shed light on these dynamics, it is useful to considered two theoretical benchmarks.

Benchmark (a): "No-news". In this case, σ_ν tends to infinity and thus the signal is completely uninformative. Given the random walk assumption (6), BLR are

$$x_{t+\infty|t} = a_t \quad ,$$

and so, under the conditions of Proposition 1, consumption is equal to productivity:

$$c_t = a_t, \quad \forall t \quad .$$

Thus, the ratio of productivity-to-consumption is flat. As illustrated by Figure 10 (left panel), this clearly fails to fit the data.

Figure 10: Benchmarks for the Productivity-to-consumption Ratio

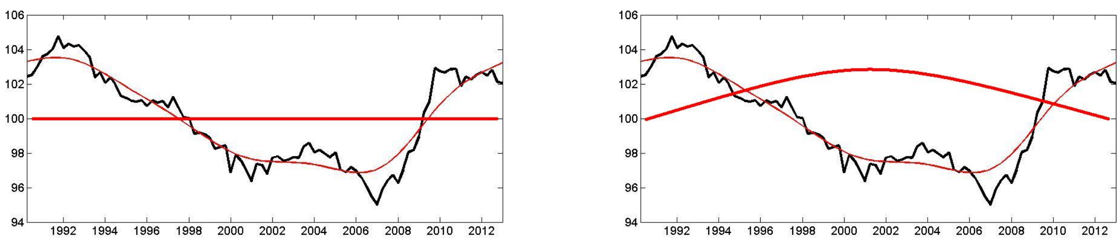

(a) No-news

(b) Perfect Foresight

Notes: Ratio for the U.S. 1990–2013, HP-filter trend ($\lambda = 800$), and theoretical benchmarks.

Benchmark (b): Perfect Foresight. Under perfect foresight, agents have knowledge of all future shocks right from 1990:Q1. Under the conditions of Proposition 1, consumption jumps immediately to the long-run level of productivity, say $x_{t+\infty}$, and remains there. As a result of the positive and then negative permanent shocks, productivity first increases and then decreases, and then stays there. The ratio of productivity and consumption, thus, has the same dynamics: it increases, then decreases, and then stay there. As illustrated by Figure 10 (right panel), this, again, fails to fit the data.

To conclude, in both the "no-news" and the perfect foresight benchmarks, the model has a strongly counterfactual prediction for the behavior of the productivity-to-consumption ratio. Indeed, in the data, the ratio finishes in a U-shaped motion. As explained above, noisy signals ($\sigma_\nu > 0$ but finite) imply a delay in learning that helps accommodate this behavior of the ratio, i.e. its decline as consumption catches up with the productivity increase, and rise when consumption growth slows down.

Japan. Figure 11 plots the same ratio for Japan. In this case we can see a more gradual increase in the ratio from its average over the period considered,

Figure 11: Productivity-to-consumption ratio, in logs (Japan, 1975–2005), and Trend

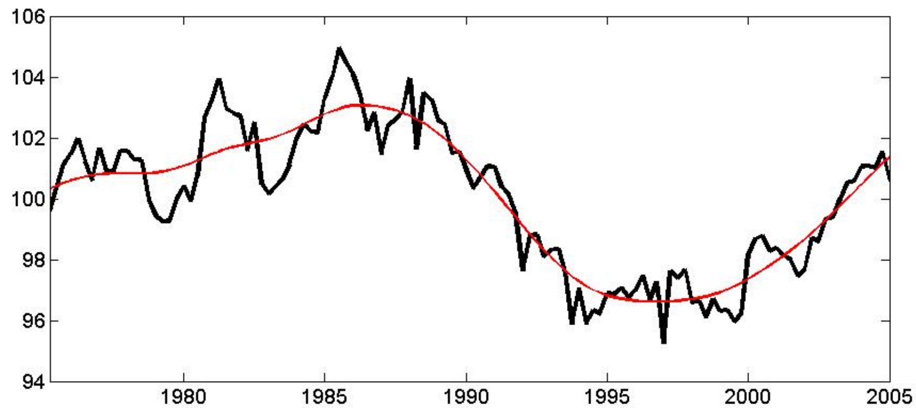

Notes: Productivity is real GDP divided by employment. Consumption is NIPA consumption divided by population. Neither series is detrended. The trend is computed with an HP-Filter ($\lambda = 800$).

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 up-and-down wave.

Great Depression. Figure 12 plots the ratio for the Great Depression. Due to data availability, we look at this data starting 1920. 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.

To summarize, this reduced-form analysis complements the results obtained through structural estimation. In the three cases considered, the productivity-to-consumption ratio appears to follow similar medium-term dynamics. Together with the evidence on productivity growth rates presented in Section 3, the overall conclusion is that in the three cases there was a slow-moving boom of aggregate productivity, followed by a slowdown. Furthermore, consumption features similar dynamics, but adjusts with a significant lag.[33]

[33]Flemming and L'Huillier (2014) briefly analyze the relationship of the productivity-to-consumption ratio to new assets agency issuance (MBSs and CMOs) and find a positive correlation between the two.

Figure 12: Productivity-to-consumption ratio, in logs (U.S., 1920–1935), and Trend

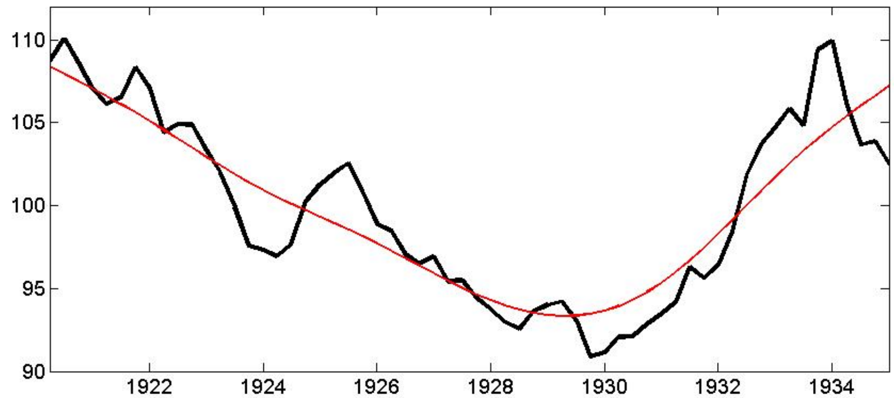

Notes: Productivity is real GDP divided by employment. Consumption is NIPA consumption divided by population. Neither series is detrended. The trend is computed with an HP-Filter ($\lambda = 800$). The sample starts in 1920 due to data availability.

5 Characterization of the Dynamics of Debt

In this section we study the model-predicted dynamics of debt, that is, the dynamics implied by the estimated permanent shocks shown in Figure 3. For brevity, we do this only for the case of the Great Recession.[34]

Figure 13 plot these dynamics. The left panel shows productivity, the center panel shows net exports and the right panel shows the debt-to-output ratio. Vertical axes are percentage deviations for steady state. Productivity increases and then decreases, the peak happening in the early 2000s. Net exports are first slightly positive, then turn negative, and turn positive around 2008. When net exports are negative the economy accumulates debt, with the debt-to-output ratio reaching its highest point around 2008.[35]

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. The higher ρ, the larger the long-run effect of a shock ε_t, and the larger the income effect. The larger the income effect, the larger the accumulation of debt. Second, the dynamics of debt depend on the relative size of the standard deviations σ_ε, σ_η, and σ_ν, because these determine the

[34] It is possible to write a closed economy model with borrowing and lending based on Iacoviello (2005) or Iacoviello and Pavan (2013) that features the same information structure. However, given our emphasis on permanent income consumption, a characterization of debt dynamics in such a framework is out of the scope of this paper.

[35] A close inspection of equation (16) 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 13: Model-predicted Productivity, Net Exports, and Debt-to-output Ratio

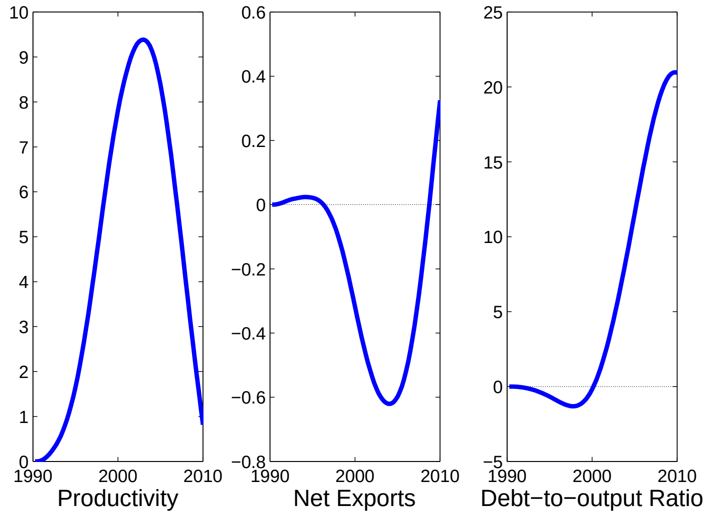

Notes: Historical effect of permanent shocks on productivity a_t, the net exports-to-output ratio nx_t, and the debt-to-output ratio b_t. On the left panel, vertical axis' units are relative percentage deviations from the steady state. On the mid- and right panels, vertical axes' units are absolute percentage deviations from the steady state.

informativeness of a_t and s_t as signals about the permanent component x_t, and thus the speed of learning. The smaller σ_ε with respect to the other two, the less informative a_t and s_t, the slower learning, and the longer it takes for beliefs and consumption to adjust. Third, the dynamics are also 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".

We stop this simulation in 2010 to make the following qualitative point. At this point the state of the U.S. economy has three adverse features: high debt, low productivity growth, and pessimistic expectations. Debt is high because it took some time for agents to recognize the productivity slowdown of the early 2000s. Productivity growth is low because of the negative permanent shocks that hit the economy after 1998. Expectations are pessimistic because agents have learnt about the decline in the trend of productivity. The latter of these ingredients implies deleveraging. Given the medium-run perspective of our exercise, it is interesting to quantify the length of the deleverage by simulating the model forward from 2010, assuming all shocks are equal to zero after the end of our sample (2013:Q1). Figure 14 shows the results. According to this simulation, it would take 7 years after 2010 for the debt-to-output ratio to return to steady state. Of course, the model is too simple to take this quantitative prediction seriously, for instance, the aggressive monetary easing after 2008 is not taken into account, among other factors. But, at least qualitatively, the medium-run nature of the deleverage predicted by these results is quite suggestive.

Figure 14: Simulation of the Debt-to-output Ratio After 2010

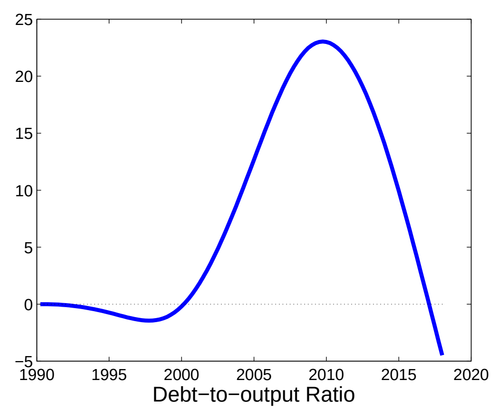

Notes: Model-produced forecast of b_t assuming all shocks after 2013 (the end of our sample) are zero.

6 Conclusion

We have explored the movements of productivity and consumption before and during the Great Recession, the Japanese crisis of the 1990s, and the Great Depression. In the three cases, productivity and consumption feature common medium-run dynamics which can be accommodated by a learning model.

Our inference is based on the permanent income logic together with rational expectations (Blundell and Preston 1998; Blundell, Pistaferri, and Preston 2008). It allows us to see through the medium-run dynamics of productivity and consumption by using structural estimation. We find similar results by looking at the data directly in reduced form.

The model features noisy news about the future. In the model, an exogenous process for productivity is the sum of a permanent and a transitory component. The decomposition is estimated by the observation of consumption, which by the permanent income hypothesis is determined by beliefs about long-run income. Having performed the decomposition of productivity, we find that consumers seem to form their beliefs with a delay. This delay accommodates an observed lagged behavior of consumption with respect to movements in the permanent component of productivity. We find the predictions of the model intuitive, and capable to provide a simple account of the behavior of consumption in these episodes. Altogether, this is also a useful exercise to understand the build-up of debt and the deleveraging process in the three cases.

References

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

Angeletos, G.-M. and J. La'O (2009). Noisy business cycles. *NBER Macroeconomics Annual 24*(319–378).

Angeletos, G.-M. and J. La'O (2013). Sentiments. *Econometrica 81*(2), 739–780.

Bansal, R. and A. Yaron (2004). Risks for the long run: a potential resolution of asset pricing puzzles. *Journal of Finance LIXI*(4), 1481–1509.

Bardou, J.-P., J.-J. Chanaron, P. Fridenson, and J. M. Laux (1982). *The automobile revolution: The impact of an industry*. University of North Carolina Press.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics 121*(3), 823–866.

Barro, R. J. and J. F. Ursua (2011). Rare macroeconomic disasters. *NBER Working Paper 17328*.

Barsky, R. B. and E. Sims (2012). Information, animal spirits, and the meaning of innovations in consumer confidence. *American Economic Review 102*(2), 1343–1377.

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 (2013). News, noise, and fluctuations: An empirical exploration. *American Economic Review 103*(7), 3045–3070.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review 98*(5), 1887–1921.

Blundell, R. and I. Preston (1998). Consumption inequality and income uncertainty. *Quarterly Journal of Economics 113*(2), 603–640.

Boz, E. (2009). Can miracles lead to crises? The role of optimism in emerging market crises. *Journal of Money, Credit, and Banking 41*(6), 1189–1215.

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

Campbell, J. Y. (1987). Does saving anticipate declining labor income? an alternative test of the permanent income hypothesis. *Econometrica 55*(6), 1249–1273.

Cao, D. and J.-P. L'Huillier (2014a). The behavior of consumption in the RBC open economy model. *Mimeo*.

Cao, D. and J.-P. L'Huillier (2014b). Technological revolutions and debt hangovers: Is there a link? *Mimeo*.

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.

Comin, D. and M. Gertler (2006). Medium-term business cycles. *American Economic Review 96*(3), 523–551.

David, P. and G. Wright (2000). *General purpose technologies and surges in productivity: Historical reflections on the future of the ICT revolution.* In P. A. David and M. Thomas, (Eds.). *The economic future in historical perspective, Oxford University Press for the British Academy.*

Eggertsson, G. B. (2008). Great Expectations and the end of the Great Depression. *American Economic Review 98*(4), 1476–1516.

Fernald, J. (2012a). Productivity and potential output before, during, and after the great recession. *Federal Reserve Bank of San Francisco Working Paper 2012–18*.

Fernald, J. (2012b). A quarterly, utilization-adjusted series on total factor productivity. *Federal Reserve Bank of San Francisco Working Paper 2012-19*.

Field, A. J. (2003). The most technologically progressive decade of the century. *American Economic Review 93*(4), 1399–1413.

Flemming, J. and J.-P. L'Huillier (2014). Financial innovation: Fundamental cause of propagation mechanism in the 2008 financial crisis? *Mimeo*.

Forni, M., L. Gambetti, M. Lippi, and L. Sala (2013a). Noise bubbles. *Mimeo*.

Forni, M., L. Gambetti, M. Lippi, and L. Sala (2013b). Noisy news in business cycles. *Mimeo*.

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 stoch market. *American Economic Review Papers and Proceedings 89*(2), 116–122.

Guvenen, F. (2007). Learning your earning: are labor income shocks really very persistent? *American Economic Review 97*(3), 687–712.

Hobijn, B. and B. Jovanovic (2001). The information-technology revolution and the stock market: Evidence. *American Economic Review 91*(5), 1203–1220.

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

Iacoviello, M. (2005). House prices, borrowing constraints and monetary policy in the business cycle. *American Economic Review 95*(3), 739–764.

Iacoviello, M. and M. Pavan (2013). Housing and debt over the life cycle and over the business cycle. *Journal of Monetary Economics* (60), 221–238.

Jaimovich, N. and S. Rebelo (2008). News and business cycles in open economies. *Journal of Money, Credit and Banking 40*(8), 1699–1711.

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

Justiniano, A., G. Primiceri, and A. Tambalotti (2010). Investment shocks and business cycles. *Journal of Monetary Economics 57*(2), 132–145.

Justiniano, A., G. Primiceri, and A. Tambalotti (2014a). Credit supply and the housing boom. *Mimeo*.

Justiniano, A., G. Primiceri, and A. Tambalotti (2014b). Household leveraging and deleveraging. *NBER Working Paper No. 18941*.

Kendrick, J. W. (1961). *Productivity Trends in the United States*. UMI.

Mian, A. R. and A. Sufi (2012). What explains high unemployment? The aggregate demand channel. *Chicago Booth Research Paper No. 13-43*.

Midrigan, V. and T. Philippon (2011). Household leverage and the recession. *Working Paper, NYU*.

Nakamura, E., J. Steinsson, R. J. Barro, and J. Ursua (2013). Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics 5*(3), 35–74.

Neumeyer, P. A. and F. Perri (2004). Business cycles in emerging economies: the role of interest rates. *Journal of Monetary Economics 52*(2), 345–380.

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

Piazza, R. Growth and crisis: unavoidable connection? *Review of Economic Dynamics, forthcoming*.

Pintus, P. A. and J. Suda (2013). Learning leverage shocks and the Great Recession. *Mimeo*.

Pintus, P. A. and Y. Wen (2013). Leveraged borrowing and boom-bust cycles. *Review of Economic Dynamics 16*(4), 617–633.

Reinhart, C. M. and V. R. Reinhart (2010). After the fall. *Macroeconomic challenges: the decade ahead*. Federal Reserve Bank of Kansas City Jackson Hole Symposium, August 26-28, 2010.

Reinhart, C. M. and K. S. Rogoff (2008). This time is different: A panoramic view of eight centuries of financial crises. *NBER WP 13882*.

Reinhart, C. M. and K. S. Rogoff (2011). *This time is different: Eight centuries of financial folly.* Princeton: Princeton University Press.

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

Syverson, C. (2013). Will history repeat itself? Comments on "Is the information technology revolution over?". *International Productivity Monitor 25.*

A Beliefs About the Long Run

Proof. To prove (8), we make use of equations (1), (2), and (3):

$$\mathbb{E}_t\left[a_{t+\tau}\right] = \mathbb{E}_t\left[x_{t+\tau} + z_{t+\tau}\right] \quad .$$

It follows that

$$\begin{aligned}
\mathbb{E}_t\left[x_{t+\tau}\right] &= \mathbb{E}_t\left[x_t + \sum_{\tau'=1}^{\tau-1}\left(x_{t+\tau'+1} - x_{t+\tau'}\right)\right] \\
&= \mathbb{E}_t\left[x_t + \sum_{\tau'=1}^{\tau-1}\left\{\rho^{\tau'+1}\left(x_t - x_{t-1}\right) + \sum_{\tau''=1}^{\tau'}\rho^{\tau''}\varepsilon_{t+\tau''}\right\}\right] \\
&= \mathbb{E}_t\left[x_t + \left(x_t - x_{t-1}\right)\sum_{\tau'=1}^{\tau-1}\rho^{\tau'+1}\right]
\end{aligned}$$

where the last inequality comes from the fact that $\mathbb{E}_t\left[\varepsilon_{t+\tau''}\right] = 0$ for all $\tau'' \geq 1$. The geometric sum $\sum_{\tau'=1}^{\tau-1}\rho^{\tau'+1}$ simplifies to $\rho\frac{1-\rho^{\tau}}{1-\rho}$. So

$$\mathbb{E}_t\left[x_{t+\tau}\right] = \mathbb{E}_t\left[x_t + \rho\frac{1-\rho^{\tau}}{1-\rho}\left(x_t - x_{t-1}\right)\right] \quad .$$

Now taking the limit of τ to infinity and noticing that $\lim_{\tau\longrightarrow\infty}\rho^{\tau} = 0$ we obtain

$$\begin{aligned}
\lim_{\tau\longrightarrow\infty}\mathbb{E}_t\left[x_{t+\tau}\right] &= \mathbb{E}_t\left[x_t + \rho\frac{1}{1-\rho}\left(x_t - x_{t-1}\right)\right] \\
&= \frac{\mathbb{E}_t\left[x_t - \rho x_{t-1}\right]}{1-\rho}.
\end{aligned}$$

Similarly,

$$\mathbb{E}_t\left[z_{t+\tau}\right] = \rho^{\tau}\mathbb{E}_t\left[z_t\right] \quad .$$

So

$$\lim_{\tau\longrightarrow\infty}\mathbb{E}_t\left[z_{t+\tau}\right] = 0 \quad .$$

Combining the two limits for $\mathbb{E}_t\left[x_{t+\tau}\right]$ and $\mathbb{E}_t\left[z_{t+\tau}\right]$, we obtain equality (8). ■

B Closed-form Solution and Limit Result for Consumption

In this section we solve the model in closed form. Let

$$\widehat{b}_t = b_t + \frac{1 - C/Y}{1-\beta} a_t \quad .$$

From the intertemporal budget constraint (16), together with the budget constraint (15), we have:

$$\begin{aligned}
\widehat{b}_t &= b_t + \frac{1 - C/Y}{1-\beta} a_t \\
&= \frac{1}{\beta} b_{t-1} - \frac{1}{\beta}\frac{C}{Y}(-c_t) + \frac{1}{\beta}\frac{1 - C/Y}{1-\beta}(-\Delta a_t + \beta r_t) \\
&\quad + \frac{1 - C/Y}{1-\beta} a_t \\
&= \frac{1}{\beta}\widehat{b}_{t-1} - \frac{1}{\beta}\frac{C}{Y}(-c_t) + \frac{1}{\beta}\frac{1 - C/Y}{1-\beta}(-a_t + \beta r_t) \\
&\quad + \frac{1 - C/Y}{1-\beta} a_t \\
&= \frac{1}{\beta}\widehat{b}_{t-1} + \frac{1}{\beta}\frac{C}{Y} c_t - \frac{1 - C/Y}{\beta} a_t + \frac{1 - C/Y}{1-\beta} r_t
\end{aligned}$$

Substituting r_t from (14) into the last equality, and also using the definition of $\widehat{c}_t$, we arrive at

$$\begin{aligned}
\widehat{b}_t &= \frac{1}{\beta}\widehat{b}_{t-1} + \frac{1}{\beta}\frac{C}{Y} c_t - \frac{1 - C/Y}{\beta} a_t + \frac{1 - C/Y}{1-\beta} \psi b_t \\
&= \frac{1}{\beta}\widehat{b}_{t-1} + \frac{1}{\beta}\frac{C}{Y} c_t - \frac{1 - C/Y}{\beta} a_t \\
&\quad + \frac{1 - C/Y}{1-\beta} \psi \left(\widehat{b}_t - \frac{1 - C/Y}{1-\beta} a_t \right) \\
&= \frac{1}{\beta}\widehat{b}_{t-1} + \frac{1}{\beta}\frac{C}{Y} \widehat{c}_t - \frac{1}{\beta} a_t \\
&\quad + \frac{1 - C/Y}{1-\beta} \psi \left(\widehat{b}_t - \frac{1 - C/Y}{1-\beta} a_t \right) \quad .
\end{aligned}$$

So

$$
\begin{aligned}
\widehat{b}_t \left(1 - \frac{1 - C/Y}{1 - \beta}\psi\right) &= \frac{1}{\beta}\widehat{b}_{t-1} + \frac{1}{\beta}\frac{C}{Y}\widehat{c}_t \\
&\quad - \left(\frac{1}{\beta} - \psi\left(\frac{1 - C/Y}{1 - \beta}\right)^2\right) a_t
\end{aligned}
$$

From the Euler equation (13), we have

$$
\begin{aligned}
\widehat{c}_t &= -\psi b_t + \mathbb{E}_t[\widehat{c}_{t+1}] \\
&= -\psi \widehat{b}_t + \psi \frac{1 - C/Y}{1 - \beta} a_t + \mathbb{E}_t[\widehat{c}_{t+1}] \quad .
\end{aligned}
$$

Again we conjecture that

$$
\widehat{c}_t = D_b \widehat{b}_{t-1} + D_k \mathbf{X}_t \quad ,
$$

where the state variable $\mathbf{X}_t$ is defined in the proof above and solve for the coefficients D_b and D_k using the method of undetermined coefficients.
Indeed, from the Euler equation:

$$
\begin{aligned}
\widehat{c}_t &= -\psi \widehat{b}_t + \psi \frac{1 - C/Y}{1 - \beta} a_t + \mathbb{E}_t[\widehat{c}_{t+1}] \\
&= -\psi \widehat{b}_t + \psi \frac{1 - C/Y}{1 - \beta} a_t + \mathbb{E}_t[D_b \widehat{b}_t + D_k \mathbf{X}_{t+1}] \\
&= (D_b - \psi)\, \widehat{b}_t + \psi \frac{1 - C/Y}{1 - \beta} a_t + \mathbb{E}[D_k \mathbf{X}_{t+1}] \\
&= (D_b - \psi)\, \frac{1}{1 - \frac{1 - C/Y}{1 - \beta}\psi} \left(\begin{array}{c} \frac{1}{\beta}\widehat{b}_{t-1} + \frac{1}{\beta}\frac{C}{Y}\widehat{c}_t \\ - \left(\frac{1}{\beta} - \psi\left(\frac{1 - C/Y}{1 - \beta}\right)^2\right) a_t \end{array} \right) \\
&\quad + \psi \frac{1 - C/Y}{1 - \beta} a_t + D_k \mathbf{A} \mathbf{X}_t \quad .
\end{aligned}
$$

Where the second equality comes from applying the conjectured solution for c_{t+1}, the dynamics of shocks, and the formula for the Kalman filter presented in BLL Appendix 5.1, from which we have

$$
\mathbb{E}_t[\mathbf{X}_{t+1}] = \mathbf{A} \mathbf{X}_t \quad ,
$$

where

$$\mathbf{A}=\begin{pmatrix}0 & 1+\rho & -\rho & \rho\\ 0 & 1+\rho & -\rho & 0\\ 0 & 1 & 0 & 0\\ 0 & 0 & 0 & \rho\end{pmatrix}.$$

So

$$\begin{aligned}&\left(1-(D_b-\psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\frac{C}{Y}\right)\widehat{c}_t\\ =\quad&(D_b-\psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\widehat{b}_{t-1}\\ &-(D_b-\psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\left(\frac{1}{\beta}-\psi\left(\frac{1-C/Y}{1-\beta}\right)^2\right)a_t\\ &+\psi\frac{1-C/Y}{1-\beta}a_t+D_k\mathbf{A}\mathbf{X}_t\quad.\end{aligned}$$

Comparing coefficient-by-coefficent to the initial conjecture of $\widehat{c}_t$, we obtain the system of equations on D_b and D_k:

$$(D_b-\psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}=\left(1-(D_b-\psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\frac{C}{Y}\right)D_b$$

and

$$\begin{aligned}&(D_b-\psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\left(\frac{1}{\beta}-\psi\left(\frac{1-C/Y}{1-\beta}\right)^2\right)\begin{pmatrix}1 & 0 & 0 & 0\end{pmatrix}\\ &+\left(1-(D_b-\psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\beta\psi}\frac{1}{\beta}\frac{C}{Y}\right)D_k\\ =\quad&\psi\frac{1-C/Y}{1-\beta}\begin{pmatrix}1 & 0 & 0 & 0\end{pmatrix}\\ &+D_k\mathbf{A}\end{aligned}$$

The first equation is a quadratic equation in D_b:

$$D_b^2+\left(\frac{1}{C/Y}-\left(1-\frac{1-C/Y}{1-\beta}\psi\right)\beta\frac{1}{C/Y}-\psi\right)D_b-\psi\frac{1}{C/Y}=0\quad.$$

This equation has two roots, but we pick the negative root to ensure the stability of the dynamic system:

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

Given D_b, we solve for the coefficients D_k using the second equation. First, the coefficient on a_t:

$$\begin{aligned}
&(D_b - \psi)\frac{1}{1 - \frac{1-C/Y}{1-\beta}\psi}\left(\frac{1}{\beta} - \psi\left(\frac{1-C/Y}{1-\beta}\right)^2\right) \\
&+ \left(1 - (D_b - \psi)\frac{1}{1 - \frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\frac{C}{Y}\right)D_{k,1} \\
= \; &\psi\frac{1-C/Y}{1-\beta}
\end{aligned}$$

So

$$D_{k,1} = \frac{\psi\frac{1-C/Y}{1-\beta} - (D_b - \psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\left(\frac{1}{\beta} - \psi\left(\frac{1-C/Y}{1-\beta}\right)^2\right)}{\left(1 - (D_b - \psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\frac{C}{Y}\right)}.$$

The coefficient on $z_{t|t}$:

$$\begin{aligned}
&\left(1 - (D_b - \psi)\frac{1}{1 - \frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\frac{C}{Y}\right)D_{k,4} \\
= \; &\rho D_{k,1} + \rho D_{k,4}
\end{aligned}$$

so

$$D_{k,4} = \frac{\rho D_{k,1}}{1 - (D_b - \psi)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\frac{C}{Y} - \rho}.$$

The coefficients on $x_{t|t}$ and $x_{t-1|t}$:

$$\begin{aligned}
&\left(1 - (D_b - \psi)\frac{1}{1 - \frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\frac{C}{Y}\right)D_{k,2} \\
= \; &(1+\rho)D_{k,1} + (1+\rho)D_{k,2} + D_{k,3}
\end{aligned}$$

and

$$\begin{aligned}
&\left(1-\left(D_b-\psi\right)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\frac{C}{Y}\right)D_{k,3} \\
= &\quad -\rho D_{k,1}-\rho D_{k,2} \quad .
\end{aligned}$$

So

$$\begin{bmatrix}\rho+\widetilde{x} & 1 \\ \rho & 1-\widetilde{x}\end{bmatrix}\begin{pmatrix}D_{k,2} \\ D_{k,3}\end{pmatrix}=-\begin{pmatrix}1+\rho \\ \rho\end{pmatrix}D_{k,1}$$

where $\widetilde{x}=\left(D_b-\psi\right)\frac{1}{1-\frac{1-C/Y}{1-\beta}\psi}\frac{1}{\beta}\frac{C}{Y}$. Thus

$$\begin{aligned}
\begin{pmatrix}D_{k,2} \\ D_{k,3}\end{pmatrix} &= -\frac{1}{\left(1-\rho-\widetilde{x}\right)\widetilde{x}}\begin{bmatrix}1-\widetilde{x} & -1 \\ -\rho & \rho+\widetilde{x}\end{bmatrix}\begin{pmatrix}1+\rho \\ \rho\end{pmatrix}D_{k,1} \\
&= -\frac{D_{k,1}}{\widetilde{x}}\frac{1}{\left(1-\rho-\widetilde{x}\right)}\begin{pmatrix}1-\widetilde{x}\left(1+\rho\right) \\ -\rho+\rho\widetilde{x}\end{pmatrix}.
\end{aligned}$$

B.1 Limit Result

We provide the proof of Proposition 1 for two cases $C/Y=1$ and $C/Y\neq 1$.

Proof of Proposition 1 when $C/Y=1$. Given that $C/Y=1$,

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

and

$$\begin{aligned}
D_{k,1} &= -\frac{x}{1-x} \\
D_{k,2} &= \frac{\left(1-x\right)\left(1+\rho\right)-\rho}{\left(1-x\right)\left(1-\rho-x\right)} \\
D_{k,3} &= -\frac{\rho}{1-\rho-x} \\
D_{k,4} &= \frac{D_{k,1}\rho}{1-\rho-x}
\end{aligned}$$

where $x=\left(D_b-\psi\right)\frac{1}{\beta}$. We have $\lim_{\psi\longrightarrow 0}D_b=-\left(1-\beta\right)$ 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 limit dynamics of consumption are

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

■

Proof of Proposition 1 when $C/Y \neq 1$. From the closed form expressions, it is easy to verify that as β goes to 1 and $\frac{\psi}{(1-\beta)^2}$ goes to 0: $D_b, D_{k,1}, D_{k,4}$ go to 0 and $\begin{pmatrix} D_{k,2} \\ D_{k,3} \end{pmatrix}$ goes to $\frac{1}{C/Y} \frac{1}{1-\rho} \begin{pmatrix} 1 \\ -\rho \end{pmatrix}$. ■

Notice, then, that the limit result requires that ψ goes to 0 faster than $1-\beta$.

C 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 (11) is modified to:

$$R_t = R_t^* + \psi \left\{ e^{\frac{B_t}{Y_t} - b} - 1 \right\} \tag{21}$$

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 (21).

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 (14), by:

$$r_t = r_t^* + \psi \cdot b_t \quad . \tag{22}$$

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

$$mb_t + m^* b_t^* = 0 \quad . \tag{23}$$

It is straightforward to show that Proposition 1 generalizes to this model. Therefore, for the standard parametrization in the literature, our main results can also be obtained in a two country model.

D Data Appendix

In the case of the Great Recession, the series for **productivity** is constructed by dividing GDP by the labor input and taking logs. 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 **net exports** is constructed by dividing net exports by population. Net exports are measured by the difference between Real Exports and Real Imports from the St. Louis Fed database (series IDs EXPGSC96 and EXPGSC96 respectively). Population is from the BLS (series IDs LNS10000000Q). The series for **consumption** is constructed by dividing Real Personal Consumption Expenditures by Population and taking logs. The series for Real Personal Consumption Expenditures is from the St. Louis Fed database (series ID PCEC96). The series for **TFP** was downloaded from John Fernald's website ("A Quarterly, Utilization-Adjusted Series on Total Factor Productivity", Fernald 2012b, supplement, series dtfp_util).

In the case of Japan, the series for **productivity** and **net exports** were constructed in the same way. All series come from the OECD website. GDP, Exports and Imports 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, the series for **productivity** is constructed by dividing per capita GDP by the labor input and taking logs. The labor input series was obtained from Kendrick 1961, Appendix A, Table XXIII, 2nd column ("Persons Engaged"). (Gordon 2000 uses the same measure.) The series for **net exports** is constructed by the difference between exports and imports. Per capita GDP, consumption, exports and imports were obtained from Robert Gordon's website.

Our data set is available upon request.

[FOR ONLINE PUBLICATION]

Supplementary Material for "Technological Revolutions and the Three Great Slumps: A Medium-Run Analysis"

Dan Cao and Jean-Paul L'Huillier

July 2014

Preamble: For printing convenience, we have split the original 51-pages paper into the main document and this one. This document contains additional results.

E Normalization and Log-linearization

E.1 Steady State.

We look for a steady state in which the following variables (normalized and non-normalized) are constant: $c = C$, $b = B/Y$, R, and Q. We assume that the steady state level of normalized debt b is determined exogenously.[36]

From the intertemporal condition (12), we have

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

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

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

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

$$1 = \beta R \frac{A}{A_+} \quad . \tag{24}$$

Since $A + /A = 1$,

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

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

The resource constraint (9) gives

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

or

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

So

$$c + bN = N + \beta bN,$$

this implies

$$c = N\left(1 - (1-\beta)\, b\right) \quad .$$

E.2 Log-linearization.

We log-linearize the intertemporal condition

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

to obtain (13). Log-linearizing the interest-elasticity equation (11) immediately gives (14).

Approximating the resource constraint delivers

$$\frac{C}{Y}\left(c_t + 1\right) + \frac{NX}{Y} + nx_t = 1 \quad ,$$

which leads to (15).

Net exports are

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

and therefore, approximating

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

to obtain (16).

F VAR Representation of the Limiting Model

In this section we derive (17) and (18). We know that

$$\begin{aligned} a_t - \rho a_{t-1} &= x_t + z_t - \rho\left(x_{t-1} + z_{t-1}\right) \\ &= x_t - \rho x_{t-1} + \eta_t \quad . \end{aligned}$$

At the limit

$$\widehat{c}_t = \frac{1}{C/Y}\frac{1}{1-\rho}\mathbb{E}_t\left[x_t - \rho x_{t-1}\right] \quad .$$

Notice that

$$x_t - \rho x_{t-1} = x_{t-1} - \rho x_{t-2} + \epsilon_t \quad ,$$

so

$$\begin{aligned} &\mathbb{E}_{t-1}\left[\widehat{c}_t\right] \\ = &\frac{1}{C/Y}\frac{1}{1-\rho}\mathbb{E}_{t-1}\left[\mathbb{E}_t\left[x_t - \rho x_{t-1}\right]\right] \\ = &\frac{1}{C/Y}\frac{1}{1-\rho}\mathbb{E}_{t-1}\left[x_t - \rho x_{t-1}\right] \\ = &\frac{1}{C/Y}\frac{1}{1-\rho}\mathbb{E}_{t-1}\left[x_{t-1} - \rho x_{t-2} + \epsilon_t\right] \\ = &\frac{1}{C/Y}\frac{1}{1-\rho}\mathbb{E}_{t-1}\left[x_{t-1} - \rho x_{t-2}\right] \\ = &\widehat{c}_{t-1} \quad , \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}_{t-1}\left[a_t - \rho a_{t-1}\right] &= \mathbb{E}_{t-1}\left[x_t - \rho x_{t-1}\right] \\ &= \frac{C}{Y}\left(1-\rho\right)\widehat{c}_{t-1} \quad . \end{aligned}$$

Therefore we have the VAR representation

$$\begin{aligned} \widehat{c}_t &= \widehat{c}_{t-1} + u_t^c \\ a_t &= \rho a_{t-1} + \frac{C}{Y}\left(1-\rho\right)\widehat{c}_{t-1} + u_t^a \quad . \end{aligned}$$

Equation (20) is obtained by induction in j. We just showed it holds for

Figure 15: Smoothed Transitory and Noise Shocks (U.S., 1990–2013)

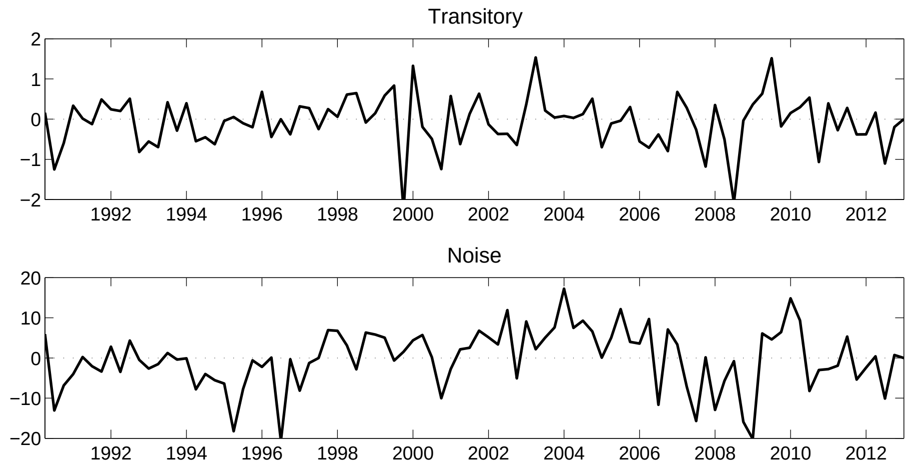

Notes: Shocks estimated using a Kalman smoother on the U.S. 1990–2013 sample. The data is composed by the first differences of the logarithm labor productivity and the ratio of net exports-to-GDP. The time unit on the x-axis is percentages. Shocks are scaled by their ML estimated standard deviation.

$j = 0$. If it holds for j, then $\mathbb{E}_t\left[a_{t+j}\right] = \rho^j a_t + C/Y \cdot (1-\rho^j)\,\widehat{c}_t$. Taking expectations at time $t-1$ on both sides yields

$$\begin{aligned}\mathbb{E}_{t-1}[a_{t+j}] &= \rho^j \mathbb{E}_{t-1}[a_t] + C/Y \cdot (1-\rho^j)\mathbb{E}_{t-1}[\widehat{c}_t] \\ &= C/Y \cdot (1-\rho^j)\widehat{c}_{t-1} + \rho^j(\rho a_{t-1} + C/Y \cdot (1-\rho)\widehat{c}_{t-1}) \\ &= \rho^{j+1} a_{t-1} + C/Y \cdot (1-\rho^{j+1})\widehat{c}_{t-1} \quad ,\end{aligned}$$

the second equality follows from (17) and (18), the third from rearranging.

G Estimated Temporary Productivity and Noise Shocks

For completeness, in this section we report our estimated transitory and noise shocks.

Figure 15 plots these shocks for the case of the Great Recession. In contrast to the estimated permanent shocks shown in the body of the paper (p. 19), transitory and noise shocks do not have any particular pattern. Figure 16 plots these shocks for Japan and the Great Depression. Similarly, these shocks do not have any particular pattern either.

Figure 16: Smoothed Transitory and Noise Shocks (Japan 1975–2005, and U.S. 1920–1935)

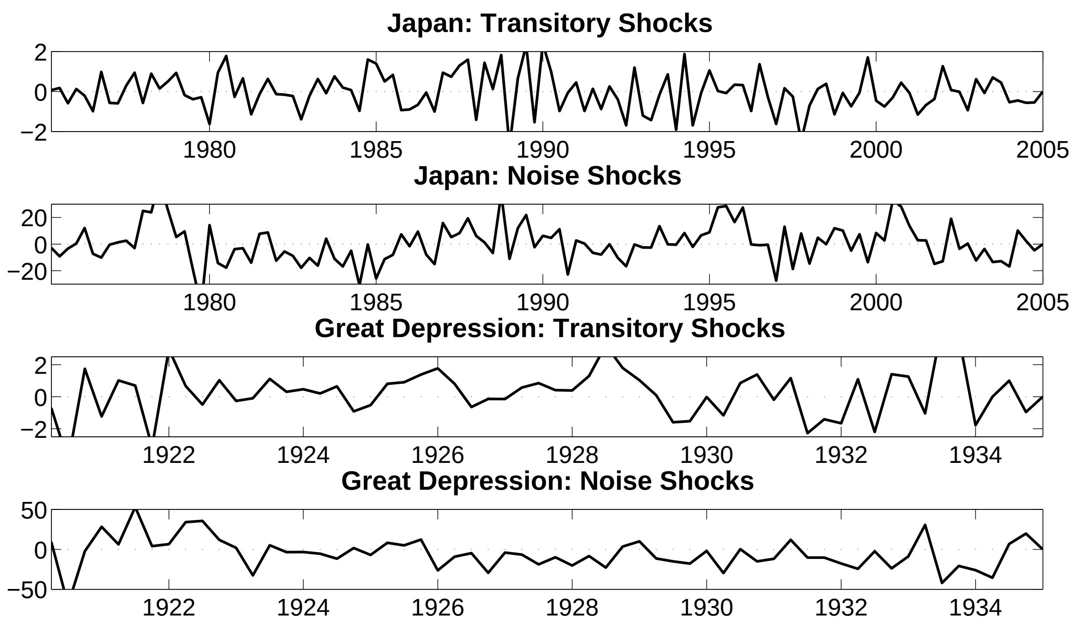

Notes: Shocks estimated using a Kalman smoother on the Japanese 1975–2005 sample, and on the U.S. 1920–1935 sample. The latter is restricted by data availability. The data is composed by the first differences of the logarithm labor productivity and the ratio of net exports-to-GDP. The time unit on the x-axis is percentages. Shocks are scaled by their ML estimated standard deviation.