Journal of Monetary Economics 96 (2018) 93–108

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmoneco

Technological revolutions and the Three Great Slumps: A medium-run analysis[☆]

Dan Cao[a], Jean-Paul L'Huillier[b,*]

[a] Georgetown University, 3700 O Street NW, Washington, DC 20001, USA
[b] Brandeis University and EIEF, 415 South Street, Waltham, MA 02453, USA

ARTICLE INFO

Article history:
Received 3 March 2015
Revised 15 February 2018
Accepted 8 April 2018
Available online 10 April 2018

JEL classification:
E21
E27
E32

Keywords:
Aggregate productivity
Permanent income
Learning
Secular stagnation
Financial crises

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. We estimate a model with noisy news. We find that beliefs about long-run income adjust to permanent shifts in productivity with an important delay. The estimation tells 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.

1. Introduction

A medium-run look at the three deepest recessions in developed economies reveals that they were all preceded by periods of great technological innovation and economic transformation. Specifically, the Great Recession in the United States was preceded by a technological revolution, happening in the mid- to 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 products that placed Japan at the global centerstage in electronics.[1] We view this period as containing the elements of a technological revolution with effects concentrated in Japan. Finally, before the Great

☆ This is a revised version of the draft entitled "Technological Revolutions and Debt Hangovers: Is There a Link?" Date of first draft: December, 2012. We benefited from useful feedback from David Backus, Philippe Bacchetta, Susanto Basu, Paul Beaudry, Emine Boz, Claire Brennecke, Robert Chirinko, Chris Faulkner-MacDonagh, Virgiliu Midrigan, Ezra Oberfield, Franck Portier, Giorgio Primiceri, Vincenzo Quadrini, Daniele Terlizzese, Nicholas Trachter, and seminar and conference participants at the Federal Reserve Banks of Chicago and Philadelphia, Boston College, Northwestern University, Princeton University, New York University, the 1st Mannheim Workshop in Quantitative Macroeconomics, the 5th Joint French Macro Workshop, and the 22nd European Summer Symposium in Economic Theory. Dan thanks EIEF for hospitality, and Jean-Paul thanks NYU.

* Corresponding author.
E-mail addresses: dc448@georgetown.edu (D. Cao), jpl@brandeis.edu, jean-paul.lhuillier@eief.it (J.-P. L'Huillier).

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

https://doi.org/10.1016/j.jmoneco.2018.04.003
0304-3932/

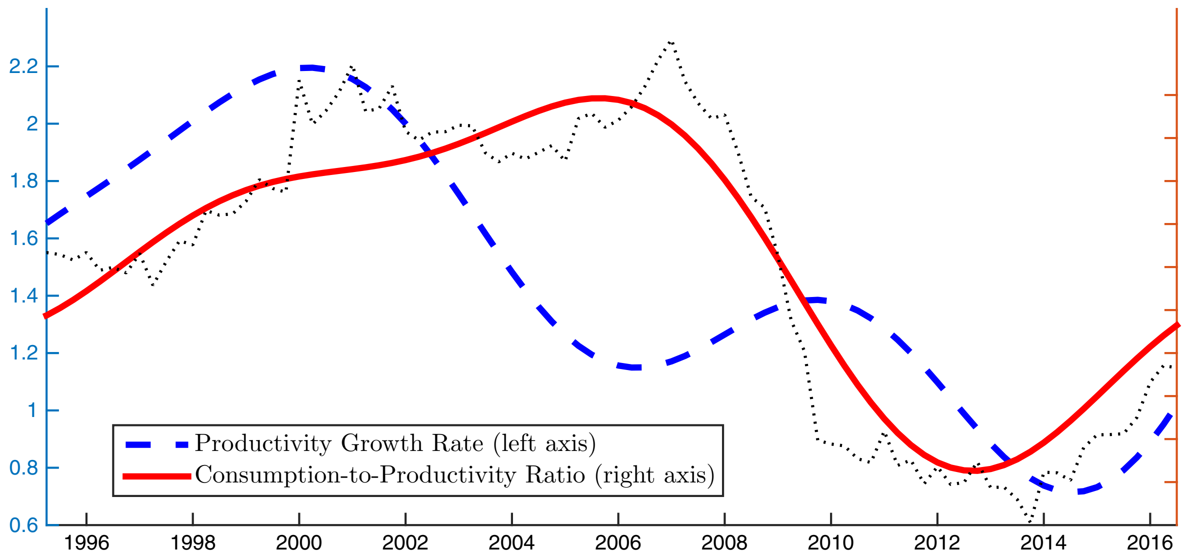

Fig. 1. Productivity growth and delayed adjustment of consumption. *Notes:* Annualized labor productivity growth rates (dashed-blue, smoothed using a band pass filter at 32–200 frequencies) with scale in the left axis (percentages); and consumption-to-productivity ratio (thin dotted-black for raw data and solid-red for smoothed data using a band pass filter at 32–200 frequencies) with scale in the right axis (average normalized to 1). The band pass filter is used to isolate the medium-run dynamics. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Depression, roughly between 1915 and 1925, the United States witnessed the so-called 2nd Industrial Revolution (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 rationalize them using a simple framework. We take a simple permanent income model in which a representative agent learns slowly about his future income. Future income is determined by permanent shifts in productivity, assumed to embody technological progress. These permanent shifts can be gauged by persistent productivity growth. However, in real time they can be difficult to tell apart from just transitory shifts. Due to imperfect information, detecting permanent shifts can be a challenging task for the agent. Consumers, who update their beliefs about future income on the basis of noisy signals, adjust their behavior gradually. This allows us to fit the dynamics of spending present in the data.[3]

The reason we need to introduce learning in the model is clear evidence of the long delays in the adjustment of spending. This, together with the gradual adjustment of productivity growth around these waves of innovation, gives rise to a slow-moving cycle. This cycle can be summarized by the following sequence of events: First, there is an initial and persistent increase in productivity growth. Anecdotically, this shift seems to coincide with the waves of innovation mentioned previously. Second, the increase in productivity growth generates, with a delay of several years, a (rational) "wave of optimism". This optimism increases consumer spending. Third, there is a large and very persistent a decline in productivity growth. This decline usually starts years *before* the economic slump. Arguably, it is caused by a slowdown in the pace of innovation, which can be attributed to the exhaustion of the low-hanging fruits of the new technology. Fourth, a "wave of pessimism" arrives. This pessimism persistently decreases consumer spending.

In order to visualize these facts, consider Fig. 1, plotting the (smoothed) growth rates of U.S. labor productivity and the ratio of consumption-to-productivity before and during the recent Great Recession.[4] According to the permanent income hypothesis, the consumption-to-productivity ratio carries information about consumers' view of their *future income*, a point developed fully in Section 2 below. Fig. 1 shows that the growth rates of productivity increased by more than 0.70% a year between roughly 1994 to 2000, with growth peaking around the turn of the century.[5] However, consumption increased

[2] The 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).

[3] To be clear, ours is a model of consumption, not of output (other than exogenous productivity shocks). It is relatively straightforward to embed our model structure in a demand-determined DSGE model to generate endogenous fluctuations in output, but for clarity and better focus we do not carry out that exercise in the present paper.

[4] Using TFP instead of labor productivity delivers a similar picture.

[5] Labor productivity annual growth rates averaged 1.46% from 1993 to 1995, and increased to an average of 2.19% from 1996 to 2002.

(relative to productivity) with a lag of several years, peaking around 2006 or 2007. Productivity growth had started to decline several years earlier. This decline was sustained.[6] Consumers took a while to revise spending downwards, with the ratio of consumption-to-productivity starting to decline only in 2007 (dotted-black line: raw data, full-red: smoothed). Overall, it is remarkable how long it took for consumers to revise their views about their long-run income.

We attempt to make sense of these observations within a standard model. Our model has two main ingredients. The first is the presence of both permanent and transitory shocks to productivity. The second is the presence of news about the future (Barsky and Sims, 2012; Beaudry and Portier, 2006; Jaimovich and Rebelo, 2009). Similar to Blanchard et al. (2013), we allow news to be noisy. Our focus though is on the effect of permanent shifts in productivity and the slow adjustment of expectations, instead of the effect of noise shocks.[7]

In order to estimate permanent shocks, we use a tractable framework 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 the productivity trend by looking at consumption. Here, we borrow the basic idea of an important body of work on household income dynamics (see Blundell and Preston, 1998, among others.)

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.

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 (shutting down all other shocks) and simulate. We do this in order to match a collection of medium-run moments, such as the magnitude of the productivity growth increase, and the magnitude of the subsequent decrease. We also focus on "timing" moments, such as the date of the peak of productivity growth, and the date of the peak of the ratio of consumption-to-productivity. Lastly, we look at the number of years elapsed between these two peaks, which is a measure of the delay in the adjustment of consumption. Our estimates of the permanent shocks that hit the U.S. economy over this period match this set of stylized facts quite well. We also repeat this exercise for the cases of Japan and the Great Depression, and conclude that the model is consistent with the data in these two cases as well.

We perform a number of out-of-sample checks by simulating other endogenous variables. In terms of beliefs, we compare model-generated beliefs to survey evidence for the U.S. economy, 1994–2010. We find that according to both the model and the survey, the U.S. consumer was most optimistic about his long-run income around 2004. We perform a similar exercise for net exports.

As pointed out in the influential work by Beaudry and Portier (2004), Pigou was an early proponent of the idea that bursts of optimism may anticipate recessions.[8] We share the same motivation, but our work emphasizes the medium run.

Little attention has been devoted to the study of medium-term aggregate consumption dynamics. The bulk of the empirical DSGEs literature 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 effect of learning. Other work in this vein includes Blanchard (1997), which focuses on unemployment and capital shares. Evans et al. (1998) resonates with our findings of high and low growth episodes, and how this interacts with agents' expectations. Relatedly, our paper provides an empirical basis for the optimism preceding crises of emerging countries in the theoretical contribution by Boz (2009).

The rest of the paper proceeds as follows. We present the model in Section 2. We take a preliminary but useful look at the data in Section 3. We present the model estimation results in Section 4. We conclude in Section 5. Appendix A contains a detailed description of our data. Other appendices present further theoretical and empirical results.

2. The model

2.1. Productivity process and information structure

We model an open economy similar to Schmitt-Grohe and Uribe (2003), adding a "news and noise" information structure (Blanchard et al., 2013, henceforth BLL).[9] Specifically, productivity a_t (in logs) is the sum of two components, permanent, x_t, and transitory z_t:

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

[6] The decline bottoms in the 2010s. (Labor productivity annual growth rates averaged 0.61% from 2010 to 2014.)

[7] Edge et al. (2007) explore learning about shifts in long-run productivity growth using U.S. data up to 2005.

[8] For similar ideas, see also Beaudry et al. (2018).

[9] Boz et al. (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. 9).

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

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

The transitory component follows the stationary process

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

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

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

and that the variances satisfy

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

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

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

This assumption is analytically convenient. Moreover, it is broadly in line with productivity data.[10] 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.[11]

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

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

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

We think of ε_t as the "news" shock because it builds up gradually 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 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 the productivity trend by looking at the behavior of consumption. As mentioned in the introduction, this connects our paper to the work of Blundell and Preston (1998). (Our identification strategy is discussed in detail below.)[12]

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.[13] 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

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

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 second equality comes directly from the definition of $x_{\tau|t}$. To prove the first equality, we make use of Eqs (1), (2), and (3).

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 permanent 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 **BLR-delay** by the time it takes BLR to reach half of the long-run level.

[10] In a similar exercise, BLL (working paper version) relax this assumption and show that for U.S. data this does not change the empirical inference about (1), (2) and (3).

[11] See BLL for the proof.

[12] Forni, Gambetti, Lippi, and Sala (2017) also use the term "noisy news", but they use a different specification of the information structure.

[13] The construction of the filter is standard. We refer the interested reader to BLL for more details.

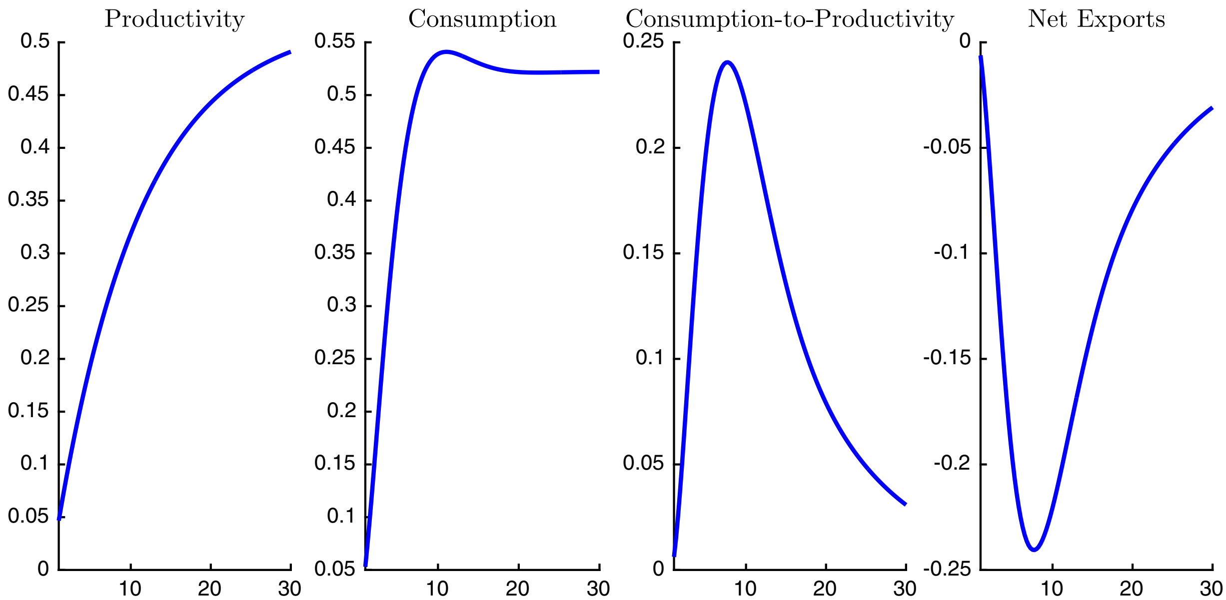

Fig. 2. Impulse response functions to a permanent technology shock. *Notes:* Impulse response functions of, from left to right, a_t, $\hat{c}_t$, c_t, and nx_t. Units in the vertical axis are percentage relative deviations from steady state in the case of a_t, $\hat{c}_t$, c_t, and absolute percentage deviations from steady state in the case of nx_t. The time unit on the x-axis is one year (four quarters).

2.2. Production and consumption

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

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

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:

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

where $A_t = e^{a_t}$. We abstract from fluctuations on employment, i.e. the consumer supplies labor N inelastically.[14] We normalize $N = 1$. The resource constraint is

$$C_t + NX_t = Y_t,$$

where NX_t stands for net exports. 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:

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

where $\bar{b}$ represents the steady state level of the debt-to-output ratio.[15]

The only first-order condition is:

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

[14] 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. 9).

[15] 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 B.7 (supplementary material).

In order to examine the dynamics of consumption, we define the ratio of consumption-to-productivity. This is simply the logarithm of normalized consumption:

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

Also, we define the variable $\hat{c}_t$, which the log-deviation of consumption

$$\hat{c}_t = c_t + a_t$$

(see the supplementary material for the remaining details of the model normalization.)

In the standard parametrization of the log-linearized model (a discount factor β close to 1 and an elasticity of the interest rate ψ close to 0)[16], the effects of productivity shocks on consumption $\hat{c}_t$ are mainly determined by BLR, as established by the following proposition.

Proposition 1. *As $\beta \longrightarrow 1$, and $\psi \longrightarrow 0$, consumption $\hat{c}_t$ is only a function of BLR. Specifically,*

$$\hat{c}_t = x_{t+\infty|t} \quad .$$

The proof is in the supplementary material. Given this formal result, and our goal of relying on the permanent income behavior of consumption for our inferences, throughout the paper we shall calibrate β and ψ to be close to this limit. We are also able to prove a version of this theorem for a more general model that includes labor supply and capital. Therefore, for this calibration, including those ingredients in our framework does not change our results. Section 4 presents an extension with investment and capital.[17]

The rest of the parameters is taken from the estimation of the model for the U.S. 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, denoted σ_a, is set at 0.53%. These values for ρ and σ_a yield standard deviations of the two technology shocks, σ_ε and σ_η, equal to 0.01% and 0.51%, respectively. The standard deviation of the noise shock, σ_ν, is set to 1.22%, implying a fairly noisy signal.

Fig. 2 shows a simulation of the model for these parameter values. The figure shows the responses of productivity a_t, consumption $\hat{c}_t$, the consumption-to-productivity ratio c_t, and net exports nx_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, consumption, and the ratio is relative percentage deviations from steady state. The scale of net exports is absolute percentage deviations from the steady state value of net exports-to-output, NX/Y.

In response to a one-standard-deviation increase in ε_t, the permanent technology shock, productivity increases slightly on impact, and then gradually continues to increase until it reaches a new long-run level. This sustained increase is slow; in fact, half of the productivity increases are reached only after 7 years. In the perfect information benchmark consumption would rise immediately to its long-run level, but here this happens gradually. In this parametrization consumption actually overshoots the long-run level, and then goes back down. This is again a consequence of noisy information, together with the high persistence of the permanent component.

A key response for our purposes is the one of the ratio of consumption-to-productivity. The behavior of this variable in the data will be the objective of the analysis of the next section. Because in the perfect information benchmark consumption would immediately jump to its long-run level, the ratio would have a declining shape. However, in this parametrization with significant noise, the ratio first increases, peaks 8 years after the shock, and then goes back to zero. Initially, net exports rise (not visible on the figure), because productivity increases faster than beliefs about long-run productivity. This is a reflection of the amount of noise in this simulation. After 3 quarters 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, together with consumption, reach a new level (at 0.53%). The ratio and net exports go back to zero.

2.3. Bivariate representation of the simple model

Another striking implication of Proposition 1 is that the model admits a bivariate representation on productivity and consumption, which provides useful insights into the behavior of consumption and allows to clearly discuss parameter identification. Because the proposition is also valid in the model that includes investment and labor supply, this representation is also valid quite generally, justifying the use of the simple model in our benchmark estimations below.

The representation is given by the following two equations:

$$\hat{c}_t = \hat{c}_{t-1} + u_t^c \tag{13}$$

$$a_t = \rho a_{t-1} + (1-\rho)\hat{c}_{t-1} + u_t^a \quad , \tag{14}$$

[16] See Schmitt-Grohe and Uribe (2003), Aguiar and Gopinath (2007); Boz et al. (2011), Hoffmann et al. (2013), among others.

[17] Using the common value in the DSGE literature of $\beta = 0.99$, or the same value for ψ used in the literature ($\psi = 0.0010$) does not qualitatively affect our conclusions. (The reason is that these values are close to the limits of interest.)

where u_t^c and u_t^a are innovations. The derivation of all expressions discussed here not presented previously can be found in Appendix B.6 (supplementary material). According to (13), consumption in the limiting model is a random walk, which simply follows from the law of iterated expectations. Equation (14) 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 $\hat{c}_{t-1}$. The reason is as follows. Past consumption $\hat{c}_{t-1}$ 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.

Not only the limit model admits a simple representation as (13) and (14), we also know that the parameters of interest are identified. The parameter σ_a, defined by the standard deviation of Δa_t, is directly identified in the data. Identification of ρ comes from Eq. (14), which can be estimated by OLS in the following form:

$$\Delta a_t = (1-\rho)(\hat{c}_{t-1} - a_{t-1}) + u_t^a \quad . \tag{15}$$

The intuition provided by Eq. (15) 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.[18] 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 consumption-to-productivity 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 and productivity changes one quarter ahead is low. Notice, this does not reflect a failure of consumers to forecast the productivity trend, because at longer horizons the equation is

$$a_{t+j} - a_t = (1-\rho^j)(\hat{c}_{t-1} - a_{t-1}) + u_{t+j}^a \quad , \tag{16}$$

and thus for long horizons the coefficient in front of the ratio goes to 1. In other words, the longer the horizon, for a given variance of the consumption-to-productivity ratio, the higher the correlation between the consumption-to-productivity ratio and the productivity trend. Notice also that these equations are valid for any degree of noise in the signal. In particular, the relationship between the productivity trend and the consumption-to-productivity ratio (15) holds on average and takes into consideration the extra volatility of the ratio coming from the noise in the signal.[19] Having identified ρ, the sizes of the permanent and transitory shocks σ_ε and σ_η can be derived from $\sigma_\varepsilon = (1-\rho)\sigma_a$ and $\sigma_\eta = \sqrt{\rho}\sigma_a$, 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.[20] The relation is monotonic and uniquely pins down σ_ν.

Given the importance of the ratio of consumption-to-productivity to evaluate consumers' views about the future, the next section fully focuses on this ratio and presents some novel stylized facts.

3. The ratio of consumption-to-productivity

Before showing the results from estimation, we shall zoom into a transformation of the data $c_t = \hat{c}_t - a_t$, because according to (15), it delivers insights into consumers' beliefs about the future. Intuitively, a high (low) ratio $\hat{c}_t - a_t$ is an indicator of 'over- (under-) consumption' with respect to current income a_t, which can be rationalized by optimistic (pessimistic) beliefs. We will look at the shape of this ratio in the three cases and then discuss two theoretical benchmarks. This will turn out to be a useful lead into the structural estimation section below.

Data. Our baseline data set includes series on labor productivity and consumption. We use quarterly data. The series for the Great Recession is from the Bureau of Economic Analysis and the Bureau of Labor Statistics. The series for Japan is from the OECD.

In the case of the Great Depression, we have data for the components of GDP from the Gordon-Krenn data set. 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. Gordon and Krenn (2010) use 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 produce a series for labor productivity, we obtain an estimate for GDP from the Gordon-Krenn data set, and we use the Kendrick (1961, Appendix A, Table XXIII, 2nd column) data set for employment, using a linear interpolation out of the annual series.

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

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

[19] Notice that in the model productivity follows a random walk, and therefore productivity by itself is not useful to identify ρ.

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

The Ratio of Consumption-to-Productivity. Fig. 3 plots the logarithm of the ratio of consumption-to-productivity around the Great Recession, the Japanese crisis, and the Great Depression. The vertical axis is centered around the average of the ratio over the period considered (normalized to 1.) The medium-run dynamics of these series are isolated using a band pass filter (between 32 and 200 quarters, following Comin and Gertler (2006).)

In all three cases, the ratio follows a slow moving "wave" that results in a hump-shaped path (after a short initial fall). This slow-moving path takes about 20 to 25 years (in case of the Great Depression our sample starts in 1920 for data availability reasons.) As the top panel shows, in the case of the Great Recession the ratio has relatively low values in the early 1990s, with a slight decreasing portion between 1990 and 1992. This is because during this period productivity is growing at a higher rate than consumption. The ratio starts to increase around 1992, and this increase becomes more pronounced starting in 1997, where consumption grows at a considerably stronger rate than productivity. The ratio reaches its highest 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.

The middle panel plots the same ratio for Japan. The ratio starts again at relatively low values, which indicate a stronger growth of productivity. We can then see an increase in the ratio, which is when consumption grows faster. The highest point of the ratio is reached in 1997, after which a downward movement brings the ratio back down, suggesting that similarly to the previous case, the ratio followed a slow-moving up-and-down wave. The bottom panel 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 hump-shaped pattern as in the two other panels. It starts at low values, then increases, reaches a highest point at the onset of the Great Depression in 1929, and then reverts back to its level of 14 years before.

To shed light on these dynamics, it is useful to considered two extreme 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 consumption-to-productivity is flat. A flat ratio clearly fails to fit the data.

Benchmark (b): Perfect Foresight. Under perfect foresight, agents have knowledge of all future shocks right from the start. Under the conditions of Proposition 1, consumption jumps immediately to the long-run level of productivity $x_{t+\infty}$ and remains there. As a result of the positive and then negative permanent shocks, productivity first increases and then decreases. The ratio thus inherits the opposite dynamics: it decreases and then increases. This, again, fails to fit the data, where the ratio increases and then decreases.

To conclude, in both extremes of "no-news" and the perfect foresight, the model has a strongly counterfactual prediction for the behavior of the consumption-to-productivity ratio. As demonstrated below, noisy signals (finite $\sigma_\nu > 0$) imply a delay of consumption that allows the ratio to slowly increase and then decrease as consumption catches up with the increase and decrease of productivity growth, resulting in a hump shape.

4. Estimation

In this section we first explain how we estimate the model. We first show the results for the Great Recession, and we perform a number of exercises with the estimated model to study which facts can be matched. We then execute a similar application to Japan and the Great Depression.

Estimation Procedure. The state-space form of the model can be estimated through Maximum Likelihood (ML).[21] Because of the bivariate representation of the model, in our baseline estimations we include the demeaned first differences of the logarithm of labor productivity Δa_t and of consumption Δc_t as observable variables.[22] We estimate the parameters ρ, σ_a and σ_ν. (Notice, given the random walk Assumption (6) for a_t, σ_ε and σ_η are determined by ρ and σ_a.)

4.1. Great Recession

Here we present our baseline estimates for the U.S. Great Recession. Table 1 contains the parameter estimates obtained from a 1985Q1–2016Q3 sample. 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 0.53% in the case of the Great Recession. Given the random walk Assumption (6) for productivity, this high value of ρ imply

[21] 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 general representative-agent models with signal extraction.

[22] Using the ratio of net exports-to-GDP nx_t instead of consumption does not qualitatively change the results.

Table 1
Parameter Estimates, Great Recession.

Parameter	Description	Value	s.e.
ρ	Persistence tech. shocks	0.97	0.01
σ_a	Std. dev. productivity	0.53	0.02
σ_ε	Std. dev. permanent tech. shock (implied)	0.01	–
σ_η	Std. dev. transitory tech. shock (implied)	0.51	–
σ_ν	Std. dev. noise	1.22	0.49

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 consumption. 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.

a standard deviation for permanent technology shocks that is fairly small, 0.01%, and a fairly large standard deviation for the transitory technology shock, 0.51%. The standard deviation of noise shocks is large, 1.22%.[23] Notice also that permanent shocks 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 this estimates. By implication, the BLR-delay in learning is quite long, computed to 3 years and 1 quarter for the parameters above.

4.2. Estimated permanent shocks

Fig. 8 in the supplementary material presents the variance decomposition of BLR at different horizons. For brevity, we comment here only the main feature of this decomposition for our purposes: Starting from the medium horizon onwards (after, say, 5 years), the largest share of the forecast error of BLR is accounted for by the permanent technology shocks. The opposite holds at shorter horizons: In this case, the forecast error of BLR is mostly accounted for by transitory and noise shocks. Thus, given our emphasis on the medium run, we focus on the effect of permanent shocks throughout the paper.

The state-space representation of the estimated model can be used in order to estimate the state and shocks using a Kalman smoother. Fig. 4 shows our estimated permanent shocks in the case of the Great Recession.[24] We estimate positive shocks from roughly 1989 to 1999, and negative shocks later on. The serial correlation of our estimated permanent shocks is not a violation of the i.i.d. assumption on these shocks, 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 is autocorrelated, and thus 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 holds starting in 1999.[25]

In order to assess the ability of the model to recover the right shocks, we compare the implied productivity growth rates to the data. This is a standard historical decomposition exercise, as follows. We feed into the model the series of estimated permanent shocks shown in Fig. 4, top panel, setting the other two shocks η_t and ν_t to zero. We then simulate the implied productivity growth rates. We superimpose these model implied growth rates to their data counterpart, using a 10-year centered moving average in order to isolate the medium-run movements. The result is presented in the bottom panel of the figure, showing that the model does a remarkable job at estimating the shocks responsible for the medium-run variation in productivity growth.

Our finding of a persistent increase, and then a decrease, of productivity growth finds support in other independent research. The estimated permanent shocks imply that we should have observed a productivity acceleration in the mid-90s, and a subsequent slowdown, arriving several years before the start of the Great Recession. First, using different techniques, Kahn and Rich (2007) also find evidence of a permanent increase in productivity growth that started in the mid-1990s. Furthermore, Fernald (2014) documents detailed evidence, at different levels of aggregation, that the growth of both labor and total-factor productivity slowed down at around 2004 in most industries. (The slow down was most pronounced in IT-intensive industries.)

4.3. Medium-Run facts

We now turn to a quantitative assessment. We ask: Given the ML-estimated model, what medium-run facts of interest are these permanent shocks able to reproduce?

[23] These estimates are of the same order of magnitude and qualitatively similar to those in BLL who estimate $\rho = 0.89$, $\sigma_a = 0.67\%$ and $\sigma_\nu = 0.90\%$ using a 1970–2008 sample (p. 3056).

[24] For brevity we do not show the estimated transitory and noise shocks here, see Fig. 9 in the supplementary material.

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

Table 2
Moments, Great Recession (Non-Targeted).

Moment	Model	Data
Magnitudes		
Increase in Productivity Growth	0.62%	0.68%
Decrease in Productivity Growth	−0.49%	−0.97%
Timing		
Date of Peak of Productivity Growth	1998Q2	2000Q1
Date of Peak of Ratio $c = \hat{c} - a$	2002Q3	2003Q4
Gap Between Peaks:	4.25 years	3.75 years

Notes: The moments from the model are calculated using the time series of the endogenous variables (productivity growth, consumption-to-productivity ratio) generated by the smoothed permanent shocks (upper panel in Fig. 4). The moments for the data are calculated using the time series for the same variables after taking centered 10-year moving averages.

We start by looking at a collection of key non-targeted moments, documented in Table 2. The first moment we consider is the peak or maximum increase in productivity growth in the 1985 to 2016 years (away from the average trend of 1.44% per annum.). The second is the bottom or maximum *decrease* in productivity growth in the same years (again away from the trend).[26] In the data, productivity growth increases by 0.68% and then declines by -0.97%, with a total decline of -1.66%. We feed the smoothed permanent shocks through the model and simulate. We find that the model, with only the help of the permanent shocks, does a remarkable job at capturing the increase in growth, and accounts for a bit more than half of the decline.

Given our emphasis on the slow-moving cycle and the delays in consumption present in the data, we also present a number of "timing" moments. These include, first, the respective dates of the peaks of productivity growth and the ratio. Here, the model performs well. According to the permanents shocks run through model, productivity growth peaked in 1998Q2; in the data, productivity growth peaks in 2000Q1. According to the model, the ratio peaked in 2002Q3; in the data, the ratio peaked in 2003Q4. So, the model's implied timing of these peaks using the permanent shocks is less than two years apart of the actual timing of the peaks, which is a small difference from a medium-run perspective. We also focus on the time gap between these two peaks, which is another intuitive measure of the delay in the reaction of consumption. In this case, it is interesting to note that the model delivers a slightly longer delay than in the data, which is of 3 years and 3 quarters.

We now look at the path of the consumption-to-productivity ratio generated by the permanent shocks (Fig. 5). As claimed in Section 3, the model delivers a ratio that increases initially and then decreases, in a slow-moving hump shape.

We finally present two supplementary out-of-sample checks of our estimation. We simulate the path of key endogenous variables using the permanent shocks only and compare these to the data. We focus on BLR and net exports.

Our beliefs data come from a survey published by Consensus Forecasts[27] This survey was used in the paper by Hoffmann et al. (2013). The survey includes a question of participants' expectations about GDP growth up to 10 years ahead, and therefore it should be comparable to our notion of BLR (this is the longest horizon in the survey.) Fig. 6, left panel, compares the evolution of growth expectations according to the survey (the bottom panel of Figure 1 in Hoffmann et al., 2013) and the BLR generated by our model. According to both measures U.S. agents seem to have been relatively most optimistic between 00 and 05.[28] The right panel of Fig. 6 compares actual data on net exports to those generated by the model. As a reminder, the model was estimated using data on productivity and consumption only. The model produces a path of net exports consistent with the data, with net exports being most negative in the mid-2000s.

4.4. Japan and Great Depression

For reasons of space, here we briefly present our results for the Japan crisis and the U.S. Great Depression. Our sample in the former episode spans 1980–2000, in the latter 1920–1935. In both cases both technology processes are, again, estimated to be very persistent. Also, there is quite a bit of noise in consumers' inference about long-run productivity. The parameter estimates are presented in detail in the supplementary material.

Table 3 shows the key non-targeted moments in these two cases. In the case of Japan, the model correctly predicts an increase and a decrease of both productivity growth and the ratio, but shows limited impact of the permanent shocks, accounting for roughly one third to one half of the increase and decrease of productivity growth. For instance, the model

[26] We use a 10-year centered moving average on the data in order to isolate the medium-run movements.

[27] In the limit model BLR and consumption are the same object, but this is of course not the case in the data.

[28] Estimating our model using net exports instead of consumption delivers the same result.

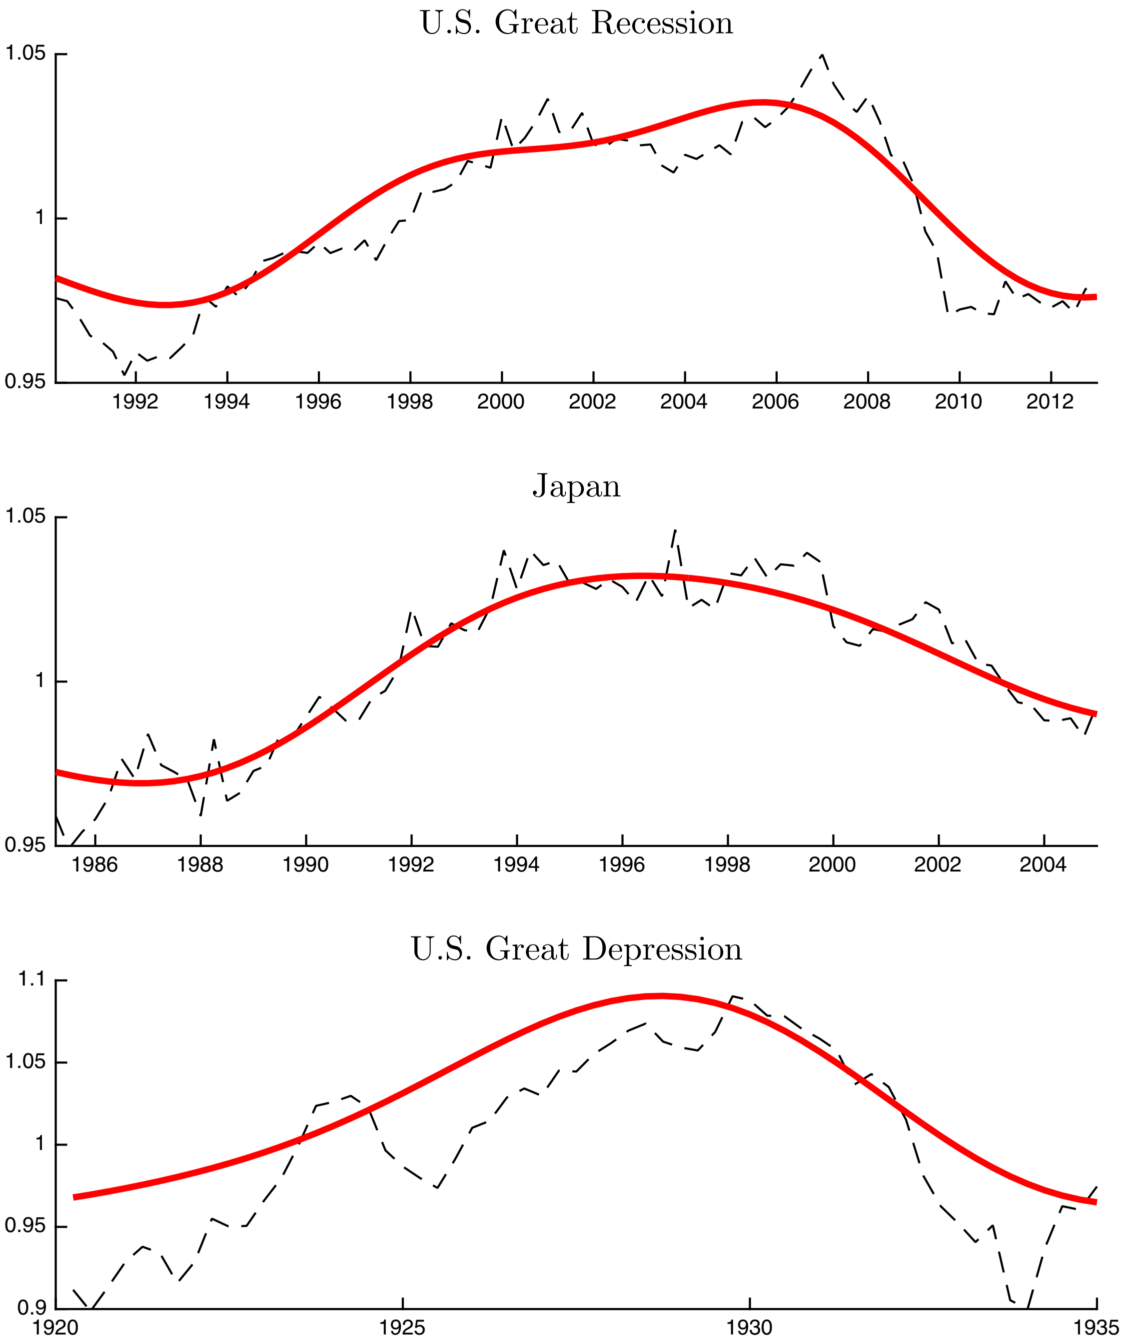

Fig. 3. Consumption-to-productivity ratio. *Notes:* Productivity is real GDP divided by employment. Consumption is NIPA consumption divided by population. In the model the ratio is $c_t = \hat{c}_t - a_t$. Black-dashed line: raw data. Red-solid line: smoothed data using a band pass filter at 32–200 frequencies. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

predicts that productivity increases by 0.64% away from the trend, wheras the increase in the data is larger, of 1.43%. Similarly, the model predicts this increase to be followed by a decrease of 0.59% away from the trend (a total decrease of 1.23%), whereas in the data the decrease is of 1.39% (total of 2.82%). In terms of the dates, the model does well in predicting the timing of the peak of productivity growth, which happened in 1985Q2 according to the data. The gap between the peaks of productivity growth and the ratio are very long in the data (13 years and 1 quarter), and the model significantly underpredicts this gap (to 5 years and a half). Accordingly, the model predicts the peak of the ratio to happen in 1993Q2, whereas in the data the ratio peaks later, in 1998Q4.

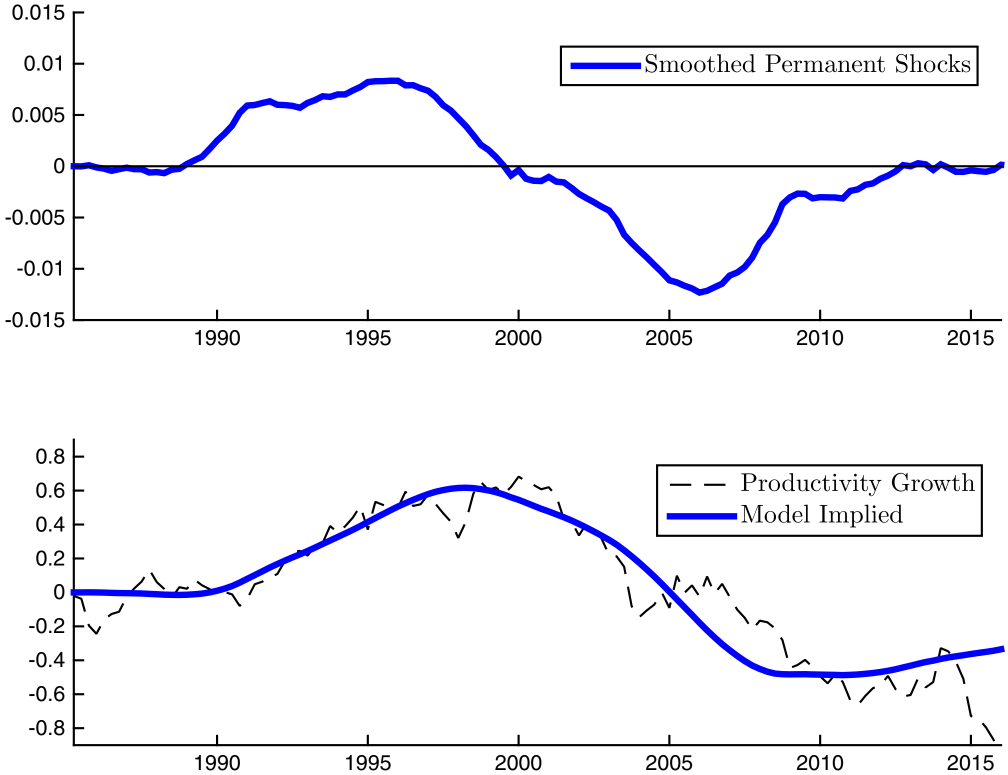

Fig. 4. Smoothed Permanent Shocks and Implied Productivity Growth, Great Recession. *Notes:* In the upper panel, smoothed permanent shocks are estimated using a Kalman smoother on the Great Recession sample. In the lower panel, the dotted-blue line represents productivity growth implied by the smoothed permanent shocks and the dashed-black line represents the centered 10-year moving average of the (demeaned) first differences of the logarithm of labor productivity. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

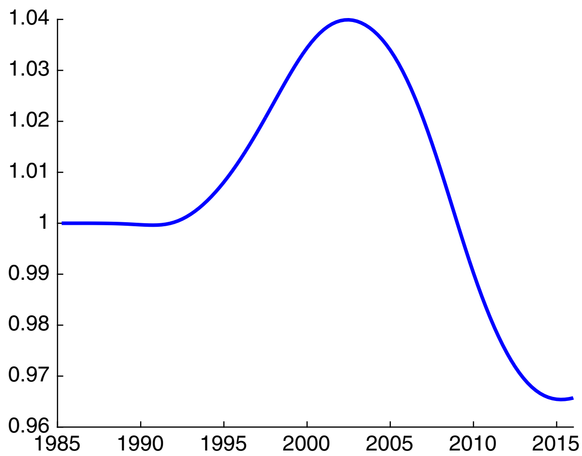

Fig. 5. Model-implied ratio of consumption-to-productivity, Great Recession.

In the case of the Great Depression, data availability restricts our sample to start in 1920. We observe productivity growth declining steadily from the start of the sample, which suggests that we are not able to observe the peak in productivity growth.[29] We observe a decrease of productivity growth of 1.61%, and the model is able to account a bit more that a third of this decline (-0.61%). Importantly, and consistent with our model of slow learning, most of this slowdown happens *before* the start of the Great Depression in 1929 (which is visible in Fig. 10 right panel, in the supplementary material), and the

[29] This observation is consistent with the implementation timing of key innovations. For instance the Ford Model-T was introduced in 1908. This suggests that one would like to have a sample for quarterly consumption and productivity starting at least 10 years before 1920.

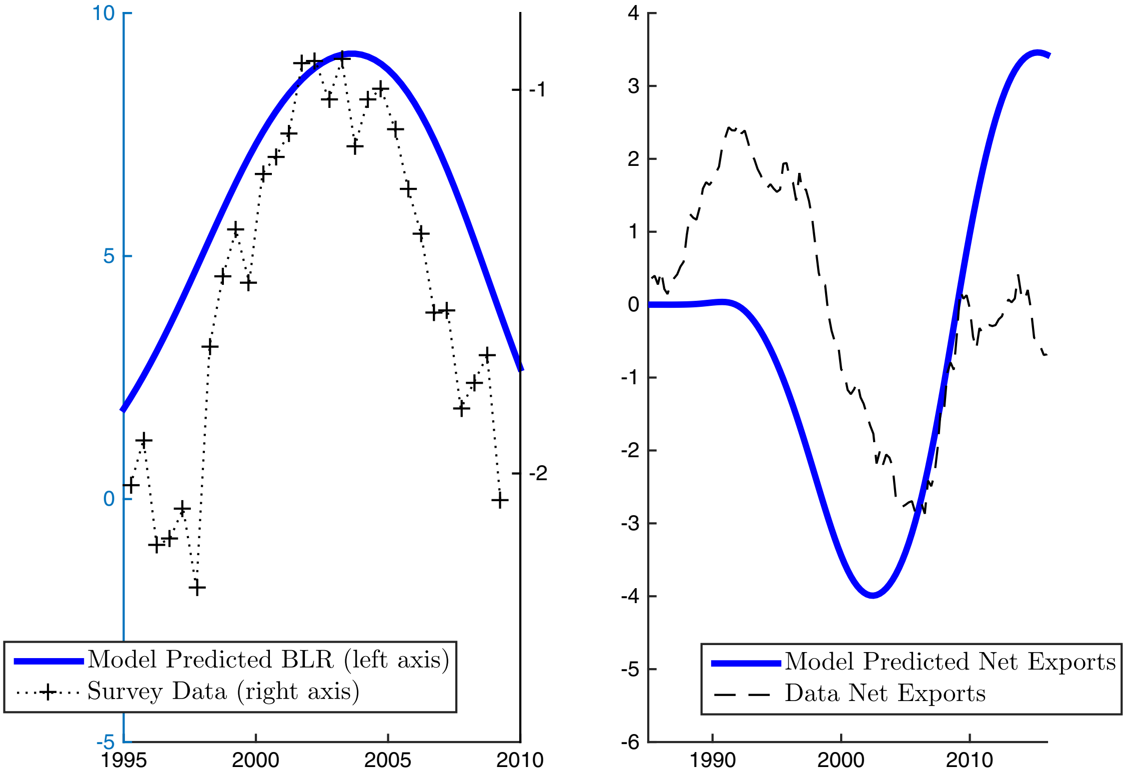

Fig. 6. Out-of-sample checks: BLR and net exports. *Notes:* In the left panel, the dotted-black line represents survey data, reproduced from Hoffmann et al. (2013) (Fig. 1, bottom panel), and the solid-blue line represents model-implied BLR, generated by the permanent shock estimates (upper panel in Fig. 4). In the right panel, the dashed-black line represents moving-average of net exports (centered 10-year), and the solid-blue line represents model implied net exports. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

peak of the ratio also happens before 1929. The model does a good job matching the date at which the ratio peaks (1926Q4 in the model, 1927Q3 in the data). The conclusion is that the gap between these two peaks must be greater than 6 years according to the model (5 years and 3 quarters according to the data).[30]

For brevity, we also report the implied path of the ratio of consumption-to-productivity in these two cases in the supplementary material (Fig. 12). In both cases the ratio initially increases and then decreases, again in a slow-moving hump shape. (In the case of the Great Depression the increase is less pronounced, presumably due to the lack of data pre-1920.)

4.5. Extension: A Model with investment

We extend the model to include capital and investment. The goal is to study whether the estimated permanent shocks imply realistic dynamics of investment.

Our specification of capital dynamics is standard. Capital accumulation is subject to the adjustment cost function

$$G(I_t, K_{t-1}) = I_t + \frac{\chi}{2}\frac{(I_t - \delta K_{t-1})^2}{K_{t-1}} \quad ,$$

where K_{t-1} stands for aggregate capital at time t and I_t stands for investment. In addition, the aggregate production function uses both capital and labor:

$$Y_t = K_{t-1}^{\alpha}(A_t N)^{1-\alpha} \quad ,$$

where N is normalized to 1.

The budget constraint is now

$$C_t + B_{t-1} + G(I_t, K_{t-1}) = Y_t + Q_t B_t$$

[30] Different from the case of Japan and U.S. Great Recession, productivity growth features a very strong recovery when the Great Depression ends (say 1933). This fact was first noted by Field (2003). See also Figure 1 in Bergeaud et al. (2016).

Table 3
Moments, Japan and Great Depression (Non-Targeted).

Moment	Model	Data
Japan		
Magnitudes		
Increase in Productivity Growth	0.64%	1.43%
Decrease in Productivity Growth	−0.59%	−1.39%
Timing		
Date of Peak of Productivity Growth	1987Q4	1985Q2
Date of Peak of Ratio $c = \hat{c} - a$	1993Q2	1998Q4
Gap Between Peaks:	5.50 years	13.25 years
Great Depression		
Magnitudes		
Increase in Productivity Growth	–	–
Decrease in Productivity Growth	−0.61%	−1.61%
Timing		
Date of Peak of Productivity Growth	–	–
Date of Peak of Ratio $c = \hat{c} - a$	1926Q4	1927Q3
Gap Between Peaks:	$>$ 6.00 years	$>$ 5.75 years

Notes: The moments from the model are calculated using the time series of the endogenous variables (productivity growth, consumption-to-productivity ratio) generated by the smoothed permanent shocks (upper panel in Fig. 4). The moments for the data are calculated using the time series for the same variables after taking centered 10-year moving averages.

and

$$K_t = (1 - \delta)K_{t-1} + I_t \quad . \tag{17}$$

With investment, net exports are given by

$$NX_t = Y_t - G(I_t, K_{t-1}) - C_t \quad .$$

The supplementary material presents the first-order conditions and the log-linearization.

We proceed as follows. We estimate this model using data on productivity and consumption. Here, we focus on the informational parameters of interest to us (ρ, σ_a, σ_ν), fixing the other parameters defining the production function and capital dynamics at standard values.[31] We then estimate the permanent shocks, and run these through the model to obtain the implied path of investment. We then compare this model-implied path of investment to the data. Because we did not use investment as an observable variable, this is a tough test of whether the permanent shocks are able to generate realistic investment dynamics.

For brevity, we do not present the parameter estimates and smoothed permanent shocks here (they can be found in the supplementary material.) Fig. 7 presents the comparison of model-implied investment to the data. The model performs well at reproducing the overall hump-shaped behavior of investment over 1985 to 2016: an increase in investment (from trend) starting around 1995, a peak some years later, and a decline after 2005. The model somewhat anticipates this path. Also, given that we are just using the permanent shocks (which are capable of generating only smooth medium-run dynamics) the model does not account for the high frequency variation quite visible in the early 2000s (a drop around the turn of the century, and a rise a few years later). However, given the simplicity of the model and the fact that investment data was not used as an input in the estimation, we find the performance of the model overall quite satisfactory.

5. Conclusion

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

At face value, the conclusion of the exercise is that the three major slumps of the developed world during the last hundred years have technological roots. But these roots are quite subtle and to see them requires taking a medium-run perspective on the data. This differentiates our paper from the bulk of the business cycles literature.

Given the obvious importance of financial frictions in the three episodes, it is quite surprising that such a simple model can account for the broad medium-run patterns in the data.[32] Considering this, we purposely decided to keep financial frictions out of the exercise, but further work will clearly enrich our understanding of these episodes by adding these frictions to a similar medium-run permanent income framework.

[31] We use $\alpha = 0.33$, $\delta = 0.025$, $\chi = 4$.

[32] A related paper here is by Gorton and Ordonez (2017), who finds support that financial cycles can be thought as a medium-run phenomenon (see the discussion on p. 4.)

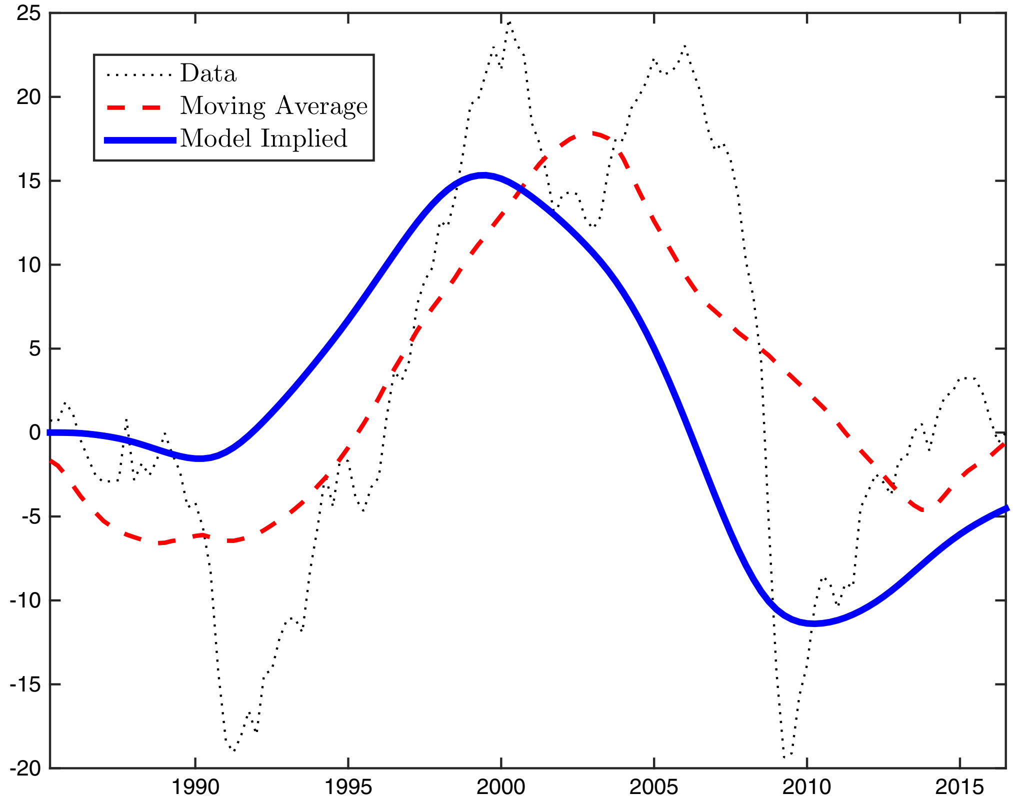

Fig. 7. Investment. *Notes:* The series for investment (dotted-black) was constructed following Justiniano et al. (2010). The dashed-red line represents its centered 10-year moving average. The solid-blue line represents the model implied investment. Units on the right axis corresponds to percentage deviations from trend or steady-state. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Finally, a question left open is what exactly makes these episodes special. Based on our findings we conjecture that the answer lies on the special nature of the observed realization of permanent technology shocks: A strong and persistent pick up of productivity growth rates, and an equally strong reversal. These seem to "sow the seeds" for trouble. Exploring this possibility further seems like a promising research avenue.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jmoneco.2018.04.003.

References

Aguiar, M., Gopinath, G., 2007. Emerging market business cycles: the cycle is the trend. J. Pol. Econ. 115 (1), 69–102.
Bardou, J.-P., Chanaron, J.-J., Fridenson, P., Laux, J.M., 1982. The Automobile Revolution: The Impact of an Industry. University of North Carolina Press.
Barsky, R.B., Sims, E., 2012. Information, animal spirits, and the meaning of innovations in consumer confidence. Am. Econ. Rev. 102 (2), 1343–1377.
Beaudry, P., Galizia, D., Portier, F., 2018. Reconciling Hayek's and Keynes' views of recessions. Rev. Econ. Stud. 85 (1), 119–156. doi:10.1093/restud/rdx008. https://oup/backfile/content_public/journal/restud/85/1/10.1093_restud_rdx008/1/rdx008.pdf.
Beaudry, P., Portier, F., 2004. An exploration into Pigou's theory of cycles. J. Monet. Econ. 51 (6), 1183–1216. https://doi.org/10.1016/j.jmoneco.2003.10.003.
Beaudry, P., Portier, F., 2006. Stock prices, news and economic fluctuations. Am. Econ. Rev. 96 (4), 1293–1307.
Bergeaud, A., Cette, G., Lecat, R., 2016. Productivity trends in advanced countries between 1890 and 2012. Rev. Income Wealth 62 (3), 420–444. doi:10.1111/roiw.12185.
Blanchard, O.J., 1997. The medium run. Brookings Pap. Econ. Act. 2.
Blanchard, O.J., L'Huillier, J.-P., Lorenzoni, G., 2013. News, noise, and fluctuations: an empirical exploration. Am. Econ. Rev. 103 (7), 3045–3070.
Blundell, R., Preston, I., 1998. Consumption inequality and income uncertainty. Q. J. Econ. 113 (2), 603–640.
Boz, E., 2009. Can miracles lead to crises? the role of optimism in emerging market crises. J. Money Credit Bank. 41 (6), 1189–1215.
Boz, E., Daude, C., Durdu, C.B., 2011. Emerging market business cycles: learning about the trend. J. Monet. Econ. 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.
Chow, G.C., Lin, A.-l., 1971. Best linear unbiased estimation, distribution, and extrapolation of time series by related series. Rev. Econ. Stat. 53 (4), 372–375.
Comin, D., Gertler, M., 2006. Medium-term business cycles. Am. Econ. Rev. 96 (3), 523–551.
David, P., Wright, G., 2000. General Purpose Technologies and Surges in Productivity: Historical Reflections on the Future of the ICT Revolution. In: David, P.A., Thomas, M. (Eds.), The economic future in historical perspective. Oxford University Press for the British Academy.

Edge, R.M., Laubach, T., Williams, J.C., 2007. Learning and shifts in long-run productivity growth. J. Monet. Econ. 54 (8), 2421–2438.
Evans, G.W., Honkapohja, S., Romer, P., 1998. Growth cycles. Am. Econ. Rev. 88 (3), 495–515.
Fernald, J., 2014. Productivity and potential output before, during, and after the great recession. NBER Macroecon. Ann. 28.
Field, A.J., 2003. The most technologically progressive decade of the century. Am. Econ. Rev. 93 (4), 1399–1413.
Forni, M., Gambetti, L., Lippi, M., Sala, L., 2017. Noisy news in business cycles. Am. Econ. J. 9 (4), 122–152. doi:10.1257/mac.20150359.
Gordon, R. J., Krenn, R., 2010. The end of the Great Depression 1939–41: Policy contributions and fiscal multipliers. NBER Working Paper No. 16380.
Gorton, G., Ordonez, G., 2017. Good booms, bad booms. NBER Working Paper 22008.
Greenwood, J., Jovanovic, B., 1999. The information-technology revolution and the stock market. Am. Econ. Rev. Papers Proc. 89 (2), 116–122.
Hobijn, B., Jovanovic, B., 2001. The information-technology revolution and the stock market: evidence. Am. Econ. Rev. 91 (5), 1203–1220.
Hoffmann, M., Krause, M., Laubach, T., 2013. The expectations driven U.S. current account. Deutsche Bundesbank Discussion Paper 10/2013.
Jaimovich, N., Rebelo, S., 2009. Can news about the future drive the business cycle? Am. Econ. Rev. 99 (4), 1097–1118.
Justiniano, A., Primiceri, G., Tambalotti, A., 2010. Investment shocks and business cycles. J. Monet. Econ. 57 (2), 132–145.
Kahn, J.A., Rich, R.W., 2007. Tracking the new economy: using growth theory to detect changes in trend productivity. J. Monet. Econ. (54) 1670–1701.
Kendrick, J.W., 1961. Productivity trends in the United States. UMI.
Pastor, L., Veronesi, P., 2009. Technological revolutions and stock prices. Am. Econ. Rev. 99 (4), 1451–1463.
Schmitt-Grohe, S., Uribe, M., 2003. Closing small open economy models. J. Int. Econ. 61 (1), 163–185.