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

Dan Cao and Jean-Paul L'Huillier*

This draft: April 2017

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 permanent shifts in productivity. This delay, together with estimated permanent shifts in the three cases, 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.

*Georgetown University, ICC 559, 37th and O Streets, NW, Washington, DC 20057, dc448@georgetown.edu; and EIEF, Via Sallustiana, 62, 00187 Rome, jean-paul.lhuillier@eief.it. 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.

"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 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 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 difficulty to tell apart from just transitory shifts. Due to imperfect information, detecting permanent shifts can be a challenging task

[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.

[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).

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.

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:

1. First, there is an initial and persistent increase in productivity growth. Anecdotically, this shift seems to coincide with the waves of innovation mentioned previously.

2. Second, the increase in productivity growth generates, with a delay of several years, a (rational) "wave of optimism". This optimism increases consumer spending.

3. 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.

4. Fourth, a "wave of pessimism" arrives. This pessimism persistently decreases consumer spending.

In order to visualize these facts, consider Figure 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.[3] 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. Figure 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.[4] However, consumption increased (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.[5] Consumers took a while to revise spending downwards, with the

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

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

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

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.

Figure 1: Productivity Growth and Delayed Adjustment of Consumption

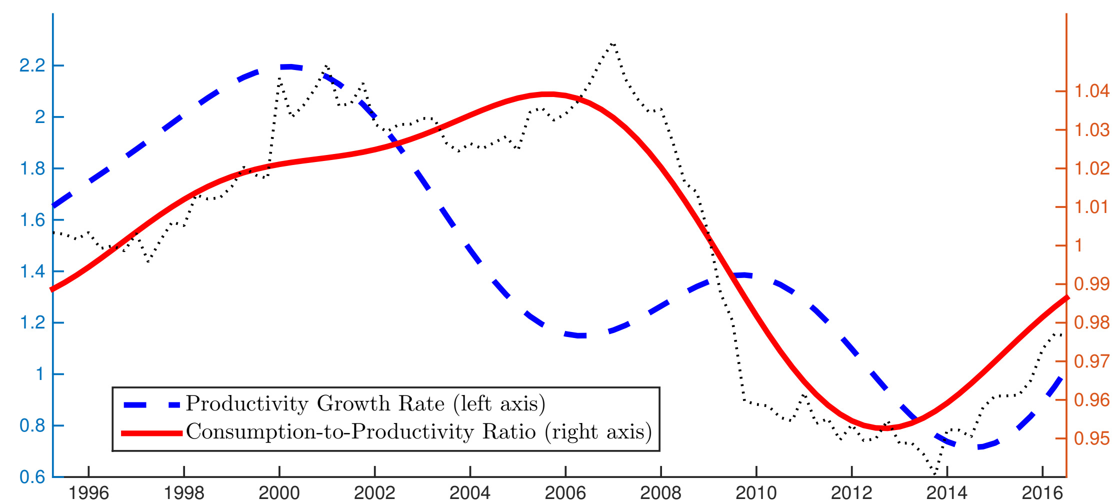

Notes: Annualized labor productivity growth rates (dashed-blue, smoothed using a band pass filter at 32-200 frequencies) with scale in the left axis (percentages); and consumption-to-productivity ratio (thin dotted-black for raw data and solid-red for smoothed data using a band pass filter at 32-200 frequencies) with scale in the right axis (average normalized to 1). The band pass filter is used to isolate the medium-run dynamics.

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 (Beaudry and Portier 2006; Jaimovich and Rebelo 2009; Barsky and Sims 2012). Similar to Blanchard, L'Huillier, and Lorenzoni (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.[6]

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

[6]Edge, Laubach, and Williams (2007) explore learning about shifts in long-run productivity growth using U.S. data up to 2005.

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.

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, Honkapohja, and Romer (1998) is an interesting paper that 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. The supplementary material presents 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, L'Huillier, and Lorenzoni 2013, henceforth BLL).[7] Specifically, productivity a_t (in logs) is the sum of two components, permanent, x_t, and transitory z_t:

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

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

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

[7] 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. 9).

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. Moreover, it is broadly in line with productivity data.[8] 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.[9]

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

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

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.)[10]

[8] 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).

[9] See BLL for the proof.

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

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.[11] 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*

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

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 equations (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.*

2.2 Production and Consumption

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

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

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

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.[12] We normalize $N = 1$. The resource constraint is

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

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:

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

where $\bar{b}$ represents the steady state level of the debt-to-output ratio.[13]

The only first-order condition is:

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

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

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

[12]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).

[13]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).

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)[14], 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. In Cao, L'Huillier, and Yoo (2016) we 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.[15]

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, 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.

Figure 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,

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

[15] 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.)

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.

Figure 2: Impulse Response Functions to a Permanent Technology Shock

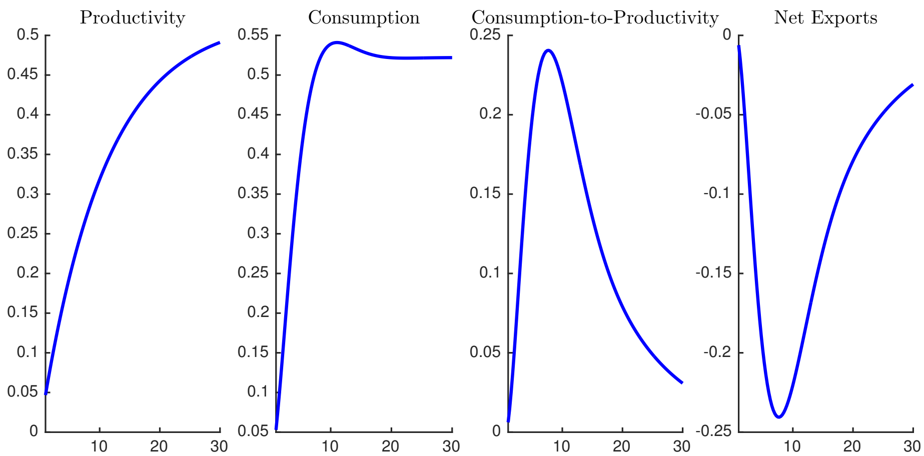

Notes: Impulse response functions of, from left to right, a_t, $\hat{c}_t$, c_t, and nx_t. Units in the vertical axis are percentage relative deviations from steady state in the case of a_t, $\hat{c}_t$, c_t, and absolute percentage deviations from steady state in the case of nx_t. The time unit on the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

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:

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

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

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 is identified by the standard deviation of the growth rates of productivity Δa_t. Identification of ρ comes from equation (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 equation (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.[16] 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.[17] 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

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

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

($\sigma_\nu \longrightarrow \infty$), 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.[18] 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 Preliminary Look at the Data: 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.

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

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

The Ratio of Consumption-to-Productivity. Figure 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.

Figure 3: Consumption-to-Productivity Ratio

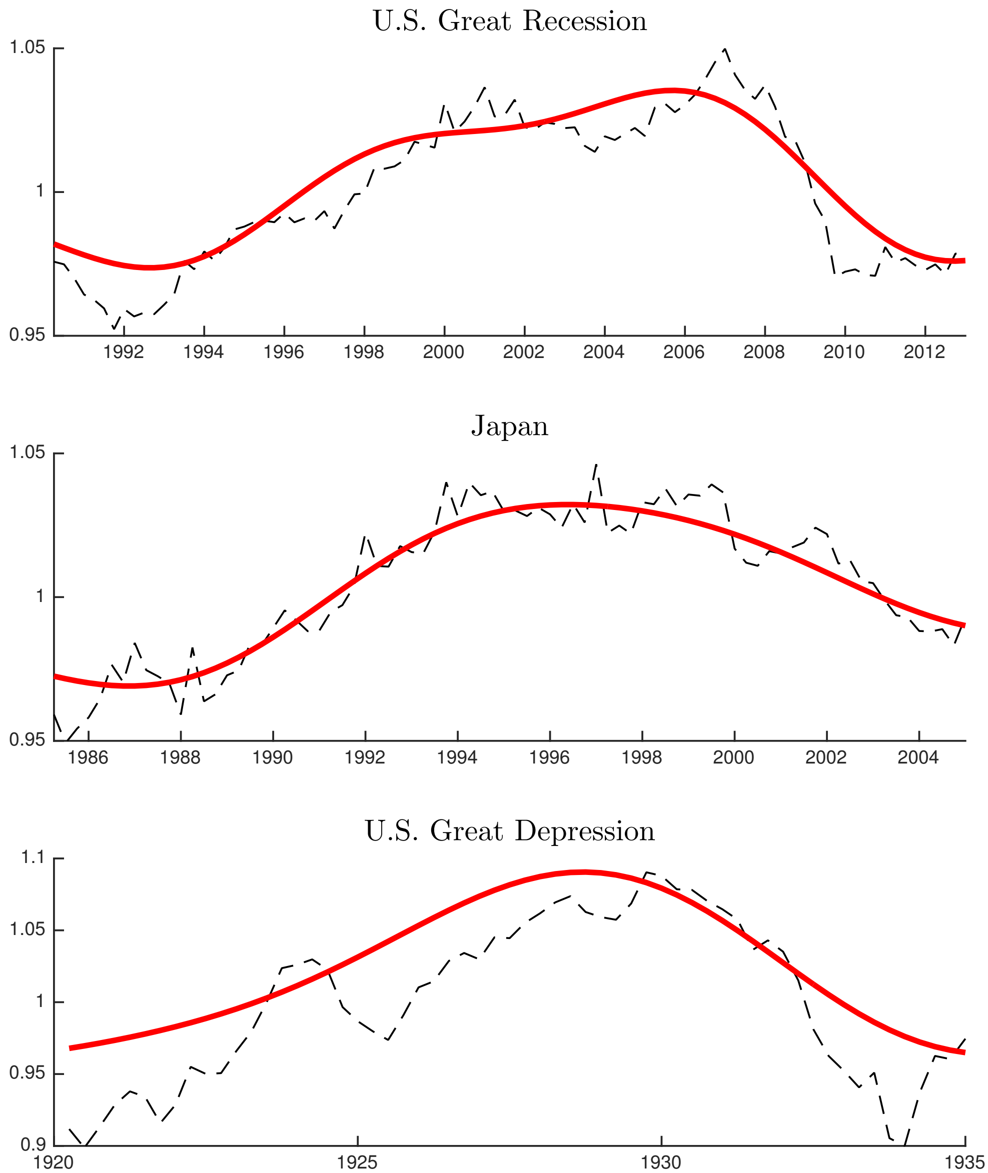

Notes: Productivity is real GDP divided by employment. Consumption is NIPA consumption divided by population. In the model the ratio is $c_t = \hat{c}_t - a_t$. Black-dashed line: raw data. Red-solid line: smoothed data using a band pass filter at 32-200 frequencies.

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).[19] 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

[19]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.

variables.[20] 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 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%.[21] 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.

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.

[20]Using the ratio of net exports-to-GDP nx_t instead of consumption does not qualitatively change the results (see our previous working paper Cao and L'Huillier 2012.)

[21]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).

4.2 Variance Decomposition and Estimated Permanent Shocks

Figure 8 of 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. Figure 4 shows our estimated permanent shocks in the case of the Great Recession.[22] 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.[23]

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 Figure 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

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

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

responsible for the medium-run variation in productivity growth.

Figure 4: Smoothed Permanent Shocks and Implied Productivity Growth, Great Recession

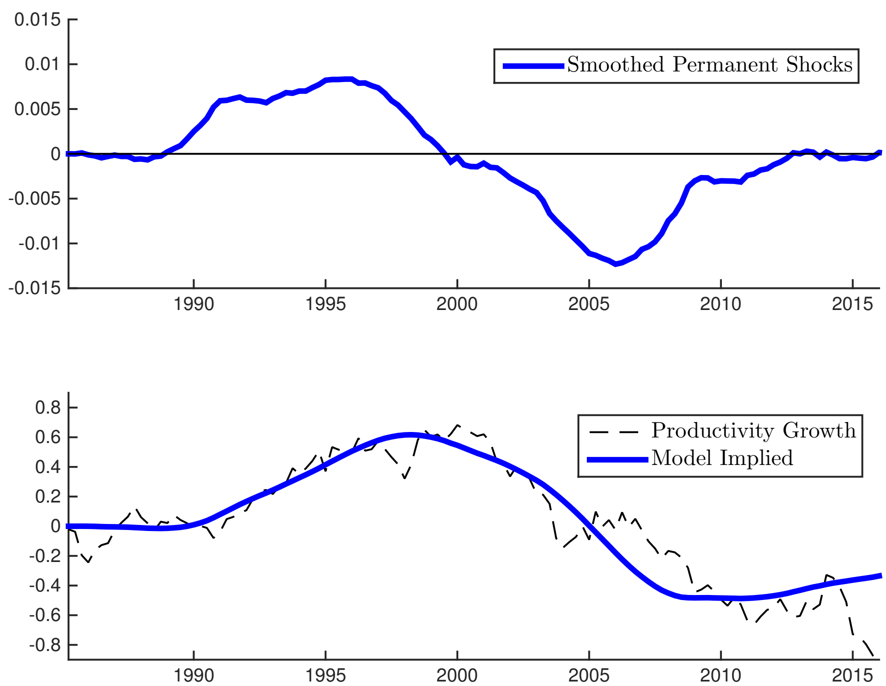

Notes: In the upper panel, smoothed permanent shocks are estimated using a Kalman smoother on the Great Recession sample. In the lower panel, the dotted-blue line represents productivity growth implied by the smoothed permanent shocks and the dashed-black line represents the centered 10-year moving average of the (demeaned) first differences of the logarithm of labor productivity.

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 Reproduced by the Estimated Permanent Shocks

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?

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).[24] 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.

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 Figure 4). The moments for the data are calculated using the time series for the same variables after taking centered 10-year moving averages.

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

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

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 (Figure 5). As claimed in Section 3, the model delivers a ratio that increases initially and then decreases, in a slow-moving hump shape.

Figure 5: Model-Implied Ratio of Consumption-to-Productivity, Great Recession

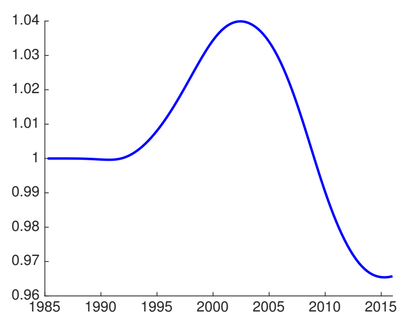

We finally present two supplementary out-of-sample checks of our estimation. We simulate the path of key endogenous variables using the permanent shocks only and compare these to the data. We focus on BLR and net exports.

Our beliefs data come from a survey published by Consensus Forecasts[25]. This survey was used in the paper by Hoffmann, Krause, and Laubach (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 sample.) Figure 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

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

optimistic between 00 and 05.[26] The right panel of Figure 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.

Figure 6: Out-of-Sample Checks: BLR and Net Exports

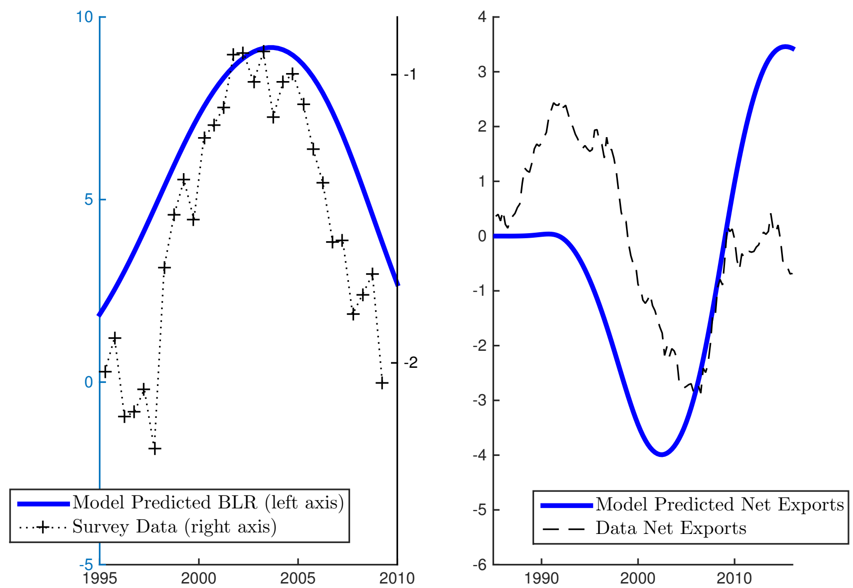

Notes: In the left panel, the dotted-black line represents survey data, reproduced from Hoffmann, Krause, and Laubach (2013) (Figure 1, bottom panel), and the solid-blue line represents model-implied BLR, generated by the permanent shock estimates (upper panel in Figure 4). In the right panel, the dashed-black line represents moving-average of net exports (centered 10-year), and the solid-blue line represents model implied net exports.

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.

[26]Estimating our model using net exports instead of consumption delivers the same result (see our previous working paper Cao and L'Huillier 2012.)

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 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.

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.[27] 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 Figure 10, right panel, in the supplementary material), and the 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).[28]

For brevity, we also report the implied path of the ratio of consumption-to-productivity in these two cases in the appendix (Figure 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.)

[27]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.

[28]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), among others.

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 Figure 4). The moments for the data are calculated using the time series for the same variables after taking centered 10-year moving averages.

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:

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

where N is normalized to 1.

The budget constraint is now

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

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, $\sigma_{\Delta a}$, σ_{ν}), fixing the other parameters defining the production function and capital dynamics at standard values.[29] 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 appendix.) Figure 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

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

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.

Figure 7: Investment

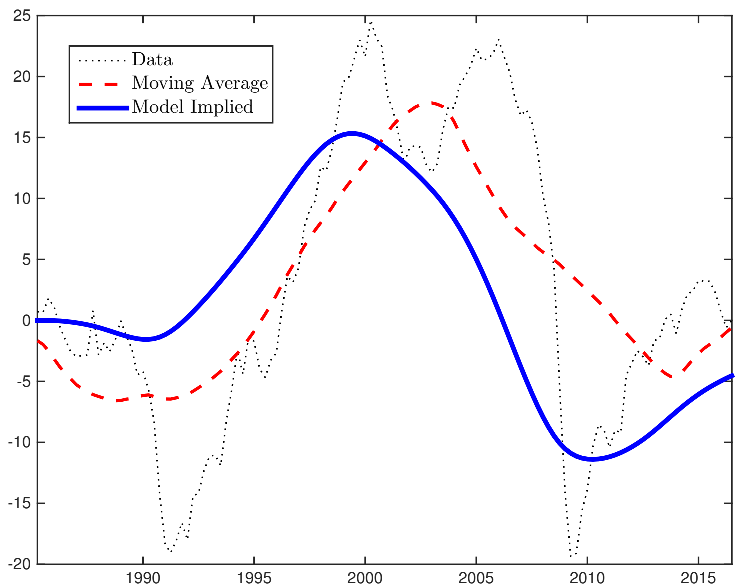

Notes: The series for investment (dotted-black) was constructed following Justiniano et al. (2010). The dashed-red line represents its centered 10-year moving average. The solid-blue line represents the model implied investment. Units on the right axis corresponds to percentage deviations from trend or steady-state.

5 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.

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. 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 (rather short-run) frictions to a similar medium-run permanent income framework.

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.

References

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

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.

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. (1997). The medium run. *Brookings Papers on Economic Activity 2.*

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. 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 (2012). Technological revolutions and debt hangovers: Is there a link? *EIEF Working Paper 16/12*.

Cao, D., J.-P. L'Huillier, and D. Yoo (2016). When is the trend the cycle? *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.*

Edge, R. M., T. Laubach, and J. C. Williams (2007). Learning and shifts in long-run productivity growth. *Journal of Monetary Economics 54*(8), 2421–2438.

Evans, G. W., S. Honkapohja, and P. Romer (1998). Growth cycles. *The American Economic Review 88*(3), 495–515.

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

Fernald, J. (2014). Productivity and potential output before, during, and after the great recession. *NBER Macroeconomic Annual 28.*

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

Forni, M., L. Gambetti, M. Lippi, and L. Sala. Noisy news in business cycles. *American Economic Journal: Macroeconomics, forthcoming*.

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

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

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 (2013). The expectations driven U.S. current account. *Deutsche Bundesbank Discussion Paper 10/2013*.

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.

Kahn, J. A. and R. W. Rich (2007). Tracking the new economy: Using growth theory to detect changes in trend productivity. *Journal of Monetary Economics* (54), 1670–1701.

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

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

Schmitt-Grohe, S. and M. Uribe (2003). Closing small open economy models. *Journal of International Economics 61*(1), 163–185.

A 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, series ID GDPC1). 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 IMPGSC96 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 PCECC96). The series for **TFP** was downloaded from John Fernald's website ("A Quarterly, Utilization-Adjusted Series on Total Factor Productivity", Fernald 2012, supplement, series dtfp_util). The series for Real Investment Expenditures is from the St. Louis Fed database (series ID GPDIC1). The series on Real Personal Durable Consumption Expenditures is from the St. Louis Fed database.

In the case of Japan, the series for **productivity** and **net exports** are 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
April 2017

B Model Solution, Proofs, and Other Results

B.1 Definitions

In order to log-linearize the model, we also define the following endogenous variables r_t, b_t, and nx_t:

$$r_t \equiv R_t - R \quad ,$$

and

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

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

NX/Y is the steady-state value of the net exports-to-output ratio.

B.2 Steady State

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

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

$$1 = \beta R \frac{A}{A_+} \quad . \tag{18}$$

Since $A + /A = 1$,

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

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

$$\bar{c} + \bar{b} = 1 + \beta \bar{b} \quad ,$$

this implies

$$\bar{c} = \left(1 - (1 - \beta)\, \bar{b}\right) \quad .$$

B.3 Log-Linearization

This equilibrium is given by the equations for the shock processes (1), (2), and (3), and other four equations:

$$c_t = -r_t + \mathbb{E}_t[c_{t+1} + \Delta a_{t+1}] \quad , \tag{19}$$

$$r_t = \psi \cdot b_t \quad , \tag{20}$$

$$c_t + \frac{1}{C/Y} nx_t = 0 \quad , \tag{21}$$

$$nx_t = b_{t-1} - \beta b_t + \frac{B}{Y}\left(-\Delta a_t + \beta^2 r_t\right) \quad . \tag{22}$$

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 (19). Log-linearizing the interest-elasticity equation (11) immediately gives (20).

Approximating the resource constraint delivers

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

which leads to (21).

Net exports are

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

and therefore, approximating

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

to obtain (22).

B.4 Closed-form Solution

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

$$\hat{b}_t = b_t + \frac{B}{Y} a_t \quad .$$

From the intertemporal budget constraint (22), together with the budget constraint (21), we have:

$$\begin{aligned}
\hat{b}_t &= b_t + \frac{B}{Y} a_t \\
&= \frac{1}{\beta} b_{t-1} - \frac{1}{\beta}\frac{C}{Y}(-c_t) + \frac{1}{\beta}\frac{B}{Y}\left(-\Delta a_t + \beta^2 r_t\right) \\
&\quad + \frac{B}{Y} a_t \\
&= \frac{1}{\beta}\hat{b}_{t-1} - \frac{1}{\beta}\frac{C}{Y}(-c_t) + \frac{1}{\beta}\frac{B}{Y}\left(-a_t + \beta^2 r_t\right) \\
&\quad + \frac{B}{Y} a_t \\
&= \frac{1}{\beta}\hat{b}_{t-1} + \frac{1}{\beta}\frac{C}{Y} c_t - \frac{B}{Y}\frac{1-\beta}{\beta} a_t + \frac{B}{Y}\beta r_t
\end{aligned}$$

Substituting r_t from (20) into the last equality, and also using the definition of $\hat{c}_t$, we arrive at

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

So

$$\hat{b}_t\left(1-\frac{B}{Y}\beta\psi\right) = \frac{1}{\beta}\hat{b}_{t-1}+\frac{1}{\beta}\frac{C}{Y}\hat{c}_t$$
$$-\left(\frac{1}{\beta}-\beta\psi\left(\frac{B}{Y}\right)^2\right)a_t$$

From the Euler equation (19), we have

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

We conjecture and verify later that

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

where the state variable $\mathbf{X}_t$ is defined as

$$\mathbf{X}_t = \begin{bmatrix} a_t & x_{t|t} & x_{t-1|t} & z_t\end{bmatrix}',$$

and solve for the coefficients D_b and D_k using the method of undetermined coefficients.

Indeed, from the Euler equation:

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

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

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

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

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

and

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

The first equation is a quadratic equation in D_b:

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

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

$$D_b = \frac{-\left(\frac{1}{C/Y} - \left(1 - \frac{B}{Y}\beta\psi\right) \beta \frac{1}{C/Y} - \psi\right) - \sqrt{\left(\frac{1}{C/Y} - \left(1 - \frac{B}{Y}\beta\psi\right) \beta \frac{1}{C/Y} - \psi\right)^2 + 4\psi \frac{1}{C/Y}}}{2} \quad . \tag{23}$$

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

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

So

$$D_{k,1} = \frac{\psi \frac{B}{Y} - (D_b - \psi) \frac{1}{1 - \frac{B}{Y}\beta\psi} \left(\frac{1}{\beta} - \beta\psi \left(\frac{B}{Y} \right)^2 \right)}{\left(1 - (D_b - \psi) \frac{1}{1 - \frac{B}{Y}\beta\psi} \frac{1}{\beta} \frac{C}{Y} \right)} \quad . \tag{24}$$

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

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

so

$$D_{k,4} = \frac{\rho D_{k,1}}{1 - (D_b - \psi) \frac{1}{1 - \frac{B}{Y}\beta\psi} \frac{1}{\beta} \frac{C}{Y} - \rho} \quad . \tag{25}$$

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

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

and

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

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 $\tilde{x} = (D_b - \psi) \frac{1}{1 - \frac{B}{Y}\beta\psi} \frac{1}{\beta} \frac{C}{Y}$. Thus

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

B.5 Limit Result for Consumption

First, we notice that from the steady-state equations,

$$\frac{C}{Y} = 1 - (1-\beta)\bar{b}$$

which goes to 1 as $\beta \to 1$.

From the expression for D_b, (23), we have

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

This implies

$$\lim_{\beta \to 1} \lim_{\psi \to 0} D_b = 0.$$

From the expression for $D_{k,1}$,(24):

$$\lim_{\psi \to 0} D_{k,1} = \frac{-\frac{\lim_{\psi \to 0} D_b}{\beta}}{1 - \frac{\lim_{\psi \to 0} D_b}{\beta} \frac{C}{Y}} = \left(1-\beta\right) \frac{1}{C/Y} \quad .$$

This implies

$$\lim_{\beta \to 1} \lim_{\psi \to 0} D_{k,1} = 0.$$

Similarly, from the expression for $D_{k,4}$, (25) we obtain:

$$\lim_{\beta \to 1} \lim_{\psi \to 0} D_{k,4} = 0.$$

From the expression for $\tilde{x}$:

$$\lim_{\psi \to 0} \tilde{x} = \lim_{\psi \to 0} D_b \frac{1}{\beta} = -\frac{1-\beta}{\beta}.$$

Therefore,

$$\lim_{\beta \to 1} \lim_{\psi \to 0} \tilde{x} = 0$$

and

$$\lim_{\beta\to 1}\lim_{\psi\to 0}\frac{D_{k,1}}{\tilde{x}} = -1.$$

Combining these limits with the expression for $D_{k,2}, D_{k,3}$, we have:

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

These limits imply the desired expression for $\hat{c}_t$.

B.6 VAR Representation of the Limiting Model

In this section we derive (13) and (14). 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

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

Notice that

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

so

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

and

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

Therefore we have the VAR representation (13) and (14).

Equation (16) is obtained by induction in j. We just showed it holds for $j = 0$. If it holds for j, then $\mathbb{E}_t\left[a_{t+j}\right] = \rho^j a_t + \left(1-\rho^j\right)\hat{c}_t$. Taking expectations at time $t-1$

on both sides yields

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

the second equality follows from (13) and (14), the third from rearranging.

B.7 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:

$$R_t = R_t^* + \psi \left\{ e^{\frac{B_t}{Y_t} - \bar{b}} - 1 \right\} \tag{26}$$

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 (26).

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 (20), by:

$$r_t = r_t^* + \psi \cdot b_t \quad . \tag{27}$$

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

$$mb_t + m^* b_t^* = 0 \quad . \tag{28}$$

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.

C Supplementary Material on Estimations

Here we present all results skipped in Section 4.

Variance Decomposition. Figure 8 shows the forecast error explained by each shock for the model estimated for the Great Recession case.

Figure 8: Variance Decomposition of BLR at Different Horizons

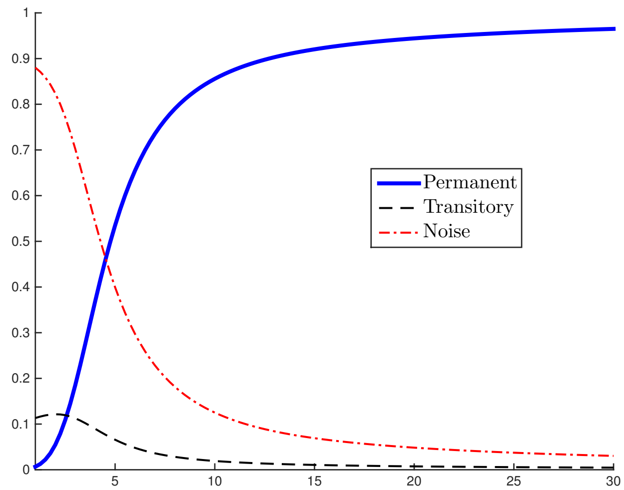

Notes: Percentage of forecast error explained by each shock.

Other shocks: Great Recession. For completeness, we report our estimated transitory and noise shocks.

Figure 9 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 11 below plots these shocks for Japan and the Great Depression. Similarly, these shocks do not have any particular pattern either.

Parameter Estimates: Japan and Great Depression. The persistence parameter ρ is estimated at 0.96 in the case of Japan, and at 0.94 in the case of the Great Depression, quite similar to the one obtained for the Great Recession. Both values imply persistent processes both for the permanent and the transitory components of productivity. The standard deviation of productivity is estimated at 1.07% in the case of Japan, and at 1.80% 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

Figure 9: Smoothed Transitory and Noise Shocks, Great Recession

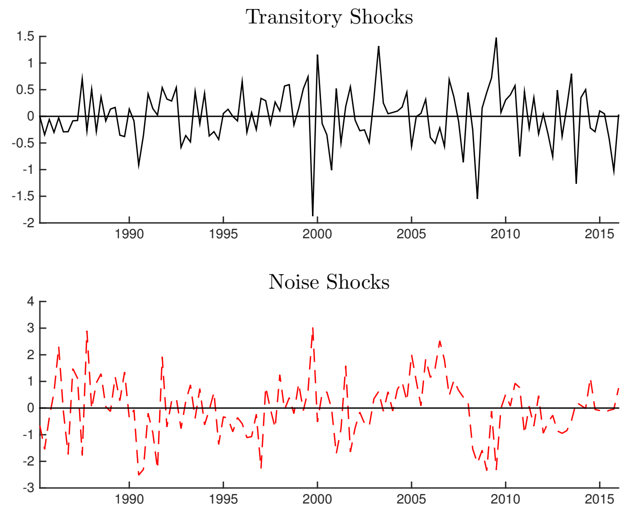

Notes: Shocks estimated using a Kalman smoother on the Great Recession sample. The observation equation is composed of the first differences of the logarithm of U.S. labor productivity and consumption. The values for σ_ε and σ_η are implied by the random walk Assumption (6) for productivity. The units on the y-axis are percentages.

deviation for permanent technology shocks of 0.06% in the case of Japan, and of 0.09% in the case of the Great Depression, and a standard deviation for the transitory technology shock of 1.05% in the case of Japan, and of 1.74% in the case of the Great Depression. The standard deviation of noise shocks is also larger than obtained before, 8.26% and 8.90% respectively, in part due to the larger standard deviation of productivity growth σ_a.

Table 4: Parameter Estimates, Japan and Great Depression

		Japan		Great Dep.	
Parameter	Description	Value	s.e.	Value	s.e.
ρ	Persistence tech. shocks	0.96	0.01	0.95	0.01
σ_a	Std. dev. productivity	1.07	0.07	1.80	0.14
σ_ε	Std. dev. permanent tech. shock (implied)	0.04	–	0.09	–
σ_η	Std. dev. transitory tech. shock (implied)	1.05	–	1.74	–
σ_ν	Std. dev. noise	8.26	3.91	8.90	3.14

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 consumption. 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.

Smoothed Permanent Shocks: Japan and Great Depression. The variance decomposition for these two cases (non reported) is similar to the one shown in Figure 8. Figure 10 plots the estimated permanent shocks. In the case of Japan we estimate, as in the case of the Great Recession. positive shocks in the first part of the samples, and negative shocks later on. In the case of Japan, the positive shocks hit roughly between 1980 and 1987. These estimated permanent shocks imply that we should have observed a productivity acceleration and deceleration. This is shown in Table 3. In the case of the Great Depression, we estimate only a few positive shocks at the beginning of the sample (this is a feature of the data that we comment on in the body.) Most importantly, we estimate negative shocks *before* the Great Depression starts, which is consistent with the pre-Great Recession productivity growth decline and the slow adjustment of beliefs about the future.

Figure 10: Smoothed Permanent Shocks, Japan and Great Depression

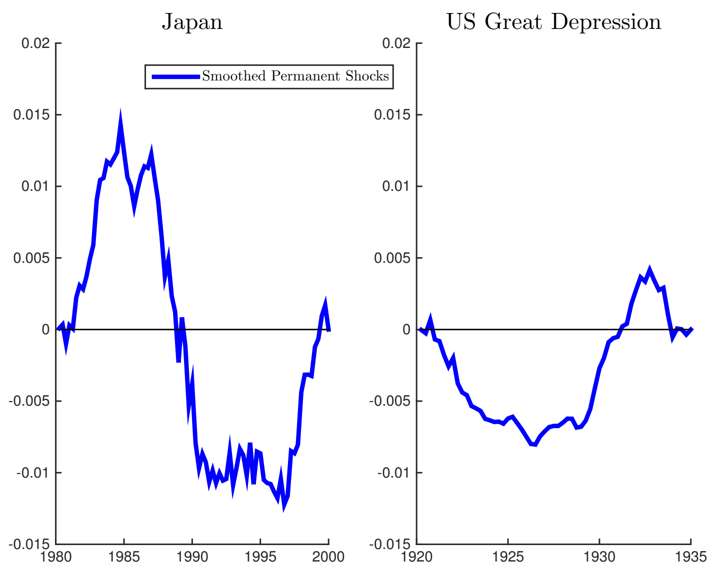

Notes: Shocks estimated using a Kalman smoother on the Japanese sample, and on the U.S. Great Depression sample. The observation equation is composed of the first differences of the logarithm of labor productivity and consumption. The units on the y-axis are percentages.

Figure 11: Smoothed Transitory and Noise Shocks, Japan and Great Depression

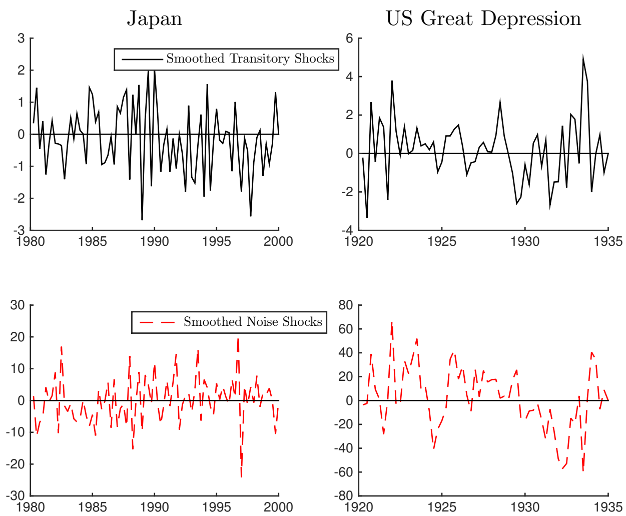

Notes: Shocks estimated using a Kalman smoother on the Japanese sample, and on the U.S. Great Depression sample. The observation equation is composed of the first differences of the logarithm of labor productivity and consumption. The units on the y-axis are percentages.

Figure 12: Model-Implied Ratio of Consumption-to-Productivity, Japan and Great Recession

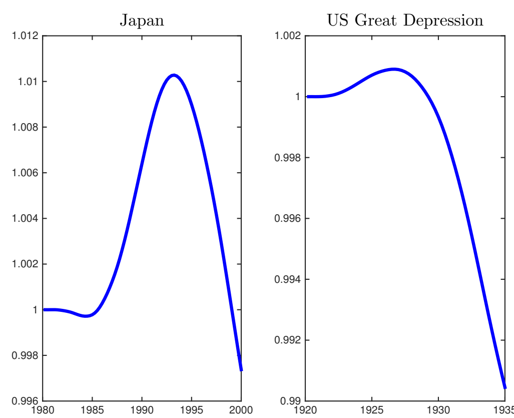

D Model with Investment

D.1 First-Order Conditions

Let Φ_t denote the Lagrange multiplier on the capital accumulation equation (17). The first-order conditions that characterize the equilibrium are:

1. F.O.C in B_t:

$$\frac{1}{C_t} = \beta R_t \mathbb{E}_t \left[\frac{1}{C_{t+1}} \right] \tag{29}$$

2. F.O.C in K_t :

$$\Phi_t = \beta \mathbb{E}\left[\frac{1}{C_{t+1}} \alpha K_t^{\alpha-1} \left(A_{t+1} N\right)^{1-\alpha} + (1-\delta)\Phi_{t+1} \right.$$
$$\left. + \frac{1}{C_{t+1}} \frac{\chi}{2} \frac{\left(I_{t+1} - \delta K_t\right)^2}{K_t^2} + \frac{1}{C_{t+1}} \frac{\chi}{2} \frac{2\delta \left(I_{t+1} - \delta K_t\right)}{K_t}\right] \tag{30}$$

3. F.O.C in I_t:

$$\Phi_t = \frac{1}{C_t}\left(1 + \chi \frac{I_t - \delta K_{t-1}}{K_{t-1}}\right). \tag{31}$$

D.2 Solution

Steady-State. The first-order conditions yields the following equations that determine the steady-state of this economy:

$$\beta R = 1$$

and

$$\alpha A K^{\alpha-1} = \beta(1-\delta)$$

and

$$\Phi = \frac{1}{C}.$$

Log-linear Approximation. As in the baseline model, we work with the following normalized variables:

$$c_t = \log\left(\frac{C_t}{A_t}\right) - \log\left(\frac{C}{A}\right)$$

and

$$k_t = \log\left(\frac{K_t}{A_t}\right) - \log\left(\frac{K}{A}\right)$$

and

$$y_t = \log\left(\frac{Y_t}{A_t}\right) - \log\left(\frac{Y}{A}\right)$$

and

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

and

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

and

$$r_t = R_t - R.$$

and

$$\phi_t = \log\left(\Phi_t A_t\right) - \log\left(\Phi A\right) \quad .$$

Using the normalized variables the production function is log-linearized as:

$$y_t = \alpha\left(k_{t-1} - \Delta a_t\right)$$

and the capital accumulation equation (17) is log-linearized as:

$$k_t = (1-\delta)\left(k_{t-1} - \Delta a_t\right) + \delta\left(d_t + i_t\right) \quad .$$

Log-linearizing the F.O.Cs yields the following linear equations.

F.O.C. (29):

$$c_t = -r_t + \mathbb{E}_t\left[\Delta a_{t+1} + c_{t+1}\right]$$

F.O.C. (30):

$$\begin{aligned}
&\Phi A \phi_t \\
&= \beta \mathbb{E}_t\left[-\alpha \Delta a_{t+1} \frac{A}{C} \alpha \left(\frac{K}{A}\right)^{\alpha-1} - c_{t+1} \frac{A}{C} \alpha \left(\frac{K}{A}\right)^{\alpha-1} + (\alpha-1)\, k_t \frac{A}{C} \alpha \left(\frac{K}{A}\right)^{\alpha-1}\right] \\
&+ \beta \mathbb{E}_t\left[(1-\delta)\Phi A\left(\phi_{t+1} - \Delta a_{t+1}\right) + \frac{A}{C}\chi\delta\left(i_{t+1} + \Delta a_{t+1} - k_t\right)\right]
\end{aligned}$$

Notice that from the steady-state equation, $\Phi = \frac{1}{C}$, this equation simplifies to

$$\begin{aligned}
&\phi_t \\
&= \beta \mathbb{E}_t\left[-\alpha \Delta a_{t+1} \alpha \left(\frac{K}{A}\right)^{\alpha-1} - c_{t+1} \alpha \left(\frac{K}{A}\right)^{\alpha-1} + (\alpha-1)\, k_t \alpha \left(\frac{K}{A}\right)^{\alpha-1}\right] \\
&+ \beta \mathbb{E}_t\left[(1-\delta)\left(\phi_{t+1} - \Delta a_{t+1}\right) + \chi\delta\left(i_{t+1} + \Delta a_{t+1} - k_t\right)\right]
\end{aligned}$$

F.O.C. in I (31):

$$\Phi A(\phi_t + d_t) = -\frac{A}{C} c_t + \frac{A}{C}\chi \frac{\frac{I}{A}\left(i_t - \Delta a_t\right) - \delta \frac{K}{A} k_{t-1}}{\frac{K}{A}} \quad .$$

Again, because $\Phi = \frac{1}{C}$:

$$\phi_t + d_t = -c_t + \chi\delta\left(i_t - \Delta a_t - k_{t-1}\right) \quad .$$

From the definition of nx_t

$$nx_t = \frac{Y_t - G(I_t, K_{t-1}) - C_t}{Y_t} - \frac{Y - G(I, K) - C}{Y} \quad ,$$

which yields:

$$\begin{aligned}
nx_t &= -G_1\left(\frac{I}{A},\frac{K}{A}\right)\frac{I}{A}i_t\frac{A}{Y} - G_2\left(\frac{I}{A},\frac{K}{A}\right)\frac{K}{A}\left(k_{t-1}-\Delta a_t\right)\frac{A}{Y} + \frac{G(I,K)}{Y}y_t - \frac{C}{Y}c_t + \frac{C}{Y}y_t \\
&= -\frac{I}{Y}i_t + \frac{G(I,K)}{Y}y_t - \frac{C}{Y}c_t + \frac{C}{Y}y_t \quad . \\
&= -\frac{I}{Y}i_t + \frac{I}{Y}y_t - \frac{C}{Y}c_t + \frac{C}{Y}y_t \quad .
\end{aligned}$$

Using NX_t we rewrite the budget constraint as

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

After dividing by Y_t, we obtain:

$$\frac{B_{t-1}}{Y_{t-1}}\frac{Y_{t-1}}{A_{t-1}}\frac{A_{t-1}}{A_t}\frac{A_t}{Y_t} = \frac{NX_t}{Y_t} + Q_t\frac{B_t}{Y_t}$$

so

$$b_{t-1} + \bar{b}\left(y_{t-1} - \Delta a_t - y_t\right) = nx_t + \beta b_t - \bar{b}\beta^2 r_t \quad .$$