

Learning and Structural Change in Macroeconomic Data

James Bullard* John Duffy†

16 August 2004‡

Abstract

We include learning in a standard equilibrium business cycle model with explicit growth. We use the model to study how the economy's agents could learn in real time about the important trend-changing events of the postwar era in the U.S., such as the productivity slowdown, increased labor force participation by women, and the "new economy" of the 1990s. We find that a large fraction of the observed variance of output relative to trend can be attributed to structural change in our model. However, we also find that the addition of learning and occasional structural breaks to the standard and widely-used growth model results in a *balanced growth puzzle*, as our approach cannot completely account for observed trends in U.S. aggregate consumption and investment. Finally, we argue that a model-consistent detrending approach, such as the one we suggest here, is necessary if the goal is to obtain an accurate assessment of an equilibrium business cycle model.

Keywords: Business cycle theory, learning, structural change, new economy, productivity slowdown.

JEL Classification Codes: E2, E3.

*Research Department, Federal Reserve Bank of St. Louis, 411 Locust Street, St. Louis, MO 63102 USA. Telephone: (314) 444-8576. Email: bullard@stls.frb.org. Any views expressed are those of the authors and do not necessarily reflect the views of the Federal Reserve Banks of St. Louis or the Federal Reserve System.

†Department of Economics, University of Pittsburgh, Pittsburgh, PA 15260. Telephone: (412) 648-1733. Email: jduffy@pitt.edu.

‡The initial version of this paper was prepared for a conference on "Economic Dynamics," at the Lorentz Center in Leiden, The Netherlands, under the title "Learning to Detrend Macroeconomic Data." We thank Cars Hommes, William Brock, and the Center for Nonlinear Dynamics in Economics and Finance (CeNDEF) at the University of Amsterdam for organizing this event. We also thank participants at subsequent seminars, especially discussant Karsten Jeske, for helpful comments.

1 Introduction

1.1 Overview

It is widely acknowledged that important structural changes occurred during the postwar era in the U.S. and other industrialized countries. A significant slowdown in productivity growth occurred beginning in the late 1960s or early 1970s, and some researchers find a significantly faster growth rate for productivity during the “new economy” era beginning in the mid- to late-1990s. Similarly, women are known to have increased their labor force participation rates beginning in the 1960s. Perron (1989) and Hansen (2001) discuss some of the econometric evidence for characterizing macroeconomic data with log-linear trends coupled with occasional structural change. They find, broadly speaking, that trend stationarity interrupted by some trend breaks provides a good empirical model for U.S. macroeconomic time series. In this paper, we take this evidence at face value and try to build models that are consistent with it.

Much of equilibrium business cycle analysis abstracts from permanent changes in trend growth paths (and, indeed, from growth itself). This includes a wide class of models ranging from the original real business cycle papers to the more recent New Keynesian macroeconomics. In nearly all of this work, the economy is viewed as essentially following a given balanced growth path, deviating from that path only because of temporary shocks which drive the business cycle. The path itself never changes. If it did, the agents in the model would want to react to such movements. In this paper we build a model that takes account of important trend-changing events in a model-consistent way. We provide one method of understanding the influence of structural change on business cycle fluctuations.

1.2 Model summary

We study a version of a simple and standard equilibrium business cycle model, namely, King, Plosser, and Rebelo (1988a), in which we explicitly allow for growth driven by two exogenous sources: productivity improvements and increases in labor input. We replace the rational expectations assumption with a recursive adaptive learning assumption following the methodology of Evans and Honkapohja (2001). Our assumption involves a “constant gain” learning algorithm, which discounts past data and allows the agents to remain alert to the possibility of structural change. We verify that the economy is stable under this learning assumption, meaning that, if there are no changes in the underlying parameters for a period of time, the economy will remain in a small neighborhood of the balanced growth path as if all agents had rational expectations all the time.

We then subject the economy under learning to two kinds of shocks, the standard business cycle shocks to total factor productivity as well as a few unexpected and perfectly persistent shocks to the factors driving growth; the latter shocks correspond to postwar U.S. events such as changing attitudes concerning women in the workforce, the “productivity slowdown,” and the “new economy.” These perfectly persistent shocks occur only once or twice in fifty years, and so it is reasonable to think that they are completely unanticipated and that agents must learn about them. When these shocks occur, the agents adjust to a new balanced growth path and learn the new rational expectations equilibrium. Thus in our model, agents are able to track a balanced growth path that is sometimes changing, while simultaneously reacting to ordinary business cycle shocks. When the ordinary business cycle shock variance is reduced to a negligible level, we are able to trace out the multivariate trend implied by the model with learning. We then remove this same multivariate, broken trend from the actual data as well as from the data generated by the model. We therefore provide a model-consistent approach to detrending the macroeconomic data. We calculate business cy-

cle statistics and discuss related issues concerning the performance of the model.

1.3 Trend-cycle decomposition via statistical filters

Trend-cycle decomposition is an issue that has plagued equilibrium business cycle research, and our model-consistent approach can address some of the issues in this area. When comparing models to the data, the discipline implied by the assumption that the economy is following a balanced growth path is often discarded. Instead, atheoretic, statistical filters are typically employed to detrend the actual data, and render it stationary.¹ This approach has been widely criticized, for instance by Cogley and Nason (1995a), Harvey (1997) and Canova (1998a). The criticisms are not hard to digest: (1) Statistical filters do not remove the same trend from the data that the balanced growth path of the model implicitly requires; (2) The “business cycle facts” are not independent of the statistical filter employed; (3) The data are often detrended one variable at a time while the model implies a multivariate trend—thus the methodology does not respect the cointegration of the variables that the model requires; (4) The filtered trends imply that trend growth rates sometimes change, but the agents in the model are not allowed to react to these trend movements.

Our methodology goes some way towards addressing these concerns. Under our model-consistent method, the trends we remove from the data will be exactly the same ones that are implied by our model. We allow the agents to react to changes in trend growth rates and we respect the cointegration of the variables that the model implies. We do this in the simplest context available for this issue, but we think our methodology has wide applicability

¹For a detailed recent discussion, see King and Rebelo (1999). Examples of statistical filtering techniques include the Hodrick–Prescott (1980) (HP) filter, Beveridge and Nelson’s (1981) decomposition procedure and the band pass filters used by Stock and Watson (1990), Baxter and King (1999), and Christiano and Fitzgerald (1999). See Canova (1998a) for a more exhaustive list.

across a range of growth and business cycle models.

1.4 Main findings

Adding structural change to the standard equilibrium business cycle model means that a new type of shock, albeit a rare one, has been included. We find that a large fraction of the observed variance of output relative to trend can be attributed to this shock. Prescott (1986) and Kydland and Prescott (1991) have argued that models closely related to the one we analyze can explain 70 to 75 percent of the business cycle variation in real output. Our analysis suggests that the remainder of the variation may be due, not to monetary or fiscal policy, but to structural change.

We also identify a *balanced growth puzzle*. According to our analysis, the balanced growth path dictated by productivity growth and growth in aggregate labor hours should have been characterized by more consumption and less investment over the period 1985 to 2001, compared to what was actually observed. This is in addition to changes in investment and consumption that might have occurred because of an increase in the growth rate of productivity, a “new economy,” which is already included in our model. We suggest a number of avenues we think would be interesting to investigate in future research regarding this puzzle.

Finally, we show that our model-consistent methodology allows us to detrend the data in a relatively smooth fashion. The trends we calculate are in some respects quite similar to those that would be calculated using available statistical filtering techniques. In this sense, we are able to provide some microfoundations for current practices in the equilibrium business cycle literature. We also show how business cycle statistics for both the model and the data are broadly consistent with the statistics which are commonly reported, when the data are detrended using the trends dictated by our model. There are some important differences, however, and we conclude that the detrending methodology is not innocuous for understanding fluctuations

in the data. A model-consistent approach like the one we suggest is necessary to accurately evaluate equilibrium business cycle models.

1.5 Recent related literature

The literature on detrending and the evaluation of equilibrium business cycle models is large. For critiques of the ability of technology-shock-driven equilibrium business cycle models to reproduce the data and a discussion of related detrending issues, see Cogley and Nason (1995ab) and Rotemberg and Woodford (1996). The debate between Canova (1998ab) and Burnside (1998) concerned the finding that different statistical filters in general yield a different set of business cycle facts. Canonical discussions of the business cycle facts can be found in Cooley and Prescott (1994), Stock and Watson (1999), and King and Rebelo (1999). King, Plosser, and Rebelo (1988a,b) discuss model-consistent detrending in the same spirit as we do. They investigate a model-consistent, linear trend in their Essay I; we essentially introduce trend breaks and learning into a similar model. Perron (1989) and Hansen (2001) discuss the econometric evidence for characterizing macroeconomic data with log-linear trends coupled with occasional structural change. The macroeconomics learning literature is summarized in Evans and Honkapohja (2001). Packalén's (2000) thesis studies expectational stability, or learnability, in business cycle models like the one we use. His main focus was on the theoretical stability of irregular equilibria. See Rotemberg (2003) for a recent discussion of the plausibility of assuming shocks to trends are independent of shocks that drive the business cycle. Rotemberg employs a "slow technological diffusion" assumption on the former shocks, an assumption we do not make use of here. For applications of learning about trends to issues in monetary policy, see Lansing (2000, 2002), Collard and Dellas (2004) and Bullard and Eusepi (2003). The effects of a change in trend productivity growth in a rational expectations environment are discussed in Pakko (2002).

2 Environment

2.1 Overview

We study a version of an equilibrium business cycle model with exogenous growth. We stress that our methodology could be applied to a wide variety of models in this general class.

Time is discrete and indexed by $t = 0, 1, 2, \dots$. The economy consists of many identical households, and the number of households is growing over time. These households make identical decisions, and so we will analyze them as if there was only one decisionmaker. We work in terms of aggregate variables, as opposed to per capita variables. We use capital letters to denote aggregates. Because we have growth explicitly in the model, the aggregate variables output, Y_t , consumption, C_t , investment I_t , and capital, K_t , will be nonstationary. We will transform these variables into their stationary counterparts in order to solve the model. When we do so, we denote the stationary variable by a small case, hatted letter, such as \hat{c}_t . With this notation in mind, we write the household problem as maximization of

$$E_t \sum_{t=0}^{\infty} \beta^t \eta^t \left[\ln C_t + \theta \ln (1 - \hat{\ell}_t) \right] \quad (1)$$

by choice of consumption and leisure at each date subject to constraints which apply at every date t :

$$C_t + I_t \leq Y_t, \quad (2)$$

$$I_t = K_{t+1} - (1 - \delta) K_t, \quad (3)$$

$$Y_t = \hat{s}_t K_t^\alpha \left(X_t N_t \hat{\ell}_t \right)^{1-\alpha}, \quad (4)$$

$$X_t = \gamma X_{t-1}, \quad X_0 = 1, \quad (5)$$

$$N_t = \eta N_{t-1}, \quad N_0 = 1, \quad (6)$$

and

$$\hat{s}_t = \hat{s}_{t-1}^\rho \epsilon_t, \quad \hat{s}_0 = 1, \quad (7)$$

where \hat{s}_t is the technology shock. The household has a time endowment of 1 at each date t , and $\hat{\ell}_t$ is the fraction of this endowment which is supplied to the labor market. The variable X_t is the level of labor-augmenting productivity, or number of efficiency units, in the economy; the growth in this variable will drive real per capita income higher over time. The variable N_t is the size of the labor force, or number of households, where the date 0 size is normalized to unity. The parameter $\beta \in (0, 1)$ is the household's discount factor, $\theta > 0$ controls the relative weight in utility placed on leisure, $\delta \in (0, 1)$ is the net depreciation rate, $\alpha \in (0, 1)$ is the capital share, $\gamma \geq 1$ is the gross rate of growth in productivity, $\eta \geq 1$ is the gross rate of labor force growth, and $\rho \in (0, 1)$ controls the degree of serial correlation in the technology shock. The standard expectations operator is denoted E_t . The stochastic term ϵ_t is *i.i.d.*, with mean of unity and variance of σ_ϵ^2 .

By combining constraints (2) and (3), and using constraint (4), we can write a Lagrangian for the household's problem. Using the first order conditions for this problem, we can write our system in terms of four equations determining C_t , $\hat{\ell}_t$, K_t , and Y_t (along with the definitions of \hat{s}_t , X_t , and N_t). In particular, combining (2) and (3) yields

$$K_{t+1} = Y_t + (1 - \delta) K_t - C_t, \quad (8)$$

output is produced according to

$$Y_t = \hat{s}_t \left[(K_t)^\alpha \left(X_t N_t \hat{\ell}_t \right)^{1-\alpha} \right], \quad (9)$$

and the first order conditions yield

$$C_t = \frac{1 - \alpha}{\theta} Y_t \left(\frac{1 - \hat{\ell}_t}{\hat{\ell}_t} \right), \quad (10)$$

as well as

$$\frac{1}{C_t} = \beta \eta E_t \left\{ \frac{1}{C_{t+1}} \left[\alpha Y_{t+1} K_{t+1}^{-1} + 1 - \delta \right] \right\}. \quad (11)$$

Our system is given by (8) through (11), along with (5), (6), and (7).

2.2 A linear representation

We now wish to transform equations (8) through (11) along with their definitional counterparts (5), (6), and (7) into a stationary, linearized system so that we may apply the techniques developed by Evans and Honkapohja (2001). We sketch the transformation here, which involves three main steps, and provide the details in Appendix A.

First, we transform equations (8) through (11) into a stationary system by replacing C_t , Y_t , and K_t as appropriate with variables of the form $\hat{c}_t = C_t/(X_t N_t)$, and so on. The hatted variables are therefore in per total efficiency unit terms. The resulting system has a nonstochastic steady state which can be calculated directly. We can denote the steady state vector as $(\hat{c}_t, \hat{k}_t, \hat{\ell}_t, \hat{y}_t) = (\bar{c}, \bar{k}, \bar{\ell}, \bar{y})$, $\forall t$. An important feature of the steady state values is that they depend on all parameters of the system, in general, and in particular on the parameters γ and η . Thus for example, a change in the gross growth rate of productivity, γ , will alter the nonstochastic steady state of the system, as well as important ratios such as the consumption-output ratio or the capital-output ratio.

Next, we linearize about the steady state, using a differences in logarithms approach with variables of the form $\tilde{c}_t = \ln(\hat{c}_t/\bar{c})$, and so on. This step requires additional, standard, approximations which are given in detail in Appendix A. However, the linearized system, written in terms of logarithmic deviations from steady state, is not satisfactory for our purposes. The tilde variables involve steady state values, such as \bar{c} , which, as we have noted above, depend on the growth rates of productivity and the labor input. If we allow agents to learn by estimating a VAR using $(\tilde{c}_t, \tilde{k}_t, \tilde{\ell}_t, \tilde{y}_t)$, then we would in effect be telling them when a change in the steady state had occurred, which is inconsistent with our wish to allow them to learn about such unexpected changes.

Consequently, as a final step we decompose the tilde variables by defining variables of the form $c_t = \ln \hat{c}_t$ and $c = \ln \bar{c}$, and so on. We then collect

all terms involving c , k , ℓ , and y into constant terms in each of the four equations. We then require that agents estimate these constant coefficients together with the coefficients on the endogenous variables of the model as discussed below; thus, agents will have to learn the new steady state values of the system that change whenever the growth rates γ or η change unexpectedly. Finishing up, we reduce the four equations down to two, defined in terms of c_t and k_t .

Following these transformations, the system can be written as

$$c_t = \mathcal{B}_{10} + \mathcal{B}_{11}E_t c_{t+1} + \mathcal{B}_{12}E_t k_{t+1} + \mathcal{B}_{13}E_t s_{t+1}, \quad (12)$$

$$k_t = \mathcal{D}_{20} + \mathcal{D}_{21}c_{t-1} + \mathcal{D}_{22}k_{t-1} + \mathcal{D}_{23}s_{t-1}, \quad (13)$$

$$s_t = \rho s_{t-1} + \vartheta_t, \quad (14)$$

with $\vartheta_t = \ln \epsilon_t$, and where the coefficients $\mathcal{B}_{i,j}$, $\mathcal{D}_{i,j}$, $i = 1, 2$; $j = 0, 1, 2, 3$; are agglomerations of the underlying parameters of the model described in detail in Appendix A.²

3 Learning

3.1 The system under recursive learning

We study the system (12)-(14) under a recursive learning assumption, as discussed in Evans and Honkapohja (2001). We imagine that initially, agents have no specific knowledge of the economy in which they operate, other than the perceived law of motion with which they are endowed (which is given below). The agents we study will be able use this perceived law of motion to learn the rational expectations equilibrium of (12)-(14)—there is precisely one parameterization of this perceived law of motion that corresponds to the rational expectations equilibrium of the system under any parameterization of the model. We close the model under a learning assumption rather than rational expectations because our environment is prone to infrequent shocks

²See Packalén (1999) for similar representations of equilibrium business cycle models.

to growth factors that agents must learn about—permanent changes in the growth rates of productivity, γ , or the labor input, η . We view such shocks as occurring infrequently, perhaps only once or twice in fifty years. This lends plausibility to our assumption that such shocks are largely unanticipated and that agents must learn about them when they occur. Our model, then, is one where the economy follows a balanced growth path buffeted by the usual business cycle shocks, s_t , but where the balanced growth path itself changes course infrequently. The latter assumption, together with the assumption that agents are learning, is what differentiates our model from other equilibrium business cycle models. We think such a model is consistent with the time-series econometric evidence of Perron (1989) and others, namely, that postwar macroeconomic U.S. data can be rendered “trend stationary” with just a few changes to the exogenous, deterministic trend component of the model.

We begin our development of the model under learning by writing the linearized model in equation form as

$$c_t = \mathcal{B}_{10} + \mathcal{B}_{11}E_t^*c_{t+1} + \mathcal{B}_{12}E_t^*k_{t+1} + \mathcal{B}_{13}E_t^*s_{t+1} + \Delta_t \quad (15)$$

$$k_t = \mathcal{D}_{20} + \mathcal{D}_{21}c_{t-1} + \mathcal{D}_{22}k_{t-1} + \mathcal{D}_{23}s_{t-1} \quad (16)$$

$$s_t = \rho s_{t-1} + \vartheta_t \quad (17)$$

In this system, we have added a small shock, Δ_t , to the first equation. While one can think of Δ as a small shock to preferences, the primary role of this shock is to prevent perfect multicollinearity in the regressions run by the agents using capital and consumption data generated by the model; in equilibrium, consumption is a perfect linear combination of the capital stock and the productivity shock.³ The operator E_t^* indicates (possibly nonrational) expectations taken using the information available at date t .

³We will keep the standard deviation of the Δ shock three orders of magnitude lower than that of the technology shock. Because it is so small, this shock does not disturb the dynamics we discuss in a quantitatively important way.

We endow the households with a perceived law of motion given by

$$c_t = a_{10} + a_{11}c_{t-1} + a_{12}k_{t-1} + a_{13}s_{t-1}, \quad (18)$$

$$k_t = a_{20} + a_{21}c_{t-1} + a_{22}k_{t-1} + a_{23}s_{t-1}. \quad (19)$$

This perceived law of motion is a good one for the agents to use, because it corresponds in form to the equilibrium law of motion for the economy. Furthermore, it represents the minimal state variable (MSV) representation of the rational expectations solution.⁴ By repeatedly calculating the coefficients in this vector autoregression as new data become available, the agents may be able to correctly infer the equilibrium. The presence of constant terms in the model (15)-(17) and in the perceived law of motion (18)-(19) is effectively saying that the agents must learn the steady state values of variables instead of being given those values. This is important for our results, because it allows the trends we calculate to be smooth.

To obtain the mapping from the perceived law of motion to the actual law of motion, we use the perceived law of motion to obtain expected values and we substitute these into (15)-(17) in place of rational expectations. Consistent with much of the discussion in Evans and Honkapohja (2001), we consider the case where the information available to agents at time t is dated $t - 1$ and earlier. The expectations are then given by

$$E_t c_{t+1} = a_{10} + a_{11}E_t c_t + a_{12}E_t k_t + a_{13}E_t s_t \quad (20)$$

$$E_t k_{t+1} = a_{20} + a_{21}E_t c_t + a_{22}E_t k_t + a_{23}E_t s_t \quad (21)$$

$$E_t s_{t+1} = \rho E_t s_t \quad (22)$$

⁴We have written the perceived law of motion so that agents estimate both equations (18)-(19) with knowing the REE coefficients of those equations. Since the law of motion for capital (16) does not depend on expectations, we could assume (as in Packalén (2000)) that agents are perfectly informed of the coefficients of this law of motion; the only equation that matters for our learning analysis is the equation for consumption, where expectations play a role. Still, in keeping with the notion that agents are learning, it seems more natural to imagine that agents must learn the coefficients on the capital accumulation equation as well, and this is the route we choose to follow.

where

$$E_t c_t = a_{10} + a_{11}c_{t-1} + a_{12}k_{t-1} + a_{13}s_{t-1} \quad (23)$$

$$E_t k_t = a_{20} + a_{21}c_{t-1} + a_{22}k_{t-1} + a_{23}s_{t-1} \quad (24)$$

$$E_t s_t = \rho s_{t-1} \quad (25)$$

Substituting appropriately and collecting terms leads to the following actual law of motion for consumption:

$$c_t = T_{10} + T_{11}c_{t-1} + T_{12}k_{t-1} + T_{13}s_{t-1} + \Delta_t \quad (26)$$

where

$$\begin{aligned} T_{10} &= \mathcal{B}_{10} + \mathcal{B}_{11} [a_{10} + a_{11}a_{10} + a_{12}a_{20}] \\ &\quad + \mathcal{B}_{12} [a_{20} + a_{21}a_{10} + a_{22}a_{20}], \end{aligned} \quad (27)$$

$$T_{11} = \mathcal{B}_{11} [a_{11}^2 + a_{12}a_{21}] + \mathcal{B}_{12} [a_{21}a_{11} + a_{22}a_{21}], \quad (28)$$

$$T_{12} = \mathcal{B}_{11} [a_{11}a_{12} + a_{12}a_{22}] + \mathcal{B}_{12} [a_{21}a_{12} + a_{22}^2], \quad (29)$$

$$\begin{aligned} T_{13} &= \mathcal{B}_{11} [a_{11}a_{13} + a_{12}a_{23} + a_{13}\rho] \\ &\quad + \mathcal{B}_{12} [a_{21}a_{13} + a_{22}a_{23} + a_{23}\rho] + \mathcal{B}_{13} [\rho^2]. \end{aligned} \quad (30)$$

We write the system under learning as

$$\begin{bmatrix} c_t \\ k_t \\ s_t \end{bmatrix} = \begin{bmatrix} T_{10} \\ \mathcal{D}_{20} \\ 0 \end{bmatrix} + \begin{bmatrix} T_{11} & T_{12} & T_{13} \\ \mathcal{D}_{21} & \mathcal{D}_{22} & \mathcal{D}_{23} \\ 0 & 0 & \rho \end{bmatrix} \begin{bmatrix} c_{t-1} \\ k_{t-1} \\ s_{t-1} \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta_t \\ 0 \\ \vartheta_t \end{bmatrix}. \quad (31)$$

A stationary MSV rational expectation solution solves

$$T_{1i} = a_{1i}, \quad (32)$$

for $i = 0, 1, 2, 3$, with all eigenvalues of the matrix

$$\begin{bmatrix} T_{11} & T_{12} & T_{13} \\ \mathcal{D}_{21} & \mathcal{D}_{22} & \mathcal{D}_{23} \\ 0 & 0 & \rho \end{bmatrix} \quad (33)$$

inside the unit circle. For the calibrations we study, there is only one such solution.

3.2 Expectational stability

We can calculate expectational stability conditions for this system. Evans and Honkapohja (2001) provide general conditions under which expectational stability governs the stability of the system under a wide variety of real time, recursive learning assumptions. Expectational stability is determined by the following matrix differential equation

$$\frac{d}{d\tau} (a_{i,j}) = T(a_{i,j}) - (a_{i,j}), \quad (34)$$

for $i = 1, 2$; $j = 0, 1, 2, 3$. This differential equation describes a process in notional time by which beliefs, or forecasts, concerning the parameter vector a deviate from realizations, represented by the T-mapping, $T(a)$. The fixed points of equation (34) give us the MSV solution. A particular MSV solution $(\bar{a}_{i,j})$ is said to be E -stable if the MSV fixed point of the differential equation (34) is locally asymptotically stable at that point.

The nontrivial part of the T -map involves only the coefficients in the consumption equation. Let $T_1(a)$ describe this system as given by equations (27-30). The Jacobian matrix required for evaluating expectational stability is given by

$$DT_1(\bar{a}) - I = \begin{bmatrix} \mathcal{E}_{11} & \mathcal{B}_{11}\bar{a}_{10} & \mathcal{B}_{11}\bar{a}_{20} & 0 \\ 0 & \mathcal{E}_{22} & \mathcal{B}_{11}\bar{a}_{21} & 0 \\ 0 & \mathcal{B}_{11}\bar{a}_{12} & \mathcal{E}_{33} & 0 \\ 0 & \mathcal{B}_{11}\bar{a}_{13} & \mathcal{B}_{11}\bar{a}_{23} & \mathcal{E}_{44} \end{bmatrix}, \quad (35)$$

where

$$\mathcal{E}_{11} = \mathcal{B}_{11}(1 + \bar{a}_{11}) + \mathcal{B}_{12}\bar{a}_{21} - 1, \quad (36)$$

$$\mathcal{E}_{22} = 2\mathcal{B}_{11}\bar{a}_{11} + \mathcal{B}_{12}\bar{a}_{21} - 1, \quad (37)$$

$$\mathcal{E}_{33} = \mathcal{B}_{11}(\bar{a}_{11} + \bar{a}_{22}) + \mathcal{B}_{12}\bar{a}_{21} - 1, \quad (38)$$

$$\mathcal{E}_{44} = \mathcal{B}_{11}(\bar{a}_{11} + \rho) + \mathcal{B}_{12}\bar{a}_{21} - 1. \quad (39)$$

The conditions for E -stability of the MSV solution applicable to the model we consider are given in Proposition 10.3 of Evans and Honkapohja (2001).

According to this proposition, E -stability obtains if the real parts of the eigenvalues of $DT_1(\bar{a})$ are less than unity, or equivalently, if the eigenvalues of $DT_1(\bar{a}) - I$ have negative real parts. We verified that the eigenvalues of the above matrix are indeed always real and negative for the baseline model calibration we describe below. We note that this finding holds for all values of η and γ used in our analysis. Thus, for all parameter values we consider in this paper, the system under learning is always expectationally stable. This suggests stability in the real-time learning dynamics under weak conditions.⁵ We therefore proceed to real time learning.

3.3 Real time learning

When the agents are learning in real-time, the parameters $a_{i,j}$ in the recursive updating scheme are time-varying. This means that the T -mapping now becomes

$$T_{10}(\xi_{t-1}) = \mathcal{B}_{10} + \mathcal{B}_{11} [a_{10,t-1} + a_{11,t-1}a_{10,t-1} + a_{12,t-1}a_{20,t-1}] + \mathcal{B}_{12} [a_{20,t-1} + a_{21,t-1}a_{10,t-1} + a_{22,t-1}a_{20,t-1}], \quad (40)$$

$$T_{11}(\xi_{t-1}) = \mathcal{B}_{11} [a_{11,t-1}^2 + a_{12,t-1}a_{21,t-1}] + \mathcal{B}_{12} [a_{21,t-1}a_{11,t-1} + a_{22,t-1}a_{21,t-1}], \quad (41)$$

$$T_{12}(\xi_{t-1}) = \mathcal{B}_{11} [a_{11,t-1}a_{12,t-1} + a_{12,t-1}a_{22,t-1}] + \mathcal{B}_{12} [a_{21,t-1}a_{12,t-1} + a_{22,t-1}^2], \quad (42)$$

and

$$T_{13}(\xi_{t-1}) = \mathcal{B}_{11} [a_{11,t-1}a_{13,t-1} + a_{12,t-1}a_{23,t-1} + a_{13,t-1}\rho] + \mathcal{B}_{12} [a_{21,t-1}a_{13,t-1} + a_{22,t-1}a_{23,t-1} + a_{23,t-1}\rho] + \mathcal{B}_{13} [\rho^2]. \quad (43)$$

⁵The interested reader is referred to Evans and Honkapohja (2001) for the details of this connection.

The actual law of motion is therefore

$$\begin{bmatrix} c_t \\ k_t \\ s_t \end{bmatrix} = \begin{bmatrix} T_{10}(\xi_{t-1}) \\ \mathcal{D}_{20} \\ 0 \end{bmatrix} + \begin{bmatrix} T_{11}(\xi_{t-1}) & T_{12}(\xi_{t-1}) & T_{13}(\xi_{t-1}) \\ \mathcal{D}_{21} & \mathcal{D}_{22} & \mathcal{D}_{23} \\ 0 & 0 & \rho \end{bmatrix} \begin{bmatrix} c_{t-1} \\ k_{t-1} \\ s_{t-1} \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta_t \\ 0 \\ \vartheta_t \end{bmatrix}. \quad (44)$$

The coefficients ξ_t are updated according to a recursive least squares estimation

$$\xi_t = \xi_{t-1} + t^{-1} R_t^{-1} z_{t-1}' z_{t-1} [T(\xi_{t-1}) - \xi_{t-1}], \quad (45)$$

$$R_t = R_{t-1} + t^{-1} [z_{t-1}' z_{t-1} - R_{t-1}], \quad (46)$$

where

$$z_{t-1} = \begin{bmatrix} 1 \\ c_{t-1} \\ k_{t-1} \\ s_{t-1} \end{bmatrix} \quad (47)$$

and

$$T(\xi_{t-1}) = \begin{bmatrix} T_{10}(\xi_{t-1}) \\ T_{11}(\xi_{t-1}) \\ T_{12}(\xi_{t-1}) \\ T_{13}(\xi_{t-1}) \end{bmatrix}. \quad (48)$$

When we study constant gain learning, we replace t^{-1} with a small positive constant g in equations (45) and (46).

In order to simulate this system, we begin with initial, $t - 1$, values of capital and consumption. We then obtain k_t from the second equation of (44). Using the third equation of (44), we draw ϑ_t and obtain s_t . Next, we draw a value Δ_t . Then we use equation (46) to obtain time t values for $r_{i,j}$, and equation (45) to obtain time t values for ξ_t . Finally, we use the first equation of (44) to obtain the time t value for c_t . This process is then repeated to generate time series on c_t , k_t , and other variables of interest.

As we have shown, this system is expectationally stable in notional time, which implies that it is stable under a real-time recursive least squares

scheme in which the agents employ (45) and (46). Rather than studying least-squares learning, we follow Sargent (1999) in considering a more general, constant-gain learning system in which the t^{-1} gain in equation (45) is replaced by a small positive constant value, g . A small constant gain, as opposed to the $1/t$ gain of recursive least squares implies that past data is discounted and that the system never settles down perfectly to a rational expectations equilibrium. Instead, it will achieve an approximate equilibrium centered around the rational expectations equilibrium path.⁶ Thus, under a constant gain updating scheme, we can no longer be assured that the stability properties of the system will hold. However, if the gain is sufficiently small and the system is in a sufficiently small neighborhood of the rational expectations equilibrium, then we may expect the system to remain in that neighborhood. Moreover, by contrast with the least squares $1/t$ gain, the small constant gain allows the system to respond immediately in the event that an underlying parameter of the model changes unexpectedly. This ever-vigilant property of the constant-gain learning system is essential to avoiding long periods of systematic forecast errors that might lead agents to conclude that their perceived law of motion was misspecified. Indeed, the constant gain assumption implies that agents recognize that their model is potentially prone to structural changes in the trend growth rate and may therefore become misspecified. The constant gain allows agents to quickly react should the balanced growth path change from the one they were previously tracking. Based on these considerations, the constant gain assumption seems reasonable given the environment we consider.

In principle, we could now ask how this system would react to any (small enough) change in any parameter of the model, not just changes in the growth rates γ and η . Suppose, for instance, that people became more patient, or that the share of capital in national income increased. Such

⁶These differences turn out to be quite small empirically, and so we do not discuss them further.

changes would alter the balanced growth path of the economy (through level effects, for these parameter changes). But the agents in the model would be able to learn the new rational expectations equilibrium implied after changes in those parameters had taken place.⁷

We now turn to comparing the model with U.S. postwar data.

4 Application to postwar U.S. data

4.1 Overview

We now illustrate how our model can be used to understand post war U.S. data. Since the model is quite simple and does not have some of the important categories of national income that exist in the data, this exercise cannot be completely satisfactory. However, since the model is also a variant of a widely-known benchmark, we can begin to assess how important structural change is for determining the nature of the business cycle in the data as well as for the performance of the model relative to the U.S. data.

4.2 Calibration

We employ a standard calibration for this model under the assumption that each period represents one quarter. For this purpose, we turn to Cooley and Prescott (1994). They suggest the following calibration. In preferences, the discount factor, $\beta = .987$, and the weight on leisure, $\theta = 1.78$. For technology, capital's share $\alpha = .4$ and the depreciation rate, $\delta = .012$. The serial correlation of the business cycle shock $\rho = .95$, and the shocks have a standard deviation of .007. Cooley and Prescott (1994) also calibrate growth

⁷One could think of rational expectations versions of our system. Completely unanticipated shocks are inconsistent with the rational expectations assumption, but one could develop a model with regime-switching, say, in productivity growth, and then proceed to analyze the dynamics of that model following switches. Such an approach has been pursued by Kahn and Rich (2004) and Andolfatto and Gomme (2003). That approach puts more structure on the nature of the trend-changing shocks than we have here, and requires agents to understand the number of dimensions on which alternative regimes might occur.

rates of labor and technological change, but since we allow changes in these growth rates, the calibration of these features is undertaken separately.

In the learning algorithm we have outlined, the gain sequence would normally be set to $1/t$ to correspond to recursive least squares. However, for the reasons noted above, we have chosen to set the gain to a small positive constant, $g = .00025$. Based on our experience with simulations, this is close to the largest value of the gain that still remains consistent with stability under recursive learning. Quantitatively, the choice of the gain does not seem to have a large impact on our results, so long as it produces a stable system.

Because the model economy does not have all of the major categories of national income that the U.S. national accounts have, a direct comparison between the model and the data is not a simple matter.⁸ All the data we use are quarterly from 1948:Q1 to 2002:Q1. The data are in real terms, 1996 dollars, seasonally adjusted, and chain-weighted. Our model has predictions for aggregates, and so we focus on them. We are quite concerned that the aggregates in the model add up, so that the trends in the labor input and productivity can be viewed as driving the trends in the other variables of interest. We have no government sector in the model, and so we subtract real government purchases from real GDP in the data we use. We also subtract real farm business product from real GDP. This gives us a measure of nonagricultural private sector output. We have a consistent private sector nonagricultural total hours series, from the Bureau of Labor Statistics Establishment Survey, for this measure of output. We use this hours series to represent our labor input. Productivity is then quarterly output divided by quarterly aggregate hours. Our model has no international sector, but net exports comprises a nontrivial component of GDP in the data. We add

⁸Consistency between the model and the actual data that the model data are compared with does not seem to be the rule. For example, King and Rebelo (1999) make no effort to remove government from their measure of output even though the model they consider does not have a government sector.

the services portion of net exports to our measure of consumption, and the goods portion of net exports to our measure of investment. In the data where sub-categories of exports and imports are available, capital goods, industrial supplies, and automobiles make up a substantial fraction of goods exports, and so we call this investment for the purposes of our study. Our measure of investment is then gross private domestic investment plus net exports of goods, plus personal consumption expenditures on consumer durables. Our measure of consumption is personal consumption expenditures on services and nondurable goods, plus net exports of services, less farm business product, which is presumably mainly consumption-oriented.

Because of chain weighting, consumption plus investment still may not add up to output. We checked this and found that any discrepancy was negligible after 1980. Before that, the discrepancy can be larger, as much as two percent of output. We therefore allocated any discrepancy to consumption and investment using the consumption-to-output ratio for that year. Thus we end up with time series in which output is indeed equal to consumption plus investment.

4.3 Breaks in the balanced growth path

It is well-known that there was a slowdown in measured productivity growth in the U.S. economy beginning sometime in the late 1960s or early 1970s. The state of the econometric evidence on this question is reviewed in Hansen (2001). A key paper in the literature is Perron (1989), who argued that for postwar quarterly real U.S. GDP, a time series model with a change in the slope coefficients of a time trend allows one to reject the random walk hypothesis in favor of trend stationarity around the broken trendline. Perron associated the 1973 slowdown in growth with the oil price shock, but this date is also associated with a slowdown in labor productivity.⁹

⁹Later authors, such as Zivot and Andrews (1992), extended the analysis to the case where the break date was viewed as unknown.

Another recent attempt to date a structural break during this period is Bai, Lumsdaine, and Stock (1998). Their analysis is multivariate and suggests a trend break sometime between 1966:Q2 and 1971:Q4, with a most likely date of 1969:Q1.

We have designed our model to allow the economy to adapt to changes of this type. We can alter the growth rate of productivity in the model at a given point in time, and, provided the change is not too large, we can expect the economy to adjust to the new balanced growth path.

How can we go about choosing break dates for our economy? We use the following approach. Our model says that the nature of the balanced growth path—the trend—is dictated by increases in productivity units $X(t)$ and increases in the labor input $N(t)$. For ease of reference, let us call these the “actual” productivity and labor input series. When the growth rates of these variables, γ and η , change, the economy must adjust to a new balanced growth path. The model also produces *measured* productivity and a *measured* labor input series. If there were never a trend break, these measured series would have the same trend as the actual series. However, since it takes some time for the economy to adjust to the new balanced growth path, in general there will be differences in the trends of the actual and the measured productivity and labor input series. In the data, we have measured increases in productivity and measured increases in the labor input. Thus it seems quite clear that we need the trends in measured productivity and measured labor input from the model to be comparable to the measured productivity and measured labor input trends we have from the data in order to have a satisfactory calibration.

One approach to calibrating the model would be to only allow trend breaks where clear econometric evidence is available. This would probably lead one to posit a single trend break in productivity sometime before 1973 (such as the one suggested by Bai, Lumsdaine, and Stock (1998)) and then require the balanced growth path to be log-linear at all other times. We

TABLE 1. SEARCH RANGES

Growth Factor	Break Date Search Range	Pre-Break Growth Rate	Post-Break Growth Rate
Labor input	1955,Q1 to 1964,Q4	-0.72 to 3.32 %	-0.12 to 3.94 %
Productivity 1	1965,Q1 to 1974,Q4	0.05 to 4.11 %	-0.75 to 3.29 %
Productivity 2	1991,Q1 to 1997,Q1	-0.75 to 3.29 %	-0.45 to 3.60 %

Table 1: We chose these search ranges for possible break dates in trend labor input and trend productivity, as well as for the possible growth rates between the trend breaks. Growth rates are in annual terms.

think this may not be the most interesting way to proceed. There could easily be smaller changes in growth rates, economically significant from the standpoint of judging business cycles, but not substantial enough to cause a rejection of a null hypothesis of log-linear growth. One example of this is the greater entry of women into the labor force beginning in the 1960s, which is often cited as one of the major changes in the U.S. economy during the postwar era. For the hours series we employ,¹⁰ a univariate test based on Andrews (1993) cannot reject the null hypothesis of no change in the growth rate of hours across the entire postwar era. A look at the data clarifies the source of this result: The hours series before the 1960s is short and relatively volatile, and any change in the growth rate, if it occurred, is relatively small. Another example of this possibility is the idea of a “new economy” in the 1990s, which is not easy to defend with statistical tests.

Instead of relying on econometric evidence alone, we used a simulated method of moments search procedure, described in more detail in Appendix B, to choose break dates for the growth factors $X(t)$ and $N(t)$, as well as for growth rates of these factors γ and η , based on the principle that the trend in measured productivity from the model should match the trend in measured productivity from the data. We began by specifying some ranges over which we wish to search for trend breaks, as well as ranges for possible

¹⁰Total nonagricultural private sector hours from the establishment survey, quarterly, 1948:Q1 to 2002:Q1. We thank Jeremy Piger for conducting this test on the hours series.

growth rates between the break dates. These ranges are described in Table 1, and reflect our “priors” on when we think reasonable dates for breaks in log-linear trends might have occurred. When trend breaks occur in our model, the agents must learn about them, and so we might expect $X(t)$ and $N(t)$ to begin growing at a different rate at a date somewhat before a trend break becomes apparent in the measured series. For this reason, we included years before apparent trend breaks in the data (such as 1965 for the productivity slowdown) as possible trend break dates in our model. We allowed two breaks for productivity, corresponding to a productivity slowdown circa 1970 and a new economy circa 1995. We allowed one break for the labor input, corresponding to changing attitudes toward women in the workforce circa 1960. The growth rate ranges are calculated as the mean quarterly growth rates for hours and productivity in the data for the appropriate time period, which we allow to possibly be higher or lower by one-half of one percent per quarter.¹¹

We have two factors driving trend growth in the model, along with three break dates and therefore five distinct periods of different growth rates (three for productivity and two for the labor input). This means there is a vector of eight objects we must choose. We begin with a set of candidate solutions. For each candidate solution, we let our model generate a trend. This involved simulating our learning model but “turning off” the standard business cycle shock, s_t . In practice, this meant reducing the standard deviation of the business cycle shock s_t by a factor of 1000 (from the Cooley and Prescott (1994) calibration of .007 to .000007), so that effectively this shock process is not important in the output generated by our model.¹² During this simulation, we leave in the trend changes indicated by the candidate vector

¹¹The table has these in annualized terms for ease of interpretation. The appropriate time period is calculated as if the trend break were dated at the midpoint of the ranges in Table 1.

¹²We require a small amount of noise in the system so that our VAR systems can still be estimated.

TABLE 2. OPTIMAL CHOICES

Growth Factor	Break Date Search Range	Pre-Break Growth Rate	Post-Break Growth Rate
Labor input	1961,Q2	1.20 %	1.91 %
Productivity 1	1973,Q3	2.47 %	1.21 %
Productivity 2	1993,Q3	1.21 %	1.86 %

Table 2: Optimal choices of trend break dates and growth rates for the two growth factors in the model. Growth rates are in annual terms.

of break dates. We then evaluate each candidate solution according to a fitness criterion. The fitness measure is the sum of mean squared deviations of measured productivity in the data from the implied trend, plus the sum of mean squared deviations of measured hours in the data from the implied trend. Hours and productivity thus receive equal weight in this calculation. We then update the set of candidate solutions in the direction of those that tended to generate better fitness scores using standard genetic operators, as discussed further in Appendix B. The process continued until no further fitness improvements could be found.

Table 2 reports our findings. For productivity growth, the break dates are consistent with those that appear often in the literature. Productivity (that is, $X(t)$) grows at a net annual rate of 2.47 percent until 1973:Q3, then slows to an annual growth rate of 1.21 percent until 1993:Q3, before accelerating to an annual rate of 1.86 percent through the end of the sample. For the labor input (that is, $N(t)$), trend breaks are much less pronounced. The labor input series grows at an annual rate of 1.20 percent initially, before accelerating to 1.91 percent in 1961:Q2.

Our first task is to show that the breaks in growth rates we have determined imply reasonable trends for the measured labor input and for measured productivity. Figures 1 and 2 combine the trends calculated using our *model* with the *actual data* on hours and productivity for the U.S. economy. The trends are generally very smooth and are what many economists would

have in mind when they say there is a “trend in the data.”

We stress that our procedure has been to use our theoretical framework to fit trends for measured productivity and measured labor input only. But the trends in these growth factors in turn *imply* trends for output, investment, and consumption. We have allowed the latter trends to be freely determined by the model, i.e. we have not sought to fit trends for output, investment and consumption to the data as we did for productivity and labor hours. In addition, the business cycle shock occurs in conjunction with the rare changes in trend we have modelled. We now turn assessing the performance of the model.

4.4 The balanced growth puzzle

While we have fit trends for productivity and hours, we are letting the trends in growth factors dictate the remaining trends in the model. Figures 3, 4, and 5 show how the trends we have calculated using the model compare to the level of output, consumption, and investment, respectively, in the U.S. data.¹³ For output, the combination of hours growth and productivity growth with some trend breaks provides a reasonable account of growth, so reasonable in fact that one might think that the trend line was simply drawn through the data by a student of business cycles. It is well known that without the trend breaks, a purely log-linear trend does not provide as reasonable of an account of this data.¹⁴

Figure 3 gives us confidence that a two-factor exogenous growth model is a good one for disentangling trend from cycle in the data.

¹³For all of the trends we report, comparison to the data requires a units normalization. We accomplish this normalization by assuming that the model is following a balanced growth path during the initial portion of the sample, before the first trend break occurs.

¹⁴See, for instance, Figure 3 in King et al. (1988a, pp. 227–228). In that figure, consumption, investment, output and other variables are shown to persistently lie either above or below the deterministic trend for many years at a time. King et al. conclude from their figure that the U.S. time series data indicate the “possibility of a low frequency component not captured by the deterministic trend” (p. 231).

The division of output between private sector consumption and investment is also dictated by the model. For these variables, the trend lines tend to run through the data in the earlier and middle portions of the sample. In the latter portion of the sample, actual consumption tends to run below trend, while investment tends to run noticeably above trend.¹⁵ It was widely reported that there was an “investment boom” in the 1990s, and the data we have seem to bear this out. Since consumption is the only other component of output here, it must run below trend to accommodate the boom.¹⁶

The consumption and investment trends are what we label the *balanced growth puzzle*. Our model is a standard one, and we expect that it can provide a reasonable account of growth during the postwar era. The model does accomplish this over much of the sample. But during the latter portion of the sample, investment booms and consumption lags, relative to what the model suggests the trends should have been. This suggests that a two-factor exogenous growth model, even with trend breaks included, is too simple to account for consumption and investment trends. (For the output trend, it seems to work well.) There are many possibilities that could be explored to explain the puzzle. There were, for instance, important tax changes during the 1980s, while our model abstracts completely from taxes. We have a one-sector model, but perhaps a multisector model is required. The increasing prevalence of new types of capital during the 1980s and 1990s suggests depreciation rates may be increasing during this period. These are just some possibilities, and we think all of these as well as others may provide a portion of the explanation.

¹⁵We considered a few alternative data arrangements to see if this feature of the analysis was robust to changes in the interpretation of “consumption” and “investment”. For instance, we considered including consolidated government spending data, allocating using available figures on government consumption versus government investment. We also considered including consumer durable purchases as consumption instead of investment. These types of changes did not alter the qualitative results.

¹⁶See Cogley (2003) for one approach to using consumption as the basis for determining trend growth changes. Cogley comes to the conclusion that trend growth has been only modestly faster in the 1990s than during the productivity slowdown era.

The balanced growth puzzle notwithstanding, we think that these trends are reasonable judgements of what the “actual” trends look like in the data. However, our point is not so much to say that the fit is good, but that we lay bare our assumptions about the growth process that allow us to detrend the data in this manner. Other authors are welcome to provide alternative assumptions on models like this one, or provide alternative growth models, in order to detrend the data in a different manner. Our hope is that constructive work can be done along these lines.

We now take the calculated trends as the prediction of our model, so that the deviations from trend are the business cycle components in the data. We turn to evaluating the properties of these business cycle components.

4.5 Business cycle statistics

The reaction of the economy to changes in the balanced growth path will depend in part on what business cycle shocks occur in tandem with the growth rate changes. In part because of this, we average over a large number of economies in order to calculate business cycle statistics for artificial economies. To generate the artificial data, we simulated the calibrated economy for a large number of periods to verify that the estimated coefficients in the agents’ regressions were close to the rational expectations values. We then collected an additional 217 observations, corresponding to the 217 quarters of actual U.S. data we have. During this latter part of the exercise, we allowed the trend breaks as discussed at quarters corresponding to the dates from Table 2, so that the agents in the economy had to also react to the trend breaks as they were coping with the business cycle shock. The trend that is taken from the artificial data is exactly the same one that is

TABLE 3. BUSINESS CYCLE STATISTICS

	Volatility		Relative Volatility		Contemporaneous Correlations	
	Data	Model	Data	Model	Data	Model
Output	3.25	3.50	1.00	1.00	1.00	1.00
Consumption	3.40	2.16	1.05	0.62	0.60	0.75
Investment	14.80	8.86	4.57	2.53	0.65	0.92
Hours	2.62	1.54	0.81	0.44	0.65	0.80
Productivity	2.52	2.44	0.77	0.70	0.61	0.92

Table 3: Business cycle statistics, model-consistent detrending.

taken from the U.S. data.¹⁷

In assessing the behavior of equilibrium business cycle models like this one, authors have typically compared volatility and contemporaneous correlation measures from the model to those suggested by the data. We do the same, using our model-consistent trends to calculate percentage deviations of all variables from their trend values. We average our statistics across 500 economies each run for 217 periods with identical trend breaks.

We begin with overall volatility, which is measured by the standard deviation of the actual and artificial data series, and displayed in first column of Table 3. These standard deviations are often more than twice the size of those reported by others, for example, King and Rebelo (1999). The reason for this is simple. The trends we use are essentially piecewise log-linear, and so do not attribute a portion of every data movement to the trend component, as many statistical filters do. Thus the portion of the variability in the data that is attributed to business cycle volatility is likely to be larger under our methodology. In this sense, the business cycle shock has to explain more under our approach than under traditional approaches to the

¹⁷An interesting question is whether an econometrician considering the productivity data generated by one of these economies would detect the breaks in trend growth rates that are built into the model. Another interesting question is whether the data generated by the model would be consistent with a random walk hypothesis in the eyes of an econometrician. We hope to investigate these issues in future work.

detrending question.

A key question for this line of research has been: How much of the variability in the data can be explained by a model of this type? That is, how much variance can we generate by simply assuming a single shock to the production technology along with occasional breaks in trend growth rates? One of our more interesting findings is that for the model, the average standard deviation for output is 3.50 according to Table 3, while for the data it is 3.25. That suggests that more than 100 percent of the variance of output about the balanced growth path can be explained with a model of this type! That is a high number even compared to other exercises along this line. It suggests that shocks to the technology coupled with the important movements in trend we have observed during the postwar era provide a promising lead on accounting for all of the variability of output around the balanced growth path during the postwar era. If anything, the model generates too much volatility.¹⁸

Since the trends are piecewise log-linear in our model, they tend to be less accommodating to the data than those computed using most statistical filters. We stress that the higher volatility implied by our method applies equally to both the model and the data. This is why the model can still explain a large fraction of the variance in the data, even when that variance has increased substantially relative to commonly reported statistics.

The volatilities in the data and for the model relative to output volatility are given in the relative volatilities column of Table 3. There are several interesting aspects of the results reported in this section. First, consumption is about as volatile as output in the data, but only two-thirds as variable as

¹⁸Recent research has argued that the technology portion of the Solow residual may be less volatile than we have calibrated it, by perhaps a factor of five. If we reduce the standard deviation of the shock to technology, the business cycle volatility of this model will fall proportionately. Again, ours is only an example, which we mainly want to keep comparable to previous research. Interested readers can consult King and Rebelo (1999) for an alternative equilibrium business cycle model that generates similar data with less volatile shocks. That model is still in the balanced growth framework and so our methods would still apply.

output in the model. The source of this finding is quite clear from Figure 4, where the U.S. consumption data tends to drift below trend later in the sample. This tends to increase the volatility of the consumption data if it is measured as deviation from trend. The relative volatility of investment is only about half as large in the model as it is in the data; however, in both the model and the data investment is much more volatile than output. Again, the investment boom of the 1990s seems to have contributed quite a lot to the variance of investment in the data.

Hours worked in the data is about 80 percent as volatile as output, somewhat lower than the one-to-one ratio that is often reported in the literature. But the relative volatility of hours in the model is still only about half what it is in the data, that is, .44 in the model versus .81 in the data. Thus one of the key findings of the original equilibrium business cycle literature, that the labor market portion of the model is not satisfactory, holds up in this example.

The contemporaneous correlations with output for both the model and the data are given in the last columns of Table 3. All variables are procyclical, both in the model and in the data. These statistics tend to be lower than their counterparts reported in the literature, for instance in King and Rebelo (1999), for both the model and the data. The model predicts too much procyclicality across all of the variables, but still, the statistics reported are noticeably lower than those typically reported. One statistic is not lower than typically reported, and that is the correlation of productivity with output in the data, which is .61. Productivity is more strongly procyclical than suggested by Cooley and Prescott (1994) or King and Rebelo (1999). Thus hours and productivity more or less move together both in the model and in the data. Using alternative techniques for detrending, this has not always been true, and in fact was judged to be a problem with the model.

5 Conclusion

The concept of a balanced growth path has had an enormous influence on macroeconomists. In this paper we have taken this concept, which underlies nearly all macroeconomic models in use today, to the data. Of course, growth rates of important macroeconomic time series are well-known to be inconsistent with purely log-linear growth through the postwar period. For this reason, we have allowed permanent trend breaks where appropriate, and we have used learning via the methodology of Evans and Honkapohja (2001) as a “glue” that holds the resulting various balanced growth paths together. In particular, learning enables us to deal with the transition from one balanced growth path to another in a smooth manner. The result is a piecewise, log-linear trend, like the ones discussed in the empirical literature on structural change. We remove this same trend from the data as our method of detrending the data. In this sense we have a model-consistent method of detrending.

We have also included an application to the postwar U.S. data. Structural change itself is a new type of shock in this model, and we find that it contributes substantially to the variance of output. We have also identified a *balanced growth puzzle*, in that we cannot completely account for observed trends in U.S. aggregate consumption and investment beginning in the mid-1980s using the simple, two-factor exogenous growth model augmented with structural change. This puzzle stems from our consideration of the multivariate nature of the trend as implied by the model; indeed, the trend for output generated by our model fits the actual data series rather well. Finally, we have shown how to calculate business cycle statistics using model consistent detrending methods. Approximating these “true” statistics via atheoretic, statistical filtering of artificial and actual data may lead the researcher to misjudge the model’s successes and failures.

References

- [1] ANDREWS, DONALD W. K. 1993. "Tests for Parameter Instability and Structural Change with Unknown Change Point." *Econometrica*, 61(4): 821-56.
- [2] ANDOLFATTO, D., AND P. GOMME. 2003. Monetary policy regimes and beliefs. *International Economic Review*, 44(1): 1-30.
- [3] BAI, J., R. LUMSDAINE, AND J. STOCK. 1998. "Testing for and Dating Common Breaks in Multivariate Time Series." *Review of Economic Studies*, 64(3): 395-432.
- [4] BAXTER, M., AND R. KING. 1999. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *Review of Economic and Statistics*, 81(4): 575-93.
- [5] BEVERIDGE, S. AND C. NELSON. 1981. "A New Approach to Decomposition of Economic Time Series Into Permanent and Transitory Components with Particular Attention to Measurement of the Business Cycle." *Journal of Monetary Economics*, 7: 151-174.
- [6] BULLARD, J. AND S. EUSEPI. 2003. "Did the Great Inflation Occur Despite Policymaker Commitment to a Taylor Rule?" Working paper, Federal Reserve Bank of St. Louis.
- [7] BURNSIDE, C. 1998. "Detrending and Business Cycle Facts: A Comment." *Journal of Monetary Economics*, 41(3): 513-532.
- [8] CANOVA, F. 1998a. "Detrending and Business Cycle Facts." *Journal of Monetary Economics*, 41(3): 475-512.
- [9] CANOVA, F. 1998b. "Detrending and Business Cycle Facts: A User's Guide." *Journal of Monetary Economics*. 41(3): 533-540.
- [10] CHRISTIANO, L., AND T. FITZGERALD. 1999. "The Band-Pass Filter." NBER Working Paper #7257, July.

- [11] COGLEY, T. 2003. "How Fast Can the New Economy Grow? A Bayesian Analysis of the Evolution of Trend Growth." Working paper, UC Davis.
- [12] COGLEY, T., AND J.M. NASON. 1995a. "Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research." *Journal of Economic Dynamics and Control*, 19(1-2): 253-278.
- [13] COGLEY, T., AND J.M. NASON. 1995b. "Output Dynamics in Real-Business-Cycle Models." *American Economic Review*, 85(3): 492-511.
- [14] COLLARD, F. AND H. DELLAS. 2004. "The Great Inflation of the 1970s." Working paper, University of Bern.
- [15] COOLEY, T., AND E. PRESCOTT. 1994. "Economic Growth and Business Cycles," in T. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton, NJ: Princeton University Press.
- [16] EVANS, G.W., AND S. HONKAPOHJA. 2001. *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.
- [17] HANSEN, B. 2001. "The New Econometrics of Structural Change: Dating Breaks in U.S. Labor Productivity." *Journal of Economic Perspectives*, 15(4): 117-128.
- [18] HARVEY, A.C. 1997. "Trends, Cycles and Autoregressions." *Economic Journal*, 107: 192-201.
- [19] HODRICK, R., AND E. PRESCOTT. 1980. "Post-War U.S. Business Cycles: An Empirical Investigation." Working paper, Carnegie-Mellon University.
- [20] KAHN, J.A., AND R.W. RICH. 2004. "Tracking the New Economy: Using Growth Theory to Detect Changes in Trend Productivity." Working Paper, Federal Reserve Bank of New York.

- [21] KING, R., C. PLOSSER, AND S. REBELO. 1988a. "Production, Growth, and Business Cycles, I: The Basic Neoclassical Model." *Journal of Monetary Economics*, 21: 195-232.
- [22] KING, R., C. PLOSSER, AND S. REBELO. 1988b. "Production, Growth, and Business Cycles II: New Directions." *Journal of Monetary Economics*, 21: 309-341.
- [23] KING, R., AND S. REBELO. 1999. "Resuscitating Real Business Cycles," in J. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: Elsevier.
- [24] KYDLAND, F., AND E. PRESCOTT. 1991. "The Econometrics of the General Equilibrium Approach to Business Cycles." *Scandinavian Journal of Economics*, 93: 161-178.
- [25] LANSING, K. 2000. "Learning About a Shift in Trend Output: Implications for Monetary Policy and Inflation." Working Paper 2000-16, Federal Reserve Bank of San Francisco.
- [26] LANSING, K. 2002. "Real-Time Estimation of Trend Output and the Illusion of Interest Rate Smoothing." *Federal Reserve Bank of San Francisco 2002 Review*, 17-34.
- [27] PACKALÉN, M. 2000. "On the Learnability of Rational Expectations Equilibria in Three Business Cycle Models." Research Report No. 87:2000, University of Helsinki.
- [28] PAKKO, M. 2002. "What Happens When the Technology Growth Trend Changes? Transition Dynamics, Capital Growth, and the 'New Economy'." *Review of Economic Dynamics*, 5(2): 376-407.
- [29] PERRON, P. 1989. "The Great Crash, the Oil Price Shock, and the Unit-Root Hypothesis." *Econometrica*, 57(6): 1361-1401.
- [30] PRESCOTT, E. 1986. "Theory Ahead of Business Cycle Measurement." *Quarterly Review*, Federal Reserve Bank of Minneapolis, 10: 9-22.

- [31] ROTEMBERG, J.J. 2003. “Stochastic Technical Progress, Smooth Trends and Nearly Distinct Business Cycles.” *American Economic Review*, 93: 1543-1559.
- [32] ROTEMBERG, J.J., AND M. WOODFORD. 1996. “Real Business Cycle Models and the Forecastable Movements in Output, Hours, and Consumption.” *American Economic Review*, 86: 71-89.
- [33] SARGENT, T.J. 1999. *The Conquest of American Inflation* Princeton, N.J.: Princeton University Press.
- [34] STOCK, J.H., AND M.W. WATSON. 1990. “Business Cycle Properties of Selected U.S. Economic Time Series, 1959-1988.” NBER Working Paper # 3376, June.
- [35] STOCK, J.H., AND M.W. WATSON. 1999. “Business Cycle Fluctuations in U.S. Macroeconomic Time Series.” In J. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: Elsevier.

A Linear representation of the model

We wish to analyze the system (8)-(11) in which the nonstationary variables, namely capital, consumption, and output, are rendered stationary via

$$\hat{k}_t = \frac{K_t}{X_t N_t}, \quad \hat{y}_t = \frac{Y_t}{X_t N_t}, \quad \hat{c}_t = \frac{C_t}{X_t N_t}. \quad (49)$$

If there was no growth in productivity over time, these variables would simply be in per capita terms; with productivity growth they are measured in per total efficiency unit terms. By dividing equations (8) through (11) by $X_t N_t$ appropriately, we can write them in terms of stationary variables as

$$\gamma \eta \hat{k}_{t+1} = \hat{y}_t + (1 - \delta) \hat{k}_t - \hat{c}_t, \quad (50)$$

$$\hat{y}_t = \hat{s}_t \left(\hat{k}_t \right)^\alpha \left(\hat{\ell}_t \right)^{1-\alpha}, \quad (51)$$

$$\hat{c}_t = \frac{(1 - \alpha)}{\theta} \hat{y}_t \left(\frac{1 - \hat{\ell}_t}{\hat{\ell}_t} \right), \quad (52)$$

and

$$\frac{1}{\hat{c}_t} = \frac{\beta}{\gamma} E_t \left\{ \frac{1}{\hat{c}_{t+1}} \left[\frac{\alpha \hat{y}_{t+1}}{\hat{k}_{t+1}} + 1 - \delta \right] \right\}. \quad (53)$$

A nonstochastic steady state of this transformed system corresponds to a balanced growth path of the original system. The gross rate of growth along the balanced growth path is $\gamma\eta$. We denote the nonstochastic steady state values by $\hat{c}_t = \bar{c}$, $\hat{y}_t = \bar{y}$, $\hat{k}_t = \bar{k}$, $\hat{\ell}_t = \bar{\ell}$, and $\hat{s}_t = \bar{s} = 1$, $\forall t$. These equations can be solved explicitly. Define φ by

$$\varphi = (1 + \theta) (\alpha - 1) \beta (\delta - 1) + \gamma [(\alpha - 1) + \theta (\alpha \beta \eta - 1)]. \quad (54)$$

Then

$$\bar{y} = \varphi^{-1} (\alpha - 1) \alpha \beta \left(\frac{\beta (\delta - 1) + \gamma}{\alpha \beta} \right)^{2 + \frac{1}{\alpha - 1}}, \quad (55)$$

$$\bar{k} = \varphi^{-1} (\alpha - 1) \alpha \beta \left(\frac{\beta (\delta - 1) + \gamma}{\alpha \beta} \right)^{\frac{\alpha}{\alpha - 1}}, \quad (56)$$

$$\bar{\ell} = \varphi^{-1} (\alpha - 1) [\beta (\delta - 1) + \gamma], \quad (57)$$

and

$$\bar{c} = \varphi^{-1} (1 - \alpha) \left(\frac{\beta (\delta - 1) + \gamma}{\alpha \beta} \right)^{\frac{\alpha}{\alpha - 1}} [(\alpha - 1) \beta (\delta - 1) - \gamma + \alpha \beta \gamma \eta]. \quad (58)$$

We can deduce that the capital to output ratio along a balanced growth path will be equal to

$$\frac{\bar{k}}{\bar{y}} = \frac{\alpha \beta}{\gamma - \beta (1 - \delta)}, \quad (59)$$

that the consumption to output ratio will be

$$\frac{\bar{c}}{\bar{y}} = \frac{\gamma - \beta (1 - \delta) - \alpha \beta (\gamma \eta - 1 + \delta)}{\gamma - \beta (1 - \delta)}, \quad (60)$$

and that the capital-labor ratio will be

$$\frac{\bar{k}}{\bar{\ell}} = \left(\frac{\gamma - \beta (1 - \delta)}{\alpha \beta} \right)^{\frac{1}{\alpha - 1}}. \quad (61)$$

Since the growth rates γ and η enter these expressions, growth matters for the calibration of models in this class.¹⁹ Many models that have been studied abstract from growth but calibrate to growth facts such as a constant capital to output ratio.

In order to apply the Evans and Honkapohja (2001) methodology to this problem, we need a linear system. Accordingly, we now proceed with a well-known linearization of this model, expressed in terms of logarithmic deviations from steady state. For this purpose we define

$$\tilde{c}_t = \ln\left(\frac{\hat{c}_t}{\bar{c}}\right), \quad \tilde{k}_t = \ln\left(\frac{\hat{k}_t}{\bar{k}}\right), \quad \tilde{\ell}_t = \ln\left(\frac{\hat{\ell}_t}{\bar{\ell}}\right), \quad (62)$$

$$\tilde{y}_t = \ln\left(\frac{\hat{y}_t}{\bar{y}}\right), \text{ and } \tilde{s}_t = \ln\left(\frac{\hat{s}_t}{\bar{s}}\right). \quad (63)$$

By noting that for any of these variables, $\hat{x}_t = e^{\tilde{x}_t} \bar{x}$, using the approximation $e^x \approx 1 + x$, and using the fact that $\bar{y} = (\gamma\eta - 1 + \delta)\bar{k} + \bar{c}$, we can write equation (50) as

$$\gamma\eta\tilde{k}_{t+1} = \frac{\bar{y}}{\bar{k}}\tilde{y}_t + (1 - \delta)\tilde{k}_t - \frac{\bar{c}}{\bar{k}}\tilde{c}_t. \quad (64)$$

For equation (51), we can write

$$\tilde{y}_t = \tilde{s}_t + \alpha\tilde{k}_t + (1 - \alpha)\tilde{\ell}_t. \quad (65)$$

Using the approximation $\tilde{c}_t\tilde{\ell}_t \approx 0$ and the fact that $\bar{c} = \frac{1-\alpha}{\theta}\bar{y}\frac{1-\bar{\ell}}{\bar{\ell}}$ allows us to write equation (52) as

$$\tilde{c}_t = \tilde{y}_t - \left(\frac{1}{1-\bar{\ell}}\right)\tilde{\ell}_t. \quad (66)$$

And finally, for equation (53), we use the fact that $\beta\gamma^{-1}(1-\delta) = 1 - \beta\gamma^{-1}\alpha\bar{y}\bar{k}^{-1}$ as well as approximations of the form $\tilde{x}\tilde{y} \approx 0$ to deduce

$$\tilde{c}_t = E_t\tilde{c}_{t+1} - \beta\gamma^{-1}\alpha\bar{y}\bar{k}^{-1}E_t\tilde{y}_{t+1} + \beta\gamma^{-1}\alpha\bar{y}\bar{k}^{-1}E_t\tilde{k}_{t+1}. \quad (67)$$

An important aspect of our analysis is that we want our agents to learn the new value of the steady state (that is, the vector $(\bar{c}, \bar{y}, \bar{k}, \bar{\ell})$) when a change

¹⁹See for instance the discussion in Cooley and Prescott (1994).

in growth occurs. With the system in the form of equations (64) through (67), one is in effect assuming that the steady state values are known, and so we cannot leave the system in this form. Instead, we let $c_t = \ln \hat{c}_t$, $k_t = \ln \hat{k}_t$, $y_t = \ln \hat{y}_t$, $\ell_t = \ln \hat{\ell}_t$, and $s_t = \ln \hat{s}_t$, and also $c = \ln \bar{c}$, $k = \ln \bar{k}$, $y = \ln \bar{y}$, $\ell = \ln \bar{\ell}$, and $s = \ln \bar{s} = 0$, and then rewrite equation (64) as

$$k_{t+1} = \kappa_0 + \kappa_1 y_t + \kappa_2 k_t + \kappa_3 c_t, \quad (68)$$

where

$$\kappa_0 = \left(1 - \frac{(1-\delta)}{\gamma\eta}\right) k - \frac{1}{\gamma\eta} \frac{\bar{y}}{\bar{k}} y + \frac{1}{\gamma\eta} \frac{\bar{c}}{\bar{k}} c, \quad (69)$$

$$\kappa_1 = \frac{1}{\gamma\eta} \frac{\bar{y}}{\bar{k}}, \quad (70)$$

$$\kappa_2 = \frac{(1-\delta)}{\gamma\eta}, \quad (71)$$

and

$$\kappa_3 = -\frac{1}{\gamma\eta} \frac{\bar{c}}{\bar{k}}. \quad (72)$$

Equation (65) can be written as

$$y_t = \alpha k_t + (1-\alpha) \ell_t + s_t. \quad (73)$$

For equation (66) we have

$$c_t = \pi_0 + \pi_1 y_t + \pi_2 \ell_t, \quad (74)$$

where

$$\pi_0 = c - y + \frac{\ell}{1-\bar{\ell}}, \quad (75)$$

$$\pi_1 = 1, \quad (76)$$

and

$$\pi_2 = \frac{-1}{1-\bar{\ell}}. \quad (77)$$

Next, equation (67) can be written as

$$c_t = \mu_0 + \mu_1 E_t c_{t+1} + \mu_2 E_t y_{t+1} + \mu_3 E_t k_{t+1}, \quad (78)$$

where

$$\mu_0 = \alpha\beta\gamma^{-1}\frac{\bar{y}}{k}(y - k), \quad (79)$$

$$\mu_1 = 1, \quad (80)$$

$$\mu_2 = -\alpha\beta\gamma^{-1}\frac{\bar{y}}{k}, \quad (81)$$

and

$$\mu_3 = \alpha\beta\gamma^{-1}\frac{\bar{y}}{k}. \quad (82)$$

And finally, the equation for the business cycle shock, (7), can be written as

$$s_t = \rho s_{t-1} + \vartheta_t, \quad (83)$$

where $\vartheta_t = \ln \epsilon_t$.

We now wish to reduce the system to three equations instead of five. Accordingly, we solve equation (74) for ℓ_t , substitute it into equation (73), solve the resulting equation for y_t , and substitute that solution into equations (68) and (78).

This gives the system described in the text,

$$c_t = \mathcal{B}_{10} + \mathcal{B}_{11}E_t c_{t+1} + \mathcal{B}_{12}E_t k_{t+1} + \mathcal{B}_{13}E_t s_{t+1}, \quad (84)$$

$$k_t = \mathcal{D}_{20} + \mathcal{D}_{21}c_{t-1} + \mathcal{D}_{22}k_{t-1} + \mathcal{D}_{23}s_{t-1}, \quad (85)$$

$$s_t = \rho s_{t-1} + \vartheta_t, \quad (86)$$

with $\vartheta_t = \ln \epsilon_t$, and where

$$\mathcal{B}_{10} = \mu_0 + \frac{\mu_2(\alpha - 1)\pi_0}{\pi_2 + (1 - \alpha)\pi_1}, \quad (87)$$

$$\mathcal{B}_{11} = \mu_1 + \frac{\mu_2(1 - \alpha)}{\pi_2 + (1 - \alpha)\pi_1}, \quad (88)$$

$$\mathcal{B}_{12} = \mu_3 + \frac{\mu_2\alpha\pi_2}{\pi_2 + (1 - \alpha)\pi_1}, \quad (89)$$

$$\mathcal{B}_{13} = \frac{\mu_2\pi_2}{\pi_2 + (1 - \alpha)\pi_1}, \quad (90)$$

$$\mathcal{D}_{20} = \kappa_0 + \frac{\kappa_1 (\alpha - 1) \pi_0}{\pi_2 + (1 - \alpha) \pi_1}, \quad (91)$$

$$\mathcal{D}_{21} = \kappa_3 + \frac{\kappa_1 (1 - \alpha)}{\pi_2 + (1 - \alpha) \pi_1}, \quad (92)$$

$$\mathcal{D}_{22} = \kappa_2 + \frac{\kappa_1 \alpha \pi_2}{\pi_2 + (1 - \alpha) \pi_1}, \quad (93)$$

and

$$\mathcal{D}_{23} = \frac{\kappa_1 \pi_2}{\pi_2 + (1 - \alpha) \pi_1}. \quad (94)$$

B Search methodology

A *string* is a list of economy characteristics that need to be chosen by the search algorithm. We used an eight-element string. The eight elements are the three trend break dates (one in the labor input, and two in productivity), along with the five growth rates for the periods between the break dates (two for the labor input, and three for productivity). The values of all of these elements were coded as real numbers. The program begins with a set of 50 candidate strings chosen randomly from the ranges given in Tables 1 and 2. For each of these strings, we simulate our model economy with the parameters given in the string. This simulation occurs with a low value for the business cycle shock variance (the calibrated standard deviation divided by 1000). We then record the implied trend in productivity and hours for the candidate string. To calculate the fitness of the string, we compute the mean sum of squared deviations of the actual data from the implied trend for both productivity and hours, and we add the two sums together. Strings that get low fitness scores have a better fit to the data under this metric. We then rank all of the strings based on the fitness scores.

The essence of genetic search is to update the population of strings using genetic operators. We used three classes of operators, namely, selection, crossover, and mutation. For selection, we simply kept the top 25 strings in the population to compete in the next iteration of the search. The bottom

25 strings were discarded. To keep the population constant, we created 25 new strings. Each of the 25 new strings was created as follows. We selected two strings from the top 25, and subjected them to one of three crossover routines, selected with equal probability. One routine, shuffle crossover, has each element of the two strings chosen with equal probability to create a new string. Another routine, arithmetic crossover, takes a random value for each element of the string chosen to be between the values held by the parents. The final method is to cut the strings at a randomly chosen element and swap the elements to the right in the string. Once crossover has been repeated 25 times there are 50 strings available for the next round of the search. We subjected all but the very best string in this set to a possible mutation. Mutation occurs element by element with small probability. If it occurs on a given element early in the search, then the program selects a random replacement for the existing element from the domains defined in Tables 1 and 2. If mutation occurs later in the search, then this type of mutation can be destructive to highly fit strings. Accordingly, we restricted mutation to chose new elements closer to existing elements as the search gets closer to completion.

We executed the genetic search for 500 iterations and reported the best fit string at iteration 500. Subsequent runs of the program produced results that were similar across searches.

Trend measured productivity versus U.S. data




Figure 1: The calculated trend in measured productivity implied by the model, as compared to the U.S. productivity data. The calculated trend is relatively smooth and not dissimilar to those suggested by statistical filters.

Model trend hours versus the U.S. data




Figure 2: The trend for hours is also relatively smooth. Our calculations indicate a change in trend in 1962, but otherwise the hours trend has remained approximately log-linear.

Model trend output versus the U.S. data




Figure 3: The calculated output trend compared to the U.S. data.

Model trend consumption versus the U.S. data




Figure 4: The calculated consumption trend versus the U.S. data. Consumption tends to fall below the calculated trend in the latter portion of the sample.

Model trend investment versus the U.S. data




Figure 5: Investment in the U.S. data, plotted against the calculated trend. Investment boomed in the latter portion of the sample.