6

Is the Macroeconomy Locally Unstable and Why Should We Care?

Paul Beaudry, *Vancouver School of Economics, University of British Columbia, and NBER*
Dana Galizia, *Carleton University*
Franck Portier, *Toulouse School of Economics and CEPR*

I. Introduction

There are two polar views about the functioning of a market economy. On the one hand, there is the view that such a system is inherently stable, with market forces tending to direct the economy toward a smooth growth path. According to this belief, most of the fluctuations in the macroeconomy result either from individually optimal adjustments to changes in the environment or from improper government interventions, with market forces acting to prevent the economy from being unstable. On the other hand, there is the view that the market economy is inherently unstable, and that left to itself it will repeatedly go through periods of socially costly booms and busts. According to this view, macroeconomic policy is needed to help stabilize an unruly system.

Most modern macroeconomic models, such as those used by large central banks and governments, are somewhere in between these two extremes. However, they are by design generally much closer to the first view than the second. In fact, most commonly used macroeconomic models have the feature that, in the absence of outside disturbances, the economy is expected to converge to a stable path. In this sense, these models are based on the premise that a decentralized economy is both a globally and locally stable system, and that market forces, in and of themselves, do not tend to produce booms and busts. The only reason why we see economic cycles in most mainstream macroeconomic models is due to outside forces that perturb an otherwise stable system. While such a framework is very tractable and flexible, the ubiquitous and recurrent feature of cycles in most market economies requires one

978-0-226-39560-9/2017/2017-0601$10.00

to question whether the market economy may, by its very nature, be inherently unstable and feature recurrent booms and busts, with a bust sowing the seeds of the next boom.[1]

There are at least two reasons why much of the macroeconomic profession adheres to the idea that the market economy is best described as a system with a unique stable steady state, where fluctuations are generated only by exogenous shocks. First, the majority of modern macroeconomic models based on optimizing behavior support the view of such a stable economic system. Second, when looking at the time series behavior of many labor market variables (such as the employment rate, the unemployment rate, or the job-finding rate), the estimated impulse response functions indicate that these variables respond to shocks in a manner suggestive of a stable system. For example, if one estimates a simple AR model for labor market variables, the roots of the system are generally well below one, which is consistent with the view that the system is stable.[2]

In this paper, we question this consensus. We begin by arguing that the local stability of the macroeconomic system should not be evaluated using linear time series methods, even if nonlinearities are thought to be very minor and only relevant rather far away from the steady state. Instead, we show why it is essential to allow for the possibility of nonlinearities (even if these may be very small) when exploring whether a dynamic system is locally stable. We then derive a simple class of time series models that we use to explore the stability properties of a number of macroeconomic aggregates. Using the results, we discuss why local instability should be treated as a relevant theoretical possibility when thinking about macroeconomic dynamics.[3] The main body of the text focuses on estimating the inherent dynamics of labor market variables and other macroeconomic aggregates. As we show, when we allow for simple nonlinearities in estimation, we generally find that this significantly changes the local properties of the system; in particular, it often switches from being locally stable when the nonlinear terms are excluded to being locally unstable when they are included.

After establishing that the macroeconomic system may be locally unstable, we then turn to examining the nature of the implied dynamics. This can be done by looking at how the system, with its stochastic elements turned off, evolves when it starts away from the steady state. If the steady state is unique and locally unstable, the system will not converge to a point.[4] Instead, in such a case there are three possible outcomes. One is that the system is globally unstable; that is, that the system will explode outward until it hits the economy's underlying capacity or nonnegativity constraints. This is very unlikely to be the case for labor market variables, since it would imply, for example, that we should see the unemployment rate approaching either 100% or 0%. This leaves two remaining possibilities, both involving endogenous fluctuations. The first of these is that the system may converge to a limit cycle; that is, the system settles into a recurrent pattern of booms and busts. The second possibility is that the system exhibits chaotic dynamics; that is, the system exhibits seemingly random nonrecurrent fluctuations (despite being fully deterministic) and is sensitive to initial conditions.[5] As we shall show, when we find that the system exhibits local instability, we generally find that there is a unique steady state and that the (nonstochastic) dynamics converge to a limit cycle. In this sense, our results suggest that a significant part of macroeconomic fluctuations may reflect forces that create endogenous boom-and-bust phenomena. Further, we find no evidence of chaotic behavior; that is, we do not find evidence that the deterministic part of the system exhibits sensitivity to initial conditions, a property that would render forecasting particularly difficult.

In the last section of the paper we discuss why our findings may be relevant for policy. In particular, we explore how the effects of stabilization policy may change in the presence of local instability and limit cycles. For example, we show that reducing the impact of shocks on the economy may not always help stabilize the system in such cases, and in particular, it may only change the frequency of fluctuations without necessarily decreasing their overall amplitude.

II. A Framework for Exploring Local Stability and Endogenous Cycles

In order to discuss the issue of local stability, it is helpful to focus on a variable that reflects cyclical fluctuations, but does not exhibit secular growth. Simple examples in macroeconomics are some labor market variables, such as the unemployment rate. Let us denote such a variable, in deviation from its mean, by x_t. One generally views x_t as being locally stable if there are endogenous forces present that tend to push it back toward its steady state when slightly perturbed. Suppose we believe that the process for x_t is approximately linear near its steady state. Then one may explore the local stability properties for x by first estimating an ARMA model for x, as given by

$$x_t = \tilde{A}(L)x_{t-1} + B(L)\varepsilon_t, \tag{1}$$

where $\tilde{A}(L)$ and $B(L)$ are polynomials in the lag operator and where ε_t is an i.i.d. process.[6] Let us denote by A the companion matrix associated with $\tilde{A}(L)$.[7] If we find that the largest eigenvalue of A is less than 1,[8] then the steady state is stable. However, even if we believe that this system may be close to linear near the steady state, and that the variance of ε may be small, such an approach can provide a misleading answer regarding whether the steady state is locally stable or unstable. This can be shown with the following example:

Suppose the data-generating process (DGP) is actually

$$x_t = \tilde{A}(L)x_{t-1} + F(x_{t-1}) + \varepsilon_t, \tag{2}$$

where $F(\cdot)$ is a nonlinear function with $F(0) = F'(0) = 0$, which would be the case, for example, if $F(\cdot)$ is a polynomial function with no constant or linear terms. In this case, whether the system is locally stable at the zero steady state depends only on the eigenvalues of the companion matrix A, and not on the function $F(\cdot)$. However, this does not necessarily mean that one can disregard $F(x_{t-1})$ in the estimation of equation (2). If the eigenvalues of A are all less than 1, and if the variance of ε is small, then omitting $F(x_{t-1})$ in the estimation of equation (2) is unlikely to alter any conclusions drawn about the local stability of the system. However, if the largest eigenvalue of A is greater than 1, then omitting $F(x_{t-1})$ can easily lead one to conclude that the system is locally stable when it is in fact locally unstable.[9] To see this most clearly, consider the situation where $\tilde{A}(L) = -\alpha$ and $F(x_{t-1}) = \beta x_{t-1}^3$, with $\alpha,\beta > 0$. Since the covariance between x_{t-1} and x_{t-1}^3 is positive but these terms enter equation (2) with opposite signs, omitting x_{t-1}^3 from the econometric specification will tend to bias the estimate of α toward zero. If the system is locally unstable ($\alpha > 1$), but β is large enough relative to $\alpha - 1$,[10] then one would nonetheless incorrectly infer that the system was locally stable. Note also in this case that the closer α is to 1, the more likely it is that even a small value of β will be enough to arrive at the wrong conclusion about the stability of the system.

We have explored the behavior of such biases in various model specifications using Monte Carlo methods. A typical example is the following: We simulated 5,000 1,000-period samples of the DGP

$$x_t = \alpha x_{t-1} - 0.6x_{t-2} - 0.3x_{t-3} - 0.01x_{t-1}^3 + .25\varepsilon_t, \tag{3}$$

where ε_t is i.i.d. $\mathcal{N}(0, 1)$, and where α takes values in [0.5,1.5]. For this DGP, local stability depends on whether $|\lambda|_{\max}$ is greater or smaller than 1, where $|\lambda|_{\max}$ is the maximum modulus of the eigenvalues of the companion matrix to $\tilde{A}(L) = \alpha - 0.6L - 0.3L^2$. When $\alpha \in [0.5, 1.5]$, this process implies that $|\lambda|_{\max} \in [0.94, 1.4]$. When $|\lambda|_{\max} \in [0.94, 1)$, the steady state $\bar{x} = 0$ is locally (and globally) stable, while when $|\lambda|_{\max} > 1$ it is locally unstable. In this latter case, the DGP is such that, in the absence of any shocks, x_t would be attracted toward a limit cycle (as long as $x_0 \neq 0$). Figure 1 plots, as a function of the DGP $|\lambda|_{\max}$ (which is in turn a function of α), the average estimated $|\lambda|_{\max}$ for both the well-specified model and a misspecified model in which the nonlinear term is dropped. The first thing to notice is that the estimate of $|\lambda|_{\max}$ for the well-specified model is unbiased, as indicated by the fact that it lies directly on the 45° line. Next, note that the difference between the estimated values of $|\lambda|_{\max}$ for the well-specified and misspecified models provides an indication of the importance of including the nonlinear term. As one can see in the figure, the bias in estimating the misspecified model is very small when $|\lambda|_{\max} < 1$ in the DGP. Thus, in this case, whether or not one includes the nonlinear term is virtually irrelevant for determining whether the steady state is locally stable. However, the same degree of nonlinearity becomes much more important when the $|\lambda|_{\max}$ in the DGP becomes greater than 1, as evidenced by the fact that the misspecified model continues to indicate that the zero steady state is stable (i.e., $|\lambda|_{\max} < 1$ for this model), when it is in fact locally unstable in the DGP. One inference that can be drawn from this exercise is that, when estimates from a linear model yield a maximum eigenvalue that is close to 1 in modulus, this may be an indication that omitting nonlinear terms has made the steady state appear to be locally stable when in fact it is not. Moreover, the degree of nonlinearity in the DGP may be very modest, but may nevertheless be important for assessing local stability. We return to this issue below in the context of our empirical results.

The above discussion suggests a simple protocol for exploring local stability. One can start by estimating a linear model for the variable of interest and examine its local stability properties. Then one can add nonlinear terms (e.g., higher-order polynomial terms) to the specification, estimate it, and check the local stability properties of the resulting nonlinear model. If the addition of the nonlinear terms causes the largest eigenvalue of the linear approximation of the system to go from below 1 to above 1, this is an indication that the system may in fact be locally unstable. This is the procedure we will follow. In following such a procedure, it is generally important to include in the nonlinear specification terms that are at least of order three, since such terms are more

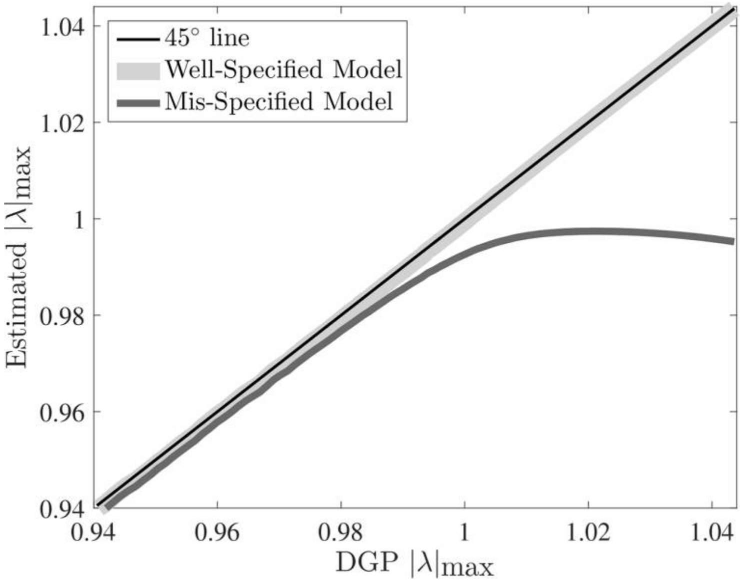

Fig. 1. Potential pitfall in assessing local stability when the DGP is nonlinear

Notes: For each of 100 values of α evenly distributed in the interval [0.5,1.5], we simulated 5,000 times a 1,000-period sample of the DGP in equation (3). From the resulting simulated data, we estimated by OLS either the well-specified equation (3) or a misspecified equation that omits the cubic term in equation (3). For each model and each α, we then computed the average maximum modulus of the eigenvalues of the companion matrix to $\tilde{A}(L) = \alpha - 0.6L - 0.3L^2$, which we denote $|\lambda|_{max}$. Figure plots, as a function of the DGP $|\lambda|_{max}$, the average estimated $|\lambda|_{max}$ for the well- and misspecified models.

likely to capture distant forces that may favor stability and whose omission would be more likely to bias in favor of inferring local stability.

A. How to Interpret Local Instability

A common property of most modern macromodels is that fluctuations reflect the effects of exogenous shocks around a stable steady state. Macroeconomic models with unstable steady states are more of a rarity. While it is true that a Walrasian equilibrium model is unlikely to have an unstable steady state, models with strategic complementarities across agents can easily give rise to such configurations, even when the complementarities are not large enough to generate multiple equilibria. To see this, consider the following extremely simple environment describing the determination of an aggregate outcome. We have a collection of n agents indexed by i who make a decision x_{it} at each date t, with the average outcome being $x_t = n^{-1}\sum_i x_{it}$. Suppose that in the absence of any interactions across agents, their actions could be described by a stable decision rule of the form $x_{it} = \tilde{A}(L)x_{it-1} + \varepsilon_t$, where the largest eigenvalue of A (where A is the companion matrix to $\tilde{A}(L)$ as defined above) is less than one. In this case, the dynamics for x_t are given by the stable system $x_t = \tilde{A}(L)x_{t-1} + \varepsilon_t$. Now suppose we modify this environment slightly to allow for strategic complements across agents, such that the decision rule of agent i now includes a term that reflects their beliefs about the average behavior of others. In particular, suppose

$$x_{it} = \tilde{A}(L)x_{it-1} + \gamma E_{it}(x_t) + \varepsilon_t, \qquad 0 < \gamma < 1$$

where $E_{it}(x_t)$ reflects agent i's expectation of others' behavior,[11] and the parameter γ governs the degree of strategic complementarity. The dynamics of x_{it} now reflect the effects of both the stable individual-level behavior, captured by $\tilde{A}(L)$, and the strategic complementarities across agents, captured by having $0 < \gamma < 1$. Assuming a rational expectations, symmetric Nash equilibrium outcome, the dynamics for x_t are then given by

$$x_t = \frac{\tilde{A}(L)}{1 - \gamma} x_{t-1} + \varepsilon_t.$$

What are the stability properties of x_t in this case? If γ is sufficient small, then x_t will remain stable. However, if γ is sufficiently large (while remaining below 1), then one may verify that the largest eigenvalue of the companion matrix for $\tilde{A}(L) \ / \ (1 - \gamma)$ will be greater than 1, in which case the system will be unstable. Note also that if the largest eigenvalue of A is below but close to 1, then only a very small degree of strategic complementarity will be enough to produce instability in the system. The above example illustrates that in environments where agents interact, and where these interactions take the form of strategic complementarities, then the appearance of a nonstable steady state can arise quite easily. This gives rise to three possibilities: (1) the economy has more than one steady state and one or more of these other steady states is an attractor, (2) the economy has more than one steady state and all are unstable, or (3) the economy has only one steady state and therefore the economy cannot converge to a steady state. While we explore all three possibilities in our empirical work, we show that the data favor the third one, so let us focus on this possibility here. If the economy has only one steady state and it is unstable, what will the dynamics of the system look like?

If the system is globally linear, then the economy must necessarily explode. However, it is unlikely that the assumption of linearity, which may be appropriate near the steady state, remains so as one moves far away from the steady state. For example, complementarity between individuals' actions, which may be present near the steady state, is likely to give way to forces of strategic substitutability when one moves far away from the steady state. This is likely to arise since, if the economy is on an exploding path, eventually resource or nonnegativity constraints will become binding. Hence, when one considers global behavior in such environments, it is necessary to recognize that some countervailing forces are likely to appear and stop the system from exploding. Such alternative forces will be manifested as some sort of nonlinearity. If this is the case, then the local instability of the steady state will create dynamics that endogenously favor economic fluctuations, as the system will neither explode nor converge to its only steady state. This outcome can take the form of a limit cycle, wherein in the absence of shocks the economy would undergo a regular and predictable boom-and-bust pattern. Alternatively, it can produce chaotic behavior, in which the economy will fluctuate in an irregular and seemingly unpredictable fashion even in the absence of shocks, and in which the deterministic dynamics would exhibit sensitivity to initial conditions, making forecasting very difficult. Note that, while chaos is an interesting theoretical possibility, as documented below and in contrast to the limit-cycle case, we do not find any evidence to suggest that it may be empirically relevant.

Our goal in this paper will be to examine the time series behavior of several cyclical indicators of macroeconomic activity while allowing for nonlinearities. We will examine (a) whether these systems exhibit unique or multiple steady states, (b) whether the implied steady states are locally stable, and (c) if the dynamics exhibit limit cycles or chaotic behavior. In principle, the empirical exercise we perform is straightforward. We begin by estimating univariate processes allowing for nonlinear forces. We then examine whether the implied dynamics are stable or explosive near the steady state(s) by looking at the root structure of the local linear approximation. Finally, we examine the behavior of the variable if we start it away from the steady state, while shutting down the effects of any stochastic terms. In particular, if we find that a variable exhibits a unique steady state and local instability, we want to examine whether it then converges to cyclical behavior. The main difficulty with this procedure is determining the class of nonlinear models to consider, as nonlinearities can take on many forms.

B. Baseline Univariate Model

Our approach to exploring local stability and possible endogenous cyclical behavior will be to estimate univariate models of the form

$$x_t = a_0 + \tilde{A}(L)x_{t-1} + F(x_{t-1}, x_{t-2}, \ldots) + \varepsilon_t, \tag{4}$$

where x_t is a cyclical variable of interest, and $F(\cdot)$ is a multivariate polynomial function with no constant or linear terms.[13] In general, we would like to allow all the lags of x that appear in $\tilde{A}(L)x_{t-1}$ to potentially enter nonlinearly through $F(\cdot)$. As we have discussed, we also wish to allow F to be of at least order three in order to avoid potentially important biases in the estimation of $\tilde{A}(L)$. When $F(\cdot)$ is of minimum order three, however, the number of terms it contains grows very fast with the number of lags of x. For example, if F is a third-order polynomial in three lags of x, then with no further restrictions the resulting regression will feature 16 nonlinear terms, while adding a fourth lag would add another 15 terms on top of this, and so on. Since we would like to allow more distant lags to enter the regression without simultaneously having an explosion of coefficients to estimate, we impose the following two restrictions: (a) F is a polynomial of order three, and (b) lags of x beyond the second enter (4) only through an accumulation term $X_{t-1} \equiv \delta\sum_{j=0}^{\infty}(1-\delta)^j x_{t-1-j}$. Under these restrictions, we can equivalently write (4) as

$$x_t = a_0 + a_1 x_{t-1} + a_2 x_{t-2} + a_3 X_t + F(x_{t-1}, x_{t-2}, X_{t-1}) + \varepsilon_t \tag{5}$$

where $F(\cdot)$ is a multivariate polynomial containing second- and third-order terms in its arguments.[14] There are two important things to note about the econometric specification (5). First, it embeds the standard $AR(2)$ process, which is known to be able to capture well many of the dynamics of macroeconomic variables. Second, it allows for more distant lags to play a role (through X_{t-1}) without requiring the estimation of too many extra parameters. We should emphasize that our inclusion of an accumulation term of this form is motivated by the class of models presented in Beaudry et al. (2015b), where the emergence of limit cycles is shown to be a possibility. See appendix for more details.

Handling Low-Frequency Movements

The specification we laid out in equation (5) is hopefully a sufficiently flexible way to allow us to offer new insight regarding the stability and cyclical properties of major macro aggregates. One remaining question is how to treat very low-frequency movements in the variable x_t. It is well known that most macroeconomic aggregates (even labor market variables) exhibit important low-frequency movements that are not generally thought to be related to the business cycle. For example, a variable such as total hours worked per capita, which according to many macro models should be stationary, often exhibits important low-frequency movements that are most often attributed to demographics and social change. We would like to remove such movements from our data, as they are likely to confound the cyclical properties we are interested in. In particular, in what follows we will interpret the framework we laid out above as modeling the cyclical element of a cyclical-trend unobserved component model. That is, we want to interpret the observed data as having been generated by a process of the form

$$x_t^O = x_t + x_t^T,$$

where x_t^O is the observed variable, x_t is the unobserved cyclical component of interest, x_t^T is a latent low-frequency trend, x_t and x_t^T are orthogonal series, and the process for x_t^T can be written

$$x_t^T = \int_{-\bar{\omega}}^{\bar{\omega}} [a(\omega)\cos(\omega t) + b(\omega)\sin(\omega t)] d\omega, \qquad \bar{\omega} \leq \pi. \tag{6}$$

Here, $a(\omega)$ and $b(\omega)$ together determine the phase and amplitude associated with trend fluctuations of frequency ω, and $\bar{\omega}$ is the maximum frequency associated with trend fluctuations.

Within this framework, in order to implement our estimation of an equation such as (5) on the cyclical component of an observed macroeconomic aggregate x_t^O, it is necessary for us to take a stance on how to remove any low-frequency movements in the data that we believe are unrelated to business-cycle behavior; that is, on how to remove x_t^T. In most of our exploration, we will do this simply using a high-pass filter that removes fluctuations with periods longer than 20 years. With quarterly data, this corresponds to setting $\bar{\omega} = 2\pi / 80$ in equation (6), and making the additional identifying assumption that, not only are trend fluctuations made up entirely of frequencies below $\bar{\omega}$ (as imposed by equation [6]), but these low frequencies are *only* associated with the trend component; that is, we also assume that the cyclical component x_t is associated only with frequencies *above* $\bar{\omega}$.[15] We will include in our analysis a Monte Carlo evaluation of the potential biases in this approach to detect local instability and limit cycles in the extracted cyclical component. For robustness, we will also report results obtained using other detrending methods.

C. *Three Embedded Models*

Expanding equation (5), we get the following expression for the evolution of x_t, which we refer to as the "full" model:

$$\begin{aligned} x_t = {} & \beta_0 + \beta_x x_{t-1} + \beta_{x'} x_{t-2} + \beta_X X_{t-1} \\ & + \beta_{x^2} x_{t-1}^2 + \beta_{x'^2} x_{t-2}^2 + \beta_{X^2} X_{t-1}^2 \\ & + \beta_{xx'} x_{t-1} x_{t-2} + \beta_{xX} x_{t-1} X_{t-1} + \beta_{x'X} x_{t-2} X_{t-1} \\ & + \beta_{x^3} x_{t-1}^3 + \beta_{x'^3} x_{t-2}^3 + \beta_{X^3} X_{t-1}^3 \\ & + \beta_{x^2 x'} x_{t-1}^2 x_{t-2} + \beta_{x^2 X} x_{t-1}^2 X_{t-1} \\ & + \beta_{x'^2 X} x_{t-2}^2 X_{t-1} + \beta_{xx'^2} x_{t-1} x_{t-2}^2 \\ & + \beta_{xX^2} x_{t-1} X_{t-1}^2 + \beta_{x'X^2} x_{t-2} X_{t-1}^2 \\ & + \beta_{xx'X} x_{t-1} x_{t-2} X_{t-1} + \varepsilon_t. \end{aligned} \tag{7}$$

This full model (7) allows for a rich nonlinear departure from the linear model. However, the drawback of such a specification is that it involves 20 parameters, making some properties hard to illustrate. For this reason, we will consider at the other extreme a very parsimonious version of equation (5), which consists of the inclusion of only one nonlinear term. We choose this one nonlinear term to be a cubic term in x_{t-1}. This choice, while somewhat arbitrary, is made for two reasons. First, from an intuitive standpoint, including odd-order nonlinear terms—of which the third is the lowest order—allows for a symmetric treatment of large positive and negative deviations from the steady state. Second, if there is only to be a single nonlinear term, having that term be solely in x_{t-1} seems to us to be the simplest choice. Our "minimal" model is therefore given by

$$x_t = \beta_0 + \beta_x x_{t-1} + \beta_{x'} x_{t-2} + \beta_X X_{t-1} + \beta_{x^3} x_{t-1}^3 + \varepsilon_t. \tag{8}$$

As we shall see, we get very similar results based on either the minimal model or the full model, but the minimal model allows for simpler illustrations.

Finally, we will also consider an "intermediate" model that lies between these two extremes. One possibility in order to find the best intermediate specification is, for each variable under study, to start from the full model and sequentially remove insignificant variables. As this procedure is somewhat arbitrary, we instead choose as our intermediate model the case where we keep most third-order terms from the full model, while eliminating all second-order terms. This has the feature of significantly reducing the number of parameters, while still allowing for relatively rich nonlinearities. In particular, of the seven third-order terms in equation (7), we keep the three cubic terms and the two cross-terms involving x_{t-1} and X_{t-1}, so that our intermediate model becomes

$$\begin{aligned} x_t &= \beta_0 + \beta_x x_{t-1} + \beta_{x'} x_{t-2} + \beta_X X_{t-1} \\ &\quad + \beta_{x^3} x_{t-1}^3 + \beta_{x'^3} x_{t-2}^3 + \beta_{X^3} X_{t-1}^3 \\ &\quad + \beta_{x^2 X} x_{t-1}^2 X_{t-1} + \beta_{xX^2} x_{t-1} X_{t-1}^2 + \varepsilon_t. \end{aligned} \tag{9}$$

Our three models—full, intermediate, and minimal—together offer a tractable nonlinear framework for examining whether certain macroeconomic variables may exhibit local instability and limit cycles. We will also compare the behavior of these nonlinear models with two specifications that do not feature any nonlinear terms: the "linear" model, given by

$$x_t = \beta_0 + \beta_x x_{t-1} + \beta_{x'} x_{t-2} + \beta_X X_{t-1} + \varepsilon_t, \tag{10}$$

and a simple $AR(2)$ model,

$$x_t = \beta_0 + \beta_x x_{t-1} + \beta_{x'} x_{t-2} + \varepsilon_t. \tag{11}$$

D. Estimation with Total Hours

Data Treatment and Estimation Results

The motivation we used to derive our reduced-form model is one based on a cyclical indicator, such as hours worked per capita. Accordingly, the first measure we examine is the (log of) BLS total hours worked, deflated by total population. The mechanisms behind our motivating model were not meant to explain low-frequency fluctuations such as those related to demographic and sociological change (e.g., aging, female labor market participation, etc.). As can be seen in panel (a) of figure 2, total hours (gray line) has exhibited important low-frequency

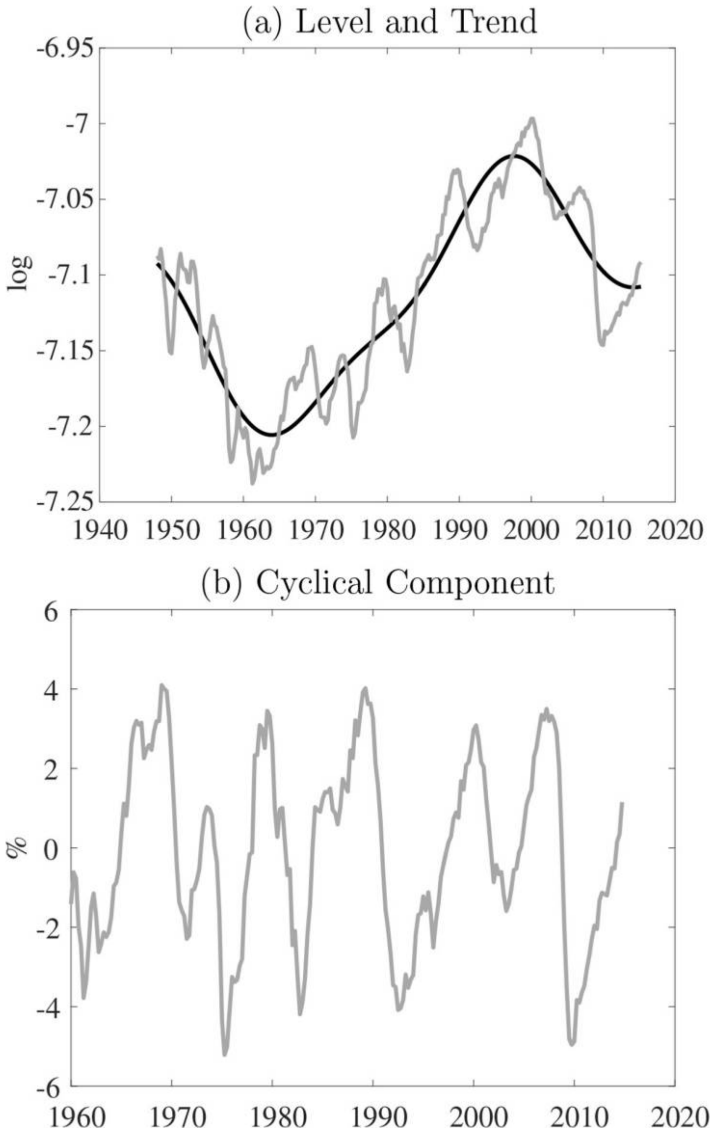

Fig. 2. Trend and cycle (high-pass filter, 80 quarters), total hours

Notes: Total hours are measured as the log of BLS total economy hours worked, deflated by total population. The trend of that series is obtained by removing from the level series its filtered component, where the filter is an 80-quarter high-pass filter, over the sample 1948:Q1–2015:Q2. The cyclical component is expressed in percentage deviation from the level series.

movements over the last 50 years. For this reason, we begin by filtering the data in order to remove these movements. We do this by using a (linear) high-pass filter that retains only fluctuations associated with periods of less than 20 years.[16] The trend (black line) and its filtered component are displayed in panels (a) and (b), respectively, of figure 2.

Once we have a series for $x = h$ (filtered hours), we then need to construct the accumulated hours series analogous to X_t to above, denoted H_t. To construct this series, we truncate the accumulation at N quarters, so that $H_t = \delta\sum_{j=0}^{N-1}(1 - \delta)^j h_{t-j}$. We set $N = 40$ so that, for example, observations of h from 1950:Q2 to 1960:Q1 are used to compute H_{1960Q1}. We assume for now that $\delta = 0.05$. Later we show that our results are robust to different values of δ.

For each of our two linear and three nonlinear models, OLS parameter estimates obtained using our total hours series are displayed in table 1, with t-statistics in parentheses. One can verify that in the nonlinear models there is always at least one significant nonlinear term. More informative are the Wald tests (see table 2) that compare the fit of the different models. For example, when testing the $AR(2)$ model against the full model, we test the joint nullity of all the coefficients except β_h and $\beta_{h'}$ (the coefficients on h_{t-1} and h_{t-2}, respectively) in the full model. The first row of table 2 shows that we reject the simple linear $AR(2)$ model against all alternatives. The second row reaches the same conclusion for the linear model. The third row shows that, at a 1% level of confidence, the minimal model is not rejected against any of the nonlinear alternatives. We first discuss the results obtained with this minimal model, since it is easiest to understand and illustrate. We will show later that results are mostly the same with the intermediate and full models.

Model Selection Using LASSO

An alternative way of selecting the model specification is to use statistical methods developed in the data-mining literature. One particular method, the LASSO (Least Absolute Shrinkage and Selection Operator), allows for both model selection and shrinkage. The method, proposed by Tibshirani (1996), consists of minimizing the residual sum of squares of the full model subject to the restriction that the sum of the absolute values of the coefficients be less than a constant. Because of the nature of this constraint, it tends to produce some coefficients that are exactly

Table 1
Estimates of Various Models, Total Hours

	Value				
Variable	*AR*(2)	Linear	Minimal	Intermediate	Full
Constant	–0.00	–0.01	–0.02	–0.01	–0.07
	(–0.02)	(–0.24)	(–0.53)	(–0.29)	(–0.85)
h_{t-1}	1.42	1.31	1.39	1.28	1.49
	(23.75)	(20.80)	(19.73)	(12.70)	(10.02)
h_{t-1}^2	—	—	—	—	–0.10
	(—)	(—)	(—)	(—)	(–1.51)
h_{t-1}^3	—	—	–0.01	–3e-03	–0.09
	(—)	(—)	(–2.39)	(–0.76)	(–1.85)
h_{t-2}	–0.48	–0.34	–0.34	–0.24	–0.44
	(–8.05)	(–4.98)	(–2.43)	(–5.12)	(–2.75)
h_{t-2}^2	—	—	—	—	–0.06
	(—)	(—)	(—)	(—)	(–0.76)
h_{t-2}^3	—	—	—	–4e-03	0.11
	(—)	(—)	(—)	(–0.83)	(1.38)
H_{t-1}	—	–0.25	–0.27	–0.03	S0.23
	(—)	(–4.11)	(–4.38)	(–0.20)	(1.09)
H_{t-1}^2	—	—	—	—	0.25
	(—)	(—)	(—)	(—)	(2.03)
H_{t-1}^3	—	—	—	–0.27	–0.40
	(—)	(—)	(—)	(–2.45)	(–2.10)
$h_{t-1}h_{t-2}$	—	—	—	—	0.17
	(—)	(—)	(—)	(—)	(1.15)
$h_{t-1}H_{t-1}$	—	—	—	—	0.25
	(—)	(—)	(—)	(—)	(1.89)
$h_{t-2}H_{t-1}$	—	—	—	—	–0.28
	(—)	(—)	(—)	(—)	(–1.88)
$h_{t-1}^2h_{t-2}$	—	—	—	—	0.27
	(—)	(—)	(—)	(—)	(1.57)
$h_{t-1}^2H_{t-1}$	—	—	—	6e-03	–0.32
	(—)	(—)	(—)	(0.42)	(–2.08)
$h_{t-2}^2H_{t-1}$	—	—	—	—	–0.30
	(—)	(—)	(—)	(—)	(–1.46)
$h_{t-1}h_{t-2}^2$	—	—	—	—	–0.48
	(—)	(—)	(—)	(—)	(–2.40)
$h_{t-1}H_{t-1}^2$	—	—	—	0.04	–0.37
	(—)	(—)	(—)	(0.87)	(–1.64)
$h_{t-2}H_{t-1}^2$	—	—	—	—	0.47
	(—)	(—)	(—)	(—)	(1.79)
$h_{t-1}h_{t-2}H_{t-1}$	—	—	—	—	0.77
	(—)	(—)	(—)	(—)	(2.25)

Note: The total hours series has been filtered with an 80-quarter, high-pass filter over the sample 1948:Q1–2015:Q2. Estimation is then done over the sample 1960:Q1–2014:Q4.

Table 2
Wald Test for the Different Models, Total Hours

	H_1			
H_0	Linear (%)	Minimal (%)	Intermediate (%)	Full (%)
$AR(2)$	0.006	0.002	0.004	0.003
Linear	—	1.788	1.544	0.343
Minimal	—	—	7.63	1.20
Intermediate	—	—	—	2.81

Notes: The five models correspond to equations (11), (10), (8), (9), and (7). Wald test corresponds to the test of joint nullity of the coefficients of those variables that are in H_1 and not in H_0.

zero, and hence act as a model-selection device. The value of the constant in the constraint is chosen to minimize the Akaike Information Criterion. The estimated equation is given by

$$\begin{aligned} x_t = & -0.07 + 1.3x_{t-1} - 0.23x_{t-2} - 0.03x_{t-1}^2 + 0.21X_{t-1}^2 \\ & + 0.04x_{t-1}x_{t-2} + 0.29x_{t-1}X_{t-1} - 0.30x_{t-2}X_{t-1} \\ & - 0.01x_{t-1}^3 - 0.21X_{t-1}^3 - 0.06x_{t-2}^2X_{t-1} + 0.03x_{t-1}X_{t-1}^2 \\ & + 0.05x_{t-1}x_{t-2}X_{t-1} + \varepsilon_t. \end{aligned} \tag{12}$$

One can check that only 12 of the 18 variables of the full model are present in equation (12), the others being assigned a coefficient of zero. We will explore the robustness of the results by comparing results obtained with minimal, intermediate, full, and LASSO models.

Local Instability and Limit Cycles

For the sake of clarity, let us write down the estimated $AR(2)$, linear, and minimal models after shutting down the stochastic component:

$$\begin{cases} h_t = -0.00 + 1.42h_{t-1} - 0.48h_{t-2}, \\ h_t = -0.01 + 1.31h_{t-1} - 0.34h_{t-2} - 0.25H_{t-1}, \\ h_t = -0.02 + 1.39h_{t-1} - 0.34h_{t-2} - 0.27H_{t-1} - 0.01h_{t-1}^3. \end{cases}$$

Note that, even though our minimal model is nonlinear, it is easy to check that it has only one steady state at $h = -0.09$. In order to examine the local stability of each of these three equations, we compute the larg-

Table 3
Some Statistics for the Different Models, Total Hours

	AR(2)	Linear	Minimal	Intermediate	Full
R^2	0.94	0.94	0.94	0.94	0.95
Adj. R^2	0.94	0.94	0.94	0.94	0.95
DW	2.22	2.09	2.12	2.09	2.09
Max eig. modulus	0.86	0.96	1.01	1.01	{1.02,1.2,1.03}

Notes: The five models correspond to equations (11), (10), (8), (9), and (7). Adj. R^2 is the adjusted R^2. DW stands for the Durbin-Watson statistics for the test of autocorrelation of the residuals. $DW = 2$ corresponds to no autocorrelation. The last line gives the value of the maximum eigenvalue for the local dynamics of each model in the neighborhood of the existing steady states.

est eigenvalue of each system when linearized around its steady state. For the minimal model, this linear approximation is simply[17]

$$h_t = 1.39h_{t-1} - 0.34h_{t-2} - 0.27H_{t-1}.$$

Maximum eigenvalues are displayed in table 3.

For the *AR*(2) model, the maximum eigenvalue is 0.86.[18] For the linear model, it is higher (0.96), but still less than 1. The steady states implied by the point estimates of the *AR*(2) and linear models therefore indicate local stability, which is not surprising. For the nonlinear minimal model, on the other hand, the maximum eigenvalue is 1.01, so that the steady state is locally unstable. This may seem to conflict with the visual impression given by the data (figure 2, panel [b]) that the economy does not explode. In fact, it is precisely the third-order term h_{t-1}^3 that prevents explosion, since it enters negatively in the estimated equation. The steady state is thus locally unstable, but as h moves away from the steady state, the term h^3 acts as a centripetal force that pushes the economy back toward the steady state if it strays too far. As a result, the economy will oscillate without explosion and without convergence to a fixed point. This can be seen by using the minimal model to construct a deterministic forecasted path for h, conditional on information at date T_0. Formally, this deterministic forecast is computed as:

$$\begin{cases} \tilde{h}_{T_0} & = h_{T_0}, \\ \tilde{h}_{T_0-1} & = h_{T_0-1}, \\ \tilde{H}_{T_0-1} & = H_{T_0-1}, \\ \tilde{h}_t & = \hat{\beta}_0 + \hat{\beta}_h\tilde{h}_{t-1} + \hat{\beta}_{h'}\tilde{h}_{t-2} + \hat{\beta}_H\tilde{H}_{t-1} + \hat{\beta}_{h^3}\tilde{h}_{t-1}^3 \quad \forall\, t > T_0, \\ \tilde{H}_t & = (1-\delta)\tilde{H}_{t-1} + \delta\tilde{h}_t \quad \forall\, t \geq T_0. \end{cases}$$

We forecast h over 1,000 periods with the estimated model. We start (arbitrarily) from the first trough of the sample (i.e., 1961:Q3) and plot the forecasted path in the model state space (h_t, h_{t-1}, H_t). As we can see in panel (a) of figure 3, the trajectory is quickly attracted to a limit cycle. Figure 4 shows that this convergence holds more generally: starting from two initial conditions—one inside the limit cycle and one outside—the system converges to the limit cycle, suggesting that it is indeed attractive. Note that, as h is highly autocorrelated, there is little loss of information if one projects the limit cycle onto the (h_t, H_t) plane (panel [b] of figure 3) instead, and we will often use that representation of the limit cycle in what follows.

It is quite interesting that such a simple specification, which was motivated from our previous theoretical analysis (see Beaudry et al. 2015b), displays a limit cycle. We will document later using Monte Carlo methods that this result is unlikely to be an artifact of our data treatment. Before exploring further the properties of the minimal model, it is useful to show that we obtain similar results with the other nonlinear models. Table 3 shows that the largest eigenvalue is also greater than 1 for the intermediate and full models. As shown in figure 5, these two models (as well as the LASSO model) also generate limit cycles, with all specifications producing cycles of similar amplitudes for total hours.[19]

Let us now focus on the minimal model and look at the size and frequency of the limit cycle. As illustrated in figure 6, the deterministic forecast for the minimal model displays a cycle whose frequency and amplitude is close to the ones observed in the data. Figure 6 also displays 95% confidence bands around the deterministic forecast, and shows that after 60 quarters, the confidence bands have the same amplitude as the deterministic cycle. Panel (a) of figure 7 compares the deterministic forecast of the minimal model to the $AR(2)$ and linear model, while panel (b) compares the deterministic forecast of the intermediate, full, and LASSO models.

Other Labor Market Variables

We now explore the robustness of the results we obtained using total hours to using other labor market variables instead. In particular, we examine the behavior implied by estimating our minimal model using, sequentially, nonfarm business hours, the job-finding rate, and the rate of unemployment. For each of these variables, we report in figure 8 the deterministic forecast path in (x_t, X_t)-space, as well as x_t over time,

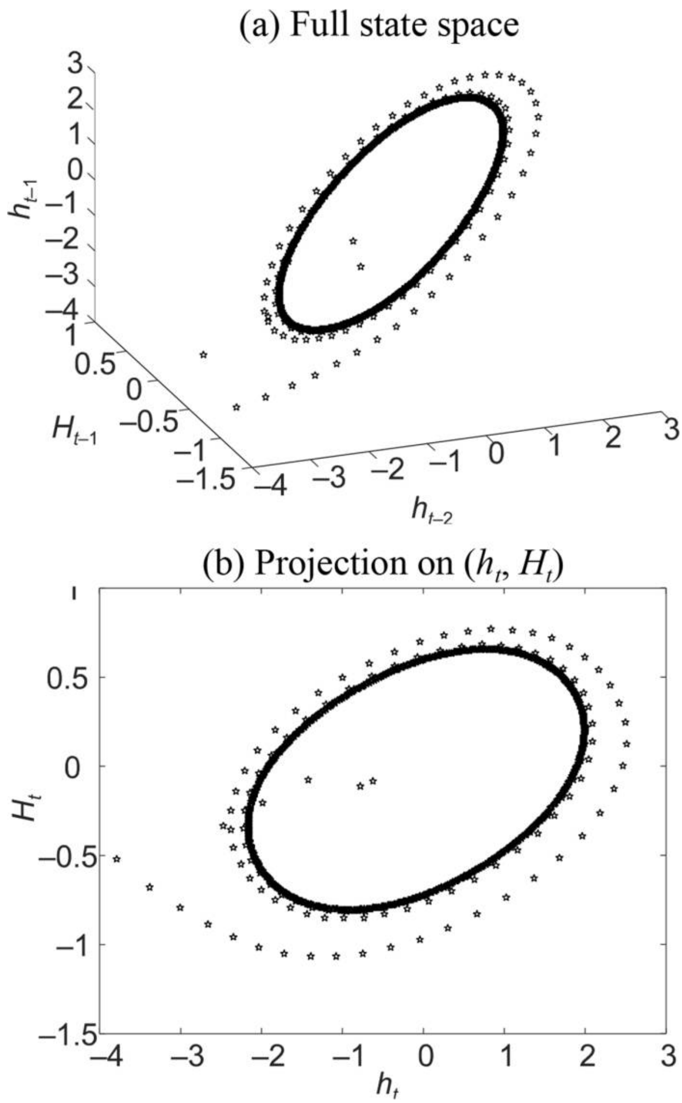

Fig. 3. The limit cycle in the minimal model, total hours

Note: This corresponds to the deterministic simulation of the minimal model (8), starting in 1961:Q3. The model has been estimated from 1960 to 2014 using 80-quarter, high-pass-filtered total hours.

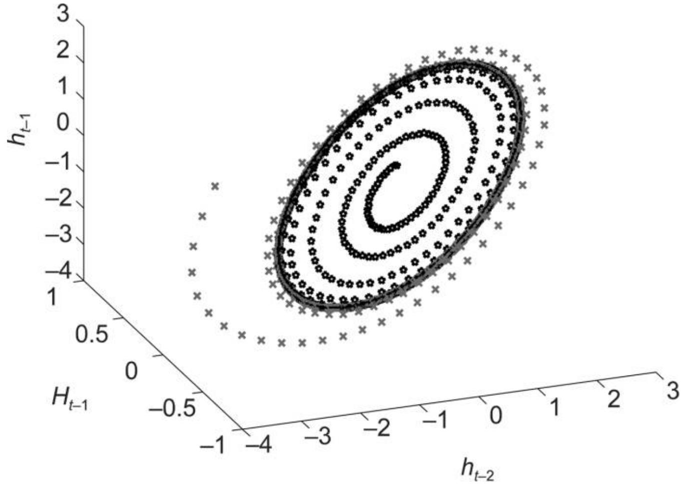

Fig. 4. Convergence to the limit cycle in the minimal model, total hours

Notes: This corresponds to the deterministic simulation of the minimal model (8), starting from two points arbitrarily chosen inside and outside the limit cycle. Those two points are, respectively, $(h_0, h_{-1}, H_{-1}) = (0.1, 0.1, 0.1)$ and $(h_0, h_{-1}, H_{-1}) = (6, 4, 10)$. The models have been estimated from 1960 to 2014 using 80-quarter, high-pass-filtered total hours. For graphical reasons, the first 13 points of each path are not shown.

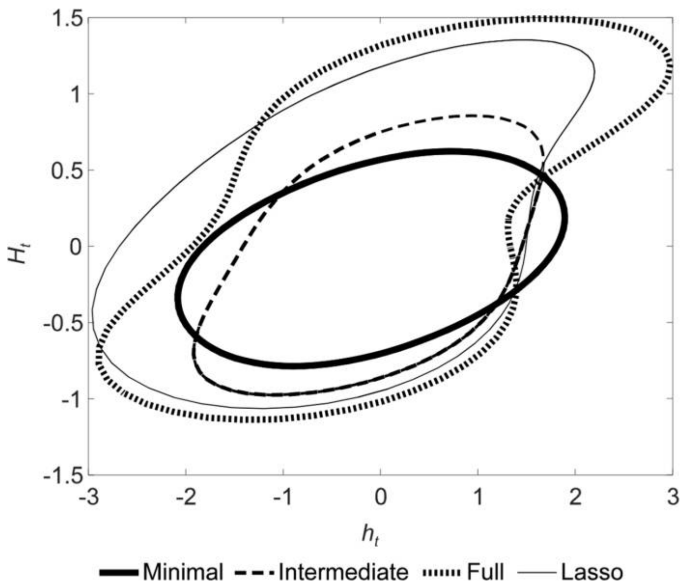

Fig. 5. The limit cycle in the four models in (h_t, H_t)-space, total hours

Note: This corresponds to the deterministic simulation of the minimal (8), intermediate (9), LASSO (12), and full (7) models, starting (arbitrarily) in 1961:Q3. The model has been estimated from 1960 to 2014 using 80-quarter, high-pass-filtered total hours.

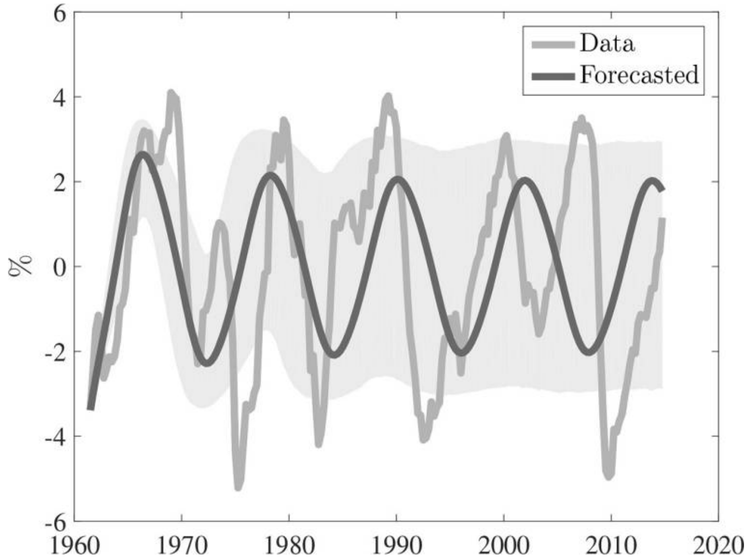

Fig. 6. Forecasted path as of 1961:Q3 with the minimal model, total hours

Notes: The dark gray line labeled "Forecasted" is the deterministic simulation of the minimal (8) model, starting (arbitrarily) in 1961:Q3, and the light gray line represents observed filtered hours. The gray area represents the 95% confidence band for the deterministic forecast, as obtained from 10,000 Monte Carlo simulations using draws from the posterior distribution over the model parameters.

where x is the variable under consideration and X its accumulated level (defined as before by $X_t \equiv \delta\sum_{i=0}^{\infty}(1-\delta)^i x_{t-i}$). Panels (a) and (b) of figure 8 correspond to the forecasted paths for nonfarm business hours, panels (c) and (d) are those associated with the job-finding rate, and finally panels (e) and (f) correspond to the rate of unemployment. In all three cases, we start the forecast from 1961:Q3. As is clear from these figures, a limit cycle appears in all three cases, and these cycles have durations close to nine years, with amplitudes similar to the actual data.

How Significant Are Limit Cycles?

Up to now, we have checked for the existence of limit cycles using only the models' point estimates. Beyond the either/or results implied by these point estimates, it is of interest to use the sampling variability to quantify how strongly the data support the presence of instability and limit cycles. To do so, we derive estimated parameter distribu-

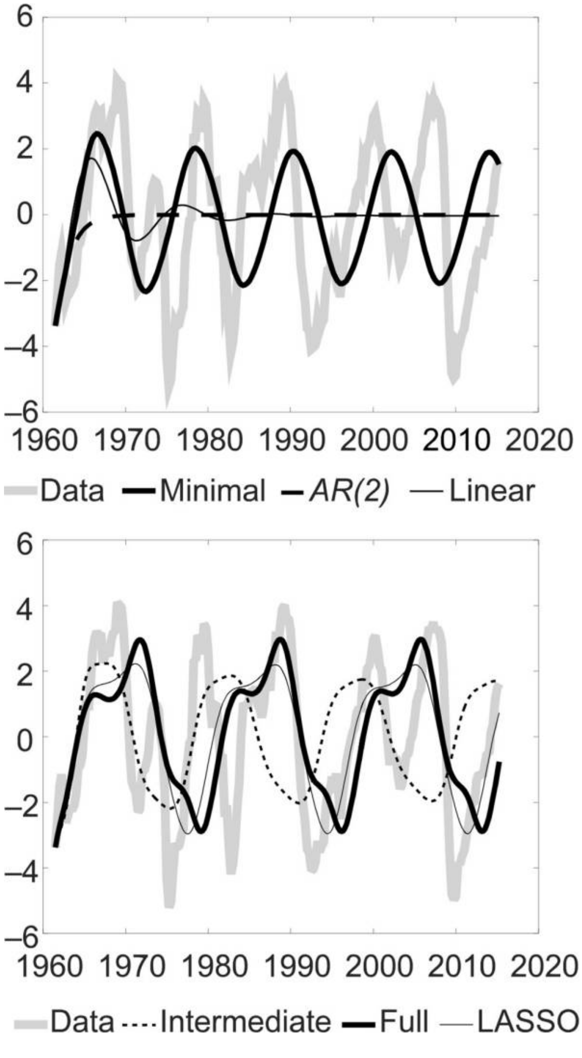

Fig. 7. The limit cycle in the four models, time series of h_t, total hours

Notes: This corresponds to the deterministic simulation of the minimal (8), intermediate (9), LASSO (12), and full (7) models, starting (arbitrarily) in 1961:Q3. The model has been estimated from 1960 to 2014 using 80-quarter, high-pass-filtered total hours.

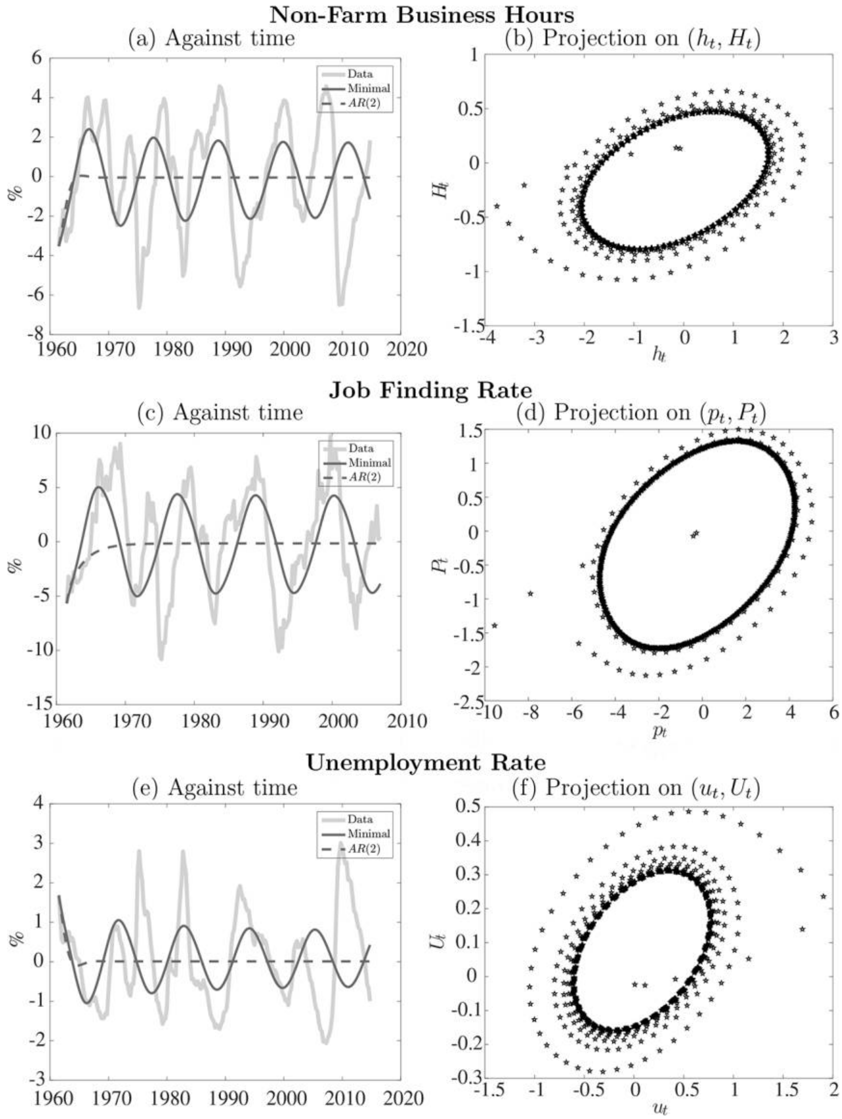

Fig. 8. The limit cycle in the minimal model, other labor market variables

Notes: Figure plots deterministic simulation of minimal model (8) starting in 1961:Q3. Model was estimated from 1960 to 2014 (–2007 for job-finding rate) using 80-quarter, high-pass-filtered series of nonfarm business hours, job-finding rate (from Shimer 2012), and unemployment rate. P and U stand for "cumulated" rates, constructed in the same way as X_t above.

Table 4
Frequency of Limit Cycles for Labor Market Variables

	Minimal (%)	Intermediate (%)	Full (%)
Total hours	74	73	78
Nonfarm bus. hours	70	55	70
Job-finding rate	76	69	59
Unemployment rate	60	32	50*

Notes: Bootstrap method used to generate 1,000 replications of the minimal, intermediate, and full model estimates for each labor market variable. For example, first cell indicates that a limit cycle was found in 74% of replications of minimal model. Data are always detrended with 80-quarter, high-pass filter before estimation. See appendix for details of samples and variable definitions. A * indicates that in more than half of replications, simulation was explosive. In such cases, we redrew bootstrap innovations until we obtained 1,000 nonexplosive replications.

tions for the minimal, intermediate, and full models using a bootstrap procedure. For each bootstrapped data set, we check for the existence of a limit cycle, which we define as meeting the following conditions: (a) starting in the simulation period corresponding to 1961:Q3,[20] the deterministic simulation of the nonlinear model estimated on this data set converges to a limit cycle; and (b) the Wald test rejects the linear model against the nonlinear one at 5%. Results are displayed in table 4. As can be seen from the table, the detection of limit cycles ranges mainly between 50 and 78%. This indicates that the null hypothesis of the absence of limit cycles in the data cannot be rejected at conventional levels of significance. Nonetheless, a person with a diffuse prior would in most cases infer that limit cycles are more likely than not.

Next, rather than focusing on our holistic definition of a limit cycle, we consider how strong the results are with regard specifically to the local instability that we have identified. To this end, we have computed—for the linear and minimal models estimated on total hours—the bootstrap distribution of the maximum eigenvalue. For the minimal model, this eigenvalue is computed for the linear approximation at the steady state.[21] Figure 9 plots the PDF for the bootstrap distributions. As can be seen, the distribution is almost entirely below one for the linear model. In contrast, allowing for a single cubic term in the estimation shifts this distribution significantly to the right, with 79% of the mass now being concentrated above 1. Again, this pattern is insufficient to reject the null hypothesis of stability at conventional levels of significance. Nonethe-

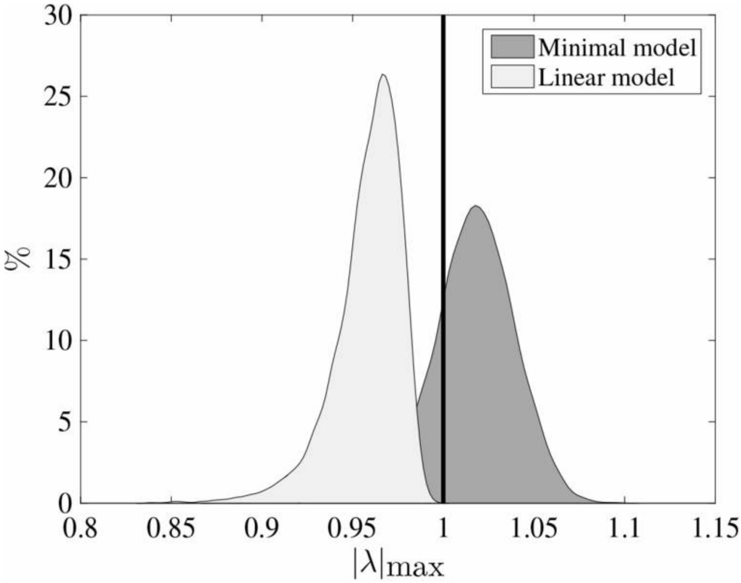

Fig. 9. Maximum eigenvalue in the linear and minimal model, total hours
Notes: Figure plots the estimated distribution of maximum eigenvalue for the linear and minimal model estimated on total hours. The densities are estimated using 1,000 bootstrap replications of each model. The models have been estimated from 1960 to 2014 using 80-quarter, high-pass-filtered series of total hours.

less, the results do clearly suggest that local instability is a relevant possibility.

Finally, as one further check on the strength of our results, we attempt to answer the following question. Suppose the DGP is indeed stable and linear. Accounting for parameter uncertainty, how likely are we to erroneously conclude that the steady state is unstable simply by allowing for nonlinearities in the estimation? To answer this question, for total hours we have computed a bootstrap distribution for the maximum eigenvalue of each nonlinear model estimated on data generated by either the $AR(2)$ or linear model, where we account for uncertainty in the parameters of the latter models. Table 5 reports instability p-values from these bootstrap distributions associated with the actual observed maximum eigenvalues (i.e., those reported in table 3). The results suggest that, accounting for sampling variability, we are unlikely to see the magnitude of instability we have observed in the data for the nonlinear models if the true DGP is one of the linear models. For example, if the true DGP were the linear model, we would obtain maximum eigenvalues

Table 5
p-Values for Maximum Eigenvalues When DGP Is Linear / Stable, Total Hours

	Estimated Model		
DGP	Minimal (%)	Intermediate (%)	Full (%)
AR(2)	0.08	4.63	9.06
Linear	1.91	14.13	13.91

Notes: In this table, the DGP is either the (stable) *AR*(2) or linear model, with parameters drawn from a bootstrapped distribution. For each of 10,000 parameter draws, a new data set was generated (again via bootstrap), and the nonlinear models were estimated on this data, yielding a distribution for the maximum eigenvalue for each nonlinear model when the DGP is linear and stable. Table reports instability *p*-values—that is, the probability for maximum eigenvalues of nonlinear models estimated on actual data (see table 3) to be larger than one in modulus.

for the intermediate and full models at least as great as the ones we observed in the data less than 15% of the time, and less than 2% of the time for the minimal model.

Multiplicity of Steady States

Since we are examining dynamics in a nonlinear framework, we need to recognize the potential for a multiplicity of attractors in the system. First, let us note that all the estimated equations are backward looking, so that there is (trivially) never indeterminacy in the solution to our equations.[22] In contrast, we do need to check for the existence of multiple steady states or multiple limit cycles for our estimated models. Consider first the case of multiple steady states, and focus on the minimal model when shocks are turned off:

$$h_t = \beta_0 + \beta_h h_{t-1} + \beta_{h'} h_{t-2} + \beta_H H_{t-1} + \beta_{h^3} h_{t-1}^3,$$
$$H_t = (1 - \delta)H_{t-1} + \delta h_t.$$

Estimation typically gives β_0 very close to zero, β_h larger than one, $\beta_{h'}$ and β_H negative, and β_{h^3} negative and small. Assume for simplicity that β_0 is exactly zero, such that $h = 0$ is a steady state. Nonzero steady states of the minimal model, if they exist, are the two opposite real numbers $\bar{h}$ and $-\bar{h}$, with[23]

$$\bar{h} = \sqrt{\frac{1 - \beta_h - \beta_{h'} - \beta_H}{\beta_{h^3}}},$$

Restricting to the case where $\beta_{h^3} < 0$, we have multiplicity of steady states if and only if

$$\beta_h + \beta_{h'} + \beta_H > 1. \tag{13}$$

Therefore, although the model is nonlinear, it can generically be in a parameter configuration with a unique steady state.

Note that, as β_H approaches zero, the condition for local instability of the zero steady state is,

$$\beta_h + \beta_{h'} > 1. \tag{14}$$

Thus, for β_H close to zero, equations (13) and (14) coincide: the occurrence of limit cycles, which requires local instability, necessarily implies multiplicity of steady states. However, when the economy has a longer memory (i.e., β_H is negative but not arbitrarily close to zero), the two conditions no longer coincide, and in particular it is possible for the steady state to be both unique and locally unstable. This is typically what the data will suggest in the minimal model.

In the case of multiplicity, the question is whether the nonzero steady states are stable or unstable when the zero steady state is unstable. Taking a first-order Taylor expansion of equation (8) around , we have the following dynamics (omitting constant terms):

$$h_t = \underbrace{(\beta_h + 3\bar{h}^2\beta_{h^3})}_{\bar{\beta}_h}h_{t-1} + \beta_{h'}h_{t-2} + \beta_H H_{t-1}. \tag{15}$$

Note that the same equation holds for the local dynamics around $-\bar{h}$. As we are looking at situations where $\beta_{h^3} < 0$, we will have $\bar{\beta}_h < \beta_h$, such that it is possible that the two nonzero steady states are stable when the zero steady state is unstable. In such a case, although there exists a limit cycle in the neighborhood of zero, it would be possible to have trajectories that would start close (but not arbitrarily close) to zero and converge to $\bar{h}$ or $-\bar{h}$. The same argument can also be made for the intermediate and full models. The occurrence of additional, attractive steady states depends ultimately on the parameters of the model. Our findings suggest, however, that such a situation is unlikely to be relevant. In particular, given the coefficients as estimated on total hours, the minimal and intermediate models both feature unique steady states. Moreover, while there are three steady states for the full model (located at $h =$ –0.77, 0.22, and 1.23), all three of these steady states are locally unstable, which implies that the system cannot converge to a steady state in any of our models estimated on total hours.

Note that, for the full model, it is in principle possible that two or even all three of the steady states could be surrounded by distinct attractive limit cycles. It is not possible to prove analytically that this is not the case, but it can be checked numerically. Figure 10 displays (in (h_t, H_t)-space) four deterministic forecasts from the full model: two starting near the intermediate ($h = 0.22$) steady state, and one each starting near the high and low steady states. Steady states are indicated by stars in the figure. One observes that all trajectories apparently converge to the same closed orbit, suggesting that the existence of multiple attractive limit cycles is unlikely.

The Power of Our Limit-Cycle-Detection Procedure

One may worry that our procedure, especially our detrending procedure, might generate spurious limit cycles even if the DGP is actually

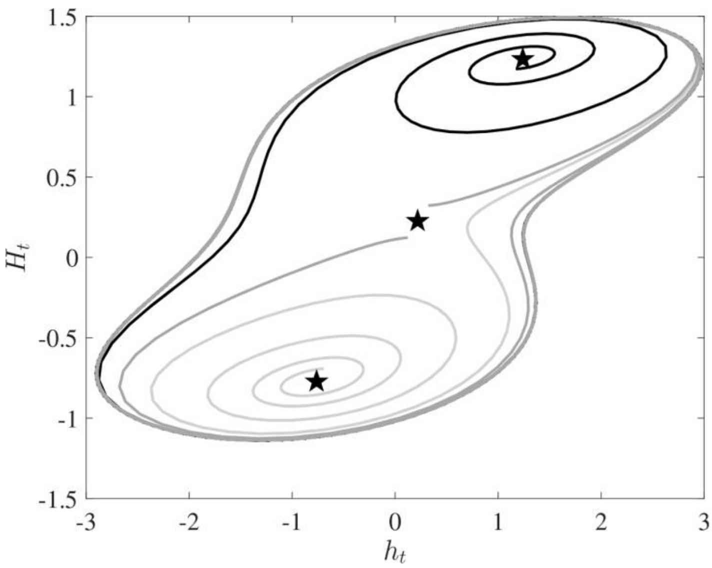

Fig. 10. Convergence to the limit cycle from the neighborhood of each of the three steady states, full model, total hours.

Notes: Figure shows deterministic forecasts starting from the neighborhoods of each of the three steady states (represented by stars). All four trajectories converge to the same limit cycle. The models have been estimated from 1960 to 2014 using 80-quarter, high-pass-filtered total hours.

locally stable, or alternatively that it may fail to detect a limit cycle when there is indeed one. We investigate these issues here using Monte Carlo analysis. First, we check that we do not spuriously find limit cycles when there are none present. To do so, we first filter (logs of) total hours with an 80-quarter, high-pass filter, so that

$$h_t^L = h_t^T + h_t$$

where h_t^L stands for levels, h_t^T for the trend and where h_t is the cyclical component. We then estimate an $AR(2)$ on the filtered data. This $AR(2)$, which is the one displayed in table 1, will serve as the DGP. We then simulate the estimated (stable) $AR(2)$ to obtain an artificial cyclical series $\hat{h}_t$, then add it to the actual trend series h_t^T to generate an artificial level series

$$\hat{h}_t^L = h_t^T + \hat{h}_t$$

By construction, these simulated level series do not feature limit cycles. We may then apply our limit-cycle-detection procedure to these simulated series in order to verify that it does not spuriously detect limit cycles where there are none. That is, we treat each simulated level series in the same way we have treated the actual level series: first filter, then estimate the minimal, intermediate, and full models, then check for the existence of a limit cycle. For each model and simulation, we check for the existence of a significant limit cycle, which we define as meeting the following conditions: (a) starting in 1961:Q3,[24] the deterministic simulation of the model converges to a limit cycle; and (b) the Wald test rejects the linear model against the nonlinear one at 5%. We perform 10,000 simulations, and report the results in the first row of table 6. The three models spuriously detect limit cycles less than 5% of the time, suggesting that our procedure is highly unlikely to detect a limit cycle if one does not exist in the DGP. The second row of table 6 shows similar results for the case where the DGP is the estimated linear model instead of the $AR(2)$.[25]

Next, we test for the possibility of missing a limit cycle with our procedure when it *does* exist in the data. In order to do so, we take the estimated minimal, intermediate, or full model as the DGP and follow the same procedure as above. The results are displayed in the bottom three rows of table 6, and show that limit cycles are relatively hard to detect using our procedure. Typically, when the estimated model is well specified (meaning that it has the same form as the DGP), our procedure detects limit cycles about half of the time.

Table 6
Percentage of Limit Cycle Detection in Monte-Carlo Simulations Simulated Hours

	Estimated Model		
DGP	Minimal (%)	Intermediate (%)	Full (%)
AR(2)	1.9	2.5	2.3
Linear	5.8	4.5	3.3
Minimal	54	34	20
Intermediate	43	51	40
Full	57	43	58

Note: In this table, the DGP is alternatively the estimated *AR*(2), linear, minimal, intermediate, or full model. The models have been estimated from 1960 to 2014 using 80-quarter, high-pass-filtered total hours, and are then simulated 10,000 times.

From these Monte Carlo simulations, we conclude that it is unlikely that we will spuriously detect limit cycles, and that correctly detecting a limit cycle when one is present occurs at best 58% of the time.

Robustness

Here we briefly discuss the robustness of our main results. We maintain focus on the case where our measure of activity is total hours per capita. We begin by examining the robustness of our results with respect to our choice of δ in the computation of H and to our detrending procedure. Recall that, in order to construct the cumulated series H, we need to make a choice for the value of δ. Table 7 shows that a limit cycle is detected generically for the minimal model, as long as δ is not too large for the intermediate and full models, and as long as δ is also not too small in the Intermediate model.

Next, table 8 considers filters other than the 80-quarter, high-pass one used so far, and also considers the model estimated with the LASSO method. Let us first note that the high-pass filter is a linear filter and therefore should not be spuriously introducing nonlinearities, so that in general we should not expect to find cycles where there are none simply by choosing a particular parameter for the filter. On the other hand, if there exists a limit cycle in the data but the filter parameter we choose removes the frequencies associated with that cycle, our procedure would in general fail to identify this cycle. Further, given the

Table 7
Robustness to δ, Total Hours

	Minimal (%)	Intermediate (%)	Full (%)
δ = .001	LC (58)	–(31)	LC (89)
δ = .01	LC (59)	–(40)	LC (85)
δ = .1	LC (76)	LC (67)	LC (66)
δ = .2	LC (40)	–(40)	–(40)

Notes: In this table, we estimate the minimal, intermediate, and full models under various values of δ. "LC" means that the estimated model displays a limit cycle. The number in parenthesis indicates the fraction of the 1,000 bootstrap replications for which we found a limit cycle with the linear model being rejected at 5% against the nonlinear one. All the models have been estimated from 1960–2014 using 80-quarter, high-pass-filtered total hours.

Table 8
Robustness to Detrending, Total Hours

	Minimal (%)	Intermediate (%)	Full (%)
High-pass 100	LC (58)	LC (35)	LC (72)
High-pass 90	LC (78)	LC (72)	LC (68)
High-pass 70	LC (80)	LC (80)	LC (80)
High-pass 60	LC (57)	LC (51)	LC (82)
High-pass 50	–(41)	LC (38)	–(67)
Band-pass(6,3,2)	–(12)	–(41)	–(37)
Hodrick-Prescott (λ = 1600)	–(28)	–(30)	–(20)
Polynomial trend (3rd-order)	–(18)	LC (29)	LC (75)
Polynomial trend (4th-order)	–(19)	LC (30)	LC (76)
Polynomial trend (5th-order)	–(53)	LC (49)	LC (76)
No detrending	–(.1)	–(21)	–(40)

Notes: In this table, we estimate the minimal (8), intermediate (9), and full (7) models using various detrending methods. "LC" means that the estimated model displays a limit cycle. The number in parenthesis indicates the fraction of the 1,000 bootstrap replications for which we found a limit cycle with the linear model being rejected at 5% against the nonlinear one. All the models have been estimated from 1960 to 2014 using total hours.

strong implicit restrictions imposed on the data when estimating our three models—which are strongest in the minimal model and weakest in the full model—we should in general expect to find that our choice of filter affects our ability to identify limit cycles even if they exist in the data, and that this sensitivity to the filter should be strongest in the minimal model and weakest in the full model. The results presented in table 8 largely support these predictions. In particular, we report results using different high-pass filters from 100 to 50 quarters, as well as for commonly used band-pass (6,32) and Hodrick-Prescott (λ = 1600) filters. We also report results for when we detrend using polynomial time trends of order three, four, and five.

Three conclusions emerge from the table. First, the detrending procedures that remove only the lowest frequencies (i.e., the 60 quarters or more high-pass filters and the polynomial time trends) are generally associated with limit cycles, while the procedures that also remove the middle range of frequencies (i.e., the band-pass [6,32] and Hodrick-Prescott [λ = 1600] filters) are not. Given the considerations discussed above, this is consistent with the view that the data features a medium-frequency limit cycle, and that two of the most commonly used filters in the business-cycle literature remove those frequencies by design. This in turn suggests the need to focus on fluctuations that are longer than traditionally thought to be associated with business cycles.

Second, the minimal model is quite sensitive to the choice of filter, with no clear pattern to this sensitivity, while the intermediate and full models are relatively insensitive. This is consistent with the view that it may be the strong model restrictions underlying the minimal model that drive much of its observed sensitivity to the filter (as opposed to spurious patterns generated by the filter itself). In particular, the fact that the more flexible specifications typically support the existence of a medium-frequency limit cycle, while the very restrictive specification may or may not depending on the filter used, is supportive of the hypothesis that a medium-frequency limit cycle is present in the data but the minimal model is too restrictive to detect it consistently. Finally, note that limit cycles are at best only found 40% of the time when we do not filter the data.

E. Behavior in Other Countries

Here we extend our exploration of macroeconomic dynamics using data on the unemployment rate in other countries.[26] These results are

Table 9
Limit Cycles in Other Countries

Unemployment rate	Minimal (%)	Intermediate (%)	Full (%)
Australia	–(.7)	–(2)	–(11)*
Austria	LC (98)	–(38)	LC (76)
Canada	–(11)	–(22)	–(11)
Denmark	–(12)	LC (45)	LC (88)
France	–(3)	–(26)	LC (52)*
Germany	LC (79)	–(49)	LC (88)
Japan	–(7)	–(33)	–(50)
Netherlands	LC (62)	LC (78)	–(94)
Sweden	LC (65)	–(64)	LC (66)
Switzerland	–(16)*	–(52)	–(27)
United Kingdom	–(4)	–(3)	LC (78)

Notes: In this table, we estimate the minimal (8), intermediate (9), and full (7) models, considering various labor market variables. "LC" means that the estimated model displays a limit cycle. The number in parenthesis indicates the fraction of the 1,000 bootstrap replications for which we found a limit cycle with the linear model being rejected at 5% against the nonlinear one. Data are always detrended with an 80-quarter, high-pass filter before estimation. See appendix for details of samples and variable definitions. A star (*) indicates that in more than half of the replications, the simulation with bootstrap innovations was explosive, so that the pseudodata sample could not be used for estimation. In this case, we have drawn again bootstrap innovations until we obtained 1,000 nonexplosive replications.

presented in table 9. We estimate our three models for 11 countries. As can be seen in the table, we find some evidence of limit cycles in about half of the countries. Although this is somewhat low, interestingly this ratio is similar to the detection rate we have found in our Monte Carlo exercise. Looking closely at each country makes it clear why we do find limit cycles in some and not in others. As an example, let us consider Switzerland and the Netherlands. Data and deterministic forecasts for these two countries are displayed in figure 11.

Switzerland and the Netherlands both have marked low-frequency dynamics. In the case of Switzerland (figure 11, panels [a] and [c]), there appears to have been a change of regime around 1990; there was almost zero unemployment in the 1970s and 1980s, but a 3.5% average since 1990. Besides this low-frequency variation, what emerges from figure 11, panel (c), is the big recession of 1991–1993, and the progressive recovery from 1993 to 2000. Instead of a relatively regular cycle, as we have observed in the United States, we see one big cycle and four smaller ones, which is a pattern that cannot be easily captured by our reduced-form models. Hence, the minimal model estimated on this data does not exhibit a limit cycle, and the dynamics of forecasted

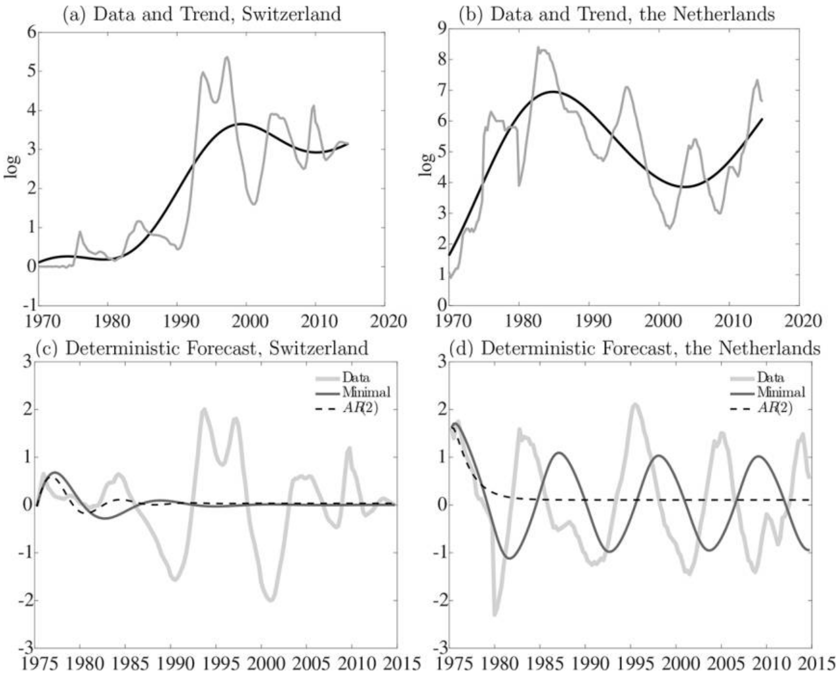

Fig. 11. Unemployment dynamics in Switzerland and the Netherlands

Notes: Unemployment rate is observed over the sample 1970:Q1–2014:Q4. Unemployment rate is measured as "Registered Unemployment Rate" for Switzerland and as "Harmonized Unemployment Rate: All Persons" for the Netherlands. The H series is constructed assuming $N = 20$, so that estimation starts in 1975:Q2. The trend of these series is obtained by removing from the level series its filtered component, where the filter is an 80-quarter, high-pass filter. The cyclical component is expressed in percentage deviation from the level series. Deterministic forecasts are done starting from 1975:Q1, with the estimated $AR(2)$ (11) and minimal (8) models.

unemployment as of 1975 are not very different from the ones obtained with an $AR(2)$ model.[27] Things are very different for the Netherlands (figure 11, panels [b] and [d]). The low-frequency dynamics feature a marked increase from the early 1970s to the late 1980s, followed by a prolonged decline through to early in the first decade of the twenty-first century. The key determinants of those trend movements are the first and second oil shocks, followed by important labor market reforms in the 1990s. Fluctuations around this trend, however, have been very regular—and in particular we do not observe one big cycle dwarfing the others, as was the case for Switzerland—and thus the estimated minimal model features a limit cycle.

F. *Quantity Variables*

Difficulties

We now turn to the estimation of our nonlinear reduced-form equation using quantity variables. Before performing this exercise, we want to highlight the added difficulties associated with quantity variables. To this end, consider the simple case in which output is given by $y_t = \theta_t + h_t$, where θ is TFP, and in which hours follows as before a process of the type

$$h_t = a_0 + a_1 h_{t-1} + a_2 h_{t-2} + a_3 H_t + F(h_{t-1}, h_{t-2}, H_{t-1}) + \varepsilon_t. \tag{16}$$

In such a case, the dynamic equation for output y is given by

$$\begin{aligned} y_t &= a_0 + a_1 y_{t-1} + a_2 y_{t-1} + a_3 \Sigma_{j=1}^{\infty}(1-\delta)^{j-1}(y_{t-j} - \theta_{t-j}) + \varepsilon_t \\ &\quad + F(y_{t-1} - \theta_{t-1}, y_{t-2} - \theta_{t-2}, \Sigma_{j=1}^{\infty}(1-\delta)^{j-1}(y_{t-j} - \theta_{t-j})) \\ &\quad + \theta_t - a_1\theta_{t-1} - a_2\theta_{t-2} - a_3 + \varepsilon_t. \end{aligned} \tag{17}$$

If we had a good measure of θ_t, then the estimation of equation (17) would not be any more difficult than our estimation of equation (16) using hours worked. However, in the absence of a good measure of θ, the estimation of equation (17) becomes difficult. Simply using a filtered measure of y_t as a representation of $y_t - \theta_t$ is unlikely to perform well. This can be seen in practice by running a Monte Carlo experiment as follows. We take our estimates of the minimal, intermediate, or full model obtained when run on the 80-quarter, high-pass-filtered total hours series. As shown previously, all three of these estimated models feature a limit cycle. We then use this estimated model to create an artificial cyclical series of hours, to which we add the log of labor productivity taken from the data. This gives us a log-level series for artificial output. We then treat this level series in the same way we will treat the actual level series of output: filter, estimate the minimal, intermediate, and full models, then check for the occurrence of a significant limit cycle. We perform 10,000 simulations, and report in table 10 the percentage of simulations for which a limit cycle is (correctly) detected. The three models hardly detect limit cycles in these simulated output series, even if by construction they contain limit cycle forces. Detection rates are around 15% when the estimated model embeds these DGP specifications. Despite this low detection rate, we nonetheless proceed

Table 10
Percentage of Limit-Cycle Detection in Monte Carlo Simulations, Simulated Output

	Estimated Model		
DGP	Minimal (%)	Intermediate (%)	Full (%)
Minimal	14	13	5
Intermediate	12	15	6
Full	12	11	11

Note: In this table, the DGP is alternatively the estimated minimal, intermediate, and full models. The models have been estimated from 1960 to 2014 using 80-quarter, high-pass-filtered total hours, and are then simulated 10,000 times.

and examine whether our procedure detects limit cycle in the actual quantity data. As we will see, the detection rate is much lower when using quantity variables than when using labor market variables—as could be expected given the Monte Carlo exercise.

Estimation Results

We estimate the three models (7), (8), and (9) on output, total consumption, durable goods expenditures, fixed investment, various components of investment, and the capacity utilization rate. Each of these variables is again first detrended using an 80-quarter, high-pass filter. Estimation results are summarized in table 11. While we do not find much evidence of limit cycles using the minimal model, interestingly, we find in the full model limit cycles detected at a frequency of over 50% using all variables except equipment investment.

III. Should We Care?

The previous section explored whether cyclical movements in macroeconomic variables exhibited characteristics suggestive of local instability and limit cycles. In this section, we examine whether one should care about identifying/differentiating such a possibility. There are at least two potential reasons why one might care. First, it is possible that it could lead to improved forecasting ability. Second, it is possible that policy changes could have differing effects on stable and locally unstable limit-cycle models. Our explorations suggest that, for forecasting purposes, differentiating between the two cases is likely of

Table 11
Existence of a Limit Cycle for Quantity Variables

	Minimal (%)	Intermediate (%)	Full (%)
Output	–(11)	LC (60)	–(53)
Nonfarm bus. output	—	(20) LC (64)	–(54)
Consumption	–(3)	LC (84)	LC (85)
Fixed investment	–(28)	–(12)	LC (54)*
Structures	LC (80)	LC (84)	LC (75)
Durables	–(32)	LC (75)	LC (85)
Residential	–(15)	LC (34)	–(62)*
Equipment	–(19)	–(14)	–(27)
Utilization	LC (86)	LC (80)	LC (53)*

Notes: In this table, we estimate the minimal, intermediate, and full models, considering various goods market variables. "LC" means that the estimated model displays a limit cycle. The number in parenthesis indicates the fraction of the 1,000 bootstrap replications for which we found a limit cycle with the linear model being rejected at 5% against the nonlinear one. Data are always detrended with an 80-quarter, high-pass filter before estimation. See appendix for details of samples and variable definitions. A star (*) indicates that in more than half of the replications the simulation with bootstrap innovations was explosive, so that the pseudo-data sample could not be used for estimation. In this case, we have drawn again bootstrap innovations until we obtained 1,000 nonexplosive replications.

second-order importance. This is most easily observed by recognizing that the variance of the one-step-ahead forecast error for the linear model is almost identical to that for any of the nonlinear models. However, as we argue in this section, we believe that, for policy purposes, identifying whether macroeconomic dynamics may contain limit cycles forces can be of first-order importance. In particular, we discuss how a policy aimed at countering the effects of shocks can lead to very different outcomes depending on whether the model exhibits limit cycles, even though the two models are similar in terms of standard forecast performance.

A. Theoretical Results

The goal of stabilization policy is to reduce inefficient macroeconomic volatility. In practice, this is often equated to the goal of reducing the cyclical volatility of output and employment. One stabilization approach is to direct policy toward countering or nullifying the impact of certain exogenous forces hitting the economy. In linear models, such a strategy will generally be effective, as a reduction in the variance of exogenous shocks translates directly in to a reduction in the variance of endogenous variables. However, in an environment where limit-cycle forces may be present, this strategy may no longer be effective. In particular, in this section we will show that when limit-cycle forces are present, (a) a reduction in the variance of shocks may actually increase the variance of endogenous variables; (b) even if reducing the variance of shocks reduces the variance of the outcomes, the relationship between the shock variance and outcome variance is likely to be quite weak; and (c) the principal effect of reducing the shock variance is likely to dampen higher-frequency movements while accentuating longer cyclical movements (whereas such a change in the shape of the spectrum would not arise in a linear model).

In order to look at these issues, we will focus again on a univariate setup.[28] For example, consider an environment where the dynamics of a state variable of interest x_t may we written as

$$x_t = F(x_{t-1}, x_{t-2}, \ldots) + \varepsilon_t,$$

where ε_t is an exogenous shock with variance σ_ε^2. If the function $F(\cdot)$ is linear and is such that x is stationary,[29] then the variance of x will be strictly increasing in the variance of ε. However, if the function $F(\cdot)$ is nonlinear, then the link between the variance of x,[30] denoted σ_x^2, and the variance of ε may be much more complicated. In particular, even in the case where $F(\cdot)$ is such that σ_x^2 exists and is finite, σ_x^2 will not necessarily be an increasing function of σ_ε^2. More to the point, when $F(\cdot)$ is such that the deterministic system $x_t = F(x_{t-1}, x_{t-2}, \ldots)$ admits a limit cycle, then σ_x^2 will often be decreasing in σ_ε^2 over some range. To illustrate this point, consider an order three data-generating process similar to those we have estimated in section II, but in which we remove both the accumulation and the second autoregressive terms for analytical tractability:

$$x_t = -\beta_x x_{t-1} + \beta_{x^3} x_{t-1}^3 + \varepsilon_t, \tag{18}$$

Under the restrictions

$$1 < \beta_x < 2, \tag{$\mathcal{R}_1$}$$

and

$$\beta_{x^3} = 1 + \beta_x, \tag{$\mathcal{R}_2$}$$

one may verify that when $\sigma_\varepsilon^2 = 0$ the system (18) possesses a limit cycle of period two[31] in which x alternates between $\bar{x}$ and $-\bar{x}$, where

$\bar{x} \equiv \sqrt{\beta_x - 1 / \beta_{x3}}$. Moreover, this two-cycle attracts all orbits for which $0 < |x_0| < 1$. Now consider an increase in σ_ε^2. How does this affect the variance of x? Proposition 1 indicates that an increase in σ_ε^2 can lead to a decrease in σ_x^2.

Proposition 1. *If the random variable x_t evolves according to equation (18) with restrictions $\mathcal{R}_1$ and $\mathcal{R}_2$, then the variance of x_t will be decreasing in σ_ε^2 when σ_ε^2 is sufficiently small.*

Proof. See appendix.

The intuition for Proposition 1 reflects the nature of the nonlinear forces generating the limit cycle. Near the limit cycle, a shock that moves x outside of the cycle (i.e., so that $|x_t| > \bar{x}$) is countered by stronger forces pushing it back toward the cycle than is a shock that moves x *inside* the cycle (i.e., so that $|x_t| < \bar{x}$). Hence, shocks that move x inward are more persistent, and therefore, on average, for σ_ε^2 sufficiently small the system spends more time inside the cycle than outside of it. This effect causes the overall variance of x to be decreasing in σ_ε^2. Although the scope of Proposition 1 is quite limited, since it covers only a small class of data-generating processes that support limit cycles, it nevertheless makes clear that the relationship between shock variance and output variance can be negative when the underlying data-generating process admits a limit cycle.

B. *Illustrating the Effects of Changes in Shock Variance*

Proposition 1 suggests that, in models featuring limit cycles, there may be a more complicated relationship between shock variance and outcome variance than the simple positive one usually encountered in stable models. To follow up on this, we have explored numerically whether this feature may arise in the richer environments of the previous section. In all cases we found that, for small values of σ_ε^2, an increase in the shock variance has either a small but noticeable negative effect or no discernible effect at all on the variance of h, while for larger values the effect is positive but substantially weaker than in a linear environment. In this sense, a robust take-away appears to be that, in environments that support limit cycles of the size we believe may be in macroeconomic data, the link between shock variance and outcome variance is likely to be very muted relative to its linear counterpart. To illustrate this, in panel (a) of figure 12 we plot the relationship between σ_ε and σ_h for the intermediate and the linear models as estimated on

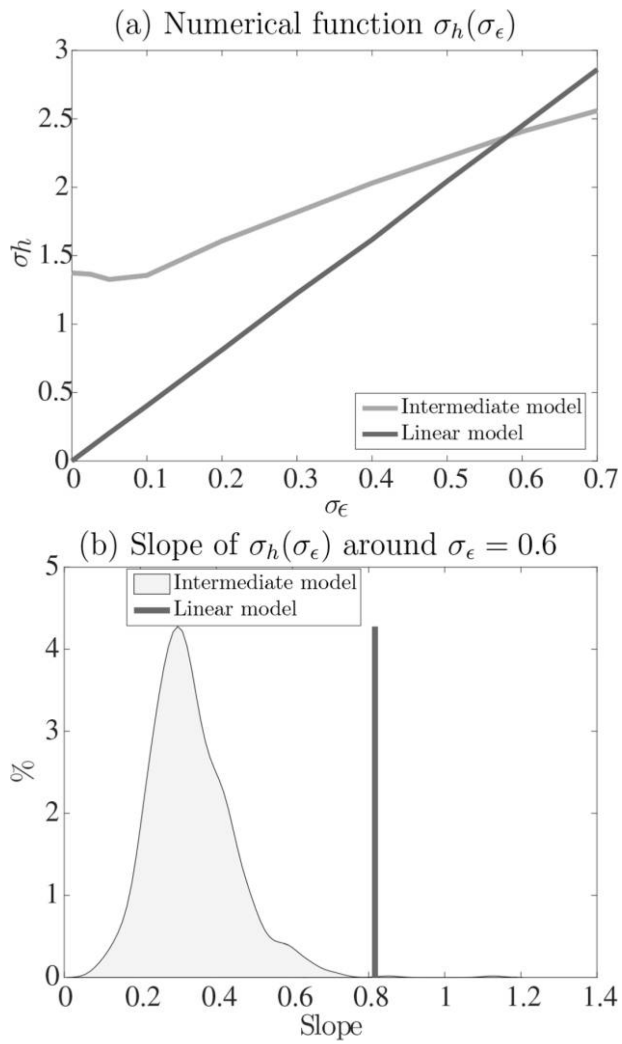

Fig. 12. Relationship between σ_h and σ_ε

Notes: For each model, we derive parameter distributions by bootstrapping with 1,000 replications. Then, for each value of σ_ε, we simulate $T_{tot} = T_{brn} + T$ periods of data, with T_{brn} = 1,000 and T = 100,000. After discarding the first T_{brn} periods of data, we compute σ_h^2 as the time average $T^{-1} \sum(h_t - \mu_h)^2$, where $\mu_h = T^{-1}\sum h_t$. For the sake of clarity, we report only the point estimates in panel (a). In panel (b), we report the bootstrap distribution of the slope of $\sigma_h(\sigma_\varepsilon)$, evaluated for the observed level $\sigma_\varepsilon \approx 0.6$, for the intermediate model. For the linear one, we report the slope at the point estimates for the parameters. For both models, the slope is computed as $[\sigma_h(0.7) - \sigma_h(0.5)] / 0.2$.

total hours worked per capita.[32] As was shown in the previous section, in the absence of shocks the intermediate model exhibits limit cycles. On the contrary, the estimated linear model has a stable steady state, so that when $\sigma_\varepsilon = 0$ we have $\sigma_h = 0$.

The are two notable features that emerge from panel (a) of figure 12. The first is that, as in the simple example of Proposition 1, the relationship between σ_ε and σ_h is actually negative for low values of σ_ε. Let us stress again, however, that this is not the pattern we think is most relevant. The second—and, we believe, more important—feature that emerges from the figure is the extent to which the relationship between σ_ε and σ_h is much weaker in the case of the nonlinear model as compared to the linear model. For example, increasing σ_ε from 0.1 to 0.5 causes σ_h to increase by 400% in the linear model, while it increases by only 62% in the intermediate model. The fact that the plotted relationship has a nonzero intercept for the limit-cycle model is not surprising, but the fact that the slope of this relationship is so much weaker on average is telling. This result is further documented in panel (b) of the same figure. Taking into account sample uncertainty, we plot the bootstrap distribution of the slope of $\sigma_h(\sigma_\varepsilon)$, evaluated for the observed level $\sigma_\varepsilon \approx 0.6$, together with the slope-point estimate for the linear model. This panel shows that the slope is lower in the intermediate model at any level of significance. This pattern suggests that stabilization policy aimed at mitigating the effects of shocks may be of more limited value than suggested by estimation of the linear model. The goal of stabilization policy in such a case may instead be to focus on identifying and mitigating the underlying forces that generate and sustain limit cycles.[33]

To further emphasize how changes in input volatility affect outcome behavior in models with or without limit cycles, in Figure 13 we plot the spectrum of the outcome variable h for several different values of σ_ε. Panel (a) shows spectra for the intermediate nonlinear specification, while panel (b) shows spectra for the linear one.

For the linear model, panel (b) illustrates that the effect of changing the volatility of ε is very simple: a fall (rise) in σ_ε decreases (increases) the spectrum substantially and uniformly. In contrast, as shown in panel (a), the effect of changing σ_ε in the model that features a limit cycle is primarily to change the shape of the spectrum. In particular, a fall in σ_ε decreases the importance of higher frequencies and accentuates the importance of lower frequencies, while a rise in σ_ε does the reverse. For example, as we can see in the figure, when we reduce σ_ε by

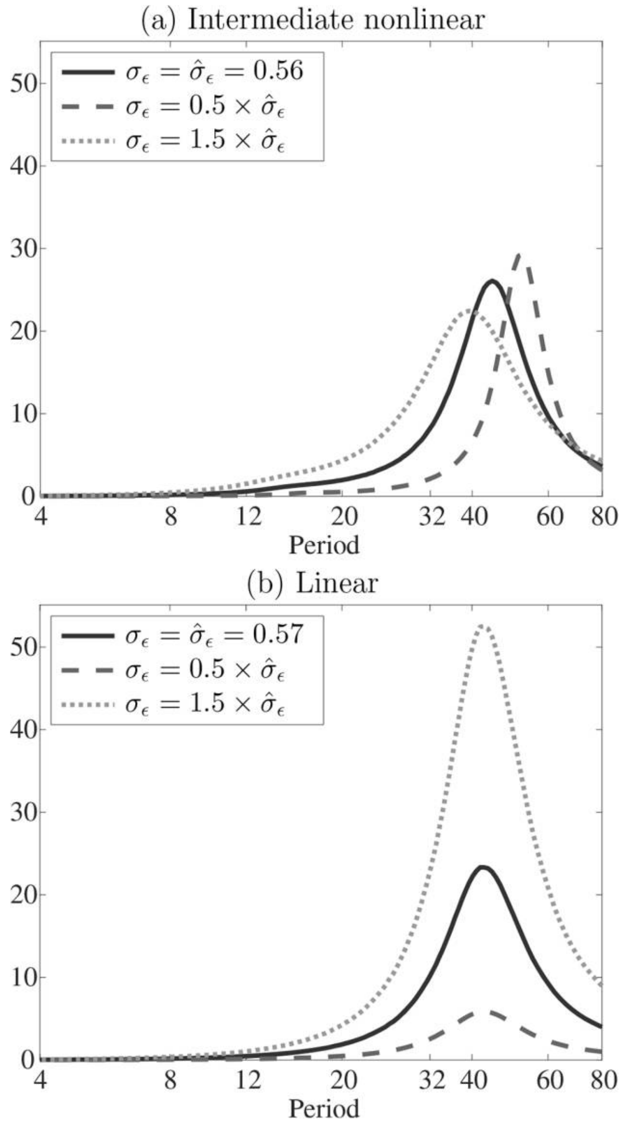

Fig. 13. Spectrum of h

Notes: For each model and each value of σ_ε, we simulate 100,000 data sets, each of length $T_{tot} = T_{brn} + T$ quarters, with $T_{brn} = 50{,}000$ and $T = 1{,}000$. For each simulated data set, we obtain the spectrum of h after discarding the first T_{brn} periods, then average the result across all data sets.

a third, from 0.84 (dotted line) to 0.56 (solid line)—where the latter is the value estimated for the data—the spectrum at periodicities below 40 quarters decreases substantially, while the importance of periodicities above 40 quarters increases substantially. Such observations are interesting, as they provide a potentially different perspective on how a reduction in shock volatility—such as is often thought to have arisen during the period of the great moderation—can affect the economy. If the underlying system exhibits limit cycles, a great moderation in input volatility would result in a less volatile economy in the short term, but would tend to *increase* the importance of the lower-frequency movements associated with cycles lasting around 10 years.

In summary, much of the focus of traditional macroeconomic stabilization policy is based on the view that fluctuations are primarily driven by shocks. If instead limit-cycle forces play an important part in generating fluctuations, this warrants a rethinking of how best to conduct policy. In particular, in this section we have emphasized that aiming to counter or to nullify shocks to the macroeconomy is likely to be much less effective at reducing economic volatility if limit cycles are present.[34] Such policies may somewhat reduce volatility, but the main effect may be to simply change the frequency at which fluctuations arise, with higher-frequency movements being dampened by the policy at the cost of amplifying lower-frequency movements. As a consequence, in such a case it could be best to de-emphasize a framework focused on countering shocks and move toward a stabilization framework that aims to reduce the forces that may be causing the limit cycles. For example, in our previous work (Beaudry et al. 2015b), we have emphasized the role of strategic complements in producing limit cycles. An effective stabilization policy in such a situation is one that reduces these complementarities. If the complementarity is related to precautionary saving associated with unemployment risk as in Beaudry et al. (2015a, 2015b), then policy should focus on mitigating the effect of unemployment risk on individuals by offering better unemployment insurance. Policies aimed at countering shocks in such an environment will not, in general, be hitting the important margins.

IV. Conclusion

The first aim of this paper has been to examine whether fluctuations in macroeconomic aggregates may be best described as reflecting the effects of shocks to an otherwise stable system, or whether such fluctuations may instead—to a large extent—reflect an underlying instability in the macroeconomy that repeatedly gives rise to sustained boom and bust phenomena. To examine this question, we have proposed a simple nonlinear time-series framework that has the potential to capture instability and limit-cycle behavior if it is present in the data. The framework we adopted was motivated by previous work that emphasized the role of strategic complementarities in creating local instability in environments with accumulated goods. We first focused on the behavior of labor market variables, and especially the behavior of total hours worked per capita, and found intriguing evidence in favor of the view that the macroeconomy may be locally unstable and produce limit-cycle forces. For example, we found that aggregate labor market variables often indicate the presence of forces that favor recurrent business cycles with a duration of close to nine years and an amplitude of 4–5%. When looking at aggregate output measures, we found less robust evidence.

One of the main challenges in this endeavor was to find a powerful tool for identifying local instability if it is present. As we showed using Monte Carlo methods, the approach we adopted is not very powerful, and this is especially true when looking at output measures. This suggests that it would be fruitful in future work to develop a more powerful method. Given this difficulty, it is interesting that we have been able to find as much evidence supporting the notion of local instability and limit cycles as we have. Moreover, even when we have not found local instability, we have almost always found that the system is very close to being unstable.[35] This suggests to us that the macroeconomy should be viewed as either locally unstable or very close to being locally unstable.

In the second section of the paper, we have briefly examined what the presence of local instability and limit-cycle forces could imply in terms of stabilization policy. Our main observation on this front is that stabilization policy aimed at countering the shocks hitting the economy may turn out to be very ineffective at stabilizing the economy. In particular, we have shown that a reduction in shock variances may have little effect on overall macroeconomic volatility if the economy is locally unstable. Since this analysis was performed using a reduced-form model, more structural analysis is necessary before coming to clear policy implications, but we nevertheless believe that these insights are highly suggestive of why the presence of local instability may matter for policy.

Appendix

Motivation for Baseline Specification

Consider the following demand-driven macroeconomic environment with two aggregate variables: per capita output denoted y_t, and per capita capital denoted k_t. These variables are the result of individual-level decisions by agents $i = 1$ to n, with n large. The capital stock of each agent i accumulates according to a standard accumulation equation, with depreciation rate δ, where we assume that a fixed fraction γ of agent i's purchases today, denoted y_{it}, contributes to his capital stock tomorrow as follows:

$$k_{it+1} = (1 - \delta)k_{it} + \gamma y_{it}, \qquad 0 < \gamma < 1 - \delta.$$

For ease of discussion, we want to think of agent i's capital stock as his holding of durable goods (including possibly housing) for which he receives utility directly. Output in this economy is assumed to be demand-determined, so that per capita output is simply the average of the individual-level purchases, that is, $y_t = n^{-1}\Sigma_i\, y_{it}$. This implies that the per capita capital stock, $k_t \equiv n^{-1}\Sigma_i\, k_{it}$, accumulates according to

$$k_{t+1} = (1 - \delta)k_t + \gamma y_t. \tag{A1}$$

To begin, consider the case where individual output demands are allowed to be affected by only three forces as follows:

$$y_{it} = \alpha_0 - \alpha_1 k_{it} + \alpha_2 y_{it-1} + \varepsilon_{it}. \tag{A2}$$

In equation (A2), output demand is allowed to be affected by the agent's holding of consumer capital. In particular, let us focus on the case where $\alpha_1 > 0$, which implies that more consumer capital leads to less desire for new purchases, as would be the case, for example, in the presence of diminishing marginal utility. Output demand is also allowed to exhibit inertia through $\alpha_2 \geq 0$, where this force may reflect, for example, habit or adjustment costs. Finally, we allow a random idiosyncratic term ε_{it}, which may be correlated across agents but is i.i.d. across time. A decision rule of the form given in equation (A2) can be derived from individual-level optimization, in which case the coefficients α_1 and α_2 will typically be smaller than 1. Aggregating equation (A2), we may obtain

$$y_t = \alpha_0 - \alpha_1 k_t + \alpha_2 y_{t-1} + \varepsilon_t, \tag{A3}$$

where $\varepsilon_t \equiv n^{-1}\sum_i \varepsilon_{it}$. Under the assumption that α_1 and α_2 are smaller than 1, equations (A1) and (A3) form a simple stable linear dynamic system. It is interesting to recognize that in this case if α_2 and $1 - \delta$ are close to 1, then the roots of this system can be very close to 1, while nonetheless staying stable. We take equations (A1) and (A3) as a simple representation of basic forces that could be driving macroeconomic fluctuations.

Now let us slightly extend the above system by introducing an additional term in the determination of individual-level demand. Here we want to add the possibility that agents' expectations about other agents' decisions may affect their own behavior. Such interaction effects can arise for many reasons. It is not our goal to take a stand on one particular reason here, but instead to argue that such interactions can drastically change the dynamics of the system, even if the effects remain small and close to linear. To capture such an effect, let us generalize the individual decision rule as follows:

$$y_{it} = \alpha_0 - \alpha_1 k_{it} + \alpha_2 y_{it-1} + G(y_{it}^e) + \varepsilon_{it}, \tag{A4}$$

where y_{it}^e is agent i's expectation about aggregate economic activity, and the function $G(\cdot)$ captures how this expectation affects his behavior. If G' is positive, agents' individual-level demands play the role of strategic complements, while if it is negative they play the role of strategic substitutes. Let us start with the extreme case where all agents base their expectations only on the past realization of y_t, and that this expectation is simply backward looking with $y_{it}^e = y_{t-1}$. In this case, the aggregate determination of output is given by

$$y_t = \alpha_0 - \alpha_1 k_t + \alpha_2 y_{t-1} + G(y_{t-1}) + \varepsilon_t. \tag{A5}$$

How does the introduction of the strategic interaction term $G(y_{t-1})$ affect the dynamic system? This depends on whether, near the steady state of the system, actions are strategic substitutes or complements. If they are strategic substitutes, then the system will tend to maintain stability and nothing very interesting is likely to happen. However, if they are strategic complements, then the dynamics may change considerably. In particular, consider the case where α_2 and $1 - \delta$ are close to 1. Then one may show that even a small degree of strategic complementarity near the steady state can render the system locally unstable.[36] In such a case, the global properties will depend on the nature of the nonlinearities in $G(\cdot)$. If the complementarities grow as the system moves away from its steady state, then it will tend to exhibit global instability, in which case the system will explode. In contrast, if the complementarities die out as the system moves away from the steady state, then the system will tend to produce either a limit cycle or chaotic behavior. Again, all this can arise even if the system remains close to linear and complementarities are rather weak. For example, suppose that G takes the form $G(y) = \mu_0 + \mu_1 y + \mu_2 y^3$. Then equation (A5) can be rewritten in univariate form using equation (A1) to eliminate k_t. Moreover, if we close the model by assuming that $y_t = h_t$, where h_t is hours worked per capita, then we can write the determination of hours as a univariate process in h_t as

$$h_t = a_0 + \mu_0 - a_1\gamma\sum_{j=1}^{\infty}(1-\delta)^{j-1}h_{t-j} + (a_2 + \mu_1)h_{t-1} + \mu_2 h_{t-1}^3 + \varepsilon_t. \tag{A6}$$

Equation (A6) is a special case of the type of univariate time-series model we estimate in section II. We emphasize here the possibility of interpreting this model as one determining hours worked as opposed to determining output, as the formulation in terms of hours worked can be shown to be more robust allowing for fluctuations in productivity (see section II.F). Note also that equation (A6) features the presence of an accumulation term, which motivates our inclusion of such a term in our econometric specification.

Data

- US Population: Total Population: all ages including Armed Forces Overseas, obtained from the FRED database (POP) from 1952:Q1 to 2015:Q2. Quarters from 1947:Q1 to 1952:Q1 are obtained from linear interpolation of the annual series of National Population obtained from US Census, where the levels have been adjusted so that the two series match in 1952:Q1.
- US total output, consumption, and the various investment types are obtained from the Bureau of Economic Analysis National Income and Product Accounts. Real quantities are computed as nominal quantities (table 1.1.5) over prices (table 1.1.4.). Sample is 1947:Q1–2015:Q2, and we do not use the observations of 2015.
- US nonfarm business output, nonfarm business hours, total hours and unemployment rate (16 years and over) are obtained from the Bureau of Labor Statistics. Sample is 1947:Q1–2015:Q2 (1948:Q1–2015:Q2 for total hours), and we do not use the observations of 2015.
- Capacity utilization: manufacturing (SIC), percent of capacity, quarterly, seasonally adjusted, obtained from the FRED database, (CUMFNS). Sample is 1948:Q1–2015Q3 and we do not use the observations of 2015.

- The series of job-finding rate was constructed by Robert Shimer. For additional details, please see Shimer (2012). The data from June 1967 and December 1975 were tabulated by Joe Ritter and made available by Hoyt Bleakley. Sample is 1948:Q1–2007:Q4.
- The various unemployment rates are obtained from the FRED database, except for France:
  - Australia: Unemployment Rate: ages 15 and over: all persons for Australia (LRUNTTTTAUQ156S), 1966:Q3–2014:Q4.
  - Austria: Registered unemployment rate for Austria (LMUNRRTTATQ156S), 1955:Q1–2014:Q4.
  - Canada: Harmonized unemployment rate: all persons for Canada (CANURHARMQDSMEI), 1956:Q1–2012:Q1.
  - Denmark: Registered unemployment rate for Denmark (LMUNRRTTDKQ156S), 1970:Q1–2014:Q4.
  - France: ILO unemployment rate, total, metropolitan France and overseas departments (001688527), Inséé Macro-economic database (BDM), 1975:Q1–2015:Q1.
  - Germany: Registered unemployment rate for Germany (LMUNRRTTDEQ156S), 1969:Q1–2014:Q1.
  - Japan: Harmonized unemployment rate: all persons for Japan (JPNURHARMQDSMEI), 1955:Q1–2012:Q1.
  - Netherlands: Harmonized unemployment rate: all persons for Netherlands (NLDURHARMQDSMEI until 2012:Q1 and harmonized unemployment: total: all persons for the Netherlands LRHUTTTTNLQ156S) after 2012:Q1 (the second-series level has been adjusted to match the first one in 2012:Q1), 1970:Q1–2014:Q4.
  - Sweden: Harmonized unemployment rate: all persons for Sweden (SWEURHARMQDSMEI), 1970:Q1–2012:Q1.
  - Switzerland: Registered unemployment rate for Switzerland (LMUNRRTTCHQ156S), 1970:Q1–2014:Q4.
  - United Kingdom: Registered unemployment rate for the United Kingdom (LMUNRRTTGBQ156S), 1956:Q1–2012:Q1.

Proof of Proposition 1

Assume that σ_ε^2 is small, so that $|x_0|$ is close to $\bar{x}$. Without loss of generality, take x_0 close to $\bar{x}$,[37] and define

$$y_t \equiv (-1)^t x_t - \bar{x}.$$

To understand y_t, note that, since x_0 is close to $\bar{x}$ and σ_ε^2 is small, we should have $x_t > 0$ for t even and $x_t < 0$ for t odd, that is, the sign of x_t

should switch every period.[38] Thus, $|y_t|$ captures the absolute deviation of x_t from the nonstochastic sequence $\bar{x}_t \equiv (-1)^t\bar{x}$ (the two-cycle), while $\text{sgn}(y_t)$ captures the direction of that deviation relative to zero: if $y_t < 0$ then x_t is "inside" the two-cycle (i.e., $|x_t| < \bar{x}$), while if $y_t > 0$ then x_t is "outside" the two-cycle (i.e., $|x_t| > \bar{x}$).

Next, it is straightforward to verify that

$$y_t = (3 - 2\beta_x)y_{t-1} - 3\sqrt{(\beta_x + 1)(\beta_x - 1)}y_{t-1}^2 - (\beta_x + 1)y_{t-1}^3 + \nu_t$$

where $\nu_t \equiv (-1)^t\varepsilon_t \sim \text{i.i.d.}N(0, \sigma_\varepsilon^2)$. Since $(3 - 2\beta_x) \in (-1, 1)$, the fixed point of this system at $y_t = 0$ is stable, and thus for σ_ε^2 small we may restrict attention to the second-order approximation to this system,

$$y_t \approx (3 - 2\beta_x)y_{t-1} - 3\sqrt{(\beta_x + 1)(\beta_x - 1)}y_{t-1}^2 + \nu_t \tag{A7}$$

y_t is clearly stationary, so that we have

$$\begin{aligned}\mathbb{E}[y_t] &\approx (3 - 2\beta_x)\mathbb{E}[y_t] - 3\sqrt{(\beta_x + 1)(\beta_x - 1)}\mathbb{E}[y_t^2] \\ &= -\frac{3}{2\bar{x}}\,\mathbb{E}[y_t^2].\end{aligned}$$

For $\sigma_\varepsilon^2 > 0$, we will have $\mathbb{E}[y_t^2] > 0$, and thus $\mathbb{E}[y_t] < 0$. Since $x_t = (-1)^t(y_t + \bar{x})$, we may obtain

$$\begin{aligned}\sigma_x^2 &= \mathbb{E}[y_t^2] + 2\bar{x}\mathbb{E}[y_t] + \bar{x}^2 \\ &\approx \bar{x}^2 - 2\mathbb{E}[y_t^2]\end{aligned}$$

where we have used the approximation from above to replace $\mathbb{E}[y_t]$. Thus, as σ_ε^2 increases from zero, x_t becomes less volatile, which completes the proof.

Endnotes

The authors thank the two discussants Laura Veldkamp and Ivan Werning, as well as Jonathan Parker for useful comments. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see http://www.nber.org/chapters/c13772.ack.

1. Although the idea of endogenous boom-and-bust cycles has a long tradition in the economics literature (Kalecki 1937; Kaldor 1940; Hicks 1950; Goodwin 1951), it is not present in most modern macromodels. Recent exceptions are Myerson (2012), Matsuyama (2013), Gu et al. (2013), and Beaudry, Galizia, and Portier (2015b).

2. In the 1980s, following Blanchard and Summers (1986), a branch of the literature explored the case of an exact unit root in unemployment series as an indication of hys-

teresis. However, in practice, estimations gave persistent but stable dynamics, as pointed out by Blanchard and Summers (1986, p. 17, fn. 1): "We shall instead use 'hysteresis' more loosely to refer to the case where the degree of dependence on the past is very high, where the sum of coefficients is close but not necessarily equal to 1."

3. A reader mainly interested in the theoretical rationale for local instability in macroeconomic models may want to consult Beaudry, Galizia, and Portier (2015b).

4. We discuss in section II.D the possibility of multiple steady states and its implications.

5. A deterministic dynamic system is considered sensitive to initial conditions when arbitrarily small differences in initial conditions can lead to significant differences in outcomes over time. In such a case, forecasting long-run outcomes is considered problematic.

6. For ease of exposition, we consider here processes without a constant term, which amounts to assuming that the steady state of the process is at zero. In our empirical exercise below, we relax this assumption.

7. If $\tilde{A}(L) = a_1 + a_2L + \cdots + a_nL^{n-1}$, then we define the associated $n \times n$ companion matrix as follows: the first row is given by $[a_1, a_2, \ldots, a_n]$, the lower-left $(n-1) \times (n-1)$ block is given by I_{n-1}, and the remaining entries are zero.

8. Throughout this paper, for the sake of brevity the term "largest eigenvalue" will be used to refer to the eigenvalue that has the largest modulus, and when we say that an eigenvalue is greater (less) than 1, unless otherwise indicated we mean that its *modulus* is greater (less) than 1.

9. Fundamentally, if the largest eigenvalue of A is greater than 1, then the system will tend to move away from the steady state (even if the variance of ε is very small), which causes higher-order terms to become more relevant in the DGP. It is for precisely this reason that using an econometric specification containing only first-order terms is more likely to yield misleading results in such a case than in the case where the largest eigenvalue of A is less than 1.

10. In particular, if $\beta > (\mathbb{E}[x_{t-1}^2] \,/\, \mathbb{E}[x_{t-1}^4])(\alpha - 1)$.

11. We assume for simplicity that n is large enough that agents ignore the effect of their own choice x_i on x.

12. Since we will be examining behavior within a limited class of nonlinear models, one must be careful about which inferences can be made. For example, if we do not find evidence of local instability in this class, this does not imply that the macroeconomic environment is necessarily stable. However, if we do find that local instability and limit cycles appear in this limited class, it will provide some reason to question the consensus view that macroeconomic fluctuations reflect mainly the effects of shocks within an otherwise stable system.

13. Even if the cyclical variables we consider generally have a zero mean, we nevertheless allow for a nonzero intercept in our baseline specifications. Results are robust to omitting the intercept.

14. Note that the F in equation (4) and the F in equation (5) are not technically the same functions.

15. This identification assumption is quite strong, but allows for a simple two-step approach to the data: first detrend, then estimate. We leave for future research the exploration of decomposition methods that do not require this assumption.

16. We discuss the robustness of our result to the filter below. The quarterly series is filtered (80-quarter, high-pass filter) over the sample 1948:Q1–2015:Q2.

17. Note that the coefficient on h_{t-1} in this approximation is in fact different from the one reported above, though the two are equal to two decimal places.

18. Note that the maximum eigenvalue is not the autocorrelation coefficient because the process is an $AR(2)$ not an $AR(1)$.

19. Note that, except under very special circumstances that are unlikely to be encountered in our case, existence of and convergence to a limit cycle cannot be proven mathematically. As a result, throughout this paper we rely on numerical simulations to check for these properties.

20. One could repeat the deterministic simulation starting from all dates of the sample. This would make the convergence to a limit cycle unchanged or more frequent.

21. We checked that this steady state was unique in each of our bootstrap replications.
22. This does not mean that microfounded models capable of generating limit cycles never feature indeterminacy.
23. This analysis can be extended to the intermediate model when $\beta_0 = 0$. Nonzero steady states, if they exist, are the two opposite real numbers $\pm\bar{h}$ with

$$\bar{h} = \sqrt{\frac{1 - \beta_h - \beta_{h'} - \beta_H}{\beta_{h^3} + \beta_{h'^3} + \beta_{H^3} + \beta_{h^2 H} + \beta_{hH^2}}}.$$

24. One could repeat the deterministic simulation starting from all dates of the sample. This would make the convergence to a limit cycle unchanged or more frequent.
25. When we set the threshold level to 1% in the Wald tests (instead of the estimated p-value), the test fails to reject the (spurious) existence of a limit cycle 0.9% (resp. 1.5%) of the time for the minimal model when the DGP is the $AR(2)$ model (resp. the linear model), and 0.7% (resp. 1.3%) and 0.6% (reps. 8%) for the intermediate and full models.
26. See appendix for details about the source and time span of each of the series.
27. A similar pattern is found using the intermediate and full models (not shown in the figure).
28. In this section we will be discussing the effects of changes in shock variances on the variances of endogenous variables. Accordingly, results from this section need to be interpreted cautiously, given the potential for the Lucas Critique to apply.
29. That is, the roots of the polynomial $1 - F(\lambda, \lambda^2, \ldots)$ lie outside the unit circle.
30. With some abuse of terminology, we define the mean and variance of x by time averages rather than ensemble averages, that is, by $\mu_x \equiv plim_{T\to\infty}(1 / T)\sum_{t=1}^{T} x_t$ and $\sigma_x^2 \equiv plim_{T\to\infty}(1 / T)\sum_{t=1}^{T}(x_t - \mu_x)^2$, respectively, assuming these probability limits exist and are independent of the initial state x_0. Thus, according to our definition, we will have $\sigma_x^2 > 0$ even in a nonstochastic environment as long as the system does not converge to a single point.
31. See May (1979).
32. Estimated parameters are given in table 1.
33. In our previous work (Beaudry et al. 2015b) we presented a structural model where limit cycles arose as the result of a complementarity in consumption decisions generated by imperfect unemployment insurance. In that framework, the policy prescription to help stabilize the economy would therefore be to reduce the complementarity in agents' decisions by improving consumption insurance through, for example, automatic stabilizers.
34. More precisely, a reduction in input volatility may have very little impact on the overall volatility of the system when the deterministic version of the system exhibits limit cycles.
35. For most cases we found the modulus of the largest eigenvalue of the estimated system to be very close to 1.
36. See Beaudry et al. (2015b) for details. One of the insights from this example is to show that the presence of strategic complementarities in demand is likely to create local instability.
37. It is straightforward to extend the following reasoning to the case where x_0 is close to $-\bar{x}$.
38. This stems from the fact that, with x_0 close to $\bar{x}$ and σ_ε^2 small, the dynamic behavior of x_t is dominated by the same forces that generate the two-cycle.

References

Beaudry, P., D. Galizia, and F. Portier. 2015a. "Reconciling Hayek's and Keynes' Views of Recessions." NBER working paper w20101, 2014, forthcoming in the *Review of Economic Studies*.

———. 2015b. "Reviving the Limit Cycle View of Macroeconomic Fluctuations." NBER working paper w21241, 2015.

Blanchard, O. J., and L. H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual 1986*, vol. 1, ed. S. Fischer, 15–90. Chicago: University of Chicago Press.

Goodwin, R. 1951. "The Nonlinear Accelerator and the Persistence of Business Cycles." *Econometrica* 19 (1): 1–17.

Gu, C., F. Mattesini, C. Monnet, and R. Wright. 2013. "Endogenous Credit Cycles." *Journal of Political Economy* 121 (5): 940–65.

Hicks, J. 1950. *A Contribution to the Theory of the Trade Cycle*. Oxford: Clarendon Press.

Kaldor, N. 1940. "A Model of the Trade Cycle." *Economic Journal* 50 (197): 78–92.

Kalecki, M. 1937. "A Theory of the Business Cycle." *Review of Economic Studies* 4 (2): 77–97.

Matsuyama, K. 2013. "The Good, the Bad, and the Ugly: An Inquiry into the Causes and Nature of Credit Cycles." *Theoretical Economics* 8 (3): 623–51.

May, R. M. 1979. "Bifurcations and Dynamic Complexity in Ecological Systems." *Annals of the New York Academy of Sciences* 357:267–81.

Myerson, R. B. 2012. "A Model of Moral-Hazard Credit Cycles." *Journal of Political Economy* 120 (5): 847–78.

Shimer, R. 2012. "Reassessing the Ins and Outs of Unemployment. *Review of Economic Dynamics* 15 (2): 127–48.

Tibshirani, R. 1996. "Regression Shrinkage and Selection via the Lasso." *Journal of the Royal Statistical Society, Series B* 58 (1): 267–88.