Is the Macroeconomy Locally Unstable and Why Should We Care?

Paul Beaudry[*], Dana Galizia[†] and Franck Portier[‡]

January 2016
Incomplete - Please do not circulate without permission

Abstract

In most modern macroeconomic models, the steady state (or balanced growth path) of the system is a local attractor in the sense that in the absence of shocks, the economy would converge to the steady state. In this paper we examine whether the times series behavior of macroeconomic aggregates (especially labor market aggregates) are in fact supportive of this locally stability view of macroeconomic dynamics or if they instead favor an alternative interpretation whereby the macro-economy may be better characterized as a locally unstable system with deterministic forces that support limit cycle behavior. To do this, we extend standard AR representations of the data to allow for smooth non-linearities using polynomials. Our main finding is that, even using a procedure which may have low power, the data provide considerable support to the view that the macro-economy may be locally unstable. An interesting finding is that the amount of non-linearity we detect in the data is rather minor but it is enough to alter the description of macroeconomic behavior. We also discuss the extent to which these two different views about the inherent dynamics of the macro-economy may matter for policy.

1 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 towards a smooth growth path. According to this belief, most of the fluctuations in the macro-economy result from either individually optimal adjustments to changes in the environment or due to improper government interventions. In such a case, the role of macroeconomic policy should be to do no

[*]Vancouver School of Economics, University of British Columbia and NBER
[†]Department of Economics,Carleton University
[‡]Toulouse School of Economics, Université Toulouse 1 Capitole and CEPR

harm: if policy makers hold back from actively influencing the economy, market forces would take care of the rest and foster desirable outcomes. 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 a 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 *stable-with-shocks* framework is very tractable and flexible, the ubiquitous and recurrent feature of cycles in most market economies incites one to ask whether the market economy, by its very nature, may inherently create recurrent boom and bust, with busts sowing the seed of the next boom.[1]

Within a linear set-up, a dynamic system is either stable or unstable (where we interpret saddle-path stability as stability). In contrast, when departing from the strict linear setup, a system can be globally stable (i.e., starting from any initial conditions, it does not explode) while simultaneously being locally unstable (i.e., even in the absence of shocks, it does not converge to a fixed point). It is this configuration that we believe is of potential of interest to macroeconomics. The goal of this paper is to highlight why macro-economists should be aware of and interested in this possibility. Accordingly, we propose to explore two issues. First, we will examine the question: Is there simple evidence supportive of the notion that the macro-economy may be globally stable but locally unstable? We address this issue here from a reduced-form perspective, as we believe that such evidence would be particularly compelling to many readers, and more

[1] Although the idea of endogenous boom-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 macro-models. Recent exceptions are Myerson [2012], Matsuyama [2013], Gu, Mattesini, Monnet, and Wright [2013] and Beaudry, Galizia, and Portier [2015b].

likely to favor future research, than evidence that is tied to a specific structural model. Second, we discuss how the possibility of local instability affects the tradeoffs associated with policy.

2 Time Series Evidence with a Reduced-Form Model

Our goal in this section is to examine the time series behavior of several macroeconomic aggregates to see whether they display features of local instability and possible limit cycle forces. Recall that if a dynamic system with a unique steady state is locally unstable but not explosive, then it will either converge to a limit cycle or it will exhibit chaos. In principal the exercise is straightforward. We will begin by estimating univariate processes allowing for non-linear forces. We will then examine whether the implied dynamics are stable or explosive near the steady state. Finally, we examine the behavior of the variable if we start it away from the steady state, while shutting down the effects of the stochastic term. In particular, if we find that a variable exhibits local instability, we want to examine whether it then converges to cyclical behavior. The main difficulty with this procedure is determining the class of non-linear models to consider. While we want to adopt a reduced-form approach, we draw on some of our past theoretical work to help us choose a class of non-linear models to explore. Since we will be examining behavior within a limited class of non-linear 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.

Before providing a rationale for the particular non-linear framework we chose, let us first present it. Consider a stationary measure of economic activity, which could be the unemployment rate (possibly detrended) or the number of hours worked per capita, and denote this variable by x_t. As is well known, the time-series properties of such variables are often well captured by a linear $AR(2)$ model or a simple ARMA model as given below.

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

where $A(L)$ and $B(L)$ are polynomials in the lag operator and where ϵ_t is an iid process,. Now consider the following generalization

$$x_t = A(L)x_{t-1} + F(x_{t-1}, x_{t-1}, \dots) + B(L)\epsilon_t, \tag{2}$$

where $F(\cdot)$ is some non-linear function which we will restrict to be a polynomial function. Equation (2) obviously embeds a simple $AR(2)$ for example, but it allows for many possible non-linearities. To make our analysis tractable, we focus primarily on estimating a restricted version of Equation (2), as given by

$$x_t = a_1 x_{t-1} + a_2 x_{t-2} + F\left(x_{t-1}, x_{t-2}, \delta \sum_{i=0}^{\infty} (1-\delta)^i x_{t-1-i}\right) + \epsilon_t, \tag{3}$$

where the polynomial function $F(\cdot)$ is restricted to be at most of order 3. [2] In particular, by letting $F(\cdot)$ be a polynomial of order 3 we are allowing for a rich set of non-linearities, and by allowing $\delta \sum_{i=0}^{\infty}(1-\delta)^i x_{t-1-i}$ to be an explanatory variable we are allowing a longer memory type term to play a role. In subsection 2.1 below we explain what led us to choose this particular class of models as a baseline to explore the idea that the macro-economy could be locally unstable and exhibit limit cycles. This subsection can be skipped if the reader is satisfied with exploring equation (3) without further motivation.

2.1 Motivation behind our Reduced-Form Model

In Beaudry, Galizia, and Portier [2015a] and [2015b], we showed that local instability and the emergence of limit cycles can easily arise in dynamic economic systems with strategic complementarities if individual-level decisions have the following properties: (i) decreasing returns to accumulation, so that when the stock of an accumulated factor is high, there is less incentive to accumulate it further; (ii) sluggishness, meaning that people tend to change their behavior slowly; and (iii) strategic complementarities in expenditures decisions, that is, if everyone one else is spending, this tends to increase one's incentive to spend.

[2] As shown in Appendix A.2, we obtain similar results when we allow for a moving average error term.

To fix ideas, consider the simple reduced-form model presented in Beaudry, Galizia, and Portier [2015b]. All variables are in logs. There is one produced good in the economy given by y_t. The flow production of this good is given by $y_t = \theta_t + h_t$, where h is the level of aggregate hours worked and θ a productivity index that follows a random walk, so that $\theta_t = \theta_{t-1} + \eta_t$. When an individual purchases y_t this gives him current services, and a fraction of these purchases also enter the household's capital stock, denoted by X_{jt} (which could represent durable goods, housing or both), such that the stock X_{jt} held by agent j accumulates according to

$$X_{jt} = (1-\delta)X_{jt-1} + \delta y_{jt}, \tag{4}$$

with $\delta \in (0,1)$ and where y_{jt} is the amount of expenditure made. Equation (4) should be thought as a log-linearization of a standard accumulation equation.

Suppose now that the expenditure decision for individual j, y_{jt} takes the following form:

$$\widetilde{y}_{jt} = \alpha_1 \widetilde{y}_{jt-1} - \alpha_2 \widetilde{X}_{jt-1} + E_{t-1} F\left(h_t\right), \tag{5}$$

where $\widetilde{y}_{jt} = y_{jt} - \theta_t$ and $\widetilde{X}_{jt} = X_{jt} - \theta_t$ are technology-adjusted variables with $\alpha_1, \alpha_2 \in (0,1)$. In equation (5), α_1 captures the sluggishness aspect, while α_2 governs the effect of decreasing returns to accumulation. Complementarity in decisions is captured by the term $E_{t-1}F\left(h_t\right)$, which we allow to be non-linear.[3], then individual households will spend less as they fear to be unemployed. The model is then closed by having aggregate production equal to aggregate expenditures. When written in logs around symmetrical allocations, this implies

$$\frac{1}{N}\sum_{j=1}^{N} y_{jt} = y_t. \tag{6}$$

This model can then be rewritten in terms only of hours worked. To do so, let's define

$$H_{jt} = (1-\delta)H_{jt-1} + \delta h_{jt},$$

[3] To keep a simple structure for the model's solution, we have assumed that agents must form one-period-ahead expectations of aggregate expenditures. See Appendix A.3 for similar results for the case when agents form expectations of aggregate expenditures contemporaneously.

and rewrite (4) as

$$
\begin{aligned}
X_t &= \delta \sum_{i=0}^{\infty} (1-\delta)^i) y_{t-i} \\
&= \delta \sum_{i=0}^{\infty} (1-\delta)^i h_{t-i} + \delta \sum_{i=0}^{\infty} \left[(1-\delta)L\right]^i \theta_t \\
&= H_t + \frac{\delta}{1-(1-\delta)L} \theta_t .
\end{aligned}
$$

where L is the lag operator. Plugging into (5), we obtain

$$h_t = \alpha_1 h_{t-1} - \alpha_2 H_{t-1} + E_{t-1} F(h_t) + \frac{1}{1-(1-\delta)L} \varepsilon_t, \tag{7}$$

where $\varepsilon_t \equiv \alpha_2 (1-\delta) \eta_{t-1}$, which can then be written

$$h_t = A(L) h_{t-1} + B(L) H_{t-1} + C(L) E_{t-1} F\left(h_t\right) + \varepsilon_t. \tag{8}$$

In Beaudry, Galizia, and Portier [2015b], we studied the perfect foresight version of (7)

$$h_t = \alpha_1 h_{t-1} - \alpha_2 H_{t-1} + F\left(h_t\right), \tag{9}$$

$$H_t = (1-\delta) H_{t-1} + \delta h_t, \tag{10}$$

with $1 > \alpha_1 > 0$, $1 > \alpha_2 > 0$, and $F(0) = 0$. We also assume that $F'(h) < 1 \; \forall h \in \mathbb{R}$ so that complementarities are insufficiently strong to create static multiple equilibria. We showed for this deterministic system that (i) $\overline{h} = 0$ is the unique steady state; (ii) the steady state is globally stable when there is no strategic interaction (i.e., when $F \equiv 0$); (iii) the steady state is locally stable when there is strategic substitutability at the steady state (i.e., when $F' < 0$);[4] (iv) the steady state becomes locally unstable if strategic complementarities become sufficiently strong, and a Hopf bifurcation occurs at this point as long as α_1 is not too small (that is, the combination of complementarities and sluggish behavior can produce a limit cycle in the neighborhood of the steady state when we maintain the assumption that $F' < 1$); and (v) if the F function is S-shaped,[5] then the limit cycle is locally attractive.

These results suggested to us that the estimation of a specification like (9) may be a good starting point to examine in a new light the behavior of macroeconomic

[4] More formally, we consider here a class of functions $F(h; \lambda)$ parametrized by the real number $\lambda \in (-\infty, 1)$. Those function are such that $\frac{\partial F(0;\lambda)}{\partial h} = \lambda$, $\frac{\partial F(h;\lambda)}{\partial h} < 1 \; \forall \lambda, h$ and $F(0; \lambda) = 0 \; \forall \lambda$. We then compare the properties of the dynamical system for various values of $F'(0)$, which indeed corresponds to varying the parameter λ.

[5] More precisely, if F''' is negative enough.

aggregates, since it is a simple framework that could capture local instability if it is present. We could build directly off (9), but this would not embed as a special case a simple $AR(2)$ model,which is most often the favored auto-regressive specification for our variables of interest. For this reason, we chose to add h_{t-2} to (9), so that our baseline reduced-form equation is

$$h_t = \alpha_1 h_{t-1} + \alpha_2 h_{t-2} + \alpha_3 H_{t-1} + E_{t-1} F\left(h_t\right) + \varepsilon_t, \tag{11}$$

together with the definition of H given in (10). We now need to be more explicit about the term $E_{t-1}F\left(h_t\right)$. Given the model, we can express this conditional expectation as a non-linear function Φ of the model history, which is summarized by the triplet $(h_{t-1}, h_{t-2}, H_{t-1})$:

$$E_{t-1}F\left(h_t\right) = \Phi(h_{t-1}, h_{t-2}, H_{t-1})$$

When we replace the expectation in (11), we now get an equation of the form given by equation (3). The final step is to discuss the form of $\Phi(\cdot)$ or $F(\cdot)$. From our previous work, we know that limit cycles will occur and be stable when the F function is S-shaped, so that we need $\Phi(\cdot)$ or $F(\cdot)$ to be at least a polynomial of order three. As an order-three polynomial already involves twenty terms, we also take this to be the maximum order we consider.

2.2 Three embedded models

If we allow the function $F(\cdot)$ in (3) to be a third-order polynomial, we get the following expression, which we refer to as the "Full" model:

$$\begin{aligned}
h_t = {} & \beta_h \ h_{t-1} + \beta_{h'} \ h_{t-2} + \beta_H \ H_{t-1} \\
& + \beta_{h^2} \ h_{t-1}^2 + \beta_{h'^2} \ h_{t-2}^2 + \beta_{H^2} \ H_{t-1}^2 \\
& + \beta_{hh'} \ h_{t-1} h_{t-2} + \beta_{hH} \ h_{t-1} H_{t-1} + \beta_{h'H} \ h_{t-2} H_{t-1} \\
& + \beta_{h^3} \ h_{t-1}^3 + \beta_{h'^3} \ h_{t-2}^3 + \beta_{H^3} \ H_{t-1}^3 \\
& + \beta_{h^2h'} \ h_{t-1}^2 h_{t-2} + \beta_{h^2H} \ h_{t-1}^2 H_{t-1} \\
& + \beta_{h'^2H} \ h_{t-2}^2 H_{t-1} + \beta_{hh'^2} \ h_{t-1} h_{t-2}^2 + \beta_{hH^2} \ h_{t-1} H_{t-1}^2 + \beta_{h'H^2} \ h_{t-2} H_{t-1}^2 \\
& + \beta_{hh'H} \ h_{t-1} h_{t-2} H_{t-1} + \epsilon_t.
\end{aligned} \tag{12}$$

The attractive feature of this "Full" model is that it allows for a rich non-linearity departure from a simple $AR(2)$ model. However, the drawback of such a specification is that it involves twenty parameters, which is quite demanding for standard macroeconomic time series. For this reason, we will consider at the other extreme a very parsimonious version of (3) which consists of the inclusion of only one non-linear term given by the cubic term in h_{t-1}. Our "Minimal" model is therefore given by

$$\begin{aligned} h_t = \beta_h \ h_{t-1} + \beta_{h'} \ h_{t-2} + \beta_H \ H_{t-1} \\ + \beta_{h^3} h_{t-1}^3 + \epsilon_t. \end{aligned} \tag{13}$$

Finally, we will also consider an "Intermediate" model which 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 choose as our intermediate model the case where we keep most third-order terms from the full model but eliminate the second-order terms. This has the feature of significantly reducing the number of parameters, while still allowing for non-linearities capable of supporting limit cycles if they are present. In particular, from the seven third-order terms, we keep the three cubic terms and the two cross-terms involving h_{t-1} and H_{t-1}, so that our Intermediate model model becomes

$$\begin{aligned} h_t &= \beta_h \ h_{t-1} + \beta_{h'} \ h_{t-2} + \beta_H \ H_{t-1} \\ &+ \beta_{h^3} \ h_{t-1}^3 + \beta_{h'^3} \ h_{t-2}^3 + \beta_{H^3} \ H_{t-1}^3 \\ &+ \beta_{h^2 H} \ h_{t-1}^2 H_{t-1} + \beta_{hH^2} \ h_{t-1} H_{t-1}^2 + \epsilon_t. \end{aligned} \tag{14}$$

Our three models — the Full model, the Intermediate model and the Minimal model — together offer a tractable non-linear framework for examining whether certain macroeconomic variables may exhibit local instability and limit cycles. As a point of comparison, we will compare the behavior of these models with that of a simple $AR(2)$ model given by

$$h_t = \beta_h \ h_{t-1} + \beta_{h'} \ h_{t-2} + \epsilon_t. \tag{15}$$

2.3 Estimation with Total Hours

Data Treatment and Estimation Results

The motivation we used to derive our reduced-form model is one base on cyclical fluctuations in hours worked. Accordingly, the first measure of hours worked 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 Figure 1, Total Hours worked per capita have exhibited important low frequency movements over the last 50 years. For this reason, we begin by filtering this data in order to remove the low frequency movements. We do this by using a (linear) high-pass filter that retains only fluctuations of period less than 20 years.[6] The series, its trend and its filtered component are displayed in Figure 1. Once we have a series for h (filtered hours), we need to construct the accumulated series H_t, defined in (10). To do so, we assume for now that $\delta = .05$ and truncate the permanent inventory equation (10) at $N = 40$, so that, for example, observations of h from 1949Q2 to 1960Q1 are used to compute H_{1960Q1}.[7] Later we show that our results are robust to different values of δ.

Figure 1: Trend and Cycle (High-pass filter, 80 quarters), Total Hours

(a) Level and Trend (b) Cyclical Component

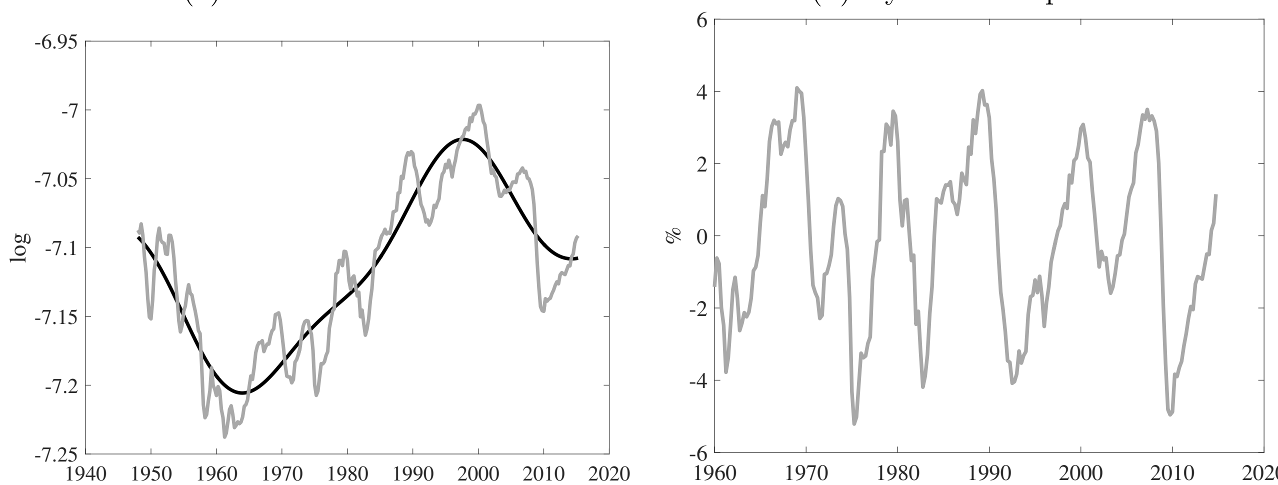

Notes: Total Hours are measured as the log of BLS total economy hours worked, deflated by total population. The trend of that series is obtained by removing from the level series its filtered component, where the filter is an 80-quarter high-pass filter, over the sample 1948Q1-2015Q2. The cyclical component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

Parameter estimates for each of our three models estimated on Total Hours worked are displayed in Table 1. One can check that in the models with non-linear terms

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

[7] That is, the series of observed H is constructed according to the equation $H_t = \delta \sum_{j=0}^{N-1} (1-\delta)^j h_{t-j}$.

there is always at least one significant 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 $\beta_{h'}$ 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 dhows that, at a 1% level of confidence, the Minimal model is not rejected against any of the non-linear alternatives. We first discuss the results obtained with this Minimal model since it is easiest to understand. We will show later that results are mostly the same with the Intermediate and Full models.

Local Instability and Limit Cycles

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

$$\begin{cases} h_t = 1.42\ h_{t-1} - .48\ h_{t-2}, \\ h_t = 1.39\ h_{t-1} - .34\ h_{t-2} - .27\ H_{t-1} - .01\ h_{t-1}^3. \end{cases}$$

First note that, in accordance with the ideas motivating our reduced-form model, we find that both H_{t-1} and h_{t-1}^3 enter negatively in the equation. Second, even though our Minimal model is non-linear, it is easy to check that it has only one steady state, which is at zero. In order to examine the local stability of each of these two equations, we compute the maximum eigenvalue (in modulus) of the system linearized around this steady state (and using equation (10) for the Minimal model). For the Minimal model, this linear approximation is simply

$$h_t = 1.39\ h_{t-1} - .34\ h_{t-2} - .27\ H_{t-1}.$$

Maximum eigenvalues are displayed in Table 3. For the $AR(2)$ model, the maximum eigenvalue is .86.[8] The steady state implied by the $AR(2)$ is therefore locally stable, which is not surprising. For the non-linear Minimal model, on the other hand, the maximum eigenvalue is 1.02, so that the steady state is locally unstable. This may seem to conflict with the visual impression given by the data (Figure 1, panel (b)) that the economy does not explode. In fact, it is precisely the third-order term h_{t-1}^3

[8] Note that the maximum eigenvalue is not the autocorrelation coefficient because the process is an $AR(2)$ and not an $AR(1)$.

Table 1: Estimates of Various Models, Total Hours

Variable	Value			
	$AR(2)$	Minimal	Intermediate	Full
h_{t-1}	1.42	1.39	1.28	1.48
	(23.80)	(19.75)	(12.73)	(10.00)
h^2_{t-1}				-0.12
				(-1.85)
h^3_{t-1}		-0.01	-3e-03	-0.09
		(-2.34)	(-0.76)	(-1.83)
h_{t-2}	-0.48	-0.34	-0.24	-0.43
	(-8.07)	(-2.46)	(-5.14)	(-2.69)
h^2_{t-2}				-0.09
				(-1.08)
h^3_{t-2}			-4e-03	0.11
			(-0.80)	(1.34)
H_{t-1}		-0.27	-0.02	0.19
		(-4.36)	(-0.11)	(0.91)
H^2_{t-1}				0.19
				(1.90)
H^3_{t-1}			-0.28	-0.37
			(-2.59)	(-1.97)
$h_{t-1}h_{t-2}$				0.20
				(1.44)
$h_{t-1}H_{t-1}$				0.23
				(1.77)
$h_{t-2}H_{t-1}$				-0.25
				(-1.73)
$h^2_{t-1}h_{t-2}$				0.27
				(1.54)
$h^2_{t-1}H_{t-1}$			5e-03	-0.31
			(0.34)	(-2.06)
$h^2_{t-2}H_{t-1}$				-0.29
				(-1.43)
$h_{t-1}h^2_{t-2}$				-0.47
				(-2.38)
$h_{t-1}H^2_{t-1}$			0.04	-0.37
			(0.93)	(-1.64)
$h_{t-2}H^2_{t-1}$				0.47
				(1.77)
$h_{t-1}h_{t-2}H_{t-1}$				0.76
				(2.24)

Notes: The Total Hours series has been filtered with an 80-quarter high-pass filter over the sample 1948Q1-2015Q2. Estimation is then done over the sample 1960Q1:2014Q4.

Table 2: Wald Test for the Different Models, Total Hours

H_0 \ H_1	Minimal	Intermediate	Full
$AR(2)$	0.002 %	0.004 %	0.003%
Minimal	-	6.96 %	1.32%
Intermediate	-	-	3.33%

Notes: The four models correspond to equations (15), (13), (14) and (12). 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.

Table 3: Some Statistics for the Different Models, Total Hours

	$AR(2)$	Minimal	Intermediate	Full
R^2	0.94	0.94	0.94	0.95
Adj. R^2	0.94	0.94	0.94	0.95
DW	2.22	2.12	2.09	2.09
Max eig. modulus	0.86	1.01	1.04	1.16

Notes: The four models correspond to equations (15), (13), (14) and (12). 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 0, which corresponds to considering only the linear terms of each model.

that prevents explosion (but has no effect on the eigenvalues), since it enters negatively in the estimated equation. The steady state is thus locally unstable, but as h moves aways from the steady state, the term in h^3 acts as a centripetal force that pushes the economy back towards the steady state if it strays too far. As a result, the economy will oscillate without explosion or convergence to a fixed point, instead eventually converging to a limit cycle. 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:

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

We forecast h over 1000 periods with the estimated model. We start (arbitrarily) from the first trough of the sample (i.e., 1961Q3) 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 2, the trajectory is quickly attracted to a limit cycle. Figure 3 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 2) 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 guided from our previous theoretical analysis, 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 Intermediate and Full models. Table 3 shows that the maximum eigenvalue modulus is also greater than one in these models. As shown in Figure 4, all three models generate limit cycles with very similar amplitudes for Total Hours h. Convergence to the limit cycle can only be formally proven in the Minimal model at the bifurcation, so that we rely on numerical simulations to check that the limit cycle is indeed attractive.

Figure 2: The Limit Cycle in the Minimal Model, Total Hours

(a) Full state space

(b) Projection onto (h_t, H_t)

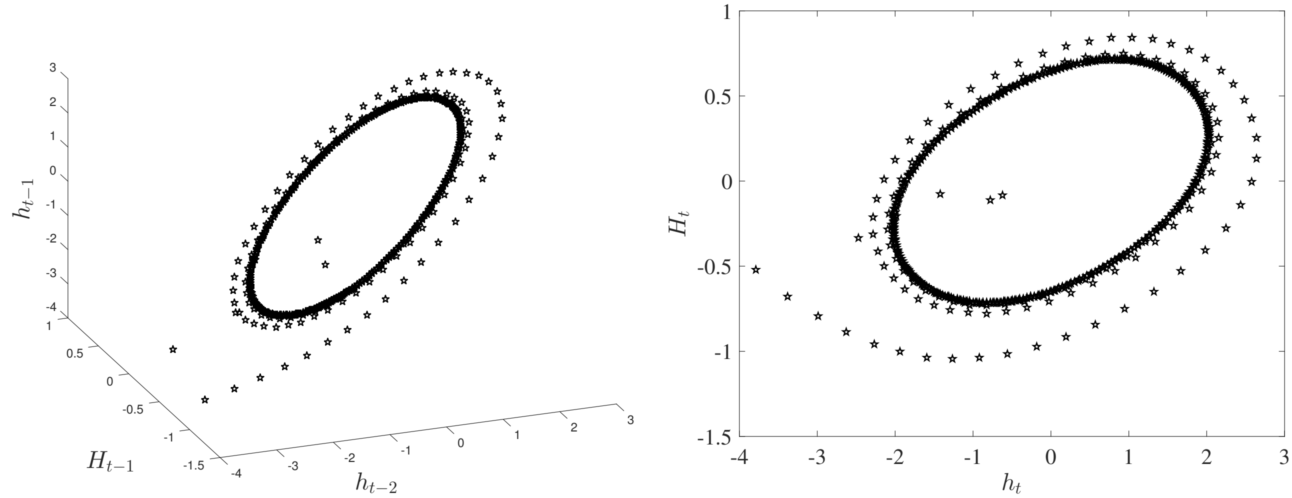

Notes: This corresponds to the deterministic simulation of the Minimal model (13)*, starting in 1961Q3. The model has been estimated over 1960-2014 using 80-quarter high-pass-filtered Total Hours.*

Forecasted Path

Let us now focus on the Minimal model andlook at the size and frequency of the limit cycle. As illustrated in panel (b) of Figure 4, the deterministic forecast for the Minimal model displays a cycle whose amplitude is close to the data one. In comparison, the forecasted path for the $AR(2)$ model quickly converges to the steady state. Panel (a) of Figure 5 displays 95% confidence bands around the deterministic forecast, and shows that after 60 quarters, the confidence bands have the same amplitude than the deterministic cycle.

The deterministic forecast path for our non-linear model might be considered misleading when compared to the actual series. In particular, the non-linear model we estimated does not imply that fluctuations are deterministic and therefore easily predictable, since shocks to the system cause cycles to be permanently dephased. As a result, unlike in a linear model, the deterministic forecast for our non-linear model in general differs from the mathematical expectation of the future state of the economy. To obtain the latter, we simulate the model (with shocks) forward a number of times, always starting from T_0, then take an average across these simulated paths. Each

Figure 3: Convergence to the Limit Cycle in the Minimal Model, Total Hours

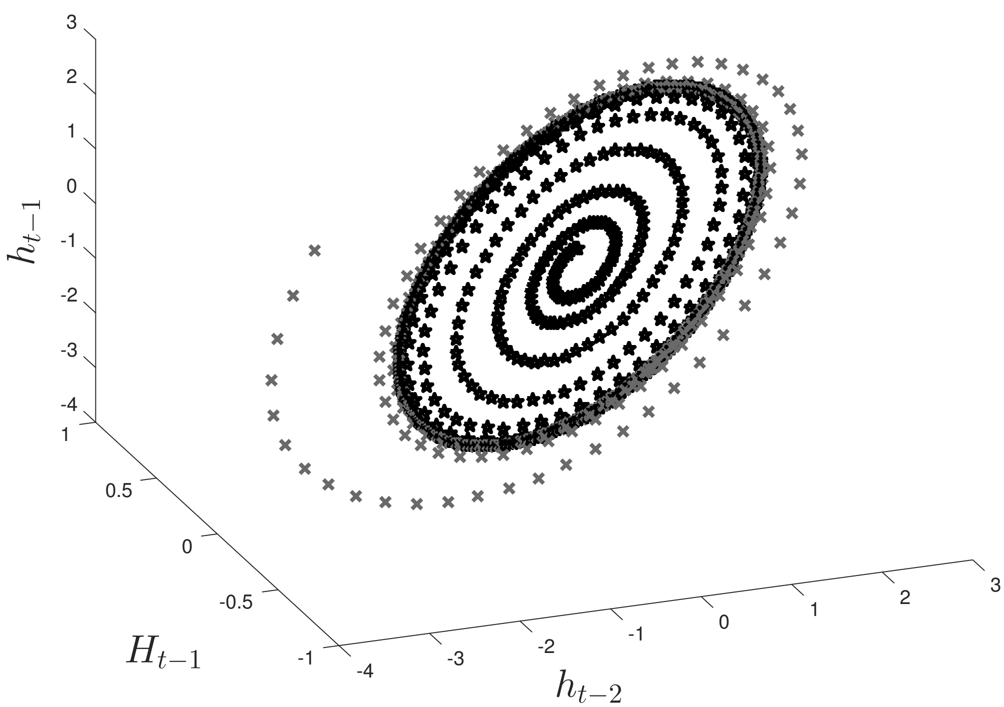

Notes: This corresponds to the deterministic simulation of the Minimal model (13), starting from two points arbitrarily chosen inside and outside the limit cycle. Those two points are, respectively, $(h_0, h_{-1}, H_{-1}) = (.1, .1, .1)$ and $(h_0, h_{-1}, H_{-1}) = (6, 4, 10)$. The models have been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

Figure 4: The Limit Cycle in the Three Models, Total Hours

(a) In (h_t, H_t)-space (b) Time series of h_t

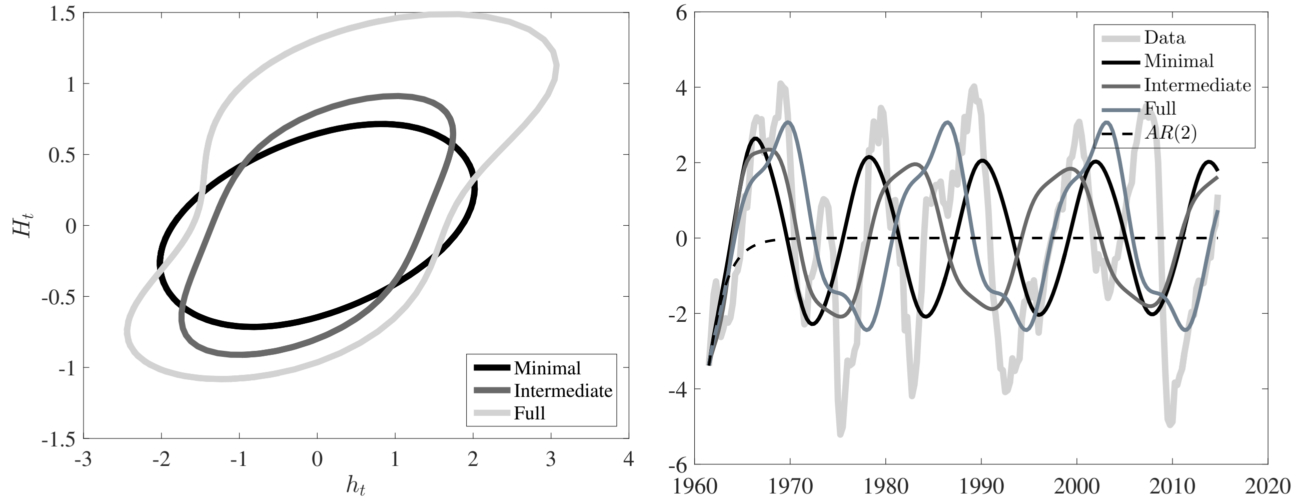

Notes: This corresponds to the deterministic simulation of the Minimal (13), Intermediate (14) and Full (12) models, starting (arbitrarily) in 1961Q3. The model has been estimated over 1960-2014 using 80-quarter high-pass-filtered Total Hours.

simulation is constructed as

$$\begin{cases} \widetilde{h}^G_{T_0} & = & h_{T_0}, \\ \widetilde{h}^G_{T_0-1} & = & h_{T_0-1}, \\ \widetilde{H}^G_{T_0-1} & = & H_{T_0-1}, \\ \widetilde{h}^G_t & = & \widehat{\beta}_h \; \widetilde{h}^G_{t-1} + \widehat{\beta}_{h'} \; \widetilde{h}^G_{t-2} + \widehat{\beta}_H \; \widetilde{H}^G_{t-1} + \widehat{\sigma}\nu_t & \forall\; t > T_0, \\ \widetilde{H}^G_t & = & (1-\delta)\widetilde{H}^G_{t-1} + \delta \widetilde{h}^G_t & \forall\; t \geq T_0, \end{cases}$$

where $\widehat{\sigma}$ is the estimated standard deviation of ϵ in the Minimal model and ν is drawn from a $\mathcal{N}(0,1)$ distribution. To construct our "Generalized" forecast, we average across 10,000 such simulations.

In panel (b) of Figure 5, we display this Generalized forecast based on the estimated minimal model. Although this path eventually dies out, note that it is much more persistent and cyclical than the $AR(2)$ model. Recall that, for a linear model, the deterministic and Generalized forecasts are the same. For the $AR(2)$ model we see that, conditional on being in the trough at 1961Q3, the economy converges monotonically to the steady state, while the Minimal model captures the subsequent expansion, its turning point, and the recession that follows. This suggest that the mean squared forecast errors should be smaller for the Minimal model than for the $AR(2)$ when looking at different forecasting horizons. This is true for the mean squared forecast

Figure 5: Forecasted Path as of 1961Q3 with the Minimal Model, Total Hours

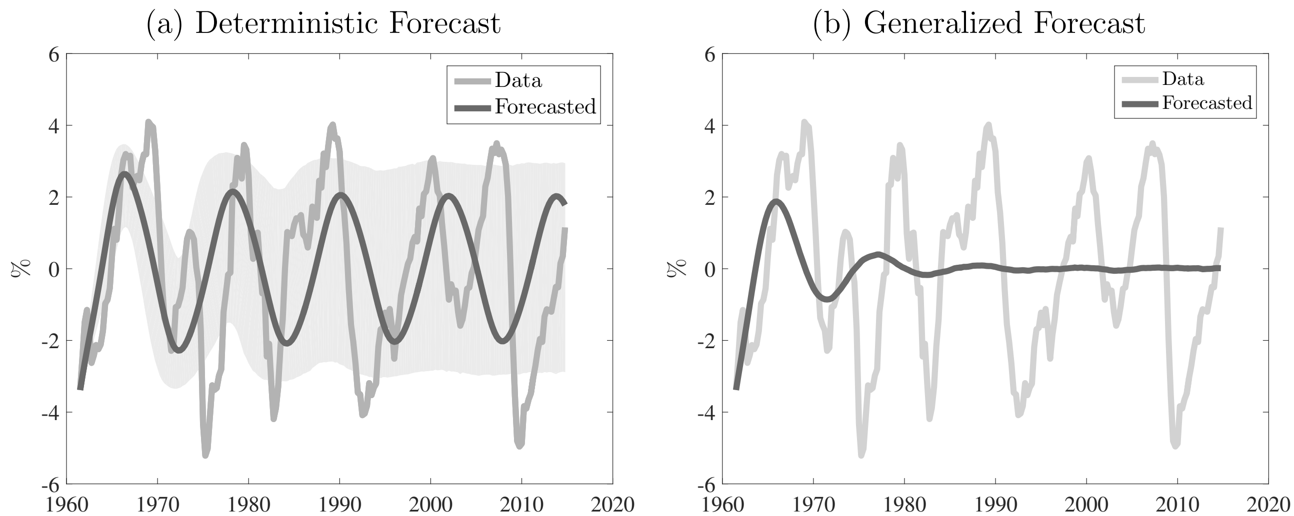

Notes: In panel (a), the dark grey line labeled "Forecasted" is the deterministic simulation of the Minimal (13) model, starting (arbitrarily) in 1961Q3, and the light grey line represents observed filtered hours. The grey area represents the 95% confidence band for the deterministic forecast, as obtained from 10,000 Monte Carlo simulations. In panel (b), the dark grey line labeled "Forecasted" is the mean of the mean of 10,000 simulations $\{\widetilde{h}_t^G\}$*, and the light grey line represents observed filtered hours.*

errors at horizon one, which is simply the R^2. But, as shown in Table 3, the increase in R^2 for the Minimal model relative to the $AR(2)$ case is extremely marginal, with the two differing only at the third decimal place. The comparison of "dynamic forecasts" at horizon one makes this point mostly clearly. These forecasts are constructed in the following way for the minimal model (and in a similar way for the $AR(2)$):

$$\widehat{h}_t = \widehat{\beta}_h \ h_{t-1} + \widehat{\beta}_{h'} \ h_{t-2} + \widehat{\beta}_H \ H_{t-1} \quad \forall \ t > T_0.$$

Figure 6: Dynamic Forecast for the Minimal and $AR(2)$ Models, Total Hours

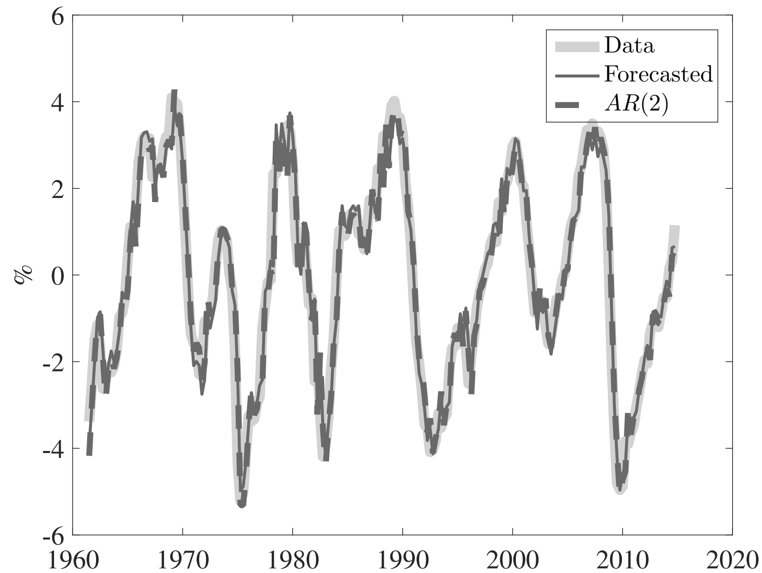

Notes: This corresponds to the dynamic forecasts $\{\widehat{h}_t\}$ computed with the Minimal (13) and $AR(2)$ (15) models.

As shown in Figure 6, there is virtually no difference between the one-step-ahead dynamic forecats for the $AR(2)$ and Minimal models since both almost perfectly fit the data. In other words, if one focuses on the one-step-ahead forecast, there is very little gain in adopting the non-linear specification. But, as illustrated in Figure 6, the non-linear and linear models differ significantly at longer horizons. Figure 7 compares the mean over 1000 simulations of the square root sum of squared k-step-ahead forecast errors for the Minimal and $AR(2)$ models. In the case of the Minimal model, we have performed generalized forecast, with 1000 simulations. [9] For horizons up to 50

[9]See the appendix for a formal description of the procesure we follow

quarters, the non-linear model performs better that the $AR(2)$, whereas the two models perform about the same for longer horizons. The predictability implied by the limit cycle appears helpful in forecasting the pattern of one business cycle.

Figure 7: Mean of the Square Root Sum of Squared k-step-ahead Forecast Errors for the Minimal and $AR(2)$ Models, Total Hours

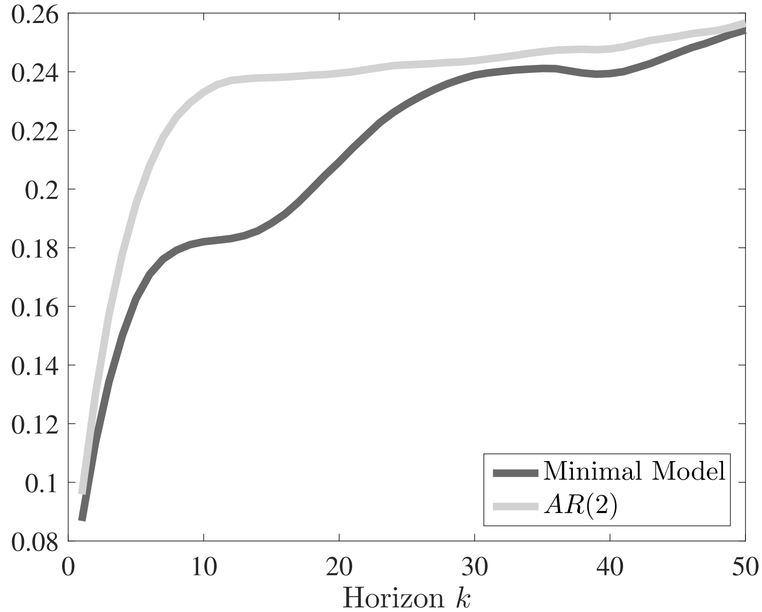

Notes: The models have been estimated over 1960-2014 using 80-quarter high-pass-filtered Total Hours.

Other Labor market variables

We now explore the robustness of our limit cycle result for other labor market variables. In particular, we examine the behavior implied by estimating our Minimal model using sequentially the Non-Farm 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 the (x_t, X_t) space as well as x_t over time, where x is the variable under consideration and X its "cumulated" level, according to (10). Panels (a) and (b) of Figure 8 correspond to the forecasted paths for Non-farm 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 system at the values associated with 1961Q3. As is clear from these figures, a limit cycle appears in all three cases, and these cycles have durations close to 9 years, with amplitudes similar to the actual data.

Non-Linearity

For a model to generate a limit cycle it must be non-linear. However, this non-linearity could be very large or it could be quite small. One can obtain a visual sense of the degree of non-linearity behind the Minimal model quite easily. With the estimated coefficients, we plot in Figure 9 the contribution to h_t of h_{t-1}, which is given by $\beta_h h_{t-1} + \beta_{h^3} h_{t-1}^3$. This term is the only non-linear component in the Minimal model and therefore allows for easy representation. In Figure 9 the light grey line corresponds to the estimated relationship $\beta_h h_{t-1} + \beta_{h^3} h_{t-1}^3$ for our Total Hours series, while the dark grey locus corresponds to the effective range of the relationship in terms of observed values of h. One can see in this Figure that the non-linearity is not very pronounced. How is such a minor non-linearity capable of generating a limit cycle? The reason is that the economy is diverging very slowly in the neighborhood of the steady state (the maximum eigenvalue modulus is 1.01), so that only a small degree of non-linearity is sufficient to prevent global explosion and instead produce a limit cycle.

Multiplicity of steady states

Since we are examining dynamics in a non-linear framework, we need to recognize the possibility of multiple steady states. First, let us note that all the estimated equations are backward-looking, so that there is (trivially) never indeterminacy of the solution to our equations.[10] 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 for ease of exposition:

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

We typically have β_h larger than one, $\beta_{h'}$ negative, and β_H and β_{h^3} negative and small. In this Mininal model, non-zero steady states, if they exist, are the two opposite real

[10] This does not mean that micro-founded models capable of generating limit cycles never feature indeterminacy.

Figure 8: The Limit Cycle in the Minimal Model, Other Labor Market Variables

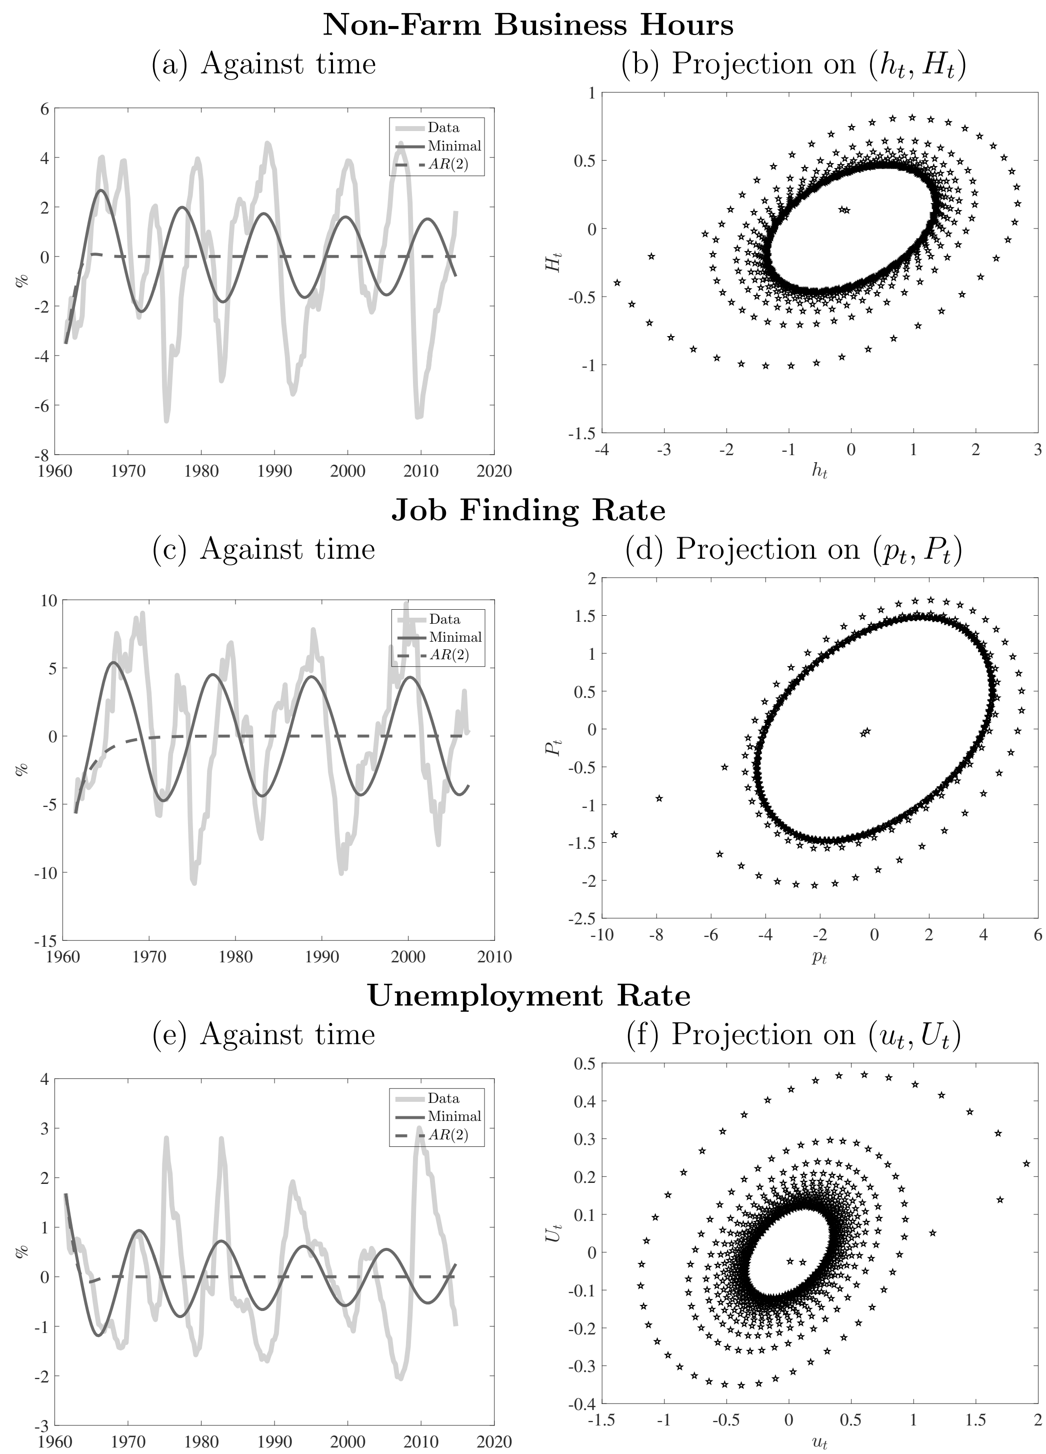

Notes: This corresponds to the deterministic simulation of the Minimal Model (13), starting in 1961Q3. The model has been estimated over 1960-2014 (2007 for the Job Finding Rate) using 80-quarter high-pass-filtered series of Non Farm Business Hours, Job Finding Rate (from [Shimer 2012]) and Unemployment Rate. P and U stand for "cumulated" rates, and are constructed in the same way than H in equation (10).

Figure 9: Non-Linearities in the Minimal Model, Total Hours

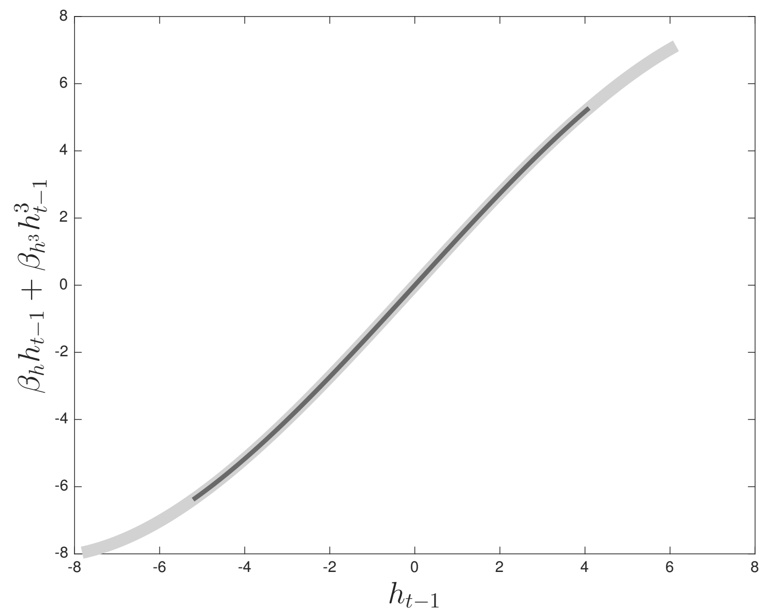

Notes: The coefficients β are taking the estimated values of the Minimal (13) *model. The dark gray line evaluates $\beta_h h_{t-1} + \beta_{h^3} h^3_{t-1}$ for all the in-sample values of h. The light gray line does it extending the range by 50% above and below. The model has been estimated over 1960-2014 using 80-quarter high-pass-filtered series of Total Hours.*

numbers $\overline{h}$ and $-\overline{h}$ with

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

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

$$\beta_h + \beta_{h'} + \beta_H > 1. \tag{16}$$

Note that when β_H goes to zero the condition for local instability of the zero steady state can be shown to be, under the sign restriction for the β's given above,

$$\beta_h + \beta_{h'} > 1. \tag{17}$$

Therefore, when β_H goes to zero, (16) and (17) are incompatible: the occurrence of attractive limit cycles, which requires local instability, excludes multiplicity of steady states.[11] Nevertheless, when β_H is below but not arbitrarily close to zero, we can see from (17) that multiple steady states are possible. In that case, are the non-zero steady states stable or unstable when the zero steady state is unstable? Taking a first-order Taylor expansion of (13) around $\overline{h}$, we have the following dynamics (omitting constant terms):

$$h_t = (\underbrace{\beta_h + 3\overline{h}^2 \beta_{h^3}}_{\overline{\beta}_h})\ h_{t-1} + \beta_{h'}\ h_{t-2} + \beta_H\ H_{t-1}. \tag{18}$$

Note that the same equation holds for the local dynamics around $-\overline{h}$. As we are looking at situations where $\beta_{h^3} < 0$, we will have $\overline{\beta}_h < \beta_h$, such that it is possible that the two non-zero 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 $\overline{h}$ or $-\overline{h}$. The same argument can also be made for the Intermediate and Full model. Whether or not such a case is actually relevant depends on parameter estimates. As we shall see, the data do not point in that direction.

Given the estimated coefficients, zero is the unique steady state in the Minimal model. However there are multiple steady states for the other two models. In fact, in

[11] This analysis can be extended to the Intermediate model. Non-zero steady states, if they exist, are the two opposite real numbers $\pm\overline{h}$ with

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

addition to the zero steady state, the Intermediate model has steady states at .32 and -.32, while the full model has steady states at 1.23 and -.65 (in addition to the zero steady state). Interestingly, in the two last models, the non zero steady states are also locally unstable, which implies that the dynamics cannot converge to a steady state in any of our estimated model when focusing on the behavior of Total Hours worked .

It is not possible to prove analytically that trajectories that start in the neighborhood of a non-zero steady state converge to the same limit cycle as those starting from the neighborhood of the zero steady state, but it can be checked numerically. Figure 10 displays the deterministic forecast of the Intermediate and Full models, starting from the neighborhood of each of their three steady states (represented with a star), in the plane (h, H). Although the convergence is much lower in the Intermediate model than in the Full one, both models show convergence towards a unique limit cycle. A further test is to compute a deterministic forecast, taking each of the observed quarters of the sample as the initial condition. In all cases, we found convergence to the same limit cycle.

Figure 10: Convergence to the Limit Cycle From the Neighborhood of Each of the Three Steady States, Total Hours

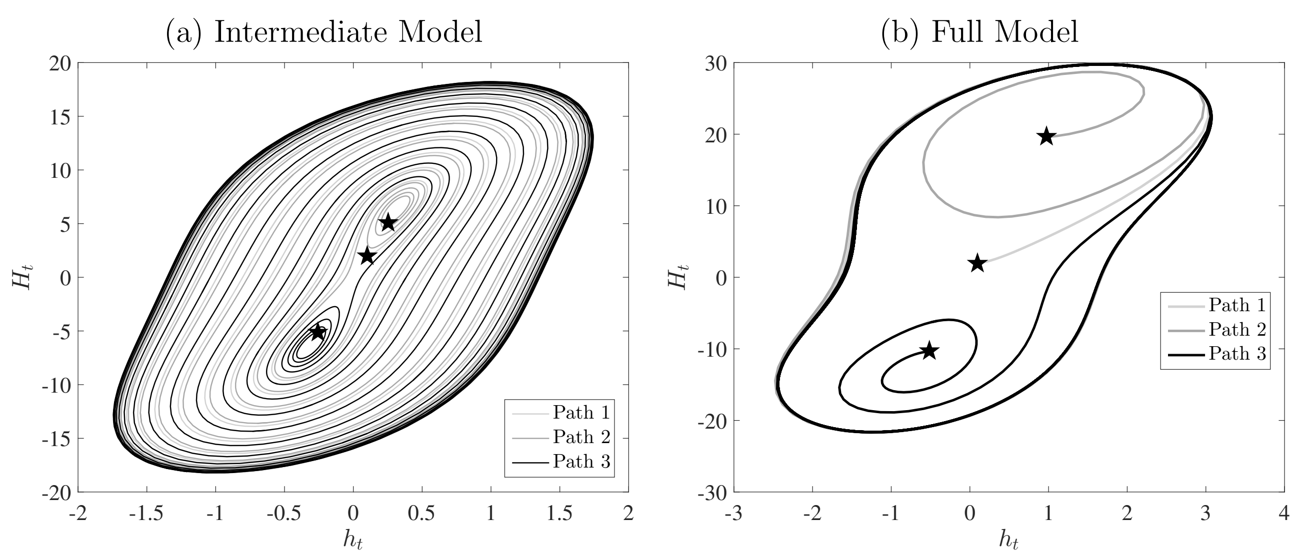

Notes: The figures for the Intermediate model (panel (a)) and Full model (panel (b)) show a deterministic forecast starting from the neighborhoods of each of the three steady states, represented by stars. The three trajectories converge to the same limit cycle. The models have been estimated over 1960-2014 using 80-quarter high-pass-filtered Total Hours.

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 Data Generating Process (DGP) is actually locally stable, or alternatively 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 $\widehat{h}_t$, then add it to the actual trend series h_t^T to generate an artificial level series

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

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: (i) the steady state is unstable, meaning that the modulus of the maximal eigenvalue is larger than one; (ii) the Wald test rejects the $AR(2)$ model against the non-linear one at the level of significance equal to the p-value found for the corresponding Wald tests in the data (i.e., the first line of Table 2); and (iii) the model converges to the limit cycle. We perform 10,000 simulations, and report the results in the first row of Table 4. The three models spuriously detect limit cycles less than .6% of the time, suggesting that it is highly unlikely to detect a limit cycle with our procedure if it does not exist in the DGP.[12]

[12] When we set the threshold level to 1% in the Wald tests (instead of the estimated p-value), the test accepts the (spurious) existence of a limit cycle 1.9% of the time for the Minimal model, and respectively

Table 4: Percentage of Limit Cycle Detection, Total Hours

	Estimated model:		
	Minimal	Intermediate	Full
DGP:			
$AR(2)$	.4	.6	.1
Minimal	53	47	17
Intermediate	43	48	26
Full	52	40	50

Notes: In this table, the DGP is alternatively the estimated $AR(2)$ (15)*, Minimal* (13)*, Intermediate* (14) *and Full* (12) *models. The models have been estimated over 1960-2014 using 80-quarter high-pass-filtered Total Hours, and are then simulated 10,000 times.*

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 rows of Table 4, and show that limit cycles are relatively hard to detect using our procedure, because of filtering and the small sample size. 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.

From theses Monte Carlo simulations, we conclude that it is unlikely that we detect spurious limit cycles, and that detection of limit cycles works at best one times out of two. Those results strengthen our previous finding of limit cycles in the dynamics of Total Hours.

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 worked per capita. We begin by examining the robustness of our results with respect to our choice of δ in the building 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 (see equation (4)). Table 5 shows that a limit cycle is detected generically for the Minimal model, at as long as δ is not too large for the Intermediate and Full models, and as long as δ is not too small in the

6.6% and 4.6% for the Intermediate and Full models.

Table 5: Robustness to δ, Total Hours

	Minimal	Intermediate	Full
$\delta = .001$	LC	-	LC
$\delta = .01$	LC	-	LC
$\delta = .1$	LC	LC	LC
$\delta = .2$	LC	-	-

Notes: In this table, we estimate the Minimal (13)*, Intermediate* (14) *and Full* (12) *models under various values of δ. "LC" means that the estimated model displays a limit cycle. All the models have been estimated over 1960-2014 using 80-quarter high-pass-filtered Total Hours.*

Intermediate model.

Next, Table 6 considers filters other than the 80-quarter high-pass one used thus far. Let us first note that the high-pass filter is a linear filter, and is therefore incapable of spuriously introducing non-linearities, so that in general we should not expect to find cycles where there are none simply by choosing the right filter. On the other hand, if there exists a limit cycle in the data but the filter we choose removes the frequencies associated with that cycle, our procedure would in general fail to identify this cycle. Further, given the 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 6 largely support these predictions. In particular, we report results using different high-pass filters from 100 to 60 quarters, as well as for commonly-used band-pass (6,32) and Hodrick-Prescott ($\lambda = 1600$) filters. We also report results for when we detrend using polynomial time trends of order 3, 4 and 5.

Two things emerge from the Table. First, the detrending procedures that remove only the lowest frequencies (i.e., the 60+-quarter 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 ($\lambda = 1600$) 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

Table 6: Robustness to Detrending, Total Hours

	Minimal	Intermediate	Full
High-pass 100	LC	LC	LC
High-pass 90	-	LC	LC
High-pass 70	LC	LC	LC
High-pass 60	-	LC	-
Band-pass(6,32)	-	-	-
Hodrick-Prescott ($\lambda = 1600$)	-	-	-
Polynomial trend (3^{rd}-order)	-	LC	LC
Polynomial trend (4^{th}-order)	-	LC	LC
Polynomial trend (5^{th}-order)	LC	LC	LC

Notes: In this table, we estimate the Minimal (13)*, Intermediate* (14) *and Full* (12) *models using various detrending methods. "LC" means that the estimated model displays a limit cycle. All the models have been estimated over 1960-2014 using Total Hours.*

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.

2.4 Behavior in Other Countries

In this subsection, we extend our exploration of macroeconomic dynamics using data on the unemployment rate in other countries.[13] These results are presented in Table 7. 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

[13] See Appendix B for details about the source and time span of each of the series.

Table 7: Limit Cycles in Other Countries

	Minimal	Intermediate	Full
Unemployment rate :			
Australia	-	-	-
Austria	LC	LC	LC
Canada	-	-	-
Denmark	-	LC	LC
France	-	-	LC
Germany	LC	LC	LC
Japan	-	-	-
Netherlands	LC	LC	-
Sweden	LC	LC	-
Switzerland	-	-	-
United Kingdom	-	-	LC

Notes: In this table, we estimate the Minimal (13)*, Intermediate* (14) *and Full* (12) *models, considering various labor market variables. "LC" means that the estimated model displays a limit cycle. Data are always detrended with an 80-quarter high-pass filter before estimation. See Appendix B for details of samples and variable definitions.*

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, with an almost zero level unemployment in the 70's and 80's, but a 3.5% average since 1990. Besides this low-frequency variation, what emerges from Figure 11 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 U.S., 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 does not exhibit a limit cycle, and the dynamics of forecasted unemployment as of 1975 is not very different from the one obtained with an $AR(2)$ model.[14] Things are very different for

[14] A similar pattern is found using the Intermediate and Full models (not shown in the Figure).

Figure 11: Unemployment Dynamics in Switzerland and the Netherlands

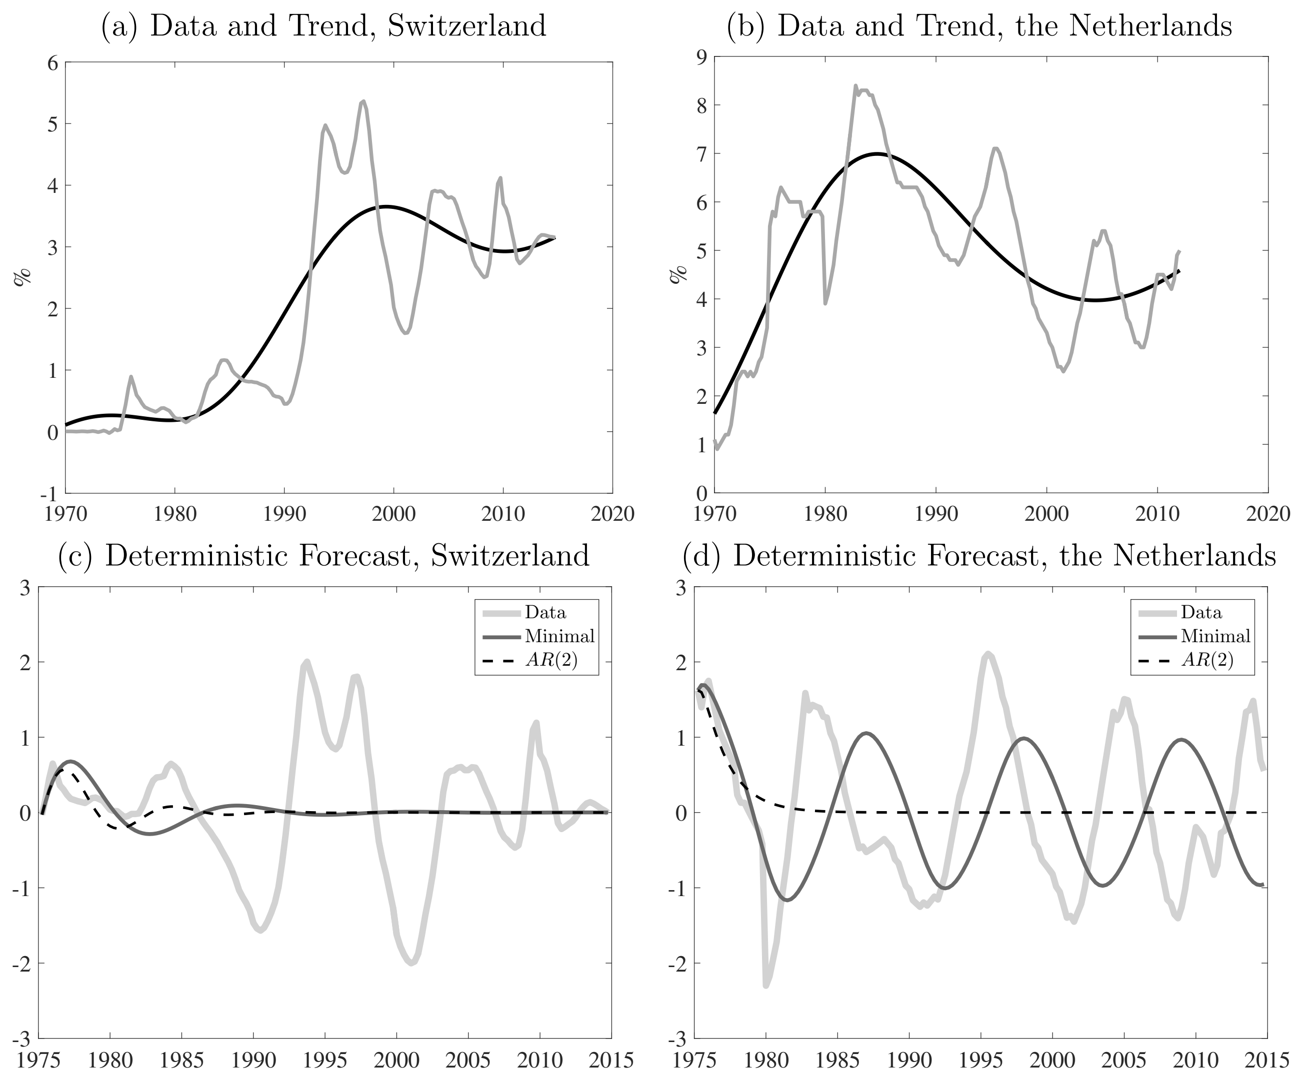

(a) Data and Trend, Switzerland

(b) Data and Trend, the Netherlands

(c) Deterministic Forecast, Switzerland

(d) Deterministic Forecast, 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 1975Q2. 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 1975Q1, with the estimated $AR(2)$ (15) and Minimal (13) models.

the Netherlands (Figure 11, panels (b) and (d)). The low-frequency dynamics feature a marked increase from the early 70's to the late 80's, followed by a decrease thereafter. The key determinants of those trend movements are the first and second oil shocks, followed by important labor market reforms in the 90's. 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. The deterministic forecast for the Netherlands as of 1975 is therefore very different from the one obtained with an $AR(2)$.

2.5 Quantity Variables

Difficulties

We now turn to the estimation of our non-linear reduced-form model using quantity variables. The equivalent of the dynamic system (7) and (10) in hours can now be rewritten in terms of output y as

$$y_t = \alpha_1 y_{t-1} - \alpha_2 Y_{t-1} + E_{t-1} F\left(y_t - \theta_t\right) + \theta_t - \left(\alpha_1 - \alpha_2\right)\theta_{t-1} \tag{19}$$

$$Y_t = (1-\delta)Y_{t-1} + \delta y_t, \tag{20}$$

Without an accurate measure of TFP, θ_t, we cannot estimate these equations in reduced form as we could with our labor-market variables: the reduced-form models (12)-(14) with filtered y_t in place of h_t are likely to be misspecified if the true model is (19) and (20). This can be seen in practice by running the following Monte Carlo experiment. We estimate a Minimal, Intermediate or Full model on 80-quarter high-pass filtered Total Hours. As shown previously, all three of these estimated models feature a limit cycle. We then simulate the estimated model to obtain an artificial cyclical series of hours, to which we add the high-pass trend that we have recovered from the data and (the log of) labor productivity, so that we obtain a log-level series of artificial output. We then treat that 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 8 the percentage of simulations for which a limit cycle is (correctly) detected. The three models hardly detect limit cycles in outputDetection rates are around 15% when the "correctly specified" model is used for estimation. Despite this low detection

Table 8: Percentage of Limit Cycle Detection, Simulated Output

	Estimated model:		
	Minimal	Intermediate	Full
DGP:			
Minimal	14	13	5
Intermediate	12	15	6
Full	12	11	11

Notes: In this table, the DGP is alternatively the estimated Minimal (13)*, Intermediate* (14) *and Full* (12) *models. The models have been estimated over 1960-2014 using 80-quarter high-pass-filtered Total Hours, and are then simulated 10,000 times.*

rate, we show in the next paragraph that we quite often find a limit cycle in the actual data, mainly with the Intermediate model.

Estimation Results

We estimate the three models on output, total consumption, durable goods expenditures, fixed investment, and on various components of investment. Estimation results are summarized in Table 9. We find a limit cycle with the Intermediate model for all variables except Non-Farm Business Output, Fixed Investment and Equipment. As an example, Figure 12 displays the deterministic forecast as of 1961Q3 with the estimated Intermediate model for both Total Output and Non-Farm Business Output. In the former case, the maximum eigenvalue is $0.998 \pm 0.08i$, with modulus 1.001, so that there is a limit cycle, whereas in the later the maximum eigenvalue is $0.995 \pm 0.08i$ with modulus .998, so that there is no limit cycle. Nevertheless, it is clear from Figure 12 that the deterministic forecasts are very similar, although the fluctuations eventually die out in the latter case.

3 Stabilization Policy in a Stochastic Limit Cycle Environment

3.1 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

Table 9: Existence of a Limit Cycle for Quantity Variables

	Minimal	Intermediate	Full
Output	-	LC	-
Non Farm Bus. Output	-	-	-
Consumption	-	LC	LC
Fixed Investment	-	-	-
Structures	LC	LC	LC
Durables	-	LC	LC
Residential	-	LC	-
Equipment	-	-	-

Notes: In this table, we estimate the Minimal (13)*, Intermediate* (14) *and Full* (12) *models, considering various goods market variables. "LC" means that the estimated model displays a limit cycle. Data are always detrended with an 80-quarter high-pass filter before estimation. See Appendix* B *for details of samples and variable definitions.*

Figure 12: Deterministic Forecast with the Intermediate Model: Total and Non-Farm Business Output

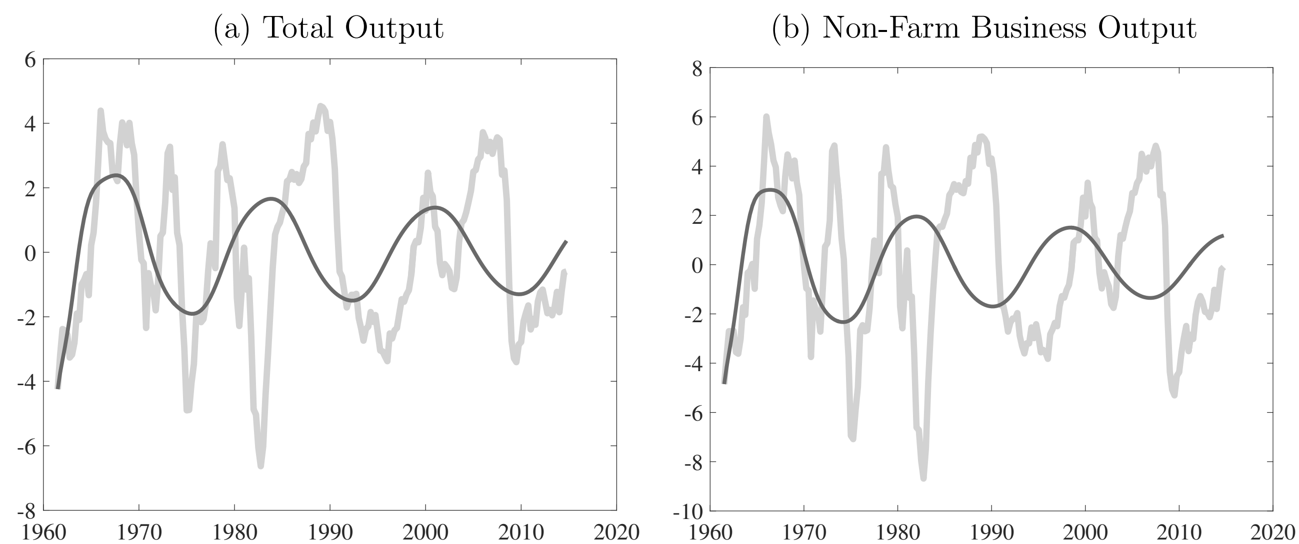

Notes: Total and Non-Farm Business Output are observed over the sample 1947:Q1-2014:Q4. The cumulated series is constructed assuming $N = 40$*, with estimation starting in 1961Q1. The trend of those series is obtained by removing from the level series its filtered component, where the filter is an 80-quarter high-pass filter, over the sample 1947Q1-2014Q4. The cyclical component is expressed in percentage deviation from the level series. Deterministic forecasts are done starting from 1961Q3 with the estimated Intermediate* (14) *model.*

employment. One stabilization approach that is commonly pursued is to direct policy towards countering or nullifying the impact of exogenous forces hitting the economy. In linear models, such a strategy can in general 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, *(i)* a reduction in the variance of shocks may actually increase the variance of endogenous variables; *(ii)* 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 *(iii)* the principal effect of reducing the shock variance is 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.[15] For example, consider an environment where the dynamics of a state variable of interest x_t may we written as

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

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,[16] then the variance of x will be strictly increasing in the variance of ϵ. However, if the function $F(\cdot)$ is non-linear, then the link between the variance of x,[17] 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

[15] In this section we will be discussing the effects of changes in shock variances on the variances of endogenous variables. In doing such an exercise, one must be aware of the Lucas Critique. Since we focus only on changing the variances of the shocks, not their time series properties, the Lucas Critique of policy evaluation does not always apply. In particular, in the case of linear models, changing the variances of the shocks will not in general change the other parameters of the DGP. Whether this also holds in non-linear models depends on the details of the model, and especially where and when expectations are taken. Accordingly, the discussion in this section needs to be interpreted cautiously given the potential for the Lucas Critique to apply.

[16] That is, the roots of the polynomial $1 - F(\lambda, \lambda^2, \dots)$ lie outside the unit circle.

[17] With some abuse of terminology, we define the mean and variance of x by time averages rather than ensemble averages, i.e., by $\mu_x \equiv \text{plim}_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} x_t$ and $\sigma_x^2 \equiv \text{plim}_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} (x_t - \mu_x)^2$, 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 non-stochastic environment as long as the system does not converge to a single point.

the deterministic system $x_t = F(x_{t-1}, x_{t-2}, \dots)$ 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 2, but in which we remove the accumulation term and the $AR(2)$ one for analytical tractability:

$$x_t = -\beta_x x_{t-1} + \beta_{x^3} x_{t-1}^3 + \epsilon_t, \tag{21}$$

Under the restrictions

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

and

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

and when $\sigma_\epsilon^2 = 0$, then one may verify that the system (21) possesses a limit cycle of period two[18] in which x alternates between $\bar{x}$ and $-\bar{x}$, where $\bar{x} \equiv \sqrt{\frac{\beta_x - 1}{\beta_{x^3}}}$. Moreover, this limit cycle is globally attractive, so that all orbits (except for those beginning at the fixed point $x = 0$) will converge to the 2-cycle. 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 (21) 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 D. □

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. The intuition for why this arises reflects the nature of the non-linear 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 towards 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 inwards are more persistent, and therefore, on average, for σ_ϵ^2 sufficiently small the system spends

[18] See May [1979].

more time inside the cycle than outside of it. This effect causes the overall variance of x to be decreasing in σ_ϵ^2.

3.2 Using the Estimated Minimal Model

To follow up on Proposition 1, we have explored numerically whether the negative relationship between shock variance and outcome variance suggested by the proposition arises in richer environments than the one it covers. In particular, we have examined this question numerically for the type of data-generating process we estimated in the previous section, and have found in all cases that the negative relationship between shock variance and output variance holds for small values of σ_ϵ^2. However, we have also generally found in these cases that the negative relationship is quite weak, so that, quantitatively, the shock variance has almost no effect on output variance at small values of σ_ϵ^2. In this sense, a more robust take-away appears to be that, in environments that support limit cycles of the size exhibited 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 Figure 13 we plot the relationship between σ_ϵ and σ_x for two models estimated using Total Hours worked per capita as the x variable.[19] The Minimal model is, as in the previous section,

$$h_t = \beta_h h_{t-1} + \beta_{h'} h_{t-2} + \beta_H H_{t-1} + \beta_{h^3} h_{t-1}^3 + \epsilon_t,$$

while the linear model simply imposes $\beta_{h^3} = 0$. The estimated parameters of these two models are given in Table 10. As was shown in the previous section, in the absence of shocks this non-linear data-generating process exhibits limit cycles. On the contrary, the estimated linear model has a stable steady state.

The relationship between the standard deviation of ϵ and the standard deviation of x for the two models is plotted in panel (a) of Figure 13. The most notable feature that emerges is the extent to which that relationship is much weaker in the case of the limit cycle model as compared to the linear model. For example, increasing σ_ϵ from 0.1 to 0.5 causes σ_x to increase by 5 times (400%) in the linear model, while it increases by only around 25% in the Minimal model. The fact that the plotted relationship has a non-zero intercept for the limit cycle model is not surprising, but the fact that the

[19] As before, these data have been filtered using an 80-quarter high-pass filter.

Table 10: Estimated Coefficients

	Minimal	**Linear**
β_h	1.39	1.31
$\beta_{h'}$	-0.34	-0.35
β_H	-0.27	-0.25
β_{h^3}	-3e-03	–
σ_ϵ	0.57	0.57

Notes: The Total Hours series have been filtered with an 80-quarter high-pass filter over the sample 1948Q1-2015Q2. Estimation is then done over the sample 1960Q1:2014Q4.

slope of this relationship is so much weaker on average is telling. This pattern suggests that directing policy to counteract the effects of exogenous shocks is potentially much less effective if the underlying system exhibits limit cycles of the type we have found in macroeconomic data.

In panel (a) of Figure 13, it is hard to tell graphically whether the slope of the relationship between σ_ϵ and σ_x is positive or negative for low values of σ_ϵ in the non-linear model. In panel (b) of the Figure we focus in on this relationship at smaller values of σ_ϵ in order to emphasize that it is in fact negative, in accordance with the intuition we derived from Proposition 1. Once again, let us stress that the pattern we think is most relevant is not that there exists a range over which the relationship between input volatility and outcome volatility is negative, but instead that the relationship can be extremely weak relative to its linear counterpart, suggesting that stabilization policy aimed at mitigating the effects of shocks may be of limited value. The goal of stabilization policy in such a case should instead be focused on identifying and mitigating the forces that generate and sustain limit cycles.[20]

To further emphasize how changes in input volatility affect outcome behavior in models with or without limit cycles, in Figure 14 we plot the spectrum of the outcome variable x for several different values of σ_ϵ. Panel (a) shows spectra for the non-linear

[20] In our previous work (Beaudry, Galizia, and Portier [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.

Figure 13: Relationship between σ_x and σ_ϵ

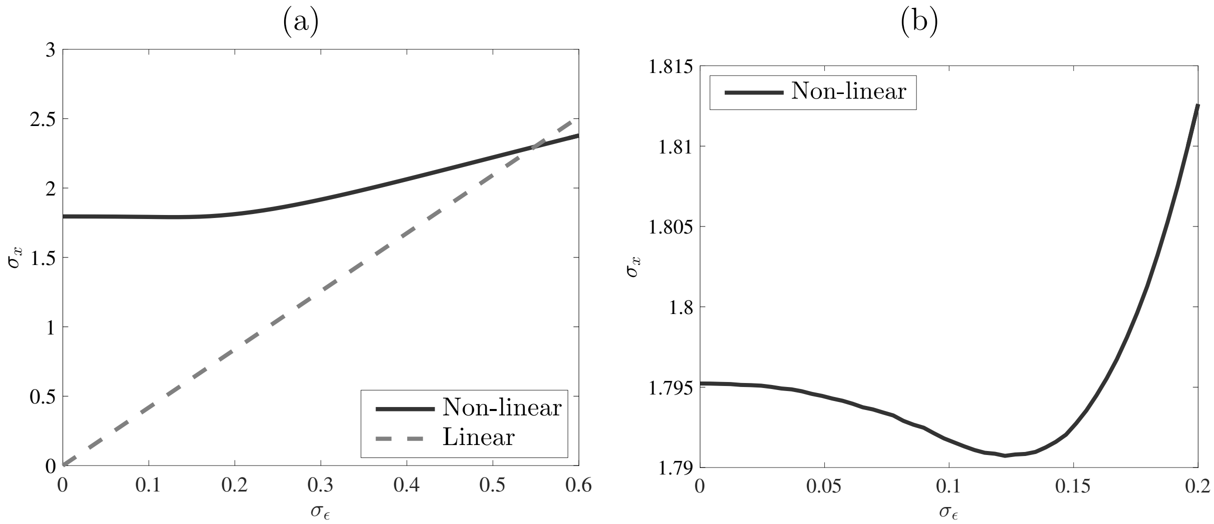

Notes: For each model and each value of σ_ϵ, we simulate $T_{tot} = T_{brn} + T$ periods of data, with $T_{brn} = 50{,}000$ and $T = 100{,}000{,}000$. After discarding the first T_{brn} periods of data, we compute σ_h^2 as the time average $T^{-1} \sum h_t^2$, where we note that $\mu_h = 0$ by symmetry.

specification, while panel (b) shows spectra for the linear one.[21] 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 half, from 0.56 (as estimated for the data) to $0.56/2 = 0.28$, the spectrum at periodicities below 40 quarters decreases substantially, while the importance of periodicities slightly 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.

[21] Spectra were generated using the parameters reported in Table 10.

Figure 14: Spectrum of x

(a) Non-linear (b) Linear

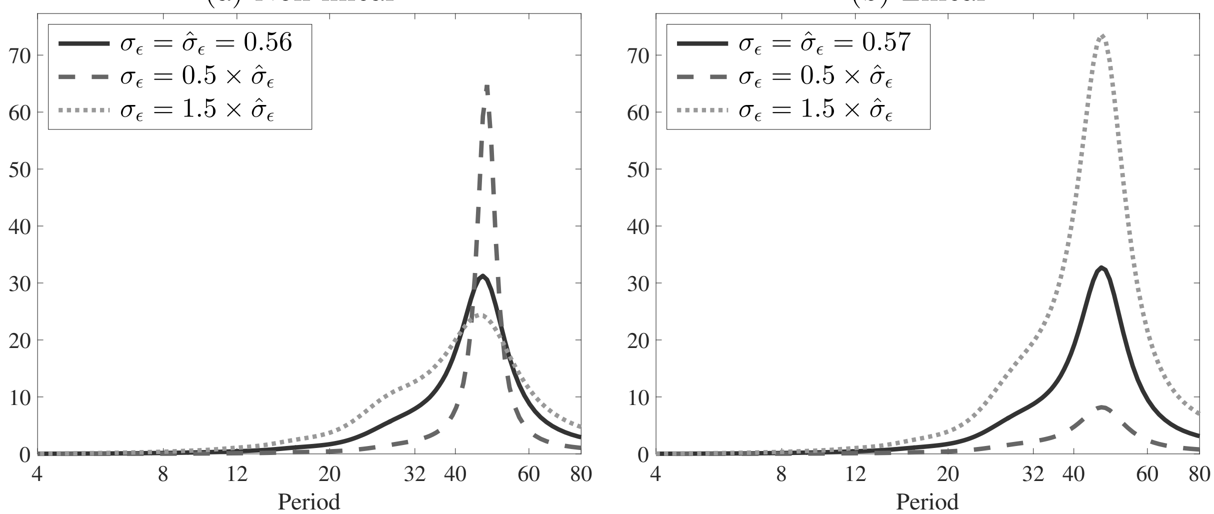

Notes: For each model and each value of σ_ϵ, we simulate 1,000,000 data sets, each of length $T_{tot} = T_{brn} + T$ quarters, with $T_{brn} = 50,000$ and $T = 1,000$. For each simulated data set, we obtain the spectrum of h after discarding the first T_{brn} periods, then average the result across all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

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 macro-economy is likely to be much less effective at reducing economic volatility if limit cycles are present.[22] 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 towards a stabilization framework that aims to reduce the forces that may be causing the limit cycles. For example, in our previous work (Beaudry, Galizia, and Portier [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

[22] More precisely, a reduction in input volatility may have very little impact on the the overall volatility of the system when the deterministic version of the system exhibits limit cycles.

these complementarities. If the complementarity is related to precautionary saving associated with unemployment risk as in Beaudry, Galizia, and Portier [2015a] and [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.

4 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 which repeatedly gives rise to sustained boom-and-bust phenomena. To examine this question we have proposed a simple non-linear 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 considerable 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 indicated the presence of forces that favor recurrent business cycles with a duration of close to 9 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 intriguing 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.[23] This suggests to us that the

[23] For most cases we found the maximum eigenvalue in modulus of the estimated system to be very close to 1.

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 locally instability may matter for policy.

References

Beaudry, P., D. Galizia, and F. Portier (2015a): "Reconciling Hayek's and Keynes' Views of Recessions," Working Paper.

——— (2015b): "Reviving the Limit Cycle View of Macroeconomic Fluctuations," Working Paper.

Cooper, R., and A. John (1988): "Coordinating Coordination Failures in Keynesian Models," *The Quarterly Journal of Economics*, 103(3), 441–463.

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 – 965.

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

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

Kalecki, M. (1937): "A Theory of the Business Cycle," *The 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), 623–651.

May, R. M. (1979): "Bifurcations and Dynamic Complexity in Ecological Systems," *Annals of the New York Academy of Sciences*, 357, 267–281.

Myerson, R. B. (2012): "A Model of Moral-Hazard Credit Cycles," *Journal of Political Economy*, 120(5), 847–878.

Shimer, R. (2012): "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 15(2), 127–148.

A Atlernative Estimation Environments

A.1 Minimal Model with AR(1) Component

See Figure A.1. Quick comment is needed.

Figure A.1: The Limit Cycle in the Minimal-$AR(1)$ and Minimal Model

(a) Against time (b) Projection on (h_t, H_t)

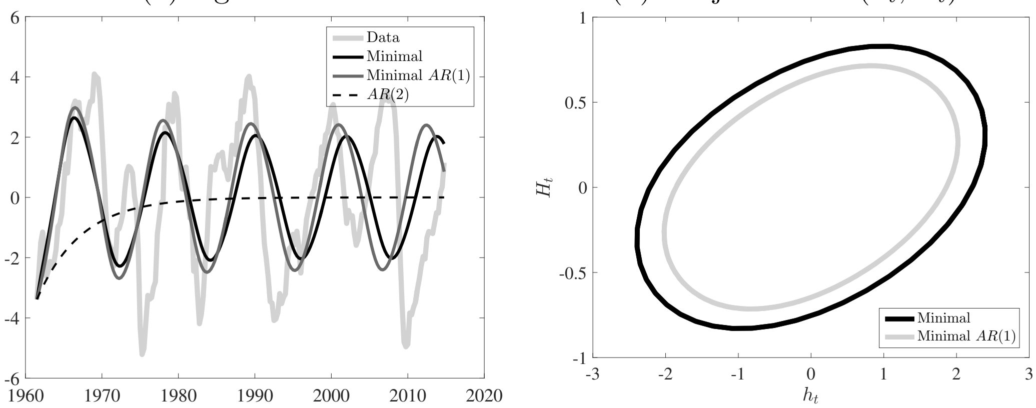

Notes: This corresponds to the deterministic simulation of the Minimal (13) *model and a version of the Minimal model with only one lag for h, that we refer to as the Minimal-*$AR(1)$ *model, starting (arbitrarily) in 1961Q3. The model has been estimated over 1960-2014 using high-passed-80-quarters filtered series of Total Hours.*

A.2 Minimal Model with MA Component

See Figure A.2. Quick comment is needed.

A.3 Contemporaneous Strategic Interactions

The econometric specifications of Section 2 assumed that agents formed their beliefs about others' actions one period in advance. In this section, we show that our results are not sensitive to this assumption by employing the strategic-interaction (SI) model of Beaudry, Galizia, and Portier [2015b], itself a dynamic version of Cooper and John [1988]. As will become clear, this model is very similar to the model of Section 2, with the key exception that the non-linearity in the dynamic system will appear in contemporaneous variables instead of lagged ones.

A.3.1 Basic set-up

Suppose agent j's expenditure is given by

$$x_{jt} = \alpha_1 x_{jt-1} + \alpha_2 x_{jt-2} - \alpha_3 X_{jt-1} + F\left(x_t\right) + \epsilon_t,$$

where F is a function satisfying $F' < 1$, $\alpha_j \in (0,1)$, $\epsilon t \sim$ i.i.d. $N\left(0, \sigma_\epsilon^2\right)$, and $X_{jt} = (1-\delta) X_{jt-1} + \delta x_{jt}$. In this decision rule, α_1 and α_2 impart sluggishness to agents'

Figure A.2: The Limit Cycle in the Minimal Model with MA Component, Total Hours

(a) Against time (b) Projection on (h_t, H_t)

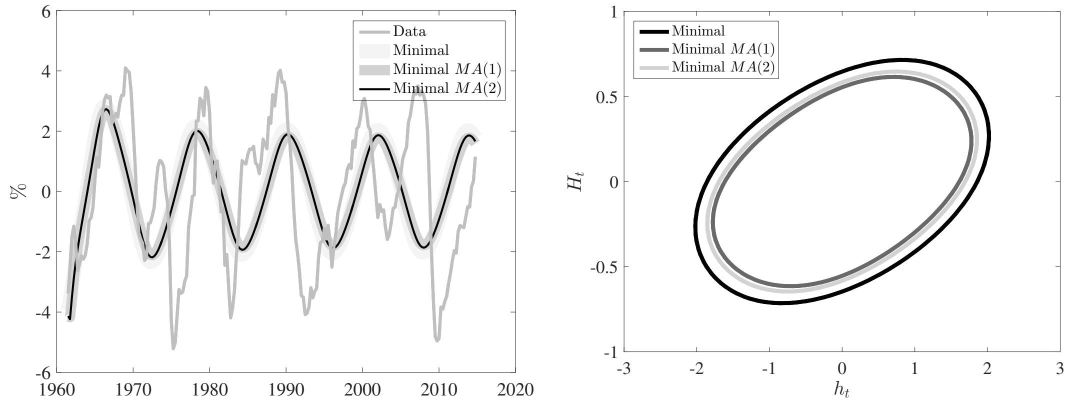

Notes: This corresponds to the deterministic simulation of the Minimal (13) model with no MA part or with a $MA(1)$ and $MA(2-)$ components , starting (arbitrarily) in 1961Q3. The model has been estimated over 196 0-2014 using high-passed-80-quarters filtered series of Total Hours.

decisions, the term X_{jt-1} governs the effect of decreasing returns to accumulation, and the function F, which may be non-linear, captures the contemporaneous strategic interaction among agents: if $F' > 0$, agents' expenditures are strategic complements, while if $F' < 0$ they are strategic substitutes. Aggregating this equation and letting $\phi(x) \equiv x - f(x)$, we have

$$\phi(x_t) = \alpha_1 x_{t-1} + \alpha_2 x_{t-2} - \alpha_3 X_{t-1} + \epsilon_t \tag{A.1}$$

Compared to the Minimal model (13) of Section 2, the only difference is that the non-linearity in (A.1) is in the date-t term (since ϕ is non-linear) rather than in a lagged term.

A.3.2 Estimation

To estimate this model, we must choose a functional form for F. Motivated by our earlier finding that we require F to be at least a third-order polynomial, we choose the simplest such function,

$$F(x) = bx - ax^3$$

with $b < 1$ and $a \geq 0$, where the latter assumption is required to ensure that the equilibrium value of x_t is always unique. We may then write (A.1) in reduced form as

$$x_t = g^{-1}\left(\beta_1 x_{t-1} + \beta_2 x_{t-2} - \beta_3 X_{t-1} + u_t\right) \tag{A.2}$$

where $\beta_j = \alpha_j/\left(1-b\right)$, $u_t \equiv \epsilon_t/\left(1-b\right) \sim$ i.i.d.$N\left(0,\sigma_u^2\right)$, $\sigma_u^2 \equiv \sigma_\eta^2/\left(1-b\right)^2$, $g\left(y\right) \equiv \gamma y^3 + y$, and $\gamma \equiv a/\left(1-b\right) \geq 0$. We may then estimate the β_j's, γ, and σ_u by maximum likelihood.

Table A.1 presents the results of the estimation where we use Total Hours for x_t which was first detrended using an 80-quarter high-pass filter.[24] The largest eigenvalue in modulus of the system (linearized around its steady state at zero) is 1.02, so that the steady state is unstable. Figure A.3 plots the deterministic forecast from 1961Q3 for the estimated model, both against time and projected onto (x_t, X_t), which verifies that the estimated SI model does indeed also produce a limit cycle.

Table A.1: Estimated Coefficients for SI Model, Total Hours

Parameter	**Value**
β_1	1.43
β_2	-0.39
β_3	0.28
γ	0.64e-02
σ_u	0.62

Notes: The Total Hours series has been filtered with an 80-quarter high-pass filter over the sample 1948Q1-2015Q2. Estimation is then done over the sample 1960Q1:2014Q4.

B Data

- U.S. Population: Total Population: All Ages including Armed Forces Overseas, obtained from the FRED database (`POP`) from 1952Q1 to 2015Q2. Quarters from 1947Q1 to 1952Q1 are obtained from linear interpolation of the annual series of National Population obtained from U.S. Census, where the levels have been adjusted so that the two series match in 1952Q1.

[24] The construction of X_t and the sample period are the same as those employed in Section 2.

Figure A.3: Deterministic Forecast for SI Model, Total Hours

(a) Against time (b) Projection on (x_t, X_t)

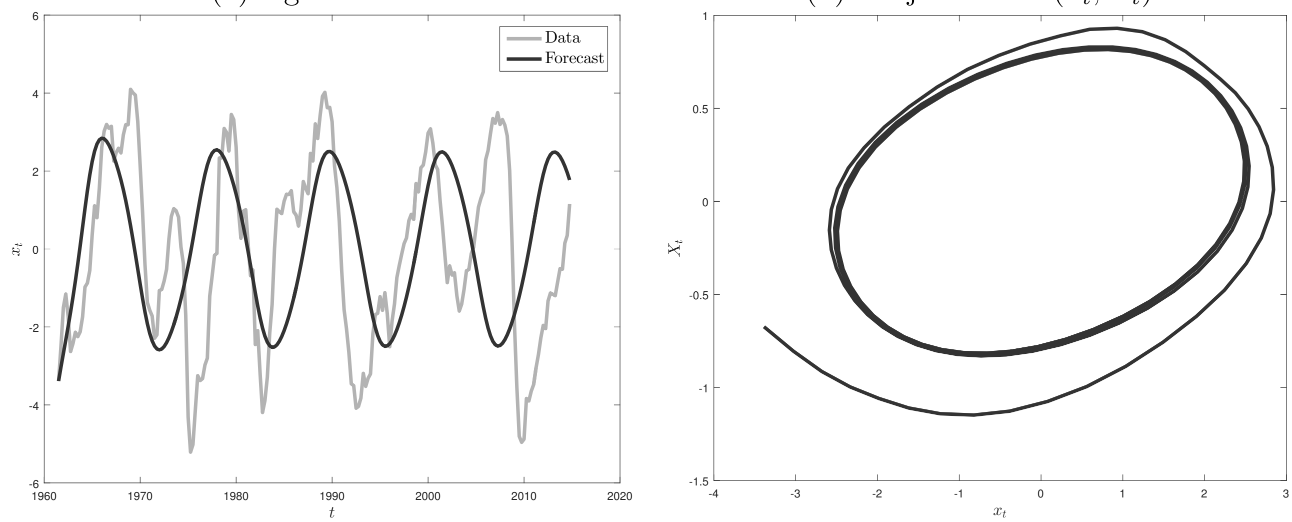

Notes: The Total Hours series has been filtered with an 80-quarter high-pass filter over the sample 1948Q1-2015Q2. Estimation is then done over the sample 1960Q1:2014Q4. Deterministic forecasts are done starting from 1961Q3 with the estimated SI (A.2) *model.*

- U.S. 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 1947Q1-2015Q2, and we do not use the observations of 2015.

- U.S. Non-Farm Business Output, Non-Farm Business Hours, Total Hours and unemployment rate (16 years and over) are obtained from the Bureau of Labor Statistics. Sample is 1947Q1-2015Q2 (1948Q1-2015Q2 for Total Hours), 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 1948Q1-2007Q4.

- The various unemployment rates are obtained from the FRED database, except for France:

  - Australia: Unemployment Rate: Aged 15 and Over: All Persons for Australia

(`LRUNTTTTAUQ156S`), 1966Q3-2014:Q4

– Austria: Registered Unemployment Rate for Austria (`LMUNRRTTATQ156S`), 1955Q1-2014Q4

– Canada: Harmonized Unemployment Rate: All Persons for Canada (`CANURHARMQDSMEI`), 1956Q1-2012Q1

– Denmark: Registered Unemployment Rate for Denmark (`LMUNRRTTDKQ156S`), 1970Q1-2014Q4

– France: ILO unemployment rate, Total, Metropolitan France and overseas departments (`001688527`), Inséé Macro-economic database (BDM), 1975Q1:2015Q1

– Germany: Registered Unemployment Rate for Germany (`LMUNRRTTDEQ156S`), 1969Q1-2014Q1

– Japan: Harmonized Unemployment Rate: All Persons for Japan (`JPNURHARMQDSMEI`), 1955Q1-2012Q1

– Netherlands: Harmonized Unemployment Rate: All Persons for Netherlands (`NLDURHARMQDSMEI` until 2012Q1 and Harmonized Unemployment: Total: All Persons for the Netherlands `LRHUTTTTNLQ156S`) after 2012Q1 (the second series level has been adjusted to match the first one in 2012Q1), 1970Q1-2014Q4

– Sweden: Harmonized Unemployment Rate: All Persons for Sweden (`SWEURHARMQDSMEI`), 1970Q1-2012Q1

– Switzerland: Registered Unemployment Rate for Switzerland (`LMUNRRTTCHQ156S`), 1970Q1-2014Q4

– United Kingdom: Registered Unemployment Rate for the United Kingdom (`LMUNRRTTGBQ156S`), 1956Q1:2012Q1

C Computing the Mean of the Square Root Sum of Squared k-step-ahead Forecast Errors for the Minimal Model

Here we describe the way we compute the mean of the square root sum of squared k-step-ahead forecast errors for the Minimal model, as it appears on Figure 7.

Take a date T_0 and compute a stochastic forecast, indexed by s, by simulating forward the Minimal model, :

$$\begin{cases} \widetilde{h}^G_{T_0}(T_0,s) & = & h_{T_0}, & \\ \widetilde{h}^G_{T_0-1}(T_0,s) & = & h_{T_0-1}, & \\ \widetilde{H}^G_{T_0-1}(T_0,s) & = & H_{T_0-1}, & \\ \widetilde{h}^G_t(T_0,s) & = & \widehat{\beta}_h \; \widetilde{h}^G_{t-1}(T_0,s) + \widehat{\beta}_{h'} \; \widetilde{h}^G_{t-2}(T_0,s) + \widehat{\beta}_H \; \widetilde{H}^G_{t-1}(T_0,s) + \widehat{\sigma}\nu^s_t & \forall\ t > T_0, \\ \widetilde{H}^G_t(T_0,s) & = & (1-\delta)\widetilde{H}^G_{t-1}(T_0,s) + \delta \widetilde{h}^{G(T_0,s)}_t & \forall\ t \geq T_0, \end{cases}$$

where $\widehat{\sigma}$ is the estimated standard deviation of ϵ in the Minimal model and ν^s is drawn from a $\mathcal{N}(0,1)$ distribution. Now define the k-step-ahead forecast when the information set is the one of T_0 and for the simulation s as

$$\mathcal{E}_k(T_0,s) = \widetilde{h}^G_{T_0+k}(T_0,s) - h_{T_0+k}$$

Denote S the number of simulations for each initial condition T_0, and T the last date in the sample. The mean of the square root sum of squared k-step-ahead forecast errors, denoted $\mathcal{M}_k$ is then given by

$$\mathcal{M}_k = \frac{1}{T-k} \sum_{T_0=1}^{T-k} \left(\sqrt{\frac{1}{S} \sum_{s=1}^{S} \mathcal{E}_k(T_0,s)^2} \right)$$

In our computation, $S = 1000$.

D Proof of Proposition 1

Assume that σ^2_ϵ is small, so that $|x_0|$ is close to $\bar{x}$. Without loss of generality, take x_0 close to $\bar{x}$,[25] 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^2_ϵ is small, we should have $x_t > 0$ for t even and $x_t < 0$ for t odd, i.e., the sign of x_t should switch every period.[26] Thus, $|y_t|$ captures the absolute deviation of x_t from the non-stochastic sequence $\bar{x}_t = (-1)^t \, \bar{x}$ (the 2-cycle), while $\operatorname{sgn}(y_t)$ captures the direction of that deviation relative to zero: if $y_t < 0$ then x_t is "inside" the 2-cycle (i.e., $|x_t| < \bar{x}$), while if $y_t > 0$ then x_t is "outside" the 2-cycle (i.e., $|x_t| > \bar{x}$).

[25] It is straightforward to extend the following reasoning to the case where x_0 is close to $-\bar{x}$.

[26] This stems from the fact that, with x_0 close to $\bar{x}$ and σ^2_ϵ small, the dynamic behavior of x_t is dominated by the same forces that generate the 2-cycle.

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 $\nu_t \equiv (-1)^t\, \epsilon_t \sim$ i.i.d. $N\left(0, \sigma_\epsilon^2\right)$. 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,

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

y_t is clearly stationary, so that we have

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

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

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

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.