

Reviving the Limit Cycle View of Macroeconomic Fluctuations

Paul Beaudry*, Dana Galizia† and Franck Portier‡

March 2015
Preliminary

Abstract

Limit cycles play very little role in the modern theory of business cycles. In this paper we present a general structure, a particular model and a DSGE extension with the hope of giving new life to this mostly dismissed view of fluctuations. In particular, we begin by showing why models with demand complementarities, which are quite ubiquitous in macroeconomics, can give rise to limit cycle under rather simple conditions. The key elements that induce limit cycles are the presence of a demand complementarity related to an accumulable good and that the demand complementarities be locally strong but globally limited. We then present a fully specified model dynamic equilibrium model where unemployment risk and precautionary savings interact to induce a limit cycle driven by the type of forces outlined in our general structure. In this model, the economy goes through expansion periods where agents accumulate durable goods and/or housing more quickly when unemployment is low because borrowing is perceived as being less risky. Subsequently the economy goes through a recession phase when agents pull back on their purchases because unemployment increases thereby causing borrowing to become more risky. While this process is shown to be recurring and compatible with rational expectation, it does not involve any indeterminacy: fundamentals fully pin-down equilibrium behavior. We then show that in macroeconomic U.S. data, there is non trivial variability at frequencies which are lower than the ones traditionally looked at by the business cycle literature, that could well be interpreted as a medium frequency limit cycle perturbed with shocks. We complete the analysis by showing how a DSGE version of our model can be estimated to have a limit cycle, and how it can shed new light on interesting features of the data once it is augmented to include some randomness.

Key Words: Business Cycle, Unemployment, Limit Cycle; JEL Class.: E3, E32, E24

*Vancouver School of Economics, University of British Columbia and NBER

†Vancouver School of Economics, University of British Columbia

‡Toulouse School of Economics, Université Toulouse 1 Capitole and CEPR.

Introduction

In most modern business cycle models, the underlying economic system is very stable. In particular, in the absence of shocks, the variables in these systems tend to converge either to a steady state or to a balanced growth path. In such frameworks business cycles are viewed as emerging from shocks – which could be either fundamental or non-fundamental – that disturb an otherwise stable system. However, it is well known that this is not the only framework that could give rise to business cycles. In particular, it may well be that economic forces naturally give rise to cyclical phenomena, that is, in the absence of any shocks (including belief shocks) economic forces by themselves may favor recurrent periods of high economic activity followed by periods of low economic activity. This type of outcome will arise for example if the underlying economic system is characterized by a limit cycle. In such a framework irregular business cycles can emerge from these underlying regular forces combined with shocks that move the system away from an attracting orbit.

The idea of self-sustaining *trade cycles*, to use the terminology of the early years of macroeconomics, can be found, although not explicitly formalized, in the dynamic version of the Keynesian theory proposed by Kalecki [1937], and more formally later on in the nonlinear versions of Samuelson’s [1939] accelerator proposed by Kaldor [1940], Hicks [1950] and Goodwin [1951]¹. Those first models were not microfounded general equilibrium models. In the 1970s and 1980s, a large literature emerged that examined the conditions under which qualitatively and quantitatively reasonable economic fluctuations might occur in a purely deterministic setting (see, e.g., Benhabib and Nishimura [1979], [1985], Day [1982], [1983], Grandmont [1985], Boldrin and Montrucchio [1986], Day and Shafer [1987]; for surveys of the literature, see Boldrin and Woodford [1990] and Scheinkman [1990]). By the early 1990s, however, this literature seemed to have largely gone dormant.

There appear to be several key reasons why interest in deterministic fluctuations may have waned, each of which are addressed in the present paper. First, the earlier literature on deterministic fluctuations can be broadly sub-divided into two categories: models with and without fully-microfounded, forward-looking agents.² The latter category, which were generally more capable of producing reasonable deterministic fluctuations than the former, likely fell out of favor as macro in general moved toward more microfounded models.

Second, in the category of models featuring forward-looking agents, the primary focus was on models with a neoclassical, competitive-equilibrium structure.³ Such models were often found to require relatively extreme parameter values in order to generate deterministic fluctuations. For example, the Turnpike Theorem of Scheinkman [1976] establishes that,

¹An earlier mention of self-sustained cycles as an interesting modeling of economic fluctuations is found in Le Corbeiller [1933], in the first volume of *Econometrica*.

²The first category includes, e.g., Benhabib and Nishimura [1979], [1985] and Boldrin and Montrucchio [1986], while the latter includes, e.g., Day [1982], [1983].

³While there are some exceptions, they are comparatively rare. Perhaps the clearest example is Hamour [1989], chapter 1, which is focused on deterministic fluctuations in an environment of increasing returns. Other exceptions include models in the search literature that are capable of generating deterministic fluctuations, such as Diamond and Fudenberg [1989], Boldrin, Kiyotaki, and Wright [1993], and Coles and Wright [1998]. Note however that these search papers were mainly concerned with characterizing the set of possible equilibria for a particular model (which for some parameterizations included deterministic cycles), rather than being focused on deterministic cycles directly.

under certain basic conditions met by these models, for a sufficiently high discount factor—i.e., for agents that are “forward-looking” enough—the steady state of the model is globally attractive, so that persistent deterministic fluctuations cannot appear.⁴ While in principle this does not rule out deterministic fluctuations completely, in practice the size of the discount factor needed to generate them was often implausibly low. For example, in a survey of deterministic-fluctuations models by Boldrin and Woodford [1990], discount factors for several of the models they discuss were on the order of 0.3 or less.⁵ As the present paper illustrates, however, if one departs from the assumptions of a neoclassical, competitive-equilibrium environment—for example, if there is a demand externality as in the model presented in section 2—then a discount factor arbitrarily close to one can relatively easily support deterministic fluctuations in equilibrium.

Third, as suggested above, models producing periodic cycles—that is, cycles which exactly repeat themselves every k periods—are clearly at odds with the data, where such consistently regular cycles cannot be found.⁶ This can be observed by looking at the spectrum of data generated by such a model, which will generally feature one or more large spikes at frequencies associated with k -period cycles. Spectra estimated on actual data generally lack such spikes,⁷ which suggests less regularity in real-world cycles. To address this issue, papers from the earlier literature largely sought to establish conditions under which such irregular cycles could emerge in a purely deterministic setting (i.e., via chaotic dynamics⁸). While in a number of cases this was found to be possible, the conditions appear to have been significantly more restrictive even than those required to generate simple periodic cycles. In contrast, rather than restricting attention to a purely deterministic setting, this paper embeds deterministic (but highly regular) cyclical mechanisms into a stochastic environment for which irregularity emerges naturally.

Finally, being inherently highly non-linear, economic models that are capable of generating deterministic fluctuations are often difficult to work with analytically beyond the very simplest of settings, and quantitative results often require computationally-expensive solution algorithms. Prior to relatively recent advances in computing technology, obtaining these quantitative results may have been infeasible and, as a result, a number of potentially fruitful areas of research—such as, for example, combining deterministic and stochastic cyclical forces—may have gone unexplored.

In this paper, we re-examine the issue of limit cycles as a foundation to a theory of business cycles by building on models with demand externalities. Our first goal is to show that

⁴See the discussion in section 2.3 for further details.

⁵It is possible in principle to rationalize such low discount factors by choosing a longer period length for the model. However, if households discount the future with a quarterly discount factor of 0.99 or greater—as is frequently the case in the business-cycle literature—a factor of 0.3 would be associated with a period length of 120+ quarters (30+ years). Since the minimum period length of a cycle is two periods, this would generate cycles on the order of 60+ years, well outside of what is normally thought of as the business cycle.

⁶In Boldrin and Woodford [1990], the authors mention that in private communication with Sir John Hicks, he has indicated that the fact that actual business cycles are far from being regular periodic motions was the reason for his loss of interest in endogenous cycle models.

⁷See Figure 14.

⁸Informally, chaotic fluctuations are deterministic fluctuations that do not converge to periodic cycles and for which the paths emanating from two different initial points cannot be made arbitrarily close by choosing those initial points sufficiently close together. See, e.g., Glendinning [1994] for a formal definition.

limit cycles tend to arise quite naturally in the presence of demand complementarity. In particular, we will clarify why demand complementarities, when they relate to an accumulable good, are more likely give rise to local instability and limit cycles as opposed to multiple equilibrium. To make this point we present a simple but rather general reduced form setup. Our second goal is to present a specific, but fully specified model, where limit cycles arise as the result of an interplay between unemployment risk and precautionary savings. We use a DSGE version of this model to show how it can explain interesting features of the data.

The outline of the paper is as follows: In the first section we present our reduced form dynamic model with demand externalities to highlight the forces that will give rise to limit cycles. In particular, we show how when a Hopf bifurcation will arise in such an environment and when it will give rise to a super-critical limit cycle. In the second section, we introduce a microfounded model that builds on Beaudry, Galizia, and Portier [2014]. We will use this setting to illustrate an intuitive economic mechanism based on demand complementarities that can give rise to limit cycles. The third section will be quantitative. We will first show that in macroeconomic U.S. data, there is non trivial variability at frequencies which are lower than the ones traditionally looked at by the business cycle literature, that could well be interpreted as a medium frequency limit cycle perturbed with shocks. We then present a fully specified dynamic general equilibrium model that incorporates our reduced form model as a special case. Using our structural model we show when how conditions for limit cycles are likely to met in an explicit dynamic general equilibrium environment. Moreover, in the fully specified model we will be able to explore the extent to which the model is capable for reproducing features of the data. In the last section we offer concluding comments.

1 Demand complementarities and limit cycles: a reduced form model

In this section we present a simple reduced form model aimed at illustrating how and when limit cycles emerge in environment with demand complementarities. We choose to begin by presenting this reduced form model as to highlight the generality of the mechanism, regardless of the source of precise micro-foundation for the demand complementarity. In the later sections we will present a structural model that can rationalize our reduced form setup. An important aspect of the analysis is to show that, even when we restrict the strength of the demand complementarities to be too weak to create static multiple equilibrium, the model will nonetheless give rise to limit cycles under fairly general conditions. We will see that the key mechanism generating limit cycles is the interplay of the demand complementarities and the dynamics associated with the fact that the good can be accumulated. That framework is more likely to give rise to limit cycles than to the type of multiple equilibria often stressed in the demand complementarity literature since the strength of the demand complementarity needed for limit cycles is less than that needed for multiple equilibrium.

1.1 The environment

The environment we want to consider is one with a large number N of agents index by i , where each agent can accumulate a good X_{it} and where X_{it} can represent either physical

capital or a durable consumption good. The accumulation equation is given by

$$X_{it+1} = (1 - \delta)X_{it} + I_{it} \quad 0 < \delta < 1 \quad (1)$$

where I_{it} is agents i 's investment in the good. Suppose initially that there is no market interactions and that the decision rule for agent i 's investment is given by

$$I_{it} = \alpha_0 - \alpha_1 X_{it} + \alpha_2 I_{it-1} \quad (2)$$

where all parameters are positive and $0 < \alpha_1 < 1$, $0 < \alpha_2 < 1$. In this decision rule, the effect of X_{it} on investment is assumed to be negative as to reflect some underlying decreasing returns to capital accumulation and the effect of past investment is positive as to reflect sluggish response due for example to adjustment costs.

When all agent behave symmetrically, The aggregate dynamics of the economy is given by the linear system:

$$\begin{pmatrix} I_t \\ X_t \end{pmatrix} = \underbrace{\begin{pmatrix} \alpha_2 - \alpha_1 & -\alpha_1(1 - \delta) \\ 1 & 1 - \delta \end{pmatrix}}_{M_L} \begin{pmatrix} I_{t-1} \\ X_{t-1} \end{pmatrix} + \begin{pmatrix} \alpha_0 \\ 0 \end{pmatrix} \quad (3)$$

The stability property of that system is established in the following proposition:

Proposition 1 *Both eigenvalues of the matrix M_L lie within the unit circle. Therefore, the system is stable.*

The proof of this proposition (as well as the other ones) is given in the appendix. According to proposition 1, the dynamics is extreme simple, with the system converging to it steady state for any starting values of $X_{it} = X_t$ and $I_{it-1} = I_{t-1}$. We now add demand complementarities to the model and study how is the dynamics affected.

1.2 Adding demand complementarities


We extend the previous setup to allow for interactions between individuals by having the investment rule be given instead

$$I_{it} = \alpha_0 - \alpha_1 X_{it} + \alpha_2 I_{it-1} + F\left(\frac{\sum I_{jt}}{N}\right) \quad (4)$$

while the law of motion of X (equation (1)) is kept unchanged. We assume that the function $F(\cdot)$ is continuous and differentiable at least three times, that $F(0) = 0$ and that $F'(\cdot) < 1$. That function represents how the actions of others, summarized but the average level of investment $\frac{\sum I_{jt}}{N}$, affect agent i investment decision I_{it} . For example, the function $F(\cdot)$ could capture price effects of inputs if $F'(\cdot) < 0$ or can capture demand complementarities if $F'(\cdot) > 0$. In this formulation we are assuming that agents are taking the average actions in the economy as given, so that (4) can be interpreted as a best response rule of an individual to the average action.


Figure 2 illustrates this best response rule for two values of the intercept, one where the equilibrium determination of investment arises with $F'\left(\frac{\sum I_{jt}}{N}\right) < 0$, and one that arises

Figure 1: Phase diagram in the model without demand complementarities



This figure represent the phase diagram of the dynamic system (3), in which there are no demand complementarities. For the sake of exposition, we have arbitrary assumed that both eigenvalues were real and positive in this graph. Note that the use of a phase diagram is only illustrative; as we are working in discrete time, the system could converge with oscillation when close to the steady state even in that configuration of the eigenvalues.

Figure 2: Best response rule for two different histories



This figure plots the best response rule (equation (4)): $I_{it} = \alpha_t + F\left(\frac{\sum I_{jt}}{N}\right)$, with $\alpha_t = \alpha_0 - \alpha_1 \sum_0^\infty (1 - \delta)^\tau I_{it-1-\tau} + \alpha_2 I_{it-1}$. The intercepts α_t and $\alpha_{t'}$ correspond to two different histories of the model.

with $0 < F' \left(\frac{\sum I_{jt}}{N} \right) < 1$. Note that in this diagram we have represented the intercept by $\alpha_t = \alpha_0 - \alpha_1 \sum_j^\infty (1 - \delta)^j I_{it-1-j} + \alpha_2 I_{it-1}$ as to make clear its dependence on past investment behavior. The dynamics of the system is induced by the fact that past investment decision determine the location of the best response rule, which in turn determines the current level of investment. Therefore today's decision feeds into the determination of next period's intercept.⁹ In order to rule out static multiple equilibrium, that is multiple solutions for I_{it} for given values of I_{it-1} and X_{it} , we assume that $F' \left(\frac{\sum I_{jt}}{N} \right) < 1$. Hence, we are restricting attention to cases where demand complementarities, if they are present, are not strong enough to produce static multiple equilibrium.

In the following, we will restrict to symmetrical equilibria, allowing us to drop the subscript i . In order to ensure the existence of a steady state, we will assume that for large enough values of I , the interaction effect becomes negative, that is, $\lim_{I \rightarrow \infty} F(I) < 0$. Combining this assumption with the restrictions that the slope of $F(\cdot)$ be always less than 1 and that $F(0) = 0$ ensures the existence of a unique steady state. We denote the steady state value of I and X by I^s and X^s . The condition $\lim_{I \rightarrow \infty} F(I) < 0$ is reasonable in most economic environments due to resource constraints. In fact, it may be appropriate in many environments to impose a stronger condition of the form $\lim_{I \rightarrow \bar{I}} F(I) = -\infty$ which imposes the upper bound \bar{I} on the feasible level of investment in this system. Note that when $F'(I^s) = 0$, the model dynamics is locally the one of the model without demand complementarities, and is therefore stable.

Our goal is to use this framework to examine how the dynamics of this system are affected by the properties of the interaction effects, and especially under what conditions on $F(\cdot)$ will give rise to limit cycles.

1.3 The local dynamics of the model with demand externalities

We now consider the dynamics given by (1) and (4). The first order approximation of this dynamic system is given by

$$\begin{aligned} \begin{pmatrix} I_t \\ X_t \end{pmatrix} &= \underbrace{\begin{pmatrix} \frac{\alpha_2 - \alpha_1}{1 - F'(I^s)} & -\frac{\alpha_1(1 - \delta)}{1 - F'(I^s)} \\ 1 & 1 - \delta \end{pmatrix}}_M \begin{pmatrix} I_{t-1} \\ X_{t-1} \end{pmatrix} \\ &+ \begin{pmatrix} \left(1 - \frac{\alpha_2 - \alpha_1}{1 - F'(I^s)}\right) I^s + \left(\frac{\alpha_1(1 - \delta)}{1 - F'(I^s)}\right) X^s \\ 0 \end{pmatrix} \end{aligned} \quad (5)$$

In order to understand the dynamics of this system, it is informative to first look at local dynamics in the neighborhood of the steady state. The eigenvalues of the M matrix are the solutions of the second order equation given by

$$Q(\lambda) = \lambda^2 - T\lambda + D = 0 \quad (6)$$

⁹Under the condition $F' \left(\frac{\sum I_{jt}}{N} \right) < 1$, the static equilibrium depicted in Figure 2 is generally viewed as stable under a *tâtonnement* type adjustment process. This stability property is not the focus of the current paper. Instead we are interested in the explicit dynamics induced by the system (1) and (4).

where T is the trace of the M matrix (and also the sum of its eigenvalues) and D is the determinant of the M matrix (and also the product of its eigenvalues). The two eigenvalues are

$$\lambda, \bar{\lambda} = \frac{T}{2} \pm \sqrt{\left(\frac{T}{2}\right)^2 - D} \quad (7)$$

where

$$T = \left(\frac{\alpha_2 - \alpha_1}{1 - F'(I^s)} + (1 - \delta) \right) \quad (8)$$

and

$$D = \frac{\alpha_2((1 - \delta))}{1 - F'(I^s)}. \quad (9)$$

As noted above, when there is no social interaction ($F'(I^s) = 0$), the roots of this system lie within the unit circle. Now consider that the F function is parametrized by $F'(I^s)$, that is allowed to go from $-\infty$ to 1.



Geometric analysis : It is useful to give a geometric analysis of the location of the two eigenvalues of the system. This is done in Figure 3, which represents the plane (T, D) . A point in this plane is couple (Trace, Determinant) of matrix M that corresponds to a particular configuration of the model parameters, including $F'(I^s)$. We have drawn three lines and a parabola in that plane. The line (B, C) corresponds to $D = 1$, the line (B, A) to the equation $Q(-1) = 0$ ($\Leftrightarrow D = -T - 1$) and the line (A, C) to the equation $Q(1) = 1$ ($\Leftrightarrow D = T - 1$). On the triangle \widehat{ABC} perimeter, at least one eigenvalue has a modulus of 1. We show in the appendix that both eigenvalues of the system are inside the unit circle when (T, D) is inside the triangle \widehat{ABC} , while at least one is outside the unit circle when (T, D) is outside the triangle \widehat{ABC} . Whether eigenvalues are really complex depends on the sign of the discriminant of the equation $Q(\lambda) = 0$, which is given by $\Delta = \left(\frac{T}{2}\right)^2 - D$. The parabola on Figure 3 correspond to the equation $\Delta = 0$ (which is $D = T^2/4$). Above the parabola, eigenvalues are complex and conjugate while they are real below. It is then possible to picture the possible configurations of local dynamics. As shown on Figure 3, the steady state can be locally stable, unstable or a saddle, with real or complex eigenvalues. Proposition 1 proves that when $F'(I^s) = 0$, the steady state corresponds to a point E that is inside the triangle \widehat{ABC} . As an example, we have put E in the region of stability with complex eigenvalues in Figure 3.

As $F'(I^s)$ varies, the eigenvalues of the system will vary, implying changes to the dynamic behavior of I and X . From equations (8) and (9), we obtain the following relation between the trace and determinant of matrix M :

$$D = \frac{\alpha_2(1 - \delta)}{\alpha_2 - \alpha_1} T - \frac{\alpha_2(1 - \delta)^2}{\alpha_2 - \alpha_1} \quad (10)$$

Therefore, when $F'(I^s)$ varies, T and D move along the line (10) in the plane (T, D) , which allows for an easy characterization of the impact of $F'(I^s)$ on the location of the eigenvalues, and therefore the stability of the steady state. We need to systematically consider the two

Figure 3: Possible configurations of the local dynamics



This figure shows the plane (T, D) , where T is the trace and D the determinant of matrix M . The point E correspond to the model without demand externalities. It is arbitrarily placed in the zone where the two eigenvalues are complex and of modulus smaller than 1. According to proposition 1, E belongs to the interior of the triangle \widehat{ABC} .

cases $\alpha_2 > \alpha_1$ and $\alpha_2 < \alpha_1$, as the line (10) slopes positively in the former case and negatively in the later.¹⁰

Let us consider the case where $F'(I^S)$ is negative, that is where the investment decisions of others have a negative effect on one own decision. In that case, it is clear from equations (8) and (9) that

$$\lim_{F'(I^S) \rightarrow -\infty} T = 1 - \delta$$


and

$$\lim_{F'(I^S) \rightarrow -\infty} D = 0$$

which corresponds to E_1 on Figure 4 (when $\alpha_2 > \alpha_1$) and on Figure 5 (when $\alpha_2 < \alpha_1$). Note that E_1 belongs to the triangle \widehat{ABC} , so that the steady state is locally stable in this case.

When $F'(I^S)$ goes from 0 to $-\infty$, the economy moves from E to E_1 following the line (10). This movement corresponds to the half-line denoted (a) on Figures 4 and 5. As E and E_1 belong to the interior of the triangle \widehat{ABC} and because the interior of the triangle \widehat{ABC} is a convex set, any point of the segment $[E, E_1]$ also belongs to the interior of the triangle \widehat{ABC} .

Figure 4: Local stability when $F'(I^S) \in]-\infty, 1]$ and $\alpha_2 > \alpha_1$




See Figure 3 for the legend.

Therefore, the following proposition holds:

¹⁰The case $\alpha_1 = \alpha_2$ is studied in the appendix.

Figure 5: Local stability when $F'(I^s) \in]-\infty, 1]$ and $\alpha_2 < \alpha_1$



See Figure 3 for the legend.

Proposition 2 *As $F'(I^S)$ varies from 0 to $-\infty$, the eigenvalues of M always stay within the unit circle and therefore the system remains locally stable.*

Proposition 2 indicates that when the actions of others play the role of strategic substitutes with one own action, this favors stability of the system. A strategic substitute property is one that often arises in Walrasian setting, and this is one reason why dynamic Walrasian environment are most often stable.

We now turn to exploring how the presence strategic complementarities affect the dynamic of the system. Considering Figures 4 and 5 an assuming that the steady state without demand externalities correspond to E , the economy will take the path (b) moving away from E as $F'(I^s)$ gets closer to one. As we see it in Figures 4 and 5, the economy will cross the perimeter of triangle \widehat{ABC} : the steady state will change from being locally stable to being unstable or saddle. In the case of half-line (b), the steady state will switch from being stable with complex eigenvalues to being unstable with complex eigenvalue again. This is one of two possible configurations. Parameters can be such that the steady state is in E' rather than in E . In that case, as shown in Figures 4 and 5, the steady state will switch from being stable with real eigenvalues to being unstable with one real eigenvalue greater than one and one smaller than one. If $\alpha_2 > \alpha_1$, both eigenvalues will be positive, while they will be negative if $\alpha_2 < \alpha_1$.

This shows that the presence of demand complementarities can radically change the dynamics of the system. In particular, if the demand complementarity is strong enough, it will cause the system to be unstable, even though the system would be stable in the absence of the complementarity. The instability arises because agents have an incentive to bunch their actions while simultaneous having an incentive exhibit some inertia due to the fact that $\alpha_2 > 0$.

1.4 Bifurcations and the occurrence of limit cycles

Such a change of local stability when a parameter varies is referred to as a bifurcation in the theory of dynamical systems. It is of particular interest since in the case of a bifurcation, limit cycles can occur in the global dynamics of the system. We formally state the occurrence of a bifurcation in proposition 3.

Proposition 3 *As $F'(I^s)$ varies from 0 towards 1, the dynamic system given by (1) and (4) will become unstable and will experience*

- a Neimark-Sacker or Hopf bifurcation if $\alpha_2 \in \left[\frac{\alpha_1}{(2-\delta)^2}, \frac{\alpha_1}{\delta^2} \right]$,
- a flip bifurcation if $\alpha_2 < \frac{\alpha_1}{(2-\delta)^2}$,
- a fold or saddle-node bifurcation if $\alpha_2 > \frac{\alpha_1}{\delta^2}$

A fold bifurcation occurs with the appearance of an eigenvalue equal to 1, a flip bifurcation with the appearance of an eigenvalue equal to -1 and a Hopf bifurcation with the appearance of two complex conjugate eigenvalues of modulus 1. The proof of the proposition is given in the appendix, and consist in checking conditions for the the economy to be in the configuration E or E' . In the case of flip and Hopf bifurcation, the system dynamics will

exhibit limit cycles. In the case of the flip bifurcation, that limit cycle will be of order two, oscillating around the steady state. Such extreme fluctuations are unlikely to be relevant for macroeconomic analysis, as fluctuations are persistent. The most interesting case from our point of view is the case where the system experience a Hopf bifurcation, as in such a case the steady state is surrounded by an isolated closed invariant curve that is unique. Contrarily to the case of the flip bifurcation, it is possible for such a limit cycle to evolve smoothly with time.

The conditions on α_2 stated in Proposition 3 under which a Hopf bifurcation will arise may at first past look rather restrictive, or at least hard to grasp. These conditions are quite restrictive when δ is very large, however they are not if δ is small. In fact, as δ approaches zero, these conditions reduce to the simple condition that $\alpha_2 > \frac{\alpha_1}{4}$, that is α_2 can't be too small. In other words, Proposition 2 could be loosely re-stated as indicating that if depreciation is not too fast, and α_2 not too small, then the system will experience a Hopf bifurcation as $F'(I^S)$ increases from 0 towards 1.

Once the existence of limit cycle is obtained, we need to check that the economy is indeed converging towards such an orbit. This orbit can be attractive, when the Hopf bifurcation is supercritical, or repulsive, when the Hopf bifurcation is subcritical. As shown on Figure 6, when the bifurcation is super critical, all orbits starting outside or inside the closed invariant curve, except at the origin, tend to the curve. The alternative case is shown on Figure 7, where the orbits that do not start from the limit cycle move away from it.

From our point of view, the emergence of a limit cycle is only of interests if it is attractive (supercritical), as in such a case, departures from the steady state will be attracted to the limit cycle. The conditions governing whether a Hopf bifurcation is supercritical or subcritical are often hard to state. However, in our setup, a simple condition can be given to ensure that the Hopf bifurcation is supercritical. This is stated in Proposition 4 where we make use of the Wan theorem..


Proposition 4 *If $F'''(I^S)$ is sufficiency negative, then the Hopf bifurcation noted in Proposition 3 will be supercritical. Therefore, the limit cycle is attractive.*

The economics for why increasing $F'(I^S)$ will cause the system to become unstable are rather clear. A high value for $F'(I^S)$ implies that agents have an incentive to accumulate all the same time. Hence people will make decision which will cause the system to go through periods of high accumulation followed by periods of low accumulation, and this will be recurrent even in the absence of any shocks. Such a behavior contrasts a steady flow of I over time which would be the natural point of rest of the system in the absence of complementarities.

The requirement that $F'''(I^S)$ be sufficiently negative for the emergence of a attractive limit cycle can also be associated with economic forces. The easiest way to interpret this condition is in terms of an S-shaped reaction function. If the reaction curve near the steady state is both positively sloped and $F'''(I^S)$ is negative, it will generally take an S-shaped form.¹¹ Note that Figure 2 was drawn with this features. The intuition for why an


¹¹A parametric example of such an S-shaped function is the sigmoid function $f(x) = \frac{1}{1+e^{-x}}$ for x on the real line.

Figure 6: An attractive limit cycle when the Hopf bifurcation is supercritical



This figure illustrates the model dynamics when the Hopf bifurcation is supercritical. The dark orbit corresponds to the limit cycle and the grey dot is the steady state of the model. Trajectories converge to the limit cycle from outside and inside.

Figure 7: An repulsive limit cycle when the Hopf bifurcation is subcritical



This figure illustrates the model dynamics when the Hopf bifurcation is subcritical. The dark orbit corresponds to the limit cycle and the grey dot is the steady state of the model. Trajectories diverges away from the limit cycle from outside and inside.


S-shaped reaction function will likely create a supercritical limit cycles is that, when the system moves into the unstable region, any perturbation from the steady state will induce explosive behavior because the strong demand complementarities close to the steady state. However, as the system gets far enough away from the steady state, the S-shaped property of the reaction function implies that the strength of the demand complementarities are fading out. Hence, as it system moves away from the steady state, the forces that were creating explosive behavior die out and instead start to favor implosive behavior. This is why an S-shaped reaction function, that is one where $F'''(I^s)$ is sufficiently negative, will favor the emergence of a supercritical limit cycle, as the limit cycle is balancing out the unstable force near the steady state, with the more stabilizing forces away from the steady state. Such a configuration is illustrated on Figure 8.

If instead, agents' reaction functions have $F'''(I^s)$ positive, then the reaction function would take the form of a flipped S, with the demand complementarities growing in strength as one moves away from the steady state.¹² This configuration is illustrated in Figure 9. In this case, when $F'(I^s)$ pushes the system to instability, the Wan theorem (reference) about Hopf bifurcations in discrete systems implies the emergence of a subcritical limit cycle, according to the dynamics shown in Figure 7.

The general insight we take away from the Wan theorem regarding Hopf bifurcations is that limit cycles are likely to emerge in our setting if demand complementarities are strong and create instability near the steady state, but tend to die out as one moves away from the steady state. We will refer to such a set up as one with strong local demand complementarities. In an economic environment, it is quite reasonable to expect that positive demand externalities are likely to die out if activity gets very large. For example, if investment demand gets sufficient large some resource constraint are likely to become binding, causing incentives akin to strategic substitutes to emerge instead of complementarities. Similarly, physical constraints, such as a non-negatively restrictions on investment or capital, are reasonable consideration in economic environments which will limit systems from diverging to zero or negative activity. Such forces will general favor the emergence of attractive limit cycles in the presence of demand complementarities. Since the Wan Theorem relates only to local behavior, it would be interesting to also examine global conditions that would ensure that the limit cycle is not nan arbitrarily small neighborhood of the steady state. This type of global analysis is much easier to carry out in a continuous time set up. Accordingly, in the appendix, we discuss global stability issues in a continuous time extension of our system (1) and (4). As is shown in this appendix, global forces that ensure that complementarities die out when activity is either very high or very low will guarantee the emergence of limit cycles in the continuous time version of the system when strong demand complementarities are present near the steady state.


¹²A parametric example of such a flipped S-shaped function is the logit function, which is the reciprocal of the sigmoid, and takes the form $g(y) = \log\left(\frac{y}{1-y}\right)$ for $y \in]0, 1[$.

Figure 8: Best response rule when $F'''(I^s) < 0$, causing the limit cycle to be supercritical



This figure plots the best response rule (equation (4)): $I_{it} = \alpha^s + F\left(\frac{\sum I_{jt}}{N}\right)$, with $\alpha^s = \alpha_0 - \alpha_1 X^s + \alpha_2 I^s$. The intercepts α^s corresponds to an economy that is at the steady state.

Figure 9: Best response rule in the case where $F'''(I^s) > 0$, causing the limit cycle to be subcritical



This figure plots the best response rule (equation (4)): $I_{it} = \alpha^s + F\left(\frac{\sum I_{jt}}{N}\right)$, with $\alpha^s = \alpha_0 - \alpha_1 X^s + \alpha_2 I^s$. The intercepts α^s corresponds to an economy that is at the steady state.

1.5 Numerical simulation

We conclude this section by sowing some numerical simulations of the model with demand complementarities. The objective here is not to calibrate the model, as it is a reduced form one with no explicit microfoundations. We aim to show that once perturbed with shocks, the limit cycle of such a model is producing realistic fluctuations. We assume that the F function is given by

$$F(I_t, u_t) = \beta_0 + \beta_1 I_t + \beta_2 I_t^2 + \beta_3 I_t^3 + \beta_4 I_t^4 + u_t.$$

u_t is an exogenous random variable, that we assume to be AR(1) :

$$u_t = \rho u_{t-1} + \varepsilon_t, \tag{11}$$

where ε is *iid* normally distributed with mean zero and variance σ^2 . Parameters are restricted such that the symmetric linear model has a globally stable equilibrium, the steady state level of I is one, $F(1) = 0$ and that there exist a stable limit cycle in the non linear model. This leads to the following choice of parameters: $\alpha_0 = 0.75$, $\alpha_1 = 0.025$, $\alpha_2 = 0.50$, $\delta = .1$ and for the nonlinear model $\beta_0 = -0.50$, $\beta_1 = 0.10$, $\beta_2 = 0.20$, $\beta_3 = 0.50$, $\beta_4 = -0.30$.

We first perform a deterministic simulation of the model, assuming that $X_0 = .8X^s$ and $I_{-1} = I^s$. Panel (a) of Figure 10 shows that the economy converges too a supercritical limit cycle in the (X, I) plane (each cross corresponds to one period). Panel (b) shows that investment goes through a deterministic cycle of booms and busts driven by endogenous economics forces. Although this cycle might share some features with economic data, it is far too regular to be an accurate description of reality. This is illustrated by the estimation of spectral density of investment and autocorrelogram of investment growth on data generated from a deterministic simulation. In the spirit of Cogley-Nason, we consider those two statistics as nicely summarizing the dynamics features of the business cycle. Results are displayed in panels (c) and (d). Although the autocorrelogram shows a lot of persistence, the spectral density is quite degenerated, with all the mass at the frequency of the limit cycle.

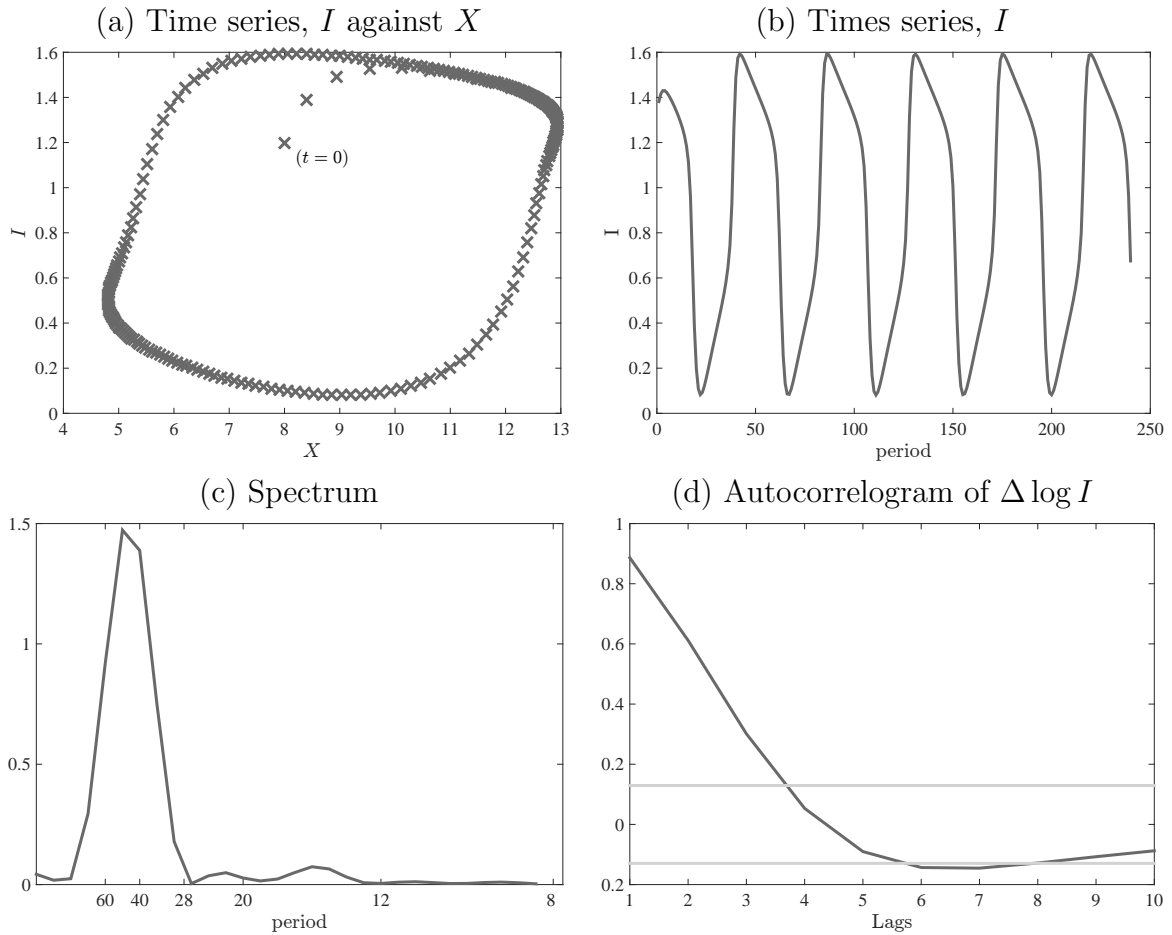
We then perform stochastic simulations of the model. Figure 11 shows in panels (a) and (b) one simulation of length 250. of the two models. In that simulations, I goes through periods of booms, with brutal bursts and even a period of protracted recession. It is of interest to notice that fluctuations now look quite different from the deterministic limit cycle. Panel(c) displays estimates of the spectral density, for 10000 simulations of length 250. The shocks “spread” the spectral density of investment, while they have little effect on autocorrelogram of investment growth.

2 A microfounded model with limit cycles

In this section we present a simple economic model that is capable of generating limit cycles. The model is an extension of the model Beaudry, Galizia, and Portier [2014]. In a nutshell, the model can be described as follows.

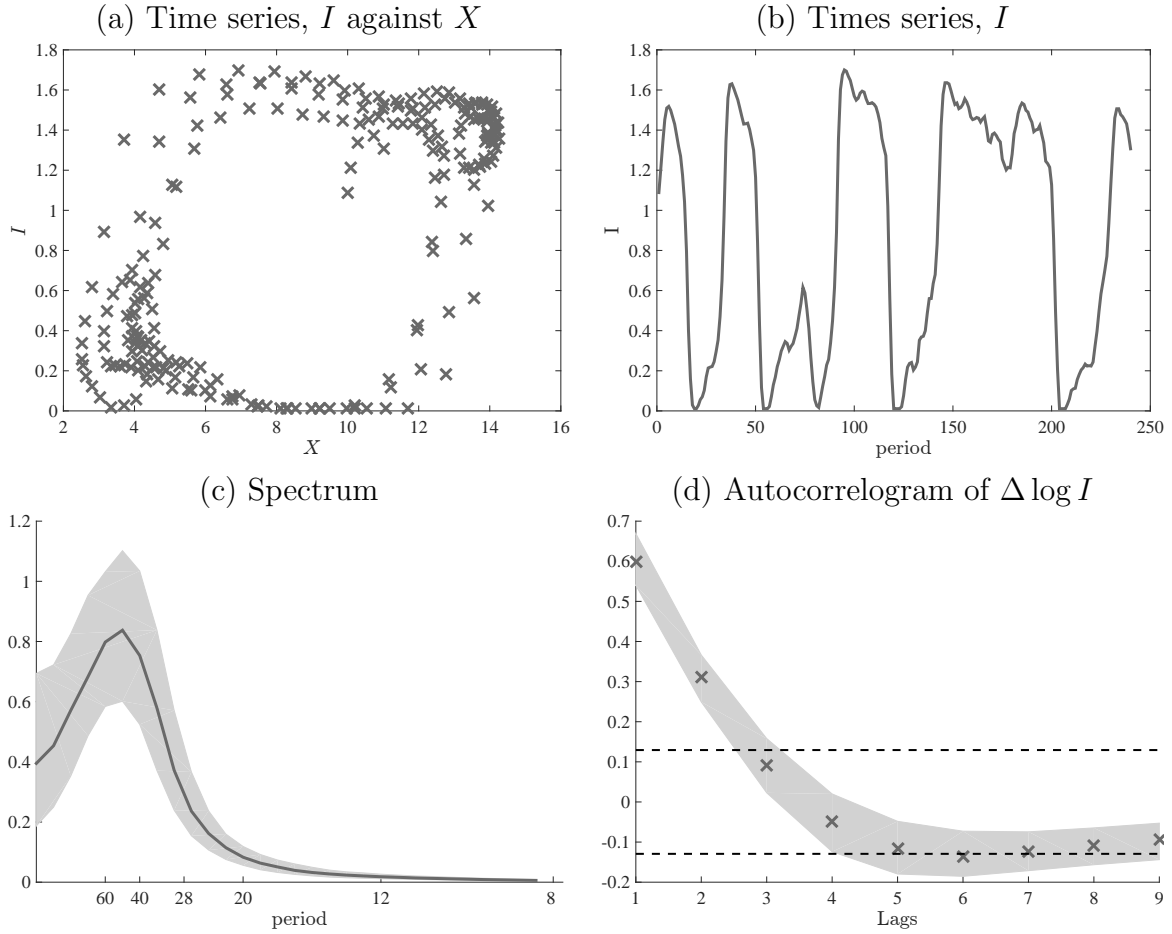
Households begin each period with a stock of durable goods and must decide how many additional goods to purchase in the goods market. Abstracting for the moment from forward-looking behavior, there are two key static factors that affect this decision. First, household

Figure 10: Deterministic simulation



This figure plots the deterministic transitional dynamics of the reduced form model in panel (a) and (b). The initial values are $X_0 = .8X^s$ and $I_{-1} = I^s$. Note that $X^s = 10$ and $I^s = 1$. In panels (b), one cross correspond to one period. For panels (c) and (d), the estimation is done 10000 times over 250 periods. The dark line shows the median estimates and the grey zone is delimited by the 18th and 82th percentiles, so that the gray band contains two-third of the simulations. This figure plots, for the numerical linear and nonlinear models, the deterministic transition from initial values $X_0 = .8X^s$ and $I_{-1} = I^s$. Note that $X^s = 10$ and $I^s = 1$. I

Figure 11: Stochastic simulations



This figure plots one simulation of the reduced form model in panel (a) and (b). Note that $X^s = 10$ and $I^s = 1$. In panels (b), one cross correspond to one period. For panels (c) and (d), the estimation is done 10000 times over 250 periods. The dark line shows the median estimates and the grey zone is delimited by the 18th and 82th percentiles, so that the gray band contains two-third of the simulations.

demand is decreasing in the size of the current stock of durables: when their existing stock of durables is low, households want to purchase more, and vice versa. Second, because of a self-insurance motive, household demand is decreasing in the unemployment rate. There are two imperfections in the model that cause this self-insurance behavior to emerge. First, there is a matching friction in the spirit of Diamond-Mortensen-Pissarides, which creates the possibility that a household may not find employment when looking for a job. Second, households are unable to perfectly insure against this idiosyncratic unemployment risk. The upshot is that an increase in the unemployment rate causes them to reduce their demand for new goods.

The combination of these two factors produces the following mechanism by which deterministic fluctuations emerge in the model: if households have a large stock of durables, they reduce their demand for new goods. This fall in demand then increases the unemployment rate, which causes households to further reduce their demand, further increasing the unemployment rate, and so on, so that, in equilibrium, output falls by significantly more than the initial fall in demand. This multiplier mechanism—which occurs because of strategic complementarity in households’ purchasing decisions—drives the excess sensitivity in the dynamic system which is a pre-condition for local instability. Once the economy reaches full or zero employment, however, the self-insurance mechanism is not operative, and thus the excessive sensitivity that creates instability disappears. In its place, inward forces—arising because of strategic substitutability, which in turn operates through the price of newly produced goods—emerge that prevent the economy from exploding. The combination of these locally-outward and globally-inward forces creates the conditions for a limit cycle to occur.

2.1 Static version

Before presenting the full dynamic model in detail, we begin by briefly presenting a simpler version of the model that is static in nature, highlighting the key properties that will be important in generating limit cycles in a dynamic setting. Further details and in-depth analysis of this static model can be found in Beaudry, Galizia, and Portier [2014].

Consider an environment populated by a mass one of households. In this economy there are two sub-periods. In the first sub-period, households purchase consumption goods and try to find employment. As there is no money in this economy, when the household buys consumption goods its bank account is debited, and when (and if) it receives employment income its bank account is credited. As we shall see, households will in general end the first sub-period with a non-zero bank account balance. Thus, in the second sub-period, households resolve their net asset positions by repaying any outstanding debts or receiving a payment for any surplus. These payments are made in terms of a second good, referred to here for simplicity as household services. Household services are also the numéraire in this economy.

Preferences for the first sub-period are represented by

$$U(c) - \nu(\ell)$$

where c represents consumption and $\ell \in [0, \bar{\ell}]$ is the labor supplied by households in the production of goods, with $\bar{\ell}$ the agent’s total time endowment. U is assumed to be strictly

increasing and strictly concave, while the dis-utility of work function ν is assumed to be strictly increasing and strictly convex, with $\nu(0) = 0$. Households are initially endowed with the same amount X of consumption goods, which they can either consume or trade. In the dynamic version of the model, X will represent a stock of durable goods and will be endogenous. Trade in consumption goods is subject to a coordination problem because of frictions in the labor market. At the beginning of the first sub-period, the household splits up responsibilities between two members. The first member, called the buyer, goes to the goods market to make purchases. The second member searches for employment opportunities in the labor market. The goods market functions in a Walrasian fashion, with both buyers and firms taking the price of these goods p (in units of household services) as given. The market for labor in this first sub-period is subject to a matching friction, with sellers of labor searching for employers and employers searching for labor. The important information assumption is that buyers do not know, when choosing how much to buy, whether the worker member of the household has secured a match. This assumption implies that buyers make purchase decisions in the presence of unemployment risk.

There is a large set of potential consumption goods firms in the economy who can decide to search for workers in view of supplying goods to the market. Each firm can hire one worker and has access to a decreasing-returns-to-scale production function $F(\ell)$, where ℓ is the number of hours worked for the firm.¹³ Being matched with a worker requires to post a vacancy at fixed cost k in terms of the output good, so that the net production of a firm hiring ℓ hours of labor is $F(\ell) - k$. Firms search for workers and, upon finding a worker, they jointly decide on the number of hours worked and on the wage to be paid. Upon a match, the determination of the wage and hours worked within a firm is done efficiently through a competitive bargaining process,¹⁴ so that in equilibrium $pF'(\ell) = w$, where w is the wage, expressed in terms of household services.¹⁵

The labor market operates as follows. All workers are assumed to search for employment. Letting n represent the vacancies, the number of matches ϕ is then given by the short side of the market, i.e., $\phi = \min\{n, 1\}$. The equilibrium condition for the goods market is then given by

$$c - X = \phi F(\ell) - nk$$

where the left-hand side is total purchases of consumption goods and the right-hand side is the total available supply after subtracting firms' vacancy posting costs. Firms enter the

¹³It is also assumed that F is such that both $F'(\ell)\ell$ and $[F(\ell) - F'(\ell)\ell]$ are strictly increasing functions of ℓ . This property is exhibited, for example, by the Cobb-Douglas function $F(\ell) = A\ell^\alpha$.

¹⁴By "competitive bargaining", We mean any bargaining process such that the equilibrium outcome satisfies (1) that workers are paid their marginal product in a match, and (2) that, conditional on being matched, workers supply and firms hire the individually-optimal number of hours at the equilibrium wage. This can be microfounded by assuming, for example, that all "matched" firms and workers meet in a secondary labor market, and that this secondary market operates in a Walrasian fashion.

¹⁵As discussed in Beaudry, Galizia, and Portier [2014], the assumption of a competitive bargaining process is for simplicity. The main mechanisms are robust to alternative bargaining protocols.

market up to the point where expected profits are zero. This condition can be written as¹⁶

$$\frac{\phi}{n} \left[F(\ell) - \frac{w}{p} \ell \right] = k$$

At the end of the first sub-period, a household's net asset position a , expressed in units of household services, is given by $a = w\ell - p(c - X)$ if the worker was employed, and $a = -p(c - X)$ if the worker was unemployed. Rather than explicitly modeling the second sub-period, for simplicity assume that the continuation value function for the second sub-period, V , is given by¹⁷

$$V(a) = \begin{cases} va & \text{if } a \geq 0 \\ (1 + \tau)va & \text{if } a < 0 \end{cases}$$

where $v, \tau > 0$ are parameters. This function is piecewise linear and concave, with a kink at $a = 0$.¹⁸ Here, the marginal value of assets is given by v when assets are positive and $(1 + \tau)v$ when assets are negative. Since buyers in general face unemployment risk when making their purchase decisions, the wedge between the marginal value of assets when in deficit and that when in surplus generates self-insurance behavior, whereby a fall in the employment rate causes buyers to reduce their purchases out of increased concern that they will end up in the costly unemployment state. This mechanism is central to the strategic complementarity that emerges in the model, which in turn is what will allow the dynamic version of the model to generate limit-cycle behavior. The strength of this mechanism, meanwhile, is governed by the parameter τ . Given the above value function V , the buyer's problem is to choose e to maximize

$$U(X + e) + \phi[-\nu(\ell) + v(w\ell - pe)] - (1 - \phi)(1 + \tau)vpe$$

subject to $e \geq 0$, where $e \equiv c - X$ is purchases of new goods. The worker's problem, meanwhile, is to choose ℓ to maximize $-\nu(\ell) + v(w\ell - pe)$.

2.1.1 Equilibrium

Letting e_j denote purchases by household j and e the average level of purchases in the economy, one may show that household j 's optimal consumption-choice decision is characterized by¹⁹

$$U'(X + e_j) = p(e)v[1 + \tau - \tau\phi(e)] \quad (12)$$

where $p(\cdot)$ and $\phi(\cdot)$ are the price of consumption goods and the employment rate, respectively, expressed as functions of aggregate purchases. The left-hand side of (12) is simply household j 's marginal utility of consumption. The right-hand side, meanwhile, captures buyer j 's expected marginal-utility cost of funds. When the economy is at full employment

¹⁶As in Beaudry, Galizia, and Portier [2014], assume that searching firms pool their ex-post profits and losses so that they each make exactly zero profits in equilibrium, regardless of whether they are matched with a worker.

¹⁷See Section 2.2 in Beaudry, Galizia, and Portier [2014] for a discussion of how to microfound such a value function.

¹⁸As noted in Beaudry, Galizia, and Portier [2014], what matters here is that the marginal value of assets be smaller in surplus than in deficit. The piecewise linearity property is assumed only for tractability.

¹⁹See Beaudry, Galizia, and Portier [2014].

($\phi(e) = 1$), this is simply equal to the price $p(e)$ of consumption goods in terms of household services, times the marginal value v of those services when assets are non-negative. When there is unemployment, however, the buyer faces some positive probability of ending up in the negative-asset state, which is associated with a higher marginal value of assets (i.e., $(1 + \tau)v$). As a result, the expected marginal-utility cost of funds is higher and, all else equal, household j would choose a lower level of purchases.

An equilibrium for this economy is given by a solution to (12) with the additional restriction that $e_j = e$. To understand how the equilibrium is affected by shifts in X , note the following properties of the equilibrium functions $p(\cdot)$ and $\phi(\cdot)$. First, one may show that $\phi(e) = \min\{e/e^*, 1\}$, where e^* is the output (net of fixed costs) produced per firm when there is a positive level of unemployment.²⁰ Second, one may show that $p(\cdot)$ is a continuous function of e , with $p'(e) = 0$ for $e < e^*$, and $p'(e) > 0$ for $e > e^*$.²¹ The consequences of these two properties for the marginal-utility cost of funds (i.e., the right-hand side of (12)) are illustrated by the curve labelled “cost of funds” in panel (a) of Figure 12. For e sufficiently small, the curve is downward-sloping: as e rises, output is increased along the extensive labor margin, lowering the unemployment rate and making purchases feel less expensive to households. Once e reaches the full-employment level e^* , however, additional increases in output come via the *intensive* labor margin, which is associated with a rising price and thus an increased cost of funds.

The two regimes—unemployment and full employment—are associated with different equilibrium responses to a rise in the endowment X .²² Panel (a) of Figure 12 shows the case for the unemployment regime. The economy is initially in equilibrium at the level e_1 of purchases, which occurs at the intersection of the cost of funds curve and the solid marginal-utility function $U'(X + e)$. A rise in the endowment by ΔX then shifts this marginal-utility function to the left by ΔX units, as represented by the dashed curve in the figure. We see that the equilibrium level of purchases falls as a result of the rise in X , and furthermore that it falls by *more* than ΔX (so that total consumption $c = X + e$ falls). This amplified response is due to the strategic complementarity that exists in the unemployment regime: a rise in the endowment causes households to reduce their demand for new goods which, via an extensive labor margin adjustment, lowers the employment rate ϕ , which in turn raises the cost of funds, causing households to reduce purchases further, further lowering the employment rate, etc.


In contrast, panel (b) of Figure 12 shows the same experiment but beginning from the full-employment regime. In this case, we again see that a rise in X is associated with a fall in

²⁰When there is unemployment, the “min” matching function and the firm’s zero-profit condition together imply $F(\ell) - F'(\ell)\ell = k$. Since k is a constant, conditional on there being unemployment this implies that $\ell = \ell^*$, where ℓ^* solves this equation. Output net of fixed costs is then $e^* \equiv F(\ell^*) - k$.

²¹Combining the household’s labor supply condition and the firm’s labor demand condition, one may obtain $p = \nu'(\ell) / [vF'(\ell)]$. As pointed out in footnote 20, when $e < e^*$ we have $\ell = \ell^*$, so that $p = p^* \equiv \nu'(\ell^*) / [vF'(\ell^*)]$. Further, once the economy achieves full employment, a rise in output must come through the intensive margin of labor (i.e., through a rise in ℓ), which causes $p(\cdot)$ to be increasing in e on $e > e^*$.

²²As shown in Beaudry, Galizia, and Portier [2014], if τ is sufficiently large there may be more than one equilibrium. While this is an interesting theoretical possibility, the evidence obtained from the quantitative exercise of section 3, though not conclusive, gives no indication that multiple equilibria are of concern. We therefore restrict attention throughout this paper to the case where the equilibrium is unique, i.e., where τ is not too large.

Figure 12: Static Equilibrium Determination



equilibrium purchases, but in this case the fall is by *less* than ΔX (so that total consumption rises). This damped response occurs as a result of the strategic substitutability that exists when the economy is at full employment: a rise in the endowment causes households to reduce their demand for new goods which, via an intensive labor margin adjustment, lowers hours-per-worker, which lowers the price p , in turn lowering the cost of funds and causing households to *increase* their purchases.

The sensitivity of purchases to changes in X in the unemployment regime because of strategic complementarity, and the corresponding *insensitivity* in the full-employment regime because of strategic substitutability, will play a crucial part in generating limit cycles in the dynamic version of the model. Note also that the sensitivity of e to X in the unemployment regime is increasing in the steepness of the slope of the cost of funds schedule in that regime. Since this steepness in turn depends positively on the parameter τ , we see that τ captures the degree of strategic complementarity in the unemployment regime.

2.2 Baseline dynamic model

Consider now a dynamic version of the above economy. Time is discrete, and each period is divided into two sub-periods, with the economy operating in each such sub-period as in the static case. The principal difference from the static model is that the stock of durable goods brought into a period is now endogenous, accumulating according to

$$X_{t+1} = (1 - \delta)(X_t + \gamma e_t) \quad (13)$$

where X_t is the stock of durables brought into period t and e_t is quantity of consumption-goods purchases in period t . For simplicity, we assume that a constant fraction $\gamma \in (0, 1]$ of

these purchases are durable.²³ $\delta \in (0, 1]$ is the depreciation rate.

The household's labor supply decision is entirely static and therefore the same as in the previous subsection (i.e., ℓ_t is chosen each period to maximize $-\nu(\ell_t) + v(w_t\ell_t - p_t e_t)$). Buyers, meanwhile, face a dynamic optimization problem, choosing c_t and e_t to maximize the objective function

$$\sum_{t=0}^{\infty} \beta^t \{U(c_t) + \phi_t [-\nu(\ell_t) + v(w_t\ell_t - p_t e_t)] - (1 - \phi_t)(1 + \tau) v p_t e_t\} \quad (14)$$

subject to $c_t = X_t + e_t$ and the accumulation equation (13), and taking ℓ_t as given.²⁴

We restrict the parameter space to cases where a steady state for this economy exists and is unique. As is the case in the static model, it can be verified that this is true as long as τ is not too large.²⁵ We further restrict parameters such that this steady state satisfies $\ell < \bar{\ell}$, so that the household's time constraint is not binding at the steady state.

2.3 Limit cycles in the dynamic model

2.3.1 The myopic case

Conditions under which limit cycles may appear in this model can be understood most easily in the myopic case where $\beta = 0$. In this case, we simply have a repeated sequence of the static economy discussed in section 2.1, with the only linkage between them being the inherited stock of durable goods. We may characterize the equilibrium evolution of the stock of durables over time as

$$X_{t+1} = (1 - \delta) [X_t + \gamma e(X_t)] \equiv g(X_t)$$

where $e(X_t)$ expresses the equilibrium level of purchases at date t as a function of the only state variable, X_t . This equilibrium is determined entirely as it was in Figure 12, with the unemployment regime characterized by strategic complementarity and the full-employment regime by strategic substitutability.

Recall the two basic conditions discussed in section 1 which are required to generate a stable limit cycle: (1) a locally unstable steady state, and (2) global non-explosiveness. Letting \bar{X} denote the steady state level of durables, these two conditions correspond mathematically to (1) $|g'(\bar{X})| > 1$, and (2) $|g'(X)| < 1$ for $|X - \bar{X}|$ sufficiently large, where

$$g'(X) = (1 - \delta) [1 + \gamma e'(X)]$$

²³In the quantitative exercise below, we will interpret “durables” as including both conventional durable goods as well as residential investment, which is conceptually similar.

²⁴In order to avoid expanding heterogeneity between individuals over time, individuals are assumed to borrow and lend via their bank account balances only within a period but not across periods. In other words, households are allowed to spend more than their income in the first sub-period of a period, but must repay any resulting debt in the second sub-period. Similar assumptions were used in Lagos and Wright [2005] and Rocheteau and Wright [2005], and more recently in Kaplan and Menzio [2014], in order to avoid having to track the asset positions of all agents in the economy over time.

²⁵See Beaudry, Galizia, and Portier [2014].

It is straightforward to verify that the second condition necessarily holds here, as follows. Suppose X is sufficiently small so that the economy is in the full-employment regime. As was shown earlier, a rise in X in this regime is associated with a fall in e , but by less than the rise in X , i.e., $-1 < e'(X) < 0$. Thus, $(1 - \delta)(1 - \gamma) < g'(X) < 1 - \delta$, and therefore $|g'(X)| < 1$ clearly holds.²⁶ Suppose instead that X is very large. In this case it can be verified that the non-negativity constraint $e \geq 0$ binds, so that $e'(X) = 0$ and therefore $g'(X) = 1 - \delta$, and thus again $|g'(X)| < 1$ holds.

Next, suppose the steady state of the system is in the unemployment regime. Then from the analysis for the static model, we know that $e'(\bar{X}) < -1$, i.e., a rise in the stock of durables leads to a more than one-for-one fall in purchases. Whether or not e falls sufficiently so that the first condition for a stable limit cycle holds will depend on the strength of the complementarity in this regime, i.e., on τ . For smaller values of τ , i.e., those for which

$$e'(\bar{X}) > -\frac{2 - \delta}{\gamma(1 - \delta)} \equiv \kappa$$

the complementarity is relatively weak, and thus $g'(\bar{X}) > -1$. In this case, the steady state is stable, so that a limit cycle will not appear. On the other hand, for larger values of τ (i.e., those for which $e'(\bar{X}) < \kappa$), we will have $g'(\bar{X}) < -1$, and thus the steady state is unstable. In combination with the fact that the system is non-explosive (as argued above), we see that in general a stable limit cycle will emerge in this case.

Figure 13 shows the formal similarity of the model determination with the one of the reduced form model. In effect, it is possible to express the expenditures e_{jt} of household j as a function of aggregate expenditures e_t and the stock of durable goods X_{t-1} , that summarizes the model history. As shown in Beaudry, Galizia, and Portier [2014], this reaction function writes

$$e_j = Z(e) - X \tag{15}$$

with

$$Z(e) \equiv U'^{-1}(Q(e))$$

and


$$Q(e) \equiv \begin{cases} \frac{\nu'(\ell^*)}{F'(\ell^*)} (1 + \tau - \tau \frac{e}{e^*}) & \text{if } 0 < e < e^* \\ \frac{\nu'(\Omega^{-1}(e))}{F'(\Omega^{-1}(e))} & \text{if } e \geq e^* \end{cases}$$

Here, $\Omega(\ell) \equiv F'(\ell)\ell$ is output net of search costs per employed worker and $e^* \equiv \Omega(\ell^*)$ is the level of output (net of firms' search costs) that would be produced if all workers were employed, with hours per employed worker equal to ℓ^* .

Figure 13 is similar to Figure 2 of the reduced form model: the model has a regime of strategic complementarities and a regime of strategic substitutability, with the location of household j optimal expenditure function (her “best response function”) depending on the model history.

²⁶It is worth emphasizing that strategic substitutability in the full-employment regime is the key property generating this relative insensitivity of e to changes in X .

Figure 13: Equilibrium determination for two different histories



This figure plots the optimal spending of household j as a function of total spendings (equation (15)): $e_{jt} = Z\left(\int_0^1 e_{it} di\right) - X_t$. The intercepts correspond to two different histories of the model as $X_t = (1 - \delta)(X_{t-1} + \gamma e_{t-1}) = \sum_1^\infty (1 - \delta)^j \gamma e_{t-j}$

2.3.2 The general case

The previous subsection showed that, when $\beta = 0$, limit cycles can emerge in the unemployment risk model. While the myopic case was useful for building intuition, of more general interest is whether limit cycles may occur for an arbitrary β . It is not immediately obvious that this should hold, and indeed, as a “Turnpike Theorem” (due to Scheinkman [1976]) below highlights, in a class of models widely used in the literature, for β sufficiently close to one limit cycles cannot occur.

In particular, consider a general deterministic dynamic economy with date- t state vector $z_t \in \mathbb{R}^n$. Let $\mathcal{W}(z_t, z_{t+1})$ denote the period- t return function when the current state is z_t and the subsequent period’s state is z_{t+1} .²⁷ The following theorem characterizes the solution to the problem of maximizing lifetime utility $\sum \beta^t \mathcal{W}(z_t, z_{t+1})$, where β is the discount factor.

Turnpike Theorem (Scheinkman [1976]) *If \mathcal{W} is concave, then there exists a $\bar{\beta} < 1$ such that if $\bar{\beta} \leq \beta \leq 1$ then the steady state is unique and globally stable.*²⁸

The key property that ensures global stability in this theorem is the assumption that \mathcal{W} is concave. Since, all else equal, fluctuations are sub-optimal when \mathcal{W} is concave, when β is sufficiently close to one it is in general optimal to take temporarily costly action in the present in order to avoid permanent fluctuations in the future. This in turn implies global convergence to the steady state, so that limit cycles cannot occur. Concavity of \mathcal{W} is a property that holds in a wide variety of economic models that have become standard in the literature, including nearly all quantitative models of the business cycle. As we shall see, however, in the unemployment-risk model discussed above, concavity of \mathcal{W} may be violated, in which case global stability may not obtain.

As a first step in establishing the potential for limit cycles in the unemployment-risk model, the following proposition verifies that the system satisfies the second condition needed for a stable limit cycle (i.e., non-explosiveness).

Proposition 5 *Given any initial endowment of durables X_0 , $\limsup_{t \rightarrow \infty} |X_t| < \infty$.*

Proposition 5 ensures that in the limit the system either exhibits deterministic fluctuations (such as a limit cycle) or converges to a fixed point. The following proposition establishes that, in contrast to models for which the Turnpike Theorem applies, local instability is possible in this model for an arbitrarily high discount factor.

Proposition 6 *There exists parameter values and functional forms such that, for some $\bar{\beta} < 1$, if $\bar{\beta} \leq \beta < 1$ then the (unique) steady state is locally unstable.*

In combination with Proposition 5, Proposition 6 confirms that there are parameter values and functional forms for which the model will generate deterministic fluctuations even if β is arbitrarily close to one. The reasons for the failure of the Turnpike Theorem to

²⁷Note that, in this formulation, \mathcal{W} implicitly incorporates any constraints and static-equilibrium outcomes, so that $\mathcal{W}(z_t, z_{t+1})$ is the equilibrium period- t return conditional on the current and next-period state being z_t and z_{t+1} , respectively.

²⁸For a proof and more formal statement of the theorem, see Scheinkman [1976] Theorem 3.

hold for this model can be clarified as follows. Suppose the steady state of the model is in the unemployment regime, and let $\mathcal{W}(X_t, X_{t+1})$ be a period- t return function such that the solution to the problem

$$\max_{\{X_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \mathcal{W}(X_t, X_{t+1}) \quad (16)$$

implements the equilibrium of the model in a neighborhood of this steady state.²⁹ If it turns out that \mathcal{W} is concave, then the Turnpike Theorem implies that the model cannot generate limit-cycle dynamics. The following proposition establishes that in fact \mathcal{W} may not be concave.

Proposition 7 *There exists parameter values and functional forms such that, in the neighborhood of an unemployment-regime steady state, \mathcal{W} is not concave.*

Intuitively, non-concavity of \mathcal{W} can arise as a result of a “bunching” mechanism in the model: because unemployment risk is low, when other agents are purchasing lots of goods it is a good time for an individual agent to purchase goods. Similarly, when other agents are purchasing few goods, it is a bad time for an individual agent to buy goods. If sufficiently strong, this bunching mechanism—which arises precisely because of the strategic complementarity in the unemployment regime—leads to a tendency to have periods of high durables accumulation alternating with periods of low durables accumulation, i.e., deterministic fluctuations.

The final proposition of this section clarifies the importance of the parameter τ in controlling the strength of this bunching mechanism, and therefore in influencing whether or not the economy will be able to generate limit-cycle dynamics.

Proposition 8 *For τ sufficiently close to zero, the steady state is stable.*

Proposition 8 thus confirms that, if τ is not sufficiently large, the degree of strategic complementarity is too small to produce an unstable steady state.

3 A DSGE model with limit cycles

This section presents the main quantitative results of the paper. We estimate an augmented version of the dynamic model discussed above, with the primary goal of establishing that it is capable of matching the key quantitative features of the hours data. We first present the data properties we are aiming to match, before describing the DSGE model and the estimation results.

3.1 Data

As our objective is to assess the possibility of a model with shocks and limit cycle to be a good approximation of the data generating process, it seems natural to look at the data in the

²⁹An example of such a \mathcal{W} is found in the proof of Proposition 7.

frequency domain. As discussed before, one of the main criticisms of earlier deterministic-fluctuations models was that they produced cycles that were far too regular. This regularity shows up clearly as a large spike in the spectrum. The ability for the model to match a much flatter spectrum as the one we will show is found in the data will thus be an important test of its ability to generate realistic data.

In order to remove non-stationarity before estimating spectral densities, we first use a band pass filter that keeps cycles with periods 2 to 100 quarters. We then display only spectral density for periods 80 to 4 quarters, so as to remove seasonality, high frequency noise and cycle above 20 years.³⁰ On those spectral density figures, the light gray zone corresponds to the traditional measure of the business cycle (8 to 32 quarters) while the dark gray zone corresponds to periods 32 to 60 quarters. The picture that emerges from Figure 14 is that much of the business cycle variability in main macroeconomic aggregates is not at periods 8 to 32 quarters, but rather above, mainly between periods 32 to 60 quarters.³¹ In the business cycle literature, in order to account for fluctuations at lower frequencies than the ones isolated by the BP(8,32) of the Hodrick-Precott(1600) filters, models are needing either very persistent shocks or a lot of propagation mechanism. We will show in this section that presence of a limit cycle can be an alternative explanation for this. Before we present the quantitative model, it is useful to show that the comovements between the main macroeconomic aggregates are not different at lower frequencies than those generally considered. As a check, we have computed the coherence between hours on the one side and output, consumption and investment on the there side. Coherence is analogous to a regression R^2 , giving the proportion of the variance of hours that can be linearly predicted by output at a given periodicity. A coherence of one would thus indicate that hours and output are perfectly correlated at that periodicity, while a coherence of zero would indicate that hours and output are orthogonal. The results displayed in Figure 15 show indeed that the same business cycle pattern of correlations is present at all periodes between 8 and 80 quarters.

In order to mainly focus on the endogenous fluctuations when estimating the model, we will focus on hours spectrum in the estimation procedure. As compared with other data series, hours is arguably less likely to be directly impacted by various exogenous shocks. For example, while GDP is directly affected by things like shocks to total factor productivity, over the business cycle we may expect hours to largely respond only indirectly to exogenous shocks. To the extent that this is true, variation in hours is more likely to be caused by the endogenous mechanisms that are the primary focus of this paper. Figure 16 displays the time series of filtered hours. Visual inspection confirms spectral analysis. For example, since the beginning of the eighties, one observes three cycles of similar amplitude and of length round 40 quarters.

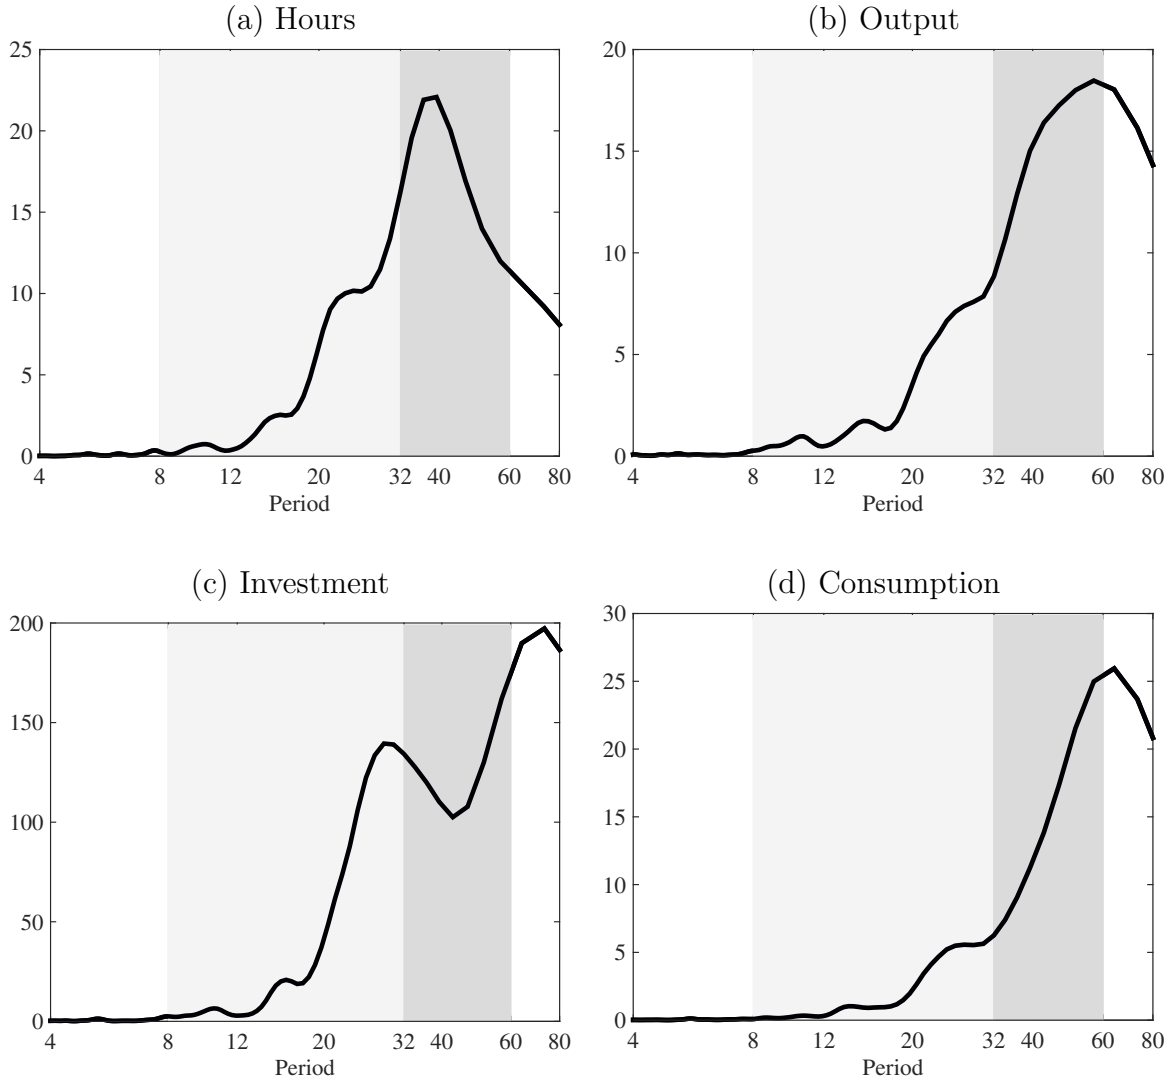
3.2 The quantitative model

The baseline dynamic model presented in section 2 was constructed with an eye toward analytical tractability. As a result, that model lacks many of the features which are known

³⁰Note that this spectral density measure is robust to changes in the bandpass filters with higher (120, 150) or lower (80) upper limit.

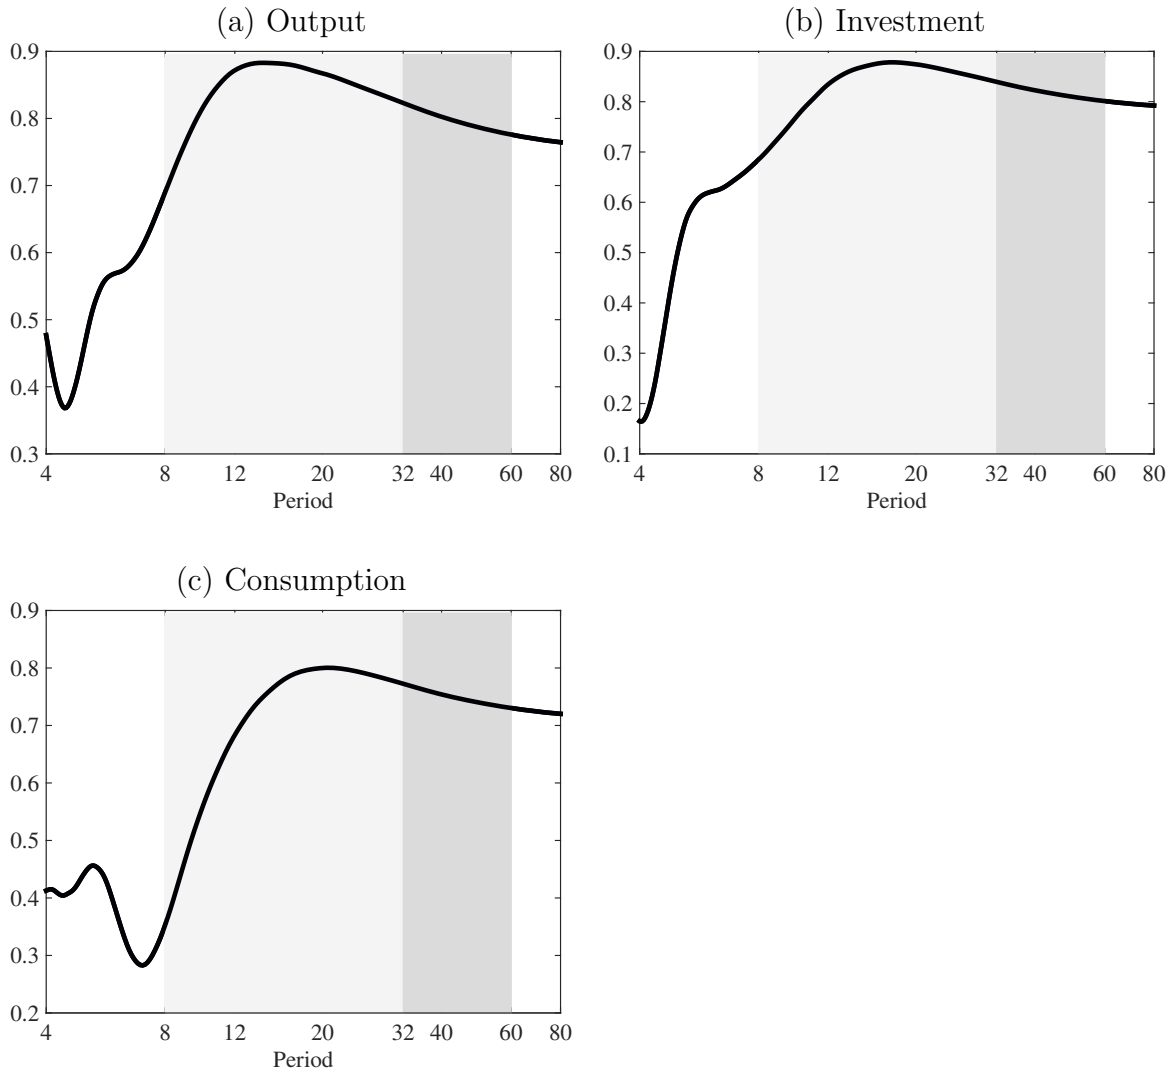
³¹See Comin and Gertler [2006] and Pancrazi [2015] for similar observations.

Figure 14: Spectral density for the main macroeconomic aggregates




All series are first detrended with a BP filter to remove fluctuations with periods greater than 100 quarters. Raw spectrum is obtained as the squared modulus of the discrete Fourier transform of the data series (scaled so that the integral with respect to angular frequency over the interval $[-\pi, \pi]$ equals the variance of the series). Spectrum in figure is kernel-smoothed raw spectrum. Kernel is a Hamming window with bandwidth parameter 11.

Figure 15: Coherence with hours for the main macroeconomic aggregates



All series are first detrended with a BP filter to remove fluctuations with periods greater than 100 quarters. Raw coherence at a periodicity p is given by $|s_{L,x}(p)|^2 / [s_L(p) s_x(p)]$, where s_L is the spectrum of hours, s_x is the spectrum of the other series, and $s_{L,x}$ is the cross-spectrum. Coherence was then kernel-smoothed using a Hamming window with bandwidth parameter 51.

Figure 16: Hours Worked Data



Notes: Hours Worked series is the log of BLS nonfarm hours worked divided by population, detrended with a BP filter to remove fluctuations with periods greater than 100 quarters. Shaded areas are NBER-dated recessions. The sample runs from 1960Q1 to 2012Q4.

to be helpful in quantitatively matching the data, and includes several others which, while not central to the key mechanisms, turn out to be restrictive in a quantitative setting. Since the main purpose of the exercise in this section is quantitative in nature, we make several adjustments to the model designed to help it in that regard.

First, as is well known, dynamic systems with a single state variable have considerable difficulty in producing deterministic fluctuations with the basic qualitative properties that we observe in macroeconomic aggregates. In particular, deterministic fluctuations in such models tend to be erratic, with the system often jumping back and forth from one side of the steady state to the other every few periods or less. Thus, if the unemployment-risk model is to have any chance of successfully replicating key features of the data, it will require the addition of at least one other state variable. To this end, and following much of the quantitative business cycle literature, we now assume that the household exhibits internal habit-formation in consumption,³² so that its period utility for consumption is now given by

$$U(c_t - hc_{t-1})$$

Here, $h \in [0, 1)$ is a parameter controlling the degree of habit persistence.

Second, the relatively simple structure of the baseline model produces a stark dichotomy, whereby in the unemployment regime all output adjustments occur along the extensive labor margin, while in the full-employment regime all adjustments occur along the intensive margin. In order to relax this stark dichotomy, in the quantitative version of the model we allow firms to be heterogeneous in terms of their fixed costs. That is, rather than assuming that all firms have fixed cost k , we assume that the n -th firm has fixed cost $k(n) \geq 0$, where $k(\cdot)$ is a non-decreasing function. This will allow for the possibility of there being regions where both extensive and intensive labor margin adjustments may occur.^{33,34}

Third, as discussed earlier and in contrast to what is observed in the data, purely deterministic models of economic fluctuations tend to yield cycles of a constant length. This can be observed either as a very regular pattern in a plot of time series data generated from the model, or as one or more large spikes in the spectrum estimated from that data.³⁵ One of the key contributions of this paper is to show that by introducing a relatively small amount

³²The key desirable property for a second state variable here is that it introduces momentum into the dynamics of X , so that movements from a high to a low level of X and back are gradual, rather than rapid as they are when X is the only state variable. Consumption habit exhibits this property by reducing period-to-period fluctuations in household demand, with the added advantages that it maintains tractability and keeps the model as close as possible to the baseline version discussed earlier. Nonetheless, there are likely a number of other choices (e.g., adjustment costs in investment or employment) that could have been made instead and that would have delivered similar qualitative dynamics.

³³To see this, note that the marginal firm entrant must earn zero expected profit, which in the unemployment regime is equivalent to the condition $F(\ell) - F'(\ell)\ell = k(n)$, where n is the index of the marginal entrant. A rise in the employment rate is associated with a rise in n , which (weakly) increases the right-hand side of this expression. Since the left-hand side of this expression is strictly increasing in ℓ , this then implies that a rise in the employment rate is in general also associated with a rise in hours-per-worker, i.e., both extensive and intensive labor margin adjustments occur.

³⁴The functional form chosen for this $k(\cdot)$ (discussed below) will nest the baseline case of a constant fixed cost. Since the parameters of this function will be estimated, the data will ultimately choose the degree to which $k(\cdot)$ is non-constant.

³⁵One may show that the spectrum associated with any limit cycle is infinitely high at a countable number of points (i.e., a countable sum of Dirac delta functions), and zero everywhere else.

of randomness into a limit-cycle model it becomes possible to produce realistically irregular fluctuations. To this end, we also include in the model an exogenous random TFP process, $\tilde{\theta}_t$.³⁶

3.3 Functional forms, calibration and estimation

Production is assumed to be of the Cobb-Douglas form

$$F(\ell) = A\ell^\alpha$$

Utility over consumption (net of habit) is assumed to be of the form

$$U(C) = aC - \frac{b}{2}C^2$$

while disutility of labor is taken to be of the form

$$\nu(\ell) = \frac{\nu_1}{1+\omega}\ell^{1+\omega}$$

The fixed cost of the n -th firm is assumed to be given by

$$k(n) = \begin{cases} 0 & n \leq n_0 \\ \frac{n-n_0}{\eta}\bar{k} & n_0 < n < n_0 + \eta \\ \bar{k} & n \geq n_0 + \eta \end{cases}$$

where n_0 , η and \bar{k} are parameters. This function is piecewise linear with three regimes: low- n firms have fixed cost zero, high- n firms have fixed cost \bar{k} , and over the intermediate range k rises linearly from zero to \bar{k} .³⁷ Finally, we assume the TFP process is given by

$$\theta_t \equiv \log(\tilde{\theta}_t) = \rho\theta_{t-1} + \epsilon_t, \quad \epsilon_t \sim N\left(0, \left(\frac{\sigma}{100}\right)^2\right)$$

Several of the model parameters were directly calibrated. In particular, we set the labor share at a standard value of $\alpha = 2/3$. The inverse Frisch elasticity was calibrated at the widely used level $\omega = 1$. We set the depreciation rate and discount factor at standard values of $\delta = 0.025$ and $\beta = 0.99$, respectively, and normalize the maximum fixed cost at $\bar{k} = 1$. Finally, the fraction of purchases entering the durables stock was calibrated at $\gamma = 0.192$, which is the average ratio of durables to total consumption in the National Income and Product Accounts data.³⁸ The remaining parameters were estimated.

³⁶For convenience, in order to retain certain analytical properties that are helpful in a computational setting, we assume that firms' fixed costs and households' second-sub-period value functions also fluctuate with the TFP process. Output, fixed costs, and the value function are thus given by $\tilde{\theta}_t F(\cdot)$, $\tilde{\theta}_t k(\cdot)$, and $\tilde{\theta}_t^{-1} V(\cdot)$, respectively.

³⁷Quadratic utility and the piecewise-linear form for $k(\cdot)$ were assumed for tractability and computational efficiency. None of the key properties of the model rely on these assumptions.

³⁸As noted above, we include the conceptually-similar residential investment under the heading of "durables". The figure of 0.192 can thus be obtained from NIPA data as the average of (Durable goods + Residential investment)/(Consumption + Residential investment) over the sample period 1960Q1-2012Q4.

Solving the model for a particular parameterization was done using the parameterized expectations (PE) approach.³⁹ Given this solution, a large data set ($T = 100,000$ periods in length) was simulated and, after taking logs of the resulting hours series and detrending it with the same BP filter as used for the data, the spectrum of log-hours was estimated. The non-calibrated parameters were then estimated so as to minimize the average squared difference between the model spectrum and the spectrum estimated from the data. Further details of the solution and estimation procedure are presented in Appendix C.

Estimated parameter values are reported in Table 1. Several things should be noted. First, the TFP process is close to the process that would be estimated directly from productivity data. For example, using John Fernald’s [2014] measure of business-sector labor productivity growth over the sample period (1960Q1-2012Q4),⁴⁰ after cumulating, linearly detrending, and fitting an AR(1) process, one obtains a persistence estimate of 0.974 and an innovation standard deviation of 0.713%, yielding an unconditional productivity standard deviation of 3.16%.⁴¹ The corresponding parameters estimated for the unemployment-risk model, meanwhile, are $\rho = 0.969$ and $\sigma = 0.570$, respectively, which yields an unconditional standard deviation of 2.30%. The fact that the model only features a single shock, and that the variance of that shock in the model is, if anything, smaller than its data counterpart highlights the more general observation that models featuring deterministic fluctuations may not require the presence of large amounts of exogenous variation in order to generate empirically reasonable business cycles.

The only other parameter with a clear comparator in the data or literature is habit persistence, which is estimated here to be $h = 0.76$, well within the range of standard estimates obtained elsewhere in the literature. For example, Smets and Wouters [2007] report a 90% confidence interval for habit of (0.64, 0.78), while Justiniano, Primiceri, and Tambalotti [2010] report a 90% confidence interval of (0.72, 0.84).

The remaining parameters in Table 1 are composed mainly of uninteresting scale parameters, and parameters for which few if any precedents exist. The parameter τ , which captures the strength of the household’s desire to reduce spending in response to a rise in unemployment risk, falls into the latter category. Given its central role in the model, however, it deserves some comment. If interpreted narrowly as a one-period financial premium on debt vis-à-vis saving, the estimate of $\tau = 0.27$, or 27%, clearly exceeds typical borrowing-lending spreads as reported in the literature. However, there are several reasons to think this view of τ may be overly restrictive. First, in order to avoid significantly complicating the model, conditional on the employment rate an individual worker’s probability of being employed is assumed to be independent from quarter to quarter. If the actual employment state of an individual exhibits persistence, then considering only one-period financial costs may understate households’ desire to reduce spending in response to an increase in unemployment. Second, borrowing-lending spreads that reflect average borrowing rates faced by all households may not accurately reflect rates faced by unemployed individuals, which are likely to be higher.

³⁹See, for example, den Haan and Marcet [1990] and Marcet and Marshall [1994]. Details can be found in Appendix C.

⁴⁰Available at http://www.frbsf.org/economic-research/economists/jferald/quarterly_tfp.xls.

⁴¹Similar values are obtained when using Fernald’s TFP or utilization-adjusted TFP measures instead of labor productivity.

Table 1: Parameter Values

Parameter	Value	Description
<i>Estimated Parameters</i>		
a	12.535	Marginal utility of consumption, intercept
b	2.247	Maginal utility of consumption, slope
h	0.761	Habit persistence
ν_1	13.274	Labor disutility scaling factor
τ	0.270	Premium on debt
A	3.199	Constant productivity factor
n_0	0.843	Measure of firms with zero fixed cost
η	0.091	Measure of firms over which fixed cost is rising
ρ	0.969	Persistence of TFP
σ	0.570	$100 \times$ s.d. of innnovation to TFP
<i>Calibrated Parameters</i>		
α	0.667	Labor share
ω	1	Inverse Frisch elasticity
δ	0.025	Depreciation of durables
β	0.99	Discount factor
\bar{k}	1	Maximum firm fixed cost
γ	0.192	Fraction of purchases entering durables stock

Third, many unemployed individuals may in fact be unable to access financial markets at all, instead being forced to rely on costly asset liquidations and/or reduced consumption levels in order to meet their obligations, the potential for either of which may cause households to strongly reduce their desired spending. To the extent that any or all of these factors should be subsumed into τ , the value estimated here may not be unreasonable.

3.4 Main Results


To illustrate the deterministic mechanisms, we first report results obtained when shutting down the TFP shock (i.e., setting $\sigma = 0$).⁴² Panel (a) of Figure 17 plots a simulated 212-quarter sample⁴³ of log-hours generated from this deterministic model. Two key properties should be noted. First, the model is clearly capable of generating cycles of a reasonable length, which in this case is approximately 30 quarters. The apparent inability of models of deterministic flutcutations to generate cycles of quantitatively reasonable lengths appears to have been one of the factors leading to the abandonment of this literature. As this exercise

⁴²In particular, we first obtained the PE coefficients from the full stochastic model. The simulation results for the deterministic model were then generated using these stochastic PE coefficients, but feeding in a constant value $\theta_t = 0$ for the TFP process. In other words, agents in the deterministic model implicitly behave as though they live in the stochastic world. As a result, any differences between the deterministic and the stochastic results in this section are due exclusively to differences in the realized sequence of TFP shocks, rather than differences in, say, agents' beliefs about the underlying data-generating process.

⁴³This is equal to the length of the sample period of the data.

demonstrates, however, unreasonable cycle lengths are by no means an unavoidable property of these models. Second, notwithstanding the reasonable cycle length, it is clear when comparing the simulated data in Figure 17 to the actual data in Figure 16 that the fluctuations in the deterministic unemployment-risk model are far too regular,⁴⁴ a shortcoming shared by many earlier models of deterministic fluctuations.

Figure 17: Deterministic Model



Note: Panel (a) shows 212-quarter simulated sample (same size as data set) of BP-filtered log(hours worked) ($\phi_t \ell_t$) generated from the deterministic model. Initial simulated series was 252 quarters long, with first and last 20 quarters discarded after BP-filtering. Details for computation of model spectrum in panel (b) can be found in Appendix C.

These properties of the deterministic model—i.e., a highly regular 30-quarter cycle—can also be seen clearly in the frequency domain. Panel (b) of Figure 17 plots the spectrum for the deterministic model (dashed line), along with the spectrum for the data (solid line) for comparison.⁴⁵ Consistent with the pattern in the time domain, the spectrum exhibits a peak

⁴⁴Note that the cycles clearly do not exactly repeat themselves. This is a technical issue here that we sidestep throughout this paper. Certain types of deterministic fluctuations share many of the basic qualitative features of a limit cycle, but never exactly repeat themselves. For example, in a bivariate discrete-time system characterized by rotation around the unit circle by θ radians per period, if θ/π is irrational then the system will never return to the same point twice. This property is due to the discrete-time formulation of the model. In a continuous-time version of the model, the cycles would necessarily repeat themselves, a direct consequence of the Poincaré-Bendixson Theorem (see, e.g., Guckenheimer and Holmes [2002], p. 44).

⁴⁵Note that the model was not re-estimated after shutting down the TFP shock. As such, there may be alternative parameterizations of the deterministic model that are better able to match the spectrum in the

at around 30 quarters. Further, the regularity of the cycle is manifested as a large spike in the spectrum. In contrast, the spectrum estimated from the data is much flatter.

Re-introducing the TFP shock into the model, we see a markedly different picture in both the time and frequency domains. Panel (a) of Figure 18 plots a 212-quarter sample of log-hours generated from the stochastic model. While clear cyclical patterns are evident in the figure, it is immediately obvious that the inclusion of the TFP shock results in fluctuations that are significantly less regular than those generated in the deterministic model, appearing qualitatively quite similar to the fluctuations found in Figure 16 for actual data. This is confirmed by the spectrum, which is plotted in panel (b) of Figure 18 alongside the data spectrum. Also plotted is a pointwise 90% simulated confidence interval from the model for data sets of the same length as the data (i.e., 212 quarters).⁴⁶ The stochastic model clearly matches the data quite well in this dimension, including possessing a peak near 40 quarters and, as compared to the deterministic model, lacking any large spike. The good fit of the model can also be seen by looking at the autocovariance function (ACF) of hours, i.e., $Cov(L_t, L_{t-k})$, where k is the lag (in quarters). Panel (a) of Figure 19 plots the result for the first 40 lags for both the data and model, along with pointwise 90% confidence intervals. As the figure shows, the curves lie nearly on top of one another, indicating that the model matches the data very well in this dimension also.⁴⁷

To verify that the good fit of the spectrum is not driven by the choice of filter, Figure 20 plots the data and model spectra for hours under four alternative filtering choices.⁴⁸ Panels (a)-(c) present results for three alternative band-pass filters with different upper bounds (100, 60, and 40 quarters, respectively), while panel (d) plots spectra using a Hodrick-Prescott filter with parameter 1600. As the figure shows, the model fits the data very well in all cases.

Next, it should be emphasized that the exogenous shock process in this model primarily accelerates and decelerates the endogenous cyclical dynamics, causing significant random fluctuations in the length of the cycle while only modestly affecting its amplitude. For example, in the deterministic version of the model the standard deviation of log-hours is 0.026, while in the stochastic model it is 0.033, implying that 79% of the standard deviation of hours is due to deterministic mechanisms. In contrast, if this TFP process were the only shock process operating in the widely-cited model of Smets and Wouters [2007], for example, it would generate a standard deviation of log-hours of only 0.005. This again suggests the more general point that, if one is willing to consider the class of models capable of generating deterministic fluctuations, then a very parsimonious set of shocks that are small in magnitude can potentially yield qualitatively and quantitatively reasonable fluctuations.

As a final exercise in this section, it is worth briefly further comparing the above results to those of Smets and Wouters [2007]. Their model has received much attention in the literature for its ability to fit well a number of key macroeconomic data series. Panel (a) of Figure 21 shows the spectrum for hours worked as generated by the Smets and Wouters [2007] model

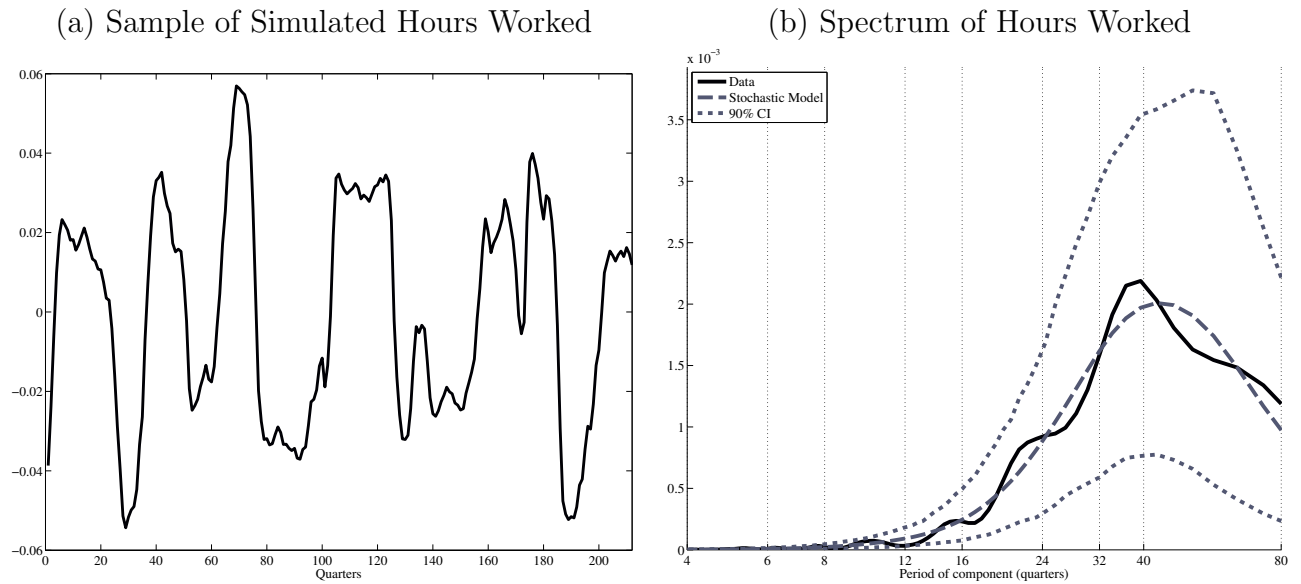
data.

⁴⁶That is, if the model were the true data-generating process, then at each periodicity the spectrum estimated from the data would lie inside the confidence interval 90% of the time.

⁴⁷Note that the ACF is simply the inverse Fourier transform of the spectrum. Since the spectrum of the model and data are similar, we would expect the ACF to be similar as well, a property clearly verified in Figure 19.

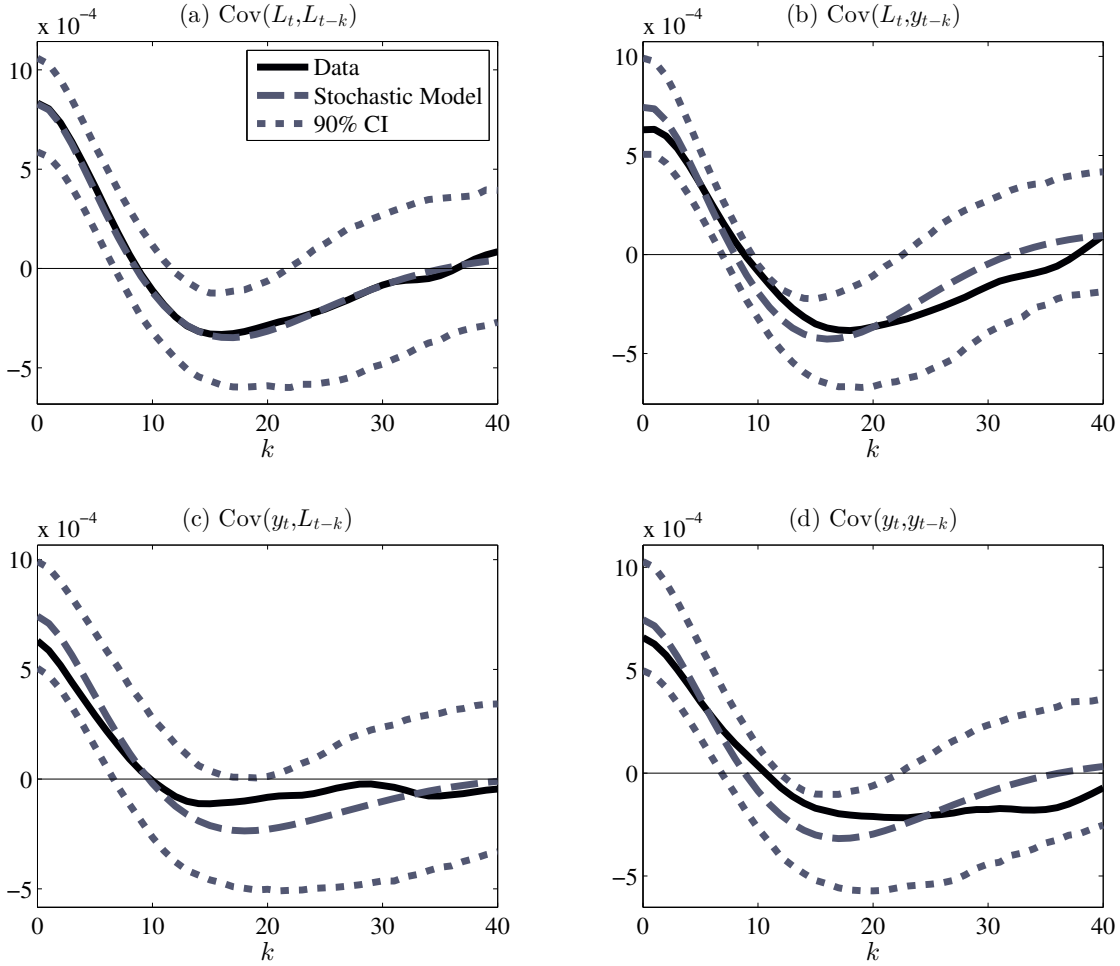
⁴⁸Note that the model spectra were obtained using the baseline model parameters as reported in Table 1.

Figure 18: Stochastic Model



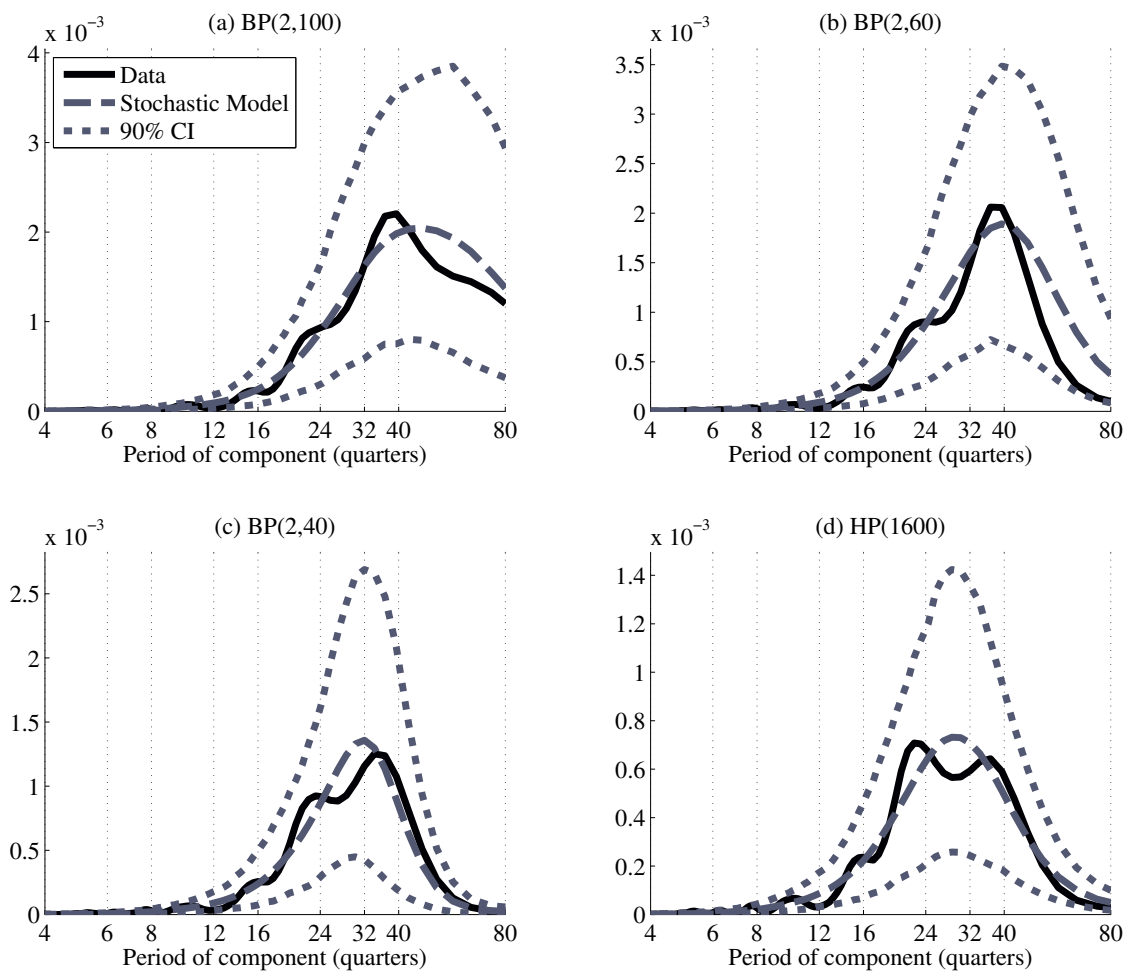
Note: Panel (a) shows 212-quarter simulated sample (same size as data set) of BP-filtered $\log(\text{hours worked})$ ($\phi_t \ell_t$) generated from the stochastic model. Initial simulated series was 252 quarters long, with first and last 20 quarters discarded after BP-filtering. Details for computation of model spectrum in panel (b) can be found in Appendix C. Dotted lines show a pointwise 90% confidence interval for the spectrum that would be estimated from a model-generated data set of the same length as the actual data set (i.e., 212 quarters).

Figure 19: Autocovariance: Hours Worked (L) and Output (y)



Note: Figure shows autocovariances of BP(2,80)-filtered hours and output in the data and stochastic model. k is the lag in quarters. Data series for output is the log of nominal GDP, deflated by population and the GDP deflator. Output in the model is the sum of wage earnings and firm profits, which is equal to total production net of fixed costs, i.e., $\tilde{\theta}_t [\phi_t F(\ell_t) - \int_0^{n_t} k(x) dx]$, where n_t is the number of firm entrants at date t . Dotted lines show pointwise 90% confidence intervals for the autocovariance functions that would be estimated from a model-generated data set of the same length as the actual data set (i.e., 212 quarters).

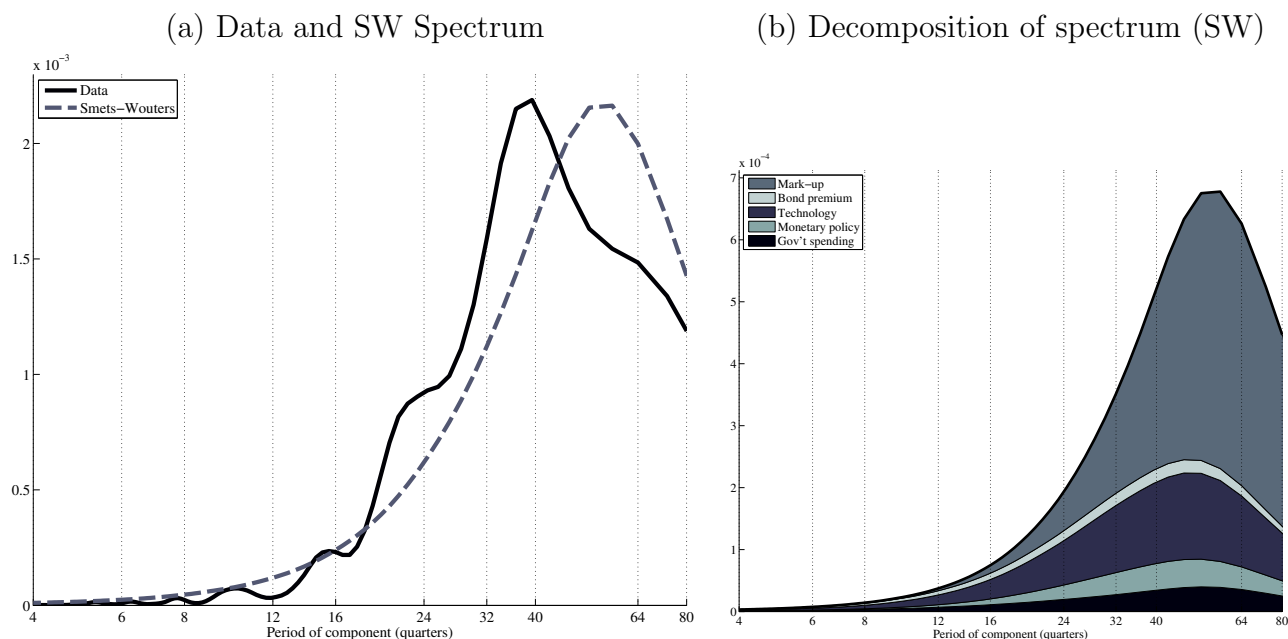
Figure 20: Spectrum: Hours Worked (Alternative Filters)



Note: Each panel plots corresponding data (solid) and model (dashed) spectrum using the reported filter instead of the baseline BP(2,80) filter. Dotted lines show pointwise 90% confidence intervals for the spectrum that would be estimated from a model-generated data set of the same length as the actual data set (i.e., 212 quarters).

at the reported median posterior parameter values. As suggested by the relatively close fit,

Figure 21: Hours Worked in Smets-Wouters (SW)



Notes: Data spectrum is as in Figure 16. Spectrum for Smets-Wouters (SW) obtained by simulating 10,000 data sets of the same size as the actual data series. For each simulation, the data was detrended and the spectrum estimated using the same procedures as for the actual data. A point-wise average was taken across all simulated spectra. Because the hours series used by SW for their estimation differs somewhat from the series used here, for purposes of comparability, in panel (a) the SW spectrum was scaled by a constant so that the total variance is the same as in the data. Panel (b) shows portion of variance at each periodicity attributable to each of the following shock groupings: “Mark-up” – price and wage mark-up shocks; “Bond Premium” – bond premium shock; “Technology” – TFP and investment-specific technology shocks; “Monetary policy” – monetary policy shock; “Gov’t spending” – government spending shock.

their model also matches patterns in the hours data reasonably well, though not quite as well as the unemployment-risk model.⁴⁹

More insight into the drivers of fluctuations in the Smets and Wouters [2007] model can be obtained by looking at a spectral variance decomposition; that is, by decomposing the total variance at each individual periodicity into the portions that are attributable to each of the shocks in that model. Panel (b) of Figure 21 presents such a decomposition. It is

⁴⁹This should not be too surprising, as the unemployment-risk model was estimated to match only the hours series, while the Smets and Wouters [2007] was estimated to simultaneously match seven different data series (including hours).

clear from the figure that, in the range of periodicities responsible for the bulk of the variance of hours, the two mark-up shocks (price and wage) in the Smets and Wouters [2007] model account for by far the largest portion. In fact, the proportion of the total hours variance that is explained by the mark-up shocks rises monotonically with periodicity, explaining around a third of the variance of hours by the 24-quarter periodicity and over half by the 36-quarter periodicity.⁵⁰ In contrast, the unemployment-risk model presented here is equally capable of matching the spectrum in hours, but does so with only a reasonably-sized TFP shock and without relying on poorly motivated mark-up shocks.

3.5 Additional Results

To this point, we have focused on the fit of the model with respect to the target series, hours worked. In this subsection, we evaluate how well the model performs in several other dimensions that were not directly targeted.

Panel (a) of Figure 22 compares the spectrum of output for the data and the stochastic model.⁵¹ As shown in the figure, the model spectrum matches the data reasonably well, though it is somewhat too large (indicating too much output variance in the model), and the average periodicity is somewhat too low. The second observation should not be too surprising, as the model does not include capital as a factor of production. Since productive capital tends to exhibit lower-frequency fluctuations than labor (the other factor of production), all else equal its omission from the model will cause the average periodicity of output to be too small. Panel (d) of Figure 19, meanwhile, plots the ACF for output, which confirms the first observation: the variance of output in the model (i.e., the autocovariance at lag $k = 0$) is slightly larger than in the data. Notwithstanding this, however, the spectrum and ACF for output in the data lies well within a 90% confidence interval for the model, suggesting a relatively good overall fit.

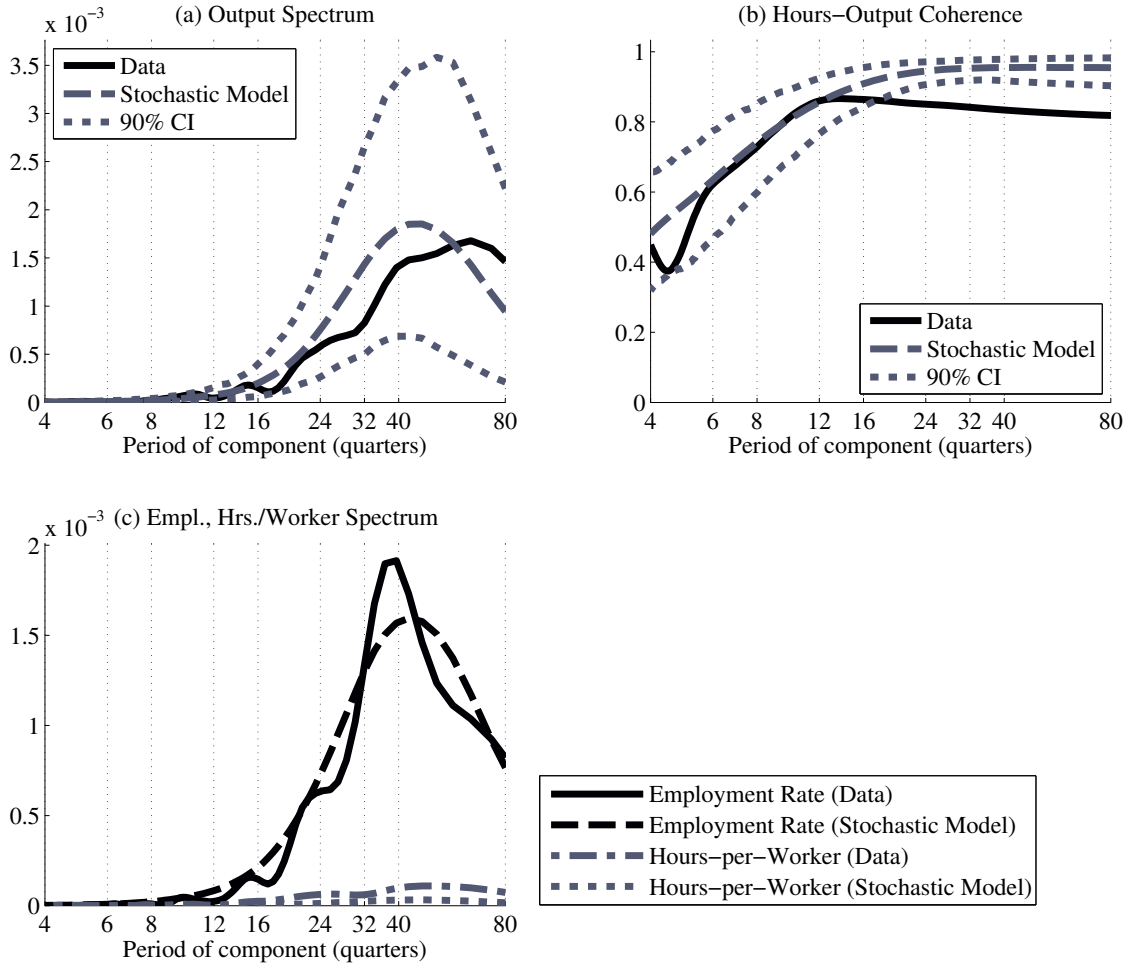
Next, panel (b) of Figure 22 plots the coherence between hours and output for the data and for the stochastic model.⁵² Coherence is analogous to a regression R^2 , giving the proportion of the variance of hours that can be linearly predicted by output at a given periodicity. A coherence of one would thus indicate that hours and output are perfectly correlated at that periodicity, while a coherence of zero would indicate that hours and output are orthogonal. In the data (solid line in the figure), we see that at the lowest periodicities hours and output are modestly correlated, with coherence around 0.4-0.5. As the periodicity rises, the coherence initially increases relatively rapidly, reaching a peak of 0.87 at around 13 quarters. Over this range, as indicated by the dashed line in the figure the model coherence matches the data very well. Beyond the 13-quarter periodicity, however, the data and model begin to diverge somewhat. The data coherence largely flattens out, with a gradual downward

⁵⁰The importance of the mark-up shocks is not exclusive to hours within the Smets and Wouters [2007] model. For example, as reported in that paper, at a 40-quarter horizon the mark-up shocks together account for over half of the forecast-error variance (FEV) of output and over 80% of the FEV of inflation.

⁵¹Data series for output is the log of nominal GDP, deflated by population and the GDP deflator, then de-trended using a BP(2,80) filter using the same procedure as with hours worked. Output in the model is the sum of wage earnings and firm profits, which is equal to total production net of fixed costs, i.e., $\tilde{\theta}_t [\phi_t F(\ell_t) - \int_0^{n_t} k(x) dx]$, where n_t is the number of firm entrants at date t .

⁵²The coherence at a periodicity P is given by $|s_{L,y}(P)|^2 / [s_L(P) s_y(P)]$, where s_L is the spectrum of hours, s_y is the spectrum of output, and $s_{L,y}$ is the cross-spectrum.

Figure 22: Spectrum: Output (Data and Stochastic Model)



Notes: Data series for output is the log of nominal GDP, deflated by population and the GDP deflator. Data series for the employment rate is the log of the BLS's index of nonfarm business employment divided by population. Data series for hours-per-worker is the log of nonfarm business hours divided by nonfarm business employment. All series were de-trended using a BP(2,80) filter using the same procedure as with hours worked. Output in the model is the sum of wage earnings and firm profits, which is equal to total production net of fixed costs, i.e., $\tilde{\theta}_t [\phi_t F(\ell_t) - \int_0^{n_t} k(x) dx]$, where n_t is the number of firm entrants at date t . Spectrum for data and model computed as with hours. Raw coherence at a periodicity p is given by $|s_{L,y}(p)|^2 / [s_L(p) s_y(p)]$, where s_L is the spectrum of hours, s_y is the spectrum of output, and $s_{L,y}$ is the cross-spectrum. Coherence was then kernel-smoothed using a Hamming window with bandwidth parameter 51. In panels (a) and (b), dotted lines show pointwise 90% confidence intervals for the spectrum and coherence, respectively, that would be estimated from a model-generated data set of the same length as the actual data set (i.e., 212 quarters).

slope, reaching 0.82 at the 80-quarter periodicity. The model coherence, meanwhile, rises somewhat over this range. As with the spectrum of output, the discrepancy between the data and model coherences at higher periodicities can be explained by the lack of productive capital in the model.⁵³ Notwithstanding this discrepancy, however, the basic qualitative properties of the relationship between hours and output in the data—namely, moderate correlation at higher frequencies but significant correlation at medium-to-low frequencies (including the range of frequencies in which the bulk of variation occurs)—are well-captured by the model.

While coherence measures the strength of the relationship between two series at a given periodicity, it provides no information about the sign of this relationship or whether one series tends to lead the other. To address how well the model fits in these dimensions, panels (b) and (c) of Figure 19 plot the cross-covariance function (CCF) for hours and output. Two things should be noted from these plots. First, hours and output are positively correlated in both the model and data. Second, in the model hours and output are in phase (i.e., the peak of the CCF occurs at a lag of $k = 0$), while in the data the peak occurs at the point where output leads hours by one quarter. Nonetheless, the CCF is close to flat in the data between its peak and $k = 0$,⁵⁴ suggesting that any lead of output is weak at best. Further, as suggested by the reported 90% confidence intervals, over all the cross-covariance between output and hours is well-captured by the model.

Finally, while we have established that the model does a good job of matching patterns in total hours, consider the model’s implications for its two component parts, the employment rate, ϕ_t , and hours-per-worker, ℓ_t . Panel (c) of Figure 22 shows spectra for the data and stochastic model for these two series.⁵⁵ From the figure, we see that the spectrum of the employment rate from the model matches fairly well the one from the data, and in particular the employment rate exhibits an overall level of volatility that is close to the volatility in the data. Thus, this model addresses one of the frequent criticisms of many models of unemployment in the literature, which is that they generate too little employment volatility.⁵⁶

On the other hand, the model does a relatively poor job of matching behavior in hours-per-worker. In particular, while the basic pattern of the model spectrum is close to that in the data, the model spectrum is in most places too small, especially beyond the lowest periodicities. This suggests that the model features too little in the way of movements along the intensive labor margin.⁵⁷ To understand why, recall that when the economy moves

⁵³Including capital would tend to reduce the coherence between output and hours by introducing another factor of production which is imperfectly correlated with hours. Since fluctuations in capital tend to be much more important at higher periodicities, the coherence would tend to fall by more at the upper end of the range of periodicities.

⁵⁴The peak of the data CCF is only 0.28% greater than it is at $k = 0$.

⁵⁵Data series for the employment rate is the log of the BLS’s index of nonfarm business employment divided by population. Data series for hours-per-worker is the log of nonfarm business hours divided by nonfarm business employment. Both series were de-trended using a BP(2,80) filter using the same procedure as with hours worked.

⁵⁶See for example Shimer [2005].

⁵⁷As Figure 22 shows, extensive-margin fluctuations are an order of magnitude larger than intensive-margin fluctuations in both the model and the data. As a result, even though the model does not capture well the intensive-margin fluctuations, this has little impact on the fit of total hours, which is driven primarily by extensive-margin fluctuations.

into a region where the fixed-cost function $k(\cdot)$ is increasing, forces come into play which cause output fluctuations to occur on both intensive and extensive labor margins. Recall also that the former are associated with strategic substitutability (through changes in the price of goods), while the latter are associated with strategic complementarity (through changes in unemployment risk). If a given change in output occurs too much along the intensive margin (as is the case for this parameterization of the model), the associated strategic substitutability tends to push the economy back towards the steady state quickly, so that any change in hours-per-worker is relatively small and short-lived.⁵⁸

3.6 Multiple Equilibria and Indeterminacy

In the estimation exercise conducted above, we only considered parameter combinations for which (a) there exists a unique steady state, and (b) the probability of having multiple static equilibria (i.e., multiple equilibria in a period, conditional on the current state and on agents' beliefs about the future) was negligible. As mentioned briefly above, these two constraints can be expressed as upper bounds on τ . Intuitively, multiple steady states and multiple equilibria may arise in this model if the strategic complementarity between agents' actions is too strong. Since τ governs the strength of this complementarity, ruling out multiple equilibria is equivalent to limiting the size of τ .

In particular, define

$$\tau^* \equiv \frac{\alpha \bar{k} b}{(1 - \alpha) v p^*}$$

$$\bar{\tau} \equiv \frac{[1 - \beta(1 - \gamma)(1 - \delta)][(1 - \delta)\gamma + \delta](1 - \beta h)(1 - h)}{[1 - \beta(1 - \delta)]\delta} \tau^*$$

The following proposition characterizes sufficient (though not necessary) conditions under which the steady state and static equilibria are unique.

Proposition 9 *The steady state of the unemployment-risk model is unique if $\tau < \bar{\tau}$. The period- t static equilibrium is unique if $\tau < \tilde{\theta}_t^2 \tau^*$.*

Ex ante, it is not clear whether imposing the constraints on τ from Proposition 9 is restrictive in practice. The results from the estimation reported above, however, give no indication that these constraints are binding. In particular, at the parameter values reported in Table 1, we have $\bar{\tau} = 2.61$ and $\tau^* = 0.82$, both well above the value of $\tau = 0.27$. Clearly, the constraint ensuring a unique steady state is not binding at the optimal parameter values. The constraint ensuring a static equilibrium, meanwhile, depends on the level of productivity $\tilde{\theta}_t$, which can in principle be arbitrarily small, and thus the constraint may be violated with strictly positive probability. Nonetheless, given the size of the estimated TFP shock, this probability is negligible in practice. For example, in 100,000 simulated periods, the smallest value of $\tilde{\theta}_t^2 \tau^*$ that occurred was 0.69, still more than twice the value of τ .

⁵⁸One way to increase the variance of hours-per-worker is thus to have the upward-sloping part of the fixed-cost function be less steep. Since hours-per-worker was not a target of the estimation algorithm, however, there is no reason why it should have favored a flatter $k(\cdot)$. Improving the fit of the model in this dimension by including hours-per-worker information as part of the estimation objective function is a task for future work.

While there are relatively simple analytical conditions that can be obtained to ensure uniqueness of the steady state and of static equilibrium, verifying dynamic determinacy—that is, the presence of a unique path converging to the limit cycle for a given initial state—is more challenging, since no analytical results are available in general. Nonetheless, in numerical simulations we were unable to find any evidence of indeterminacy. In particular, given initial values for the state variables and arbitrary initial values for the jump variables (chosen in practice from some neighborhood of the PE solution), one may simulate a non-stochastic version of the model forward.⁵⁹ If, for a given initial state, the system were to converge to the limit cycle for multiple combinations of initial jump variables, this would indicate the presence of indeterminacy. However, performing this experiment many times beginning from different initial conditions, in all cases the system eventually exploded, which suggests to me that indeterminacy is not likely to be an issue here.

Conclusion

The objective of this paper was to show that the limit cycle view of the business cycle was neither a theoretical curiosity nor an quantitatively implausible theory of macroeconomic fluctuations. To that end, we proposed in a first section a general reduced form framework with strategic complementarities between agents decisions. In this framework, we have shown that limit cycles emerge, not as a knife-edge case but quite generally under mild conditions on the degree of strategic complementarities. In particular, we show that the strength of strategic complementarities that is needed is less than what is needed to generate multiple equilibria and sunspot fluctuations. Then we have proposed in a second section a general equilibrium model that builds on Beaudry, Galizia, and Portier [2014]. In that model, unemployment risk and precautionary savings interact to induce a limit cycle driven by the type of forces outlined in our general structure. In a third section, we have highlighted some interesting properties of data spectral density. First, the spectral density is not singular with all its mass at a given frequency, as would suggest a purely deterministic limit cycle model. Second, there is a non negligible mass in the spectrum at medium frequencies that are lower than the ones generally considered by the business cycle literature. Given these observations, we propose a DSGE version of our model that introduces habit persistence and shocks to total factor productivity. The model is estimated to fit the spectral density of hours worked. Estimation favors a syncretic view according to which the data generating process are well approximated with a limit cycle perturbed by TFP shocks. That model is then showed to also fit the spectral density of output. We believe that those results convincingly show that the limit cycle view of the business cycle is a worth pursuing avenue for future quantitative research.

⁵⁹See Appendix D for details.

References

- BEAUDRY, P., D. GALIZIA, AND F. PORTIER (2014): “Reconciling Hayek’s and Keynes’ views of recessions,” Working Paper.
- BENHABIB, J., AND K. NISHIMURA (1979): “The Hopf Bifurcation and the Existence and Stability of Closed Orbits in Multisector Models of Optimal Economic Growth,” *Journal of Economic Theory*, 21(3), 421–444.
- (1985): “Competitive Equilibrium Cycles,” *Journal of Economic Theory*, 35(2), 284–306.
- BOLDRIN, M., N. KIYOTAKI, AND R. WRIGHT (1993): “A dynamic equilibrium model of search, production, and exchange,” *Journal of Economic Dynamics and Control*, 17(5), 723–758.
- BOLDRIN, M., AND L. MONTRUCCHIO (1986): “On the Indeterminacy of Capital Accumulation Paths,” *Journal of Economic Theory*, 40(1), 26–39.
- BOLDRIN, M., AND M. WOODFORD (1990): “Equilibrium models displaying endogenous fluctuations and chaos,” *Journal of Monetary Economics*, 25(2), 189–222.
- COLES, M. G., AND R. WRIGHT (1998): “A Dynamic Equilibrium Model of Search, Bargaining, and Money,” *Journal of Economic Theory*, 78(1), 32–54.
- COMIN, D., AND M. GERTLER (2006): “Medium-Term Business Cycles,” *American Economic Review*, 96(3), 523–551.
- DAY, R. H. (1982): “Irregular Growth Cycles,” *The American Economic Review*, 72(3), 406–414.
- (1983): “The Emergence of Chaos from Classical Economic Growth,” *The Quarterly Journal of Economics*, 98(2), 201–213.
- DAY, R. H., AND W. SHAFER (1987): “Ergodic Fluctuations in Deterministic Economic Models,” *Journal of Economic Behavior and Organization*, 8(3), 339–361.
- DEN HAAN, W. J., AND A. MARCET (1990): “Solving the Stochastic Growth Model by Parameterizing Expectations,” *Journal of Business & Economic Statistics*, 8(1), 31–34.
- DIAMOND, P., AND D. FUDENBERG (1989): “Rational expectations business cycles in search equilibrium,” *Journal of Political Economy*, 97(3), 606–619.
- FERNALD, J. (2014): “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Federal Reserve Bank of San Francisco Working Paper 2012-19.
- GLENDINNING, P. (1994): *Stability, Instability and Chaos: an introduction to the theory of nonlinear differential equations*. Cambridge University Press, New York, NY.
- GOODWIN, R. (1951): “The Nonlinear Accelerator and the Persistence of Business Cycles,” *Econometrica*, 19(1), 1–17.
- GRANDMONT, J.-M. (1985): “On Endogenous Competitive Business Cycles,” *Econometrica*, 53(5), 995–1045.
- GUCKENHEIMER, J., AND P. HOLMES (2002): *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer, New York.

- HAMMOUR, M. L. (1989): “Endogenous Business Cycles: Some Theory and Evidence,” PhD Thesis, Massachusetts Institute of Technology.
- HICKS, J. (1950): *A Contribution to the Theory of the Trade Cycle*. Clarendon Press, Oxford.
- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2010): “Investment shocks and business cycles,” *Journal of Monetary Economics*, 57(2), 132–145.
- 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.
- KAPLAN, G., AND G. MENZIO (2014): “Shopping Externalities and Self-Fulfilling Unemployment Fluctuations,” Working Paper.
- LAGOS, R., AND R. WRIGHT (2005): “A Unified Framework for Monetary Theory and Policy Analysis,” *Journal of Political Economy*, 113(3), 463–484.
- LE CORBEILLER, P. (1933): “Les systemes autoentretenus et les oscillations de relaxation,” *Econometrica*, 1(3), 328–332.
- MAR CET, A., AND D. A. MARSHALL (1994): “Solving Nonlinear Rational Expectations Models by Parameterized Expectations: Convergence to Stationary Solutions,” Federal Reserve Bank of Minneapolis Discussion Paper 91.
- PANCRAZI, R. (2015): “The heterogeneous Great Moderation,” *European Economic Review*, 74, 207–228.
- ROCHETEAU, G., AND R. WRIGHT (2005): “Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium,” *Econometrica*, 73(1), 175–202.
- SAMUELSON, P. (1939): “Interaction Between the Multiplier Analysis and the Principle of Acceleration,” *The Review of Economic Studies*, 21(2), 75–78.
- SCHEINKMAN, J. (1976): “On Optimal Steady States of n-Sector Growth Models when Utility is Discounted,” *Journal of Economic Theory*, 12(1), 11–30.
- (1990): “Nonlinearities in Economic Dynamics,” *The Economic Journal*, 100(400), 33–48, Conference Papers.
- SHIMER, R. (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *The American Economic Review*, 95(1), 25–49.
- SMETS, F., AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *The American Economic Review*, 97(3), 586–606.
- WAN, Y.-H. (1978): “Computations of the stability condition for the Hopf bifurcation of diffeomorphisms on R^2 ,” *SIAM Journal of Applied Mathematics*, 34(167-175).
- WIKRAN, A. (2013): *Discrete Dynamical Systems with an Introduction to Discrete Optimization Problems*. Bookboon.com, London, U.K.

Appendix

A Proofs of section 1

A.1 Proposition 1

The two eigenvalues of matrix M_L are the solution the equation

$$Q(\lambda) = \lambda^2 - T\lambda + D = 0 \quad (\text{A.1})$$

where T is the trace of the M_L matrix (and also the sum of its eigenvalues) and D is the determinant of the M_L matrix (and also the product of its eigenvalues). The two eigenvalues are therefore given by

$$\lambda, \bar{\lambda} = \frac{T}{2} \pm \sqrt{\left(\frac{T}{2}\right)^2 - D} \quad (\text{7})$$

where

$$T = (\alpha_2 - \alpha_1 + (1 - \delta)) \quad (\text{A.2})$$

and

$$D = \alpha_2((1 - \delta)). \quad (\text{A.3})$$

From (A.3), we have that $\lambda\bar{\lambda} \in]0, 1[$. Therefore, if the eigenvalues are complex, their modulus is in between zero and one so that they are both inside the unit circle. If the two eigenvalues are real, they have the same sign. From (A.2), we have that $\lambda + \bar{\lambda} \in]-1, 2[$. Therefore, if eigenvalues are negative, they are both inside the unit circle. If they are both positive, let's assume that λ is the largest eigenvalue. If $\lambda > 1$, then $\bar{\lambda} < 1$ because $\lambda + \bar{\lambda} = T < 2$. Given that $\lambda\bar{\lambda} > 1$, we have $\bar{\lambda} < \frac{1}{\lambda}$ so that $\lambda + \frac{1}{\lambda} < 2$ which implies $(1 - \lambda)^2 < 0$. This is not possible and hence, $\lambda < 1$. As λ is the largest eigenvalue and that both are real and positive, both eigenvalues are inside the unit circle.

A.2 Proposition 2

With demand complementarities, the trace and determinants of matrix M are given by

$$T = \left(\frac{\alpha_2 - \alpha_1}{1 - F'(I^s)} + (1 - \delta) \right) \quad (\text{A.4})$$

and

$$D = \frac{\alpha_2((1 - \delta))}{1 - F'(I^s)}. \quad (\text{A.5})$$

From equations (A.4) and (A.5), we have the following relation between the trace and determinant of matrix M :

$$D = \frac{\alpha_2(1 - \delta)}{\alpha_2 - \alpha_1} T - \frac{\alpha_2(1 - \delta)^2}{\alpha_2 - \alpha_1} \quad (\text{A.6})$$

Therefore, when $F'(I^s)$ varies, T and D move along the line (A.6) in the plane (T, D) . We have shown that when $F'(I^s) = 0$, (T, D) belongs to the triangle \widehat{ABC} , meaning that

both eigenvalues of M are inside the unit circle. This corresponds to point E or point E' (depending of the configuration of parameters) on Figure 4.

When $F'(I^s)$ tends to $-\infty$, D tends to 0 and T tends to $1 - \delta$, which corresponds to point E_1 on Figure 4. As this point is inside the triangle \widehat{ABC} , both eigenvalues are inside the unit circle. When $F'(I^s)$ goes from 0 to $-\infty$, (T, D) moves along the segment $[E, E_1]$ or $[E', E_1]$. Because both belongs to \widehat{ABC} and because the interior of triangle \widehat{ABC} is a convex set, both eigenvalues of matrix M stay inside the unit circle when $F'(I^s)$ goes from 0 to $-\infty$.

A.3 Proposition 3

A fold bifurcation occurs with the appearance of an eigenvalue equal to 1, a flip bifurcation with the appearance of an eigenvalue equal to -1 and a Hopf bifurcation with the appearance of two complex conjugate eigenvalues of modulus 1. From (A.5) and (A.4), we see that when $F'(I^s)$ tends to 1 from below, D tends to $+\infty$ and T tends to $\pm\infty$ depending on the sign of $\alpha_2 - \alpha_1$. Therefore, starting from point E or E' for which $F'(I^s) = 0$, (T, D) will exit the triangle \widehat{ABC} , so that at least on eigenvalue will have a modulus one when the half-line (A.6) that starts at E (or E') will cross the triangle \widehat{ABC} , as shown in Figures 4 and 5.

Consider first the case $\alpha_2 > \alpha_1$. This corresponds to Figure 4. In that case, the line (A.6) has a positive slope, and will cross either segment $[AC]$ ((a) on the Figure) or segment $[BC]$ ((b) on the Figure). In case (a), we will have a fold bifurcation has eigenvalues are real and one will be equal to 1 when crossing the triangle \widehat{ABC} . In cas (b), both eigenvalues are complex, and will both have modulus 1 when crossing \widehat{ABC} , so that we will have a Hopf bifurcation. We will be in case (b) when $D = 1$ and $T < 2$. $D = 1$ implies $F'(I^s) = 1 - \alpha_2(1 - \delta)$. Plugging into the expression of T , the condition $T < 2$ writes $1 - \delta + \frac{\alpha_2 - \alpha_1}{\alpha_2(1 - \delta)} < 2$ which can be simplified to $\alpha_2 < \frac{\alpha_1}{\delta^2}$. Therefore, if $\alpha_1 < \alpha_2 < \frac{\alpha_1}{\delta^2}$, we have a Hopf bifurcation and if $\alpha_2 > \frac{\alpha_1}{\delta^2}$, we have a fold bifurcation.

Consider now the case $\alpha_2 < \alpha_1$. This corresponds to Figure 5. In that case, the line (A.6) has a negative slope, and will cross either segment $[AB]$ ((a) on the Figure) or segment $[BC]$ ((b) on the Figure). In case (a), we will have a flip bifurcation has eigenvalues are real and one will be equal to -1 when crossing the triangle \widehat{ABC} . In cas (b), both eigenvalues are complex, and will both have modulus 1 when crossing \widehat{ABC} , so that we will have a Hopf bifurcation. We will be in case (b) when $D = 1$ and $T > -2$. $D = 1$ implies $F'(I^s) = 1 - \alpha_2(1 - \delta)$. Plugging into the expression of T , the condition $T > -2$ writes $1 - \delta + \frac{\alpha_2 - \alpha_1}{\alpha_2(1 - \delta)} > -2$ which can be simplified to $\alpha_2 > \frac{\alpha_1}{(2 - \delta)^2}$. Therefore, if $\alpha_2 < \frac{\alpha_1}{(2 - \delta)^2}$, we have a flip bifurcation and if $\alpha_1 > \alpha_2 > \frac{\alpha_1}{(2 - \delta)^2}$, we have a Hopf bifurcation.

Finally, in the case $\alpha_1 = \alpha_2$, we always have $T = 1 - \delta$, so that D increases with $F'(I^s)$ along a vertical line that necessarily crosses the segment $[BC]$, so that we have a Hopf bifurcation.

Putting all those results together gives the conditions stated in Proposition 3.

A.4 Proposition 4

For this proposition, we make use of Wan's [1978] theorem and of the formulation given by Wikran [2013]. For symmetric allocations, our non-linear dynamical system is given by

$$\begin{pmatrix} I_t - F(I_t) \\ X_t \end{pmatrix} = \begin{pmatrix} \alpha_2 - \alpha_1 & -\alpha_1(1 - \delta) \\ 1 & 1 - \delta \end{pmatrix} \begin{pmatrix} I_{t-1} \\ X_{t-1} \end{pmatrix} + \begin{pmatrix} \alpha_0 \\ 0 \end{pmatrix} \quad (\text{A.7})$$

To study the stability of the limit cycle in case this system goes through a Hopf bifurcation, we need to write the system in the following "standard form"

$$\begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \end{pmatrix} + \begin{pmatrix} f(y_{2t-1}, y_{2t-1}) \\ g(y_{1t-1}, y_{2t-1}) \end{pmatrix} \quad (\text{A.8})$$

where y_1 and y_2 are functions of I and X . Let μ be the bifurcation parameter ($\mu = F'(I^s)$ in our case) and μ_0 the value for which the Hopf bifurcation occurs. Define

$$d = \frac{d|\lambda(\mu_0)|}{d\mu}$$

and

$$a = -\text{Re} \left(\frac{(1 - 2\lambda)\bar{\lambda}^2}{1 - \lambda} \xi_{11} \xi_{20} \right) - \frac{1}{2} |\xi_{11}|^2 - |\xi_{02}|^2 + \text{Re}(\bar{\lambda} \xi_{21})$$

where

$$\begin{aligned} \xi_{20} &= \frac{1}{8} ((f_{11} - f_{22} + 2g_{12}) + i(g_{11} - g_{22} - 2f_{12})) \\ \xi_{11} &= \frac{1}{4} ((f_{11} + f_{22}) + i(g_{11} + g_{22})) \\ \xi_{02} &= \frac{1}{8} ((f_{11} - f_{22} - 2g_{12}) + i(g_{11} - g_{22} + 2f_{12})) \\ \xi_{21} &= \frac{1}{16} ((f_{111} + f_{122} + g_{112} + g_{222}) + i(g_{111} + g_{122} - f_{112} - f_{222})) \end{aligned}$$

According to Wan [1978], the Hopf bifurcation is supercritical if $d > 0$ and $a < 0$.

We first write (A.7) under the standard form (A.8). Denoting $i_t = I_t - I^s$ and $x_t = X_t - X^s$ and $\widehat{F}(i_t) = F(i_t + I^s)$, and recalling that $F(I^s) = \widehat{F}(0) = 0$, we can rewrite (A.7) as

$$\begin{pmatrix} i_t - \widehat{F}(i_t) \\ x_t \end{pmatrix} = \begin{pmatrix} \alpha_2 - \alpha_1 & -\alpha_1(1 - \delta) \\ 1 & 1 - \delta \end{pmatrix} \begin{pmatrix} i_{t-1} \\ x_{t-1} \end{pmatrix} \quad (\text{A.9})$$

Define $H(i_t) = i_t - \widehat{F}(i_t)$. Under the restriction $F'(\cdot) < 1$, H is a monotonous increasing function, and is therefore invertible. Denote $G(\cdot) = H^{-1}(\cdot)$. Adding and subtracting to the right hand side of the first equation of (A.9) a first order approximation of G around zero, we obtain

$$\begin{pmatrix} i_t - \widehat{F}(i_t) \\ x_t \end{pmatrix} = \underbrace{\begin{pmatrix} \frac{\alpha_2 - \alpha_1}{1 - F'(I^s)} & -\frac{\alpha_1(1 - \delta)}{1 - F'(I^s)} \\ 1 & 1 - \delta \end{pmatrix}}_M \begin{pmatrix} i_{t-1} \\ x_{t-1} \end{pmatrix} + \begin{pmatrix} m(i_{t-1}, x_{t-1}) \\ 0 \end{pmatrix}, \quad (\text{A.10})$$

with

$$m(i_{t-1}, x_{t-1}) = G(\alpha_1(1-\delta)x_{t-1} + (\alpha_2 - \alpha_1)i_{t-1}) + \frac{\alpha_1(1-\delta)}{1-F'(I^s)}x_{t-1} - \frac{\alpha_2 - \alpha_1}{1-F'(I^s)}i_{t-1}.$$

The eigenvalues of M are the solution of equation

$$Q(\lambda) = \lambda^2 - T\lambda + D = 0$$

where T is the trace of the M matrix and D its determinant, with $T = \left(\frac{\alpha_2 - \alpha_1}{1-F'(I^s)} + (1-\delta)\right)$ and $D = \frac{\alpha_2((1-\delta))}{1-F'(I^s)}$. At the Hopf bifurcation, $D = 1$ and the two eigenvalues are $\lambda = \cos \theta \pm i \sin \theta$ with $\theta = \arctan\left(\frac{\sqrt{(T/2)^2 - 1}}{T/2}\right)$. Let λ be the eigenvalue with positive imaginary part and $\bar{\lambda}$ its conjugate, and denote D and C the two matrices

$$D = \begin{pmatrix} \lambda & 0 \\ 0 & \bar{\lambda} \end{pmatrix}$$

and

$$C = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}.$$

By construction, λ and $\bar{\lambda}$ are the eigenvalues of C . We introduce matrices V_C and V_A which are such that $C = V_C D V_C^{-1}$ and $M = V_M D V_M^{-1}$. Those matrices can be computed to be

$$V_C = \begin{pmatrix} \sin \theta & \sin \theta \\ -i \sin \theta & i \sin \theta \end{pmatrix}$$

and

$$V_M = \begin{pmatrix} \lambda + \delta - 1 & \bar{\lambda} + \delta - 1 \\ 1 & 1 \end{pmatrix}$$

We therefore have $C = V_C D V_C^{-1} = V_C V_M^{-1} M V_M V_C^{-1} = B M B^{-1}$ with

$$B = V_C V_M^{-1} = \begin{pmatrix} 0 & \sin \theta \\ -1 & \cos \theta - (1-\delta) \end{pmatrix}$$

Let make the change of variable $(y_{1t}, y_{2t})' = B \times (i_t, x_t)'$ to obtain the “standard form” of (A.7)

$$\begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \end{pmatrix} + \begin{pmatrix} f(y_{2t-1}, y_{2t-1}) \\ g(y_{1t-1}, y_{2t-1}) \end{pmatrix} \quad (\text{A.8})$$

with

$$\begin{aligned} f(y_{2t-1}, y_{2t-1}) &= 0 \\ g(y_{1t-1}, y_{2t-1}) &= -G \left(\frac{\gamma_1}{\sin \theta} y_{1t-1} - \gamma_2 y_{2t-1} \right) + \frac{1}{1-F'(I^s)} \left(\frac{\gamma_1}{\sin \theta} y_{1t-1} - \gamma_2 y_{2t-1} \right) \end{aligned}$$

and $\gamma_1 = -\alpha_2(1-\delta) + (\alpha_2 - \alpha_1) \cos \theta$, $\gamma_2 = \alpha_2 - \alpha_1$.

We can now check the conditions for the Hopf bifurcation to be supercritical, namely

$$d = \frac{d|\lambda(\mu_0)|}{d\mu} < 0$$

and

$$a = -\operatorname{Re} \left(\frac{(1-2\lambda)\bar{\lambda}^2}{1-\lambda} \xi_{11}\xi_{20} \right) - \frac{1}{2}|\xi_{11}|^2 - |\xi_{02}|^2 + \operatorname{Re}(\bar{\lambda}\xi_{21}) < 0$$

Denote $\mu = F'(I^s)$ the bifurcation parameter. $|\lambda| = \det(M) = \frac{\alpha_2(1-\delta)}{1-\mu}$, so that

$$d = \frac{d|\lambda|}{d\mu} = \frac{(1-\mu) + \alpha_2(1-\delta)}{(1-\mu)^2} > 0$$

as $\mu = F'(I^s) < 1$.

Consider now the expression for a . As $G(I)$ is the reciprocal function of $I - F(I)$, we have

$$G''' = \frac{F'''(1-F')^2 + 2F''^2(1-F')}{(1-F')^4},$$

with $F' < 1$. This shows that G''' is an increasing function of F''' . When F''' becomes large in absolute terms and negative, so does G''' . In the expression for a , the first three terms, $-\operatorname{Re} \left(\frac{(1-2\lambda)\bar{\lambda}^2}{1-\lambda} \xi_{11}\xi_{20} \right) - \frac{1}{2}|\xi_{11}|^2 - |\xi_{02}|^2$, are not function of F''' . The last term is

$$\begin{aligned} \operatorname{Re}(\bar{\lambda}\xi_{21}) &= \frac{\alpha_2(1-\delta)}{16} \left(\frac{\gamma_1^2}{\sin^2 \theta} + \gamma_2^2 \right) G''' \\ &= \kappa G''' \end{aligned}$$

with $\kappa > 0$. If F''' is sufficiently negative, then will be G''' , and therefore $\operatorname{Re}(\bar{\lambda}\xi_{21})$ and a .

Therefore, $d > 0$ and under the condition $F''' \ll 0$, we have $a < 0$. By Wan's [1978] theorem, the limit cycle is supercritical.

A.5 Global analysis of reduced form dynamics in continuous time

TO BE WORKED

In order to get a sense of the forces that drive the global dynamics in a system with demand complementarities, it is easiest to consider a continuous time setup. To this end, we generalize here our discrete time setup so as to be able to easily move it to its continuous time analogue. In particular, let us consider the following discrete time generalization of our previous system given by 1 and 4

$$H_t = h \cdot I_t + (1 - \delta_H)H_{t-1}, \quad 0 < \delta_H \leq 1 \quad (\text{A.11})$$

and

$$X_{t+1} = I_t + (1 - \delta)X_{t-1} \quad (\text{A.12})$$

$$I_t = \alpha_0 - \alpha_1 X_t + \alpha_2 H_{t-1} + F(I_t), \quad F'(I) < 1 \quad (\text{A.13})$$

In this system, we now have two accumulation equation for two state variables X_t and H_{t-1} . We also have the the decision variable I_t that effects the accumulation of both state variables. Once again the potential for demand complementarities arise in the determination of I_t . Note that our previous set up corresponds to the special case where $h = \delta_I = 1$, that is, the case where the second state variable depreciates fully after one period. We introduce this more general specification as to be able to extend it to continuous time while maintaining two state variables. In the structural example given later, I_t will represent the purchase of a durable good, X_t will be the stock of the durable good, and H_t will present a habit term that affects agents utility for the goods. Alternatively, this type of system can be derived in an environment where there are intertemporal linkages in the production of the durable due to something akin to adjustment costs to investment. The continuous time analogue to A.11 and A.12 is given by

$$\dot{H} = h \cdot I_t - \delta_H I_t, \quad 0 < \delta_I \quad (\text{A.14})$$

and

$$\dot{X} = I_t - \delta X_t \quad (\text{A.15})$$

For our analysis of the system given by A.13, A.14 and A.15, we will restrict attention to the case where $\alpha_2 < \frac{1}{h} \min[\alpha_1 + \delta_H + \delta, (\alpha_1 + \delta)\delta_H]$ as this guarantees that the system is stable when $F(I) = 0$. It is relevant to note that this continuous time system should mimic quite closely to our discrete time system when δ_I is sufficiently large. For the system given by A.13, A.14 and A.15, we could again analyze local dynamics and explore what happens as we vary $F'(I^s)$. Analogues to Prop 1 2 and 3 can be derived in such a case.⁶⁰ However, what interests us here is the global behavior of this system. In particular, we want use the Poincarre-Bendixson Theorem to help us understand when limit cycles are likely to arise in this set-up. This theorem states, in loose terms, that there are only three possible outcome configurations in a two dimensional system of differential equations: either the system is globally stable, the system is globally unstable, or the system exhibit at least one attracting limit cycle. Moreover, an important aspect of the Poincarre-Bendixson Theorem is the use of bounded sets to show the existence of limit cycles. In particular, if it can be shown that there exists a bounded open set from which the system will never leave, and in this set there is a unstable steady state, then the system exhibits at least one limit cycle. As noted above, in economic environments there are often forces that may create bounds for investment decisions. For example, there may be an upper bound on the number of workers who can work in the sector. Suppose this upper bound on available work creates an upper bound on I denoted \bar{I} . This upper bound can be captured in our framework by assuming the social interact term has the following property.

$$\lim_{I \rightarrow \bar{I}} F(I) = -\infty$$

⁶⁰ The condition for a Hopf bifurcation to arise in this continuous time system as $F'(I^s)$ moves towards one is that $\alpha_2 \in [\frac{\alpha_1}{h}, \frac{\alpha_1 \delta_H^2}{h \delta^2}]$.

This condition can be interpreted as reflecting the fact that as aggregate investment goes towards the upper bound, prices are going to infinity creating a strong dis-incentive for any one individual to invest. This condition will thereby ensure that it will not be an equilibrium outcome to choose an level of investment greater or equal to that \bar{e} . Once we have an upper bound on investment, we immediately get as a by product an upper bound on the two stock variables X_t and H_{t-1} .

In terms of a lower bound on I , in economic environments this may come for at least two sources. First, it may be the case that as investment gets sufficiently low or negative, then some prices fall to zero creating a strong incentive for people to invest even if others are not investing. If the lower bound is given by \underline{I} , then the condition $\lim_{I \rightarrow \underline{I}} F(I) = \infty$ will guarantee that the equilibrium choice of investment at any point in time will be above \underline{e} .

Proposition A.1 *If interaction effect $F(\cdot)$ is such that the system given by A.13, A.14 and A.15 is locally unstable, and both limit conditions $\lim_{I \rightarrow \bar{I}} F(I) = -\infty$ and that $I_t \geq 0$ hold, then the system will have at least one attracting limit cycle.*

Proposition A.1 is a direct application of the Bendixson-Poincarre theorem. The important element we want to emphasize from this Proposition is its conceptual similarity with our previous interpretation of Proposition 4. Recall that Proposition 4 indicated that a key element for the emergence of a (local) limit cycle was the S-shaped form of the function $F(\cdot)$ near the steady state. Proposition A.1 can be loosely interpreted as indicating that a limit cycle will emerge when the function $F(\cdot)$ looks s-shaped when looking from a global perspective.

While we believe that the upper bound condition $\lim_{I \rightarrow \bar{I}} F(I) = -\infty$ is quite natural in many economic applications, the existence of a lower bound driven by the condition $\lim_{I \rightarrow \underline{I}} F(I) = \infty$ is not as easy to justify. An alternative source of lower bound I is simply that it may be physically impossible to have negative investment. This type of constraint can be added to our system A.13, A.14 and A.15 by including the constraint that whenever $\alpha_0 + \alpha_2 I_t - \alpha_1 X_t < 0$, then $I_t = 0$, $\dot{I} = -\delta_1$ and $\dot{X} = -\delta$.

B Proofs of section 2

B.1 Proposition 5

Recall that

$$X_{t+1} = (1 - \delta)(X_t + e_t)$$

Since $e_t \geq 0$, if $\limsup_{t \rightarrow \infty} |X_t| = \infty$ then $\limsup_{t \rightarrow \infty} X_t = \infty$. Suppose then that

$$\limsup_{t \rightarrow \infty} X_t = \infty$$

Since $\delta \in (0, 1]$, this necessarily implies that $\limsup_{t \rightarrow \infty} e_t = \infty$. But e_t is bounded above by the level of output, the maximum feasible level of which occurs when $\phi_t = 1$ and $\ell_t = \bar{\ell}$, in which case total output is given by $F(\bar{\ell}) < \infty$. Thus we clearly cannot have $\limsup_{t \rightarrow \infty} e_t = \infty$, and thus we cannot have $\limsup_{t \rightarrow \infty} |X_t| = \infty$.

B.2 Proposition 6

The proof proceeds by example, showing that, for the case where $\gamma = 1$ and $U(c) = ac - \frac{b}{2}c^2$, there exists parameter values and functional forms such that for β close enough to one the steady state is unstable.

With $\gamma = 1$ and $U(c) = ac - \frac{b}{2}c^2$, we may characterize the evolution of this system by the conditions

$$a - b(X_t + e_t) = vp(e_t)[1 + \tau - \tau\phi(e_t)] - \beta(1 - \delta)vp(e_{t+1})[1 + \tau - \tau\phi(e_{t+1})] \quad (\text{B.16})$$

$$X_{t+1} = (1 - \delta)(X_t + e_t) \quad (\text{B.17})$$

where $p(\cdot)$ and $\phi(\cdot)$ are as in the static model. For a given state X_t and a given anticipated level of e_{t+1} , a sufficient condition to ensure that (B.16) has a unique solution is given by

$$b > vp^* \frac{\tau}{e^*} \equiv b_0 \quad (\text{B.18})$$

where e^* is output per firm (net of fixed costs) when the economy is in the unemployment regime and p^* is the price in the unemployment regime, as described in section 2.1 for the static model (see footnote 21 regarding p^*). I henceforth assume that (B.18) holds.

Next, the steady-state level of e is given by the solution \bar{e} to

$$a - \frac{b}{\delta}\bar{e} = [1 - \beta(1 - \delta)]vp(\bar{e})[1 + \tau - \tau\phi(\bar{e})]$$

with the steady-state level of X then given by

$$\bar{X} = \frac{1 - \delta}{\delta}\bar{e}$$

Note that a sufficient condition for the steady state to be unique is given by

$$b > \delta[1 - \beta(1 - \delta)]b_0$$

which is clearly implied by (B.18).

Next, note that, for any $e \in (0, e^*)$ (i.e., in the unemployment regime), the level of a that implements $\bar{e} = e$ is given by

$$\frac{b}{\delta}e + [1 - \beta(1 - \delta)]vp^* \left(1 + \tau - \tau\frac{e}{e^*}\right)$$

Note also that \bar{e} is continuous in β . Thus, choose some $\bar{e}_1 \in (0, e^*)$, and let $a = a_1$, where a_1 is the value of a that would implement $\bar{e} = \bar{e}_1$ when $\beta = 1$, i.e.,

$$a_1 \equiv \frac{b}{\delta}\bar{e}_1 + \delta vp^* \left(1 + \tau - \tau\frac{\bar{e}_1}{e^*}\right)$$

Thus, if $\beta = 1$ the steady state is in the unemployment regime by construction, and by continuity of \bar{e} in β the steady state is also necessarily in the unemployment regime for β sufficiently close to one. This implies the existence of a $\underline{\beta} < 1$ such that the steady state is

in the unemployment regime when $\beta > \underline{\beta}$. Assume henceforth that $\beta \in (\underline{\beta}, 1)$ and note that this implies that $p'(\bar{e}) = 0$ and $\phi'(\bar{e}) = 1/e^*$.

Next, linearizing equations (B.17)-(B.18) around this steady state and solving, we may obtain in matrix form

$$\begin{pmatrix} \hat{X}_{t+1} \\ \hat{e}_{t+1} \end{pmatrix} = \begin{pmatrix} 1 - \delta & 1 - \delta \\ -\frac{b}{\beta(1-\delta)b_0} & -\frac{b-b_0}{\beta(1-\delta)b_0} \end{pmatrix} \begin{pmatrix} \hat{X}_t \\ \hat{e}_t \end{pmatrix} \equiv A \begin{pmatrix} \hat{X}_t \\ \hat{e}_t \end{pmatrix}$$

Thus, the steady state is locally stable if and only if at least one of the two eigenvalues of A lies inside the complex unit circle. These eigenvalues are given by

$$\lambda_i = \frac{\left[1 - \delta - \frac{b-b_0}{\beta(1-\delta)b_0}\right] \pm \sqrt{\left[1 - \delta - \frac{b-b_0}{\beta(1-\delta)b_0}\right]^2 - 4\beta^{-1}}}{2}$$

Note that $\lambda_1\lambda_2 = \beta^{-1} > 1$, so that if the eigenvalues are complex then both must lie outside the unit circle. Suppose

$$b = [1 + q(1 - \delta)^2] b_0 \quad (\text{B.19})$$

for some $q > 0$, and note that as long as $\delta < 1$, which I henceforth assume, such a value of b satisfies (B.18). One may then show that the eigenvalues are complex as long as

$$(1 - \delta)^2 (\beta - q)^2 < 4\beta$$

Clearly, for β close enough to q this condition necessarily holds, and thus, if q is close enough to one (e.g., if $q = 1$), then for β arbitrarily close to one the eigenvalues are complex and therefore outside the unit circle, in which case the steady state is unstable.

B.3 Proposition 7

Let

$$\mathcal{V}(e_t; X_t) \equiv U(X_t + e_t) - vp^* \left[(1 + \tau) e_t - \frac{1}{2} \tau \frac{e_t^2}{e^*} \right]$$

where e^* is output per firm (net of fixed costs) when the economy is in the unemployment regime and p^* is the price in the unemployment regime, as described in section 2.1 for the static model (see footnote 21 regarding p^*). It can be verified that maximizing

$$\sum_{t=0}^{\infty} \beta^t \mathcal{V}(e_t; X_t)$$

subject to (13) implements the de-centralized equilibrium outcome in the neighborhood of an unemployment-regime steady state. Thus, using

$$\mathcal{W}(X_t, X_{t+1}) \equiv \mathcal{V} \left(\frac{1}{\gamma(1-\delta)} X_{t+1} - \frac{1}{\gamma} X_t; X_t \right)$$

in problem (16) satisfies the desired properties. Next, we may obtain

$$\mathcal{W}_{11}(\bar{X}, \bar{X}) = \frac{(1-\gamma)^2}{\gamma^2} U''(\bar{X} + \bar{e}) + \frac{1}{\gamma^2} \frac{vp^*\tau}{e^*}$$

Thus, $\mathcal{W}_{11}(\bar{X}, \bar{X}) > 0$ if

$$\frac{vp^*\tau}{e^*} > -(1-\gamma)^2 U''(\bar{X} + \bar{e})$$

This condition can clearly hold for certain parameter values (e.g., for γ sufficiently close to one), in which case \mathcal{W} is not concave.

B.4 Proposition 8

We show that the steady state is locally stable when $\tau = 0$. By continuity of all relevant functions, it then follows that the steady state is locally stable for $\tau > 0$ sufficiently small.

When $\tau = 0$, equilibrium is characterized by the equations

$$\begin{aligned} U'(X_t + e_t) - vp(e_t) &= \beta(1-\delta)(1-\gamma)U'(X_{t+1} + e_{t+1}) - \beta(1-\delta)vp(e_{t+1}) \\ X_{t+1} &= (1-\delta)(X_t + \gamma e_t) \end{aligned}$$

Assume the steady state is in the unemployment regime, so that $p(e_t) = p^*$ in a neighborhood of the steady state.⁶¹ Linearizing around this steady state, assuming $\beta(1-\delta)(1-\gamma) > 0$ we may obtain in matrix form

$$\begin{pmatrix} \hat{X}_{t+1} \\ \hat{e}_{t+1} \end{pmatrix} = \begin{pmatrix} 1-\delta & (1-\delta)\gamma \\ \frac{[1-\beta(1-\delta)^2(1-\gamma)]}{\beta(1-\delta)(1-\gamma)} & \frac{[1-\beta(1-\delta)^2(1-\gamma)\gamma]}{\beta(1-\delta)(1-\gamma)} \end{pmatrix} \begin{pmatrix} \hat{X}_t \\ \hat{e}_t \end{pmatrix} \equiv A \begin{pmatrix} \hat{X}_t \\ \hat{e}_t \end{pmatrix}$$

The steady state is locally stable if at least one of the eigenvalues of A lies inside the unit circle. It is straightforward to show that the smallest eigenvalue of A is given by $\lambda_1 = (1-\delta)(1-\gamma)$, which is clearly less than one in modulus. Thus, the steady state is locally stable. If instead $\beta(1-\delta)(1-\gamma) = 0$, then $\hat{e}_t = -\hat{X}_t$ and thus $\hat{X}_{t+1} = \lambda_1 \hat{X}_t$, which is clearly a stable system as well.

C Solution and estimation

C.1 Solution

To solve the model for a given parameterization, letting $\tilde{e}_t \equiv e_t/\tilde{\theta}_t$ equilibrium in the economy is characterized by the following equations:

$$a - b \left(X_t + \tilde{\theta}_t \tilde{e}_t - hc_{t-1} \right) + (1-\delta)\gamma\lambda_t = \tilde{\theta}_t^{-1} \frac{\nu_1}{\alpha A} [\ell(\tilde{e}_t)]^{\omega+1-\alpha} [1 + \tau - \tau\phi(\tilde{e}_t)] + \mu_t \quad (\text{C.20})$$

$$\mu_t = \mathbb{E}_t \left\{ \beta h \left[a - b \left(X_{t+1} + \tilde{\theta}_{t+1} \tilde{e}_{t+1} - hc_t \right) \right] \right\} \quad (\text{C.21})$$

$$\lambda_t = \mathbb{E}_t \left\{ \beta \left[a - b \left(X_{t+1} + \tilde{\theta}_{t+1} \tilde{e}_{t+1} - hc_t \right) + (1-\delta)\lambda_{t+1} - \mu_{t+1} \right] \right\} \quad (\text{C.22})$$

$$c_t = X_t + \tilde{\theta}_t \tilde{e}_t \quad (\text{C.23})$$

⁶¹It is straightforward to verify that a full-employment-regime steady state must be stable.

$$X_{t+1} = (1 - \delta) \left(X_t + \gamma \tilde{\theta}_t \tilde{e}_t \right) \quad (\text{C.24})$$

Here, $\phi(\tilde{e})$ and $\ell(\tilde{e})$ are the equilibrium levels of the employment rate and hours-per-worker conditional on total purchases \tilde{e} , and are given by

$$\phi(\tilde{e}) \equiv \begin{cases} \frac{1}{2} \left(n_0 + \sqrt{n_0^2 + 4\eta \frac{\tilde{e}}{\bar{e}^*}} \right) & \text{if } 0 < \tilde{e} \leq \bar{e} \\ \frac{\tilde{e}}{\bar{e}^*} & \text{if } \bar{e} < \tilde{e} < e^* \\ 1 & \text{if } \tilde{e} \geq e^* \end{cases}$$

$$\ell(\tilde{e}) \equiv \begin{cases} \left[\frac{2\tilde{e}}{\alpha A \left(n_0 + \sqrt{n_0^2 + 4\eta \frac{\tilde{e}}{\bar{e}^*}} \right)} \right]^{\frac{1}{\alpha}} & \text{if } 0 < \tilde{e} \leq \bar{e} \\ \left(\frac{e^*}{\alpha A} \right)^{\frac{1}{\alpha}} & \text{if } \bar{e} < \tilde{e} < e^* \\ \left(\frac{\tilde{e}}{\alpha A} \right)^{\frac{1}{\alpha}} & \text{if } \tilde{e} \geq e^* \end{cases}$$

where $e^* \equiv \frac{\alpha}{1-\alpha} \bar{k}$ and $\bar{e} \equiv (n_0 + \eta) e^*$. Meanwhile, μ_t and λ_t are the Lagrange multipliers on the definition of consumption and the durables accumulation equations ((C.23) and (C.24)), respectively.

Conditional on the state variables X_t , c_{t-1} and θ_t , and on values of the Lagrange multipliers μ_t and λ_t , equation (C.20) can be solved for \tilde{e}_t . To obtain values of μ_t and λ_t , I employ the method of parameterized expectations as follows. Let $Y_t \equiv (X_t - \bar{X}, c_{t-1} - \bar{c}, \theta_t)'$ denote the vector of state variables (expressed as deviations from steady state). The expectations in equations (C.21) and (C.22) are assumed to be functions only of Y_t , i.e.,

$$\mathbb{E}_t \left\{ \beta h \left[a - b \left(X_{t+1} + \tilde{\theta}_{t+1} \tilde{e}_{t+1} - h c_t \right) \right] \right\} = g_\mu(Y_t)$$

$$\mathbb{E}_t \left\{ \beta \left[a - b \left(X_{t+1} + \tilde{\theta}_{t+1} \tilde{e}_{t+1} - h c_t \right) + (1 - \delta) \lambda_{t+1} - \mu_{t+1} \right] \right\} = g_\lambda(Y_t)$$

I parameterize the functions $g_j(\cdot)$ by assuming that they are well-approximated by N -th-degree multivariate polynomials in the state variables. In particular, let $Y_t^{(N)}$ denote the vector whose first element is 1 and whose remaining elements are obtained by collecting all multivariate polynomial terms in Y_t (e.g., X_t , c_{t-1} , θ_t , X_t^2 , $X_t c_{t-1}$, $X_t \theta_t$, c_t^2 , $c_t \theta_t$, etc.) up to degree N . I assume that

$$g_j(Y_t) = \Theta_j' Y_t^{(N)}$$

where Θ_j is a vector of coefficients on the polynomial terms. Thus, given Θ_μ , Θ_λ and the state Y_t , μ_t and λ_t are obtained as

$$\mu_t = \Theta_\mu' Y_t^{(N)}$$

$$\lambda_t = \Theta_\lambda' Y_t^{(N)}$$

These values and values for the state variables can be plugged into (C.20) to yield a solution for \tilde{e}_t , which can then be replaced in (C.23) and (C.24) to obtain values for the subsequent period's state. In practice, I use $N = 2$.⁶²

⁶²I experimented with larger values of N and found that it resulted in a substantial increase in computational time without significantly affecting the results.

To obtain Θ_μ and Θ_λ , I proceed iteratively as follows. Begin with some initial guesses $\Theta_{\mu,0}$ and $\Theta_{\lambda,0}$,⁶³ and generate a sample of length $T = 100,000$ of the exogenous process θ_t . Next, given $\Theta_{\mu,i}$ and $\Theta_{\lambda,i}$, assume that $g_j(Y_t) = \Theta'_{j,i} Y_t^{(N)}$ and simulate the path of the economy for T periods. Given this simulated path, let $\mathbf{Y}^{(N)}$ denote the matrix whose t -th row is given by $Y_t^{(N)'$, and construct T -vectors \tilde{g}_μ and \tilde{g}_λ , the t -th elements of which are given respectively by

$$\beta h \left[a - b \left(X_{t+1} + \tilde{\theta}_{t+1} \tilde{e}_{t+1} - hc_t \right) \right]$$

and

$$\beta \left[a - b \left(X_{t+1} + \tilde{\theta}_{t+1} \tilde{e}_{t+1} - hc_t \right) + (1 - \delta) \lambda_{t+1} - \mu_{t+1} \right]$$

i.e., the terms inside the conditional-expectation operators in equations (C.21) and (C.22). Then update the guesses of Θ_j via

$$\Theta_{j,i+1} = (\mathbf{Y}^{(N)'} \mathbf{Y}^{(N)})^{-1} \mathbf{Y}^{(N)'} \tilde{g}_j$$

and iterate until convergence.

C.2 Estimation

As discussed in section 3.3, estimation was done by searching for parameters to minimize $\overline{S^2}$, the average squared difference between the model spectrum and the spectrum estimated from the data.

To obtain $\overline{S^2}$ given a solution to the model for a parameterization, $T = 100,000$ periods of data were simulated. This simulated sample was then subdivided into $N_{sim} = 1,000$ overlapping subsamples. For each subsample, the log of hours was BP-filtered, after which 20 quarters from either end of the subsample were removed, leaving a series of the same length as the actual data sample. The spectrum was then estimated on each individual subsample in the same way as for the actual data, and the results then averaged across all subsamples to yield the spectrum for the model.

D Solving the model forward

In the non-stochastic case, we may re-arrange equation (C.20) to yield

$$a - bX_t + bhc_{t-1} + (1 - \delta) \gamma \lambda_t - \mu_t = \frac{\nu_1}{\alpha A} [l(e_t)]^{\omega+1-\alpha} [1 + \tau - \tau \phi(e_t)] + be_t \equiv H(e_t)$$

Thus, given the state variables X_t and c_{t-1} and current values of μ_t and λ_t , we may obtain

$$e_t = H^{-1}(a - bX_t + bhc_{t-1} + (1 - \delta) \gamma \lambda_t - \mu_t)$$

⁶³In practice, I set the first elements of $\Theta_{\mu,0}$ and $\Theta_{\lambda,0}$ to the steady-state values $\bar{\mu}$ and $\bar{\lambda}$, respectively, and the remaining elements to zero. This corresponds to an initial belief that the g_j 's are constant and equal to their steady-state levels.

where the conditions in Proposition 9 ensure that H is an invertible function. This value of e_t then gives c_t and X_{t+1} via equations (C.23) and (C.24), respectively. From equations (C.21) and (C.22) we can then solve for μ_{t+1} and λ_{t+1} as

$$\mu_{t+1} = \frac{a - bX_{t+1} + bhc_t - H\left(\frac{1}{b}\left(a - bX_{t+1} + bhc_t - \frac{1}{\beta h}\mu_t\right)\right) - \gamma\left(\frac{1}{\beta h}\mu_t - \frac{1}{\beta}\lambda_t\right)}{1 - \gamma}$$

$$\lambda_{t+1} = \frac{a - bX_{t+1} + bhc_t - H\left(\frac{1}{b}\left(a - bX_{t+1} + bhc_t - \frac{1}{\beta h}\mu_t\right)\right) - \left(\frac{1}{\beta h}\mu_t - \frac{1}{\beta}\lambda_t\right)}{(1 - \delta)(1 - \gamma)}$$