

Reviving the Limit Cycle View of Macroeconomic Fluctuations

Paul Beaudry*, Dana Galizia† and Franck Portier‡§

July 2015

Abstract

There is a long tradition in macroeconomics suggesting that market imperfections may explain why economies repeatedly go through periods of booms and busts, with booms sowing the seeds of the subsequent busts. This idea can be captured mathematically as a limit cycle. For several reasons, limit cycles play almost no role in modern quantitative business cycle analysis. In this paper we present both a general structure and a particular model with the aim of giving new life to this mostly dismissed view of fluctuations. We begin by showing why and when models with strategic complementarities—which are quite common in macroeconomics—can give rise to unique-equilibrium dynamics characterized by a limit cycle. We then develop and estimate a fully-specified dynamic general equilibrium model that embeds a demand complementarity to see whether the data favors a configuration supportive of a limit cycle. Booms and busts arise endogenously in our setting because agents want to concentrate their purchases of goods at times when purchases by others are high, since in such situations unemployment is low and therefore taking on debt is perceived as being less risky. A key feature of our approach is that we allow limit-cycle forces to compete with exogenous disturbances in explaining the data. Our estimation results indicate that US business cycle fluctuations in employment and output can be well explained by endogenous demand-driven cycles buffeted by technological disturbances that render those fluctuations irregular.

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

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

†Carleton University

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

§The authors thank Jess Benhabib and Kiminori Matsuyama for helpful discussions. The authors would also like to thank seminar participants at Banque de France, Pompeu Fabra-Toulouse “skiminar”, NBER Summer Institute, UCLA, UCSD, the University of Pennsylvania and Northwestern University for comments.

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 produce cyclical phenomena; that is, in the absence of any shocks (including shocks to agents’ beliefs) economic forces by themselves may favor recurrent periods of high economic activity followed by periods of low economic activity. This sort of outcome will arise, for example, if the underlying economic system generates a *limit cycle*.¹ In such a framework, irregular business cycles can emerge from these underlying regular forces when 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 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].² In the 1970s and 1980s, a larger literature emerged that examined the conditions under which qualitatively and quantitatively reasonable economic fluctuations might occur in purely deterministic settings (see, e.g., Benhabib and Nishimura [1979] and [1985], Day [1982] and [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, the interest in such models for understanding business cycle fluctuations greatly diminished and became quite removed from the mainstream research agenda.³

¹Informally, limit cycles are indefinitely repeating fluctuations that are capable of emerging in some dynamic systems.

²An earlier mention of self-sustaining cycles as a model for economic fluctuations is found in Le Corbeiller [1933] in the first volume of *Econometrica*.

³There are at least two strands of the macroeconomic literature that has productively continued to pursue ideas related to limit cycles: a literature on innovation cycles and growth (see, for, example Shleifer [1986] and Matsuyama [1999]), and a literature on endogenous credit cycles in an OLG setting (see, for example, Azariadis and Smith [1998], Matsuyama [2007] and [2013], Myerson [2012] and Gu, Mattesini, Monnet, and Wright [2013]).

There appear to be several key reasons why interest in limit cycles 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 optimizing, forward-looking agents.⁴ 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 micro-founded models.

Second, in the category of models featuring optimizing forward-looking agents, the primary focus was on models with a neoclassical, competitive-equilibrium structure.⁵ 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, under certain basic conditions met by these models and holding all other parameters constant, 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.⁶ 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 such models by Boldrin and Woodford [1990], discount factors for several of the models they discuss were on the order of 0.3 or less.⁷ 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—then a discount factor arbitrarily close to one can relatively easily support de-

⁴The first category includes, for example, Benhabib and Nishimura [1979] and [1985], and Boldrin and Montrucchio [1986], while the latter includes, for example, the pre-1970s examples cited above, as well as Day [1982] and [1983].

⁵While 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], Coles and Wright [1998], Lagos and Wright [2003], Rocheteau and Wright [2013] and Gu, Mattesini, Monnet, and Wright [2013]. 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 limit cycles), rather than being focused on cycles directly.

⁶See the discussion in section 2.3 for further details.

⁷It 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.

terministic fluctuations in equilibrium.

Third, 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.⁸ This can be observed by looking at the spectrum of data generated by such models, which will generally feature one or more large spikes at frequencies associated with k -period cycles. Spectra estimated on macroeconomic data generally lack such spikes,⁹ which suggests much less regularity in real-world cycles. To address this issue, papers from the earlier literature frequently sought to establish conditions under which such irregular cycles could emerge in a purely deterministic setting via chaotic dynamics.¹⁰ While in a number of cases this was found to be possible, the conditions are often more restrictive than those required to generate simple periodic cycles. In contrast, rather than restricting attention to a purely deterministic setting, this paper embeds a limit cycle mechanism into a stochastic environment in which irregular fluctuations emerge as the interplay between exogenous shocks and endogenous cycles.

Finally, being inherently 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 forces¹¹—have gone largely 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 limit cycles tend to arise quite naturally in the presence of demand complementarities.

⁸In Boldrin and Woodford [1990], the authors mention private communication with Sir John Hicks in which Hicks indicated that he lost interest in endogenous cycle models because actual business cycles were far from being regular periodic motions.

⁹See Figure 5.

¹⁰Roughly, chaotic fluctuations appear in systems where the orbits emanating from two different initial points typically cannot be made arbitrarily close by choosing those initial points sufficiently close together. Chaotic systems “appear” random, despite being entirely deterministic.

¹¹One notable paper that combines these forces is Benhabib and Nishimura [1989], who consider the theoretical question of how such a combination would affect the joint probability distribution of the variables of the system.

In particular, we clarify why demand complementarities, when they relate to a good that can be accumulated, are more likely to give rise to local instability and limit cycles than to multiple equilibrium. To make this point, we first present a simple but fairly general reduced-form setup. Our second goal is to present a particular but fully specified model wherein limit cycles can arise as the result of demand complementarities stemming from an interplay between unemployment risk and precautionary savings. We estimate a DSGE version of this model to see how it chooses to explain features of the data.

The outline of the paper is as follows. In Section 1, we present a simple reduced-form dynamic model to highlight what type of interactions between agents tends to give rise to limit cycles. In particular, we show how and when demand complementarities can cause the steady state of the system to become unstable and a limit cycle to appear around it as part of a Hopf bifurcation,^{12,13} and further establish a simple condition under which this limit cycle will be attractive. In Section 2, we introduce an optimization-based model which incorporates unemployment risk that in turn causes agents’ consumption decisions to be strategic complements as a result of precautionary behavior. The model builds on Beaudry, Galizia, and Portier [2015]. We will show that this explicit model exhibits the main features emphasized by the reduced-form model. Section 3 takes the model to the data to see whether estimation will reveal the presence of a limit cycle. Since our estimation framework allows the limit cycle to compete with exogenous disturbances in explaining the data, we will be able to assess the extent to which it can reduce the reliance on exogenous disturbances in explaining business cycle fluctuations. Finally, in the last section we offer concluding comments.

¹²Since our model is formulated in discrete time, the bifurcation we consider is more appropriately referred to as a Neimark-Sacker (rather than Hopf) bifurcation. Nonetheless, we will typically follow convention in applying the term “Hopf bifurcation” to both continuous and discrete environments.

¹³Informally, a Hopf bifurcation—which may occur in both continuous and discrete formulations—is characterized by a loss of stability in which the resulting limit cycle involves rotation around the steady state in two-dimensional phase space. Discrete (but not continuous) systems may also produce limit cycles in a *one*-dimensional setting via a “flip” bifurcation, in which case the system “jumps” back and forth over the steady state. As discussed further in Section 1, for several reasons we will largely focus our discussion around the dynamics associated with Hopf bifurcations.

1 Demand Complementarities and Limit Cycles: A Reduced-Form Model

In this section we present a simple reduced-form dynamic model aimed at illustrating how and when limit cycles can emerge in an environment with demand complementarities. As we show, the key mechanism that allows the model to generate limit cycles is the interplay between the demand complementarities and the dynamic forces associated with the fact that the good in the model can be accumulated. We begin with a reduced-form setup so as to highlight the generality of this mechanism, regardless of the source of the precise micro-foundations that drive agents' interactions. In the next section we will present a structural model that rationalizes this set-up, and in which this key mechanism emerges naturally.

In addition to establishing that the model can, under fairly general conditions, produce a Hopf bifurcation associated with an attractive limit cycle, an important goal of the analysis in this section is to show that this happens *even when we restrict the strength of the demand complementarities to be too weak to create static multiple equilibria*. In fact, as we make clear, the strength of the demand complementarities necessary for a limit cycle to appear in this environment is *always* less than that needed to generate multiple equilibrium.

1.1 The Environment

Consider an environment with a large number N of agents indexed by i , where each agent can accumulate a good X_{it} , which can be 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 are no interactions between agents and that the decision rule for agent i 's investment is given simply by

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

where all parameters are strictly 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 so as to reflect some underlying

decreasing returns to capital accumulation, while the effect of past investment is positive so as to reflect sluggish response that may be due, for example, to adjustment costs.¹⁴

When all agents behave symmetrically, the aggregate dynamics of the economy are 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 of this system is established in the following proposition.

Proposition 1. *Both eigenvalues of the matrix M_L lie strictly inside the unit circle. Therefore, system (3) is stable.*

All proofs are given in the appendix. According to Proposition 1, the dynamics are extreme simple, with the system converging to its steady state for any starting values of $X_{it} = X_t$ and $I_{it-1} = I_{t-1}$. We now add agent interactions to the model and study how is the dynamics are affected.

1.2 Adding Interactions Between Agents


To generalize the previous setup in order to allow for interactions between individuals, we modify the investment decision rule to

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

while keeping the law of motion for X (equation (1)) unchanged. We assume that the function $F(\cdot)$ is continuous and differentiable at least three times and that $F(0) = 0$. The function $F(\cdot)$ captures how the actions of others, summarized by the average level of investment $I_t \equiv \sum I_{jt}/N$, affect agent i 's investment decision I_{it} . For example, the function $F(\cdot)$ can capture implicitly the price increase induced by the demand of others if $F'(\cdot) < 0$, or can capture demand complementarities if $F'(\cdot) > 0$. In this formulation, we are assuming that agents take the average actions in the economy as given, so that (4) can be interpreted as agent i 's best-response rule to the average action.

¹⁴We focus here on a system with two state variables. We do this since we need a system with at least two state variables for the possibility of a Hopf bifurcation to arise. It is our conjecture that limit cycles that emerge from Hopf bifurcations are more likely to be relevant for business cycle analysis than those arising from flip bifurcations (which can arise in one-dimensional systems).

Figure 1: Best-Response Rule for Two Different Histories



Notes: This figure plots the best-response rule (equation (4)), $I_{it} = \alpha_t + F(I_t)$, with $\alpha_t \equiv \alpha_0 - \alpha_1 \sum_{j=0}^{\infty} (1 - \delta)^j I_{it-1-j} + \alpha_2 I_{it-1}$. The intercepts α_t and α'_t correspond to two different histories of the model.

Figure 1 illustrates this best-response rule for two different values of the intercept. Note that in this diagram the intercept is given by $\alpha_t = \alpha_0 - \alpha_1 \sum_{j=0}^{\infty} (1-\delta)^j I_{it-1-j} + \alpha_2 I_{it-1}$, which is a function of the history of the economy prior to date t . Thus, past investment decisions determine the location of the best-response rule, which in turn determines the current level of investment, which then feeds into the determination of the next period's intercept, and so on. In order to rule out static multiple equilibria—that is, multiple solutions for I_t for given values of I_{it-1} and X_{it} —we assume that $F'(I_t) < 1$.¹⁵ Thus, we are restricting attention to cases where demand complementarities, if they are present, are never strong enough to produce static multiple equilibrium.

In what follows we consider only symmetric equilibria, so that we may henceforth drop the subscript i . We make the additional assumption that $\alpha_2 < \alpha_1/\delta$ so that a steady state necessarily exists and is unique, and let I^s and X^s denote the steady state values of I and X . Note that the condition $\alpha_2 < \alpha_1/\delta$ will always be satisfied when δ is sufficiently small, since both α_1 and α_2 are strictly positive.

Our goal is to examine how the dynamics of the system (1) and (4) are affected by the properties of the interaction effects, and especially what conditions on $F(\cdot)$ will give rise to a Hopf bifurcation that is associated with the emergence of an attractive limit cycle.

1.3 The Local Dynamics of the Model with Interactions Between Agents

We now consider the dynamics implied by the bivariate system (1) and (4). In order to understand those dynamics, it is useful to first look at local dynamics in the neighborhood of the steady state. The first-order approximation of this dynamic system is given by

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

The eigenvalues of the matrix M are the solutions to the quadratic equation

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

¹⁵Under the condition $F'(I_t) < 1$, the static equilibrium depicted in Figure 1 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 M (and also the sum of its eigenvalues) and D is the determinant of M (and also the product of its eigenvalues). The two eigenvalues are thus given by

$$\lambda, \bar{\lambda} = \frac{T}{2} \pm \sqrt{\left(\frac{T}{2}\right)^2 - D} \quad (6)$$

where

$$T \equiv \frac{\alpha_2 - \alpha_1}{1 - F'(I^s)} + 1 - \delta \quad (7)$$

and

$$D \equiv \frac{\alpha_2(1 - \delta)}{1 - F'(I^s)}. \quad (8)$$

When $F'(I^s) = 0$, the model dynamics are locally the same as in the model without demand complementarities, and in particular, as noted in Proposition 1, the roots of the system lie inside the unit circle in this case. More generally, we may (locally) parameterize F by $F'(I^s)$ and ask what happens as F' varies.

Geometric analysis: It is informative to consider a geometric analysis of the location of the two eigenvalues of the linearized system (5). This is done in Figure 2, which presents the plane (T, D) . A point in this plane is a pair (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 \overline{BC} corresponds to $D = 1$, the line \overline{AB} to the equation $Q(-1) = 0$ ($\Leftrightarrow D = -T - 1$) and the line \overline{AC} to the equation $Q(1) = 0$ ($\Leftrightarrow D = T - 1$). On the perimeter of triangle \widehat{ABC} , at least one eigenvalue has a modulus of one. We review in the appendix why both eigenvalues of the system are inside the unit circle when (T, D) is inside \widehat{ABC} , while at least one eigenvalue is outside the unit circle when (T, D) is outside \widehat{ABC} . Whether the eigenvalues are real or complex depends on the sign of the discriminant of the equation $Q(\lambda) = 0$, which is given by $\Delta \equiv \left(\frac{T}{2}\right)^2 - D$. The parabola in Figure 2 corresponds to the equation $\Delta = 0$ ($\Leftrightarrow D = T^2/4$). Above the parabola, the eigenvalues are complex and conjugate, while on or below the parabola they are real. The possible configurations of the local dynamics can then be easily characterized by considering the location of (T, D) within this diagram.

Without any restrictions on $F'(I^s)$, the steady state can be locally stable, unstable, or a saddle, with either real or complex eigenvalues. Proposition 1 proves that when $F'(I^s) = 0$,

the steady state values correspond to points E or E' (depending on the parameters) that are inside \widehat{ABC} . Furthermore, it is easy to verify that D is always greater than zero, so that the steady state must lie in the top part of \widehat{ABC} . As an example, in Figure 2 we have put E and E' in the region of stability with complex eigenvalues.

As $F'(I^s)$ varies, the eigenvalues of the system will vary, implying changes to the dynamic behavior of I and X . From equations (7) and (8), assuming $\alpha_1 \neq \alpha_2$,¹⁶ we may obtain the following relationship between the trace and determinant of matrix M :

$$D = \frac{\alpha_2(1-\delta)}{\alpha_2 - \alpha_1}T - \frac{\alpha_2(1-\delta)^2}{\alpha_2 - \alpha_1}. \quad (9)$$

Therefore, when $F'(I^s)$ varies, T and D move along the line (9) in (T, D) -space, which allows for an easy characterization of the impact of $F'(I^s)$ on the location of the eigenvalues, and therefore also of the stability of the steady state. We need to systematically consider the two cases $\alpha_2 > \alpha_1$ and $\alpha_2 < \alpha_1$, as the line (9) slopes positively in the former case and negatively in the latter.

Let us consider first the strategic substitutability case where $F'(I^s)$ is negative; that is, where the investment decisions of others have a negative effect on one's own decision. In that case, it is clear from equations (7) and (8) that

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

and


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

We will denote the point $(1 - \delta, 0)$ by E_1 on Figure 2. Note that E_1 lies inside the “stability” triangle \widehat{ABC} .

Let's assume that the steady state without strategic interactions is E and $\alpha_2 > \alpha_1$. When $F'(I^s)$ goes from 0 to $-\infty$, the point (T, D) moves from E to E_1 along the line given by equation (9). This movement corresponds to the half-line denoted (a) on Figure 2. As E and E_1 both belong to the interior of \widehat{ABC} , and because the interior of \widehat{ABC} is a convex set, any point of the segment $[E, E_1]$ also belongs to the interior of triangle \widehat{ABC} . The same argument applies if parameters are such that the steady state corresponds to E' and $\alpha_2 < \alpha_1$, with the relevant half-line being (a') . Thus, the following proposition holds:

¹⁶The case where $\alpha_1 = \alpha_2$ is addressed in the appendix.

Figure 2: Local Stability when $F'(I^s) \in (-\infty, 1)$



Notes: This figure shows the plane (T, D) , where T is the trace and D the determinant of matrix M . The points E and E' correspond to two possible configurations of the model without demand externalities. According to Proposition 1, those points are inside the triangle of stability \widehat{ABC} . They are arbitrarily placed in the zone where the two eigenvalues are complex. E and arrows (a) and (b) correspond to the case where $\alpha_2 > \alpha_1$; E' and arrows (a') , (b') and (b'') to the case where $\alpha_2 < \alpha_1$. The case $\alpha_1 = \alpha_2$ is investigated in the appendix.

Proposition 2. *As $F'(I^s)$ varies from 0 to $-\infty$, the eigenvalues of M always stay within the complex 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's own action, the system always maintains stability. Since Walrasian settings are typically characterized by strategic substitutability, this is one reason why dynamic Walrasian environments are generally stable.

We now turn to exploring how the presence of strategic complementarities (i.e., when $F'(I^s) > 0$) affects the dynamics of the system. In Figure 2, a rise in $F'(I^s)$ beginning from $F'(I^s) = 0$ corresponds to movement along the line given by equation (9), starting from the point E or E' and in the opposite direction to E_1 . This movement is denoted by the half-lines (b), (b') or (b'') in Figure 2. It can also be easily verified that, as $F'(I^s)$ gets closer to one, (T, D) will necessarily cross the perimeter of triangle \widehat{ABC} , thereby causing the steady state to change from being locally stable to being unstable. The location of the crossing depends on parameters.

Consider first the case where $\alpha_2 > \alpha_1$, and assume that the steady state corresponds to E on Figure 2. Under our earlier assumption that $\alpha_2 < \alpha_1/\delta$ (which helped guarantee a unique steady state), it can be verified that the half line (b) will never cross line segment \overline{AC} , which is the line segment associated with the largest of two real eigenvalues being equal to one. Thus, when $\alpha_2 > \alpha_1$, (b) must cross line segment \overline{BC} ; that is, the point at which the system loses stability as $F'(I^s)$ increases must be associated with complex eigenvalues. Now consider the case where $\alpha_2 < \alpha_1$, and assume that the steady state corresponds to E on Figure 2. In this case, (T, D) can cross the perimeter of triangle \widehat{ABC} under two different possible configurations. If the slope of line (9) is sufficiently negative, it will cross line segment \overline{BC} , at which point the eigenvalues will be complex. This is the case drawn in the figure as half line (b'). On the other hand, if the slope of line (9) is negative but sufficiently flat, the point $F'(I^s) = 0$ will be associated with a point in (T, D) -space like E' rather than in E . In that case, as shown by half-line (b'') in Figure 2, as $F'(I^s)$ increases (T, D) will cross line segment \overline{AB} , which is the line segment associated with the smaller of two real eigenvalues being equal to negative one.

1.4 Bifurcations and the Occurrence of Limit Cycles

Our graphical analysis of the previous subsection shows that the presence of demand complementarities can change the qualitative dynamics of the system. In particular, if the complementarities become strong enough the system will transition from being locally stable to being locally unstable. In the theory of dynamical systems, a change in local stability when a parameter varies is referred to as a bifurcation. Bifurcations are of interest since there is a close relationship between the nature of a bifurcation and the emergence of limit cycles. We formalize the nature of possible bifurcations for our bivariate system (1) and (4) in the following proposition.

Proposition 3. *As $F'(I^s)$ increases from zero towards one, the dynamic system given by (1) and (4) will eventually become locally unstable. In particular:*

- *If $\alpha_2 > \alpha_1/(2 - \delta)^2$, then a Hopf (Neimark-Sacker) bifurcation will occur.*
- *If $\alpha_2 < \alpha_1/(2 - \delta)^2$, then a flip bifurcation will occur.*

A flip bifurcation occurs with the appearance of an eigenvalue equal to negative one, and a Hopf bifurcation with the appearance of two complex conjugate eigenvalues of modulus one. The proof of Proposition 3 is given in the appendix, and involves establishing conditions under which line (9) crosses line segment \overline{BC} (for the Hopf bifurcation) or line segment \overline{AB} (for the flip bifurcation) as $F'(I^s)$ increases from zero. In either case, the bifurcation that occurs as $F'(I^s)$ increases will be associated with the emergence of a limit cycle. In the case of a flip bifurcation, the limit cycle that emerges close to the bifurcation point will be of period two, i.e., will involve jumps back and forth over the steady state every period. Such extreme fluctuations are unlikely to be very relevant for business cycle analysis. We therefore henceforth focus on the more interesting case from our point of view, which is the case where the system experiences a Hopf bifurcation. In this case, close to the bifurcation point there will emerge around the steady state (in (X, I) -space) a unique isolated closed invariant curve.¹⁷ Beginning from any point on this closed curve, the system will remain

¹⁷We have to this point side-stepped a minor technical issue, which is that, in discrete time, the bounded non-convergent orbits appearing near a Hopf bifurcation will possess the basic qualitative features of a limit cycle, but may never exactly repeat themselves and thus may not strictly speaking be limit cycles. For

on it thereafter, neither converging to a single point nor diverging to infinity, but instead rotating around the steady state along that curve indefinitely. Further, in contrast to the flip bifurcation, the cycle that emerges near a Hopf bifurcation may be of any period length and hence the resulting dynamics appear more promising for understanding business cycles.

The condition on α_2 under which a Hopf bifurcation (rather than a flip bifurcation) will arise according to Proposition 3 may, at first pass, look rather restrictive. In fact, as $\delta \rightarrow 1$ this condition approaches $\alpha_2 \geq \alpha_1$, which necessitates a fairly large amount of sluggishness in investment. However, if δ is small, the condition becomes significantly less restrictive. For example, as $\delta \rightarrow 0$ the condition becomes $\alpha_2 > \alpha_1/4$, a simple lower bound on the degree of sluggishness. Given these considerations, one could loosely re-state Proposition 2 as indicating that if depreciation is not too fast, and sluggishness not too small, then the system will experience a Hopf bifurcation as $F'(I^s)$ increases from 0 towards 1.¹⁸

Having established conditions under which a Hopf bifurcation emerges, we turn now to the question of whether such a limit cycle is attractive; that is, whether the economy would be expected to converge towards such an orbit given an arbitrary starting point. To use language from the theory of dynamical systems, a bifurcation may be either *supercritical* or *subcritical*. In a supercritical bifurcation, the limit cycle emerges on the “unstable” side of the bifurcation and attracts nearby orbits, while in a subcritical bifurcation the limit cycle emerges on the “stable” side of the bifurcation and repels nearby orbits.

The emergence of a limit cycle is mainly of interest to us if it is attractive, so that departures from the steady state will approach the limit cycle over time. The conditions governing whether a Hopf bifurcation is supercritical or subcritical are often hard to state. However, in our setup, a simple sufficient 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 [1978] theorem.

example, in a bivariate discrete-time system characterized by rotation along the unit circle by θ radians per period, if θ/π is irrational then the system will never return to the same point twice, despite the clear sense in which the dynamics are cyclical. We will ignore this uninteresting technicality and apply the term “limit cycle” broadly to also include the entire isolated closed invariant curve (e.g., the entire unit circle in the example of this footnote) that emerges.

¹⁸Note also that it is precisely when δ is small that our earlier condition on α_2 guaranteeing a unique steady state (i.e., $\alpha_2 < \alpha_1/\delta$) is also likely to be satisfied.

Proposition 4. *If $F'''(I^s)$ is sufficiently negative, then the Hopf bifurcation noted in Proposition 3 will be supercritical.*

The economics for why increasing $F'(I^s)$ will cause the system to become unstable is rather intuitive. A high value of $F'(I^s)$ implies that an individual agent has a large incentive to accumulate more capital at times when other agents increase their accumulation. This leads to a feedback effect whereby any initial individual desire to have high current investment—due to some combination of a low current capital stock (the decreasing returns channel) and a high level of investment in the previous period (the sluggishness channel)—becomes amplified in equilibrium through a multiplier-type mechanism. When this feedback effect is strong enough, it will cause small initial deviations from the steady state to grow over time, pushing the system away from the fixed point. As a result, the economy will tend to go through repeated episodes of periods of high accumulation followed by periods of low accumulation, even in the absence of any exogenous shocks. Such behavior contrasts sharply with the steady flow of I over time that 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 an attractive limit cycle can also be related to economic forces. If the best-response function, F , is positively sloped near the steady state and $F'''(I^s)$ is negative, then it will take an S-shaped form.¹⁹ Note that Figure 1 was drawn with these features. The intuition for why an S-shaped best-response function favors the emergence of an attractive limit cycle can be understood as follows. As noted above, when the system is locally unstable, the demand complementarities are strong enough near the steady state that any perturbation from that point will tend to induce outward “explosive” forces. If the best-response function is S-shaped, however, then as the system moves away from the steady state the demand complementarities will eventually fade out (i.e., $F'(I)$ eventually falls), so that the explosive forces that are in play near the steady state are gradually replaced with inward “stabilizing” ones. As long as $F'''(I^s)$ is sufficiently negative, these stabilizing forces will emerge quickly enough, and an attractive limit cycle will appear at the boundary between the inner explosiveness region

¹⁹A 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.

and the outer stability region. Such a configuration for F is the one we have assumed when drawing Figure 1.

If instead the best-response function has $F'''(I^s) > 0$, then instead of dying out, the demand complementarities would tend to *grow* in strength as the system moves away from the steady state,²⁰ so that inward stabilizing forces do not appear. In this case, when $F'(I^s)$ is large enough for the system to become unstable, the Wan [1978] theorem implies the presence of a subcritical Hopf bifurcation in which a repulsive limit cycle appears just before the system becomes unstable.²¹

The general insight we take away from the Wan [1978] theorem regarding Hopf bifurcations is that attractive 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 may refer to such a setting 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 becomes sufficient large, some resource constraints are likely to become binding, causing strategic substitutability to emerge in place of complementarities. Similarly, physical constraints, such as a non-negativity restrictions on investment and capital or Inada conditions implying that the marginal productivity of capital tends to infinity at zero are reasonable considerations in economic environments that will limit systems from diverging to zero or to negative activity. Such forces will in general favor the emergence of attractive limit cycles in the presence of demand complementarities.

²⁰A parametric example of such a 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 $y \in (0, 1)$.

²¹Note that subcriticality of a bifurcation does not necessarily imply global explosiveness on the unstable side of the bifurcation, nor does it rule out the emergence of an attractive limit cycle in that region. Rather, the results of this section (including those based on Wan [1978]) are inherently about the local behavior of the system, where “local” in this case means “to a third-order Taylor approximation on some sufficiently small neighborhood of the steady state”. Conclusions about the *global* behavior of the system cannot in general be inferred from these local results. In particular, subcriticality only implies that if an attractive limit cycle does emerge on the unstable side, then it must involve terms higher than third order. For example, if we impose the additional assumptions that $\lim_{I \rightarrow 0} F(I) = \infty$ and $\lim_{I \rightarrow \bar{I}} F(I) = -\infty$ for some $\bar{I} \in (0, \infty)$, then the system never becomes explosive, even if the Hopf bifurcation of Proposition 4 is subcritical.

2 Unemployment Risk and Precautionary Saving as a Source of Demand Complementarities

In this section we present a model with unemployment risk and precautionary savings, where equilibrium behavior will exhibit the main features emphasized in our reduced-form model. The model builds on Beaudry, Galizia, and Portier [2015]. Before presenting the model's formal structure, we first discuss its main components informally in order to help motivate the setup and assumptions.

Our aim is to choose a setup that captures a common notion of demand complementarities that is often associated with macroeconomic fluctuations.²² In effect, we want to capture the idea that agents may be hesitant to make large purchases when unemployment is high because they fear becoming unemployed. Thus, we require the model to feature unemployment risk that cannot be fully insured. Simultaneously, we want to allow for the possibility of a feedback effect whereby unemployment can be high in part because agents are holding back on their purchases out of fear of becoming unemployed. This feedback effect will be the source of the demand complementarity in the model. While in general multiple equilibria could arise if this effect is strong enough, we will focus on situations where this is not the case, i.e., where the equilibrium is unique. Finally, we want to explore an environment where the emergence of limit cycles does not rely on the conduct of monetary policy, since we view boom-bust phenomena as being ubiquitous across monetary regimes. To achieve all of these goals, we consider a flexible-price environment where agents can buy goods using credit, but where not all trades are coordinated in centralized markets. In particular, this will allow agents to buy goods even if they are unsure whether they will manage to sell their labor during the period. Meanwhile, the unemployment risk and precautionary behavior in the model come from the assumption that some segments of the labor market are characterized

²²The model of this section shares key features with the long tradition of macro models emphasizing strategic complementarities, aggregate demand externalities and multipliers, such as Diamond [1982] and Cooper and John [1988], though we do not emphasize multiple equilibria. Unemployment risk and its effects on consumption decisions is at the core of the model. The empirical relevance of precautionary saving in the presence of unemployment risk has been documented by many, including Carroll [1992], Carroll and Dunn [1997], Carroll, Sommer, and Slacalek [2012] and Alan, Crossley, and Low [2012]. Recently, there have been several papers that explore the role of unemployment risk and precautionary saving in business cycle models; see, for example, Challe and Ragot [2013], Heathcote and Perri [2012], Ravn and Sterk [2012] and den Haan, Rendahl, and Riegler [2014].

by search frictions. As a result, ex-ante identical agents will have heterogeneous outcomes in terms of debt due to the fact that their labor income will have an idiosyncratic component. Since we want to avoid the complications associated with full-fledged heterogeneous-agent models (i.e., where we would need to track the distribution of debt across periods), we will adopt a sequence-of-markets approach wherein costly debt can be incurred in a first sub-period and then repaid in a second sub-period. This will allow us to maintain many of the attractive analytical features of a representative-agent setup while simultaneously allowing for sufficient heterogeneity in employment outcomes to create a role for precautionary saving, and for a path for this precautionary saving to feed back in to unemployment risk. The desire to coordinate activity will arise in the model as a result of agents having an incentive to purchase higher levels of goods when others' purchases are high, since unemployment would be low and therefore buying on credit less risky. Because these goods will be partially durable, however, any boom driven by purchases will eventually come to an end since, at some point, the rising stock of durables will cause the marginal utility of new purchases to fall enough that agents stop buying. These two forces—one favoring the bunching of durable purchases and the other limiting booms by the diminishing marginal utility of goods—will allow for the emergence of limit cycles. Note that the model is highly stylized and omits many elements known to be relevant for business cycle analysis. For example, the model abstracts from the accumulation of productive capital and focuses only on consumer capital. Despite these omissions, we will use the model as a description of the functioning of the aggregate economy.

2.1 The Environment

Time is discrete. There is a $[0, 1]$ continuum of households indexed by j who live forever with discount factor $\beta \in (0, 1)$. At the beginning of time, households are endowed with the same level of durable goods $X_{j0} = X_0$. Each period is divided into morning and afternoon sub-periods.²³ Agents consume goods and supply labor in both sub-periods, with preferences

²³As will become clear, the alternation of decentralized mornings and centralized afternoons will be useful in obtaining a degenerate distribution of wealth at the end of each period, which helps in solving the model. In a monetary model, Lagos and Wright [2005] use a similar assumption to obtain a degenerate distribution of money holdings. See also Rocheteau and Wright [2005], and more recently Kaplan and Menzio [2014].

given by $U(c_j) - \nu(\ell_j) + \tilde{U}(\tilde{c}_j) - \tilde{\nu}(\tilde{\ell}_j)$, where c_j and $\ell_j \in [0, \bar{\ell}]$ (respectively, \tilde{c}_j and $\tilde{\ell}_j \in [0, \tilde{\bar{\ell}}]$) are consumption and labor in the morning (respectively, in the afternoon). Intertemporal preferences are therefore given by

$$\sum_{t=0}^{\infty} \beta^t \left[U(c_{jt}) - \nu(\ell_{jt}) + \tilde{U}(\tilde{c}_{jt}) - \tilde{\nu}(\tilde{\ell}_{jt}) \right]. \quad (10)$$

The afternoon good is chosen to be the numéraire.

Mornings: In the morning, households can purchase consumption goods and search for employment. We will refer to this good as the morning good, and it will be at least partially durable. There is no money in this economy, but there is credit. When the household buys morning 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. In order to limit the state space of the system, we assume that households must balance their asset positions in the afternoon by repaying any outstanding debts or receiving a payment for any surplus.²⁴ These payments are made in terms of the afternoon good, which is the numéraire in this economy.

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 accumulate a stock of the morning good, which they can either consume or trade. This stock evolves according to

$$X_{jt+1} = (1 - \delta)(X_{jt} + \gamma e_{jt}), \quad (11)$$

where X_{jt} is the stock brought into morning t and e_{jt} is quantity of morning-good purchases in morning t . For simplicity, we assume that a constant fraction $\gamma \in (0, 1]$ of these purchases are durable.²⁵ $\delta \in (0, 1]$ is the depreciation rate.

Trade in the morning good is subject to a coordination problem because of frictions in the labor market. At the beginning of the morning sub-period, the household splits up

²⁴This assumption, together with the symmetric initial endowment of durable goods and the fact that morning purchases are not employment-state-contingent (see below), will imply that all households exit the period- t afternoon with the same levels of both durable goods and financial assets (although they are in general heterogeneous at the end of the morning period). This degeneracy in the wealth distribution makes the analysis much more tractable.

²⁵In the quantitative exercise below, we will interpret “durables” as including conventional durable goods as well as residential investment.

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 the good p_t (in units of the afternoon good) as given. The morning market for labor is subject to a matching friction, with sellers of labor searching for employers and employers searching for labor. The important 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 firms in the economy who can decide to search for workers in view of supplying the morning good to the market. Each firm can hire one worker and has access to a decreasing-returns-to-scale production function $F(\ell_t)$, where ℓ_t is the number of hours supplied by the worker in a match.²⁶ In order to match with a worker, a firm is required to post a vacancy at fixed cost Φ in terms of the morning good, so that the net production of a firm hiring ℓ_t hours of labor is $F(\ell_t) - \Phi$. Firms search for workers and, upon finding a worker, they jointly negotiate the terms of trade—which consists of a wage and a level of worked hours, (w_t, ℓ_t) —according to some protocol that we assume to be a model primitive, following Gu, Mattesini, Monnet, and Wright [2013]. In our baseline estimation of this model, we will assume Nash bargaining as the bargaining protocol. However, in the next subsection, we will assume a simpler bargaining protocol as to clarify intuition.

The labor market operates as follows. All workers are assumed to search for employment. Letting N_t represent the number of firms searching for workers, the number of matches M_t is then given by the constant-returns-to-scale matching function $M(\cdot)$, with $M_t = M(N_t, 1)$. Note that M_t is also the employment rate, so that $u_t = 1 - M_t$ is the unemployment rate. The resource constraint for the morning good is then given by

$$\int_0^1 (c_{jt} - X_{jt})dj = M_t F(\ell_t) - N_t \Phi,$$

where the left-hand side is total net purchases of consumption goods and the right-hand side

²⁶We also assume that F is such that both $F'(\ell_t)\ell_t$ and $[F(\ell_t) - F'(\ell_t)\ell_t]$ are strictly increasing functions of ℓ_t , and that $F(0) = 0$. These properties are exhibited, for example, by the Cobb-Douglas function $F(\ell_t) = A\ell_t^\alpha$.

is the total available supply after subtracting firms' vacancy-posting costs. Firms enter the market up to the point where expected profits are zero. This condition can be written as²⁷

$$\frac{M_t}{N_t} \left[F(\ell_t) - \frac{w_t}{p_t} \ell_t \right] = \Phi. \quad (12)$$

At the end of the morning, a household's individual state is summarized by its net financial asset position, a_{jt} (in units of the afternoon good), which is given by $a_{jt} = w_t \ell_{jt} - p_t (c_{jt} - X_{jt})$ if the worker was employed, and $a_{jt} = -p_t (c_{jt} - X_{jt})$ if the worker was unemployed.

Afternoons: The afternoon market is set up so as to allow households to pay back any debt incurred in the morning market. In fact, the afternoon economy will be modeled in such a way that if there were no frictions in the morning then there would be no trade between agents in the afternoon. In particular, we assume that $\tilde{U}(\cdot)$ is increasing and strictly concave in \tilde{c}_{jt} , but assume that the dis-utility of work is linear in the afternoon, i.e.,

$$\tilde{v}(\tilde{\ell}_{jt}) = v \tilde{\ell}_{jt}.$$

To ensure that taking on debt is potentially risky (which will induce precautionary savings), we assume that, in the afternoon, households can produce a good for their own consumption using one unit of labor to produce one unit of goods. However, in order to produce goods in the afternoon that can be transferred to others in order to satisfy debts incurred in the morning, a household must supply $1 + \tau$ units of labor, $\tau > 0$. The parameter τ will play an important role in the analysis since it controls the perceived risk associated with accumulating debt in the morning sub-period.

Since a household's financial asset position a_{jt} is pre-determined when entering the afternoon sub-period, decisions for \tilde{c}_{jt} and $\tilde{\ell}_{jt}$ will satisfy

$$v = \tilde{U}'(\tilde{\ell}_{jt} + a_{jt}),$$

$$\tilde{c}_{jt} = \begin{cases} \tilde{\ell}_{jt} + a_{jt} & \text{if } a_{jt} \geq 0, \\ \tilde{\ell}_{jt} + (1 + \tau) a_{jt} & \text{if } a_{jt} < 0. \end{cases}$$

²⁷We assume that searching firms pool their ex-post profits and losses so that they each make exactly zero profits in equilibrium, regardless of whether or not they are matched with a worker.

Letting $V(a_{jt})$ denote the afternoon-period utility ($\tilde{U}(\tilde{c}_j) - \tilde{v}(\tilde{\ell}_j)$) associated with entering the afternoon with financial assets a_{jt} , it is easy to check that V is given (up to a constant) by

$$V(a_{jt}) = \begin{cases} va_{jt} & \text{if } a_{jt} \geq 0, \\ (1 + \tau)va_{jt} & \text{if } a_{jt} < 0. \end{cases}$$

This function $V(\cdot)$ is piecewise linear and concave, with a kink at $a = 0$.²⁸ Here, the marginal value of financial 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. In particular, a fall in the employment rate causes buyers to reduce their purchases, since they fear ending up in the unemployment state with debt that is costly to serve. This mechanism is central to the strategic complementarity that emerges in the model, which in turn is necessary to generate local instability and limit-cycle dynamics. The strength of this mechanism is governed to a large extent by τ , which parameterizes agents' aversion to debt.

In order to repay creditors, households exiting the morning in debt are forced to work more than non-indebted households in the afternoon, with all financial asset positions being fully resolved (i.e., reset to zero) by the end of the afternoon sub-period. Meanwhile, since the durables-accumulation decision was made in the morning prior to realization of the idiosyncratic labor market status of the households, and since all households are assumed to have a symmetric initial endowment of durables $X_{j0} = X_0$, in equilibrium they will all exit period t with the same level of durables $X_{jt+1} = X_{t+1}$. Thus, despite heterogeneity in labor market outcomes, the distribution of both financial and capital assets at the end of every period will be degenerate.

Intertemporal Equilibrium: Substituting $V(a)$ from above into the objective function (10), the buyer's problem is to choose paths for X_{jt+1} and e_{jt} to maximize

$$\sum_{t=0}^{\infty} \beta^t \left\{ U(X_{jt} + e_{jt}) + (1 - u_t) [-\nu(\ell_{jt}) + v(w_t \ell_{jt} - p_t e_{jt})] + u_t (1 + \tau)v(-p_t e_{jt}) \right\} \quad (13)$$

²⁸As noted in Beaudry, Galizia, and Portier [2015], what matters here is that the marginal value of financial assets be smaller in surplus than in deficit.

subject to the accumulation equation (11). Meanwhile, if the worker finds a match, his labor supply decision is given by the ℓ_{jt} that maximizes $-\nu(\ell_{jt}) + v(w_t\ell_{jt} - p_t e_{jt})$.

A symmetric intertemporal equilibrium for this economy is then given by sequences of relative prices of goods p_t , wage rates w_t , morning good purchases e_t , labor supplied in matches ℓ_t , vacancy postings N_t , the stock of durables X_{t+1} , and unemployment rates u_t such that:

1. $\{e_t\}$ and $\{X_{t+1}\}$ maximize the buyer's objective function (13) subject to the accumulation equation (11), taking prices, wages, the unemployment rate and ℓ_t as given;
2. The within-a-match terms of trade (w_t, ℓ_t) are jointly set by the firm and the worker according to some exogenous protocol, taking e_t and prices as given;
3. The goods market clears, i.e., $e_t = (1 - u_t)F(\ell) - N_t\Phi$;
4. Firms' entry decisions are such that the zero-profit condition (12) holds; and
5. Jobs are created according to the matching function, i.e., $1 - u_t = M(N_t, 1)$.

2.2 Understanding the Period- t Equilibrium in the Myopic Case

In order to most clearly highlight the nature of the demand complementarity in the model, it is useful to temporarily focus on the myopic case where $\beta = 0$. In this case, we simply have a repeated sequence of static decisions by households, with the only linkage between periods being the inherited stock of durable goods X_t . To further ease the presentation of the forces at play, we make two assumptions. First, the matching function takes the form $M_t = \min\{N_t, 1\}$; that is, the number of matches in the economy is given by the short side of the market. Second, the determination of the wage and hours worked within a firm-worker pair is done efficiently through a type of Walrasian bargaining process similar in spirit to that in Lucas and Prescott [1974]; that is, given a match, the wage and hours worked are determined so as to equate the demand for labor by the firm to the supply of labor by the worker.²⁹ As a result, in equilibrium workers will be paid their marginal products

²⁹We make this assumption here for analytical convenience. In the quantitative model of section 3, we replace this assumption by the more standard one of Nash bargaining. In Beaudry, Galizia, and Portier [2015] it is shown that the demand complementarity we exploit here is robust to alternative bargaining protocols.

($pF'(\ell_t) = w_t$), and further w_t will be equal to a worker's marginal dis-utility of work. It is easy to verify that this bargaining outcome lies within the bargaining set between the firm and the worker.

Letting e_t denote the average level of morning purchases in the economy, one may show that household j 's optimal consumption-choice decision can be expressed as³⁰

$$U'(X_t + e_{jt}) = p(e_t)v[1 + \tau u(e_t)] \equiv r(e_t), \quad (14)$$

where $p(\cdot)$ and $u(\cdot)$ are the price of consumption goods and the unemployment rate, respectively, expressed as functions of aggregate purchases. The left-hand side of (14) is household j 's marginal utility of consumption. The right-hand side captures buyer j 's expected marginal-utility cost of funds, that we denote $r(e_t)$. When the economy is at full employment ($u(e_t) = 0$), this is simply equal to the price $p(e_t)$ of morning goods in terms of afternoon ones, times the marginal value v of those afternoon goods 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.

A period- t equilibrium for this economy is given by a solution to (14) with the additional restriction that $e_{jt} = e_t$. To understand how the equilibrium is affected by the initial level of durables X_t , note the following properties of the equilibrium functions $p(\cdot)$ and $u(\cdot)$. First, one may show that $1 - u(e_t) = \min\{e_t/e^*, 1\}$, where e^* is output (net of vacancy-posting costs) produced per firm when there is a positive level of unemployment.³¹ Second, one may show that $p(\cdot)$ is a continuous function of e_t , with $p'(e_t) = 0$ for $e_t < e^*$, and $p'(e_t) > 0$ for $e_t > e^*$.³² The consequences of these two properties for the marginal-utility cost of funds (i.e., the right-hand side of (14)) are illustrated by the curve labeled “cost of funds” in

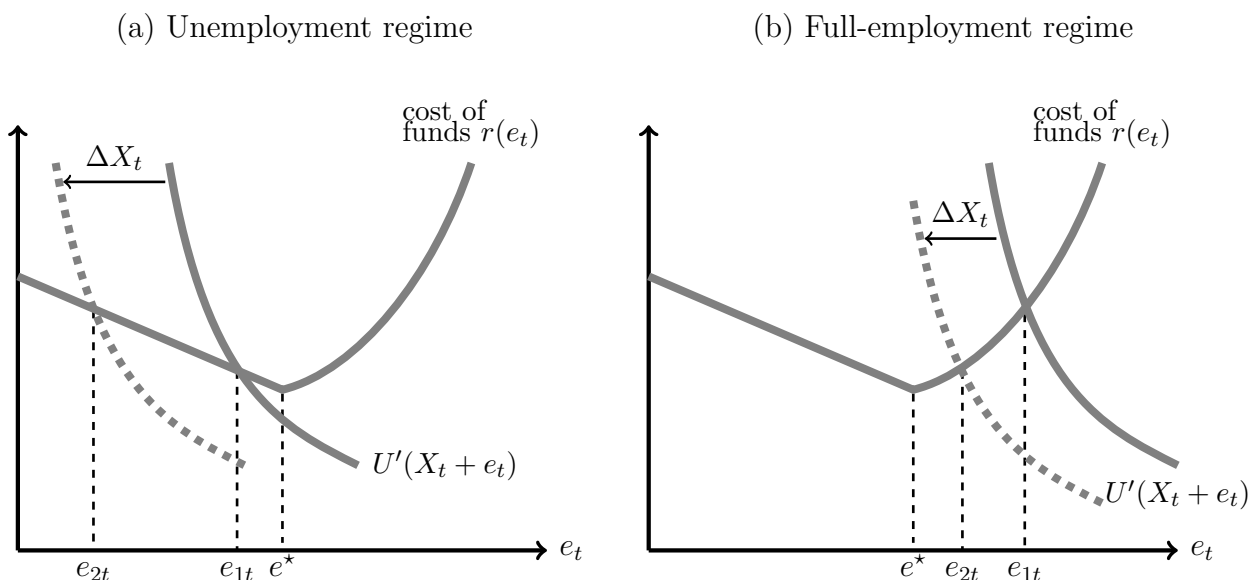
³⁰See Beaudry, Galizia, and Portier [2015].

³¹When there is unemployment, the “min” matching function and the firm's zero-profit condition together imply $F(\ell_t) - F'(\ell_t)\ell = \Phi$. Since Φ is a constant, conditional on there being unemployment this implies that $\ell_t = \ell^*$, where ℓ^* solves this zero-profit equation. Output net of vacancy costs is then $e^* \equiv F(\ell^*) - \Phi$.

³²Combining the household's labor supply condition and the firm's labor demand condition, one may obtain $p_t = v'(\ell_t)/[vF'(\ell_t)]$. As pointed out in footnote 31, when $e_t < e^*$ we have $\ell_t = \ell^*$, so that $p_t = p^* \equiv v'(\ell^*)/[vF'(\ell^*)]$. 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_t), which causes $p(\cdot)$ to be increasing in e_t on $e_t > e^*$.

panel (a) of Figure 3. For e_t 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.

Figure 3: Period t Equilibrium Determination in the Myopic Model



The two regimes—unemployment and full-employment—are associated with different equilibrium properties depending on different levels of the inherited X_t .³³ Panel (a) of Figure 3 shows the case for the unemployment regime. With low X_t , the economy is in equilibrium at the level e_{1t} of purchases, which occurs at the intersection of the cost-of-funds curve and the solid marginal-utility function $U'(X_t + e_t)$. In such a situation, if the inherited stock of durables is increased by an amount ΔX_t , the marginal-utility function shifts to the left by ΔX units, as represented by the dashed curve in the figure. We see that the equilibrium level of purchases is smaller as a result of the higher level of X_t , and furthermore that it has

³³As shown in Beaudry, Galizia, and Portier [2015], if debt aversion 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, gives no indication that multiple equilibrium 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.

fallen by *more* than ΔX_t (so that total consumption $c_t = X_t + e_t$ is lower). This amplified response is due to the strategic complementarity that exists in the unemployment regime: a rise in the inherited stock of durables causes households to reduce their demand for new goods which, via an extensive labor-margin adjustment, increases the unemployment rate u_t , 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 3 shows the same experiment but beginning from the full-employment regime. In this case, we again see that a higher X_t is associated with a fall in equilibrium purchases, but now the fall is by *less* than ΔX_t (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 inherited stock of durables 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 the initial level of X_t in the unemployment regime comes from the presence of strategic complementarity, and the corresponding *insensitivity* in the full-employment regime comes from strategic substitutability. This will play a crucial part in generating limit cycles. Note also that the sensitivity of e_t to X_t 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 debt aversion parameter τ , we see that τ controls the degree of strategic complementarity in the unemployment regime.

In order to see the link between the current model and our previous reduced-form model, Figure 4 illustrates the relationship between the household's decision e_{jt} and the decision by other households e_t . In effect, inverting the U' function in (14), one can express the expenditure e_{jt} of household j as a function of aggregate expenditure e_t and the stock of durable goods X_t as follows:

$$e_{jt} = U'^{-1}(r(e_t)) - X_t \quad (15)$$

As can be seen, Figure 4 has similarities with Figure 1 of the reduced-form model, and in particular features a regime of strategic complementarity and a regime of strategic substitutability, with the vertical location of household j 's optimal-expenditure (i.e., best-response) function depending on the history of the economy. Two best-response functions are plotted

in the figure. In the lower one—which captures a situation where X_t is high—purchases of goods play the role of strategic complements. In the higher one—which captures a situation where X_t is low—purchases play the role of strategic substitutes. Note also that as long as debt aversion τ is not too large, then the best-response function in Figure 4 necessarily crosses the 45-degree line only once, implying that the equilibrium is unique.

It is easy to verify that this myopic economy can exhibit limit cycles, with periods of high purchases followed by periods of low purchases. In particular, using the logic presented in Section 1, it can be verified that, as long as δ is sufficiently low (so that the steady state is in the unemployment regime), increasing the degree of complementarity by increasing debt aversion τ will eventually cause this system to become locally unstable. However, the system nevertheless remains globally stable, since once X_t becomes too small the economy reaches full employment, while once X_t becomes too large it reaches zero employment.³⁴ In either case, the strong local demand complementarities in effect at the steady state are replaced with complementarities, thereby preventing the system from exploding.


2.3 Understanding the Intertemporal Equilibrium When Agents Are Forward-Looking

In the previous subsection we saw that, when $\beta = 0$, the actions of individual agents will be strategic complements when X is high and strategic substitutes when X is low. While the myopic case is useful for building intuition and relating the current structural model with the reduced-form model of Section 1, of more general interest is whether limit cycles may occur in the current setting 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, limit cycles cannot occur for β sufficiently close to one (holding all else constant).³⁵

³⁴The fact that the non-negativity constraint on employment binds in some periods when the steady state is locally unstable is due to several stark assumptions made in this section for expositional purposes only. In the quantitative exercise below we will relax these assumptions, allowing for the presence of endogenous forces that prevent employment from falling to zero.

³⁵See also Brock and Scheinkman [1976] and McKenzie [1976]. It is important to understand the meaning and scope of the turnpike theorem. That theorem highlights a trade-off between the level of discounting and the “curvature” of preferences plus technology. With the preferences and technology commonly used in quantitative macroeconomics, curvature is in general low enough so that a low β is needed for limit cycles to occur. Nevertheless, as shown by Benhabib and Rustichini [1990], for any positive discount rate there exists

Figure 4: Equilibrium Determination for Two Different Histories



Notes: This figure plots the optimal spending of household j as a function of total spending, $e_{jt} = U'^{-1}(r(e_t)) - X_t$ (equation (15)). The intercepts correspond to two different histories of the model, $X_t = (1 - \delta)(X_{t-1} + \gamma e_{t-1}) = \sum_1^\infty (1 - \delta)^j \gamma e_{t-j}$.

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} .³⁶ 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.*³⁷

The key property that ensures global stability in this theorem is the assumption that \mathcal{W} is concave. Since, all else equal, fluctuations are payoff-decreasing when \mathcal{W} is concave, if 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 results in 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 many popular 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 our model satisfies a key condition needed for the existence of a stable limit cycle, 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.

a large family of standard Cobb-Douglas technologies with three sectors which have optimal growth paths of persistent cycles. Such an example economy is presented in Benhabib and Nishimura [1998]. More generally, Boldrin and Montrucchio [1986] have shown that any differentiable difference equation corresponds to the policy function of some well-behaved concave dynamic program with a sufficiently large discount rate.

³⁶Note 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.

³⁷For a proof and more formal statement of the theorem, see Scheinkman [1976] Theorem 3.

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 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.³⁸ 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: when unemployment risk is low—that is, 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 model—tends to create periods of high durables accumulation alternating with periods of low durables accumulation.

The final proposition of this section clarifies the importance of the debt aversion 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.

³⁸An example of such a \mathcal{W} is found in the proof of Proposition 7.

3 Estimating a DSGE Model That May Involve Limit Cycles

This section presents a quantitative exploration into limit cycles. In particular, we will estimate an augmented version of the dynamic model discussed above where we also allow for exogenous disturbances, so as to force limit cycles to compete with exogenous driving forces in explaining business-cycle properties of the data. Our main goals will be to examine whether the estimation procedure chooses parameter values that support limit cycles, and to examine the extent to which the reliance on exogenous driving forces is reduced if a limit cycle is present. We first present the data properties we aim to match before describing the extended DSGE model and the estimation results.

3.1 Data

One of the main criticisms of limit cycles as an explanation of macroeconomic fluctuations relates to the fact that business cycles are quite irregular in length, duration, and depth, while limit cycles tend to be quite regular. In order to quantify the degree of irregularity in a macroeconomic data series, it is helpful to look at its spectrum,³⁹ which offers a simple visual way of examining whether the cyclical properties of the series are regular or irregular. If cycles are very regular, this will show up in the spectrum as one or more large peaks at particular periodicities. In contrast, if the cycle is very irregular, the spectrum will have weight across a large range of frequencies.

In order to review the extent of irregularity in business cycles, we start by looking at the behavior of hours worked (per capita) in the economy. This is a useful macroeconomic variable to examine for our purposes since it is not too far from being stationary. Nonetheless, because of demographic changes and changes in female labor market attachment, hours worked per capita do exhibit some low-frequency movements that are unrelated to business cycles. To eliminate these low-frequency movements, we use a band-pass filter to remove fluctuations associated with periods longer than 20 years (80 quarters).⁴⁰ Panel (a) of Figure

³⁹As is well known, any stationary data series can be written as the sum of orthogonal sine waves of different frequencies. The spectrum is the function giving the variance of the corresponding sine wave for a particular frequency, and can thus be viewed as a decomposition of the variance of the series by frequency.

⁴⁰Note however that we do not remove any of the high-frequency fluctuations in the data.

5 displays the resulting times series for US hours over the period 1960-2012. Several things emerge from the plot. First, visual inspection reveals that the series corresponds closely to the standard timing of business cycles, with NBER-dated recessions (shaded areas) being associated with important falls in hours worked. Second, it is interesting to see that, with the exception of the 1970s, a typical cycle in hours is quite long. For example, since the beginning of the 1980s, one observes three cycles, each of length close to 40 quarters, while the principal cycle in the 1960s is of a similar length. This observation is echoed in panel (b) of Figure 5, where we plot the spectrum of our hours series.⁴¹ As could be expected from inspecting plot (a), the spectrum does not put all of its mass at a few periodicities—which would have been a sign of near-perfect regularity—but it does place substantial weight at periodicities between 30 and 50 quarters, with a peak at around 40 quarters. This is interesting, since business cycles are commonly defined as fluctuations between 8 and 32 quarters (shown as the light gray zone superimposed on the figure). However, to the extent that we believe it reflects business cycle forces, the behavior of this filtered hours series suggests that periodicities relevant for business cycle analysis should to be extended to include slightly lower-frequency movements corresponding to fluctuations from 32 to 60 quarters (shown as the dark gray zone superimposed on the figure).⁴²

The hours series presented in panel (a) of Figure 5 will be the target series we seek to explain with our model. In particular, we will estimate the model parameters so as to replicate the spectrum for hours as shown in panel (b). Using the estimated model, we can then evaluate how well it matches the data in other dimensions. Of particular interest to us will be to see how well it simultaneously explains the spectrum of output and the coherence between output and hours. In panel (c) of Figure 5 we report the spectrum for output after applying the same band-pass filter used to de-trend hours. As can be seen, the spectrum for output resembles that for hours, with substantial weight being placed on frequencies that are lower than traditionally associated with business cycles. Finally, panel (d) of the figure reports the coherence between output and hours, which provides information about

⁴¹The figure displays the spectral density of hours for periodicities between 4 and 80 quarters in length, after first using the band-pass filter that removes cycles with periods longer than 80 quarters.

⁴²See Comin and Gertler [2006] and Pancrazi [2015] for similar observations. Note also that these basic properties of the hours spectrum are robust to increases in the upper limit of the band-pass filter from our baseline level of 80 quarters up to as much as 150 quarters.

how co-movements between output and hours behave at different frequencies.⁴³ The result displayed in panel (d) of Figure 5 shows that hours and output are highly correlated at all periods between 8 and 80 quarters.

3.2 The Extended DSGE Model

The dynamic model with unemployment risk presented in Section 2 was constructed with an eye toward analytical tractability. As a result, that model lacks features that are known to be helpful in quantitatively matching the data, and includes several others that, while not central to the key mechanisms, turn out to be overly restrictive in a quantitative setting. Since the main purpose of the exercise in this section is quantitative in nature, we make adjustments to the model designed to increase flexibility in that regard.

First, as mentioned earlier, dynamic systems with a single state variable often have difficulty producing deterministic fluctuations with the basic qualitative properties that we observe in macroeconomic aggregates. In particular, such fluctuations 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 a reasonable chance of replicating key features of the data, it will require the addition of at least one other state variable. Further, as experience with the reduced-form model of Section 1 suggests, this state variable should be capable of inducing some degree of sluggishness.⁴⁴ To this end, we will assume that the household exhibits internal habit-formation in consumption of the morning good, c_t ,⁴⁵ so that its period utility for morning consumption is now given by

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

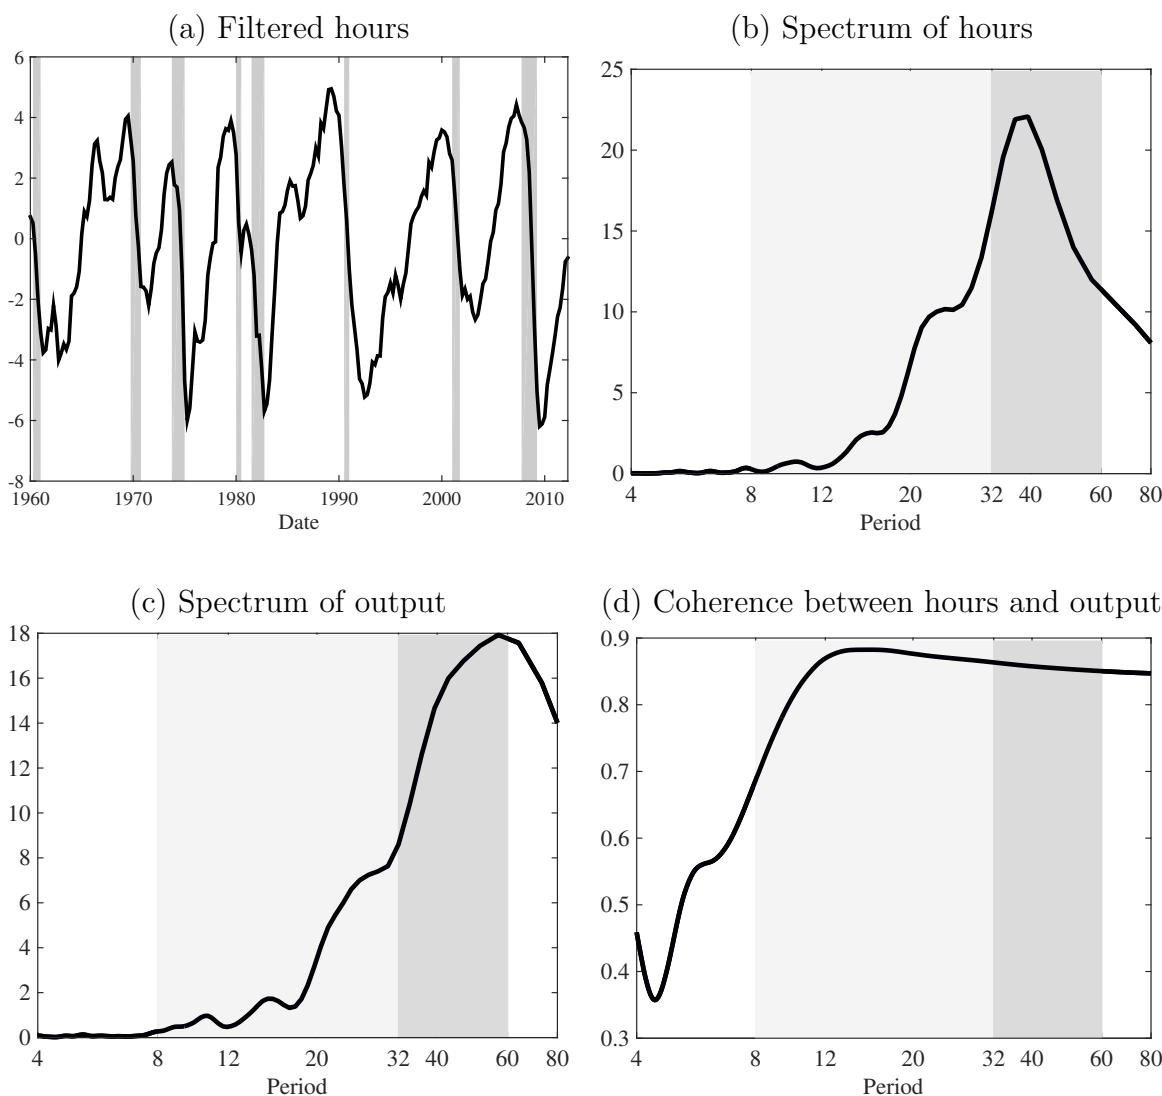
Here, $h \in [0, 1)$ is a parameter controlling the degree of habit persistence.

⁴³Coherence is analogous to a regression R^2 , giving the proportion of the variance of one series that can be linearly predicted by another at each periodicity. A coherence of one would thus indicate that two series are perfectly correlated at that periodicity, while a coherence of zero would indicate that they are orthogonal.

⁴⁴Recall that in our reduced-form model, the parameter that controls sluggishness (α_2), needed to be sufficiently high for the system to undergo a Hopf bifurcation instead of a flip bifurcation.

⁴⁵Consumption habit introduces sluggishness into the model while maintaining tractability and keeping the model as close as possible to the baseline version presented in Section 2. There are nonetheless a number of alternative means of introducing sluggishness (e.g., adjustment costs in investment or employment) that could be used instead and that would be capable delivering similar qualitative dynamics.

Figure 5: Data, spectral density and coherence



Notes: Hours Worked series is the log of BLS nonfarm hours worked divided by population, detrended with a band-pass filter to remove fluctuations with periods greater than 80 quarters. In panel (a), shaded areas are NBER-dated recessions. The sample runs from 1960Q1 to 2012Q4. Output is GDP and is also detrended with the same band-pass filter. 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 is kernel-smoothed raw spectrum. Kernel is a Hamming window with bandwidth parameter 11. 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.

Second, we relax two simplifying assumptions made earlier; namely, the assumption of a “min” matching function (introduced in subsection 2.2), and the assumption of homogeneous firms. With those two assumptions in place, we obtained a simple dichotomy featuring an unemployment regime in which all output adjustments occur along the extensive labor margin (number of workers), and a full-employment regime in which all adjustments occur along the intensive margin (hours per worker). In order to relax this stark dichotomy, we now return to a more general form for the matching function, and also allow firms to be somewhat heterogeneous. We model the latter as heterogeneity in firms’ vacancy-posting costs, so that instead of assuming that all firms have cost Φ , we assume that the N -th firm has vacancy cost $\Phi(N) \geq 0$, where $\Phi(\cdot)$ is a non-decreasing function. These two assumptions together will in general create the possibility of both extensive and intensive labor margin adjustments occurring simultaneously.⁴⁶

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.⁴⁷ One of the key contributions of this paper is to allow limit-cycle forces to compete with—and possibly complement—exogenous disturbances in explaining the data. To this end, we also include in the model an exogenous random TFP process, $\tilde{\theta}_t$, which multiplies firms’ production functions so that output is given by $\tilde{\theta}_t F(\ell_t)$.⁴⁸

Fourth, we make the more standard assumption that within-a-match terms of trade (w_t, ℓ_t) are set according to Nash bargaining.⁴⁹ In this case, hours worked are set efficiently

⁴⁶Note that the functional forms chosen for the matching and vacancy-cost functions (discussed below) will nest the baseline cases of “min” matching and constant vacancy cost. Since the parameters of these functions will be estimated, the data will thus ultimately choose the degree to which the simple dichotomy is relaxed.

⁴⁷One 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.

⁴⁸For convenience, in order to retain certain analytical properties that are helpful in a computational setting, we assume that firms’ vacancy costs and the production of the afternoon good also fluctuate with this TFP process. Output, vacancy costs, and afternoon-period utility are thus given by $\tilde{\theta}_t F(\cdot)$, $\tilde{\theta}_t \Phi(\cdot)$, and $\tilde{\theta}_t^{-1} V(\cdot)$, respectively.

⁴⁹Neither our qualitative nor our quantitative results are sensitive to the choice of Nash versus Walrasian bargaining. See Beaudry, Galizia, and Portier [2015] for further discussion of the qualitative results. See Galizia [2014] for quantitative results under the assumption of Walrasian bargaining.

with the marginal product of labor being equalized to the marginal dis-utility of labor. Wages are then set to split the match surplus according to a Nash bargaining protocol, where parameter ξ is the bargaining strength of the worker.

3.3 Functional Forms, Calibration and Estimation

Our choice of functional forms is as follows. Production is assumed to be of the Cobb-Douglas form

$$\tilde{\theta}_t F(\ell) = \tilde{\theta}_t \ell^\alpha.$$

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

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

We choose a quadratic formulation for preferences since this greatly simplifies certain computational issues. The dis-utility of labor in the morning market is taken to be of the form

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

The matching function is assumed to be of the CES form⁵⁰

$$M(N) = (1 + N^{-\chi})^{-\frac{1}{\chi}}$$

where $\chi > 0$ is a parameter. Note that this function nests the Cobb-Douglas matching function as $\chi \rightarrow 0$ and the “min” matching function as $\chi \rightarrow \infty$. The vacancy cost of the N -th firm is assumed to be given by

$$\Phi(N) = \bar{\Phi} G\left(\frac{\log(N) - \log(\bar{N})}{\sigma_N}\right)$$

where G is the standard normal CDF and \bar{N} , σ_N and $\bar{\Phi}$ are parameters. Note that this function nests the constant function $\Phi(N) = \bar{\Phi}$ as $\bar{N} \rightarrow 0$. Finally, we assume that the TFP process is given simply by

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

⁵⁰This matching function was proposed by den Haan, Ramey, and Watson [2000].

Several of the model parameters are calibrated to common values in the literature. In particular, we set $\alpha = 2/3$ so to have a reasonable amount of decreasing returns to hours worked. The inverse Frisch elasticity is set at the widely used level $\omega = 1$. As the period length is a quarter, we set the depreciation rate and discount factor at standard values of $\delta = 0.025$ and $\beta = 0.99$, respectively, and normalize the maximum vacancy cost and marginal dis-utility of afternoon labor to $\bar{\Phi} = 1$ and $v = 1$, respectively. We also calibrate the Nash bargaining share of the worker to $\xi = 0.5$.⁵¹ 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.⁵² The remaining parameters were estimated.

Solving the model for a particular parameterization was done using the parameterized expectations (PE) approach.⁵³ 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 band-pass 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 can be noted from the estimation results. The first interesting result is that the estimated TFP process is close to the process that can be obtained directly from measured productivity data. For example, using John Fernald’s [2014] measure of business-sector labor productivity growth over the sample period (1960Q1-2012Q4),⁵⁴ 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%.⁵⁵

⁵¹When also included in the set of estimated parameters, the value of ξ obtained was close to the value used to initialize the estimation algorithm, suggesting that this parameter is not identified.

⁵²As noted above, we include the conceptually-similar residential investment under the heading of “durables”. The figure of 0.192 is obtained from NIPA data as the average of (Durable goods + Residential investment)/(Consumption + Residential investment) over the sample period 1960Q1-2012Q4.

⁵³See, for example, den Haan and Marcet [1990] and Marcet and Marshall [1994]. Details can be found in Appendix C.

⁵⁴Available at http://www.frbsf.org/economic-research/economists/jferald/quarterly_tfp.xls.

⁵⁵Similar values are obtained when using Fernald’s TFP or utilization-adjusted TFP measures instead of

The corresponding parameters estimated for the unemployment-risk model, meanwhile, are $\rho = 0.967$ and $\sigma = 0.614$, respectively, which yields an unconditional standard deviation of 2.42%. 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 is a feature we will return to.

The other parameter with a clear comparator in the data or literature is habit persistence, which is estimated here to be $h = 0.75$, well within the range of standard estimates obtained elsewhere. 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).

Table 1: Parameter Values

Parameter	Value	Description
a	11.857	Marginal utility of consumption, intercept
b	2.618	Marginal utility of consumption, slope
h	0.753	Habit persistence
ν_1	12.065	Labor dis-utility scaling factor
τ	0.146	Debt aversion
A	2.935	Constant productivity factor
χ	92.299	Complementarity in matching function
\bar{N}	0.892	Mean parameter, fixed cost function
σ_N	0.012	Dispersion parameter, fixed cost function
ρ	0.967	Persistence of TFP
σ	0.614	$100 \times$ s.d. of innovation 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
$\bar{\Phi}$	1	Maximum firm vacancy cost
γ	0.192	Fraction of purchases entering durables stock
ξ	0.5	Worker share in Nash bargaining

The remaining parameters in Table 1 are composed mainly of scale parameters, and parameters for which few if any precedents exist. The debt aversion parameter τ , which labor productivity.

drives 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 $\tau = 0.146$, or 14.6%, exceeds typical borrowing-lending spreads. However, we believe this view is overly restrictive, and in particular we interpret the parameter τ as capturing in a reduced-form way all of the factors that may lead agents to perceive taking on debt as risky.⁵⁶

3.4 Main Results

To illustrate the deterministic mechanisms, we first report results obtained when shutting down the estimated TFP shock (i.e., setting $\sigma = 0$).⁵⁷ Panel (a) of Figure 6 plots a simulated 212-quarter sample⁵⁸ of log-hours generated from the deterministic version of the model, where hours are given by $L_t = (1 - u_t)\ell_t$. Two key properties should be noted. First, estimation of the model has yielded a set of parameters that generate deterministic cyclical behavior, and these cycles are of a reasonable length (approximately 32 quarters). The difficulty that some earlier models had in generating cycles of quantitatively reasonable lengths may have been one of the factors leading to limited interest in using such a framework for understanding business cycles. However, as this exercise demonstrates, reasonable-length endogenous cycles can be generated in our framework relatively easily precisely because it possesses the three key features we highlighted in Section 1: diminishing returns to capital accumulation, sluggishness, and strong local demand complementarities. Second, notwithstanding the reasonable cycle length, it is clear when comparing the simulated data in Figure

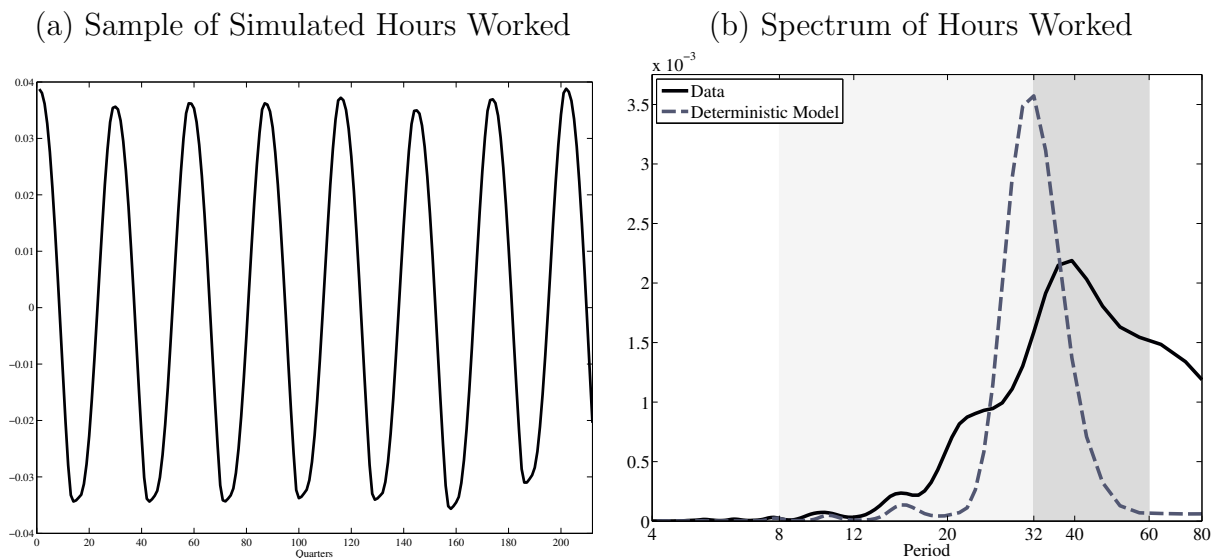
⁵⁶As has been well documented, households appear to exhibit a considerable degree of risk-aversion, well beyond what can be produced using time-separable preferences with a reasonable degree of intertemporal substitutability. This has led to the adoption of more flexible preference specifications (e.g., the specification proposed by Epstein and Zin [1989]) in which risk attitude is parameterized separately from intertemporal substitutability. The parameter τ in our model can be similarly viewed as a way of directly parameterizing household attitudes toward the risks associated with taking on debt in the presence of uncertainty about employment outcomes.

⁵⁷In 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.

⁵⁸This is equal to the length of the sample period of the data.

6 to the actual data in Figure 5 that the fluctuations in the deterministic unemployment-risk model are far too regular when compared to the data.⁵⁹

Figure 6: Deterministic Model



Notes: Panel (a) shows 212-quarter simulated sample (same size as data set) of band-pass filtered $\log(\text{hours worked})$ (where hours are computed as $L_t = (1 - u_t)\ell_t$) generated from the deterministic model. Initial simulated series was 252 quarters long, with first and last 20 quarters discarded after band-pass filtering. Details for computation of model spectrum in panel (b) can be found in Appendix C.

These properties of the deterministic version of our estimated model—i.e., a highly regular 32-quarter cycle—can also be seen clearly in the frequency domain. Panel (b) of Figure 6 plots the spectrum for the deterministic version (dashed line), along with the spectrum for the data (solid line) for comparison.⁶⁰ Consistent with the pattern in the time domain, the spectrum exhibits a peak at around 32 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.

⁵⁹While successive cycles are close to symmetric in panel (a) of Figure 6, they are clearly not exactly identical. As discussed in footnote 17, this is an artifact of the discrete-time formulation of the model. In a continuous-time formulation, each cycle would be exactly identical to the last.

⁶⁰Note 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 data.

Re-introducing the estimated TFP shock into the model, we see a markedly different picture in both the time and frequency domains. Panel (a) of Figure 7 plots a 212-quarter arbitrary 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 5 for actual data. This is confirmed by the spectrum, which is plotted in panel (b) of Figure 7 alongside the data spectrum. Also plotted is a point-wise 90% simulated confidence interval from the model for data sets of the same length as the data (i.e., 212 quarters).⁶¹ 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 auto-covariance function (ACF) of hours, i.e., $Cov(L_t, L_{t-k})$, where k is the lag (in quarters). Panel (a) of Figure 8 plots the result for the first 40 lags for both the data and model, along with point-wise 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.⁶²

To verify that the good fit of the spectrum is not driven by the choice of filter, Figure 9 plots the data and model spectra for hours under four alternative filtering choices.⁶³ 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.

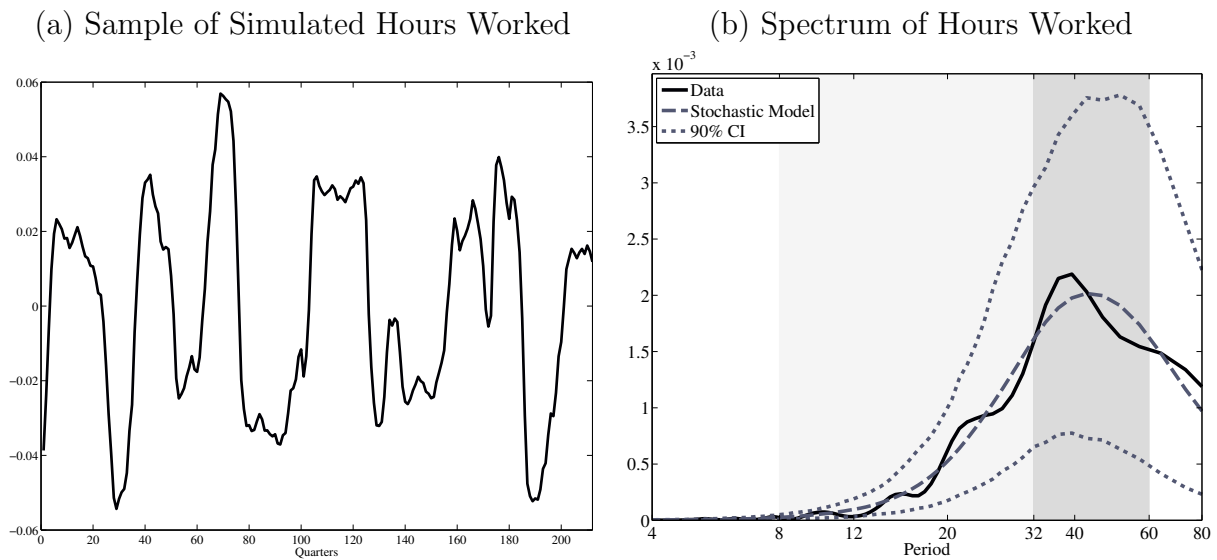
Next, it should be emphasized that the exogenous shock process in this model primarily accelerates and decelerates the endogenous cyclical dynamics, causing significant random fluctuations in the length of the cycle while only modestly affecting its amplitude. For example, in the deterministic version of the model the standard deviation of log-hours is 0.024, while in the stochastic model it is 0.033, implying that almost three quarters of the

⁶¹That 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.

⁶²Note that the ACF is simply the inverse Fourier transform of the spectrum. Since the spectra of the model and data are similar, we would expect the ACFs to be similar as well, a property clearly verified in Figure 8.

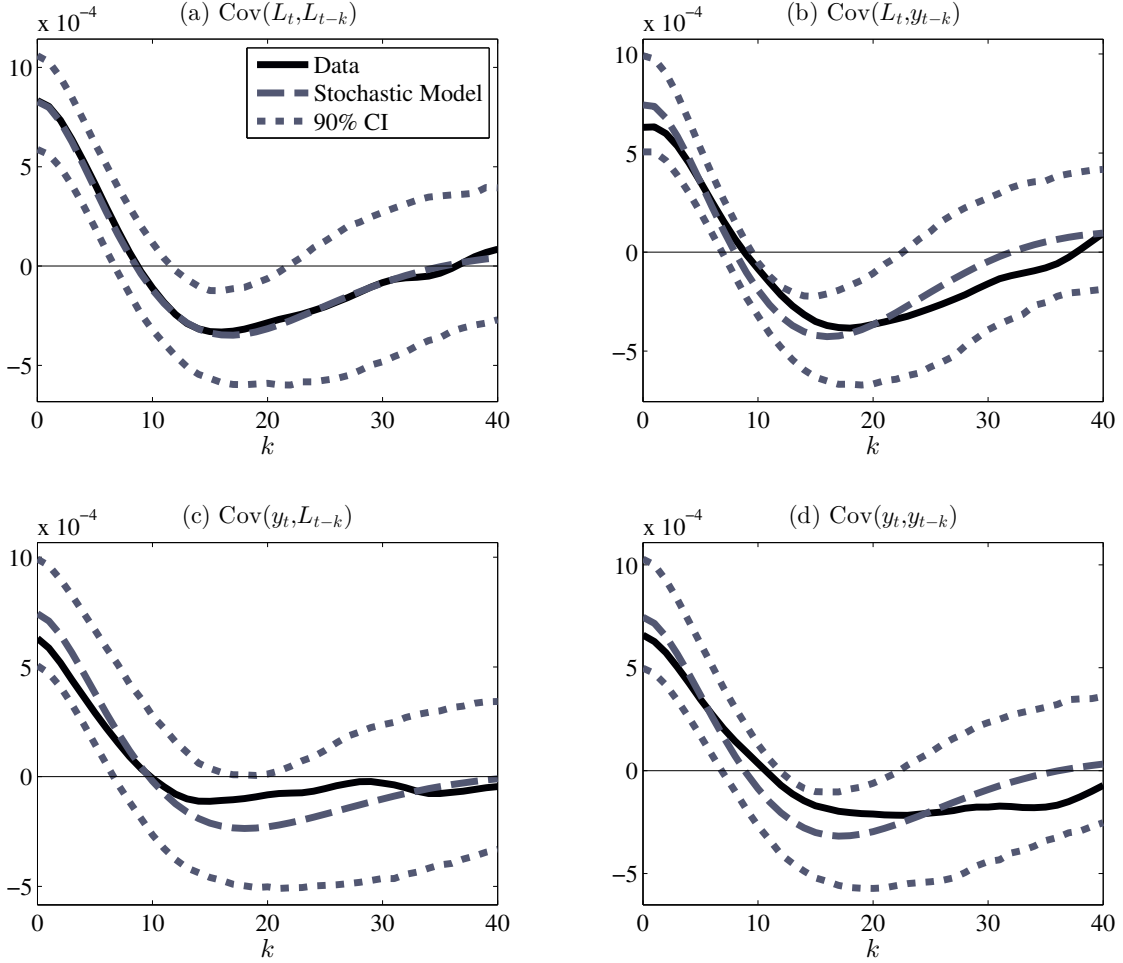
⁶³Note that the model spectra were obtained using the baseline model parameters as reported in Table 1.

Figure 7: Stochastic Model



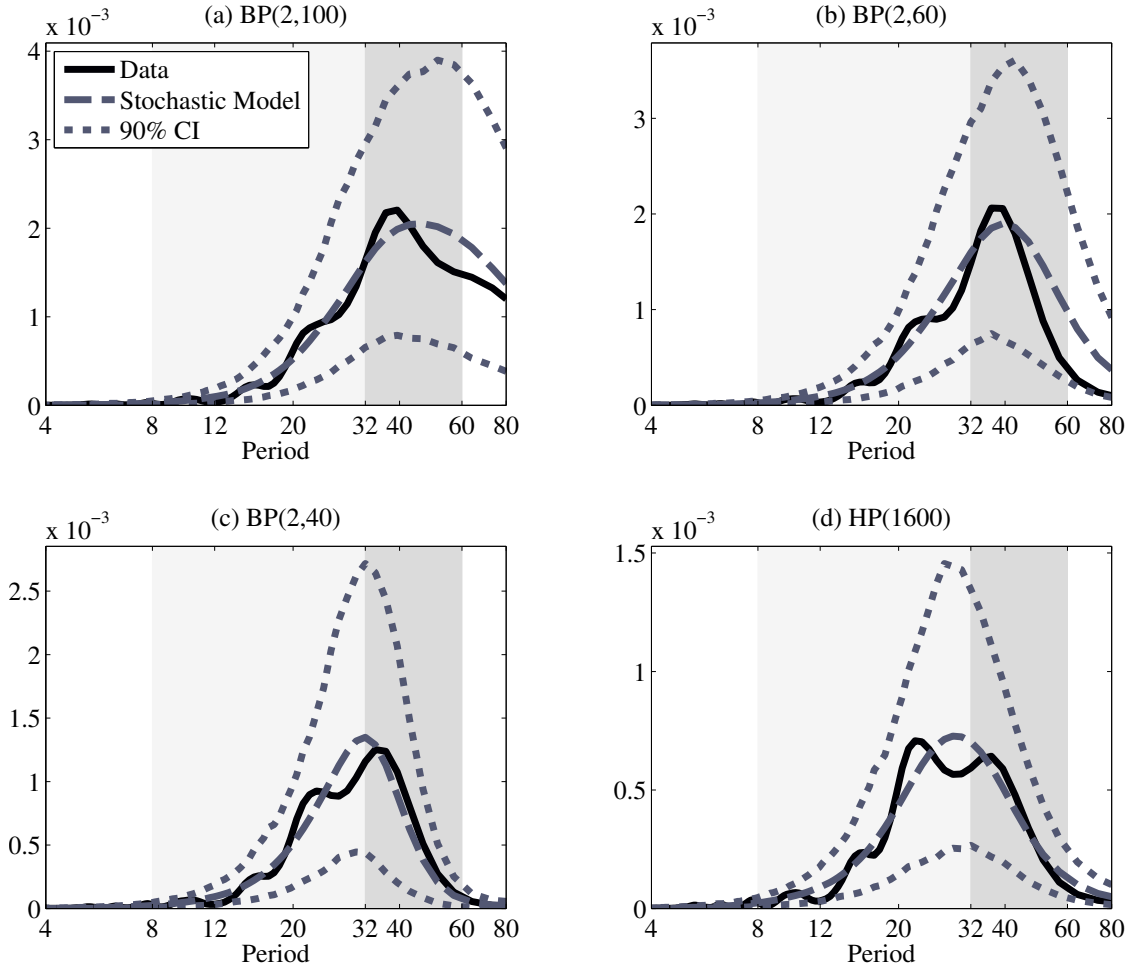
Note: Panel (a) shows 212-quarter simulated sample (same size as data set) of band-pass filtered $\log(\text{hours worked})$ (where hours are computed as $L_t = (1 - u_t)\ell_t$) generated from the stochastic model. Initial simulated series was 252 quarters long, with first and last 20 quarters discarded after band-pass 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 8: Autocovariance: Hours Worked (L) and Output (y)



Notes: 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 vacancy costs, i.e., $\tilde{\theta}_t \left[(1 - u_t)F(\ell_t) - \int_0^{N_t} \Phi(x) dx \right]$, 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 9: Spectrum: Hours Worked (Alternative Filters)



Notes: 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).


standard deviation of hours is due to deterministic mechanisms. Further, if this TFP process were the only shock process operating in, for example, the widely-cited model of Smets and Wouters [2007], it would generate a standard deviation of log-hours of only 0.006. 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 parsimonious set of shocks with moderate volatility can potentially yield qualitatively and quantitatively reasonable fluctuations.

As a final exercise in this section, it is worth 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 10 shows the spectrum for hours worked as generated by the Smets and Wouters [2007] model at the reported median posterior parameter values. As suggested by the relatively close fit, their model also matches patterns in the hours data reasonably well.

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 10 presents such a decomposition. It is 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.⁶⁴ 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. Thus, in effect, the role played by exogenous mark-up shocks in Smets and Wouters [2007] is replaced in our model by an endogenous demand cycle that emerges naturally in our environment as a result of a sufficiently strong incentive (because of a positive feedback mechanism from demand to unemployment risk) for agents to coordinate their purchase decisions.

⁶⁴The 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.

Figure 10: Hours Worked in Smets-Wouters (SW)



Notes: Data spectrum is as in Figure 5. 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 de-trended 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.

3.5 Additional Results

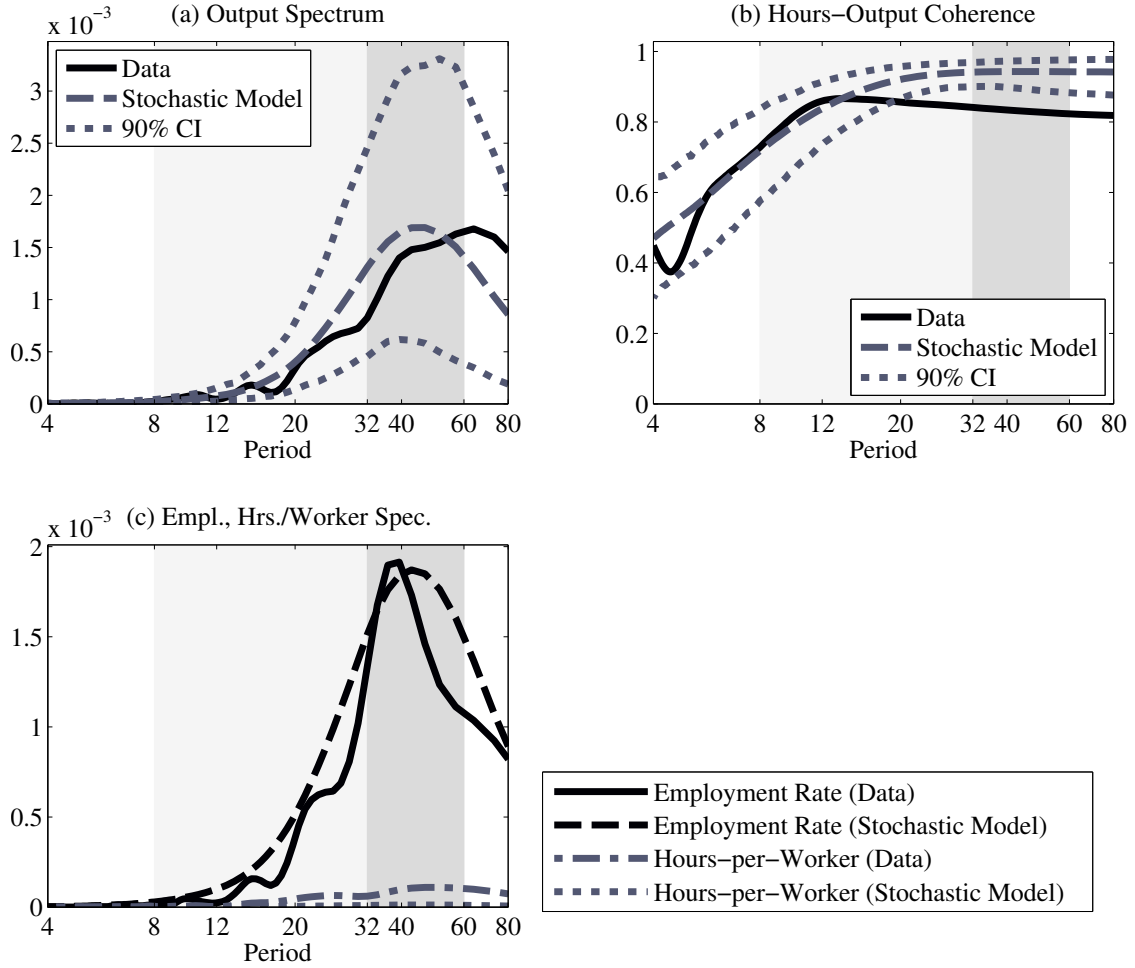
To this point, we have focused on the fit of the model with respect to the target series, hours worked. In this subsection, we evaluate how well the model performs in several other dimensions that were not directly targeted.

Panel (a) of Figure 11 compares the spectrum of output for the data and the stochastic model. 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. Panel (d) of Figure 8, 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 lie well within 90% confidence intervals for the model, suggesting a relatively good overall fit.

Next, panel (b) of Figure 11 plots the coherence between hours and output for the data and for the stochastic model. 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 slope, reaching 0.82 at the 80-quarter periodicity. The model coherence, meanwhile, rises somewhat over this range. Notwithstanding this discrepancy, 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 8 plot the cross-covariance function (CCF) for hours and output. Two

Figure 11: 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 detrended 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 vacancy costs, i.e., $\tilde{\theta}_t \left[(1 - u_t)F(\ell_t) - \int_0^{N_t} \Phi(x) dx \right]$, 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).

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 $k = 0$,⁶⁵ 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, $1 - u_t$, and hours-per-worker, ℓ_t . Panel (c) of Figure 11 shows spectra for the data and stochastic model for these two series.⁶⁶ 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.⁶⁷ 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.⁶⁸

Conclusion

Business cycle analysis is currently dominated by the pure impulse-propagation framework, whereby macroeconomic fluctuations are entirely explained as the results of exogenous shocks

⁶⁵The peak of the data CCF is only 0.28% greater than it is at $k = 0$.

⁶⁶Data 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.

⁶⁷See for example Shimer [2005].

⁶⁸As Figure 11 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.

perturbing an otherwise stable system. In this paper we have presented theory and evidence that give support to a broader perspective which includes limit cycles forces as an element inherent to fluctuations. In particular, according to this alternative view, business cycles should be thought as arising from the interplay of two different classes of forces. On the one hand, there are endogenous cyclical forces (driven by complementarities) that favor the bunching of economic activity, which in turn causes the economy to be naturally unstable, so that it would undergo boom and bust cycles even in the absence of any disturbances. On the other hand, there are also exogenous forces that interact with the natural cycle of the economy to create irregularity in the business cycle, whereby each cycle is different from the preceding one. In this framework, endogenous cyclical forces do not live in a vacuum but are instead constantly being upset by exogenous disturbances. The exogenous forces continue to play an important role in generating business cycles, but their role is slightly changed: instead of explaining the totality of business cycle fluctuations, much of the effect of exogenous disturbances is to accelerate or decelerate the equilibrium tendency of the economy to cycle.

We have proceeded in two steps to make our case in support of incorporating limit cycle ideas into mainstream business cycle analysis. First, we presented a simple reduced-form model with accumulation, sluggishness and agent interactions. Using this framework, we showed how the presence of strategic complementarities can give rise to a Hopf bifurcation not as a knife-edge case, but quite generically when demand complementarities are present locally. In addition, we showed that the strength of strategic complementarities that is needed to produce a limit cycle is less than what is needed to generate multiple equilibrium. Second, we presented a fully specified stochastic dynamic general equilibrium model. In this structural model, unemployment risk and precautionary savings interact to create a demand complementarity that is operative near the economy's steady state. Agents' decisions to purchase goods are strategic complements in the model because when others are purchasing goods it is less risky to buy on credit, since the risk of unemployment is low. We allowed fluctuations in the model to be driven by both exogenous disturbances and an endogenous limit cycle. We then estimated the model to fit the spectral density of hours worked. Based on the estimated parameters, we found that the system would exhibit a limit cycle in the absence

of any shocks. However, we also found that exogenous shocks played an important role in making observed fluctuations in the model sufficiently irregular to match the properties of the data. In this sense, our empirical analysis suggests that macroeconomic fluctuations may be understood as the result of both forces, with limit cycles playing a central role. One of the attractive features of the model is that, as we show, it is able to quantitatively replicate business cycle phenomena with only a single, moderately volatile shock precisely because of the underlying presence of a limit cycle.

References

- ALAN, S., T. CROSSLEY, AND H. LOW (2012): “Saving on a rainy day, borrowing for a rainy day,” IFS Working Papers W12/11, Institute for Fiscal Studies.
- AZARIADIS, C., AND B. SMITH (1998): “Financial Intermediation and Regime Switching in Business Cycles,” *American Economic Review*, 88(3), 516–536.
- BEAUDRY, P., D. GALIZIA, AND F. PORTIER (2015): “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.
- (1989): “Stochastic Equilibrium Oscillations,” *International Economic Review*, 30(1), 85–102.
- (1998): “Indeterminacy and Sunspots with Constant Returns,” *Journal of Economic Theory*, 81(1), 58–96.
- BENHABIB, J., AND A. RUSTICHINI (1990): “Equilibrium cycling with small discounting,” *Journal of Economic Theory*, 52(2), 423–432.
- 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.
- BROCK, W., AND J. SCHEINKMAN (1976): “Global asymptotic stability of optimal control systems with applications to the theory of economic growth,” *Journal of Economic Theory*, 12, 164–190.
- CARROLL, C., AND W. DUNN (1997): “Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets,” in *NBER Macroeconomics Annual 1997, Volume 12*, NBER Chapters, pp. 165–230. National Bureau of Economic Research, Inc.
- CARROLL, C., M. SOMMER, AND J. SLACALEK (2012): “Dissecting Saving Dynamics: Measuring Wealth, Precautionary, and Credit Effects,” IMF Working Papers 12/219, International Monetary Fund.
- CARROLL, C. D. (1992): “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, 23(2), 61–156.
- CHALLE, E., AND X. RAGOT (2013): “Precautionary Saving over the Business Cycle,” PSE Working Papers hal-00843150, forthcoming in *The Economic Journal*.

- 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.
- COOPER, R., AND A. JOHN (1988): “Coordinating Coordination Failures in Keynesian Models,” *The Quarterly Journal of Economics*, 103(3), 441–463.
- 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.
- DEN HAAN, W. J., G. RAMEY, AND J. WATSON (2000): “Job Destruction and Propagation of Shocks,” *American Economic Review*, 90(3), 482–498.
- DEN HAAN, W. J., P. RENDAHL, AND M. RIEGLER (2014): “Unemployment (Fears), Precautionary Savings, and Aggregate Demand,” Working Paper.
- DIAMOND, P., AND D. FUDENBERG (1989): “Rational expectations business cycles in search equilibrium,” *Journal of Political Economy*, 97(3), 606–619.
- DIAMOND, P. A. (1982): “Aggregate Demand Management in Search Equilibrium,” *Journal of Political Economy*, 90(5), 881–94.
- EPSTEIN, L. G., AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57(4), 937–969.
- FERNALD, J. (2014): “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Federal Reserve Bank of San Francisco Working Paper 2012-19.
- GALIZIA, D. (2014): “Can a Limit-Cycle Model Explain Business Cycle Fluctuations?,” Working Paper.
- 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.
- GU, C., F. MATTESINI, C. MONNET, AND R. WRIGHT (2013): “Endogenous Credit Cycles,” *Journal of Political Economy*, 121(5), 940 – 965.
- HAMMOUR, M. L. (1989): “Endogenous Business Cycles: Some Theory and Evidence,” PhD Thesis, Massachusetts Institute of Technology.
- HEATHCOTE, J., AND F. PERRI (2012): “Wealth and Volatility,” Meeting Papers 914, Society for Economic Dynamics.

- 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.
- KUZNETSOV, Y. A. (1998): *Elements of Applied Bifurcation Theory*, vol. 112 of *Applied Mathematical Sciences*. Springer, New-York, 2 edn.
- LAGOS, R., AND R. WRIGHT (2003): “Dynamics, cycles, and sunspot equilibria in ‘genuinely dynamic, fundamentally disaggregative’ models of money,” *Journal of Economic Theory*, 109(2), 156–171.
- (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.
- LUCAS, R. J., AND E. C. PRESCOTT (1974): “Equilibrium search and unemployment,” *Journal of Economic Theory*, 7(2), 188–209.
- MARCET, 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.
- MATSUYAMA, K. (1999): “Growing Through Cycles,” *Econometrica*, 67(2), 335–347.
- (2007): “Credit Traps and Credit Cycles,” *American Economic Review*, 97(1), 503–516.
- (2013): “The Good, the Bad, and the Ugly: An inquiry into the causes and nature of credit cycles,” *Theoretical Economics*, 8(3), 623–651.
- MCKENZIE, L. (1976): “Turnpike Theory,” *Econometrica*, 5, 841–65.
- MYERSON, R. B. (2012): “A Model of Moral-Hazard Credit Cycles,” *Journal of Political Economy*, 120(5), 847–878.
- PANCRAZI, R. (2015): “The heterogeneous Great Moderation,” *European Economic Review*, 74, 207–228.
- RAVN, M., AND V. STERK (2012): “Job Uncertainty and Deep Recessions,” Discussion paper, University College of London.
- ROCHETEAU, G., AND R. WRIGHT (2005): “Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium,” *Econometrica*, 73(1), 175–202.
- (2013): “Liquidity and asset-market dynamics,” *Journal of Monetary Economics*, 60(2), 275–294.

- 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.
- SHLEIFER, A. (1986): “Implementation Cycles,” *Journal of Political Economy*, 94(6), 1163–1190.
- 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).
- WIKAN, 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

$$\lambda, \bar{\lambda} = \frac{T}{2} \pm \sqrt{\left(\frac{T}{2}\right)^2 - D} \quad (\text{6})$$

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, then their modulus is between zero and one, and thus they are both inside the unit circle. If the two eigenvalues are real, then they have the same sign and at least one of them is less than one in absolute value. From (A.2), we have that $\lambda + \bar{\lambda} \in (-1, 2)$. Therefore, if the eigenvalues are both negative, then they are both inside the unit circle. If they are both positive, let λ be the larger eigenvalue, and suppose $\lambda \geq 1$. Given that $\lambda\bar{\lambda} < 1$, we have $\bar{\lambda} < \frac{1}{\lambda}$ and thus $\lambda + \bar{\lambda} = T < 2$ implies that $\lambda + \frac{1}{\lambda} < 2$, which in turn implies $(1 - \lambda)^2 < 0$. This is not possible and hence we must have $\lambda < 1$. Since λ is the largest of two real positive eigenvalues, both eigenvalues must lie inside the unit circle. \square

A.2 Proposition 2

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

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

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 relationship between T and D :

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

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) in Figure 2.

When $F'(I^s) \rightarrow -\infty$, we have $D \rightarrow 0$ and $T \rightarrow 1 - \delta$, which corresponds to point E_1 in Figure 2. 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 belong 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 flip bifurcation occurs 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 either from point E or E' (for which $F'(I^s) = 0$), (T, D) will eventually exit the triangle \widehat{ABC} . At the point where the half-line along (A.6) starting from E (or E') crosses \widehat{ABC} , at least on eigenvalue will have a modulus one.

Consider first the case $\alpha_2 < \alpha_1$. In this case, the line (A.6) has a negative slope, and will cross either segment \overline{AB} ((a) in the figure) or segment \overline{BC} ((b) in the figure). In case (a), we will have a flip bifurcation since the eigenvalues will be real and one of them equal to -1 when crossing the triangle \widehat{ABC} . In case (b), we will have a Hopf bifurcation since the eigenvalues will be complex and will both have modulus 1 when crossing \widehat{ABC} . We will be in case (b) when $D = 1$ and $T > -2$. $D = 1$ implies $F'(I^s) = 1 - \alpha_2(1 - \delta)$. Plugging this into the expression for T , the condition $T > -2$ becomes $1 - \delta + \frac{\alpha_2 - \alpha_1}{\alpha_2(1 - \delta)} > -2$ 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, while if $\alpha_1 > \alpha_2 > \frac{\alpha_1}{(2 - \delta)^2}$, we have a Hopf bifurcation.

Consider next the case $\alpha_2 > \alpha_1$. In this case, the line (A.6) has a positive slope, and could potentially cross either segment \overline{AC} or segment \overline{BC} . If it crosses the segment \overline{BC} , the eigenvalues will be complex with modulus 1 when crossing \widehat{ABC} , so that we will have a Hopf bifurcation. We will be in this case when $D = 1$ and $T < 2$. $D = 1$ implies $F'(I^s) = 1 - \alpha_2(1 - \delta)$. Plugging this into the expression of T , the condition $T < 2$ becomes $1 - \delta + \frac{\alpha_2 - \alpha_1}{\alpha_2(1 - \delta)} < 2$ which can be simplified to $\alpha_2 < \frac{\alpha_1}{\delta^2}$. Therefore, if $\alpha_1 < \alpha_2 < \frac{\alpha_1}{\delta^2}$, we will have a Hopf bifurcation. If $\alpha_2 > \frac{\alpha_1}{\delta^2}$, then as we increase F' we would cross the segment \overline{AC} . However, this possibility is ruled out by our assumption that $\alpha_2 > \frac{\alpha_1}{\delta}$, which was imposed to guarantee a unique steady state.

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 \overline{BC} , so that we have a Hopf bifurcation. Putting all of these results together gives the conditions stated in Proposition 3. \square

A.4 Proposition 4

For this proposition, we make use of Wan's [1978] theorem and of the formulation given by Wikan [2013] (See Kuznetsov [1998] for a comprehensive exposition of bifurcation theory).

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 (invertible) 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

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

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

We first write (A.7) in 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 strictly increasing function, and is therefore invertible. Denote $G(\cdot) \equiv H^{-1}(\cdot)$. 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

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

The eigenvalues of M are the solution of the equation

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

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

$$\Lambda \equiv \begin{pmatrix} \lambda & 0 \\ 0 & \bar{\lambda} \end{pmatrix}$$

and

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

By construction, λ and $\bar{\lambda}$ are the eigenvalues of C . We introduce matrices

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

and

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

whose columns are eigenvectors of C and M , respectively, and are thus such that $C = V_C \Lambda V_C^{-1}$ and $M = V_M \Lambda V_M^{-1}$. We therefore have $C = V_C \Lambda V_C^{-1} = V_C V_M^{-1} M V_M V_C^{-1} = B M B^{-1}$ with

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

Let us make the change of variables $(y_{1t}, y_{2t})' = B \times (i_t, x_t)'$ 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 $\gamma_1 \equiv -\alpha_2(1 - \delta) + (\alpha_2 - \alpha_1) \cos \theta$, $\gamma_2 \equiv \alpha_2 - \alpha_1$.

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

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

and

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

With $\mu \equiv F'(I^s)$ as the bifurcation parameter, we have $|\lambda| = \det(M) = \frac{\alpha_2(1-\delta)}{1-\mu}$, 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 functions of F''' , while 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 so 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 $a < 0$, in which case by Wan's [1978] theorem the limit cycle is supercritical.

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 we must have $\limsup_{t \rightarrow \infty} X_t = \infty$. 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

$$a - b(X_t + e_t) = vp(e_t)[1 + \tau u(e_t)] - \beta(1 - \delta)vp(e_{t+1})[1 + \tau u(e_{t+1})], \quad (\text{B.11})$$

$$X_{t+1} = (1 - \delta)(X_t + e_t), \quad (\text{B.12})$$

where $p(\cdot)$ and $u(\cdot)$ are defined in the main text. For a given state X_t and a given anticipated level of e_{t+1} , a sufficient condition to ensure that (B.11) has a unique solution is given by

$$b > vp^* \frac{\tau}{e^*} \equiv b_0, \quad (\text{B.13})$$

where e^* is output per firm (net of vacancy costs) when the economy is in the unemployment regime and p^* is the price in the unemployment regime, as described in section 2.2 (see footnote 32 regarding p^*). We henceforth assume that (B.13) holds.

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

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

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.13).

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.12)-(B.13) 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

$$b = [1 + q(1 - \delta)^2] b_0 \quad (\text{B.14})$$

for some $q > 0$, and note that as long as $\delta < 1$, which I henceforth assume, such a value of b satisfies (B.13). 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 vacancy costs) when the economy is in the unemployment regime and p^* is the price in the unemployment regime, as described in section 2.2 (see footnote 32 regarding p^*). It can be verified that maximizing

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

subject to (11) 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.

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:

$$a - b \left(X_t + \tilde{\theta}_t \tilde{e}_t - hc_{t-1} \right) + (1 - \delta) \gamma \lambda_t = \tilde{\theta}_t^{-1} \frac{\nu_1}{\alpha A} [\ell(\tilde{e}_t)]^{\omega+1-\alpha} [1 + \tau u(\tilde{e}_t)] + \mu_t, \quad (\text{C.15})$$

$$\mu_t = \mathbb{E}_t \left\{ \beta h \left[a - b \left(X_{t+1} + \tilde{\theta}_{t+1} \tilde{e}_{t+1} - hc_t \right) \right] \right\}, \quad (\text{C.16})$$

$$\lambda_t = \mathbb{E}_t \left\{ \beta \left[a - b \left(X_{t+1} + \tilde{\theta}_{t+1} \tilde{e}_{t+1} - hc_t \right) + (1 - \delta) \lambda_{t+1} - \mu_{t+1} \right] \right\}, \quad (\text{C.17})$$

$$c_t = X_t + \tilde{\theta}_t \tilde{e}_t, \quad (\text{C.18})$$

$$X_{t+1} = (1 - \delta) \left(X_t + \gamma \tilde{\theta}_t \tilde{e}_t \right). \quad (\text{C.19})$$

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

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

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

where $e^* \equiv \frac{\alpha}{1-\alpha} \bar{\Phi}$ 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.18) and (C.19)), 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.15) can be solved for \tilde{e}_t . To obtain values of μ_t and λ_t , we 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.16) and (C.17) are assumed to be functions only of Y_t , i.e.,

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

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

We 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 . We 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

$$\begin{aligned} \mu_t &= \Theta_\mu' Y_t^{(N)}, \\ \lambda_t &= \Theta_\lambda' Y_t^{(N)}. \end{aligned}$$

These values and values for the state variables can be plugged into (C.15) to yield a solution for \tilde{e}_t , which can then be replaced in (C.18) and (C.19) to obtain values for the subsequent period's state. In practice, we use $N = 2$.⁶⁹

To obtain Θ_μ and Θ_λ , we proceed iteratively as follows. Begin with some initial guesses $\Theta_{\mu,0}$ and $\Theta_{\lambda,0}$,⁷⁰ 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.16) and (C.17). 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.

⁶⁹We experimented with larger values of N and found that it resulted in a substantial increase in computational time without significantly affecting the results.

⁷⁰In practice, we 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.

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 band-pass 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.