

NBER WORKING PAPER SERIES

PUTTING THE CYCLE BACK INTO BUSINESS CYCLE ANALYSIS

Paul Beaudry
Dana Galizia
Franck Portier

Working Paper 22825
<http://www.nber.org/papers/w22825>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2016

The authors thank Jess Benhabib, Kiminori Matsuyama and Morten Ravn for helpful discussions. The authors would also like to thank seminar participants at Banque de France, University of Manchester, Pompeu Fabra-Toulouse "skiminar", NBER Summer Institute, Hydra workshop, UCL, UCLA, UCSD, Wisconsin, the University of Pennsylvania and Northwestern University for comments. This is a heavily revised version of a paper previously circulated under the heading "Reviving the Limit Cycle View of Macroeconomic Fluctuations." Franck Portier acknowledges financial support by the ADEMU project, "A Dynamic Economic and Monetary Union," funded by the European Union's Horizon 2020 Program under grant agreement No 649396. Paul Beaudry acknowledges support from the SSHRC of Canada. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Paul Beaudry, Dana Galizia, and Franck Portier. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Putting the Cycle Back into Business Cycle Analysis
Paul Beaudry, Dana Galizia, and Franck Portier
NBER Working Paper No. 22825
November 2016
JEL No. E10

ABSTRACT

This paper begins by re-examining the spectral properties of several cyclically sensitive variables such as hours worked, unemployment and capacity utilization. For each of these series, we document the presence of an important peak in the spectral density at a periodicity of approximately 36-40 quarters. We take this pattern as suggestive of intriguing but little-studied cyclical phenomena at the long end of the business cycle, and we ask how best to explain it. In particular, we explore whether such patterns may reflect slow-moving limit cycle forces, wherein booms sow the seeds of the subsequent busts. To this end, we present a general class of models, featuring local complementarities, that can give rise to unique-equilibrium behavior characterized by stochastic limit cycles. We then use the framework to extend a New Keynesian-type model in a manner aimed at capturing the notion of an accumulation-liquidation cycle. We estimate the model by indirect inference and find that the cyclical properties identified in the data can be well explained by stochastic limit cycles forces, where the exogenous disturbances to the system are very short lived. This contrasts with results from most other macroeconomic models, which typically require very persistent shocks in order to explain macroeconomic fluctuations.

Paul Beaudry
Vancouver School of Economics
University of British Columbia
6000 Iona Drive
Vancouver, BC V6T 1L4
CANADA
and University of British Columbia
and also NBER
paulbe@interchange.ubc.ca

Franck Portier
Toulouse School of Economics
University of Toulouse
Manufacture des Tabacs
21 Allée de Brienne
31000 Toulouse, FRANCE
and CEPR
franck.portier@tse-fr.eu

Dana Galizia
Department of Economics, Carleton University
Room B-845, Loeb Building
1125 Colonel By Drive
Ottawa, ON
K1S 5B6
Canada
dana.galizia@gmail.com

Introduction

It is well known that market economies repeatedly go through periods where, for sustained lengths of time, productive factors are used very intensively—with low rates of unemployment, high levels of hours worked per capita, and intensive use of productive capital—followed by periods where these utilization rates are reversed. A key aim of business cycle analysis is to explain such fluctuations. There are at least two potential types of explanations. According to the first, a period of low factor usage and high unemployment can be seen as mainly reflecting the effects of a negative shock to the economy, whereby an economy’s internal adjustment mechanisms have not yet managed to fully adjust to the outside disturbance. Within this class of explanations, economic booms and busts are not closely linked, since periods of high unemployment are not viewed as being caused by a previous boom period, but instead are viewed as resulting foremost from an external shock to the system. Such a framework for understanding business cycles generally results in a rich time-varying narrative in terms of the shocks driving observed fluctuations. An alternative class of explanation is one where boom and bust are tightly linked, being the natural outcome of market forces. The idea of endogenous boom-bust cycles is well captured by the expression that “a boom sows the seeds of the subsequent bust.” According to this alternative perspective, a low rate of factor usage arises as the equilibrium consequence of a prior boom period and, symmetrically, a bust period predicts the subsequent arrival of a boom period, rather than a return to normality. In this second class of explanations shocks can still play a role, but are generally less important since there are strong internal propagation mechanisms that would endogenously produce fluctuations even in their absence. Let us emphasize that, according to this second view, booms and busts are not reflective of animal spirits and/or indeterminacy, but are instead the unique equilibrium outcome of market forces.

With the risk of over-generalizing, it appears fair to say that a majority of modern macroeconomic models—especially those used by central banks—are much closer to the first type of explanation than the second.¹ There are good reasons for this state of affairs, the most

¹One way to classify the type of endogenous forces present in an estimated macro model is to *(i)* simulate the model with the autocorrelation of the exogenous shock processes set to zero (with all other parameters set to their estimated values), then *(ii)* plot the implied spectral densities of the main stationary endogenous variables. Unless these spectral densities all exhibit similar peaks at business cycle frequencies, then it

important of which is that equilibrium models that feature simple endogenous auto-regressive forces combined with persistent shocks have been shown to offer a quantitatively reasonable explanation to many aspects of business cycle data. Notwithstanding this feature, the goal of this paper is to present theory and evidence in favor of a view of fluctuations in which sustained booms and busts are the unique equilibrium outcome of a system of interacting agents. We approach the issue in three steps. First, we re-examine the cyclical properties of the US economy by focusing on the spectral densities of factor usage variables such as hours worked, unemployment, and capacity utilization over the period 1947-2015.² We will show that these variables suggest that the US economy exhibits important recurrent cyclical fluctuations at a periodicity of about nine to ten years. While this periodicity is slightly longer than traditionally associated with business cycles, it is only marginally so. For this reason, we believe these movements should be seen as part of the business cycle, as opposed to being associated with much longer medium-run cycles.³ Second, we will present a general class of models in which recurrent booms and busts emerge under fairly simple conditions. This class of models extends the work of Cooper and John [1988] to a dynamic environment, and shows how static strategic complementarities between agents, when embedded in a dynamic environment, tend to produce unique equilibrium outcomes characterized by boom-bust cycles. In particular, we will emphasize why strategic complementarities should be seen as more likely to produce a unique equilibrium featuring boom-bust cycles than multiple equilibria featuring sunspot-driven cycles. Third, we extend a simple New Keynesian model to allow for strategic complementarities and capital accumulation in a manner consistent with our general framework, and then estimate this model by indirect inference. The object of this exercise is to see whether this flexible New Keynesian model favors an interpretation of the data as featuring endogenously recurring cycles, or instead favors a predominantly shock-driven theory of cycles.

can be said that the model does not generate endogenous boom-bust cycles. According to this metric, most quantitative New Keynesian models and most real business cycle models do not generate endogenous boom-bust cycles, and can therefore be considered as part of the first class of explanations.

²In Beaudry, Galizia, and Portier [2014] we provide some complementary reduced-form evidence suggestive of recurrent boom-bust cycles in the US economy. However, the evidence presented there can be criticized as being sensitive to filtering procedures. The evidence we present here is more robust to such criticism.

³For example, we do not believe that these features should be lumped in with those emphasized in Comin and Gertler [2006], where medium-run cycles last up to 50 years.

Before proceeding further, it is necessary to clarify what we mean by an economy endogenously exhibiting boom-bust phenomena, as opposed to being driven mainly by shocks. To do so, it will be useful to draw a distinction between two different but related notions of an endogenous boom-bust cycle: a strong version, and a weak version. The strong version is the one we will focus upon in this paper, as it is less studied and allows for a much clearer distinction from most mainstream models. In the strong version, the macro-economy is seen as a locally unstable system, so that even in the absence of any shocks the economy would not settle down to a steady growth path. Instead, factor utilization would continuously fluctuate, neither exploding nor converging to a steady state. Accordingly, absent of shocks, this system would repeatedly go through booms and busts. The simplest embodiment of this idea is one where the deterministic forces of the system produce a limit cycle. This limit cycle interpretation will in fact be the operational notion we will associate with the strong version of the endogenous boom-bust cycle hypothesis.⁴

It is important to note that the idea that macroeconomic fluctuations may reflect limit cycle forces is not at all novel, having been advocated by many in the past, including early incarnations due to Kalecki [1937], Kaldor [1940], Hicks [1950] and Goodwin [1951].⁵ 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.⁶ The reasons why the notion of macroeconomic fluctuations driven by limit cycles plays little

⁴Another possibility is one where the deterministic forces of the economy produce chaotic behavior. In our empirical explorations, we have not found any evidence for such outcomes and therefore we will not focus on this possibility.

⁵An earlier mention of self-sustaining cycles as a model for economic fluctuations is found in Le Corbeiller [1933] in the first volume of *Econometrica*.

⁶There 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]).

role in modern quantitative research can be attributed to at least two difficulties. First, if the economy exhibited a deterministic limit cycle, the cycles would be very regular and predictable, which is inconsistent with the data. Second, the literature on limit cycles has generally made neither a clear empirical case nor a strong theoretical one for why they should be considered to be as or more relevant than the alternative explanations. An important contribution of this paper can therefore be seen as reviving the limit cycle view of fluctuations by offering new perspectives on these two arguments. In particular, with respect to the first argument, we directly address the criticism of the excessive regularity of limit cycles by examining instead the notion of a stochastic limit cycle, which corresponds to a system that is subject to exogenous shocks, but where the deterministic part of the system admits a limit cycle. Such systems have been studied little by quantitative macro-economists, but recent solution techniques make this a tractable endeavor. With respect to the second argument regarding the empirical relevance and theoretical plausibility of limit cycles, we address it by presenting a whole class of simple models that are capable of exhibiting stochastic limit cycles, and showing that this class of models finds support in the data. In particular, in the last section of the paper, we present a simple New Keynesian model with financial frictions and durable goods accumulation which fits into our general structure. The model is set up in a way so as to be flexible enough to allow for the possibility of limit cycles under certain parameterizations, but without imposing them. The type of cyclical behavior the model aims to capture is that generally associated with what we like to call accumulation-liquidation cycles, though others may prefer to think of them as a credit cycles. In our setup, agents make their consumption decisions to satisfy their Euler equation, where the interest rate they face reflects both the policy rate set by the central bank, as well as a risk premium. When the stock of durables (which may include housing) is sufficiently low, agents will want to make new purchases, and this will stimulate economic activity and employment. As the economy starts booming, two forces come into play. On the one hand, the policy interest rate may increase as the result of central bank action; on the other hand, the risk premium on borrowing may decrease. The net effect of these two forces drives the dynamics of the system, with the possibility of limit cycles arising. We use this model as a lens to interpret post-war macroeconomic data, and especially to see if, when estimated using observations

on hours worked and interest rate spreads, it favors a choice of parameters that interprets the data as embedding limit cycle forces.

While our aim is to provide theoretical and empirical support for the strong version of the endogenous boom-bust hypothesis—as embodied by the notion of stochastic limit cycles—it is nevertheless relevant to briefly discuss the weaker version of this hypothesis, and to highlight why we chose not to make it our focus. In the weak version, the system is such that, in the absence of shocks, the economy’s transitional dynamics feature dampened oscillations.⁷ For example, this could happen if the macro-economy can be represented as a linear system with complex eigenvalues of modulus smaller than one. Such a system can be seen as at least partially embedding an endogenous boom-bust mechanism. While it could be relevant to systematically study this weaker version of boom-bust cycles, we see it as less interesting than our current focus for two reasons. First, there are many models in the literature that satisfy this weaker definition, so that providing theoretical and empirical support for this weaker hypothesis is not very difficult, especially if we allow the weak version to be arbitrarily weak.⁸ Second, as we shall show, our main results also shed light on when the weaker version is likely to arise, and our empirical exploration does not rule out this possibility *ex ante* (though we find that the estimation nonetheless favors the strong version in the end). Hence, by exploring the strong version, we also provide a clear framework for thinking about the plausibility and relevance of the weak version.

The remaining sections of the paper are organized as follows. In Section 1 we highlight a set of spectral properties of U.S. data on hours worked, unemployment, capacity utilization, GDP, and TFP. These properties motivate our analysis and will be partly used later on in estimating our model. Note that the evidence presented in this section will not differentiate between the strong or weak versions of the endogenous boom-bust hypothesis. In Section 2 we present a simple reduced-form dynamic set-up where agents both accumulate goods and interact strategically with one another. Following Cooper and John [1988], these strategic interactions can be characterized either by substitutability or complementarity. In this

⁷Kiyotaki and Moore [1997] is one example of such a model (see Figure 3 of the paper).

⁸A very weak version of the hypothesis could correspond to any linear model with complex eigenvalues, even if those complex eigenvalues were close to zero in absolute value, and therefore unlikely to generate quantitatively important boom-bust cycles.

framework we discuss the conditions that tend to give rise to limit cycles, and in particular, we show how and when demand complementarities can cause the steady state of the system to become locally unstable and for a limit cycle to appear around it as part of a Hopf bifurcation.^{9,10} We further establish a simple condition under which this limit cycle will be attractive. We also use this section to introduce the notion of stochastic limit cycles, and to illustrate the effects of shocks in such a setup. To this point, we will have downplayed the role of forward-looking behavior so as to make the analysis as simple as possible. In Section 2.6, however, we extend the setup to include forward-looking elements, which allows us to discuss conditions under which equilibrium dynamics are determinate and the saddle path converges to a limit cycle. In Section E, we extend a standard three-equation New Keynesian model in a manner that allows for the features highlighted in our reduced-form model. For example, the model is extended to an environment where consumption services come from both newly purchased goods as well from an accumulated stock of durable goods. We introduce complementarities in this setup by allowing for an interest rate spread on household borrowing that is counter-cyclical. Section 3.3 takes the model to the data to see whether estimation will reveal the presence of a limit cycle or whether the data is better explained through more traditional channels. 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 persistent exogenous disturbances in explaining business cycle fluctuations. Finally, in the last section we offer concluding comments.

1 Motivating Observations

In this section we examine the cyclical properties of a set of quarterly U.S. macroeconomic variables covering the period 1947Q1-2015Q2.¹¹ One potential way of describing the cyclical


⁹Since 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.

¹⁰Informally, 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.

¹¹Sources for all data series are discussed in Appendix A.

properties of stationary data is to focus on the spectral density of the series, which depicts the importance of cycles of different frequencies in explaining the data. As is well known, if the spectral density of a time series displays a substantial peak at a given frequency, this is an indication of recurrent cyclical phenomena at that frequency.¹² The difficulty in using spectral analysis with macroeconomic data is that these data are frequently non-stationary, and therefore some detrending procedure must be done to make a series stationary before looking at its spectral density. This problem is especially acute for quantity variables such as GDP. However, for labor market variables, the problem is less severe since these data can be considered close to stationary. Accordingly, we begin our exploration of cyclical properties by focusing on U.S. non-farm business hours worked per capita. The series is plotted in Panel (a) of Figure 1.

Figure 1: Properties of Hours Worked per Capita



Notes: Panel (a) plots the log of Non-Farm Business Hours divided by Total Population. Panel (b) is an estimate of the spectral density of hours in levels (black line) and for 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (grey lines). A high-pass (P) filter removes all fluctuations of period greater than P .

As the figure shows, over the sample period hours worked exhibited substantial fluctuations, but with limited evidence of any long-run trend. For this reason, it may be reasonable to look directly at the spectral density of this series without any initial detrending transfor-

¹²See for example Sargent [1987].

mation. Panel (b) of Figure 1 plots several different versions of the hours worked spectral density over the range of periodicities from 4 to 80 quarters. The dark line in the figure represents the spectral density of the demeaned series (i.e., without any attempt to remove low-frequency movements that may reflect non-stationarities).¹³ Since we do not want to claim that this data is necessarily completely stationary, we also plot in the figure the spectra obtained after first detrending the data using a number of different filters designed to remove very low-frequency movements. In particular, each grey line in the figure represents the spectral density of the series after it has been detrended using a high-pass filter that removes fluctuations with periodicities greater than x quarters in length, where x ranges from 100 to 200. The results suggest that the spectral properties of hours worked at periodicities below 60 quarters—the range of periodicities we will focus on henceforth—are very robust to whether or not one first removes very low-frequency movements from the series.

What does Panel (b) of Figure 1 reveal about the cyclical properties of hours worked? To us, the dominant feature that emerges is the distinct peak in the spectral density at around 40 quarters, with this peak being much more pronounced than anything found at periodicities lower than 32 quarters. This suggests that a significant proportion of the fluctuations in hours worked are coming from some type of cyclical force with a periodicity of about 10 years. It is interesting to notice that this spectral hump is mainly contained in the range from 32-50 quarters, and therefore just slightly beyond the traditional range of periodicities (6-32 quarters) usually thought of as capturing business cycle fluctuations. Note also that this peak, while at a slightly lower frequency than traditional business cycle analysis, is not capturing the type of medium-run phenomena emphasized in work by Comin and Gertler [2006] who focus indistinctly on cycles between 32 and 200 quarters. In our opinion, this hump should be considered as part of the business cycle, suggesting that the traditional definition of the business cycle may be slightly too narrow. To make this case, in Panel (a) of Figure 2 we plot the hours worked series after having removed fluctuations longer than 60 quarters using a high-pass filter. In the figure, we also highlight NBER recessions in grey. As

¹³We obtain non-parametric power spectral density estimates by computing the discrete Fourier transform (DFT) of the series using a fast Fourier transform algorithm, and then smoothing it with a Hamming kernel. One key element is the number of points in the DFT, which determines the graphical resolution. In order to be able to clearly observe the spectral density between periodicities of 32 to 50 quarters, we use zero-padding to interpolate the DFT (see Section D in the Online Appendix for more details).

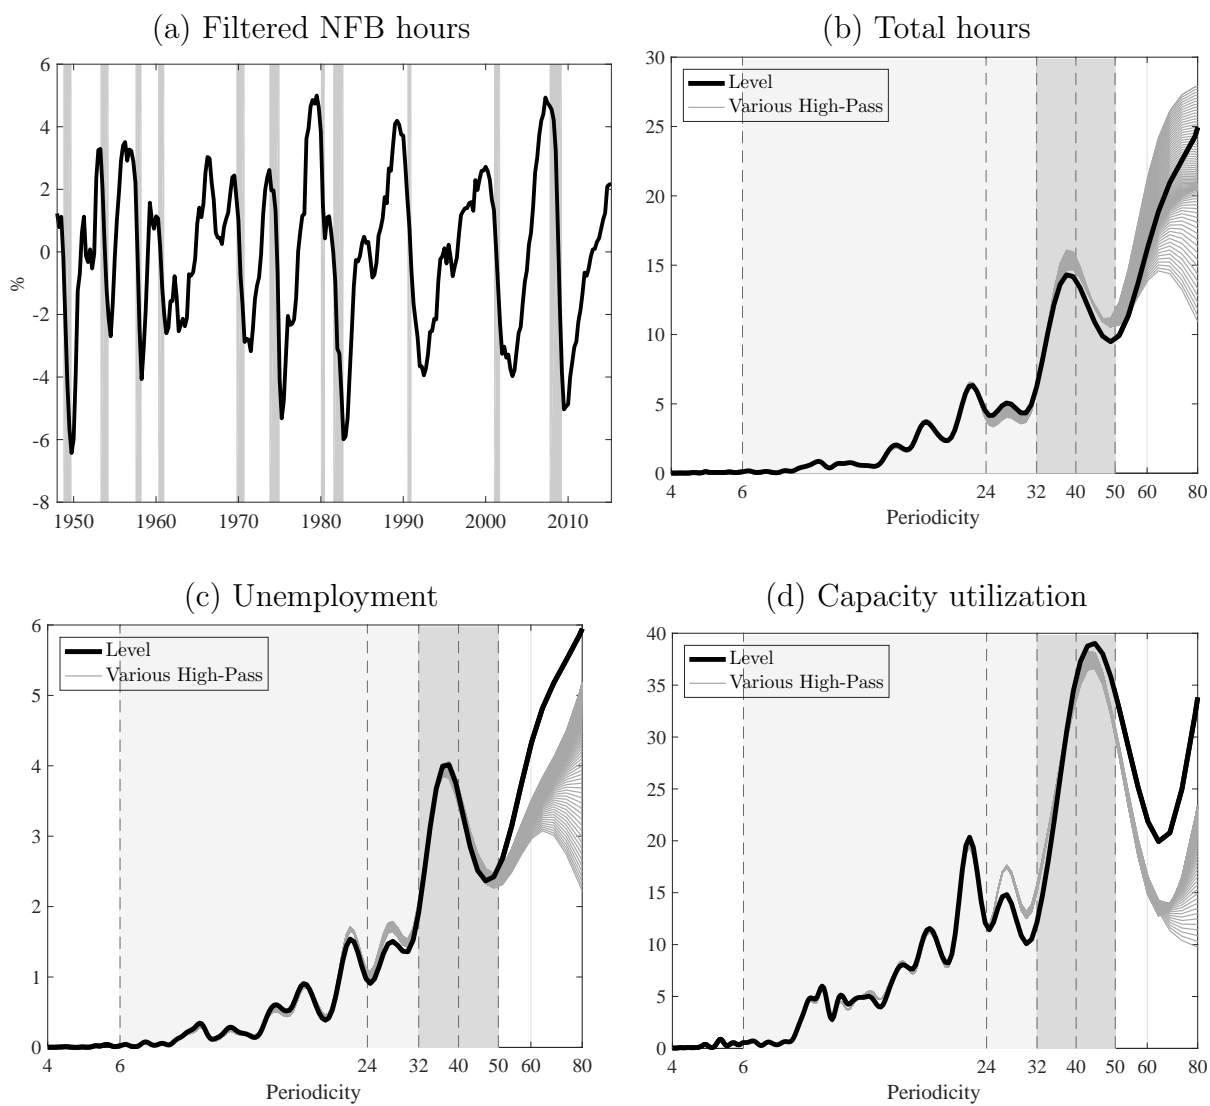
can be seen, the fluctuations in detrended hours that we observe when retaining these slightly lower frequencies match extremely well with the standard narrative of the business cycle, rather than reflecting, say, lower-frequency movements that are unrelated to the business cycle.

To further support the notion of a peak in the spectral density of work input around a periodicity of 40 quarters, in Panels (b) and (c) of Figure 2 we redo the same exercise for two other measures of work activity. Panel (b) reports the spectral density for total hours worked per capita, and Panel (c) the spectral density for the unemployment rate. In addition, in Panel (d) we report the spectral density of a capacity utilization measure. In all three cases, we plot both the spectral density for the untransformed data (dark line), as well as a set of spectra where we first remove low-frequency movements using high-pass filters as in Figure 1 (light grey lines). We highlight in dark grey the band of frequencies corresponding to periodicities from 32 to 50 quarters.¹⁴ As can be seen, in all of these plots we see distinct peaks in the spectra around 40 quarters, and these peaks are apparent regardless of whether or not we first remove very low-frequency movements. Together, Panel (b) of Figure 1 and Panels (b)-(d) of Figure 2 indicate that the aggregate utilization of workers and capital by firms in the U.S. exhibits important recurrent cyclical phenomena at approximately ten-year intervals. To explain such phenomena, one needs to rely on a theory where substantial cyclical forces are present. While there are different mechanisms that can explain such observations, we will explore in subsequent sections the plausibility and relevance of stochastic limit cycle forces as a potential candidate.

Our observations of a distinct peak in the spectral density of a set of macroeconomic variables may appear somewhat at odds with conventional wisdom. In particular, it is well known, at least since Granger [1966], that several macroeconomic variables do not exhibit such peaks, and for this reason the business cycle is often defined in terms of co-movement between variables instead of reflecting distinct cyclical timing. This view of business cycle dynamics depends on the variables one chooses to examine. We have focused on variables—which we like to call cyclically sensitive variables—where business cycle fluctuations are large

¹⁴Given our stated focus on periodicities below 60 quarters, we cut the x -axis off at this point. For higher periodicities, as we have seen in Panel (b) of Figure 1, high-pass filtering is likely to affect the spectral density estimate.

Figure 2: Cyclical Component and Spectral Densities




Notes: Panel (a) plots the high-pass (60) filtered version of Non Farm Business Hours per capita for the U.S.A. Panels (b) to (d) show estimates of the spectral density of Total Hours per capita, Unemployment rate and Capacity utilization rate in levels (black line) and for 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (grey lines). A high-pass (P) filter removes all fluctuations of period greater than P .

in relation to trend movements. For such variables, the breakdown between trend and cycle is less problematic and one can more easily detect spectral density peaks. In contrast, if one focuses on quantity variables, for example GDP, one does not easily detect any such peaks. To see this, in Figure 3 we report the same information we reported before regarding the spectral density, but in this case the series is real per capita GDP. Here we see that the spectra of the non-detrended data and of the filtered data have very little in common with each other. Since it does not make much sense to report the spectral density of non-detrended GDP (it is clearly non-stationary), in Panel (a) of Figure 4 we focus on the spectral density of GDP after removing low-frequency movements using the various high-pass filters. These spectral densities are in line with conventional wisdom: even when we have removed very low-frequency movements, we still do not detect any substantial peak in the spectral density of GDP around 40 quarters. How can this be? What explains the different spectral properties of output versus hours worked? The main explanation to this puzzle lies in the behavior of TFP. In Panel (b) of Figure 4 we plot the spectral density of TFP after having removed low-frequency movements in the same way we have done for GDP and other variables. What is noticeable about the spectral density of TFP is the quick pick-up it exhibits just above periodicities of 40 quarters. As with GDP, we do not see any marked peaks in the spectral density of TFP. An interesting aspect to note is that if we add the spectral density of hours worked to that of TFP, we get almost exactly that of GDP. This suggests that looking at the spectral density of GDP may be a much less informative way to understand business cycle phenomena than looking at the behavior of cyclically sensitive variables such as hours worked and capacity utilization. Instead, GDP likely captures two distinct processes: the business cycle process associated with the usage of factors, and a lower-frequency process associated with movement in TFP. For this reason, we believe that business cycle analysis may gain by focusing more closely on explaining the behavior of cyclically sensitive variables, such as employment, which are less likely to be contaminated by lower-frequency movements in TFP, and this is precisely what we aim to do.

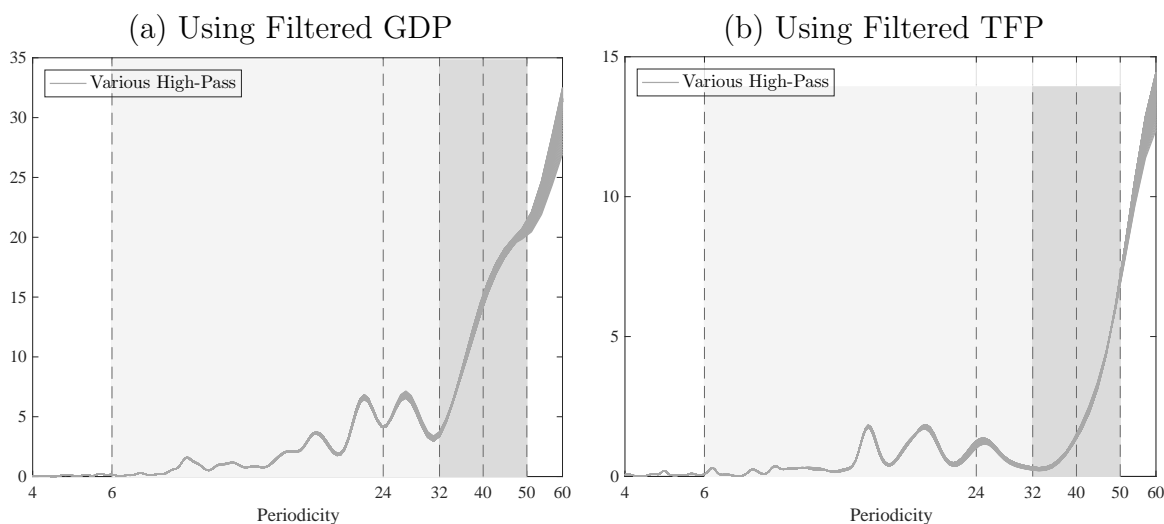
As a final data exploration step, we also examined whether the business cycle fluctuations which we have been focusing upon can be taken to be normally distributed. To this end, we perform D'Agostino and Pearson's [1973] and Jarque and Bera's [1987] omnibus tests,

Figure 3: Spectral Density of U.S. GDP per Capita, Levels and Various Bandpass Filters



Notes: This figure shows an estimate of the spectral density of U.S. GDP per capita in levels (black line) and for 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (grey lines). A high-pass (P) filter removes all fluctuations of period greater than P .

Figure 4: Spectral Density for GDP and TFP



Notes: This figure shows an estimate of the spectral density of U.S. GDP per capita (panel (a)) and Total Factor Productivity (panel (b)) for 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (grey lines). A high-pass (P) filter removes all fluctuations of period greater than P . TFP is the corrected quarterly TFP series of Fernald [2014].

which combine skewness and kurtosis into a single statistic, on our main series after using a high-pass filter to remove periodicities greater than 60 quarters, which allows us to retain all the variation that we have argued is relevant for the business cycle, while removing more medium- and long-run fluctuations. The null hypothesis for the test is normality. For non-farm business hours, the unemployment rate, and capacity utilization, the p -values we find are, respectively, 5%, 1% and smaller than 1% for the D’Agostino-Pearson test and 7%, 1% and close to 0% for the Jarque-Bera test. This indicates that linear-Gaussian models might not be an appropriate way to describe these data, and that one may need to allow for some type of non-linearity in order to explain these movements. Accordingly, we will explore a class of explanation that allows for such non-linearities. Moreover, when we estimate our model below, we will use the skewness and kurtosis properties of the data to help identify parameters.

2 Demand Complementarities as a Source of Limit Cycles: A Simple 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, and how this can create a peak in the spectral density of a variable. As we show, the key mechanism that allows the model to generate limit cycles is the interplay between demand complementarities and dynamic accumulation. We begin with a reduced-form setup so as to highlight the generality of this mechanism, regardless of the precise micro-foundations that drive agents’ interactions. In the next section, we will provide a structural model that fits into this general class.

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 equilibria.

2.1 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 productive 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 parameters α_1 , α_2 and δ are strictly between 0 and 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 a sluggish response that may be due, for example, to adjustment costs.¹⁵

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 Appendix B. According to Proposition 1, the dynamics are extremely 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 the dynamics are affected.

¹⁵We 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).

2.2 Adding Interactions Between Agents

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

$$I_{it} = \alpha_0 - \alpha_1 X_{it} + \alpha_2 I_{it-1} + \mathbb{E}_{it}[F(I_t)], \quad (4)$$


while keeping the law of motion for X (Equation (1)) unchanged. Here, $I_t \equiv \sum I_{jt}/N$ is the average level of investment in the economy and \mathbb{E}_{it} is the expectation operator. Hence, in this setup, agent actions depend on their expectations of other agents' actions. 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 , affect agent i 's investment decision I_{it} . For example, if $F'(\cdot) < 0$ then the function $F(\cdot)$ can implicitly capture an agent's optimal response to an increase in prices caused by increased demand by others, while if $F'(\cdot) > 0$ it can capture some form of demand complementarity. As we will focus on a symmetric equilibrium and since there are no stochastic driving forces, we can drop the expectation operator and treat the system as deterministic. 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.

Figure 5 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 $F'(I_t) < 1$.¹⁶ Thus, we are restricting attention to cases where demand complementarities, if they are present, are never strong enough to produce static multiple equilibria.

In what follows we consider only symmetric equilibria, so that we may henceforth drop

¹⁶Under the condition $F'(I_t) < 1$, the static equilibrium depicted in Figure 5 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).

Figure 5: 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.

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.

2.3 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 neighbourhood 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

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

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

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

and

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

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 6, 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 proofs of Appendix B 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 $\Delta = \left(\frac{T}{2}\right)^2 - D$. The parabola in Figure 6 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 6 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 $\alpha_1 \neq \alpha_2$,¹⁷ 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

¹⁷The case where $\alpha_1 = \alpha_2$ is addressed in the proofs of Appendix B.

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 6. 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 6. 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:

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 6, 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 6. It can also be easily verified that, as $F'(I^s)$ gets

Figure 6: 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 proofs of Appendix B.

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 6. Under our earlier assumption that $\alpha_2 < \alpha_1/\delta$ (which helped guaranteeing 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 6. 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 6, 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.

2.4 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 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.¹⁸ Beginning from any point on this closed curve, the system will remain 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 in general 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 3

¹⁸We 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 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.

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 1.¹⁹

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.

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

¹⁹Note 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.

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.²⁰ Note that Figure 5 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 and the outer stability region. Such a configuration for F is the one we have assumed when drawing Figure 5.

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,²¹ 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.²²

²⁰A 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.

²¹A 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 $y \in (0, 1)$.

²²Note 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 neighbourhood 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 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 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.

2.5 Stochastic Limit Cycles

As we have seen, in dynamic environments with accumulation, strategic complementarities between agents' actions can readily create limit cycles in aggregate outcomes. We showed that this can arise even when individual-level behavior favors stability, in that the system would converge to the steady state in the absence of agent interactions. Moreover, in our environment the complementarities can be considered modest, in that they imply elasticities less than one. However, if the behavior of all agents is deterministic, then the resulting cyclical dynamics will be far too regular to match the patterns observed in macroeconomic data series. For example, in Figure 7 we report an example of the limit cycle dynamics implied by a simple parameterization of Equation (4).²³ While the dynamics are not those of a perfect sine wave, the pattern is nonetheless very regular.


For the concept of limit cycles to potentially have a chance of helping to explain macroeconomic fluctuations, it is necessary to include exogenous stochastic driving forces.²⁴ For

the system never becomes explosive, even if the Hopf bifurcation of Proposition 4 is subcritical.

²³In this simulation, the function F is assumed to be piecewise-linear. See Appendix C for more details.

²⁴An alternative route, which we do not pursue here, would be to focus on chaotic dynamics.

Figure 7: Deterministic Simulation



Notes: This figure shows the deterministic simulation of a model given by equations (10) and (1) for symmetrical allocations and with $\mu_t = 0 \forall t$. The function F is assumed to be piecewise-linear. See Appendix C for more details.

example, suppose we modify our agents' decision rule (4) to

$$I_{it} = \alpha_0 - \alpha_1 X_{it} + \alpha_2 I_{it-1} + F(I_t) + \mu_t, \quad (10)$$

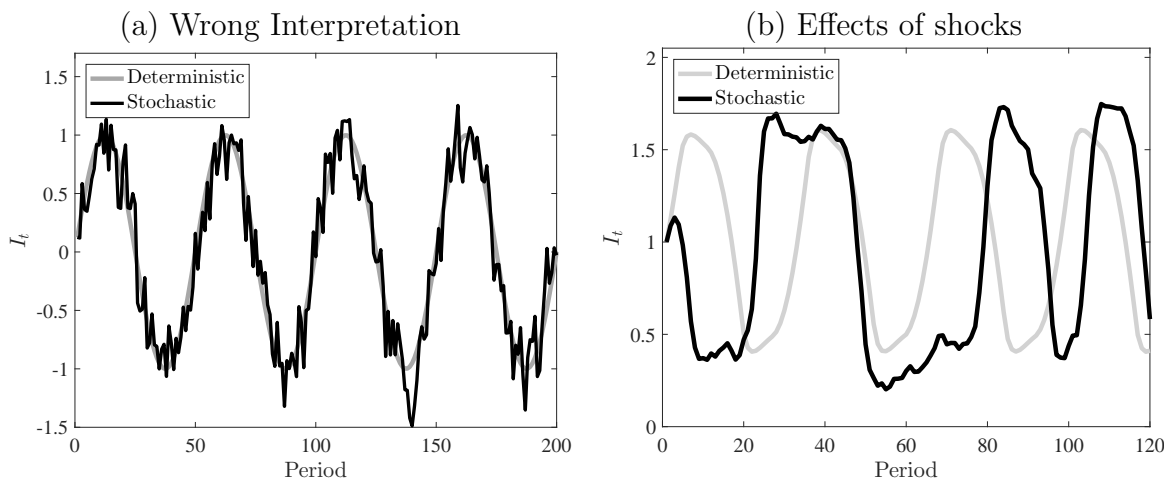
where μ_t is a stationary stochastic process representing an exogenous force affecting agents' investment decision.²⁵ How does the addition of this exogenous stochastic force affect equilibrium dynamics? In this case, we can no longer say that the dynamic system exhibits a limit cycle, but rather that it exhibits a *stochastic* limit cycle. What is important to understand is that the addition of the stochastic term μ_t does not simply add noise around an otherwise deterministic cycle, as for example would be the case in Panel (a) of Figure 8.²⁶ That is, the exogenous forces do not simply produce random amplitude shifts that temporarily perturb the system from the limit cycle. Rather, they also create random *phase* shifts that accelerate/delay the cycle itself. Furthermore, even though μ_t is stationary, and while the amplitude displacement caused by a shock is temporary, the phase displacement will nonetheless have a

²⁵We are restricting attention to the introduction of stochastic elements that enter in an additive form. Looking at the implications of allowing for stochastic elements that enter in some other form would be interesting but is left for future work. As we will focus on symmetric equilibria, and since the stochastic driving force is common across agents, we have again dropped the expectation operator.

²⁶Panel (a) of Figure 8 is generated with the model $I_t = \sin\left(\frac{\pi}{25}t\right)$ for the deterministic cycle, and with $I_t = \sin\left(\frac{\pi}{25}t\right) + \mu_t$ for the stochastic cycle, with $\mu_t = \rho\mu_{t-1} + \varepsilon_t$ and $\rho = 0.15$, $\sigma_\varepsilon = 0.2$.

permanent component.²⁷ To illustrate these effects, we return to the parameterization used in Figure 7, but now add in the shock μ_t .²⁸ The resulting dynamic system is presented in Panel (b) of Figure 8. This model would still have a deterministic cycle if no shocks were present, but now instead of observing a smooth cycle, the model generates data that look more like that observed for macroeconomic variables; namely, boom-and-bust cycles that have both stochastic amplitudes and durations. As we will later show, the addition of such stochastic elements also has the effect of changing the spectral density of a series. In the absence of stochastic elements, the spectral density implied by a limit cycle will tend to have very extreme peaks. In contrast, the presence of stochastic elements will tend to smooth out the spectral density, making peaks less pronounced but without generally removing them altogether.

Figure 8: Stochastic Limit Cycles



Notes: Panel (a) of this figure shows the deterministic and stochastic simulation of the model $I_t = \sin(\frac{\pi}{25}t) + \mu_t$. Panel (b) shows the deterministic and stochastic simulations of a model given by equations (10) and (1) for symmetrical allocations and where μ_t follows an AR(1) process. The function F is assumed to be piecewise-linear. μ_t is always equal to 0 in the deterministic simulations while it follows an AR(1) process in the stochastic ones. See Appendix C for more details.

²⁷That is, in the absence of any further shocks, the system will converge back to the limit cycle, but will be permanently either ahead of or behind where it would have been had it never been perturbed.

²⁸In this simulation, the F function is the piecewise-linear one used before, and the shock follows the autoregressive process $\mu_t = \rho\mu_{t-1} + \varepsilon_t$ and $\rho = 0.95$, $\sigma_\varepsilon = 0.03$. See Appendix C for more details.

2.6 Adding Forward-Looking Elements

In the previous section, we illustrated how agent interaction can create limit cycles when agents are accumulating goods and when there is inertia in their behavior. However, the framework we explored purposely omitted any forward-looking elements on the part of agents. This omission has the advantage of allowing us to highlight key forces that can give rise to limit cycles without needing to simultaneously address issues regarding the extra potential source of multiple equilibria that may arise when agents are forward-looking. In this section, we extend the results of the previous section to the case where individual agents' decisions take the form

$$I_{it} = \alpha_0 - \alpha_1 X_{it} + \alpha_2 I_{it-1} + \alpha_3 \mathbb{E}_{it} [I_{it+1}] + \mathbb{E}_{it} [F(I_t)] \quad (11)$$

Here, as before, the determination of X_t still obeys $X_{it+1} = (1 - \delta)X_{it} + I_{it}$, and the function $F(\cdot)$ controls the complementarity between agents' decisions. Also as before, we focus on symmetric equilibrium and assume that $0 \leq F'(\cdot) < 1$ so that the complementarity is never sufficiently strong that it creates static multiple equilibria.²⁹ For now we return to a formulation where we omit any exogenous stochastic driving forces and consider perfect foresight dynamics, so that we can remove the expectation operators in Equation (11). Note that when we consider an environment with forward-looking agents, microfounded models most often will include a transversality condition (TVC). For this reason, let us assume that our reduced form model also includes a TVC of the form $\lim_{t \rightarrow \infty} \beta^t X_t = 0$ for some parameter $0 < \beta < 1$.

With respect to our previous setup, the only difference in (11) is the inclusion of the expectation term $\mathbb{E}_t [I_{it+1}]$. We again restrict attention to situations where the parameters are such that there is a unique steady state. The condition $\alpha_1 > \delta(\alpha_2 + \alpha_3)$, which we henceforth assume holds, guarantees the absence of multiple steady states when $0 < F' < 1$. Note that as δ goes to zero, this constraint reduces simply to a non-negativity constraint on α_1 . We also assume henceforth that the parameters are such that, in the absence of any agent interaction (i.e., $F'(I) = 0$ for all I), the behavior of I_{it} is determinate and converges to the steady state. Having $\alpha_1, \alpha_2, \alpha_3 \in (0, 1)$ is almost sufficient for this.³⁰

²⁹We may also want to assume that $F''' < 0$ and that F' is largest at the steady state.

³⁰We need the additional restriction that if $(1 - \delta)\alpha_2 > \alpha_3$ then $\alpha_1 > (1 - \delta) + (\alpha_2 - \alpha_3) \left[1 + \frac{\alpha_2}{\alpha_3} (1 - \delta)^2 \right] -$

As before, let us denote the steady state of the system by I^s and let the strength of the complementarities be parameterized by $F'(I^s)$, which we will allow to vary from zero to one. Supposing initially that $F'(I^s) = 0$, we are starting from a situation where two of the eigenvalues of the dynamic system, λ_{11} and λ_{12} , are inside the unit circle, and the third eigenvalue, λ_2 , is real and outside; that is, we are starting from a saddle-path stable configuration. The following proposition establishes that we must have $\lambda_2 > 1$ in such a configuration.

Proposition 5. *Suppose the dynamic system governed by (11) and (1) has a real eigenvalue less than or equal to -1 . Then, for any value of $F'(I^s) < 1$, it also has a real eigenvalue that is greater than 1.*

Given Proposition 5, if this dynamic system undergoes a bifurcation as $F'(I^s)$ increases from zero to one, it must be one of the following four types:³¹

1. λ_2 falls below 1 (indeterminacy-inducing fold).
2. λ_{11} (or λ_{12}) is (or first becomes) real and rises above 1 (destabilizing fold).
3. λ_{11} (or λ_{12}) is (or first becomes) real and falls below -1 (flip).
4. λ_{11} and λ_{12} are (or become) complex and leave the unit circle together (Hopf).

The first type of bifurcation corresponds to the case where the system becomes locally indeterminate as complementarities increase, in the sense of generating a continuum of equilibrium paths satisfying the TVC. We will refer to this as an indeterminacy-inducing fold bifurcation. This is the type of bifurcation emphasized in the sunspot literature (see, for example, Benhabib and Farmer [1994]). The second type of bifurcation corresponds to a situation where the system no longer has any equilibrium paths that satisfy the TVC, so that the system becomes globally unstable.³² The third and fourth types of bifurcation correspond to the flip and Hopf bifurcations, which were discussed in Section 2.4, and which

³¹ $\frac{\alpha_2}{\alpha_3}(1 - \delta)$. Note also that, while we are assuming that $\alpha_3 < 1$, the results from this section extend to the case where α_3 is also allowed to be in the interval $(0, 2)$.

³¹We focus here only on the first bifurcation that the system undergoes as $F'(I^s)$ increases. In general, it may undergo further bifurcations beyond this point.

³²In particular, this is only necessarily true under our maintained assumption that the steady state is unique.

give rise (locally) to limit cycles. As noted earlier, the question that interests us most is under what conditions a Hopf bifurcation arises. However, before discussing this case, it is useful to first recognize that, in our setup with the parameter restrictions discussed above, the first two types of bifurcation never arise, as stated formally in the following proposition.

Proposition 6. *As $F'(I^s)$ increases from zero to one, the dynamic system governed by (11) and (1) will not undergo either type of fold bifurcation. Therefore, increasing the strength of the complementarities will not produce indeterminacy.*

Proposition 6 is interesting as it implies that, although we have extended our framework to include forward-looking behavior in the presence of complementarities across agents, we do not observe (self-fulfilling) multiple equilibria in our setup. This may seem surprising given that the literature on indeterminacy often seems to draw on complementarities to generate multiplicity. However, since fold bifurcations are always associated with multiple steady states (see, for example, Kuznetsov [1998], ch. 4), and since we have ruled out multiple steady states by assumption, it should not be too surprising that our system cannot undergo an indeterminacy-inducing (nor a destabilizing) fold bifurcation. Note that, as long as δ is small, the assumption made to rule out multiple steady states (i.e., $\alpha_1 > \delta(\alpha_2 + \alpha_3)$) will not be very restrictive. Given this, and given that it ensures a unique steady state, it seems reasonable for our purposes to maintain this assumption.

Proposition 6 leaves us with only two possible types of bifurcation that can arise as the strength of the complementarity increases: a Hopf bifurcation, or a flip bifurcation. Note that these were also the only two types of bifurcation that could arise in our analysis of the previous section where we omitted forward-looking behavior. As previously noted, flip bifurcations are unlikely to be very relevant in explaining macroeconomic phenomena, as they are associated with cycles of periodicity 2. In the current setup, such bifurcations can again only arise if α_2 is sufficiently small, as indicated by Proposition 7.

Proposition 7. *As $F'(I^s)$ increases from zero to one, the dynamic system governed by (11) and (1) will undergo a flip bifurcation only if $\alpha_2 < \frac{\alpha_1}{2-\delta} - \alpha_3$.*

Since flip bifurcations are not our focus, for the remaining discussion of this section we assume that $\alpha_2 > \frac{\alpha_1}{2-\delta} - \alpha_3$; i.e., we assume that individual behavior is characterized by sufficient sluggishness to rule out flip bifurcations. Under this added condition, Propositions 6 and 7 together imply that the only potential type of bifurcation in the system governed by (11) and (1) is a Hopf bifurcation.

Let us emphasize that the nature of the Hopf bifurcation in the presence of forward-looking behavior is somewhat different from the case without it. In particular, in the presence of forward-looking behavior, as the two complex stable eigenvalues become unstable due to a Hopf bifurcation, the system nonetheless retains a similar saddle-path structure to the pre-bifurcation system. Specifically, even if an attractive limit cycle appears, it will only be attractive on a two-dimensional manifold—which we refer to as the non-explosive manifold—in our three-dimensional phase space.³³ The transversality condition will then force I_t to jump—for any given initial values of X_t and I_{t-1} —onto this manifold. Figure 9 presents a phase diagram that illustrates this configuration, which we refer to as a saddle limit cycle. The light grey manifold in the figure contains an attractive limit cycle, and beginning from any point initially on this manifold (except the steady state), the system converges to the cycle (dark grey lines). The dynamics of the system beginning from any other point, however, are dominated by the remaining (real) unstable eigenvalue, causing paths emanating from such points to explode (black lines).³⁴


We are now in a position to examine under what conditions an increase in complementarity will lead to a Hopf bifurcation in our forward-looking system. Proposition 8 gives a necessary condition for this, while Proposition 9 gives sufficient one.

Proposition 8. *As $F'(I^s)$ increases from zero to one, the dynamic system governed by (11) and (1) can lead to a Hopf bifurcation only if $(1 - \delta)\alpha_2 > \alpha_3$.*

³³Formally, if A is the eigenspace associated with the two complex eigenvalues, then the non-explosive manifold is the invariant two-dimensional subspace that is tangent to A at the steady state.

³⁴The dynamic system we are considering is three-dimensional, which allows for the possibility of such a saddle limit cycle to appear as part of a Hopf bifurcation. In a two-dimensional system, on the other hand, a Hopf bifurcation would necessarily be associated with indeterminacy, since beginning from any point in the phase space the system would converge to the limit cycle (i.e., the transversality condition would be satisfied). This latter case has appeared often in the macroeconomic literature. In contrast, the type of configuration implied by Figure 9 is less common.

Figure 9: A Saddle Limit Cycle



Notes: The light grey zone is the non-explosive manifold (which is here locally drawn as a linear plane, but which is in general nonlinear). The dark grey lines are two paths that converge to the limit cycle, one from the inside and one from the outside. The black lines are two paths for which the jump variable has not jumped onto the non-explosive manifold, and which therefore diverge.


Proposition 9. *As $F'(I^s)$ increases from zero to one, the dynamic system governed by (11) and (1) will experience a Hopf bifurcation if $0 < (1 - \delta)^2 \frac{\alpha_2^2}{\alpha_3} + [1 - (1 - \delta)^2] \alpha_2 - \alpha_1 - \alpha_3 < (1 - \delta) \left(\frac{\alpha_2}{\alpha_3} - 1 \right)$.*

Proposition 8 indicates that complementarities may not be sufficient to create a Hopf bifurcation when agents are forward-looking. This is in contrast to the case with no forward-looking elements, where complementarities could always create local instability. In effect, the proposition indicates that the amount of sluggishness in the system needs to be sufficiently important (as governed by α_2) relative to the forward-looking behavior (as governed by α_3) for complementarities to potentially create a Hopf bifurcation. For example, if $\alpha_2 < \alpha_3$ then a Hopf bifurcation cannot arise. This condition will be helpful when exploring whether or not a particular model may allow for limit cycles.

Proposition 9 provides a sufficient condition for the emergence of a Hopf bifurcation due to complementarities in the presence of forward-looking agents. This condition is not easily interpretable. In order to get a better idea of its meaning, we illustrate its impact in Figure 10. The box in the figure corresponds to the set of α_1 , α_2 and $\alpha_3 \in [0, 1]$. The light grey balls correspond to the set of points where the system undergoes a Hopf bifurcation as $F'(I^s)$ increases from zero to one.³⁵ The figure is drawn for the case where $\delta = 0.01$ (similar figures arise for alternative values of δ). As can be seen in the figure, most of the points satisfying the necessary condition $(1 - \delta)\alpha_2 > \alpha_3$ actually support the emergence of a limit cycle as complementarities increase. In this sense, the condition $(1 - \delta)\alpha_2 > \alpha_3$ can be seen as necessary and *almost* sufficient for the emergence of a Hopf bifurcation, even though Proposition 9 indicates that the actual sufficient condition is more subtle. For the parameter space where the system does not experience a Hopf bifurcation as we increase $F'(I^s)$ towards 1, it can be further be shown that the eigenvalues of the system will necessarily become complex when $F'(I^s)$ is sufficiently large. In this sense, increasing complementarities in this system always tends to create cyclical forces, either of the strong form associated with limit cycles, or in the weaker form of complex eigenvalues in a stable system.

³⁵The darker grey balls in the figure represent combinations where the system undergoes a flip bifurcation, and the black balls correspond to indeterminacy-inducing fold bifurcations. Both of these case are in areas ruled out by our assumptions.

Figure 10: Nature of the Bifurcation When Complementarities Are Increased



Notes: This figure is related to the forward-looking model given by equations (1) and (11) for symmetrical allocations. For various values of the parameters $(\alpha_1, \alpha_2, \alpha_3) \in (0, 1)^3$ and for fixed depreciation $\delta = 0.01$, we increase the degree of complementarity $F'(I^s)$ from 0 to the point where the first bifurcation occurs. The light grey balls correspond to the triplets of parameters where the system undergoes a Hopf bifurcation. The darker grey balls in the figure represent combinations where the system undergoes a flip bifurcation, and the black balls correspond to indeterminacy-inducing fold bifurcations. Areas without balls correspond to parameter configurations for which the steady state is non-unique (i.e., $\alpha_1 \leq \delta(\alpha_2 + \alpha_3)$) or, if it is unique, the system is not initially in a saddle-path stable configuration (i.e., if $(1 - \delta)\alpha_2 > \alpha_3$ then $\alpha_1 > (1 - \delta) + (\alpha_2 - \alpha_3) \left[1 + \frac{\alpha_2}{\alpha_3}(1 - \delta)^2 \right] - \frac{\alpha_2}{\alpha_3}(1 - \delta)$).

3 A New Keynesian Model with Financial Frictions that Can Generate Limit Cycles

3.1 The Model

In this section, we extend the canonical three-equation New Keynesian model to include a set of elements that will allow it to capture the key forces emphasized in our reduced-form model. This involves adding to the canonical model a form of capital accumulation, some sluggishness in behavior, and most importantly some form of complementarity. In the next section, we will estimate this extended model by indirect inference and check whether the estimated parameters suggest that limit cycle forces may be an important driver of the business cycle.

Before we present the equations of the model, it is helpful to first briefly describe the mechanisms we want to capture with it. The model is aimed at capturing what we like to call an accumulation-liquidation cycle, but which others may prefer to think of as a credit cycle.³⁶ The narrative for such a cycle has many precedents in the literature, and our formulation is only one possible interpretation. Complementarity in our model will emerge through the determination of the risk premium faced by households on their borrowing, and the main capital variable will be household capital (which can be interpreted as combining durable goods and housing). The interest rate faced by households reflects the behavior of the central bank, as well as a risk premium. A limit cycle can potentially arise in our setup when the process of accumulation interacts sufficiently strongly with the determination of the risk premium. In such cases, when the economy is perturbed from its steady state, instead of returning to that point it will converge to a boom-bust cycle, wherein for several periods consumption expenditures are high, the risk premium is low, and the accumulation of capital is rapid, before the boom eventually becomes a bust as the excess stock of capital reduces demand for new durable goods/housing, which produces lower equilibrium expenditures and a higher risk premium. It is important to emphasize that the model we develop does not *a priori* impose the presence of such boom-bust cycles. Rather, our goal is simply to construct an environment that is flexible enough to capture such forces, and then estimate to the model

³⁶We are reluctant to call it a credit cycle as we adopt a setup where no credit is actually extended in equilibrium.

to see whether these forces may be present in the data. In the development of the model, we chose to keep all aspects as simple as possible in order to keep a close connection with the reduced-form model of the previous section. A more thorough derivation of a similar structure, where we are more explicit about the source of the market imperfections that can produce complementarities, can be found in a earlier version of this paper.³⁷

The formal environment we consider is one with a continuum of mass one of households indexed by i who purchase consumption services from the market. The preferences of agent i are given by $\mathbb{E}_0 \sum_t \beta^t \xi_{t-1} [U(C_{it} - \gamma C_{t-1}) + \nu(1 - L_{it})]$, with $U'(\cdot) > 0, U''(\cdot) < 0, \nu'(\cdot) > 0, \nu''(\cdot) < 0$ and $0 \leq \gamma < 1 - \delta$. C_{it} represents the consumption services purchased by household i in period t , C_t denotes the average level of consumption in the economy, β is the discount factor, and ξ_t denotes an exogenous shock to the discount factor at date t . Note that this preference structure assumes the presence of external habit, which will be the source of sluggishness in the economy.³⁸ Household i 's problem takes the form³⁹

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \xi_{t-1} [U(C_{it} - \gamma C_{t-1}) + \nu(1 - L_{it})]$$

subject to

$$P_t C_{it} + (1 + r_{t-1}) B_{it} + P_t^x D_{it} = B_{it+1} + r_t^x X_{it} + w_t L_{it} + \Gamma_{it},$$

$$X_{it+1} = (1 - \delta) X_{it} + D_{it}.$$

Here, B_{it+1} represents the borrowing by the household at time t , to be repaid with interest at $t + 1$, r_t represents the nominal interest rate faced by the household on such borrowing, $w_t L_{it}$ is labor income, X_{it} represents the stock of durable goods (or houses) held by the household at the beginning of period t , Γ_{it} are firm profits that are returned to households, and D_{it} the quantity of new durable goods purchased by the household at t . P_t and P_t^x are the (nominal) prices of consumption services and new durable goods, respectively, at date t . Households are assumed to buy all of their consumption services—including those derived from durable goods—from the market. Specifically, households do not consume the

³⁷See Beaudry, Galizia, and Portier [2015].

³⁸We assume that habit is external to the individual only for simplicity. Extending the analysis to have internal habit is not difficult.

³⁹The full details of the derivations in this section are presented in Section E of the Online Appendix.

services of their durable goods directly. Instead, they rent X_{it} out to firms each period at nominal rental rate r_t^x . Firms then combine the rented stock of durables with labor in order to produce consumption services, C , as well as new durable goods, D , in a manner to be described shortly. Under these conditions, dropping i subscripts, the household's Euler equation with respect to their optimal choice of C takes the familiar form^{40,41}

$$U'(C_t - \gamma C_{t-1}) = \beta_t \mathbb{E}_t [U(C_{t+1} - \gamma C_t) (1 + r_t - \pi_{t+1})], \quad (12)$$

where π_{t+1} is the inflation rate from t to $t + 1$, and $\beta_t \equiv \beta \frac{\xi_t}{\xi_{t-1}}$.

There are two types of firms in the model: final good firms, and intermediate good firms. The final good sector is competitive. This sector provides consumption services to households by buying a set of differentiated intermediate services, denoted C_{jt} , from intermediate good firms. We assume a measure one of intermediate good firms, indexed by j . The technology of the final good firms is

$$C_t = \left(\int_0^1 C_{jt}^\eta dj \right)^{\frac{1}{\eta}}, \quad (13)$$

with $\eta \in (0, 1)$. The objective of the final good firm is thus to solve

$$\max P_t C_t - \int_0^1 P_{jt} C_{jt} dj$$

subject to (E.22), where p_{jt} is the price of intermediate good j . This gives rise to demand for intermediate good j given by

$$C_{jt} = \left(\frac{P_{jt}}{P_t} \right)^{-\frac{1}{1-\eta}} C_t.$$

Intermediate good producers are monopolistically competitive and take the demand from final good firms as given. These firms also exhibit sticky prices, in that the arrival of the

⁴⁰We have here used the approximation $\frac{1+r}{1+\pi} = 1 + r - \pi$.

⁴¹The household's problem also produces a labor supply condition

$$\frac{\nu'(L_t)}{U'(C_t - \gamma C_{t-1})} = \frac{w_t}{P_t}$$

and an arbitrage condition between the holding of bonds and capital given by

$$\mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{P_{t+1}} \left\{ (1 + r_t) - \left(\frac{r_{t+1}^x + (1 - \delta) P_{t+1}^x}{P_t^x} \right) \right\} \right] = 0.$$

option to change their price follows a Poisson process. At this point, we could choose to close the model by specifying that the production of C_{jt} comes exclusively from currently produced goods. However, if this were the case then there would be no role for accumulation forces. Since we want to capture dynamics that may arise due to the accumulation and liquidation of household capital, we allow consumption services to come from either the production of new goods or from the service flow of existing durables. Specifically, we assume that intermediate firms produce an intermediate factor M according to the technology

$$M_{jt} = BF(\Theta_t L_{jt}),$$

with $B > 0$. A fixed fraction⁴² $(1 - \varphi)$ of that factor is transformed one-to-one into consumption services, whereas the use of the rented stock of durables also produces consumption services with a one to one technology, so that, letting X_{jt} denote the amount of capital rented by intermediate firm j , the amount of consumption services produced by this firm is

$$C_{jt} = X_{jt} + (1 - \varphi)M_{jt} = X_{jt} + (1 - \varphi)BF(\Theta_t L_{jt}) \quad (14)$$

The remaining fraction φ of the intermediate factor M is transformed one-to-one into new durable goods D , so that

$$D_{jt} = \varphi M_{jt} = \varphi BF(\Theta_t L_{jt})$$

Period- t profits of an intermediate good producer are given by

$$\Gamma_{jt} = P_{jt}C_{jt} + P_t^X D_{jt} - r_t^x X_{jt} - w_t L_{jt}.$$

Cost minimization implies that

$$w_t = [(1 - \varphi)r_t^x + \varphi P_t^x] \Theta_t BF'(\Theta_t L_{jt})$$

Hence, assuming $F'' < 0$, all intermediate good firms will hire the same amount of labor (i.e., $L_{jt} = L_t$), which implies that the aggregate capital stock will satisfy

$$X_{t+1} = (1 - \delta) X_t + \varphi BF(\Theta_t L_t). \quad (15)$$

⁴²Our modeling structure implies that agents do not separately choose how many new durable and non-durable goods to purchase. Instead, these are assumed to be in constant fixed proportions. This keeps the model very tractable and gives it a structure almost identical to what we have presented in previous sections. Allowing for a more general treatment that decouples these two decisions is not difficult, but it makes the model more complex and less directly comparable to the previous sections. We have chosen the simpler route so as to favor more clarity regarding the functioning of the model.

With the normalization $B = \frac{1}{1-\varphi}$, and letting $\psi \equiv \frac{\varphi}{1-\varphi}$, we obtain

$$C_t = X_t + F(\Theta_t L_t) \quad (16)$$

and

$$X_{t+1} = (1 - \delta) X_t + \psi F(\Theta_t L_t). \quad (17)$$

In contrast to the baseline New Keynesian model, we do not assume that the interest rate faced by households is equal to the policy rate set by the central bank. Instead we assume that it is equal to the policy interest rate (which we denote i_t) plus a risk premium r_t^p ; that is, we assume that $r_t = i_t + r_t^p$. The environment we consider is one where the risk premium r_t^p potentially changes with the state of the economy in a way that can produce complementarities between agents. In particular, we allow for the possibility that, when the labor market is tight, creditors may perceive lending to households to be less risky, leading to a fall in the risk premium. To keep the model as simple as possible, we directly allow for this counter-cyclical risk premium by positing r_t^p to be a non-increasing function of the level of employment in the economy, L_t , via

$$r_t^p = R(L_t), \quad R'(L_t) \leq 0.$$

We could alternatively have explicitly microfounded this mechanism,⁴³ but doing so would significantly complicate the presentation without, we believe, adding much insight to the properties we are focused on.

In setting the policy rate i_t , we assume that the central bank follows a slightly modified Taylor-type rule where the interest rate reacts to expected inflation and to expected labor market conditions⁴⁴ via

$$i_t = \phi_\pi \mathbb{E}_t[\pi_{t+1}] + \phi_\ell \mathbb{E}_t[G(L_{t+1})], \quad \phi_\pi, \phi_\ell, G' > 0. \quad (18)$$

Choosing to have the central bank react to expected developments in this economy—as opposed to reacting only to contemporaneous variables—will allow us to simplify the model

⁴³For example, Gourio [2013] generates a countercyclical risk premium in a model with disaster risk. See also the discussion on “risk-averse recessions” in Cochrane [2016].

⁴⁴Note that the level of employment in this model can also be viewed as a measure of the output gap. Hence, wherever we refer to employment we could alternatively have referred to the output gap.

considerably without losing the flavor of standard New Keynesian models. In particular, since our aim is to use this model to explain the dynamics of a set of real variables, and since we would like to stay away from the general debate over whether or not a Phillips curve properly calibrated to micro-economic data can explain inflation movements, we will assume that the central bank sets $\phi_\pi = 1$. This parameterization implies that the central bank varies its nominal policy rate such that the expected *real* policy rate increases with expected labor market activity. The attractiveness of this assumption is that it will allow the model to take a block-recursive structure, wherein the aggregate quantity variables are determined independently of the inflation rate. The inflation rate remains a function of these quantity variables, but it will not affect them. The degree of price stickiness then governs the extent to which inflation varies with economic activity. By varying the degree of price stickiness, inflation in the model can become arbitrarily volatile or arbitrarily stable, but this margin does not affect the other variables. This convenient assumption will allow us to focus on the implications of the model for real variables while allowing us to sidestep debates about inflation and about the degree of price stickiness that is necessary to match the data. All we need to assume regarding price-setting for our purposes is that there is *some* stickiness, so that production of consumption services remains demand-determined.

Intermediate firms in our environment are assumed to be subject to nominal rigidities. However, as we have assumed that $\phi_\pi = 1$, we can bypass the need to be explicit about the intermediate firm's pricing problem and its implications for inflation, as the realizations of inflation do not feed back into the determination of the quantity variables that are of interest to us. Specifically, substituting (E.27) (with $\phi_\pi = 1$) and (16) into the Euler Equation (E.21), and assuming the shocks are small enough that the approximation $\mathbb{E}_t [\pi_{t+1}] \approx \pi_{t+1}$ holds, we obtain⁴⁵

$$U'(X_t + F(\theta_t L_t) - \gamma(X_{t-1} + F(\theta_{t-1} L_{t-1}))) = \beta_t \mathbb{E}_t \{U'(X_{t+1} + F(\theta_{t+1} L_{t+1}) - \gamma(X_t + F(\theta_t L_t))) [1 + \phi_\ell G(L_{t+1}) + R(L_t)]\}. \quad (19)$$

Equation (E.30) and the accumulation Equation (E.26) together form a two-variable system

⁴⁵We adopt the common approximation used in these model that the aggregate production of consumption services can be represented by $C_t = X_t + F(\theta_t L_t)$; that is, we dis-regard the effect of price dispersion among intermediate firms in affecting aggregate production.

in X and L , and in particular notice that inflation π does not appear anywhere. Since C and D can be obtained from X and L , it then follows that all quantity variables are determined independently from the inflation rate.

The attractive feature of Equation (E.30) is that it is in a form very similar to Equation (11). To see this even more clearly, let us posit the following simple functional forms. Let us assume that the utility function is $U(z) = -\exp\{-\frac{z}{\omega}\}$, $\omega > 0$, the production of new goods is given by $F(\Theta L) = \ln(\Theta L)$, the interest rate policy function is $G(L) = \ln(L)$, and Θ_t is constant. Under these conditions, (E.30) can be rewritten⁴⁶

$$\ell_t = \mu_t - \hat{\alpha}_1 X_t + \hat{\alpha}_2 \ell_{t-1} + \hat{\alpha}_3 \mathbb{E}_t[\ell_{t+1}] + \hat{F}(\ell_t) \quad (20)$$

where $\ell_t = \ln L_t$, $\mu_t = \frac{1}{\kappa} [\psi (1 - \frac{\gamma}{1-\delta}) \ln \Theta - \omega \ln \beta_t]$, $\hat{\alpha}_1 = \frac{\delta}{\kappa} (1 - \frac{\gamma}{1-\delta})$, $\hat{\alpha}_2 = \frac{\gamma}{\kappa} (1 - \frac{\psi}{1-\delta})$, $\hat{\alpha}_3 = \frac{1-\omega\phi_\ell}{\kappa}$ and $\hat{F}(\ell_t) = -\frac{\omega}{\kappa} \tilde{R}(\ell_t)$, with $\kappa = 1 + \gamma - \psi > 0$.

Ignoring the stochastic component of μ_t , Equation (E.32) has a form identical to that of Equation (11) as long as $\omega\phi_\ell < 1$. In effect, one can view this Euler equation as giving a structural interpretation to the behavior explored in the previous section. Since this equation is of the form we previously analysed, we know that it can potentially generate endogenous cyclical behavior in the form of limit cycles. For this to arise, we would need—among other conditions—that \tilde{R}' be sufficiently negative near the steady state (but remaining greater than -1). In other words, for limit cycles to arise, we need the counter-cyclical sensitivity of the risk premium to be sufficiently strong. If this is the case, then this model can in principle exhibit a type of accumulation-liquidation cycle in which the model's steady state would be unstable and the forward-looking agents in the model would find it optimal to bunch their purchases in boom times. That is, during booms, agents consume and accumulate durables, which creates a strong labor market so that lending to household is less risky, which rationalizes a low risk premium on loans and high purchases by households. Eventually, however, households accumulate enough so that demand for new goods falls, causing a bust in employment, and the risk premium on loans to increase. The bust acts as a liquidation period, which comes to an end once the stock of durables has become sufficiently depleted that the purchasing of more new goods becomes desirable once again. While this narrative has a similar flavor to

⁴⁶In this derivation we have used the approximations $\mathbb{E}[\ln(z)] \approx \ln(\mathbb{E}[z])$ and $\ln(1+z) \approx z$.

the one often associated with indeterminacy and multiple equilibria, we wish to re-emphasize that in our framework such boom-bust cycles do not require expectational effects to be so strong as to create self-fulfilling expectations.

From Equation (E.32) we see that current employment—and thus consumption—depends positively on both past and future employment. These effects reflect those commonly found in consumption Euler equations. In addition to these effects, Equation (E.32) also includes the effect of the state variable X_t , whereby higher levels of X_t decrease employment, since less employment is required to achieve the same consumption level. It is through this channel that the model captures liquidation effects. Finally, employment exhibits a complementarity structure through the risk premium: increased employment lowers the risk premium on borrowing, which favors more consumption and the purchase of new goods, which in turn further stimulates employment. As long as $-\frac{\omega}{\kappa}\tilde{R}'$ is smaller than one, such a feedback effect will not create static multiple equilibria, but as we have seen, it may nonetheless cause a limit cycle to emerge.

3.2 Spectral Properties of the Interest Rates and Interest Spreads

Our extended New Keynesian model jointly determines employment, the policy rate, and the risk premium. One of the implications of this model is that interest rates and interest rate spreads should share some cyclical properties with that of employment. For example, if the dynamics of employment are determined by this model and these dynamics induce a peak in the spectral density of employment, then the policy rate set by the central bank and the interest rates faced by private agents should also exhibit similar peaks in their spectral densities. Hence, before bringing this model to the data in a formal manner, it is desirable to first examine the spectral properties of interest rates and the risk premium to see if they indeed exhibit spectral peaks somewhat similar to those observed for employment.

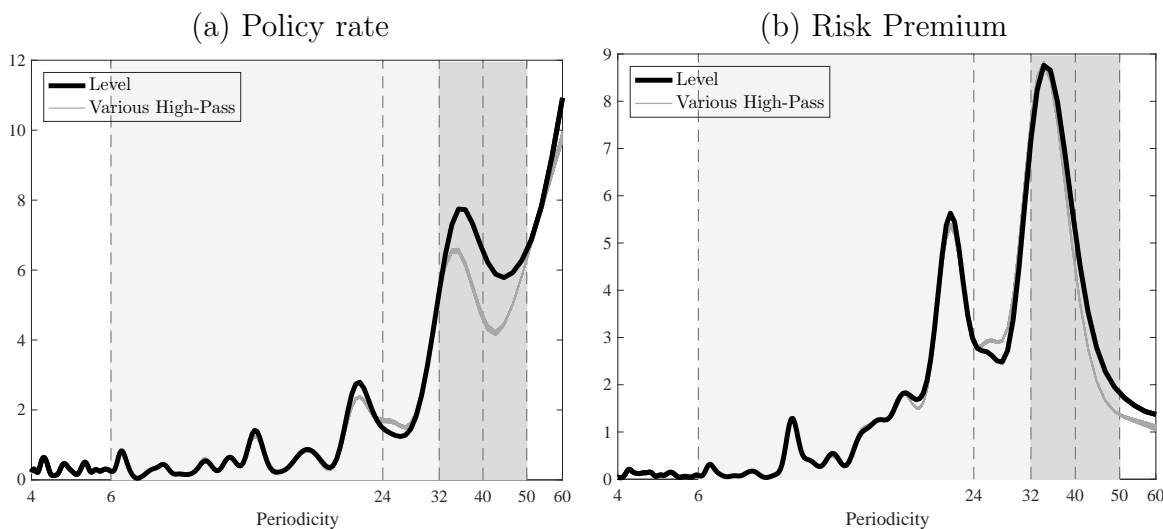
In Panels (a) and (b) of Figure 11, we report an estimate of the spectral density of the real return on U.S. federal funds⁴⁷ and that of the spread between the federal funds rate and BBA bonds, which we will take as our measure of the risk premium on borrowing.⁴⁸ In each

⁴⁷This corresponds to the interest on three-month federal funds minus realized inflation over the holding period.

⁴⁸See Appendix A for a precise description of the data.

case, we report the spectral density for the level of the series, as well as spectral densities for transformed series where we have removed very low-frequency movements using a high-pass filter as in Section 1. The interesting aspect to note in both of these figures is that these series appear to exhibit a spectral peak in the vicinity of 40 quarters, regardless again of whether or not we first remove very low-frequency movements. The peak around 40 quarters is most noticeable in our spread series, but is also visible in the behavior of the policy rate. These similarities are important since the simplicity of our extended New Keynesian model implies that all three series should have somewhat similar spectral properties. The model also implies that the policy interest rates should co-move positively with employment, while the interest spread should co-move negatively, which is what we observe in the data (see Figure 14). While these observations certainly do not validate the model on their own, they represent preconditions for such a model to be potentially relevant.

Figure 11: Spectral Densities of the Policy Rate and the Risk Premium



Notes: This figure shows an estimate of the spectral density of the real return on U.S. federal funds (policy rate) (Panel (a)) and the spread between the federal funds rate and BBA bonds (Panel (b)), for levels (black lines) and for 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (grey lines). A high-pass (P) filter removes all fluctuations of period greater than P .

3.3 Estimation

To bring our model to the data, it remains only to specify the process for μ_t and the functional form for \tilde{R} . For the former, we assume simply that μ_t follows a stationary AR(1) process⁴⁹ $\mu_t = \rho\mu_{t-1} + \epsilon_t$, where ϵ_t is a Gaussian white noise with variance σ^2 . For the latter, since we have no clear prior about the form of this relationship, we approximate it simply by a third-order polynomial.

We estimate this model using the indirect inference method of Gourieroux, Monfort, and Renault [1993], where for each parameter set the model is solved by a third-order perturbation method.⁵⁰ The solution and estimation is somewhat involved, as it allows for the possibility of limit cycles in a stochastic model with forward-looking agents. To our knowledge, such an exercise is novel. The parameters of the model are chosen so as to minimize its distance to a set of features of the data we have already emphasized. We focus on three sets of observations. The first set corresponds to the spectral density of hours worked per capita (as shown in Panel (b) of Figure 1). The second set corresponds to the spectral density of the risk premium (as shown in Panel (b) of Figure 11). In these first two cases, we aim to fit the point estimates of the spectral densities (using the non-detrended data) at periodicities between 2 and 50 quarters. The last set of observations is a set of five additional moments of the data: the correlation between hours and the risk premium, as well as the skewness and kurtosis of each of these two variables. Each of the data moments in this last set are obtained after first detrending the data series using a high-pass filter that removes fluctuations longer than 50 quarters. This is in line with our objective of using the current model to explain macroeconomic fluctuations arising at periodicities ranging from 2-50 quarters. Our estimation does not directly aim to match the real policy interest rate. Instead, we will use this series to provide a non-targeted dimension on which to judge the model's overall fit.

We estimate nine parameters: γ , ω , ψ , ϕ_ℓ , ρ , σ , and the coefficients on the first-, second-, and third-order polynomial terms for the function $\tilde{R}(\cdot)$, which we denote \tilde{R}_1 , \tilde{R}_2 and \tilde{R}_3 .

⁴⁹To avoid cluttering the notation, we henceforth omit all constant terms, with all variables now denoting the corresponding deviations from the non-stochastic steady state.

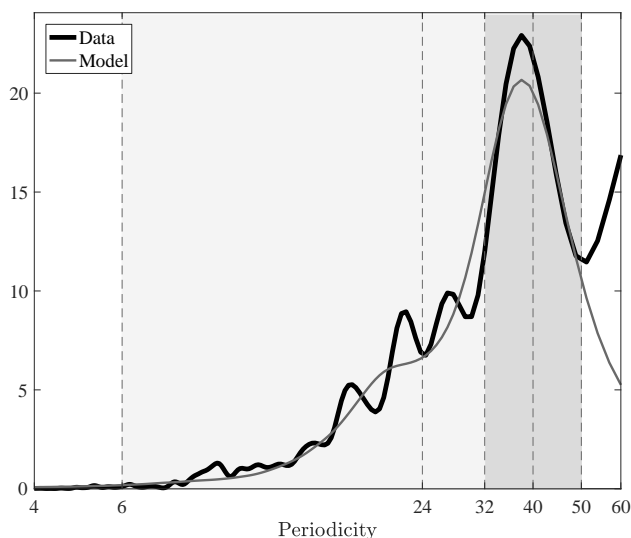
⁵⁰Details of the solution and estimation are given in Section F of the Online Appendix.

The weighting matrix used in estimation is a simple identity matrix. The depreciation rate is not estimated but set to $\delta = 0.05$ in order to match the average depreciation of houses and durable goods.

Figures 12-14 illustrate the fit of the estimated model along the targeted dimensions, as well as the fit of the (non-targeted) real policy rate spectral density. As can be seen from the figures, the model is capable of replicating many of the salient features of all three of our variables of interest, although with some exceptions. For example, as can be seen in Panel (a) of Figure 13, the model does not capture the peak in the spectral density of the risk premium observed around a periodicity of 20 quarters. As can be seen in Panel (b) of that figure, the model also over-predicts the over all variance of the policy rate at the targeted frequencies. Notwithstanding these exceptions, the model—which uses only eight parameters to attempt to fit a large set of features of the data—over all provides a reasonable fit. It is especially interesting to note from Figure 12 how well our parsimonious model is able to fit the spectral density of hours worked between 2 and 50 quarters. It does not capture all the bumps and wiggles, but it fits the overall pattern very nicely, especially the marked peak around 40 quarters. In Figure 14 we report the fit of the model with respect to the final set of moments targeted by the estimation. As can be seen in the figure, the model also does well at capturing these additional properties of the data, with the exception of the risk premium skewness, which is of the wrong sign in the model.

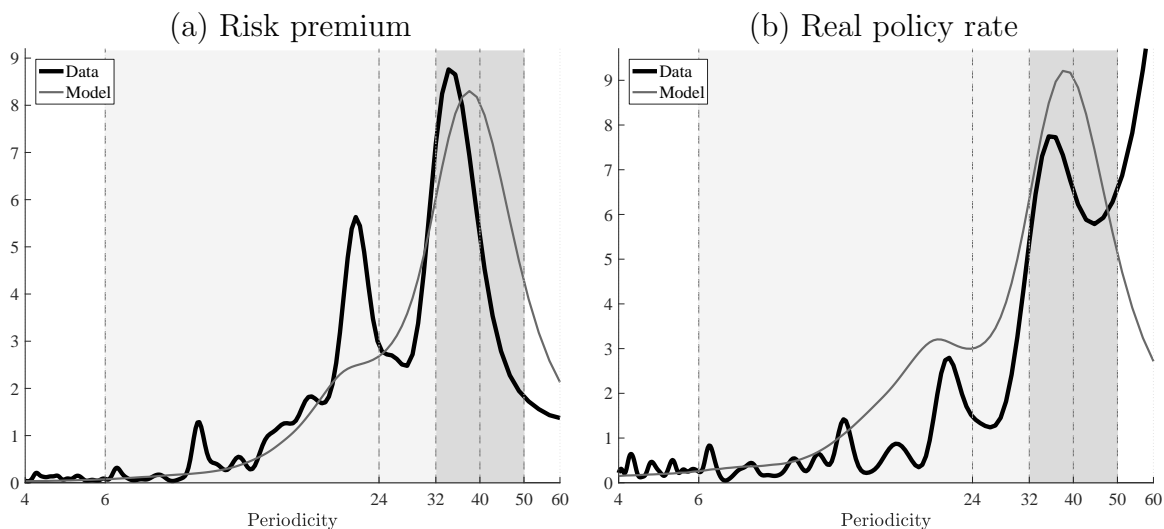
The parameter estimates for the model are presented in Table 1 (standard errors in parentheses) and the implied eigenvalues evaluated at the steady state are given in Table 2. The estimates imply a habit parameter of $\gamma = 0.53$, which is in line with the values commonly found in the literature. We also find that monetary authorities react reasonably aggressively to (expected) economic activity, with an elasticity of the policy rate to employment of $\phi_\ell = 0.19$. Interestingly, the model implies that the first-order effect of economic activity on the risk premium (\tilde{R}_1) is of the same order of magnitude as—but of the opposite sign to—that on the policy rate. The most interesting finding regarding the parameter estimates is that the shock process is estimated to have an auto-correlation that is effectively zero. This implies that the model’s dynamics are almost completely due to endogenous forces. The one parameter that may be considered somewhat large relative to the literature is our estimate

Figure 12: Fit of Hours Spectral Density




Notes: This figure compares the estimate of the spectral density of U.S. Non Farm Business hours per capita with the one obtained from our estimated model. Both series have been high-pass (50) filtered, in order to remove all fluctuations of period greater than 50 quarters.

Figure 13: Fit of Risk Premium and Real Policy Rate Spectral Densities



Notes: This figure compares the estimate of the spectral density of U.S. risk premium and real policy rate with the one obtained from our estimated model. Both series have been high-pass (50) filtered, in order to remove all fluctuations of period greater than 50 quarters.

Figure 14: Fit of Moments



Notes: This figure compares various moments estimated using U.S. data with the ones obtained from our estimated model. All series have been high-pass (50) filtered, in order to remove all fluctuations of period greater than 50 quarters.

of ω , which gives the elasticity of consumption with respect to the effective interest rate. Our estimates suggest that this elasticity is between 4 and 5, indicating a very high response of consumption expenditures. Note that, in our setup, the effective interest rate faced by households reflects not only the policy interest rate, but also a risk premium. Although the sum of these two parts moves only mildly with economic activity (see below), their combined effect on consumption is nonetheless important since ω is relatively high.

Table 1: Estimated Parameter Values

Habit persistence	γ	0.5335
Durables proportion	ψ	0.4000
Taylor rule parameter on output gap	ϕ_ℓ	0.1903
CARA utility parameter	ω	4.6248
Persistence of shocks	ρ	-0.0000
S.D. of shock innovations	σ	0.8530
Spread function parameter	\tilde{R}_1	-0.1624
Spread function parameter	\tilde{R}_2	0.00076
Spread function parameter	\tilde{R}_3	0.00027

Notes: Table displays the estimated parameters of the model.

To illustrate the deterministic mechanisms implied by the parameter estimates, Figure

Table 2: Estimated Eigenvalues at the Steady State

$Re(\lambda)$	$Im(\lambda)$	$ \lambda $
1.1030	0.2166	1.1241
1.1030	-0.2166	1.1241
1.9393	0	1.9393

Notes: Eigenvalues are computed using a first-order approximation of the model around the non-stochastic steady state.

15 reports results obtained when feeding in a constant value of $\mu_t = 0$ for the exogenous process.⁵¹ Panel (a) of Figure 15 plots a simulated 270-quarter sample⁵² of log-hours generated from the deterministic version of the model. Two key properties should be noted. First, the parameters obtained from fitting the model to the data produce endogenous cyclical behavior, with cycles of a reasonable length (around 38 quarters). This is consistent with Table 2, which indicates that the steady state is unstable and features two complex eigenvalues. 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 a limit cycle framework to understand business cycles. However, as this exercise demonstrates, reasonable-length endogenous cycles can be generated in our framework relatively easily, precisely because the model possesses the three key features we highlighted in the previous section: diminishing returns to capital accumulation, sluggishness, and local demand complementarities. Second, notwithstanding the reasonable cycle length, it is clear when comparing the simulated data in Panel (a) of Figure 15 to actual economic data that the fluctuations in the deterministic version of the model are far too regular. These properties of the deterministic version of our estimated model—i.e., a highly regular 38-quarter cycle—can also be seen clearly in the frequency domain. Panel (b) of Figure 15 plots the spectral density of hours for the deterministic model (grey line), along with the spectral density for the data (black line) for comparison.⁵³ This spectral density exhibits an extremely

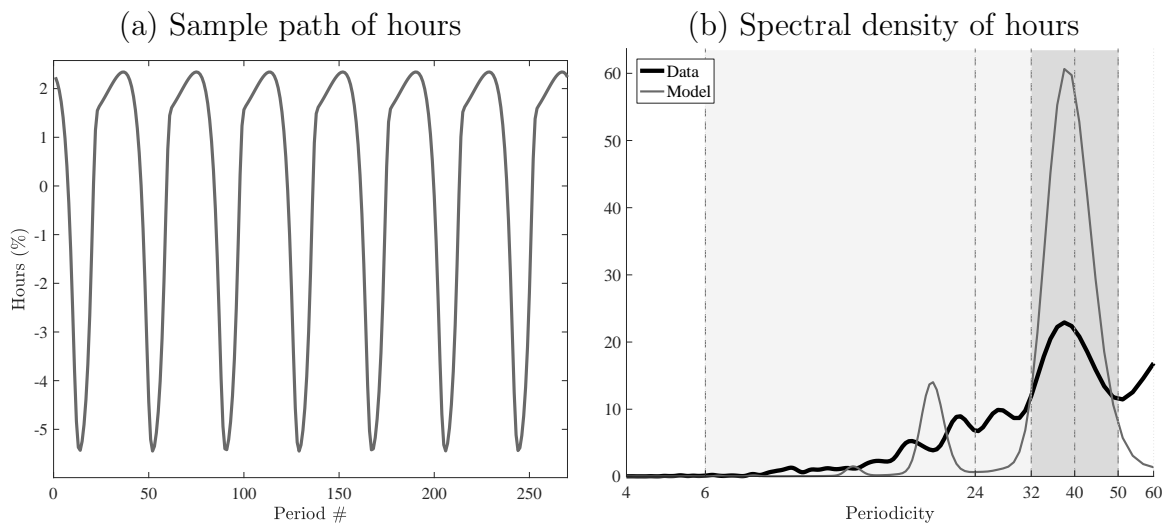
⁵¹Note that the model was solved (and simulated) using the estimated parameters presented in Table 1, and in particular we did not first re-solve the model with $\sigma = 0$. Thus, agents in this deterministic simulation implicitly behave as though they live in the stochastic world. As a result, any differences between the deterministic and stochastic results in this section are due exclusively to differences in the realized sequence of shocks, rather than differences in, say, agents' beliefs about the underlying data-generating process.

⁵²This is equal to the length of the sample period of the data.

⁵³Note that the model was not re-estimated after shutting down the exogenous shock. As such, there

large peak—characteristic of a highly regular cycle—at the 38-quarter periodicity,⁵⁴ while the spectral density of the data is much flatter.

Figure 15: Estimated Model When Shocks Are Turned Off




Notes: This figure corresponds to the counterfactual simulation of the estimated model when shocks are turned off, and initial conditions place the system on the limit cycle. Panel (a) shows the evolution of hours along the limit cycle. Panel (b) compares the spectral density of hours in the data (black line) and in the model (grey line).

Re-introducing the estimated shocks into the model, we see a markedly different picture in both the time and frequency domains. Figure 16 plots an arbitrary 270-quarter sample of log-hours generated from the full stochastic model. While clear cyclical patterns are evident in the figure, it is immediately obvious that the inclusion of shocks results in fluctuations that are significantly less regular than those generated in the deterministic model, appearing qualitatively quite similar to the fluctuations found in actual data. This is confirmed by the hours spectral density (Figure 12), which matches the data quite well. In particular, the spectral density of the stochastic model includes a distinct peak close to 40 quarters, suggesting some degree of regularity at that periodicity, but without the exaggerated peak

may be alternative parameterizations of the deterministic model that are better able to match the spectral density in the data.

⁵⁴The deterministic model spectral density also contains smaller peaks at integer multiples of the frequency of the main cycle (i.e., at around $19 = 38/2$ quarters, $12.67 = 38/3$ quarters, etc.). Such secondary peaks arise when the data exhibits a regular but not perfectly sinusoidal cycle, as is clearly the case in Panel (a) of the figure.

Figure 16: Sample Path of Hours in the Estimated Model with Shocks



Notes: This figure shows one particular 270 quarters simulation of the estimated model.

observed at this point in the deterministic model.


It should be emphasized that the exogenous shock process in this model primarily acts to accelerate and decelerate the endogenous cyclical dynamics, causing significant random fluctuations in the length of the cycle, while only modestly affecting its amplitude. In fact, somewhat counter-intuitively, when the shock is shut down (as in Figure 15), the variance of log-hours actually *increases* relative to the full stochastic case (the variance of log-hours in the stochastic case is 6.24, while in the deterministic case it is 7.35).⁵⁵ Moreover, it is worth emphasizing once again that the estimated shocks in the model are essentially i.i.d., indicating that, in an environment featuring limit cycles, macroeconomic fluctuations can be explained without the need to rely on persistent exogenous disturbances.

The role of complementarities in the model, however, is extremely important. To illustrate this, in Figure 17 we plot the spectral density implied by the model for hours worked

⁵⁵Note that, since we have simply fed a constant sequence $\mu_t = 0$ of shocks into our model without first resolving it under the assumption that $\sigma = 0$, this phenomenon is not due in any way to rational-expectations effects. The fact that a fall in shock volatility can lead to a rise in the volatility of endogenous variables in a limit cycle model was pointed out in Beaudry, Galizia, and Portier [2016]. Roughly speaking, because of the nonlinear forces at play, shocks that push the system “inside” the limit cycle have more persistent effects than those that push it “outside”. For relatively small shocks, this leads to a decrease in outcome volatility when the shock volatility increases. See Beaudry, Galizia, and Portier [2016] for a more detailed discussion of these mechanisms in the context of an estimated reduced-form univariate equation.

after shutting down the endogenous risk premium (i.e, setting $\tilde{R}_i = 0$ for $i = 1, 2, 3$, but keeping all other parameters at their estimated levels). As can be seen, in this case the model no longer mimics the properties of the data, and in fact the spectral density is very close to zero. That is, without the complementarities to amplify them, the small i.i.d. disturbances are sufficient to generate only a tiny amount of volatility in hours.

Figure 17: Spectral Density of Hours with Complementarities Turned Off




Notes: This figure corresponds to the counterfactual simulation of the estimated model when the risk premium is forced to be constant, which correspond to imposing $\tilde{R}_1 = \tilde{R}_2 = \tilde{R}_3 = 0$. The other parameters are not re-estimated. It compares the spectral density of hours in the data (black line) and in the model (the flat grey line that is almost confounded with the x-axis).

As we have already emphasized, nonlinearities are crucial in order to have a steady state that is locally unstable without also having explosive dynamics, a combination which is a pre-condition for limit cycles to emerge. To get a sense of the estimated degree of nonlinearity in our model, we plot in Figure 18 the effective real interest rate function $r(\ell) = \phi_\ell \ell + \sum_{i=1}^3 \tilde{R}_i \ell^i$. The nonlinearities coming from the nonlinear response of the risk premium to economic activity are apparent in the figure, though they are mild. Near the steady state, a one-percent rise in hours is associated with around a 65-basis-point fall in the risk premium. As we move away from the steady state, this sensitivity fades. For example, at the deepest trough and highest peak recorded in our data sample, a one-percent rise in hours would be associated with 57- and 53-basis-point falls in the risk premium (12% and

18% less sensitivity than at the steady state), respectively. Such differences in the strength of the complementarities may appear small at first pass, but they are sufficient to have large quantitative impacts on the dynamics of the system.

Figure 18: Sensitivity of the Real Interest Rate Faced by the Households to Economic Activity



Notes: This figure plots the real interest rate function $r(\ell) = \phi_\ell \ell + \sum_{i=1}^3 \tilde{R}_i \ell^i$, with estimated coefficients and over the range of hours deviations that we observe in the data. In the figure, r is annualized.

3.4 Robustness

In this section we report results from some robustness checks on the baseline estimation of our New Keynesian limit cycle model. As we noted in Section 1, the data spectral densities of several key stationary macroeconomic variables exhibit peaks around 40 quarters, declining from there before reaching a local minimum at around 50 quarters, then increasing again beyond that point. This was our motivation for choosing the range of periodicities between 2 and 50 quarters to focus on in our estimation. We turn now to evaluating how robust our results are to focusing on some alternative ranges. In particular, since the lower end of the standard business cycle range in the literature is 6 quarters, we consider what happens if we estimate the parameters of the model restricting attention to periodicities from 6 to 50 quarters. We also consider what happens if we restrict attention further to the range beyond the traditional upper bound of 32 quarters (i.e., restricting to periodicities between 32 and


50 quarters). Finally, we repeat these same exercises, but using 60 quarters as the upper bound.

Figure 19 reports the results. Column (a) of Figure 19 plots the hours, risk premium, and real policy rate spectral densities for period ranges of the form $(x,50)$, $x = 2, 6, 32$, while column (b) does the same for period ranges of the form $(x,60)$.⁵⁶ The corresponding data spectral densities are also plotted (solid black lines) for comparison. Several things emerge from the figure. First, from column (a) we see that increasing the lower bound from 2 (solid grey line) to 6 (dashed grey line) quarters has a minimal effect on the parameter estimates, and this translates into a correspondingly minimal effect on the model fit. Second, when we move to a lower bound of 32 quarters (dash-dot black line) in column (a), we unsurprisingly see a much better fit of the hours and policy rate spectral densities in this range (the model spectral density for hours in particular lies almost directly on top of the data spectral density) with roughly the same fit of the risk premium, while the estimates lead to substantial under-prediction of the volatility of hours and the risk premium at the traditional business cycle periodicities (6 to 32 quarters). Finally, moving the upper bound from 50 to 60 quarters (column (b)) appears to have only a modest effect on the fit of our target series (hours and the risk premium), while the model no longer over-predicts the over all volatility of the real policy rate.

Table 3 reports, for each of the estimations, the eigenvalues associated with the steady state of the system. In all cases, the eigenvalues retain the same basic configuration as in the baseline estimation: a pair of complex conjugate eigenvalues outside the unit circle, and a real positive eigenvalue greater than one. That is, in all cases our estimations produce parameters that generate limit cycles.

⁵⁶Note that the (2,50) plots simply reproduce Figures 12 and 13.

Figure 19: Robustness: Estimated and Data Spectral Densities for Various Estimations



This figure shows estimated and actual spectral densities of hours, risk premium and policy rate when data are filtered with various band-pass filters before estimation. These band-pass filters are either $(x,50)$ or $(x,60)$, with $x = 2, 6$ or 32 .

Table 3: Robustness: Estimated Eigenvalues at the Steady State for Various Estimations

Band-Pass filter	$\lambda_1, \bar{\lambda}_1 (\lambda_1)$	λ_2
(2,50)	$1.1031 \pm 0.2166i$ (1.1241)	1.9386
(6,50)	$1.0993 \pm 0.2069i$ (1.1185)	2.0745
(32,50)	$1.0719 \pm 0.1881i$ (1.0882)	2.3744
(2,60)	$1.0621 \pm 0.2140i$ (1.0835)	2.2491
(6,60)	$1.0745 \pm 0.2071i$ (1.0943)	2.2078
(32,60)	$1.0771 \pm 0.2021i$ (1.0959)	2.2206

Notes: Each line of this table corresponds to a different estimation of the model. In each estimation, we minimize the distance between the model and the data (a) spectral densities of hours and risk premium and (b) set of five moments. In each estimation, the data are filtered with a different band-pass filter.

Conclusion

Why do market economies repeatedly go through periods of boom and bust? There are at least two broad classes of explanations. On the one hand, it could be that the economy is inherently stable and that observed booms and busts are entirely due to outside disturbances and the resulting adjustment processes. On the other hand, it could be that the economy is locally unstable, in that it does not contain forces that tend to push it towards a stable resting position. Instead, the economy's internal forces keep it in a constant state of motion, never converging to a single point nor exploding. According to this explanation, the economy is still buffeted by outside disturbances, but these shocks are neither the sole nor primary drivers of the business cycle. Instead, they mainly serve to make fluctuations irregular and hard to predict.

In this paper our aim has been to examine the plausibility and relevance of this second, less mainstream view of macroeconomic fluctuations. To this end, we have made three contributions. First, we have documented a set of patterns in the data that should be present if limit cycle forces play a relevant role in the economy. In particular, we have shown that several macroeconomic variables exhibit distinct peaks in the spectral density, and that these peaks occur at roughly the 40-quarter periodicity. While this periodicity is somewhat long relative to the standard definition of the business cycle, we argue that cycles of this length should be thought of as an integral part of business cycle phenomena, as opposed to reflecting some more medium-run behavior. It is important to note that such observations do not necessarily imply the presence of limit cycles, but rather provide suggestive evidence for their existence. Second, we present a simple class of economic model capable of producing limit cycles. Our goal in this section was to clarify how readily limit cycles can arise in environments commonly studied by macroeconomists. The class of models we consider can be seen as a dynamic extension of the models studied by Cooper and John [1988] in their seminal article, which emphasized how forces of strategic complementarity may help to explain macroeconomic behavior. As part of our analysis, we attempted to clarify how the concept of limit cycles is distinct from the concepts of indeterminacy and multiple equilibria (though the two can sometimes coincide). Finally, in the last section of the paper, we

extended the canonical three-equation New Keynesian model in a manner that incorporates the elements emphasized in the second section. The model is aimed at capturing a simple narrative of endogenous boom-bust cycles, wherein the accumulation of capital interacts with financial conditions to sustain fluctuations. When we explore the model empirically, we find that the data favor parameter estimates that support the presence of limit cycle forces; that is, the data favor being explained through limit cycles forces instead of through more conventional mechanisms. We also find that shocks remain important in explaining macroeconomic fluctuations, but interestingly the estimated shock process is essentially i.i.d., instead of the highly persistent shock processes that are often required in more mainstream models. It should be emphasized that model we present in the last section is highly stylized and is certainly not a full description of all macroeconomic behavior. Nonetheless, it is interesting that such a parsimonious and simple model with an i.i.d. shock can capture many features of macroeconomic data once limit cycles are allowed to be part of the endogenous mechanisms. This, we believe, nicely illustrates the potential empirical relevance of such a framework.

References

- ALESSIO, S. M. (2016): *Digital Signal Processing and Spectral Analysis for Scientists*. Springer, Cham.
- 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 (2014): “Reconciling Hayek’s and Keynes’ Views of Recessions,” Working Paper 20101 (NBER) and 9966 (CEPR), forthcoming in *The Review of Economic Studies*.
- (2015): “Reviving the Limit Cycle View of Macroeconomic Fluctuations,” Working Paper 21241 (NBER) and 10645 (CEPR).
- (2016): “Is the Macroeconomy Locally Unstable and Why Should We Care?,” Working Paper 22275 (NBER) and 11305 (CEPR), forthcoming in *NBER Macroeconomic Annual 2016*.
- BENHABIB, J., AND R. E. A. FARMER (1994): “Indeterminacy and Increasing Returns,” *Journal of Economic Theory*, 63(1), 19–41.
- BENHABIB, J., AND K. NISHIMURA (1979): “The Hopf Bifurcation and the Existence and Stability of Closed Orbits in Multisector Models of Optimal Economic Growth,” *Journal of Economic Theory*, 21(3), 421–444.
- (1985): “Competitive Equilibrium Cycles,” *Journal of Economic Theory*, 35(2), 284–306.
- BOLDRIN, M., 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.
- COCHRANE, J. (2005): “Time Series for Macroeconomics and Finance,” Unpublished manuscript.
- COCHRANE, J. H. (2016): “Macro-Finance,” NBER Working Papers 22485, National Bureau of Economic Research, Inc.

- 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.
- D’AGOSTINO, R., AND E. S. PEARSON (1973): “Tests for Departure from Normality. Empirical Results for the Distributions of b_2 and $\sqrt{b_1}$,” *Biometrika*, 60(3), 613–622.
- 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.
- FERNALD, J. (2014): “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Federal Reserve Bank of San Francisco Working Paper 2012-19.
- GOODWIN, R. (1951): “The Nonlinear Accelerator and the Persistence of Business Cycles,” *Econometrica*, 19(1), 1–17.
- GOURIEROUX, C., A. MONFORT, AND E. RENAULT (1993): “Indirect Inference,” *Journal of Applied Econometrics*, 8, S85–S118.
- GOURIO, F. (2013): “Credit Risk and Disaster Risk,” *American Economic Journal: Macroeconomics*, 5(3), 1–34.
- GRANDMONT, J.-M. (1985): “On Endogenous Competitive Business Cycles,” *Econometrica*, 53(5), 995–1045.
- GRANGER, C. W. J. (1966): “The Typical Spectral Shape of an Economic Variable,” *Econometrica*, 34(1), 150–161.
- GU, C., F. MATTESINI, C. MONNET, AND R. WRIGHT (2013): “Endogenous Credit Cycles,” *Journal of Political Economy*, 121(5), 940 – 965.
- HICKS, J. (1950): *A Contribution to the Theory of the Trade Cycle*. Clarendon Press, Oxford.

- JARQUE, C., AND A. BERA (1987): “A Test for Normality of Observations and Regression Residuals,” *International Statistical Review / Revue Internationale de Statistique*, 55(2), 163–172.
- 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.
- KIYOTAKI, N., AND J. MOORE (1997): “Credit Cycles,” *Journal of Political Economy*, 105(2), 211–48.
- KUZNETSOV, Y. A. (1998): *Elements of Applied Bifurcation Theory*, vol. 112 of *Applied Mathematical Sciences*. Springer, New-York, 2 edn.
- LE CORBEILLER, P. (1933): “Les systemes autoentretenus et les oscillations de relaxation,” *Econometrica*, 1(3), 328–332.
- MARCZAK, M., AND V. GÓMEZ (2012): “SPECTRAN, a Set of Matlab Programs for Spectral Analysis,” Discussion paper 60, Universität Hohenheim.
- 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.
- MYERSON, R. B. (2012): “A Model of Moral-Hazard Credit Cycles,” *Journal of Political Economy*, 120(5), 847–878.
- PRIESTLEY, M. (1981): *Spectral Analysis of Time Series*, vol. 1 Univariate Series. Academic Press, London.
- SARGENT, T. J. (1987): *Macroeconomic Theory*. Emerald Group Publishing Limited, second edn.
- SCHEINKMAN, J. (1990): “Nonlinearities in Economic Dynamics,” *The Economic Journal*, 100(400), 33–48, Conference Papers.

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

- U.S. Population: Total Population: All Ages including Armed Forces Overseas, obtained from the FRED database (POP) from 1952Q1 to 2015Q2. Quarters from 1947Q1 to 1952Q1 are obtained from linear interpolation of the annual series of National Population obtained from U.S. Census, where the levels have been adjusted so that the two series match in 1952Q1.
- U.S. Total GDP is obtained from the Bureau of Economic Analysis National Income and Product Accounts. Real quantities are computed as nominal quantities (Table 1.1.5) over prices (Table 1.1.4.). Sample is 1947Q1-2015Q2.
- Non-Farm Business Hours, Total Hours and unemployment rate (16 years and over) are obtained from the Bureau of Labor Statistics. Sample is 1947Q1-2015Q2 (1948Q1-2015Q2 for total hours), and we do not use the observations of 2015.
- Capacity utilization: Manufacturing (SIC), Percent of Capacity, Quarterly, Seasonally Adjusted, obtained from the FRED database, (CUMFNS). Sample is 1948Q1-2015Q2.
- TFP: utilization-Adjusted quarterly-TFP series for the U.S. Business Sector, produced by John Fernald, series ID: dtfp_util, 1947Q1-2015Q2.
- Spread: Moody’s Seasoned Baa Corporate Bond Minus Federal Funds Rate, quarterly average, obtained from the FRED database, (BAAFFM). Sample is 1954Q3-2015Q2.
- Policy rate : computed from the FRED database as Moody’s Seasoned Baa Corporate Bond Yield (BAAFFM) minus the spread (BAAFFM). Sample is 1954Q3-2015Q2. The real policy rate is obtained by subtracting realized inflation, computed using the “Consumer Price Index for All Urban Consumers: All Items” (CPIAUCSL from the FRED database).

B Proofs of Propositions

B.1 Proposition 1

The two eigenvalues of matrix M_L are the solution the equation

$$Q(\lambda) = \lambda^2 - T\lambda + D = 0, \tag{B.1}$$

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

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

where

$$T = \alpha_2 - \alpha_1 + (1 - \delta) \tag{B.3}$$

and

$$D = \alpha_2(1 - \delta). \tag{B.4}$$

From (B.4), 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 (B.3), 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

B.2 Proposition 2

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

$$T = \frac{\alpha_2 - \alpha_1}{1 - F'(I^s)} + (1 - \delta) \quad (\text{B.5})$$

and

$$D = \frac{\alpha_2(1 - \delta)}{1 - F'(I^s)}. \quad (\text{B.6})$$

From equations (B.5) and (B.6), we have the following relationship between T and D :

$$D = \frac{\alpha_2(1 - \delta)}{\alpha_2 - \alpha_1} T - \frac{\alpha_2(1 - \delta)^2}{\alpha_2 - \alpha_1}. \quad (\text{B.7})$$

Therefore, when $F'(I^s)$ varies, T and D move along the line (B.7) 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 6.

When $F'(I^s) \rightarrow -\infty$, we have $D \rightarrow 0$ and $T \rightarrow 1 - \delta$, which corresponds to point E_1 in Figure 6. 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 $-\infty$.

B.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 (B.6) and (B.5), 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 (B.7) 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 (B.7) 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 (B.7) 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 \overline{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

B.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

$$\begin{pmatrix} I_t - F(I_t) \\ X_t \end{pmatrix} = \begin{pmatrix} \alpha_2 - \alpha_1 & -\alpha_1(1 - \delta) \\ 1 & 1 - \delta \end{pmatrix} \begin{pmatrix} I_{t-1} \\ X_{t-1} \end{pmatrix} + \begin{pmatrix} \alpha_0 \\ 0 \end{pmatrix}. \quad (\text{B.8})$$

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"

$$\begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \end{pmatrix} + \begin{pmatrix} f(y_{2t-1}, y_{2t-1}) \\ g(y_{1t-1}, y_{2t-1}) \end{pmatrix}, \quad (\text{B.9})$$

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 (B.8) in the standard form (B.9). 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 (B.8) as

$$\begin{pmatrix} i_t - \widehat{F}(i_t) \\ x_t \end{pmatrix} = \begin{pmatrix} \alpha_2 - \alpha_1 & -\alpha_1(1 - \delta) \\ 1 & 1 - \delta \end{pmatrix} \begin{pmatrix} i_{t-1} \\ x_{t-1} \end{pmatrix}. \quad (\text{B.10})$$

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 (B.10) a first order approximation of G around zero, we obtain

$$\begin{pmatrix} i_t - \widehat{F}(i_t) \\ x_t \end{pmatrix} = \underbrace{\begin{pmatrix} \frac{\alpha_2 - \alpha_1}{1 - F'(I^s)} & -\frac{\alpha_1(1 - \delta)}{1 - F'(I^s)} \\ 1 & 1 - \delta \end{pmatrix}}_M \begin{pmatrix} i_{t-1} \\ x_{t-1} \end{pmatrix} + \begin{pmatrix} m(i_{t-1}, x_{t-1}) \\ 0 \end{pmatrix}, \quad (\text{B.11})$$

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 $(T/2, \sqrt{D - (T/2)^2})$ 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 (B.8)

$$\begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \end{pmatrix} + \begin{pmatrix} f(y_{2t-1}, y_{2t-1}) \\ g(y_{1t-1}, y_{2t-1}) \end{pmatrix}, \quad (\text{B.9})$$

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.5 Proposition 5

Since we are considering the properties of a deterministic system, we can replace $\mathbb{E}_t [I_{t+1}]$ by I_{t+1} , and rewrite equations (11) and (1) as

$$(1 - \rho)I_t = -\alpha_1 X_t + \alpha_2 I_{t-1} + \alpha_3 I_{t+1} \tag{B.12}$$

$$X_{t+1} = (1 - \delta)X_t + I_t \tag{B.13}$$

where $\rho \in [0, 1)$ stands in for $F'(I^s)$ and we have omitted constant terms for simplicity.⁵⁷ Rewrite (B.12) as

$$I_{t+1} = b_1 I_t + b_2 I_{t-1} + b_3 X_t$$

⁵⁷Note that this formulation embeds the case where $X_{t+1} = (1 - \delta)X_t + \psi I_t$ if we replace α_3 by α_3/ψ .

where $b_1 = \frac{1-\rho}{\alpha_3}$, $b_2 = -\frac{\alpha_2}{\alpha_3}$, and $b_3 = \frac{\alpha_1}{\alpha_3}$. The characteristic equation associated with this system is

$$\lambda^3 - (1 - \delta + b_1) \lambda^2 + [b_1(1 - \delta) - b_2 - b_3] \lambda + (1 - \delta)b_2 = 0. \quad (\text{B.14})$$

Suppose d is any real root of this equation. Then this equation can be factored into the form

$$(\lambda - d) (\lambda^2 - \tau\lambda + c) = 0. \quad (\text{B.15})$$

Expanding this expression and equating the coefficients with those in (B.14), we may obtain the system

$$\tau + d = 1 - \delta + b_1 \quad (\text{B.16})$$

$$c + \tau d = b_1(1 - \delta) - b_2 - b_3 \quad (\text{B.17})$$

$$cd = -(1 - \delta)b_2. \quad (\text{B.18})$$

Suppose that this system contains a real eigenvalue that is less than or equal to -1 . Letting d denote this eigenvalue, we have from (B.18) that $c < (1 - \delta)b_2 < 0$, while from (B.16) we have that $\tau > 2 - \delta + b_1 > 1$. Since the roots of $\lambda^2 - \tau\lambda + c = 0$ are given by $\frac{\tau \pm \sqrt{\tau^2 - 4c}}{2}$, this implies that these two roots are real. The larger of these two roots is given by $\frac{\tau + \sqrt{\tau^2 - 4c}}{2} > \tau > 1$, so that we also have a real root that is greater than 1, which completes the proof.

B.6 Proposition 6

Suppose the system undergoes either type of fold bifurcation. Then, at that point, there must be a root of (B.14) equal to one. Using the factorization (B.15) with $d = 1$, we can eliminate τ and c in (B.16)-(B.18) to yield $\delta(b_2 - 1) = -\delta b_1 - b_3$, or, equivalently, $\alpha_1 = \delta(\alpha_2 + \alpha_3) - \delta(1 - \rho)$. But this latter expression implies that $\alpha_1 < \delta(\alpha_2 + \alpha_3)$, a possibility that we have ruled out in order to ensure a unique steady state. Thus, the system cannot undergo a fold bifurcation

B.7 Proposition 7

Suppose the system undergoes a flip bifurcation. Then, at that point, there must be a root of (B.14) equal to -1 . Using the factorization (B.15) with $d = -1$, we can substitute this value into (B.16)-(B.18) and eliminate τ and c to yield $\alpha_2 = \frac{\alpha_1}{2-\delta} - \alpha_3 - (1 - \rho)$, which implies that $\alpha_2 < \frac{\alpha_1}{2-\delta} - \alpha_3$. This completes the proof.

B.8 Proposition 8

Suppose the system undergoes a Hopf bifurcation. Then at that point the roots of $\lambda^2 - \tau\lambda + c = 0$ (see (B.15)) must be complex with modulus 1. The latter condition implies that $c = 1$, which implies from (B.18) that $d = -(1 - \delta)b_2$. Since we must have $d > 1$ at the bifurcation, this implies that $(1 - \delta)\alpha_2 > \alpha_3$, which completes the proof.

B.9 Proposition 9

A Hopf bifurcation will arise if there exists a $\rho \in [0, 1)$ such that $c = 1$. Substituting this value into (B.16)-(B.18) and eliminating τ and d yields $(1 - \delta)[(1 - \delta)b_2 + b_1]b_2 = 1 + [1 - (1 - \delta)^2]b_2 + b_3 - b_1(1 - \delta)$ or, equivalently,

$$1 - \rho = \frac{(1 - \delta)^2 \frac{\alpha_2^2}{\alpha_3} + [1 - (1 - \delta)^2] \alpha_2 - \alpha_1 - \alpha_3}{(1 - \delta) \left(\frac{\alpha_2}{\alpha_3} - 1 \right)}.$$

The desired result follows if the right-hand side of this expression is between 0 and 1, which is true if the condition stated in the Proposition holds.

C Simulations in Figures 7 and 8

The cycle model : Panel (a) of Figure 8 are obtained with the following reduced form model:

$$I_t = \sin(\omega t) + \mu_t$$

with $\omega = \pi/25$. In the deterministic simulations, we have $\mu_t = 0 \forall t$. In the stochastic simulation, the shock μ follows the autoregressive process $\mu_t = \rho\mu_{t-1} + \varepsilon_t$ and $\rho = .15$, $\sigma_\varepsilon = .2$.

The limit cycle model : Figure 7 and panel (b) of Figure 8 are obtained with the following reduced form model:

$$\begin{aligned} I_{it} &= \alpha_0 + u_{1t} - \alpha_1 X_{it} + \alpha_2 I_{it-1} + F\left(\frac{\sum I_{jt}}{N}\right) + \mu_t, \\ X_{it+1} &= (1 - \delta)X_{it} + I_{it}. \end{aligned}$$

when we restrict to symmetric allocations.

F is assumed to be piecewise linear and is constructed as follows:

$$F(I) = \begin{cases} a_0 + \beta_0 I & \text{if } I < \bar{I}_1 \\ a_1 + \beta_1 I & \text{if } \bar{I}_1 \leq I \leq \bar{I}_2 \\ a_2 + \beta_2 I & \text{if } I > \bar{I}_2 \end{cases}$$

Denoting by I^s the steady state level of investment, we assume $\bar{I}_1 < I^s < \bar{I}_2$. F is constrained to be continuous, which implies


$$\begin{aligned} a_0 + \beta_0 \bar{I}_1 &= a_1 + \beta_1 \bar{I}_1 \\ a_1 + \beta_1 \bar{I}_1 &= a_2 + \beta_2 \bar{I}_2 \end{aligned}$$

Finally, we assume that $F(\bar{I}_1) = \gamma_1 \bar{I}_1$ and $F(\bar{I}_2) = \gamma_2 \bar{I}_2$.

This formulation allows to choose \bar{I}_1 , \bar{I}_2 , γ_1 , γ_2 , β_1 and β_2 . The parameters β_1 , a_0 , a_1 and a_2 are then chosen to satisfy continuity of F . They are given by

$$\begin{aligned} \beta_1 &= \frac{\gamma_2 \bar{I}_2 - \gamma_1 \bar{I}_1}{\bar{I}_2 - \bar{I}_1} \\ a_0 &= (\gamma_1 - \beta_0) \bar{I}_1 \\ a_1 &= (\gamma_1 - \beta_1) \bar{I}_1 \\ a_2 &= (\gamma_2 - \beta_2) \bar{I}_2 \end{aligned}$$

Figure 20: The F function



The obtained F function is displayed on Figure 20.

Finally, given δ , α_1 and α_2 , we set α_0 such that the steady state I^s is equal to one, which implies

$$\alpha_0 = (1 + \frac{\alpha_1}{\delta} - \alpha_2 - \beta_2)I^s - a_1$$

The parameters values are $\bar{I}_1 = .5$, $\bar{I}_2 = 1.5$, $\gamma_1 = 1.2$ and $\gamma_2 = .8$, $\beta_0 = .01$ and $\beta_2 = .1$. In the deterministic simulations, we have $\mu_t = 0 \forall t$. In the stochastic limit cycle, the shock μ follows the autoregressive process $\mu_t = \rho\mu_{t-1} + \varepsilon_t$ and $\rho = .95$, $\sigma_\varepsilon = .03$.

Online Appendix

D Spectral Density Estimation

D.1 Schuster's Periodogram

We estimate the spectral density of series $\{x_t\}_{t=0}^{T-1}$ of finite length T by first computing the Discrete Fourier Transform (DFT) X_k , which results from sampling the Discrete Time Fourier Transform (DTFT) at frequency intervals $\Delta\omega = \frac{2\pi}{T}$ in $[-\pi, \pi)$:

$$X_k = X\left(e^{i\frac{2\pi}{T}k}\right) = \sum_{t=0}^{T-1} x_t e^{-i\frac{2\pi}{T}kt}, \quad (\text{D.19})$$

for $k = 1, \dots, T-1$. We then can compute samples of the Sample Spectral Density (SSD) S_k from samples of Schuster's periodogram I_k ⁵⁸ according to

$$S_k = I_k = \frac{1}{T} |X_k|^2 \quad (\text{D.20})$$

Taking advantage of the fact that X is even, this amounts to evaluating the spectral density at T frequencies equally spaced between 0 and π .⁵⁹

D.2 Zero-Padding to Increase the Graphic Resolution of the Spectrum

As we have computed only T samples of the DTFT $X(e^{i\omega})$, we might not have a detailed enough picture of the shape of the underlying function $X(e^{i\omega})$, and therefore of the spectral density $|X(e^{i\omega})|^2$. This problem is particularly acute if one is interested in the behavior of the spectrum at longer periodicities (i.e., lower frequencies). Specifically, since we uniformly sample frequencies, and since the periodicity p corresponding to frequency ω is given by $p = \frac{2\pi}{\omega}$, the spectrum is sparser at longer periodicities (and denser at shorter ones). While the degree of accuracy with which the samples of X_k describe the shape of $X(e^{i\omega})$ is dictated and limited by the length T of the data set, we can nonetheless increase the number of points at which we sample the DTFT in order to increase the graphic resolution of the spectrum. One common (and numerically efficient) way to do this is to add a number of zeros to the end of the sequence x_t before computing the DFT. This operation is referred to as zero-padding. As an example, suppose that we add exactly T zeros to the end of the length- T sequence $\{x_t\}$. One can easily check that this has no effect on the DFT computed at the original T sampled frequencies, instead simply adding another set of T sampled frequencies at the midpoints between each successive pair of original frequencies.⁶⁰

⁵⁸Another approach for obtaining the spectral density is to take a Fourier transform of the sequence of autocovariances of x . We show below that this method gives essentially the same result when applied to our hours series.

⁵⁹See Priestley [1981] for a detailed exposition of spectral analysis, Alessio [2016] for practical implementation and Cochrane [2005] for a quick introduction.

⁶⁰This is true when the number of zeros added to the end of the sample is an integer multiple of T . When instead a non-integer multiple is added, the set of frequencies at which the padded DFT is computed no longer contains the original set of points, so that the two cannot be directly compared in this way. Nonetheless, the overall pattern of the sampled spectrum is in general unaffected by zero-padding.

If one is interested in the behavior of the spectral density at long enough periodicities, zero-padding in this way is useful. We will denote by N the number of samples at which the DTFT (and thus the SSD) is sampled, meaning that $T' = N - T$ zeros will be added to the sequence $\{x_t\}$ before computing the DFT. In the main text, we have set $N = 1,024$.⁶¹

D.3 Smoothed Periodogram Estimates

We obtain the raw spectrum estimate of a series non-parametrically as the squared modulus of the DFT of the (zero-padded) data sequence, divided by the length of the data set.⁶² This estimate is called Schuster's periodogram, or simply the periodogram. It turns out that the periodogram is asymptotically unbiased, but is not a consistent estimate of the spectrum, and in particular the estimate of the spectrum at a given frequency ω_k is generally quite unstable (i.e., it has a high standard error). Notwithstanding this fact, the overall pattern of the spectrum is much more stable, in the sense that the average value of the estimated spectrum within a given frequency band *surrounding* ω_k is in fact consistent. In order to obtain a stable and consistent estimate of the spectrum, we exploit this fact by performing frequency-averaged smoothing. In particular, we obtain our estimate of the SSD $S(\omega)$ by kernel-smoothing the periodogram $I(\omega)$ over an interval of frequencies centered at ω . Since the errors in adjacent values of $I(\omega)$ are uncorrelated in large samples, this operation reduces the standard errors of the estimates without adding too much bias. In our estimations, we use a Hamming window of length $W = 13$ as the smoothing kernel.⁶³

D.4 Smoothing and Zero-Padding with a Multi-Peaked Spectral Density

To illustrate the effects of smoothing and zero-padding, in this section we compare the estimated spectral density with the known theoretical one for a process that exhibits peaks in the spectral density at periods 20, 40 and 100 quarters. We think this is a good description of the factor variables we are studying (i.e., hours worked, unemployment, capacity utilization), that display both business cycle movements and lower-frequency movements unrelated to the business cycle. We construct our theoretical series as the sum of three independent stationary AR(2) processes, denoted x_1 , x_2 and x_3 .

Each of the x_i follows an AR(2) process

$$x_{it} = \rho_{i1}x_{it-1} + \rho_{i2}x_{it-2} + \varepsilon_{it},$$

where ε_i is i.i.d. $N(0, \sigma_i^2)$. The spectral density of this process can be shown to be given by

$$S(\omega) = \sigma_i^2 \{2\pi[1 + \rho_{i1}^2 + \rho_{i2}^2 + 2(\rho_{i1}\rho_{i2} - \rho_{i1})\cos(\omega) - 2\rho_{i2}\cos(2\omega)]\}^{-1}$$

It can also be shown (see, e.g., Sargent [1987]) that for a given ρ_{i2} , the spectral density has a peak at frequency $\bar{\omega}_i$ if

$$\rho_{i1} = -\frac{4\rho_{i2}\cos(2\pi/\bar{\omega}_i)}{1 - \rho_{i2}}$$


⁶¹As is well known, standard numerical routines for computing the DFT (i.e., those based on the Fast Fourier Transform algorithm) are computationally more efficient when N is a power of 2, which is why we set $N = 1,024$ rather than, say, $N = 1,000$.

⁶²Note that we divide by the original length of the series (i.e., T), rather than by the length of the zero-padded series (i.e., N).

⁶³Using alternative kernel functions makes little difference to the results.

We set $\bar{\omega}_i$ equal to 20, 40, and 100 quarters, respectively, for the three processes, and ρ_{i2} equal to -0.9, -0.95, and -0.95. The corresponding values for ρ_{i1} are 1.802, 1.9247, and 1.9449. We set σ_i equal to 6, 2, and 1. We then construct $x_t = x_{1t} + x_{2t} + x_{3t}$. The theoretical spectral density of x is shown in Figure 21. As in the factor utilisation series we are using in the main text, the spectral density shows long-run fluctuations, but the bulk of the business cycle movements is explained by movements at the 40-quarter periodicity, although we observe another peak at periodicity 20 quarters.

Figure 21: Theoretical Spectral Density (Sum of Three AR(2))

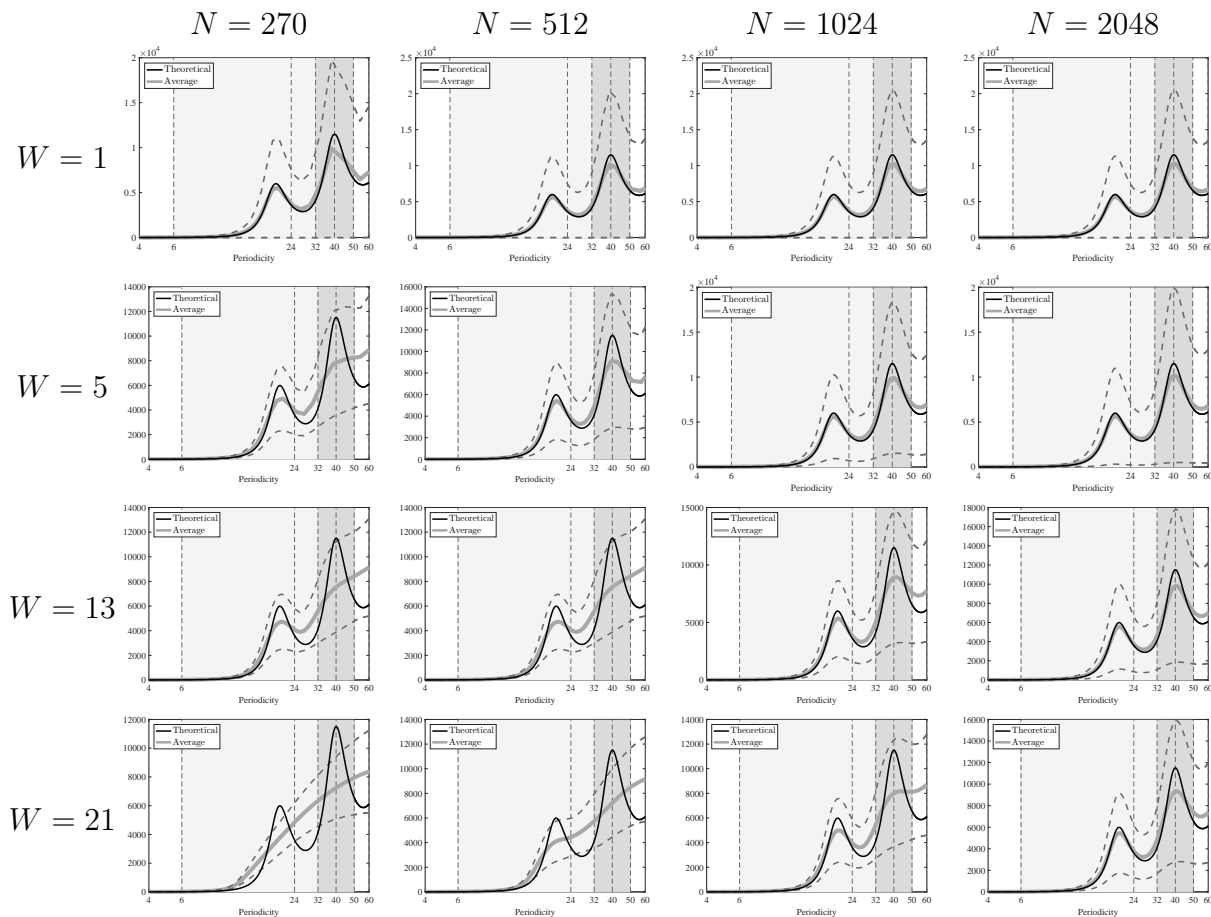


Notes: Figure shows the theoretical spectral density of the sum of three independent AR(2) processes, which have peaks in their spectral densities at, respectively, 20, 40 and 100 quarters.

We simulate this process 1,000,000 times, with $T = 270$ for each simulation, which is the length of our observed macroeconomic series. We estimate the spectral density for various values of N (zero-padding) and W (length of the Hamming window). Higher N corresponds to higher resolution, and higher W to more smoothing. On each panel of Figure 22, we report the mean of the estimated spectrum over the 1,000,000 simulations (solid grey line), the mean \pm one standard deviation (dashed lines), and the theoretical spectrum (solid black line). As we can see moving down the figure (i.e., for increasing W), more smoothing tends to reduce the error variance, but at the cost of increasing bias. Effectively, the additional smoothing “blurs out” the humps in the true spectrum. For example, with no zero-padding ($N = 270$), the peak in the spectral density at 40 quarters is (on average) hardly detected once we have any smoothing at all. Meanwhile, moving

rightward across the figure (i.e., for increasing N), we see that more zero-padding tends to reduce the bias (and in particular, allows for the humps surrounding the peaks to be better picked up on average), but typically increases the error variance. As these properties suggest, by appropriately choosing the combination of zero-padding and smoothing, one can minimize the error variance while maintaining the ability to pick up the key features of the true spectrum (e.g., the peaks at 20 and 40 quarters).

Figure 22: Effects of Smoothing and Zero-Padding (Sum of Three AR(2))

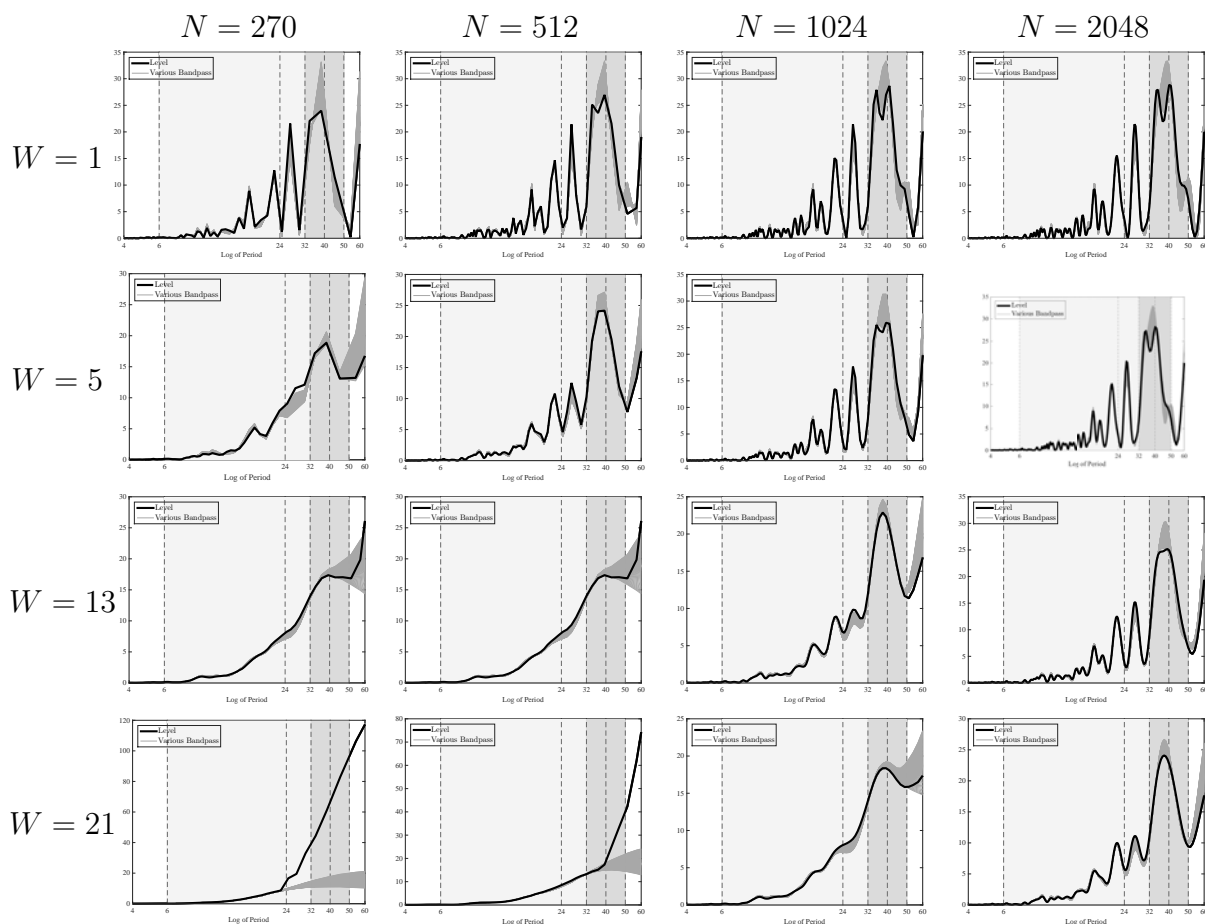


Notes: Figure shows estimates of the spectral density using simulations of the sum of three independent AR(2), which have peaks in their spectral densities at, respectively, 20, 40 and 100 quarters. The black line is the theoretical spectrum, the solid grey line is the average estimated spectrum over 1,000,000 simulations, and the dotted grey lines corresponds to that average \pm one standard deviation. W is the length of the Hamming window (smoothing parameter) and N is the number of points at which the spectral density is evaluated (zero-padding parameter).

D.5 Smoothing and Zero-Padding with Non-Farm Business Hours per Capita

Figure 23 presents estimates of the spectral density of U.S. Non Farm Business Hours per Capita for different choices of the zero-padding parameter (N) and different lengths of the Hamming window (W). The results indicate that, as long as the amount of zero-padding is not too small (i.e., N larger), we systematically observe the peak at around 40 quarters in the spectral density. In fact, it is only with minimal zero-padding (N low) and a wide smoothing window (W high) that the peak is entirely washed out. We take this as evidence of the robustness of that peak.

Figure 23: Changing Smoothing and Zero-Padding



Notes: Figure shows estimates of the spectral density of U.S. Non-Farm Business Hours per Capita over the sample 1947Q1-2015Q2. The different lines correspond to estimates of the spectral density of hours in levels (black line) and of 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (thin grey lines). W is the length of the Hamming window (smoothing parameter) and N is the number of points at which the spectral density is evaluated (zero-padding parameter).

D.6 Detrending with a Polynomial Trend

In this section, we check that detrending our hours series with a polynomial trend of degrees 1 to 5 does not affect our main finding; namely, the existence of a peak in the spectrum at a periodicity around 40 quarters. Plots confirming that our finding is robust to polynomial detrending are shown in Figure 24.

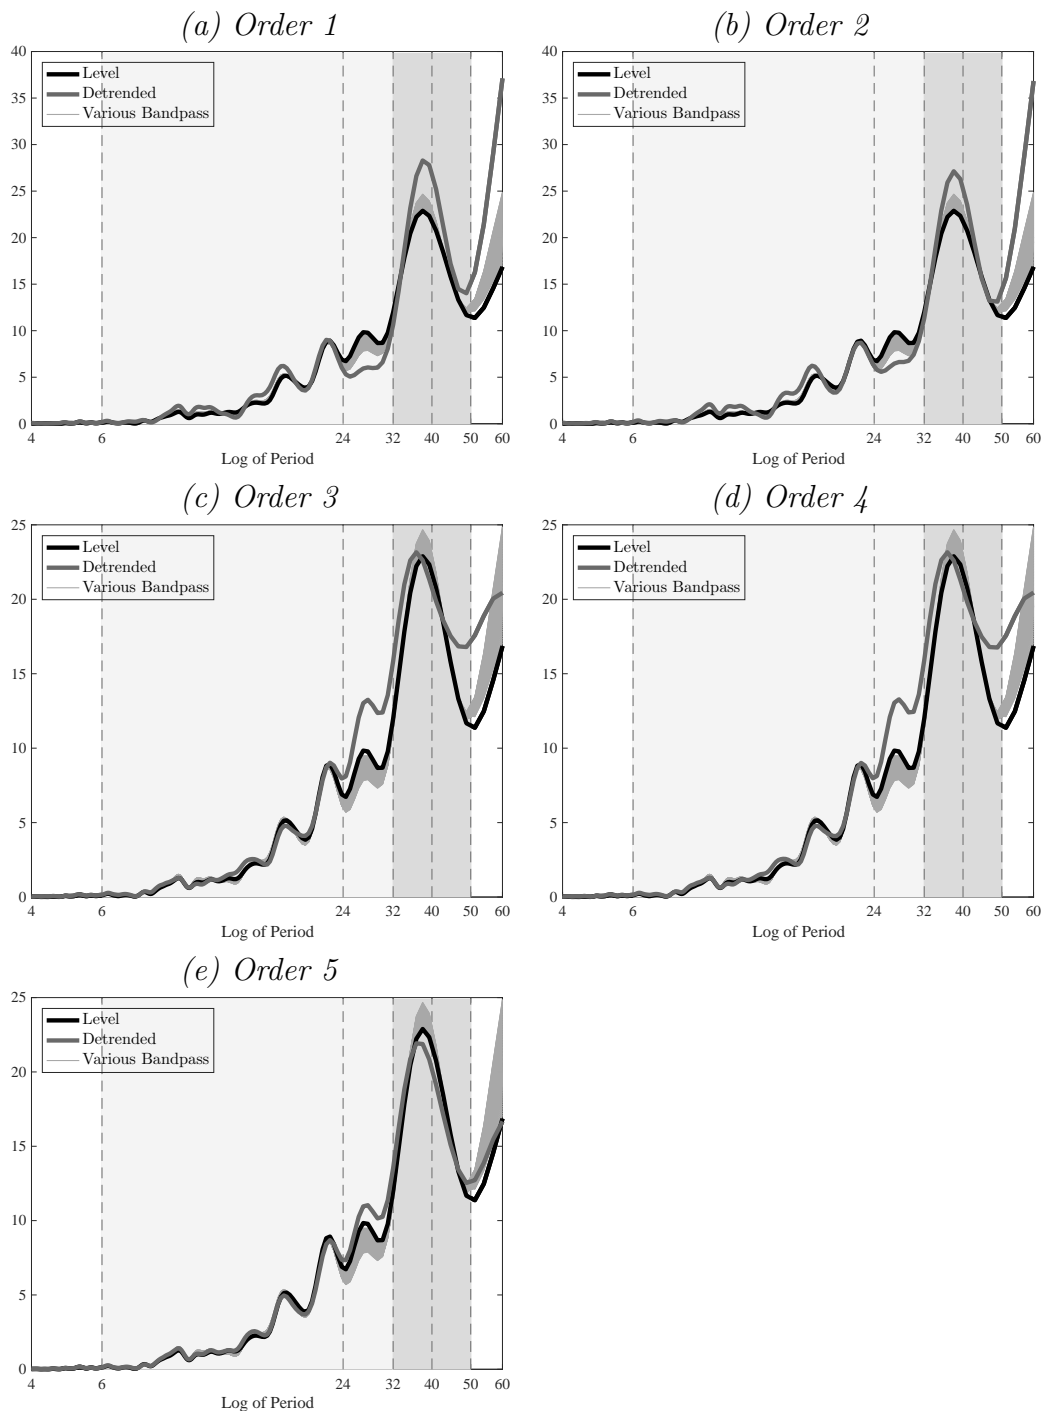
D.7 Alternative Estimators

As another robustness test, we estimate the spectrum using the `SPECTRAN` package (for `MATLAB`), which is described in Marczak and Gómez [2012]. The spectrum is computed in this case as the Fourier transform of the covariogram (rather the periodogram as we have done thus far). Smoothing is achieved by applying a window function of length M to the covariogram before taking its Fourier transform.⁶⁴ Three different window shapes are proposed: the Blackman-Tukey window, Parzen window, and Tukey-Hanning window. The width of the window used in estimation is set as a function of the number of samples of the spectrum. In the case where no zero-padding is done ($N = 270$), these “optimal” widths correspond to lengths of, respectively, $M = 68$, 89, and 89 quarters for the three methods.⁶⁵ Figure 25 shows the estimated spectrum of Non Farm Business hours for the three windows and with or without zero-padding ($N = 270$, 512, or 1024). Results again confirm the existence of a peak at a periodicity around 40 quarters, as long as there is enough zero-padding.

⁶⁴Specifically, the k -th-order sample autocovariance is first multiplied by $w(|k|)$, where the window function w is an even function with the property that $\max_k w(k) = w(0) = 1$, and the window length $M > 0$ is such that $w(|k|) \neq 0$ for $|k| = M - 1$ and $w(|k|) = 0$ for $|k| \geq M$.

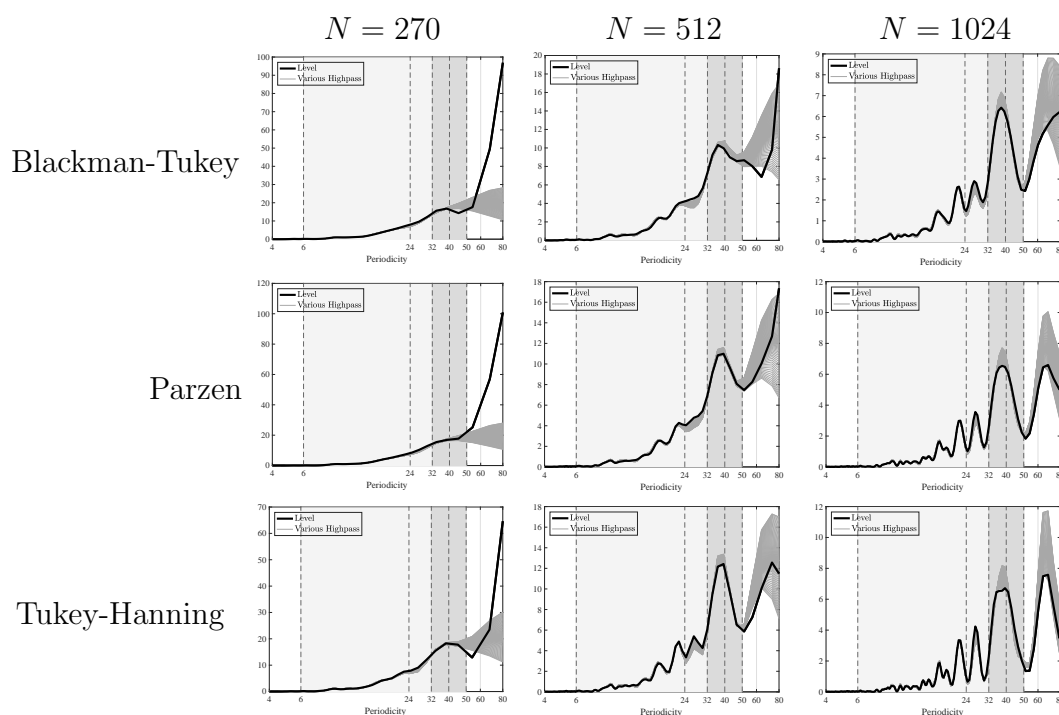
⁶⁵Note that, in contrast to the kernel-smoothing case, in this case a wider window corresponds to less smoothing.

Figure 24: Using a Polynomial Trend of Various Orders for Benchmark Smoothing ($W = 13$) and Zero-Padding ($N = 1024$)



Notes: Figure shows estimates of the spectral density of U.S. Non-Farm Business Hours per Capita over the sample 1947Q1-2015Q2, when polynomial trends of order 1 to 5 have been removed from the data. The different lines correspond to estimates of the spectral density of hours in levels (black line), of hours detrended with a polynomial trend (thick grey line) and of 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (thin grey lines).

Figure 25: Non-Farm Business Hours with Various Windows and Estimation Using the Covariogram



Notes: Notes: This Figure shows estimates of the spectral density of U.S. Non Farm Business Hours per Capita over the sample 1947Q1-2015Q2, as computed from the covariogram using the SPECTRAN package. The different lines correspond to estimate of the spectral density of hours in levels (black line) and of 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (thin grey lines). N is the number of points at which is evaluated the spectral density (zero padding).

E Steps in Deriving the Equilibrium of the New Keynesian Model with Financial Frictions and Durables

E.1 The Model

Households

There is a continuum of mass one of households indexed by i who purchase consumption services from the market. The preferences of agent i are given by $\mathbb{E}_0 \sum_t \beta^t \xi_{t-1} [U(C_{it} - \gamma C_{t-1}) + \nu(1 - L_{it})]$, with $U'(\cdot) > 0$, $U''(\cdot) < 0$, $\nu'(\cdot), \nu''(\cdot) < 0$ and $0 \leq \gamma < 1 - \delta$. C_{it} represents the consumption services purchased by household i in period t , C_t denotes the average level of consumption in the economy, β is the discount factor, and ξ_t denotes an exogenous shock to the discount factor at date t . Note that this preference structure assumes the presence of external habit. Household i 's problem takes the form

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \xi_{t-1} [U(C_{it} - \gamma C_{t-1}) + \nu(1 - L_{it})]$$

subject to

$$\begin{aligned} P_t C_{it} + (1 + r_{t-1}) B_{it} + P_t^x D_{it} &= B_{it+1} + r_t^x X_{it} + w_t L_{it} + \Gamma_{it}, \\ X_{it+1} &= (1 - \delta) X_{it} + D_{it}. \end{aligned}$$

Here, B_{it+1} represents the borrowing by the household at time t , to be repaid with interest at $t + 1$, r_t represents the nominal interest rate faced by the household on such borrowing, $w_t L_{it}$ is labor income, X_{it} represents the stock of durable goods (or houses) held by the household at the beginning of period t , Γ_{it} are firm profits that are returned to households, and D_{it} the quantity of new durable goods purchased by the household at t . P_t and P_t^x are the (nominal) prices of consumption services and new durable goods, respectively, at date t . Households are assumed to buy all of their consumption services—including those derived from durable goods—from the market. Specifically, households do not consume the services of their durable goods directly. Instead, they rent X_{it} out to firms each period at nominal rental rate r_t^x . Firms then combine the rented stock of durables with labor in order to produce consumption services, C , as well as new durable goods, D , in a manner to be described shortly. Under these conditions, dropping i subscripts, the household's Euler equation with respect to their optimal choice of C takes the familiar form:

$$U'(C_t - \gamma C_{t-1}) = \beta_t \mathbb{E}_t [U'(C_{t+1} - \gamma C_t) (1 + r_t - \pi_{t+1})], \quad (\text{E.21})$$

where π_{t+1} is the inflation rate from t to $t + 1$, and $\beta_t \equiv \beta \frac{\xi_t}{\xi_{t-1}}$. Note that we have here used the approximation $\frac{1+r}{1+\pi} = 1 + r - \pi$.

The households problem also leads a labor supply decisions represented by

$$\frac{\nu'(L_t)}{U'(C_t - \gamma C_{t-1})} = \frac{w_t}{P_t}$$

and an arbitrage condition between the holding of bonds and capital given by

$$\mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{P_{t+1}} \left\{ (1 + r_t) - \left(\frac{r_{t+1}^x + (1 - \delta) P_{t+1}^x}{P_t^x} \right) \right\} \right] = 0.$$

Final Good producers

The final good sector is competitive. This sector provides consumption services to households by buying a set of differentiated intermediate services, denoted C_{jt} , from intermediate good firms. We assume a measure one of intermediate good firms, indexed by j . The technology of the final good firms is

$$C_t = \left(\int_0^1 C_{jt}^\eta dj \right)^{\frac{1}{\eta}}, \quad (\text{E.22})$$

with $\eta \in (0, 1)$. The objective of the final good firm is thus to solve

$$\max P_t C_t - \int_0^1 P_{jt} C_{jt} dj$$

subject to (E.22), where p_{jt} is the price of intermediate good j . This gives rise to demand for intermediate good j given by

$$C_{jt} = \left(\frac{P_{jt}}{P_t} \right)^{-\frac{1}{1-\eta}} C_t.$$

Intermediate good firms

Intermediate good producers are monopolistically competitive and take the demand from final good firms as given. We assume that intermediate firms produce an intermediate factor M according to the technology

$$M_{jt} = BF(\Theta_t L_{jt}).$$

A fixed fraction $(1 - \varphi)$ of that factor is transformed one to one into consumption services, whereas the use of the rented stock of durables also produces consumption services with a one to one technology, so that the amount of consumption services produced by intermediate firm j is

$$C_{jt} = X_{jt} + (1 - \varphi)M_{jt} = X_{jt} + (1 - \varphi)BF(\Theta_t L_{jt}). \quad (\text{E.23})$$

The remaining fraction φ of the intermediate factor M is transformed one to one into durable goods D , so that

$$D_{jt} = \varphi M_{jt} = \varphi BF(\Theta_t L_{jt}).$$

Period- t profits of an intermediate good producer are given by

$$\Gamma_{jt} = P_{jt} C_{jt} + P_t^X D_{jt} - r_t^x X_{jt} - w_t L_{jt}.$$

Cost minimisation implies that

$$w_t = ((1 - \varphi)r_t^x + \varphi P_t^x) \Theta_t BF'(\Theta_t L_{jt}).$$

Hence, assuming $F'' < 0$, all intermediate good firms will hire the same amount of labor (i.e., $L_{jt} = L_t$), which implies that the aggregate capital stock will satisfy

$$X_{t+1} = (1 - \delta) X_t + \varphi BF(\Theta_t L_t). \quad (\text{E.24})$$

With the normalisation $B = \frac{1}{1-\varphi}$ and denoting $\psi = \frac{\varphi}{1-\varphi}$, we obtain

$$C_t = X_t + F(\Theta_t L_t) \quad (\text{E.25})$$

and

$$X_{t+1} = (1 - \delta) X_t + \psi F(\Theta_t L_t). \quad (\text{E.26})$$

Setting of interest rates

We assume that the interest rate faced by households is equal to the policy rate set by the central bank (which we denote i_t) plus a risk premium r_t^p ; that is, we assume that $r_t = i_t + r_t^p$. We directly allow for counter-cyclical risk premium by positing r_t^p to be a non-increasing function of the level of employment in the economy, L_t , via

$$r_t^p = R(L_t), \quad R'(L_t) \leq 0.$$

In setting the policy rate i_t , we assume that the central bank follows a slightly modified Taylor-type rule where the interest rate reacts to expected inflation and to expected labor market conditions via

$$i_t = \phi_\pi \mathbb{E}_t [\pi_{t+1}] + \phi_\ell \mathbb{E}_t [G(L_{t+1})], \quad \phi_\pi, \phi_\ell, G' > 0. \quad (\text{E.27})$$

Equilibrium in the absence of sticky prices

In the absence of sticky prices, all intermediate good producers will act in the same manner and set P_{jt} as a mark-up over the marginal cost, implying that

$$P_{jt} = \frac{r_t^x}{1 - \eta}.$$

The equilibrium values for $\{P_t^x, w_t, r_t^x, P_t, r_t, \pi_t, i_t, L_t, C_t, X_t\}$ are given as the solution to the set of equations:

$$\begin{aligned} U'(C_t - \gamma C_{t-1}) &= \beta_t \mathbb{E}_t [U(C_{t+1} - \gamma C_t) (1 + r_t - \pi_{t+1})], \\ \frac{\nu'(L_t)}{U'(C_t - \gamma C_{t-1})} &= \frac{w_t}{P_t}, \\ \mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{P_{t+1}} \left\{ (1 + r_t) - \left(\frac{r_{t+1}^x + (1 - \delta)P_{t+1}^x}{P_t^x} \right) \right\} \right] &= 0, \\ w_t &= (r_t^x + \psi P_t^x) \Theta_t F'(\Theta_t L_t), \\ P_t &= \frac{r_t^x}{1 - \eta}, \\ C_t &= X_t + F(\Theta_t L_t), \\ X_{t+1} &= (1 - \delta) X_t + \psi F(\Theta_t L_t), \\ i_t &= \phi_\pi \mathbb{E}_t [\pi_{t+1}] + \phi_\ell \mathbb{E}_t [G(L_{t+1})], \\ \pi_t &= \frac{P_t}{P_{t-1}}, \\ r_t &= i_t + R(L_t). \end{aligned}$$

In fact, this system does not determine the level of P_t , but instead gives solutions for $\left\{ \frac{P_t^x}{P_t}, \frac{w_t}{P_t}, \frac{r_t^x}{P_t}, \pi_t, i_t, r_t, L_t, C_t, X_t \right\}$.

With sticky prices

Assume now that intermediate firms are confronted to price stickiness, in that the arrival of the option to change their price follows a Poisson process. In that case, the previous equilibrium conditions are not affected except for the two equations

$$P_t = \frac{r_t^x}{1 - \eta} \quad (\text{E.28})$$

$$C_t = X_t + F(\Theta_t L_t); \quad (\text{E.29})$$

With sticky prices, the pricing equation (E.28) is replaced by a pricing equation for a firm that adjusts its price, with the result that the price is set to essentially be a markup on expected discounted marginal cost. From this we can derive a Phillips curve following standard steps.

Aggregate output in the presence of sticky prices is now be given by

$$C_t = \left(\int_0^1 (X_{jt} + F(\Theta_t L_{jt}))^\rho dj \right)^{\frac{1}{\rho}},$$

where this differs from (E.29) because of the distribution in X_{jt} . The X_{jt} will satisfy the equation

$$X_{jt} - F(\Theta_t L_{jt}) = \left(\frac{P_{jt}}{P_t} \right)^{\frac{-1}{1-\rho}} C_t$$

One important approximation step we are taking in the paper is that we disregard the impact of production dispersion on aggregate output, and are indeed assuming that (E.29) holds even in the case with sticky prices, while this is only approximatively true.

Deriving the linear reduced form

With the assumption $\phi_\pi = 1$, we can bypass the need to be explicit about the intermediate firm's pricing problem and its implications for inflation, as the realisations of inflation do not feed back into the determination of the quantity variables that are of interest to us. Specifically, substituting (E.27) (with $\phi_\pi = 1$) and (E.29) into the Euler Equation (E.21), and denoting $\varepsilon_{t+1}^\pi = \mathbb{E}_t[\pi_{t+1}] - \pi_{t+1}$, we obtain:

$$\begin{aligned} U'(X_t + F(\theta_t L_t) - \gamma(X_{t-1} + F(\theta_{t-1} L_{t-1}))) = \\ \beta_t \mathbb{E}_t \left\{ U'(X_{t+1} + F(\theta_{t+1} L_{t+1}) - \gamma(X_t + F(\theta_t L_t))) \right. \\ \left. \times (1 + \varepsilon_{t+1}^\pi + \phi_\ell G(L_{t+1}) + R(L_t)) \right\}. \quad (\text{E.30}) \end{aligned}$$

We assume the following functional forms. The utility function is $U(z) = -\exp\{-\frac{z}{\omega}\}$, $\omega > 0$, the production of new goods is given by $F(\Theta L) = \ln(\Theta L)$, the interest rate policy function is $G(L) = \ln(L)$, and Θ_t is constant. With the approximation $\ln(1+z) \approx z$, we have

$$\ln(1 + \varepsilon_{t+1}^\pi + \phi_\ell G(L_{t+1}) + R(L_t)) \approx \varepsilon_{t+1}^\pi + \phi_\ell \ln L_{t+1} + R(L_t).$$

Therefore, (E.30) can be rewritten, with the approximation $\mathbb{E}[\ln(z)] \approx \ln(\mathbb{E}[z])$:

$$\begin{aligned} \frac{1}{\omega} (X_t + \ln \Theta + \ln L_t - \gamma(X_{t-1} + \ln \Theta + \ln L_{t-1})) = \\ \ln \beta_t + \frac{1}{\omega} (\mathbb{E}_t[X_{t+1}] + \ln \Theta + \mathbb{E}_t[\ln L_{t+1}] - \gamma(X_t + \ln \Theta + \ln L_t)) \\ + \mathbb{E}_t[\varepsilon_{t+1}^\pi] + \phi_\ell \mathbb{E}_t[\ln L_{t+1}] + R(L_t). \quad (\text{E.31}) \end{aligned}$$

Given that $\mathbb{E}_t [\varepsilon_{t+1}^\pi] = 0$, this equations rewrites, using $\ell_t = \ln L_t$,

$$\ell_t = \mu_t - \hat{\alpha}_1 X_t + \hat{\alpha}_2 \ell_{t-1} + \hat{\alpha}_3 \mathbb{E}_t [\ell_{t+1}] + \hat{F}(\ell_t), \quad (\text{E.32})$$

where $\mu_t = \frac{1}{\kappa} \left[\psi \left(1 - \frac{\gamma}{1-\delta} \right) \ln \Theta - \omega \ln \beta_t \right]$, $\hat{\alpha}_1 = \frac{\delta}{\kappa} \left(1 - \frac{\gamma}{1-\delta} \right)$, $\hat{\alpha}_2 = \frac{\gamma}{\kappa} \left(1 - \frac{\psi}{1-\delta} \right)$, $\hat{\alpha}_3 = \frac{1-\omega\phi_\ell}{\kappa}$ and $\hat{F}(\ell_t) = -\frac{\omega}{\kappa} \tilde{R}(\ell_t)$, with $\kappa = 1 + \gamma - \psi > 0$.

Equation (E.32) and the accumulation Equation (E.26) together form a two-variable linear system in X and ℓ . It is worth noticing that inflation π does not appear anywhere. Since C and D can be obtained from X and L , it then follows that all quantity variables are determined independently from the inflation rate.

F Solving and Estimating the Model

F.1 Solution Method

F.1.1 Solving a Deterministic Version of the New Keynesian Model

Consider the deterministic version of our model, with all variables expressed as deviations from the steady state. We can write our dynamic economic system as

$$\begin{pmatrix} y_{t+1} \\ \ell_{t+1} \end{pmatrix} = \begin{pmatrix} f(y_t, \ell_t) \\ g(y_t, \ell_t) \end{pmatrix} \equiv h(y_t, \ell_t), \quad (\text{F.33})$$

where $y_t = (X_t, \ell_{t-1})'$ is the vector of predetermined variables, $h(0) = 0$, and y_0 is given. Letting $x_t \equiv (y_t, \ell_t)'$, a solution is a function $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that, after setting $\ell_0 = \phi(y_0)$, the resulting sequence $\{x_t\}$ obtained from (F.33) satisfies the transversality condition (TVC) $\limsup_{t \rightarrow \infty} \|x_t\| < \infty$ (i.e., the system remains bounded). Suppose this ϕ exists and is unique, and let $x(t; y) \equiv h^t(y, \phi(y))$ ⁶⁶ denote the state at date t when $y_0 = y$ and $\ell_0 = \phi(y)$ (i.e., we have $x_t = x(t; y_0)$). Letting

$$\mathcal{M} \equiv \{x(0; y) : y \in \mathbb{R}^2\} \subset \mathbb{R}^3,$$

it must be the case that $x(t; y) \in \mathcal{M}$ for all t, y .⁶⁷ That is, \mathcal{M} is the unique 2-dimensional invariant manifold⁶⁸ of h such that system (F.33), when restricted to \mathcal{M} , produces non-explosive dynamics.⁶⁹ In this case, ϕ is the function that projects x onto \mathcal{M} by choice of l . We henceforth make the following assumption about ϕ (if it exists), which is necessary if we are to solve for it using perturbation methods (as we do below): ϕ is analytic on a neighborhood of the steady state.

To find \mathcal{M} , let $A \equiv D_x h(0)$,⁷⁰ and write the linearized version of (F.33) as.

$$x_{t+1} = Ax_t. \quad (\text{F.34})$$

⁶⁶We use h^t to denote the t -th iterate of h ; i.e., h^0 is the identity function and $h^t \equiv h \circ h^{t-1}$ for $t \geq 1$.

⁶⁷To see this, let $y(t; y)$ denote the first two elements of $x(t; y)$ and $l(t; y)$ the third, and note that the trajectory passing through $x(t; y) \equiv (y(t; y)', l(t; y))'$ remains bounded by construction. Furthermore, also by construction, the trajectory passing through $x(0; y(t; y)) \equiv (y(t; y)', \phi(y(t; y)))'$ remains bounded as well. If $l(t; y) \neq \phi(y(t; y))$, then if we had $y_0 = y(t; y)$ there would be two different solutions for ℓ_0 that produce bounded trajectories, which contradicts our supposition that ϕ is unique.

⁶⁸A manifold $\tilde{\mathcal{M}}$ is said to be h -invariant if $x \in \tilde{\mathcal{M}}$ implies $h(x) \in \tilde{\mathcal{M}}$.

⁶⁹That is, \mathcal{M} is precisely the “non-explosive manifold” referred to in the text.

⁷⁰Here, the notation D_z indicates the partial derivative operator with respect to the vector z .

Let $\mathcal{E} \subset \mathbb{R}^3$ be a 2-dimensional invariant subspace of A ; that is, (a) $\mathcal{E} = \{\alpha z_1 + \beta z_2 : \alpha, \beta \in \mathbb{R}\}$ for some linearly independent basis vectors $z_1, z_2 \in \mathbb{R}^3$, and (b) if $x \in \mathcal{E}$ then $Ax \in \mathcal{E}$. Note that there are at most three possible such subspaces.⁷¹ Given such an \mathcal{E} , we look for a candidate \mathcal{M} as the 2-dimensional invariant manifold of (F.33) with the following properties: (a) \mathcal{M} is tangent to \mathcal{E} at the steady state; and (b) the associated projection function ϕ is analytic on a neighborhood of the steady state.⁷² Given a candidate \mathcal{M} and the associated ϕ , we can then obtain an expression for the evolution of y_t on \mathcal{M} as

$$y_{t+1} = \theta(y_t) \equiv f(y_t, \phi(y_t)).$$

Using this expression, we can check numerically whether indeed trajectories on the candidate \mathcal{M} remain bounded. If they do, then \mathcal{M} continues to be a candidate non-explosive manifold; otherwise, we rule it out. We can check all possible \mathcal{E} 's and their associated candidate \mathcal{M} 's in this way. If, after doing so, we have found more than one candidate \mathcal{M} , then we conclude that the solution is indeterminate (i.e., that ϕ is not unique). If, on the other hand, all candidate \mathcal{M} 's contain unbounded trajectories, then we conclude that no solution exists. Finally, if exactly one candidate \mathcal{M} is found whose trajectories are bounded, then we conclude that \mathcal{M} is indeed the non-explosive manifold we sought.⁷³

F.1.2 Finding ϕ Using Perturbation Methods

Given a subspace \mathcal{E} as defined above, the associated candidate ϕ (which fully determines \mathcal{M}) will in general not possess a closed-form solution. We thus solve for ϕ up to a k -th-order Taylor approximation around the steady state as follows. First, note that, by invariance of \mathcal{M} , ϕ must implicitly solve⁷⁴

$$\phi(f(y, \phi(y))) = g(y, \phi(y)). \quad (\text{F.35})$$

Differentiating (F.35) once with respect to y and evaluating at $y = 0$ yields

$$D_y \phi \cdot [D_y f + D_l f \cdot D_y \phi] = D_y g + D_l g \cdot D_y \phi, \quad (\text{F.36})$$

which is a quadratic system in the two elements of $D_y \phi$. In general, this equation has a multiplicity of solutions. However, $D_y \phi$ is, by definition,⁷⁵ the function that projects x onto \mathcal{E} by choice of l .

⁷¹To ensure this, we must actually impose the following additional technical requirement in constructing \mathcal{E} : if the algebraic multiplicity of an eigenvalue λ of A is greater than one, and \mathcal{E} contains a generalized eigenvector corresponding to λ , then \mathcal{E} must contain *all* of the generalized eigenvectors corresponding to λ . We henceforth impose this requirement. As a consequence, if the three eigenvalues of A are all real and distinct, then there are three \mathcal{E} 's; namely, those associated with the three possible ways of choosing two linearly independent eigenvectors of A as the basis vectors z_1 and z_2 . If instead A has an eigenvalue λ with an algebraic multiplicity of two, then there is only one \mathcal{E} ; i.e., the one given by the generalized eigenspace of λ . If instead A has a complex conjugate pair of eigenvalues, then there is also only one \mathcal{E} ; i.e., the one obtained by taking $\text{Re}(v)$ and $\text{Im}(v)$ as the basis vectors, where v is any eigenvector associated with one of the complex eigenvalues. Finally, if A has an eigenvalue with an algebraic multiplicity of three, then such an \mathcal{E} does not exist.

⁷²In general, there are an uncountably infinite number of manifolds satisfying requirement (a), but only one of these manifolds also satisfies requirement (b).

⁷³If A has either an eigenvalue with an algebraic multiplicity of two or a complex conjugate pair of eigenvalues, so that there is only one \mathcal{E} , then we also need to verify that trajectories passing through points *not* on \mathcal{M} do become unbounded. If they do not, then we also conclude in this case that the solution is indeterminate.

⁷⁴This expression can be obtained from (F.33) by using the ‘‘invariance’’ relationships $\ell_t = \phi(y_t)$ and $\ell_{t+1} = \phi(y_{t+1})$, and the fact that $y_{t+1} = f(y_t, \ell_t)$.

⁷⁵In particular, we are using here the fact that \mathcal{M} is tangent to \mathcal{E} at the steady state.

It can be verified that this projection function—which can be obtained using standard eigenvalue-eigenvector methods—is indeed a solution to (F.36). This gives us a candidate ϕ up to a first-order approximation.

To obtain the desired k -th-order approximation, we proceed iteratively as follows. Suppose we have in hand a $(j - 1)$ -th order approximation. To obtain the j -th-order approximation, obtain all j -th-order derivatives of the expression (F.35) and evaluate them at $y = 0$. We can express the result as a system of linear equations in the j -th-order derivatives of ϕ , with coefficients that are functions of the known derivatives of ϕ up to $(j - 1)$ -th order. It is thus straightforward to solve this system for the j -th-order derivatives.

F.1.3 Extension to the Stochastic Case

Write the stochastic version of the model as

$$\begin{pmatrix} y_{t+1} \\ \mathbb{E}_t[\ell_{t+1}] \end{pmatrix} = \begin{pmatrix} f(y_t, \ell_t, \mu_t; \sigma_t) \\ g(y_t, \ell_t, \mu_t; \sigma_t) \end{pmatrix} \equiv h(y_t, \ell_t, \mu_t; \sigma_t), \quad (\text{F.37})$$

with $\mu_t = \rho\mu_{t-1} + \sigma_t\eta_t$, $\eta_t \sim N(0, 1)$, and $\sigma_t = \sigma$.⁷⁶ Write the augmented state vector of this system as $z_t = (x'_t, \mu_t, \sigma_t)'$. Whereas before the solution ϕ was the function that projected the system onto a manifold \mathcal{M} in x -space, in the stochastic environment the corresponding manifold is in z -space, and in particular, $\phi : \mathbb{R}^4 \rightarrow \mathbb{R}$ is the function such that (a) $\ell_t = \phi(y_t, \mu_t; \sigma)$ for all t satisfies (F.37), and (b) the sequence generated by

$$y_{t+1} = \theta(y_t, \mu_t; \sigma) \equiv f(y_t, \phi(y_t, \mu_t; \sigma), \mu_t; \sigma)$$

satisfies the stochastic TVC, $\limsup_{t \rightarrow \infty} \mathbb{E}_0 \|y_t\| < \infty$.

Noting that $D_\sigma h(0) = 0$ in our setup, and taking date- t expectations of both sides of the linearized version of (F.37), we may obtain

$$\mathbb{E}_t[z_{t+1}] = \begin{pmatrix} 0 \\ 0 \\ \sigma \end{pmatrix} + \underbrace{\begin{pmatrix} A_x & A_\mu & 0 \\ 0 & \rho & 0 \\ 0 & 0 & 0 \end{pmatrix}}_A z_t, \quad (\text{F.38})$$

where $A_x \equiv D_x h(0)$ and $A_\mu \equiv D_\mu h(0)$. Note that ρ and 0 are eigenvalues of A , and let v_ρ and v_0 denote corresponding eigenvectors. Let \mathcal{E} be a 4-dimensional invariant subspace of A such that $v_\rho, v_0 \in \mathcal{E}$; that is, (a) $\mathcal{E} = \left\{ \sum_{j=1}^4 \alpha_j z_j : \alpha_j \in \mathbb{R}, j = 1, \dots, 4 \right\}$ for some linearly independent basis vectors $z_j \in \mathbb{R}^5$, (b) $v_\rho, v_0 \in \mathcal{E}$, and (c) if $x \in \mathcal{E}$ then $Ax \in \mathcal{E}$. Note that, as in the non-stochastic case, there are at most three possible such subspaces.

Given such an \mathcal{E} , we look for a candidate \mathcal{M} (or, equivalently, the associated projection function ϕ) in a similar way to the non-stochastic case. That is, we seek the 4-dimensional manifold \mathcal{M} with the following properties: (a) $\ell_t = \phi(y_t, \mu_t; \sigma)$ for all t satisfies the expression (F.37); (b) \mathcal{M} is tangent to \mathcal{E} at the non-stochastic steady state $z_t = 0$; and (c) the function ϕ is analytic on a neighborhood of the non-stochastic steady state. After obtaining the candidate \mathcal{M} for each possible \mathcal{E} , we can then check numerically whether the stochastic TVC is satisfied for exactly one of these candidate \mathcal{M} 's, in which case we have found the desired solution.

⁷⁶The reason we introduce σ_t as a (degenerate) state variable should become apparent shortly.

As in the non-stochastic case, we cannot in general find ϕ analytically. However, from (F.37), we may obtain that ϕ implicitly solves

$$\mathbb{E}[\phi(f(y, \phi(y, \mu; \sigma), \mu; \sigma), \rho\mu + \sigma\eta; \sigma)] = g(y, \phi(y, \mu; \sigma), \mu; \sigma), \quad (\text{F.39})$$

where the expectation on the left-hand side of (F.39) is taken over realizations of the i.i.d. $N(0, 1)$ random variable η . Note also that the other time-varying variables in this expression (i.e., y and μ) are determined independently of η . We can thus easily solve for the k -th-order Taylor approximation to ϕ around the non-stochastic steady state $z = 0$ in a manner similar to the non-stochastic case by sequentially differentiating the expression (F.39) with respect to the vector $(y, \mu; \sigma)$.

F.2 Estimation Procedure

To estimate the model, we use an indirect inference method as follows. Let $x_t \in \mathbb{R}^n$ denote a vector of date- t observations in our data set, $t = 1, \dots, T$, and let $\mathbf{x}_T \equiv (x'_1, \dots, x'_T)'$ denote the full data set in matrix form. Let $F : \mathbb{R}^{T \times n} \rightarrow \mathbb{R}^q$ be the function that generates the q -vector of features of the data we wish to match (i.e., $F(\mathbf{x}_T)$ is a vector containing all relevant spectrum values, plus the correlation, skewness and kurtosis for hours and the risk premium).

Suppose we simulate M data sets of length T from the model using the parameterization θ . Collect the m -th simulated data set in the matrix $\tilde{\mathbf{x}}_T^m(\theta) \in \mathbb{R}^{T \times n}$, $m = 1, \dots, M$. The estimation strategy is to choose the parameter vector θ to minimize the Euclidean distance between $F(\mathbf{x}_T)$ and the average value of $F(\tilde{\mathbf{x}}_T^m(\theta))$, i.e., we seek the parameter vector

$$\hat{\theta} = \underset{\theta \in \Theta}{\operatorname{argmin}} \left[F(\mathbf{x}_T) - \frac{1}{M} \sum_{m=1}^M F(\tilde{\mathbf{x}}_T^m(\theta)) \right]' \left[F(\mathbf{x}_T) - \frac{1}{M} \sum_{m=1}^M F(\tilde{\mathbf{x}}_T^m(\theta)) \right],$$

where Θ is the parameter space. In practice, we simulate $M = 3,000$ data sets for each parameter draw, and estimate $\hat{\theta}$ using MATLAB's `fminsearch` optimization function.

F.2.1 The Parameter Space

We estimate the nine parameters of the model imposing several restrictions on the parameter space Θ . First, we require that the habit parameter γ and durables-share parameter ψ be non-negative and less than one, i.e., $0 \leq \gamma, \psi < 1$. Second, we require that the policy rate reacts positively to expected hours, but not so strongly as to cause current hours to fall in response to an increase in expected hours, i.e., $0 < \phi_\ell < 1/\omega$. Third, we impose that $\tilde{R}'(0) \leq 0$ (i.e., we have complementarity near the steady state), but that the degree of complementarity is never so strong as to generate static multiple equilibria.⁷⁷ This latter property is ensured if the function $\hat{\ell} + \frac{\omega}{\kappa} \tilde{R}(\hat{\ell})$ is strictly increasing in $\hat{\ell}$ (so that it is invertible), which requires $0 \geq \tilde{R}_1 > -\frac{1+\gamma-\psi}{\omega}$ and $\tilde{R}_3 > \frac{\omega \tilde{R}_2^2}{3(1+\gamma-\psi+\omega \tilde{R}_1)}$. Fourth, we impose that the shock process is stationary, i.e., $|\rho| < 1$. Finally, we require that the parameters be such that a solution to the model exists and is unique (see Appendix F.1). None of the estimated parameters is on the boundary of the set of constraints we have imposed.

⁷⁷By static multiple equilibria, we mean a situation where, for a given \hat{X}_t , $\hat{\ell}_{t-1}$ and expectation about $\hat{\ell}_{t+1}$, there are multiple values of $\hat{\ell}_t$ consistent with the dynamic equilibrium condition.