Putting the Cycle Back into Business Cycle Analysis

By Paul Beaudry and Dana Galizia and Franck Portier*

Are business cycles mainly a response to persistent exogenous shocks, or do they instead reflect a strong endogenous mechanism which produces recurrent boom-bust phenomena? In this paper we present evidence in favor of the second interpretation and we highlight the set of key elements that influence our answer. The elements that tend to favor this type of interpretation of business cycles are: (i) slightly extending the frequency window one associates with business cycle phenomena, (ii) allowing for strategic complementarities across agents that arise due to financial frictions, and (iii) allowing for a locally unstable steady state in estimation.
JEL: E3, E32, E24
Keywords: Business Cycle, Limit Cycle

Market economies repeatedly go through periods where, for sustained periods 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. The types of forces and mechanisms that drive these fluctuations remain a highly debated subject. As an organizational framework, two conceptual views are worth distinguishing. On the one hand, there is the view that business cycles are primarily driven by persistent exogenous shocks. In models reflecting this view, such shocks are generally propagated through a variety of endogenous mechanisms, including those that may ultimately prolong their effects such as adjustment costs, capital accumulation, and habit persistence. However, it generally remains the case that, if the persistence of the exogenous shocks in such models were substantially reduced, business cycle-type fluctuations would largely disappear. As a result, such models can be viewed as supporting the notion that persistent exogenous shocks are central to understanding business cycles. On the other hand, there is

* Beaudry: Bank of Canada, 234 Wellington Street, Ottawa, ON, K1A 0G9, Canada (email: PBeaudry@bank-banque-canada.ca); Dana Galizia: Department of Economics, Carleton University, Department of Economics, Loeb Building, 1125 Colonel By Drive, Ottawa, ON, K1S 5B6, Canada (email: dana.galizia@carleton.ca); Franck Portier: Department of Economics, University College London, 30 Gordon Street, London, WC1H OAX, United Kingdom and Center for Economic Policy Research (email: f.portier@ucl.ac.uk). The authors thank Jess Benhabib, Claudio Borio, Doyne Farmer, Greg Kaplan, Vadim Marmer, Kiminori Matsuyama and Morten Ravn for helpful discussions, the participants to the many seminars and conferences where this work was presented for helpful feedback. We are grateful to the editor in charge Gita Gopinath and the referees for having provided guidance and comments that led to substantial improvements in our work. Natasha Kang and Mengying Wei at the Vancouver School of Economics provided excellent research assistance. All remaining errors are our own. Paul Beaudry acknowledges research support from the SSHRC of Canada. Franck Portier acknowledges financial support from the ADEMU project, "A Dynamic Economic and Monetary Union," funded by the European Union's Horizon 2020 Program under grant agreement No 649396. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

an alternative view wherein the bulk of business cycle fluctuations is believed to be the result of forces that are internal to the economy and that endogenously favor recurrent periods of boom and bust. According to this alternative view, even if shocks to the economy are not persistent, large and/or prolonged business cycles nevertheless arise due to the equilibrium incentives present in a decentralized economy.[1]

Even though from a theoretical point of view both of the above possibilities are reasonable, there are at least two important empirical reasons why mainstream macroeconomics has broadly coalesced around the first class of explanations. The first reason is based on the estimation of a vast array of structural models that allow for internal propagation mechanisms and exogenous driving forces to compete in explaining observed fluctuations. Overwhelmingly, the results of such estimations suggest that persistent exogenous driving forces are required to explain the data, with estimated internal propagation mechanisms being far too weak to explain business cycle fluctuations without them.[2] The second reason is based on reduced-form evidence on the spectral properties of many macroeconomic aggregates. Since Granger (1966) and Sargent (1987), it has been argued that the autocovariance patterns in the data are generally not supportive of strong internal boom-bust mechanisms, which would typically imply pronounced peaks in the spectral densities of macro aggregates at business cycle frequencies. However, such peaks, it has been argued, do not appear to be present in the data.

The objective of this paper is to provide new impetus to the second class of explanations noted above. In particular, we provide both reduced-form and structural evidence in support of this view, and further show why certain procedures that are commonly used in macroeconomic research may have biased previous research against it. To this end, we proceed in three steps. The first step is purely empirical. We examine anew the spectral density properties of many trend-less macroeconomic aggregates, such as work hours, rates of unemployment, and risk premia. We highlight a recurrent peak in several spectral densities in US macroeconomic and financial data at periodicities around 9 to 10 years. We complement this visual inspection with a set of tests aimed at documenting the statistical significance of this local peak. While the presence of such a peak does not necessarily imply strong endogenous cyclical forces,[3] it is an important first step in

[1] While we present these two frameworks as distinct, there is actually a continuum between the two. To get a sense of whether a business cycle model relies more on persistent shocks versus internal propagation mechanisms in explaining the data, one could compute the fraction of some data feature explained by the model—for example, the auto-correlation or variance of a variable—that is lost if one reduces the persistence of the shocks to zero. If this fraction is close to 100%, then the model can be said to rely primarily on persistent exogenous shocks, while if this fraction is close to zero, we can say that the model relies primarily on endogenous mechanisms. Frisch (1933) first introduced this distinction between the "*propagation* problem" and the "*impulse* problem".

[2] Certain credit cycle models, such as the seminal piece by Kiyotaki and Moore (1997), have internal propagation mechanism that can potentially be very strong. However, when such models are estimated, the implied parameters generally do not generate quantitatively meaningful endogenous cyclical behavior. Instead, the estimated versions of these models most often maintain a reliance on persistent exogenous shocks to explain business cycles features. The current paper offers insights to why this may be the case.

[3] One could in principle obtain such patterns with strongly cyclical exogenous forces.

our argument, as it runs counter to the notion that the spectral properties of the data rule out such a possibility. In addition to providing motivation for the following sections, these spectral densities play a central role in our later structural estimation exercises.

In a second step, we present a simple framework aimed at highlighting key economic elements that favor the endogenous emergence of peaks in the spectral density of equilibrium outcomes. In particular, we focus on understanding the types of market forces, i.e., the types of interactions between agents—that would give rise to cyclicality and spectral peaks in situations where the individual-level decisions on their own do not favor them. The framework emphasizes that there are two main ingredients that, in combination, favor spectral peaks: allowing for strategic complementarities across agents, and having some accumulation (i.e., stock) variable that, when high, tends to depress the individual value of accumulating more. We then discuss the extent to which such forces are present in existing models and use these properties to motivate the type of model we bring to the data in the subsequent section. We also discuss the connection between models that are capable of producing spectral peaks and models in which limit cycles may emerge.

In the third step, we present a dynamic stochastic general equilibrium model that builds on a simple New Keynesian model in a manner that incorporates the two elements discussed above in the second step. In particular, the model contains real-financial linkages that produce complementarities, and an accumulated capital stock that exhibits diminishing returns. In the model, unemployment risk, default risk, and risk premia on borrowing are jointly determined. The model aims to capture the common narrative of an accumulation-credit cycle, wherein booms are periods in which banks perceive lending to be safe and risk premia are correspondingly low, which allows households to spend more on durable goods and housing, which in turn contributes to making lending safe by keeping unemployment low.[4] The model allows for endogenous propagation forces that can potentially generate boom-bust cycles as the unique equilibrium outcome.[5] The model also features a persistent exogenous driving force, and in estimating the model we allow this stochastic force to compete with the endogenous forces in order to evaluate their respective roles. We use this model to illustrate how, depending on the estimation approach, one may draw very different conclusions regarding the relative importance of exogenous and endogenous forces in driving business cycle movements. In particular, we show that if one targets the spectral densities documented in our first step, and if one adopts an estimation method

[4]It should be emphasized that there are many possible sources of complementarity that could be embedded in the general class of models we describe in our second step and that could potentially generate similar dynamics to our accumulation-credit cycle model. While we focus on credit market frictions since we believe they likely play a role in business cycle fluctuations (see Section I.D), we by no means wish to rule out the possibility of other relevant sources of complementarity.

[5]Note that the endogenous boom-bust cycles that arise in our model do not reflect multiple equilibria or indeterminacy.

that allows for a locally unstable steady state, then the estimation results suggest that business cycles are mainly driven by internal forces buffeted by temporary shocks. In particular, our point estimates in this case actually suggest the presence of stochastic limit cycles, where the stochastic component is essentially an *i.i.d.* process.[6] In contrast, if we use more standard estimation techniques, if we focus less on explaining business cycle properties, or if we restrict the presence of complementarities, then we find more evidence in favor of persistent exogenous forces being the main driver of business cycles.

While we view our results as providing novel support for the idea that business cycles may be largely driven by endogenous cyclical forces, we do not claim that these results—which are based on the estimation of only one model—constitute a definitive answer to this question. Instead, we believe an important contribution of the paper is to illustrate why the current consensus regarding the importance of persistent exogenous shocks may reflect some arbitrary choices. In particular, we highlight how different estimation targets and methods can greatly influence one's conclusions regarding the relative contribution of internal versus external mechanisms in driving business cycles. For example, we show how, when adopting an estimation approach that allows for stochastic limit cycles, one finds little need for persistent exogenous shocks. We hope that these results will motivate exploration into other business cycle models in order to assess how sensitive inferences regarding the relative importance of internal versus external propagation are to three factors: focusing on frequencies that are slightly lower than the traditional focus of business cycle analysis, allowing for the possibility of locally unstable steady states, and finally, not unduly restricting strategic complementarities across agents. Our results provide a clear—even extreme—example of these sensitivities where, depending on how these factors are treated, the economy may appear to be driven primarily by exogenous shocks, or may instead appear to be driven primarily by endogenous mechanisms.

It is important to note that the idea that macroeconomic fluctuations may predominantly reflect strong internal propagation mechanisms, and even the possibility of 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).[7] 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

[6]While the idea that business cycles may possibly reflect stochastic limit cycle forces is not new, our empirical finding of support for such a view within the confines of a stochastic general equilibrium model with forward-looking agents appears unprecedented.

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

cycle fluctuations greatly diminished and became quite removed from the mainstream research agenda.[8] There are at least two reasons for this. First, if the economy exhibited a deterministic limit cycle, the cycles would be highly regular and predictable, which is inconsistent with the data. Second, the literature on limit cycles has generally made neither a clear empirical nor a strong theoretical case for their relevance. 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 the notion of stochastic limit cycles, wherein the system is buffeted by 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 now make this a tractable endeavor.[9]

The remaining sections of the paper are organized as follows. In Section I we document the spectral properties of U.S. data on hours worked, unemployment, and several indicators of financial conditions. These properties motivate our analysis and will be used later on in estimating our model. In Section II we present a simple mechanical 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. We use this framework to highlight when complementaries are likely to produce cyclical outcomes. In Section III, we extend a standard three-equation New Keynesian model in a manner that allows for the features highlighted in the model of Section II to be present. We the report several estimations of this model to clarify what choices and restrictions would lead one to conclude that the economy is driven primarily by persistent exogenous shocks versus inferring that it is driven by a strong endogenous propagation mechanism, including the possibility of limit cycles. Finally, in the last section we offer concluding comments.

I. Motivating Observations

The objective of this section is to re-assess certain properties of U.S. business cycles.[10] In particular, we question the common view that business cycle fluctuations are best described as being a-cyclical by providing evidence suggesting that they may embed or reflect cyclical forces. Moreover, we examine whether

[8]There 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 et al. (2013)). One should also mention a large literature on endogenous business cycles under bounded rationality and learning, following early ideas of Grandmont (1998). Hommes (2013) reviews this literature and the debate on endogenous business cycles versus exogenous shocks, and particularly the role of stochastic shocks in models with limit cycles and chaos.

[9]See Galizia (2018).

[10]Our analysis is primarily focused on the U.S., but in Section D of the online Appendix we also document that the properties we emphasize are present in several other industrialized economies.

there may be concurrent cyclical forces in financial markets, as would be implied by models of real-financial linkages. To illustrate what defines our notion of cyclicality, consider a stationary zero-mean series x_t (possibly expressed as deviations from an appropriate long-run trend), and let $\gamma_q \equiv \mathbb{E}[x_t x_{t+q}]$ be its auto-covariance function (ACF). Compare the following two cases. In the first, suppose that $\gamma_q > 0$ for all q, which would, for example, be the case if x_t follows a simple AR(1) process with positive auto-regressive parameter. In this case, positive values of x will, on average, be followed by positive values in each subsequent period. We refer to this x_t as being a-cyclical, since a large boom today is not a predictor of a large recession to come; that is, there is no sense in which a current boom is sowing the seeds of a subsequent bust. In contrast, suppose instead that there is some n for which γ_{kn} is negative when k is odd and positive when k is even. In this case, positive values of x in one period are, on average, followed by negative values n periods in the future and positive values n periods after that. We refer to such a process as being cyclical, since a large boom today is predicted to be followed by a large bust in the future, with subsequent echo effects.

As the above example illustrates, this notion of cyclicality is simply a property of the ACF: if the ACF displays oscillations, then we will say that the series exhibits cyclicality; if it does not, then we will say it is a-cyclical. It should be emphasized that, according to this definition, cyclicality does not imply the presence of deterministic cycles of constant length, nor does it imply that the cycle is highly predictable. In particular, even if a series exhibits cyclicality, both the realized amplitude and realized length of cycles will generally be random.

There are two well known practical difficulties that arise when attempting to assess whether macroeconomic phenomena exhibit cyclicality. First, there are no generally accepted macroeconomic series that are thought to reflect only business cycle forces. Instead, most macroeconomic series are believed to reflect both business cycle forces and lower-frequency forces unrelated to business cycles (e.g., demographic factors, trend growth, etc.). Therefore, in order to evaluate business cycle properties, one needs to find a way to extract properties of the data that are unlikely to be contaminated by the lower-frequency forces that are not of direct interest. For this reason, simple plots of the ACF are not usually very informative. Second, since we are only interested in evaluating cyclical forces that are potentially attributable to business cycle phenomena, one needs to define the relevant range of periodicities—that is, cycle lengths—to focus on.

To address these two issues, we proceed in steps. First, to address the issue of cycle lengths, we exploit NBER recession dates to examine the probability distribution governing the arrival of recessions; in particular, we look at the conditional probability of the economy being in recession at time $t + k$ given that it was in recession at time t. Using this information, we propose a range of periodicities within which one should look for cyclical forces.[11] Second, we look at the spec-

[11] As we shall see, the proposed range of periodicities suggested by our analysis is slightly longer than that usually associated with business cycle phenomena. However, given that our conclusions come directly

tral densities of a set of non-trending variables that can be reasonably taken to be stationary (based on unit root tests). For example, this is the case for many labor market variables—such as hours worked per capita, employment rates, job finding rates, and unemployment rates—and many financial variables. We focus on such series since it is not necessary to transform them before looking at their spectral densities. This is especially attractive since standard data transformations can spuriously induce the appearance of cyclicality in a-cyclical data. The reason we focus on spectral densities, meanwhile, is that they are a representation of the ACF that allows the cyclical properties of a series to emerge even in the presence of low-frequency confounders. Specifically, a local peak (or hump) in the spectrum of a series would be an indicator that it exhibits cyclicality at that periodicity. Thus, to evaluate whether a given series displays business-cycle-relevant cyclicality, we check whether its spectral density exhibits a peak within the range of business cycle periodicities (i.e., those suggested by the NBER recession dates). In addition to simply plotting spectral density estimates, which allows for an easy visual inspection for any peaks, we provide formal tests to evaluate whether any such peaks are statistically significant.

A. Conditional Probabilities of NBER Recession Dates

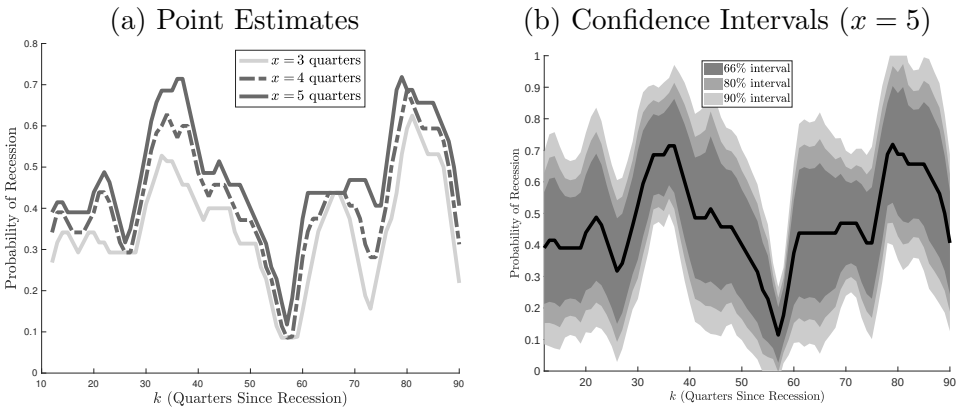
In this subsection we use NBER recession dates to suggest the range of periodicities that may be associated with cyclicality in macroeconomic data.[12] One of the attractive features of NBER recession dates is that they are a synthesis of a number of different variables that are, by design, associated with standard ideas about business cycle activity. In panel (a) of Figure 1, we plot the probability that the economy will be declared by the NBER dating committee to be in a recession at some point in an x-quarter window around time $t + k$, given that it is in a (NBER) recession at time t. The Figure plots this probability as we vary k between 12 and 90 quarters, using all NBER recession dates from 1946Q1 to 2017Q2. We look in a $2x$-quarter window around date $t+k$ since NBER recessions are rather short-lived, showing results for different window widths ranging from $x = 3$ to $x = 5$ quarters. We start at $k = 12$ to ensure that the recession under way at t is excluded from the window.

As can be seen in panel (a) of Figure 1, regardless of the window width, a rather clear pattern emerges. The probability that the economy is in a recession k quarters from now, given that it is in a recession currently, increases as k goes from 12 quarters up to around 36-40 quarters, then decreases until around 56-60 quarters, at which point it starts increasing again. This pattern, which can be roughly approximated by a sine wave with a period of around 38 quarters, suggests the possibility that business cycles may reflect cyclical elements at a periodicity

from recession dates, it would be difficult to argue that the phenomena we capture are somehow distinct from the business cycle (e.g., that they may better be thought of as some alternative "medium-run" cyclical phenomena as, for example, in Comin and Gertler (2006)).

[12]Sources for all data series are provided in Section A of the online Appendix.

FIGURE 1. CONDITIONAL PROBABILITY OF BEING IN A RECESSION



Note: Panel (a) displays the fraction of time the economy was in a recession within an x-quarter window around time $t+k$, conditional on being in a recession at time t, where x is allowed to vary between 3 and 5 quarters. Panel (b) shows confidence intervals for the $x = 5$ case. See Section B of the online Appendix for the $x = 3$ and $x = 4$ cases, as well as for details of how these confidence intervals were constructed. The figure was constructed using NBER recession dates over the period 1946Q1-2017Q2.

of around 9 to 10 years. Particularly interesting is the fall in the probability of a recession after 9-10 years, and the subsequent increase after reaching a minimum at around 14-15 years. Panel (b) of the Figure plots confidence intervals for the $x = 5$ case, which confirms that the general pattern we highlighted in panel (a)—namely, a local peak in the probability for $k \approx$ 36-40 quarters, followed by a local trough around 56-60 quarters—is unlikely to have occurred simply by chance.[13] In Section B of the online Appendix we present confidence intervals for the $x = 3$ and $x = 4$ cases, which yield similar conclusions.

We take the above observations as the basis for formulating the hypothesis that post-war business cycles may contain cyclical elements that express themselves at a periodicity roughly between 36 and 40 quarters. It should be noted that these periodicities are slightly longer than the ones commonly associated with business cycle phenomena, as first proposed by Burns and Mitchell (1946).[14] This may

[13]See Section B of the online Appendix for details of how these confidence intervals were constructed. In that Appendix, we also provide a joint test of the null hypothesis that the conditional probability is constant in k, and find that the null is rejected at conventional level of significance.

[14]The traditional definition of the business cycle focuses on movements in macroeconomic variables at periodicities between 6 and 32 quarters. According to Baxter and King (1999) and Stock and Watson (1999), the reason was that the NBER chronology lists thirty complete cycles since 1858, with the shortest full cycle (peak to peak) being six quarters, and the longest thirty-nine quarters and 90% of these cycles being no longer than 32 quarters. While this definition may have seemed appropriate 30 years ago, it appears overly restrictive now given the more recent NBER cycle dates. For example, the cycle in the 1990s lasted 43 quarters from the peak in July 1990 to the subsequent peak in March 2001. Similarly, the cycle that started from the peak in 2007 has lasted more than 40 quarters so far, having apparently not yet reached another peak. For this reason—and this will be supported by our spectral evidence below—we argue that the definition of the business cycle should include fluctuations up to periodicities of at least 40 quarters, and maybe even up to 50 quarters.

reflect the fact that Burns and Mitchell (1946) were focusing on pre-war data, where recessions were more frequent. The pattern in Figure 1 suggests that it may not be appropriate to focus business cycle analysis only on fluctuations lasting less than 8 years, as is often taken as a benchmark, but rather a wider window that includes periodicities up to at least 10 years. In section D of the online Appendix, we document that a similar pattern to that observed in Figure 1 is also present in most other G7 countries. This may not be too surprising, as this set of countries often share recessions and expansions.

B. Looking for a Peak in Spectral Densities

We now examine the spectral properties of a set of non-trending U.S. macroeconomic variables. As noted previously, one potential way of describing the cyclical properties of stationary data is to focus on the spectral density, which depicts the importance of cycles of different frequencies in explaining the data. 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. The traditional view, as expressed for example in Granger (1966) and Sargent (1987), is that most macroeconomic time series do not exhibit peaks in their spectral densities at business cycle frequencies. This view accordingly suggests that business cycle theory should not seek to explain macroeconomic fluctuations as cyclical phenomena.[15] In this section, we re-examine the validity of this consensus relative to the alternative view in which business cycles may reflect cyclical forces that express themselves at a periodicity of around 36-40 quarters (i.e., the timing motivated by our above analysis of conditional recession probabilities).

The main challenge faced in using spectral methods relates to the long-run properties of the data. In particular, spectral densities are only well-defined for stationary data, while many macroeconomic variables are non-stationary. It has therefore been common to substantially transform (i.e., de-trend) non-stationary variables before looking at their spectral properties. However, if a variable is thought to be the sum of a stationary cyclical component and a non-stationary trend component, there is no theory-free way of isolating the cyclical component. Moreover, it is well known that de-trending procedures can create spurious cycles. Thus, the use of spectral methods is most informative if it can be applied to series that can plausibly be argued to be stationary prior to any transformation.[16] For example, this may be the case for many labor market variables, such as the employment and unemployment rates. For this reason, we begin by examining

[15] In light of this, it is generally agreed upon (see e.g. Sargent (1987)) that business cycle research should focus mainly on explaining the co-movement properties of macro variables, as there is substantial co-movement across variables at business cycle frequencies. It is worth emphasizing that we do not question here the view that most macroeconomic aggregates co-move substantially over the business cycle, only the view that cyclical phenomena are not present as well.

[16] If a series is non-stationary but known to be $I(1)$, then in principle one can look at the spectrum of the first difference of the series. However, since the first-difference filter heavily emphasizes movements at the highest frequencies and de-emphasizes those at lower frequencies, doing so may substantially mask the properties of any cyclical component that may be present at somewhat lower frequencies.

the spectral densities of labor market variables. Even if we accept that such labor market indicators can be thought of as stationary, however, we nevertheless believe that such measures most likely reflect both business cycle and longer-run phenomena. For example, demographic factors, which are typically thought of as being quite distinct from business cycle movements, tend to produce low-frequency movements in employment patterns. For this reason, we focus on the shape of a spectral density at periodicities shorter than 60 quarters, with the rationale that fluctuations at periodicities longer than 60 quarters most likely reflect factors distinct from business cycle phenomena.

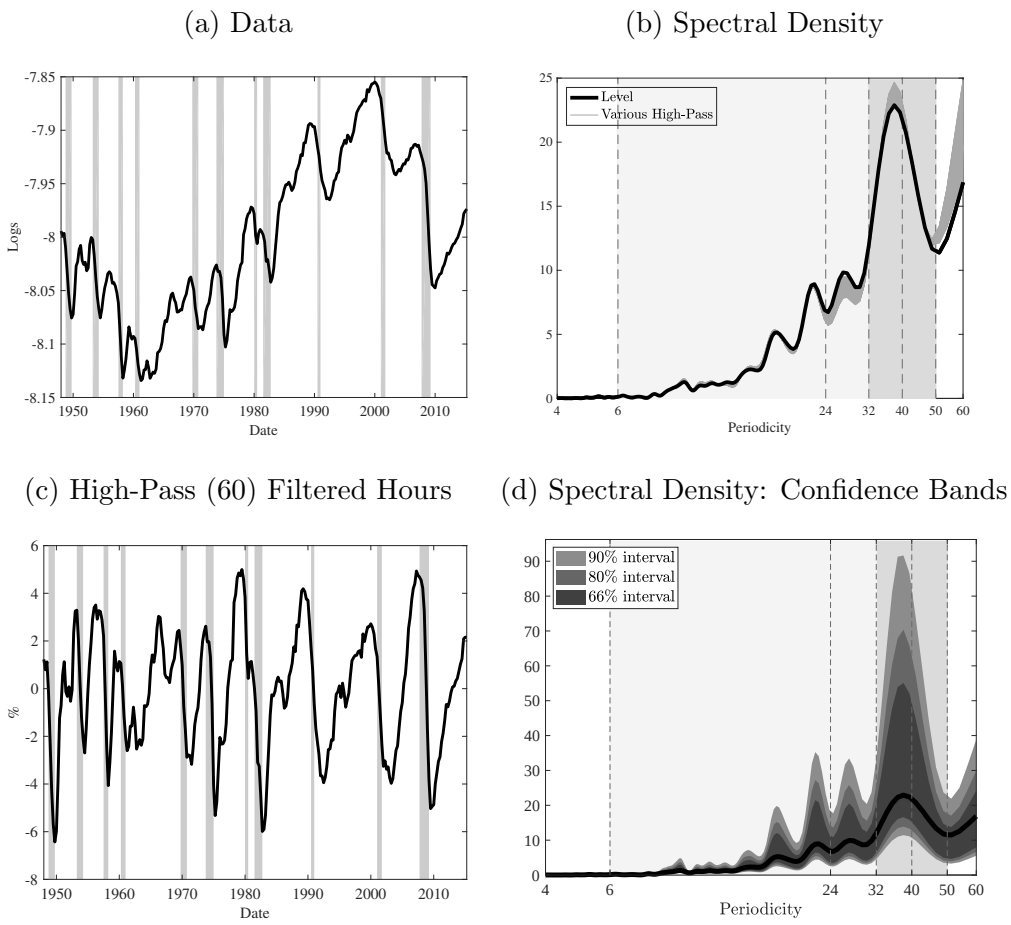
We begin by examining the properties of the (log) of U.S. non-farm business (NFB) hours worked per capita from 1948Q1-2015Q2. This series is plotted in panel (a) of Figure 2. As the Figure shows, hours exhibited substantial fluctuations over the sample period, but there is little evidence of any long-run trend. For this reason, it seems plausible to treat this series as stationary, as is also confirmed formally by Dickey-Fuller tests.[17] Accordingly, we begin by looking directly at the spectral density of this series without any prior transformation (except de-meaning). The dark line in panel (b) of Figure 2 plots this spectral density over the range of periodicities from 4 to 60 quarters.[18] Since it is common in macroeconomics to try to remove very low-frequency movements—that is, movements at frequencies much lower than business cycle frequencies—for comparison we also plot on the same axes the spectra obtained when first passing the series through a high-pass filter. In particular, each gray line in the Figure represents the spectral density of the series after it has been transformed using a high-pass filter that removes fluctuations with periodicities greater than P quarters in length, where P ranges from 100 to 200. The results suggest that the spectral properties of hours at periodicities around 36-40 quarters—the range of periodicities that interest us most—are, as we would expect, invariant to whether or not one first removes very low-frequency movements from this series.

What does panel (b) of Figure 2 reveal about the cyclical properties of hours? To us, the dominant feature is the distinct hump in the spectral density surrounding the local peak at around 38 quarters, the bulk of which is contained in the 32- to 50-quarter range. This hump is much more pronounced than anything found at periodicities less than 32 quarters, which suggests that a significant proportion of the fluctuations in hours may come from some cyclical force with a periodicity of about 9-10 years, precisely in line with our earlier analysis using conditional recession probabilities. To make it clear that this is not simply a coincidence,

[17] Specifically, we perform an augmented Dickey-Fuller test on the series. Akaike, Bayesian, and Hannan-Quinn information criteria all suggest using a single lag. With this specification, the presence of a unit root is rejected at 5% significance.

[18] We 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 C in the online Appendix for more details on spectral density estimation).

FIGURE 2. PROPERTIES OF HOURS WORKED PER CAPITA



Note: 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 (gray lines). A high-pass (P) filter removes all fluctuations of period greater than P. Panel (c) displays high-pass (60) filtered hours. Panel (d) shows bootstrapped 66%, 80% and 90% point-wise confidence bands for the spectral density.

and that the fluctuations in hours in this range of frequencies should indeed be thought of as business cycle phenomena, in panel (c) of Figure 2 we plot the hours series after having removed fluctuations longer than 60 quarters using the high-pass filter. In panels (a) and (c), we have also highlighted NBER recessions in gray. Unlike in panel (a), where there are clearly significant long-run fluctuations in hours that are unrelated to the NBER recessions, the fluctuations remaining in the de-trended hours series of panel (c) correspond very closely to the standard narrative of the business cycle.[19] In fact, to the extent that there are fluctuations in panel (c) that do not correspond to NBER recessions, they are evidently due to unrelated movements much *shorter*, not longer, than the 32-50 quarter range in which the spectral hump appears.

Testing For a Local Peak

The local peak in the spectral density of hours worked observed in panel (b) of Figure 2 suggests that the labor market may be subject to cyclical forces that play out at periodicities around 38 quarters. However, as illustrated by the point-wise confidence bands shown in panel (d) of Figure 2,[20] there is substantial uncertainty surrounding our estimates of the spectral density. For this reason, it is desirable to examine the statistical significance of this observed local peak.[21]

To this end, we test a shape restriction on the spectral density; namely, that it is hump-shaped over a certain range of frequencies. The choice of these frequencies is dictated by the pattern of conditional probabilities of NBER recessions that we highlighted in Figure 1, which suggests important cyclical forces in the 32-40 quarter range. We therefore test for a peak in the following way. We select three intervals of frequencies: a "lower trough range" interval Ω_{T_1}, which contains frequencies corresponding to cycles 16-32 quarters long; a "peak range" Ω_P for 32-40 quarters; and an "upper trough range" Ω_{T_2} for 40-60 quarters. We report results under two different null hypotheses. The first null hypothesis is that the spectral density is flat over the region in question. While we do not believe this null hypothesis is especially pertinent, we provide results for it for the sake of completeness. We refer to this as the flat null. The second and, in our view, more relevant null hypothesis that we consider is one motivated by the influential work of Granger (1966), which argued that the spectral densities of most macroe-

[19] For the typical "RBC" moments (i.e., output, consumption, investment and hours standard deviations, standard deviations relative to the output one, autocorrelations, correlations with output and hours-average labor productivity correlation), we obtain the same overall patterns with Band-Pass(6,32) or (6,50) filtered data.

[20] We use a bootstrap procedure to compute these confidence intervals. See Section C.4 in the online Appendix for details.

[21] Note that confidence bands (including those in panel (d) of Figure 2) are insufficient for this task, since "peaked-ness" is inherently a property of the size of the spectrum at one group of frequencies relative to another. For example, confidence bands can tell you about the likelihood that, due purely to sampling error, the point estimate at some frequency is high. They cannot tell you, however, about the likelihood that, due to sampling error, the point estimate is both high at one frequency and *simultaneously* low at some other frequency. Since the latter is precisely the type of event that could produce a spurious peak in the estimated spectral density, some other approach is needed in order to evaluate its likelihood.

conomic variables have shapes similar to those of persistent $AR(1)$ processes, an argument that has since become the consensus view. Accordingly, the second null hypothesis is that the spectral density is that of the $AR(1)$ process that best fits the data.[22] We refer to this as the $AR(1)$ null. We then test whether the spectral density over the interval Ω_P is higher relative to the Ω_{T_1} and Ω_{T_2} intervals than is predicted by the null.

Formally, let X be the series under consideration and $\gamma_k = \mathrm{Cov}\,(X_t, X_{t-k})$ its autocovariance function. The associated spectral density function is

$$f(\omega) \equiv \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} e^{-ik\omega}\gamma_k \ ,$$

and the periodogram for a T-period sample is given by

$$I(\omega) \equiv \frac{1}{2\pi T} \left| \sum_{t=1}^{T} X_t e^{-i\omega t} \right|^2 \ .$$

Let $\omega_j \equiv 2\pi j/T$, $j = 0, \ldots, T^*$, $T^* \equiv \lfloor T/2 \rfloor$, denote the Fourier frequencies. A useful statistical result (Hamilton (1994), p. 164) is that, asymptotically,

$$R(\omega_j) \equiv \frac{2I(\omega_j)}{f(\omega_j)} \overset{i.i.d.}{\sim} \chi^2_2 \quad , \quad \omega_j \notin \{0, \pi\} \ , \tag{1}$$

where *i.i.d.*-ness is across different j's. This result implies that

$$D_q \equiv \frac{1}{m_q} \sum_{\omega_j \in \Omega_q} R(\omega_j) \sim \frac{1}{m_q} \chi^2_{2m_q} \ ,$$

and that the statistic D_q is *i.i.d.* across $q \in \{P, T_1, T_2\}$, where m_q is the number of Fourier frequencies in Ω_q. Our shape test is then as follows: conditional on a null for the spectral density f (either flat or derived from an $AR(1)$), we compute the observed value of D_q, denoted d_q. For the frequencies $\omega_j \in \Omega_{T_1} \cup \Omega_P \cup \Omega_{T_2}$, the "flat" null corresponds to

$$f(\omega_j) = \bar{f} \ ,$$

and the $AR(1)$ null, with the autocorrelation parameter ρ and innovation variance σ^2, corresponds to

$$f(\omega_j) = \frac{\sigma^2}{2\pi g(\omega_j; \rho)} \ ,$$

[22] We estimate the $AR(1)$ parameters by OLS.

where $g(\omega;\rho) \equiv 1+\rho^2-2\rho\cos(\omega)$. The p-value we seek is then

$$p = \Pr\left\{\frac{D_P}{D_{T_1}} \geq \frac{d_P}{d_{T_1}} \text{ and } \frac{D_P}{D_{T_2}} \geq \frac{d_P}{d_{T_2}}\right\}.$$

Since the joint distribution for D_P/D_{T_1} and D_P/D_{T_2} is complicated, we use a Monte Carlo approach to compute p: we randomly draw $\tilde{D}_q = \chi^2_{2m_q}/m_q$ for each q, then compute $\tilde{D}_P/\tilde{D}_{T_1}$ and $\tilde{D}_P/\tilde{D}_{T_2}$. This is repeated $N = 1{,}000{,}000$ times, and the resulting simulated distribution is used to obtain p. When we apply this shape test to the hours series from Figure 2, we obtain p-values of 1.5% for the flat null and 3.9% for the $AR(1)$ null.[23] At standard significance levels, we therefore reject both that the spectrum is flat in the relevant region, and also that it has the Granger shape.

Other Related Variables

Above, we provide evidence that aggregate hours worked per capita appears to contain an important cyclical component, as captured by the peak in its spectrum in a window near 38 quarters. It is immediately relevant to ask whether this property is representative of other cyclical measures. To this end, in Figure 3 we present estimates of the spectral density for two other closely related labor market measures (the employment rate and the unemployment rate), as well as for the investment-to-GDP ratio, which is often viewed as a cyclical indicator. The employment rate and investment ratio cover the same post-war period as our hours series, while for the unemployment rate, in the hopes of obtaining more precise estimates, we take advantage of the longer series constructed by Ramey and Zubairy (2018) that covers 1890Q1-2015Q2. We also report results for the unemployment rate for the post-war period alone. Note that, as with the hours series, we first tested each series to check that it could plausibly be treated as stationary before obtaining its spectral density.

In each case, we again see a peak in the spectral density in the same general range as observed for hours. Note that the longer unemployment series has its peak at a slightly shorter periodicity of around 35-36 quarters.[24] The second-to-fifth rows of Table 1 report, for each series, the p-values for the same hump-shape test we performed above for hours worked (reproduced in the first row of the Table). As can be seen, the flat null is clearly rejected for all series at the 5% level, while the $AR(1)$ null is rejected for three of the series at a 5% level, and

[23] Assuming the underlying data-generating process is Gaussian, the statistic $R(\omega_j)$ in (1) is distributed as χ^2_2 even in finite samples. However, independence across j's holds only asymptotically, potentially making our p-values—which were computed under this independence assumption—inappropriate. To check this, we tried an alternative bootstrap-based method to compute p-values that explicitly accounted for the finite size of our sample, and the results were nearly identical (available upon request). Thus, possible dependence across the $R(\omega_j)$'s does not appear to be driving any of our results.

[24] The observation that the cyclical component of uneployment may have a shorter periodicity in the pre-war period is not too surprising, given that it is well known that business cycles tended to be shorter on average in the pre-war period.

6% for fourth (the employment rate). We take these results as strong additional evidence that the real sector of the economy may reflect cyclical forces.[25] In Section D of the online Appendix, we document that such cyclical features also appear to be present for unemployment rates in other countries of the G7.

TABLE 1— HUMP-SHAPE (W/ PEAK AT 32-40Q) TEST (p-VALUES)

	Flat Null	$AR(1)$ Null
Real variables:		
NFB Hours	1.5%	3.9%
Employment	2.4%	5.8%
Unemployment (post-war)	1.1%	3.6%
Unemployment (long sample)	0.2%	5.0%
Investment/GDP ratio	0.6%	0.8%
Financial variables:		
BAA/FF Spread	2.4%	6.1%
NFCI	0.1%	0.1%
AAA/T-bill Spread	0.1%	0.4%
Shiller P/E Ratio	1.9%	1.9%

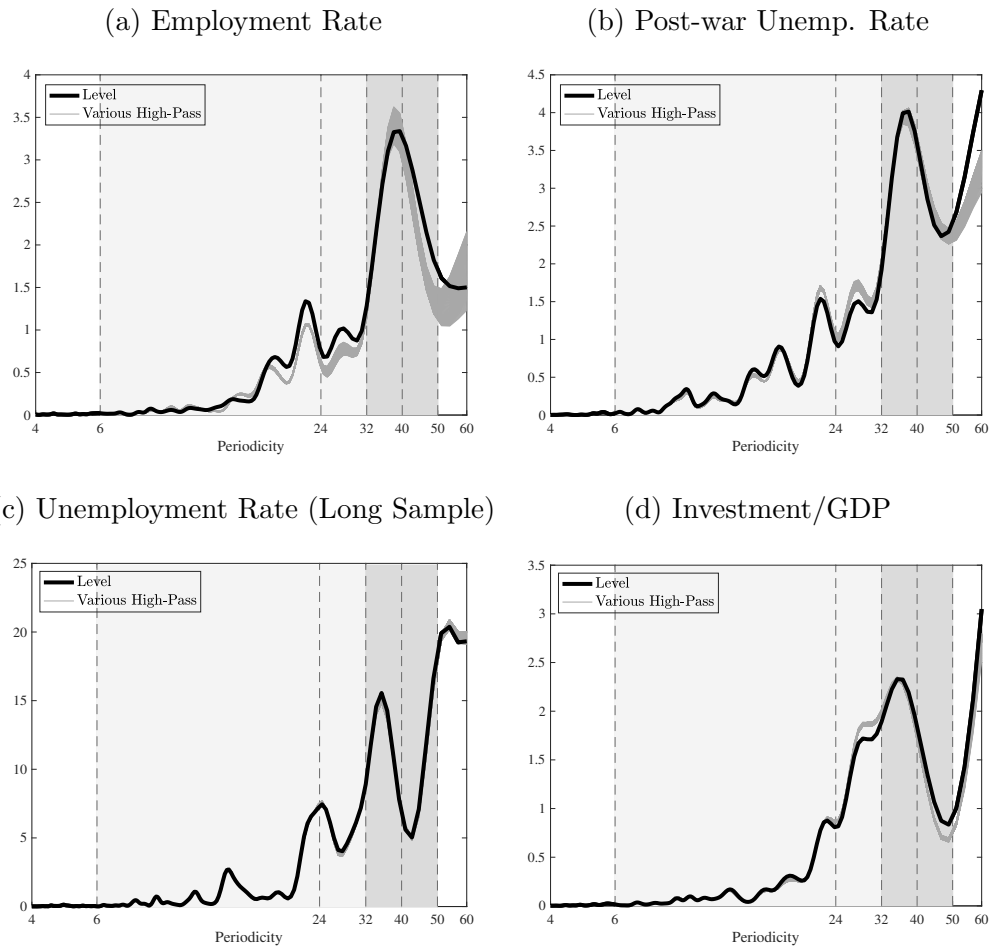
Note: Table displays simulated p-values for the hump-shape test with peak at 32-40 quarters, under both the flat and $AR(1)$ null hypotheses.

C. Spectral Implications for Hours In Standard Models

In the previous sections we have shown that several measures of labor market activity, as well as another cyclical indicator of real activity (the investment-to-GDP ratio), exhibit significant spectral peaks at a periodicity of around 38 quarters. In this subsection, we argue that this observation conflicts with the spectra implied by many modern business cycle models. To illustrate this, we report in Figure 4 the spectral density of hours implied by six models that we think span the range of quantitative models: two real models, a New Keynesian one, and three models with financial frictions. The first model is the simplest Real Business Cycle model described in Cooley and Prescott (1995). In that stripped-down model, fluctuations only come from persistent technology shocks. The second model is an RBC model augmented to include variable capital utilization, as

[25] One may also be interested in knowing whether quantity variables (such as output) also exhibit a hump in their spectral densities in a similar range to that observed for the other variables. The difficulty with such variables, however, is that they are clearly non-stationary, so that some type of transformation is needed before one can examine the spectral density. For several different measures of output per capita, we have explored using the (controversial) method of passing the series through a high-pass filter to remove the non-stationary component. For the cases we examined, we did not find much evidence of a peak in filtered output. As we discuss in Section F of the online Appendix, however, we provide evidence in support of the idea that this is likely the result of certain properties of productivity (as measured by utilization-adjusted TFP), rather than an absence of underlying cyclical forces altogether.

FIGURE 3. SPECTRAL DENSITY OF OTHER U.S. SERIES



Note: These panels display estimates of the spectral density of the employment rate (panel (a)) and unemployment rate (panel (b) and (c)) and investment/GDP ratio (panel (d)) 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 (gray lines). A high-pass (P) filter removes all fluctuations of period greater than P. For employment rate, investment/GDP ratio and post-war unemployment rate, the sample is 19489Q1-2015Q2 and it is 1890Q1-2015Q2 for the long sample. For long sample unemployment, the smoothing window width has been reduced from 13 to 9 quarters has we have roughly twice as many observations.

well as investment-specific technology shocks as a second source of fluctuations.[26] The third model is the rich New Keynesian model of Smets and Wouters (2007), which features seven shocks and a variety of real and nominal frictions. The fourth model is the financial frictions model of Carlstrom and Fuerst (1997), while the fifth is the financial-accelerator model of Bernanke, Gertler and Gilchrist (1999), as estimated on U.S. data by Christensen and Dib (2008). The sixth and final model, as proposed by Christiano, Motto and Rostagno (2014), incorporates the microeconomics of the debt-contracting framework of Bernanke, Gertler and Gilchrist (1999) into an otherwise standard monetary model of the business cycle, and features a large number of shocks, including news and risk shocks.[27] In each case, we see essentially the same thing: the spectrum does not have a peak, but rather is monotonically-increasing in the periodicity, similar to an $AR(1)$ process (i.e., it has the Granger shape). This observation should not be too surprising: as is well known, internal propagation in these models is generally quite weak, and therefore the endogenous variables largely inherit the properties of the exogenous driving forces, which in these cases are either equal to or close to $AR(1)$ processes. Although these examples are only illustrative, in our exploration of different models we have not yet found an estimated model that produces a peak in the spectral density of hours similar to the one we find in the data.

D. Is There a Related Financial Cycle?

The financial crisis of 2007 has revived interest in connecting business cycles to financial conditions.[28] As such, we now turn to looking at the spectral properties of a set of financial market indicators. In parallel to our exploration of real variables, we again restrict attention to variables that appear to be stationary (based on Dickey-Fuller tests), allowing us to directly examine their spectral densities without prior transformation.[29] In particular, we examine whether indicators of financial market conditions also exhibit peaks in their spectral densities in a range of frequencies similar to that observed for hours worked. Knowing whether this is the case will be important in helping to identify the type of model that may best explain the employment cycle we emphasized previously. To this end, we report estimates for four financial market indicators in Figure 5. Our main series of interest, which we will later also use in model estimation, is the spread between the Federal Funds (FF) rate and the rate on BAA bonds. This series runs from 1954Q3 to 2015Q2. This interest spread is a common measure of financial market risk discussed throughout the macro-financial linkage literature, and is typically interpreted as a measure of the market risk premium. The estimated spectral density for this measure of the risk premium is presented in panel (a) of Figure 5.

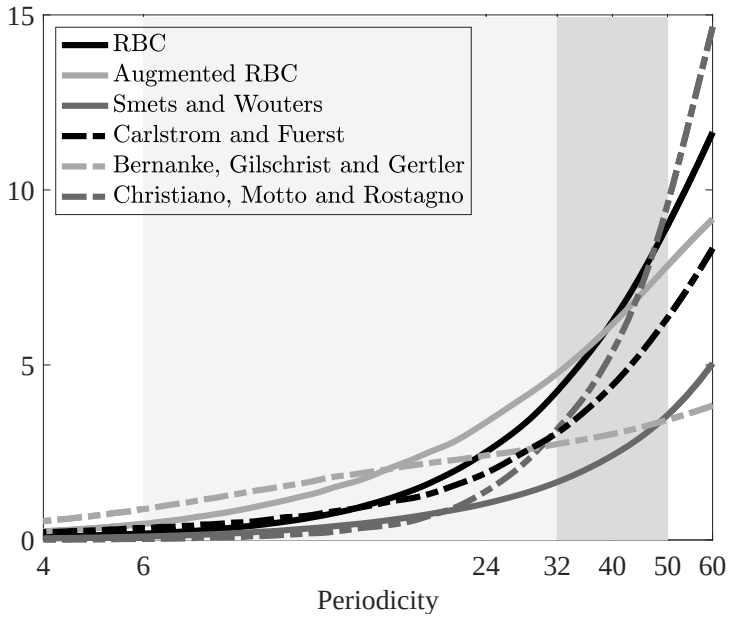
[26] This augmented model is calibrated following Fernandez-Villaverde (2016).

[27] To simulate this model, we use the DYNARE code provided in the replication package of the Macroeconomic Model Data Base (see Wieland et al. (2016)).

[28] See for example Borio (2014).

[29] We nevertheless again also show that the results are robust to first filtering the data with a high-pass filter in order to remove very low-frequency movements.

Figure 4. Spectral Densities of Hours in Some Standard Models



Note: The figure displays the mean spectral density of hours over 10,000 simulations of length 270. Models and parameters values are Cooley and Prescott (1995) for the standard RBC model, Fernandez-Villaverde (2016) for the augmented RBC (with variable capital utilization and investment specific technology shocks), Smets and Wouters (2007), Carlstrom and Fuerst (1997) and Christensen and Dib (2008) for a version of Bernanke, Gertler and Gilchrist (1999) estimated on U.S. data and Christiano, Motto and Rostagno (2014). For better visual display, all the series have been standardized to have unit variance.

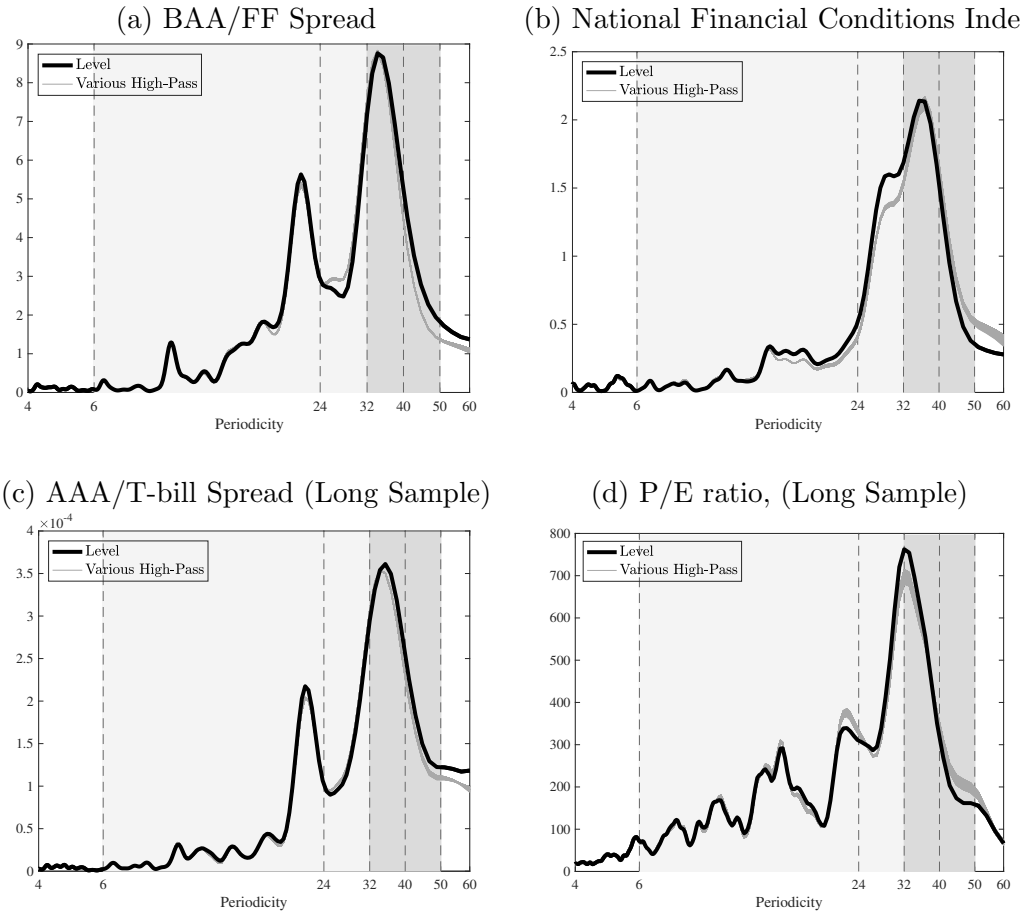
As is clear from the Figure, we again find evidence of a hump in the spectral density in the range just above the standard business cycle range, though the actual peak is located slightly below the ones found for our post-war real measures.[30]

To verify this visual impression, in the sixth row of Table 1 we report results for the same hump-shape test as performed for the real variables above. The data reject that the spectral density is flat or similar to an $AR(1)$ in the relevant range, in favor of the alternative that there is a local peak centered somewhere between 32 and 40 quarters. In Section E of the online Appendix, we also examine the co-movement between this risk premium measure and our previous hours worked measure. We find the correlation between the two to be highest (in absolute value) at precisely the frequencies emphasized by the peaks in their spectral densities. We also find that, within the 4-60 quarter range of periodicities, the spectral coherence between the two series is maximized at 38 quarters. Together, these observations suggest that the labor market and the financial market may share cyclical forces that express themselves at a periodicity of around nine years. In a later section our goal will be to explore one class of explanations for these shared cyclical patterns.

The other variables presented in Figure 5 are as follows. In panel (b) we show the spectral density for the Chicago Fed's National Financial Conditions Index (NFCI), which begins in 1971Q1. As described in Brave and Butters (2011), the NFCI, which is computed from a large sample of financial indicators, is a synthetic index between -1 and 1 that attempts to summarize financial conditions. In panel (c), we show the spectrum of the spread between the 3-month T-bill rate and the rate on AAA bonds. Relative to the series shown in panel (a), the advantage of this series is that it goes back to 1920Q1, potentially allowing for more precise estimates. In panel (d), we show the spectrum of a price-earnings ratio series that goes back to 1871. All of these series exhibit humps in their spectral densities in the range slightly beyond the standard business cycle range. The AAA risk premium and price-earnings ratio series have shapes quite similar to our baseline spread series in panel (a), while the synthetic NFCI series has a hump that is slightly more spread out than the others. In the final three rows of Table 1, we report the results of our formal peak tests for these variables. In each case, the data quite strongly reject the null hypothesis that there is no peak in the spectrum in the 32 to 40 quarter range. Thus, the overall picture that emerges from our visual and statistical evidence suggests a close link between the cycle

[30]For this interest rate spread measure there is also a second, smaller, peak near 21 quarters. This second peak, which also appears for two other spread series we consider (including the one in panel (c) of the Figure) but not in any of the other series, could be indicative of a second source of cyclicality relevant only for interest rate spreads. Alternatively, as is well known, a peak at periodicity n with a second peak at periodicity n/k for some integer k (e.g., at half the periodicity of the original) is indicative of a recurrent cyclical pattern with periodicity n, but where that pattern is not well captured by a simple sine wave. Even though 21 quarters is not exactly half of 34 quarters, one would not be able to reject the null hypothesis that, in the population spectrum, the smaller peak is at exactly half the periodicity of the larger one (i.e., that the smaller peak is simply a reflection of cyclical behavior at the larger periodicity). While this issue bears further study, for now we concentrate on the peak near 38 quarters.

FIGURE 5. SPECTRAL DENSITY OF SOME FINANCIAL SERIES



Note: Panels (a) to (d) show estimates of the spectral density for the series 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 (gray lines). A high-pass (P) filter aims to remove all fluctuations of period greater than P. All series end in 2015Q2, and start, respectively, in 1954Q3 for (a), 1973Q1 for (b), 1920Q1 for (c) and 1871Q1 for (d).

in employment and the cycle in financial market conditions. This suggests to us that, in attempting to the explain cyclical patterns in the data, one should look toward models that feature real-financial linkages.

II. Explaining Spectral Peaks: A Class of Models

In the previous section we documented the presence of a significant peak in the spectral densities of several non-trending macroeconomic variables at a periodicity near 38 quarters. In particular, we emphasized the presence of this peak in variables related to both employment and financial market conditions. The presence of such peaks suggests that the macroeconomy may embed cyclical forces that manifest themselves at a frequency slightly lower than what has traditionally been associated with business cycle fluctuations. However, as we have discussed, we believe that such forces should nevertheless be viewed as a part of the business cycle—rather than something distinct from it—since they appear intimately linked to the occurrence of recessions. The object of the next two sections is to explore mechanisms that may lie behind such phenomena.

There are at least three classes of explanations for such cyclical behavior: explanations based mainly on properties of the exogenous driving forces, explanations based primarily on the properties of individual level behavior, and, finally, explanations based on equilibrium interactions—that is, explanations where cyclical outcomes arise as the result of market interactions between individuals who, in the absence of such interactions, would not tend to make cyclical choices. Our main goal in this section is to better understand the mechanisms that may give rise to this latter possibility. To this end, we begin by examining how a hump-shaped spectral density can arise in a mechanical model where agent interactions are captured by a market-determined variable that each agent takes as given. The framework allows us to isolate simple conditions under which endogenous cyclical outcomes emerge purely as the result of equilibrium forces. We also use the framework to clarify potential challenges that may arise when attempting to explain hump-shaped spectral densities in an equilibrium framework, especially when local instability and limit cycles cannot be ruled out. In Section III below, we will explore the empirical relevance of these notions using an optimization-based forward-looking model.

A. Demand Complementarities as a Source of Cyclicality

The Environment

To understand the economic forces that may generate cyclicality, let us consider an environment with a large number N of agents indexed by j, where agent j makes a decision e_{jt}. Here, e_{jt} could represent a decision regarding an expenditure, a level of output or a level of investment. The decision rule of the individual

is assumed to take the form[31]

$$e_{jt} = \alpha_0 + \alpha_1 X_{jt} + \alpha_2 e_{jt-1} + \alpha_3 q_t + \mu_t \ , \tag{2}$$

where X_{jt} is a stock variable that could represent, for example, capital, net worth, or a habit stock, and which satisfies an accumulation equation of the form

$$X_{jt+1} = (1-\delta)X_{jt} + \psi e_{jt} \ , \qquad 0 < \delta < 1 \ . \tag{3}$$

The term μ_t in equation (2) represents an exogenous driving force,[32] and q_t represents some market-determined variable, such as a price or a matching rate. Since we allow α_1 to be positive or negative (or zero), for simplicity and without loss of generality we henceforth normalize $\psi = 1$. The inclusion of e_{jt-1} in (2), meanwhile, is intended to capture potential inertial effects. A key aspect of this setup is the interaction between individuals, as captured by the market-determined variable q_t. In particular, to capture equilibrium forces, we allow q_t to be an arbitrary function of the average behavior of others, i.e.,

$$q_t = Q\left(\frac{1}{N}\sum_{j=1}^{N} e_{jt}\right) \ , \tag{4}$$

where, following Cooper and John (1988), we will refer to the actions of agents as being (strategic) complements if $\alpha_3 Q'(\cdot) > 0$ and (strategic) substitutes if $\alpha_3 Q'(\cdot) < 0$. Compared with Cooper and John (1988), our framework embeds their static game within each period, with dynamic elements added through both the inertial forces captured by $\alpha_2 e_{jt-1}$ and the accumulation forces captured by $\alpha_1 X_{jt}$. For now we will consider Q to be a linear function of aggregate behavior, i.e., we take $q_t = \alpha_4 e_t$, where $e_t = N^{-1}\sum_j e_{jt}$. Later we will discuss the effect of allowing the function $Q(\cdot)$ to be non-linear. We also restrict attention to situations where $\alpha_3 Q'(e) < 1$ so as to rule out the possibility of multiple equilibria. Moreover, we will assume that μ_t is a stationary stochastic process of the form $\mu_t = D(L)\epsilon_t$ where $D(L)$ is a polynomial in the lag operator L, and ϵ_t are *i.i.d.* shocks with unit variance. Finally, as we are interested here in understanding the role of agent interactions in creating cyclical aggregate outcomes, we will exclude the possibility of exogenously driven cyclicality by making the following assumption:

ASSUMPTION 1: *The spectral density $s_\mu(\omega)$ of μ_t is monotonic on frequencies $\omega \in [0, \pi]$.*

Note that Assumption 1 would hold if, for example, μ followed an AR(1) pro-

[31] In Beaudry, Galizia and Portier (2016*b*), we show that most of the properties that we derive here are also present if the model includes a forward-looking component; that is, if we specify behavior as $e_{jt} = \alpha_0 + \alpha_1 X_{jt} + \alpha_2 e_{jt-1} + \alpha_5 \mathbb{E}_t e_{jt+1} + \alpha_3 q_t + \mu_t$.

[32] The analysis can be easily extended to include idiosyncratic shocks as well.

cess. We now derive conditions under which aggregate behavior in this setup could feature cyclical behavior, as captured by a hump-shaped spectral density. To see the forces at play explicitly, it is useful to write out the spectral density of e_t. Focusing on the symmetric outcome where $e_{jt} = e_t$, the evolution of e_t can be written as

$$e_t = \left(\frac{\alpha_1 + \alpha_2}{1 - \alpha_3\alpha_4} + 1 - \delta \right) e_{t-1} - \frac{\alpha_2(1-\delta)}{1-\alpha_3\alpha_4} e_{t-2} + \frac{1-(1-\delta)L}{1-\alpha_3\alpha_4}\mu_t \ ,$$

and accordingly the spectrum of e_t can be written as

$$s_e(\omega) = s_\mu(\omega) \times \frac{[1-(1-\delta)\exp\{i\omega\}]\,[1-(1-\delta)\exp\{-i\omega\}]}{(1-\alpha_3\alpha_4)^2} \times g(\omega) \ , \tag{5}$$

where $i = \sqrt{-1}$, $g(\omega) \equiv [B(\exp\{i\omega\})B(\exp\{-i\omega\})]^{-1}$ and $B(L) \equiv 1 - \left(\frac{\alpha_1+\alpha_2}{1-\alpha_3\alpha_4} + 1 - \delta \right) L + \frac{\alpha_2(1-\delta)}{1-\alpha_3\alpha_4} L^2$.

From equation (5), we see the different forces that can cause the spectrum to have a local peak. The first term represents the properties of the exogenous driving force. Assumption 1 implies that this term is monotonic, and so rules out that it can, by itself, be the source of a hump-shaped spectral density for e. It can be verified that the second term, which captures the direct impact of the capital accumulation process (3), is also monotonic on $[0, \pi]$, and therefore also cannot directly create a hump-shaped spectral density for e. This leaves the last term, $g(\omega)$, which captures the transitional dynamic forces generated by the interaction of individual decision rules and the market-determined value of q.[33] The question we now ask is: Under what conditions will agent interaction (through the endogenous determination of q) be the cause of a hump in the spectral density? To make the results as clear as possible, we focus on situations where the following assumption is also satisfied.

ASSUMPTION 2: *α_1, α_2 and δ are such that, when $\alpha_3\alpha_4 = 0$, the eigenvalues of the system[34] represented by equations* (2)-(4) *are real, positive, and smaller than 1.*

Assumption 2 implies that, if there were no forces linking agents' decisions together— as would be the case if either $\alpha_3 = 0$ or $\alpha_4 = 0$—then $g(\omega)$ would be monotonic and, accordingly, would not by itself introduce a local peak in the spectral density of e.[35] In other words, Assumption 2 rules out the possibility

[33]It should be noted that, even if each of these components were monotonic, their product could nevertheless exhibit a local peak. Such an interaction could be an alternative avenue to explain the observed pattern in the data. In our model of Section 3 that we take to the data, such a possibility will be allowed, but in the end it is not favored by the data, and so we do not focus on it here.

[34]By the eigenvalues of the system, we mean the eigenvalues of the deterministic version of the system (i.e., the eigenvalues of the system when μ_t does not appear in (2)). Equivalently, these are the eigenvalues of the AR(2) process for e_t defined by $B(L)e_t = \nu_t$ with ν_t *i.i.d.*

[35]See the proof of Proposition 1.

that individual-level choices are cyclical when q is constant. This brings us to the main proposition for this section, which indicates necessary conditions for agent interactions to create a hump-shaped spectral density.[36]

PROPOSITION 1: *Under Assumptions 1 and 2, for agent interactions to produce a hump in the spectral density—i.e., for $g(\omega)$ to have a maximum on the interior of $[0, \pi]$—it is necessary to have $\alpha_3\alpha_4 > 0$, $\alpha_1 < 0$, and $\alpha_2 > 0$.*

Proposition 1 highlights three forces that are necessary for aggregate outcomes to exhibit a spectral hump even when individual-level forces are not sufficient to produce one. The first of these conditions is that the decisions of agents act as strategic complements. This is captured by the condition $\alpha_3\alpha_4 > 0$; that is, α_3 and α_4 need to be such that the decision by one agent to increase e_{jt} causes others to do the same. In addition to this complementarity property, it is also necessary that the stock X_{jt} tends to have a dampening effect on e_{jt}. Finally, the inertial effect must be positive, i.e., there must be sluggishness in the adjustment of e_{jt} over time. The intuition for these results is most easily understood in the case where e is interpreted as an investment decision and X is interpreted as a stock of capital. Complementarity ($\alpha_3\alpha_4 > 0$) causes agents to want to invest at the same time, but as they invest together, this leads to an increase in X which, due to decreasing returns ($\alpha_1 < 0$), eventually leads them to want to dis-invest. These countervailing forces can produce oscillations in e. Finally, sluggishness ($\alpha_2 > 0$) is necessary to prevent the resulting cycles from occurring too rapidly to produce a hump in the spectrum.[37]

That the conditions highlighted in Proposition 1 may lead to cyclicality in this simple setup may not be too surprising. However, what to us is more interesting is that these conditions are in fact *necessary* here. In particular, as we discuss later, few of the commonly estimated types of models in the literature embed these forces simultaneously, and therefore many models appear, by design, limited in their capacity to endogenously explain aggregate cyclicality.

B. Pushing the Complementarities Further: Local Instability and Limit Cycles

The second issue we want to address with the model represented by equations (2)-(4) relates to how such a system behaves if the complementarity, as captured by $\alpha_3\alpha_4$, becomes strong enough. In particular, Proposition 2 indicates that, even if we restrict our focus to the case where $\alpha_3\alpha_4 < 1$, as the complementarity becomes sufficiently strong, this system will become unstable.

PROPOSITION 2: *There exists an $\alpha^\star \in (0, 1)$ such that if $\alpha^\star < \alpha_3\alpha_4 < 1$ then the system represented by (2)-(4) is unstable. Moreover, as $\alpha_3\alpha_4$ increases from*

[36] Proofs of all propositions are presented in the Appendix at the end of this article.

[37] Specifically, cycles would last two periods, and thus the "peak" in $g(\omega)$ would occur at frequency π, rather than an interior point of $[0, \pi]$ as desired.

zero, at the point where stability is lost the eigenvalues of the system are complex if $\delta^2 < -\frac{\alpha_1}{\alpha_2} < (2-\delta)^2$

Proposition 1 and Proposition 2 together suggest that one needs to be aware of certain difficulties if one builds on a structure similar to equations (2)-(4). In particular, Proposition 1 suggests that allowing for complementarities may be necessary to understand endogenously generated hump-shaped spectral densities, while Proposition 2 suggests that it is precisely in such a case that instability can arise. At first pass, one may be tempted to disregard the possibility of instability since market economies do not appear to be explosive. However, as we discuss below, this may be too dismissive. In contrast, we take Proposition 2 as suggesting that, in frameworks that embed complementarities and dynamics together, one should take seriously the possibility of local instability.

Since equations (2) and (3) are linear, and since we have assumed thus far that Q is also linear, any form of local instability in this setup necessarily implies explosiveness, and therefore could be ruled out by the observation that the economy is not explosive. But this equivalence between instability and explosiveness is a knife-edge feature that depends on not having any non-linearity in the model. In contrast, if we were to allow for some non-linearities in the model—for example by allowing Q to be non linear—then Proposition 2 suggests that, as complementarities (e.g., α_3) increase, the system can experience a Hopf bifurcation[38] resulting in the emergence of a locally unstable steady state surrounded by a stochastic limit cycle,[39] which is a situation where the internal dynamics of the system (i.e., when the stochastic elements are set to zero) support a perpetual (limit) cycle. As discussed in the introduction, the idea that macroeconomic fluctuations may reflect limit cycle forces has a long history in the literature, even though it has played a minor role in modern macroeconomics. The value of Propositions 1 and 2 is to emphasize that the equilibrium and behavioral features that may allow a linear model to produce a hump-shaped spectral density may simultaneously be features that allow a non-linear version of the same model to produce the same pattern as the result of a limit cycle. Hence, when attempting to explain hump-shaped spectral densities, it appears warranted to explore whether such a pattern may reflect limit cycle forces, or whether it can be well captured by a stable linear structure. We explore this issue further in Section III.

C. Discussion

While the reduced-form model represented by equations (2)-(4) is largely mechanical, we believe it nevertheless suggests some general modeling features that could help to explain the spectral density features we documented in Section I.

[38]Since the system under consideration is in discrete time, the bifurcation that occurs is more accurately referred to as a Neimark-Sacker bifurcation (see Kuznetsov (1998)).

[39]See Beaudry, Galizia and Portier (2015), (2016*b*) for further details about this statement including an extension to models with a forward-looking component in which a saddle-path limit cycle can arise.

In particular, Proposition 1 emphasizes that, to produce a hump-shaped spectral density as the result of equilibrium interactions, one likely needs agents' decisions to positively affect the decisions of others; that is, the decisions should be strategic complements. However, in many macroeconomic models, individual-level decisions tend to act as strategic *substitutes* (as would be captured in our model by having $\alpha_3\alpha_4 < 0$) due to standard price effects. This property remains true even in most New Keynesian models, since the central bank typically raises interest rates if one group increases their purchases, thereby causing others to decrease their purchases. Hence, to produce a hump-shaped spectral density, a model will likely need to depart sufficiently from standard neo-classical principles—which favor substitutability—to a model in which actions can act as strategic complements. Our empirical observations regarding the risk premium on borrowing suggest that this premium tends to fall when activity is high, which may represent such a source of complementarity. This suggests that one promising avenue to explain spectral peaks may be models of financial frictions[40].

It should be noted that many macro models of financial frictions do not in fact have a structure that, according to the framework presented above, is likely to generate a hump-shaped spectral density. For example, relative to, say, a simple neo-classical benchmark, workhorse models of financial frictions such as Carlstrom and Fuerst (1997) and Bernanke, Gertler and Gilchrist (1999) introduce net worth as a relevant determinant of activity. This net worth can be thought of as a stock variable X_{it} that favors—rather than dampens—activity (i.e., it features $\alpha_1 > 0$),[41] with agents' choices of current activity typically continuing to act as substitutes (through standard price effects), rather than complements. Thus, the net worth channel may amplify and prolong the effects of shocks, but is unlikely to generate a hump-shaped spectral density. Accordingly, in the next section, we explore a model in which financial frictions do allow for complementarity, and where the accumulation of household capital has a dampening effect on new purchases.[42] We use this model to explore a set of issues, namely, (*i*) how allowing for complementarities can change one's inference regarding the role played by exogenous forces in explaining the data, (*ii*) the sensitivity of such inferences to the data frequencies used in estimation, and (*iii*) the implications of allowing for limit cycles in the estimation of the model.

[40]See Beaudry, Galizia and Portier (2018) for suych a model which is a flexible price alternative to the New Keynesian model of the next section

[41]Recall that Proposition 1 emphasizes that an accumulation variable should affect activity negatively if it is to generate spectral peaks.

[42]One alternative possibility would be to pursue mechanisms more akin to those emphasized in Kiyotaki and Moore (1997). We choose not to follow this route as we want to place financial frictions that affect households, as opposed to firms, central to the mechanism. See Mian and Sufi (2018) for an exposition of the importance of financial frictions on the household side in the business cycle.

III. A New Keynesian Model

In this section, we use an extended New Keynesian model to explore how, depending on the mechanisms allowed, the frequency range targeted in the estimation, and the estimation method, inferences regarding the relative importance of endogenous and exogenous forces in explaining the data can differ. To this end, we study an environment with real-financial linkages that focuses on the household, and emphasizes how labor market risk (in particular, unemployment risk) can affect financial conditions through its effect on default risk, which in turn can affect consumer demand, thereby feeding back to the labor market. Our focus on the household is motivated in part by the work of Mian and Sufi (2018) and the main elements of the models are meant to capture ideas of the credit cycle as emphasized in the work by Minsky (1986).

A. The Model

We begin by presenting the household setup and the determination of lending rates in the banking sector. In Section I of the online Appendix, we present all the details of the model.

The Determination of Household Consumption Decisions and Risk Premia on Loans

Consider an environment with a mass one continuum of identical households, each composed of a mass one continuum of identical members (workers) and a household head. Each household, through it members, purchases consumption services on the market at nominal price P_t. Letting variables with h subscripts denote variables for household h, and those without denote aggregate variables, household h's preferences are given by

$$\mathbb{E}_0 \sum_t \beta^t \xi_{t-1} [U\left(C_{ht} - \gamma C_{t-1}\right) + \nu(1 - e_{ht})]$$

where $\mathbb{E}_t$ is the expectation operator, C_{ht} is the consumption services purchased by the household at time t, C_t is aggregate consumption services (so that there is external habit formation), e_{ht} is the fraction of employed household members, $U(\cdot)$ and $\nu(\cdot)$ are standard concave utility functions, and ξ_t represents an exogenous preference shifter. The worker-members of the household look for jobs and are ready to accept employment as long as the real wage $\frac{W_t}{P_t}$ is no smaller than the reservation value of their time to the household. In addition to purchasing consumption services on the market, the household also invests in durable goods. These durable goods could represent, for example, clothes, furniture, cars or houses. To avoid issues of indivisibility, we assume that households do not directly consume the services from their durable goods, but instead rent their durable goods to firms, who use them to produce and sell consumption services back to

the households. This is why we specify utility over consumption services instead of consumption goods. A household's holding of durable goods is denoted by X_t,[43] the nominal rental rate on durable goods is denoted by R_t^X, and the nominal price of durables is P_t^X. Durable goods accumulate according to

$$X_t = (1-\delta)X_{t-1} + I_t \tag{6}$$

where I_t is the total amount of durable goods purchased by the household at time t and δ is the depreciation rate.

In order for the financial market to play a role, we assume that members of the household need to place orders with firms at the beginning of each period before they have received any wage or rental payments. For this reason, household members take out loans at the beginning of each period, with the plan to pay them back at the beginning of next period after they have received their income payments. The uncertainty at this stage is only idiosyncratic, and comes from the fact that jobs are indivisible. All the workers inelastically supply one unit of labor, but firms will demand only $e_t \leq 1$ jobs, so that a fraction $1 - e_t$ of workers will be unemployed. The key market imperfection we introduce is that the financial link between a household and its members is imperfect, in the sense that if a household member cannot pay back its loan, it may be costly for banks to recover the loan amount from the household. In particular, if a household member is unable to pay back a loan—which will be the case when she cannot find a job—then with exogenous probability ϕ the bank can pay a cost $\Phi < 1$ (per unit of the loan) to recover the funds from the household, while with probability $1 - \phi$ it is prohibitively costly to pursue the household, in which case the bank is forced to accept a default. The variables ϕ and Φ will therefore control the degree of financial market imperfection, with $\phi = 1$ and $\Phi = 0$ creating a frictionless credit market. As we show in Section I of the online Appendix, this financial market imperfection yields a budget constraint for household h of the form

$$D_{ht+1} = \left[\bar{e}_t + (1-\bar{e}_t)\,\phi\right](1+r_t)\left(D_{ht} + P_t C_{ht} + P_t^X I_{ht}^X\right) - (1+i_t)\,Y_{ht}\,, \tag{7}$$

where D_{ht} is debt owed by the household when entering period t, r_t is the nominal interest rate charged on one-period loans by the banking sector, i_t is the risk-free interest rate banks pay on deposits, $\bar{e}_t$ is the aggregate employment rate, and $Y_{ht} \equiv e_t W_t + R_t^X X_{ht} + \Pi_t$ is total nominal household income, where Π is total firm profits (of which each household receives an equal share). Note that the effective borrowing rate for the household is $[\bar{e}_t + (1-\bar{e}_t)\,\phi](1+r_t)$, which is the loan rate times the probability that the household will have to pay the loan back. The household has an Euler equation associated with the optimal choice of

[43] Since all households will be identical, for notational simplicity we will often drop h subscripts where no confusion should arise.

consumption services given by

$$U'\left(C_t-\gamma C_{t-1}\right)=\beta\frac{\xi_t}{\xi_{t-1}}\left[\bar{e}_t+\left(1-\bar{e}_t\right)\phi\right]\left(1+r_t\right)\mathbb{E}_t\left[\frac{U'\left(C_{t+1}-\gamma C_t\right)}{1+\pi_{t+1}}\right], \tag{8}$$

where $\pi_{t+1}\equiv P_{t+1}/P_t-1$ is the inflation rate from period t to $t+1$. If $\phi=1$ then we have a standard Euler equation where the marginal rate of substitution in consumption across periods is set equal to the real rate of interest faced by households. When $\phi<1$, Equation (8) reflects the fact that the household knows that it may default on some fraction of its loans. The household will also have an Euler equation associated with the purchase of durables given by[44]

$$\begin{aligned}U'\left(C_t-\gamma C_{t-1}\right)&=\beta\frac{\xi_t}{\xi_{t-1}}\\&\times\mathbb{E}_t\left[\frac{U'\left(C_{t+1}-\gamma C_t\right)}{\left(1+\pi_{t+1}\right)P_t^X}\left\{\frac{R_{t+1}^X\left(1+i_{t+1}\right)}{\left[\bar{e}_{t+1}+\left(1-\bar{e}_{t+1}\right)\phi\right]\left(1+r_{t+1}\right)}+\left(1-\delta\right)P_{t+1}^X\right\}\right].\end{aligned} \tag{9}$$

Equation (9), when combined with (8), can be interpreted as an arbitrage condition that the return to holding a durable good must satisfy.[45]

The central bank sets the nominal interest rate for safe debt i_t, which is also the bank deposit rate, and competition will lead r_t to be such that banks make zero profits. The increased probability of loan defaults when unemployment is high will cause banks to compensate by increasing their margins over the risk-free interest rate. In particular, as shown in Section I of the online Appendix, in a zero profit equilibrium we will have

$$1+r_t=\left(1+i_t\right)\frac{1+\left(1-\bar{e}_t\right)\phi\Phi}{\bar{e}_t+\left(1-\bar{e}_t\right)\phi}. \tag{10}$$

We refer to $r_t^p\equiv\frac{1+r_t}{1+i_t}-1$ as the risk premium. Note here that the risk premium does not represent an excess return but simply corresponds to the amount needed to compensate for expected default. Using (10) to replace r_t in (8) we get

$$U'\left(C_t-\gamma C_{t-1}\right)=\beta\frac{\xi_t}{\xi_{t-1}}\left[1+\left(1-\bar{e}_t\right)\phi\Phi\right]\left(1+i_t\right)\mathbb{E}_t\left[\frac{U'\left(C_{t+1}-\gamma C_t\right)}{1+\pi_{t+1}}\right]. \tag{11}$$

From Equations (8) and (10) we see how unemployment risk, financial conditions, and purchasing decisions all become interrelated due to the fact that loans

[44]Note that the household treats the purchase of durable goods as it would any other asset.

[45]In the case where $\phi=1$ and $i_t=r_t$, equation (9) reduces to the standard asset-pricing condition

$$U'\left(C_t-\gamma C_{t-1}\right)=\beta\frac{\xi_t}{\xi_{t-1}}\mathbb{E}_t\left[\frac{U'\left(C_{t+1}-\gamma C_t\right)}{\left(1+\pi_{t+1}\right)P_t^X}\left\{R_{t+1}^X+\left(1-\delta\right)P_{t+1}^X\right\}\right],$$

to households occasionally involve default and the rate of default $(1-\bar{e}_t)\,\phi$ is endogenously determined. Equation (10) implies that as unemployment increases so does the risk premium on loans, while equation (8) indicates that a higher risk premium on loans will lead households to delay their purchases. Equation (11) gathers these two forces together indicating that higher unemployment (i.e., a fall in $\bar{e}_t$) will lead to a delay of consumption. This is the source of strategic complementarity in this model: if an agent decides to purchase more goods, this will tend to lower the unemployment rate, which in turn allows banks to charge a lower borrowing rate, thereby stimulating other agents to purchase more. Note that this effect runs through $\bar{e}$, so it is external to the household, as was the case in Section II.

The household can dictate the reservation wage to household members which implies that workers accept all jobs for which the real wage satisfies

(12)
$$\frac{W_t}{P_t} \geq \frac{v'\left(1-e_t\right)}{U'\left(C_t-\gamma C_{t-1}\right)}\frac{\left[e_t+\left(1-e_t\right)\phi\right]\left(1+r_t\right)}{\left(1+i_t\right)}+\frac{1+r_t}{1+i_t}\left(1-\phi\right)\left(C_t+\frac{P_t^X}{P_t}I_t^X\right).$$

In equilibrium, firms will offer wages that satisfy (12) with equality. Note that if $\phi=1$ and $r_t=i_t$, then (12) implies accepting all wage offers where the real wage is higher than the marginal rate of substitution between leisure and consumption.[46]

Firms

There are two types of firms in the model: final good firms, and intermediate service firms. The final good sector is competitive and provides consumption services to households by buying a set of differentiated intermediate services, denoted C_{kt}, from the unit mass of intermediate service firms, and combining them in a standard way according to a Dixit-Stiglitz aggregator.[47] Intermediate service producers, meanwhile, are monopolistically competitive in the supply of differentiated consumption services, and take the demand for such services from final good firms as given. These intermediate firms produce consumption services using durable goods, which can either be rented from households or produced anew according to the production technology $F(e_{kt},\theta_t)$, where e_{kt} is labor hired by firm k, and θ_t is exogenous productivity. We assume that newly produced durable goods can immediately produce consumption services, and that the total output of consumption services is simply linear in the total quantity of durables; that is, we assume $C_{kt}=s[X_{kt}+F(e_{kt},\theta_t)]$, $s>0$, where X_{kt} is the amount of durable goods rented by firm k from households.[48] Moreover, after using newly produced durables to produce consumption services, the firm can sell the

[46]Note that the extra term on the right-hand side of (12) reflects the fact that, by accepting a job, the household loses the possibility of being allowed to default on that worker's loan.

[47]See Section I of the online Appendix for further details.

[48]Note that in the production of consumption services, the capital stock variable X_t and the employment level e_t enter in a separable manner as opposed to entering in the more common form of

remaining un-depreciated amount to households at the market price P_t^X.[49] To increase generality, we will assume that the depreciation of new durable goods is given by $1 - \psi \geq \delta$, which allows new durable goods to potentially depreciate faster in the first period than in subsequent periods. This extension allows for the possibility of interpreting that a fraction of the new goods depreciate fully within the first period (i.e., are non-durable) and the remaining fraction depreciates at the standard durables rate δ.

Since intermediate service producers have a choice between two ways of obtaining durables for use in production (i.e., renting existing durables and producing new ones using labor), if both ways are to be used in equilibrium then the net marginal cost of an additional unit of durables must be equalized across the two methods. For a new unit of durables, this marginal cost is equal to the wage cost per additional unit produced, less the value of the un-depreciated portion sold to households. For rented durables, this marginal cost is simply the rental rate. Thus, we must have

$$R_t^X = \frac{W_t}{F_e(e_{kt}, \theta_t)} - \psi P_t^X \tag{13}$$

Following the New Keynesian literature, we assume that the market for intermediate services is subject to sticky prices *à la* Calvo (1983). This yields a standard New Keynesian Phillips curve, though as will become clear shortly, for our purposes we will not need to derive it.

The Central Bank and Equilibrium Outcomes

To close the model, we still need to specify how the central bank determines the risk-free interest rate. In order to keep the model tractable, we restrict attention to a monetary policy rule governed by only one parameter φ_e, which allows the central bank to only imperfectly control its objective of stabilizing inflation and employment. To this end, we assume that the central bank sets the nominal interest rate to induce an expected real interest rate that rises and falls with expected employment. By allowing the central bank to adjust only to expected variables, it can only imperfectly stabilize the economy. To be more precise, we assume that the central bank sets the nominal interest rate according to a rule of

complements. We adopt this formulation for two reasons. First, from a theoretical point of view, this formulation implies that when X_t is high it necessarily tends to depress the desired level of employment since utility is concave in services. In contrast, if e_t and X_t where assumed to be complements, then when X_t is high it become ambiguous whether it depresses or favors current employment. Since from section II, we know that having an accumulation variable that has a negative effect of new purchases is important for allowing a model to endogenously produce hump shaped spectral densities, this formulation is as easy way of permitting such a possibility. Second, from a conceptual point of view, when thinking of X_t as household capital, it appears more appropriate to assume that it does not directly increase the marginal product of labor in the market sector.

[49] Note that in the market for durable goods the intermediate firms are price takers, while they are price setters in the consumption service market.

the form

$$1 + i_t \approx \Theta \mathbb{E}_t \left[e_{t+1}^{\varphi_e} (1 + \pi_{t+1}) \right] \tag{14}$$

where φ_e controls the extent to which the central bank tries to stabilize inflation and employment, and Θ controls the steady state level of i. As we show shortly, the attractive feature of this monetary policy rule is that it gives the equilibrium equations a block recursive structure.[50]

The equilibrium outcomes for this model are given by a set of nine equations determining the two aggregate quantities $\{C_t, I_t\}$, the employment level e_t, the relative prices $\left\{ i_t, r_t, \frac{R_t^X}{P_t}, \frac{P_t^X}{P_t}, \frac{W_t}{P_t} \right\}$ and the inflation rate π_{t+1}.[51] Using the monetary policy rule (14), as shown in Section I of the online Appendix the equilibrium equations have a convenient block-recursive structure whereby the variables e_t, X_{t+1} and r_t^p can be solved for first using the equations[52]

$$1 + r_t^p = \frac{1 + (1 - e_t)\, \phi \Phi}{e_t + (1 - e_t)\, \phi} \tag{15}$$

$$X_{t+1} = (1 - \delta) X_t + \psi F(e_t, \theta_t)\,, \tag{16}$$

(17)

$$\begin{aligned} U' \left(s \left(X_t + F(e_t, \theta_t) \right) - \gamma s \left(X_{t-1} + F(e_{t-1}, \theta_{t-1}) \right) \right) = \beta \Theta \frac{\xi_t}{\xi_{t-1}} \left[e_t + (1 - e_t)\, \phi \right] (1 + r_t^p) \\ \times \mathbb{E}_t \left[U' \left(s \left(X_{t+1} + F(e_{t+1}, \theta_{t+1}) \right) - \gamma s \left(X_t + F(e_t, \theta_t) \right) \right) e_{t+1}^{\varphi_e} \right]. \end{aligned}$$

These three equations will provide our basis for exploring whether such a model can capture the spectral properties for hours and the risk premium that we documented in Section I.

[50] The precise form of the Taylor rule we use to obtain the block recursive property is

$$1 + i_t = \Theta \mathbb{E}_t \left[e_{t+1}^{\varphi_e} \frac{U' (C_{t+1} - \gamma C_t)}{\mathbb{E}_t \left[\frac{U' (C_{t+1} - \gamma C_t)}{1 + \pi_{t+1}} \right]} \right].$$

This deviates slightly from (14) due to Jensen's inequality, which is why (14) is expressed with an $\approx$ symbol.

[51] The relevant equilibrium equations correspond to (8), (9), (10), (12) with equality, (13), the aggregate consumption equation $C_t = s[X_t + F(e_t, \theta_t])$, the accumulation equation (6), the Phillips curve and the Taylor rule.

[52] Given the values of e_t, X_{t+1} and r_t^p obtained from this system, the remaining equations simultaneously determine the remaining variables $\left\{ C_t, \frac{R_t^X}{P_t}, \frac{P_t^X}{P_t}, \frac{W_t}{P_t}, \pi_t \right\}$. Note that, as we do not consider the implications of the model for inflation, we will not need to explicitly derive the optimal pricing behavior of firms.

Shocks

There are two exogenous forces in the model that affect the determination of hours and the risk premium: the preference shifter ξ_t and the level of technology θ_t. Note that the preference shifter could alternatively be interpreted as a monetary shock, since allowing for a monetary shock gives rise to the exact same equations for the determination of hours and the risk premium. We use the more abstract interpretation of it as a preference shifter, since it allows for several different interpretations. In our estimation, we will focus on the case where technology is constant and the only stochastic driving force is the preference shifter, so as to see whether such a minimalist exogenous structure, once embedded in an environment with a potentially rich endogenous propagation mechanism, can capture the spectral properties of the data.

Limit Cycles and Arbitrage

As will become clear when we present our estimation results below, for certain parameterizations our model will feature limit cycles. A common criticism of earlier models featuring limit cycles is that these cycles would be subject to arbitrage forces that would tend to erase them. We wish to emphasize that, while this criticism may have been valid for models that did not feature rationally optimizing forward-looking agents (e.g., Hicks (1950) and Goodwin (1951)), it does not apply to our model, since agents are rationally optimizing and forward-looking. For example, even though there may be predictable cyclical elements in the price of durable goods, so that agents could potentially borrow to purchase durables when the price is low and then sell them for a profit in the future when the price is high, the return they would earn from this strategy in equilibrium is necessarily less than the cost of servicing the associated debt, and would therefore not be optimal.

Equilibrium narrative

Before proceeding to the estimation of this model, it is useful to briefly discuss how agents in this environment would perceive macroeconomic fluctuations and why their behavior could generate endogenous cyclical outcomes. To this end, suppose the economy has been in recession for a while. Banks in this environment are reluctant to lend to people because unemployment is high. This leads them to charge a high premium on loans so as to cover expected defaults. However, if the recession has been going on for long enough, household capital will have depreciated significantly, causing the marginal utility of new purchases to be high. In fact, the marginal utility will eventually become sufficiently high that an agent will be ready to increase their borrowing in order to make new purchases, even if the risk premium on borrowing remains high. At that point the tide starts to turn. As some individuals start purchasing more, demand increases, which reduces unemployment. Banks respond to this reduced unemployment risk by

decreasing the risk premium they charge on loans, which favors more purchasing by households. As a result, a boom period emerges. The boom will continue with households continuing to accumulate capital until, at some point, the marginal utility of new purchases becomes low. The low marginal utility induces households to reduce their borrowing, even if the risk premium on loans charged by banks is low. This is how a recession would endogenously start. Importantly, agents in this economy understand these boom-bust dynamics, but this does not stop them from taking part, even if, individually, they would prefer stable consumption. In fact, knowledge of the endogenous boom-bust cycle further pushes agents to go along with the cycle, knowing that a boom is a good time to accumulate—even over-accumulate—since the risk premium is currently low and a recession is known to be coming eventually. Similarly, during a downturn, agents prefer to hold back from making purchases, since they know that financial conditions are temporarily bad. Note that this mechanism could be self-sustaining, which would be the case if it reflected a limit cycle. Alternatively, the mechanism could peter out over time, requiring exogenous shocks to sustain it.

The addition of shocks to this model eliminates the perfect predictability of the endogenous cyclical forces, but does not change the endogenous mechanisms. Agents would no longer know exactly when a boom or a bust would finish, but they would know that the system is more susceptible to a bust if it has been booming for a while, and would know that the economy is more susceptible to a boom if it has been in a prolonged downturn.[53] It is also important to emphasize that the above description of the economy does not involve indeterminacy or self full-filling beliefs.[54] A boom does not arise in this set up simply because people believe it will arise. It arises because the main capital stock has been depleted sufficiently, and this pushes people to take risks and borrow more. Changes in expectations are not the driving force; rather, they act more like an amplification mechanism. It should also be noted that the model narrative is not especially novel, being meant to capture many elements that are common to descriptions of the financial cycle.

B. Functional Forms and Estimation

To bring our model to the data, it remains to specify functional forms and the stochastic process. We assume that period utility is CRRA and given by $U(c) = (c^{1-\omega}-1)/(1-\omega)$, while the production function is given by $F(e,\theta) = \theta e^{\alpha}$. As noted above, we take technology θ as constant, and normalize it to 1.[55] We also normalize $s = 1$. As shown above, we may reduce our system to three equations in the variables X, e, and r^p. Linearizing the two dynamic equations (16)-(17)

[53]Note that these implied patterns are consistent with the properties of the conditional probability of recession described in Section I.

[54]In Beaudry, Galizia and Portier (2016*b*), it is shown that this type of model does not in general allow for self-fulfilling fluctuations unless one allows the complementarities to be sufficiently strong to allow for multiple steady states.

[55]Allowing for deterministic growth in the model does not change any results.

with respect to $\log(X)$, $\log(e)$, r^p, and $\mu_t \equiv -\Delta \log(\xi_t)$, we obtain equations of the form[56]

$$\widehat{X}^{\star}_{t+1} = (1-\delta)\,\widehat{X}^{\star}_{t} + \psi \widehat{e}_t\,, \tag{18}$$

$$\widehat{e}_t = \alpha_1 \widehat{X}^{\star}_{t} + \alpha_2 \widehat{e}_{t-1} + \alpha_3 \widehat{e}_{t+1} - \alpha_4 \widehat{r}^{p}_{t} + \alpha_4 \mu_t\,, \tag{19}$$

$$\widehat{r}^{p}_{t} = \varrho_1 \widehat{e}_t\,, \tag{20}$$

where $\widehat{X}$ and $\widehat{e}$ are log-deviations from steady state, $\widehat{r}^p$ is the deviation in levels, $\widehat{X}^{\star} \equiv \psi \widehat{X}/(\alpha\delta)$, and the α_j's are functions of the structural parameters with $\alpha_1 < 0$ and $\alpha_2, \alpha_3, \alpha_4 > 0$. Note that, since the risk premium is a negative function of employment ($\varrho_1 < 0$), the above system has a structure similar to the model of Section II. We assume that μ_t follows a stationary AR(1) process $\mu_t = \rho\mu_{t-1} + \epsilon_t$, where ϵ_t is a Gaussian white noise with variance σ^2.

We estimate four different versions of our three-equation model. In all versions we use the dynamic equations in their linearized forms (18) and (19). In one version, which we refer to as the linear risk premium (RP) model, we use as the third equation the static risk premium equation (15) in its (log-)linearized form as given by (20). In another version, which we refer to as the no friction model, we assume that all household members receive the backing of the household (i.e., $\phi = 1$) and recovery from the household is costless (i.e., $\Phi = 0$), so that the risk premium is always zero, and thus there is no complementarity. This estimation will help to illustrate the importance of the complementarity in allowing our model to match the key features of the data. In a third version, we shut down both the complementarity channel (by setting $\phi = 1$, $\Phi = 0$) *and* the accumulation channel (by setting $\psi = 0$, so that $X_t = 0$ for all t). We refer to this as the canonical model, since it corresponds closely to the canonical New Keynesian model with habit.

In the final version of the model, which we refer to as the non-linear RP model, we allow the risk premium to be a non-linear function of (log-)employment.[57] As discussed in Section II, allowing for non-linearity in the strength of the complementarity will allow us to expand the parameter space to include situations where there may be local instability and limit cycles. In particular, to allow for a greater set of possibilities, for this case we let the debt-backing probability be a function of the employment rate, i.e., $\phi_t = \phi(e_t)$, and approximate (15) to the third order as

$$\widehat{r}^{p}_{t} = \varrho_1\, \widehat{e}_t + \varrho_2\, \widehat{e}^{2}_{t} + \varrho_3\, \widehat{e}^{3}_{t}\,, \tag{21}$$

where the coefficients ϱ_1, ϱ_2 and ϱ_3 are functions of the model parameters and of

[56] See Section J of the online Appendix for details.

[57] We could also allow for non-linearities in both (18) and (19). However, we chose to allow for non-linearities in the risk premium so as to make the analysis more transparent, since it allows us to refer to results from Section II.

the three first derivatives of ϕ (ϕ', ϕ'' and ϕ''') evaluated at the steady state.[58] To facilitate comparison with the linear RP model and to avoid issues of identification, we will fix $\phi' = 0$ in the estimation. Note that we take a third-order approximation since, for the model to be capable of producing (attractive) limit cycles, we will typically need the third-order coefficient ϱ_3 to be sufficiently positive (see Kuznetsov (1998) for details).

We estimate this model using the indirect inference method of Gourieroux, Monfort and Renault (1993), where for each parameterization the model is solved by a first-order (linear RP, canonical, and no friction models) or third-order (non-linear RP model) perturbation method.[59] In the non-linear RP model, the solution and estimation are somewhat involved, as it allows for the possibility of a locally unstable steady state and limit cycles in a stochastic model with forward-looking agents. To our knowledge, such an exercise is novel.[60] For each version of the model, the parameters are chosen so as to minimize its distance to a set of features of the data that we have already emphasized. We focus on three sets of observations. The first set, which is used for all four models, corresponds to the spectral density of hours worked per capita (as shown in panel (b) of Figure 2).[61] The second set, which is used for both the linear and non-linear RP models, adds the spectral density of the risk premium (as shown in panel (c) of Figure 5).[62] For these first two sets, we aim to fit the point estimates of the spectral densities (using the non-detrended data) at periodicities between 2 and 50 quarters.[63] The last set of observations, which is used only for the non-linear RP model, 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.

We calibrate three parameters for all four models: the depreciation rate is set to $\delta = 0.05$ in order to match the average depreciation of houses and durable goods, the elasticity of the production function with respect to employment is set to $\alpha = 2/3$, and the monetary policy scale variable Θ is set so as to yield a

[58] See Section J.1 of the online Appendix.

[59] Details of the solution and estimation are given in Section J of the online Appendix.

[60] For details about the solution method, see Galizia (2018).

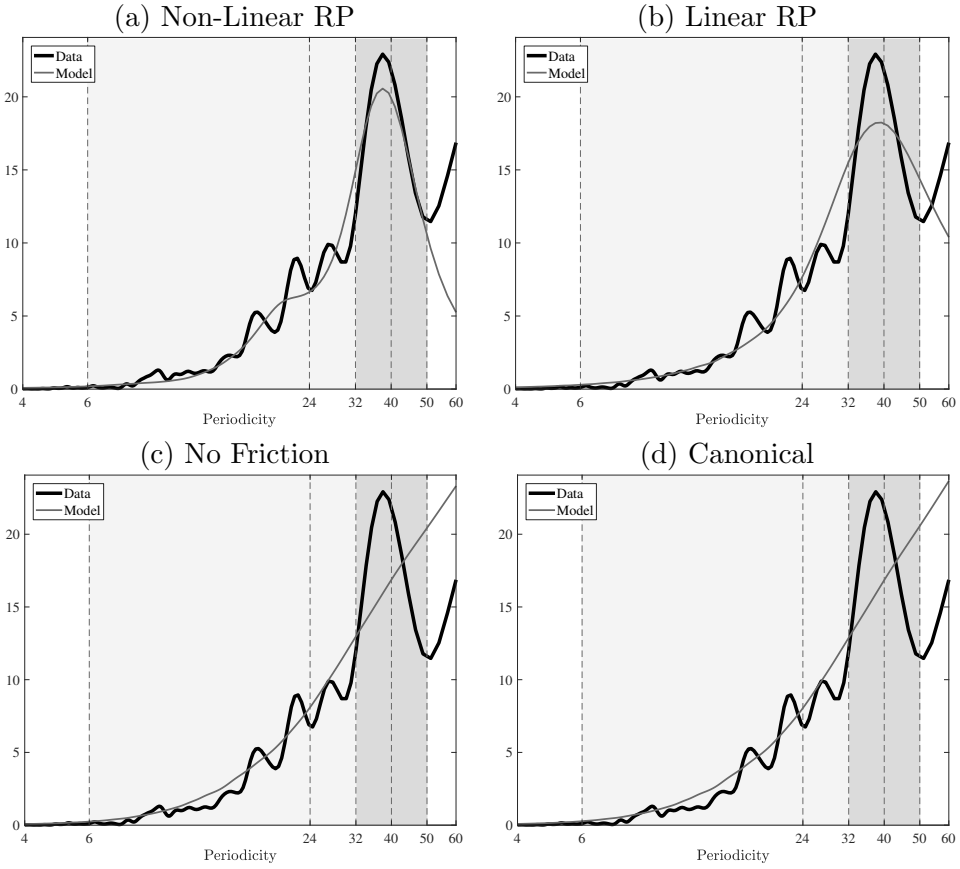
[61] Note that, in our model, since employed workers each work the same number of hours, hours is simply proportional to employment.

[62] We use the BAA Corporate Bond spread series, rather than a series that directly measures interest rates faced by households, as the former is available going back to 1954, while we have only found quarterly measures of the latter that go back to the early 1970s. The coherences between the bond spread and those consumer spread series over the 32-50 quarter range—the range straddling the spectral peaks—are around 0.8, suggesting that fluctuations in the bond spread may be a reasonable proxy for fluctuations in consumer spreads. As a check on our results, we re-estimated the model also including a measurement error process on the risk premium, and found little change in the results (available upon request).

[63] See Christiano and Vigfusson (2003), Qu and Tkachenko (2012) and Sala (2015) for previous work estimating DSGE type models in the frequency domain.

steady state unemployment rate of 0.0583 (the average over our sample period). Depending on the particular model, we then estimate as many as ten parameters. For the two risk premium models, we estimate ω, γ, ψ, φ_e, ϕ, Φ, ρ, and σ, plus for the non-linear RP model, ϱ_2 and ϱ_3 (from which we can back out ϱ_1, ϕ'' and ϕ'''). Because of the structure of the no friction and canonical models, and the fact that they are estimated only on the hours spectrum, to avoid identification issues we fix a number of the parameters in estimation. In addition to the parameters restricted by definition (ϕ and Φ, plus ψ in the canonical model), we fix ω and φ_e, plus ψ in the no friction model and γ in the canonical model, at the corresponding estimated values from the linear RP model.

FIGURE 6. FIT OF HOURS SPECTRAL DENSITY

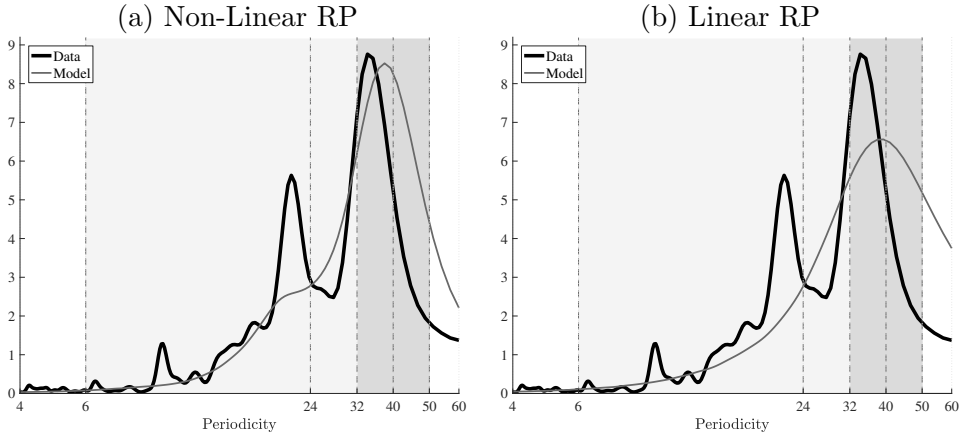


Note: This figure compares the estimate of the spectral density of U.S. Non-Farm Business hours per capita with the ones obtained from our four estimated models.

Figures 6-8 illustrate the fit of the estimated model along the targeted dimensions for the four different versions of the model. Consider first panel (b) of Figures 6 and 7, which show the estimated spectral densities of hours and the

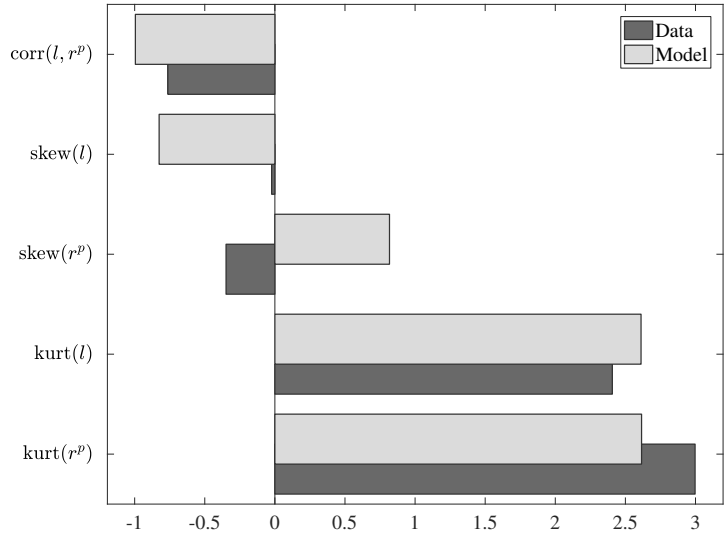
risk premium, respectively, for the linear RP model. While our parsimonious model does not capture all the bumps and wiggles in the spectrum of hours in Figure 6, it nonetheless fits the overall pattern nicely, most importantly the peak in the spectrum near 40 quarters, though that peak is noticeably flatter than the one observed in the data. Meanwhile, while the model does not capture the smaller peak in the spectral density of the risk premium in Figure 7 observed around 21 quarters, it again fits reasonably well the overall hump-shaped pattern with a peak close to 40 quarters, though again that peak is flatter than the one observed in the data. Consider next panel (c) of Figure 6, which shows the fit of the hours spectrum for the no friction model, obtained by re-estimating the model after shutting down the complementarity channel. Despite the fact that for this model the estimation is no longer constrained to simultaneously match the risk premium spectrum, the fit of the hours spectrum is significantly worse than in the linear RP model, and in particular the no friction model is unable to replicate the peak in the spectrum of hours near 40 quarters. Thus, evidently the "complementarities" part of our accumulation-with-complementarities mechanism is of fundamental importantance in allowing our model to capture the salient business cycle features of the data. As shown in panel (d) of Figure 6, the canonical model, which is obtained by re-estimating the model with both the complementarity and accumulation channels shut down, tells a similar story. Lastly, consider panel (a) of Figures 6 and 7, which show the results for the non-linear RP model. Relative to the linear RP model, both spectra fit noticeably better, and in particular while both models have peaks near 40 quarters, unlike in the linear RP model the peaks in the non-linear RP model are almost as pronounced as they are in the data.

Figure 7. Fit of Risk Premium Spectral Density



Note: This figure compares the estimate of the spectral density of U.S. risk premium with the one obtained from our estimated non-linear and linear RP models.

FIGURE 8. FIT OF OTHER MOMENTS (NON-LINEAR RP)



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

The parameter estimates for the four models are presented in Table 2, along with bootstrap standard error estimates in parentheses.[64] Comparing columns (a) and (b) of the Table, we see that the non-shock parameter estimates are broadly similar between our two preferred models (the linear and non-linear RP models). Further, the estimated habit parameter (γ) of 0.53-0.59 is well in line with the values commonly found in the literature. The two parameter estimates that may be considered somewhat low relative to the literature are our estimates of the CRRA parameter (ω) of 0.24-0.3, implying a relatively high intertemporal elasticity of substitution of around 3-4, and of the Taylor rule elasticity (φ_e) of 0.042-0.047, which implies that a one-percentage-point increase in expected employment is associated with an increase in the annualized policy rate of about 17-19 basis points. The first of these parameters implies a strong response of consumption to the interest rate faced by households. Since the household interest rate is the sum of the pro-cyclical policy rate and the counter-cyclical risk premium, the second of these parameters tends to favor counter-cyclicality in the overall household interest rate. Taken together, these parameters help to increase the effect of the complementarity in the model—which, as we have shown, is helpful in matching key features of the data—by increasing both the counter-cyclicality of the household interest rate and the response of consumption to that

[64]Standard errors were obtained by simulating $N = 200$ data sets from the corresponding model (using the point estimates for the parameters), and then re-estimating the model on each simulated data set.

counter-cyclicality.

Table 2—Estimated Parameter Values

		(a) Non-Linear RP	(b) Linear RP	(c) No Friction	(d) Canonical
ω	CRRA parameter	0.2997	0.2408	0.2408*	0.2408*
		(0.0200)	(0.0423)	—	—
γ	Habit	0.5335	0.5876	0.9405	0.5876*
		(0.0031)	(0.0512)	(0.0836)	—
ψ	One minus initial dep.	0.4000	0.2994	0.2994*	0*
		(0.0028)	(0.0644)	—	—
φ_e	Taylor rule	0.0421	0.0467	0.0467*	0.0467*
		(0.0028)	(0.0057)	—	—
ϕ	Debt backing	0.8668	0.8827	1*	1*
		(0.0067)	(0.0074)	—	—
Φ	Recovery cost	0.0421	0.0458	0*	0*
		(0.0029)	(0.0067)	—	—
ϱ_2	Risk prem. (2nd order)	0.0167	—	—	—
		(0.0008)	—	—	—
ϱ_3	Risk prem. (3rd order)	0.5929	—	—	—
		(0.0575)	—	—	—
ρ	Autocorrelation	-0.0000	0.1387	0.8541	0.8609
		(0.0000)	(0.0799)	(0.1118)	(0.1077)
σ	Innovation s.d.	0.00016	0.00027	0.00148	0.00148
		(0.00003)	(0.00010)	(0.00076)	(0.00075)
$s.d.(\mu)$	Implied uncond. s.d.	0.00016	0.00027	0.00285	0.00292

Note: Table displays the estimated parameters of the model for each of the four estimation scenarios with standard errors in parentheses. * indicates calibrated values. Estimates for the non-linear RP model imply $\varrho_1 = -0.1624$, $\phi'' = -3.1$ and $\phi''' = -227.1$, while for the linear RP model we have $\varrho_1 = -0.1506$. In the bottom row, we report the unconditional standard deviation of the shock process μ_t implied by the point estimates for ρ and σ.

The most interesting finding regarding our parameter estimates revolves around the shock process parameters ρ and σ. In particular, as one moves leftward beginning from the canonical model (column (d) of Table 2) to the no friction model, the linear RP model, and finally to the non-linear RP model, both the estimated persistence and standard deviation of the shock process monotonically decrease. As a result, the unconditional standard deviation of the shock process (reported in the bottom row of the table) also monotonically decreases. Thus, sequentially allowing for accumulation, complementarity, and then non-linearity in the complementarity lets the model better fit the data (as discussed above), while also relying less on exogenous stochastic forces. The upshot is that in the non-linear RP model, the shock is effectively *i.i.d.*, so that the model's dynamics are almost entirely due to endogenous forces. Further, the largest factor in reducing reliance on exogenous shocks comes from the introduction of the complementarity (i.e., moving from (c) to (b)), which is associated with a fall of more than 90% in the unconditional standard deviation of the estimated shock process. These re-

sults suggest that our accumulation-with-complementarities mechanism may be a promising avenue for those seeking to introduce stronger internal propagation and a lower reliance on exogenous shocks into business cycle models.

The strength and form of the internal propagation in each model can be seen more clearly in Table 3, which reports the eigenvalues of the first-order approximation to the solved model around the non-stochastic steady state, along with their moduli. In the canonical and no friction models, these eigenvalues are real, positive, and stable, with values given by 0.5617 and 0.5202, respectively. Thus, these models are characterized by monotonic convergence—which may explain their inability to capture the hump in the spectral densities near 40 quarters—and a relatively low degree of endogenous persistence. For example, in the canonical model, deviations from steady state have an endogenous half-life of slightly more than one quarter. For the no friction model, the larger eigenvalue is close to 0.95, but it can be verified that the smaller eigenvalue is the one driving most of the variance in hours, and it too is associated with an endogenous half-life of just over one quarter. In contrast, the linear RP model has a pair of complex eigenvalues, which allows it to generate spectral peaks, and these eigenvalues both have a modulus of 0.93, which is suggestive of a relatively larger degree of endogenous persistence (endogenous half-life of nine quarters). Finally, the non-linear RP model also has a pair of complex eigenvalues, but with a modulus exceeding one (1.12).[65] That is, when given the option, the data appears to favor a configuration featuring local instability and limit cycles, which generates significant internal propagation and a correspondingly lower reliance on exogenous processes to drive fluctuations. This suggests a more general point, which is that, by ruling out parameterizations that produce local instability and limit cycles—as is implicitly done by standard solution methods (e.g., standard perturbation methods implemented in DYNARE.)— one may be significantly biasing the results of estimation.

TABLE 3—EIGENVALUES AT THE STEADY STATE

	Non-Linear RP	Linear RP	No Friction	Canonical
$\lambda_{11}, \lambda_{12}$	$1.1032 \pm 0.2164i$	$0.9258 \pm 0.1372i$	0.5617, 0.9488	0.5202
$\lvert\lambda_{11}\rvert$	1.1242	0.9359	0.5617	0.5202
$\lvert\lambda_{12}\rvert$	1.1242	0.9359	0.9488	—

Note: Table reports the eigenvalues of the first-order approximation to the solved model around the non-stochastic steady state. Note that the solved canonical model has only one dimension (X is no longer a relevant state variable) and therefore has only one eigenvalue.

[65] Note that we constrained the parameter space to allow only for parameterizations producing a determinate solution; that is, where the third eigenvalue (not shown in Table 3) of the unsolved system is unstable. However, we found no indication that this constraint was binding for the linear or non-linear RP models, suggesting that the data do not favor a configuration yielding indeterminacy.

To illustrate the deterministic mechanisms implied by the parameter estimates in the non-linear RP model, Figure 9 reports results from feeding in a constant value of $\mu_t = 0$ for the exogenous process.[66] Panel (a) of Figure 9 plots a simulated 270-quarter sample[67] of hours generated from this deterministic version of the model. Two key properties should be noted. First, the estimated parameters produce endogenous cyclical behavior, with cycles of a reasonable length (around 38 quarters). This is consistent with Table 3, 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 two key features we highlighted in the previous section: complementarities, and an accumulation variable that affects current demand negatively. Second, notwithstanding the reasonable cycle length, it is clear when comparing the simulated data in panel (a) of Figure 9 to actual economic data that the fluctuations in the deterministic model are far too regular. These two properties of the deterministic model, i.e., a highly regular 38-quarter cycle—can also be seen clearly in the frequency domain. Panel (b) of Figure 9 plots the spectral density of hours for the deterministic model (gray line), along with the spectral density for the data (black line) for comparison. This spectral density exhibits an extremely large peak—characteristic of a highly regular cycle—at the 38-quarter periodicity,[68] while the spectral density of the data is much flatter.

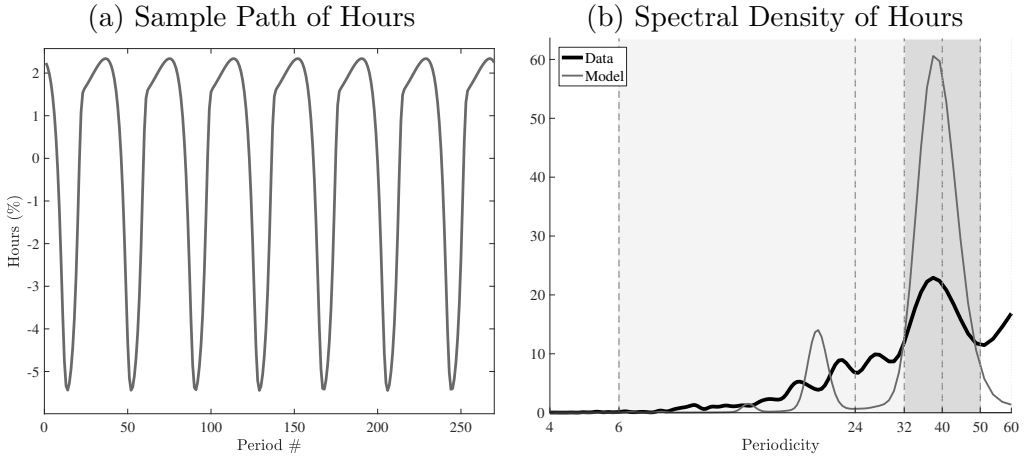
Re-introducing the estimated shocks into the non-linear RP model, we see a markedly different picture in both the time and frequency domains. Figure 10 plots an arbitrary 270-quarter sample of log-hours generated from the full stochastic model. While clear cyclical patterns are evident, it is immediately obvious that the inclusion of shocks—even the (essentially) *i.i.d.* shocks that are present in our model—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 (panel (a) of Figure 6), 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

[66]Note that, as typically done when computing transitional dynamics, the model was solved (and simulated) using the estimated parameters presented in column (a) of Table 2, 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 are due exclusively to differences in the realized sequence of shocks, rather than differences in, say, agents' beliefs about the underlying data-generating process.

[67]This is equal to the length of the sample period of the data.

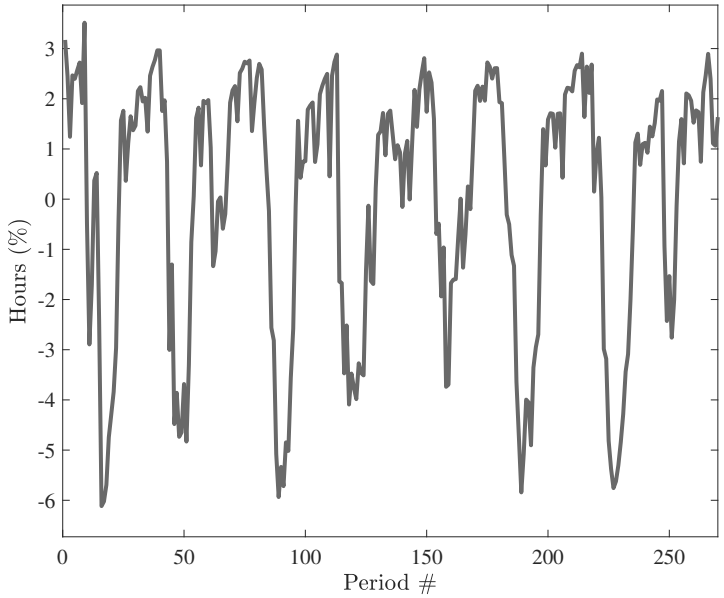
[68]The deterministic model spectral density also contains smaller peaks at integer multiples of the frequency of the main cycle (i.e., at around $19 = 38/2$ quarters, $12.67 = 38/3$ quarters, etc.). Such secondary peaks arise when the data exhibits a regular but not perfectly sinusoidal cycle, as is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

Figure 9. Non-Linear RP Model When Shocks Are Turned Off



Note: This figure corresponds to the counterfactual simulation of the estimated non-linear RP 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 estimated model when simulated without shocks (gray line).

Figure 10. Sample Path of Hours in the Non-Linear RP Model with Shocks



Note: This figure shows one particular 270-quarter simulation of the estimated non-linear RP model.

exaggerated peak observed at this point in the deterministic model.

It should be emphasized that the exogenous shock process in the non-linear RP 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 9), the variance of log-hours actually *increases* relative to the full stochastic case (the variance of log-hours in the stochastic case is 6.88, while in the deterministic case it is 7.48).[69] The role of complementarities in the model, however, is extremely important: if we shut down the endogenous risk premium (i.e, set $\phi = 1$ and $\Phi = \varrho_2 = \varrho_3 = 0$), but keep all other parameters at their estimated levels, the variance of log-hours in the model is less than 0.1 (compared with 6.88 with the complementarity). Thus, without the complementarities to amplify them, the small *i.i.d.* disturbances can only generate a tiny amount of volatility in hours.

C. *Sensitivity of Results to Target Frequency Range*

As we noted in Section I, the spectral densities of several key macroeconomic variables exhibit peaks around 38 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 to target the 2-50 quarter range in our estimation. We turn now to evaluating how this choice affects the results. First, since the lower end of the standard business cycle range in the literature is 6 quarters, we consider what happens if we restrict attention to periodicities from 6 to 50 quarters in the estimation. We also consider what happens if we restrict attention further to only the range beyond the traditional upper bound of 32 quarters (i.e., restricting to periodicities between 32 and 50 quarters), and then repeat each of these exercises using 60 quarters as the upper bound instead of 50. We do these exercises for the non-linear RP model, though the results are similar for the linear RP model (available upon request). As we will see, the above choices make little difference to our results. As a final exercise, we consider what happens when targeting the 2-100 quarter range, and show that in this case the results do fundamentally change.

The top row of plots in Figure 11 shows the hours and risk premium spectral densities for period ranges of the form $(x,50)$, $x = 2, 6, 32$,[70] while the second row shows results for ranges of the form $(x,60)$. The corresponding data spectral

[69] Note that, since we have simply fed a constant sequence $\mu_t = 0$ of shocks into our model without first re-solving 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*a*). Roughly speaking, because of the non-linear 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*a*) for a more detailed discussion of these mechanisms in the context of an estimated reduced-form univariate equation.

[70] Note that the (2,50) results simply reproduce the information in Figures 6 and 7.

densities are also plotted (solid black lines) for comparison. In all cases, the effect of changing the lower bound to 6 or 32 quarters, or the upper bound to 60 quarters, has minimal effect on the parameter values, and this translates into a minimal change in the overall fit of the spectral densities: in all cases the non-linear RP model continues to match the spectral peak near 40 quarters well. Further, as shown in the first six rows of Table 4, the eigenvalues associated with the solved system remain complex (indicating cyclicality) and outside the unit circle (indicating the presence of a limit cycle) for all six of these different period ranges.

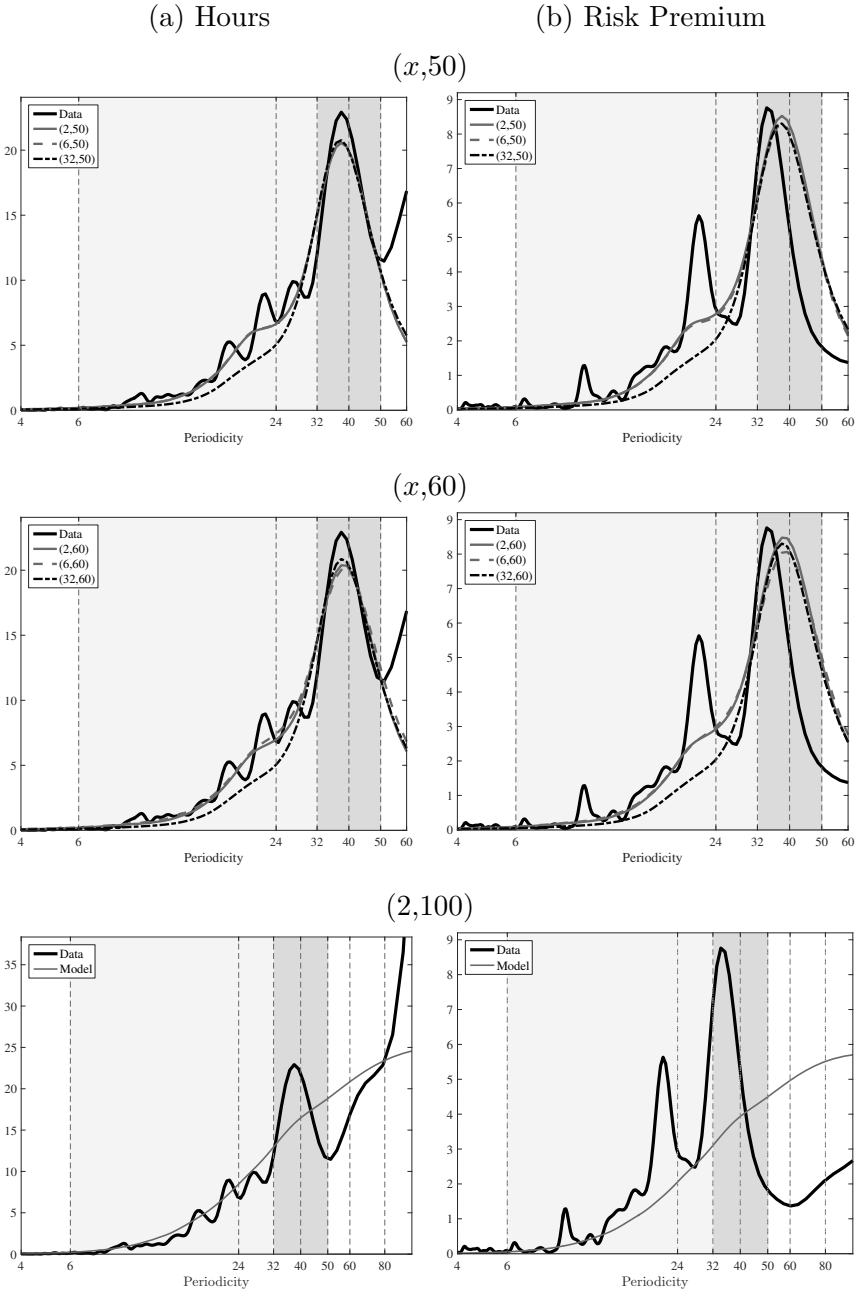
While increasing the upper bound of the target periodicities from 50 to 60 quarters has little effect on the results, the same is not true if we include even more low-frequency fluctuations. The last row of plots in Figure 11 shows the fit if we extend the upper bound to 100 quarters (note the scale change in the horizontal axis), while the last row of Table 4 reports the associated eigenvalues. The model no longer captures the peak in the hours spectrum near 40 quarters, since this is no longer the dominant feature of the data. Instead, the dominant feature is now the steep increase that occurs beyond 60 quarters, which reflects large but slow-moving forces (such as demographic changes) unrelated to the business cycle. The estimated autocorrelation of the exogenous driving force is now above 0.8, with an unconditional standard deviation of 0.0017 (an order of magnitude greater than in the baseline case), suggesting that the dynamics are mainly driven by exogenous forces. This highlights the more general point that if one attempts to simultaneously explain fluctuations in hours data at all frequencies, not just those related to the business cycle, one may likely miss important business cycle features unless one explicitly includes in the model mechanisms to explain the lower frequencies movements. This may help to explain why few modern business cycle models—which are typically implicitly estimated to simultaneously fit all frequencies—generate a peak in the spectrum near 40 quarters.

Table 4— Eigenvalues at the Steady State for Various Estimations

Est. Range	$\lambda_{11}, \lambda_{12}$	$\lvert\lambda_{1j}\rvert$
(2,50)	$1.1032 \pm 0.2164i$	1.1242
(6,50)	$1.1048 \pm 0.2175i$	1.1260
(32,50)	$1.0626 \pm 0.2258i$	1.0864
(2,60)	$1.1182 \pm 0.1987i$	1.1358
(6,60)	$1.1212 \pm 0.1959i$	1.1382
(32,60)	$1.0671 \pm 0.2230i$	1.0902
(2,100)	0.6055, 0.9453	0.6055, 0.9453

Note: Each line of this table corresponds to a different estimation of the non-linear RP model. In each estimation we target a different range of periods.

Figure 11. Estimated and Data Spectral Densities for Various Estimations



Note: This figure shows estimated and actual spectral densities of hours (column (a)) and the risk premium (column (b)) when we target different period ranges.

We have also explored the effect of estimating the model only on frequencies between 2-32 quarters (results available upon request). This tends to favor inferring—as found in much of the literature—that persistent exogenous shocks drive business cycles. For example, estimating the linear RP model to fit only the 2-32 quarter range, the resulting shock process has an autoregressive parameter of 0.99. In comparison, when estimating this same model to fit the 2-50 quarter range, we get an autoregressive parameter of only 0.14.

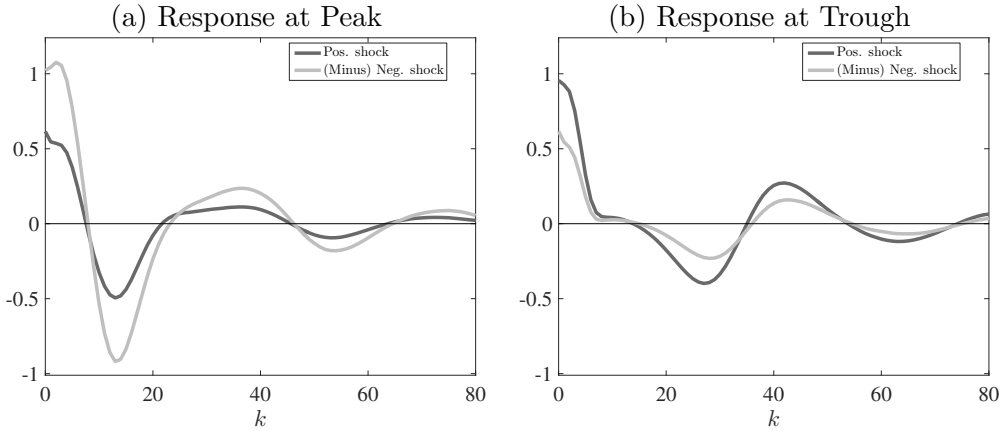
D. Is Allowing for Non-Linearities and Local Instability Important?

As noted above, the fit of the estimated linear and non-linear RP models are quite similar, as are a number of the estimated structrual parameters. One may then naturally ask: in what ways (if any) are the non-linearities important? We highlighted one way above, which is that allowing for non-linearities expands the parameter space to include parameterizations that produce limit cycles. If the dynamics of the economy does indeed feature limit cycles, by considering only a linear approximation (and the parameterizations that yield a valid rational expectations solution for this approximation), the estimation will necessarily be biased. For example, if the estimated non-linear RP model were the true model, but one employed the linear approximation to it (i.e., the linear RP model) in estimation, one would incorrectly conclude that the steady state is locally stable.

Allowing for non-linearities also has implications for (a)symmetry in the model. For example, in the estimated non-linear RP model, the business cycle is asymmetric, with booms lasting longer on average than recessions. This can be seen most clearly by looking at the deterministic component of the business cycle (i.e., the limit cycle illustrated in panel (a) of Figure 9), which features booms that last 23 quarters (trough to peak) and downturns in employment that last 15.5 quarters (peak to trough). There are also implications for the symmetry of the response of the economy to shocks. To illustrate this, panel (a) of Figure 12 plots the response of hours in the non-linear RP model to a one-standard-deviation positive shock, along with minus the response to a one-standard-deviation negative shock, conditional on initially being at the peak of the cycle. Panel (b) of the Figure plots the same except beginning from the trough of the cycle.[71] The responses are clearly different depending both on whether the shock is positive or negative, and on whether the economy is initially in a boom or a bust. One particularly interesting implication of the Figure is that when the economy is at the peak of a boom period and is hit by a negative (i.e., contractionary) shock, the magnitude of the response is substantially larger than if it is initially in a bust. This suggests that peak times could be periods where the economy is particularly sensitive to negative shocks (e.g., a financial disruption).

[71] See notes to Figure for further details.

FIGURE 12. HOURS IMPULSE RESPONSE AT CYCLE PEAK (NON-LINEAR RP)



Note: Figure shows the average response of hours in the non-linear RP model to a one-standard-deviation shock to ϵ_t as at date $t+k$, for $k = 0, \ldots, 80$, conditional on initially being at a peak (panel (a)) or trough (panel (b)) of the business cycle. Dark gray line shows response to a positive shock, light gray shows minus the response to a negative shock. Responses obtained as an average across 100,000 simulations, where for each simulation we compute the difference between a random simulation for the path of hours and the path that would occur if the first simulated shock were one standard deviation higher and the rest unchanged. For panel (a) we set the initial state as $(X_t, I_{t-1}, \mu_{t-1}) = (0.127, 1.731, 0)$, while for panel (b) we set it as $(-8.210, -6.030, 0)$, which correspond respectively to the peak and trough of the deterministic cycle.

IV. Conclusion

Why do market economies experience business cycles? There are at least two broad classes of explanations. On the one hand, it could be that market economies are inherently stable and that observed booms and busts are mainly due to persistent outside disturbances. On the other hand, it could be that the economy is locally unstable—or close to unstable—in that there are not strong forces that tend to push it towards a stable resting position. Instead, the economy's internal forces may endogenously favor cyclical outcomes, where booms tend to cause busts, and vice versa. The contribution of this paper has been to provide theory and evidence in support of this second view, while simultaneously highlighting the key elements that influence inference in this dimension.

We have emphasized several features that have led us to infer that business cycles may be generated by strong endogenous forces, as opposed to persistent exogenous shocks. However, in concluding, we would like to emphasize what we view as the most important issue for this debate: the question of what business cycle theory should aim to explain. If one adopts the conventional consensus that business cycle theory should be mainly concerned with movements in macroeconomic aggregates that arise at periodicities between 6 and 32 quarters, then standard models with weak internal propagation mechanisms can offer a reasonable explanation of the data. In contrast, if one agrees that business cycle theory should extend its focus to include slightly lower frequency movements, such as those associated with fluctuations of up to 50 quarters,[72] then the need to consider strong internal propagation mechanisms becomes much more relevant. In particular, we have documented that many macroeconomic aggregates appear to exhibit a peak in their spectral densities at periodicities between 32 and 50 quarters, and that the implied movements coincide with NBER cycle dating. Moreover, we have emphasized that such a pattern is very unlikely to have been a spurious draw from an AR(1) process. Given that cyclically sensitive variables such as unemployment and risk premia all exhibit such a peak, we believe that explaining this cyclical pattern should be a priority in business cycle analysis. If one accepts this, then, in our opinion, the existence of a strong internal propagation mechanism in the economy becomes more likely. While we have made an explicit case for this inference using one particular model, we conjecture that successfully explaining the observed humps in the spectral densities using any model will likely require it to feature complementarities that generate a strong internal cyclical mechanism.[73]

[72] A popular approach in the estimation of macroeconomic models is to include (almost) all frequencies. For example, this is the case when using likelihood-based methods to fit unfiltered (or, at most, first-differenced) data. In principle this is fine if the model is built to explain both business cycle fluctuations and the relatively large lower-frequency fluctuations associated with, for example, demographic changes. However, if the model is built to understand business cycles but is estimated using all frequencies, then the estimation may favor parameters that help to explain the large lower-frequency movements at the detriment of explaining business cycle movements. This point was illustrated in Section III.

[73] An alternative class of explanations would be one where the exogenous force itself is highly cyclical. While this is certainly possible, we have not been able to find evidence of cyclicality in variables commonly considered to be drivers of business cycles, such as technological change or oil prices.

Whether the resulting model delivers the more extreme form of endogenous propagation associated with limit cycles, or if instead it favors dampened fluctuations (as we observed when estimating a linear version of our model) will likely depend on model details. Nonetheless, in either case, we conjecture that the role played by shocks in driving the business cycle is likely to be greatly diminished if one attempts to explain the features of the data we have emphasized.

Appendix

A. Proof of Proposition 1

The proof requires showing that, under Assumption 2, $\alpha_3\alpha_4 > 0$, $\alpha_2 > 0$, and $\alpha_1 < 0$ are necessary for $g(\omega)$ to be humped-shaped on $\omega \in [0, \pi]$. Sargent (1987) (pp. 262-65) shows that for such a hump shape to arise it is necessary that the roots of $B(z) = 0$ be complex, and that if we write $B(L)$ as $1 - t_1 L - t_2 L^2$, then it must also be the case that $4t_2 + t_1(1 - t_2) < 0$.

The proof will proceed in two steps. The first step will be to show that $\alpha_2 > 0$ and $\alpha_1 < 0$ are necessary for the roots of $B(z) = 0$ be to complex under Assumption 2. The second step will be to show that $\alpha_3\alpha_4 > 0$ is necessary for $4t_2 + t_1(1 - t_2) < 0$ under Assumption 2. Before proceeding with these two steps, it is helpful to make explicit some of the implications of Assumption 2. It is straightforward to verify that Assumption 2 requires that

$$\alpha_1 < \delta\,(1 - \alpha_2)\ , \tag{A.1}$$

$$(1 - \delta + \alpha_1 + \alpha_2)^2 > 4\alpha_2\,(1 - \delta)\ , \tag{A.2}$$

which in turn imply

$$\alpha_1 + \alpha_2 < 1\ , \tag{A.3}$$

$$(1 - \delta + \alpha_1 + \alpha_2)\,[1 + \alpha_2\,(1 - \delta)] > 4\alpha_2\,(1 - \delta)\ . \tag{A.4}$$

First step:

For the roots of $B(z) = 0$ to be complex, there must exist a real $\phi \equiv \frac{1}{1-\alpha_3\alpha_4}$ such that

$$[1 - \delta + (\alpha_1 + \alpha_2)\,\phi]^2 < 4\alpha_2\,(1 - \delta)\,\phi\ . \tag{A.5}$$

Since the left-hand side of (A.5) is non-negative, it can only hold if the right-hand side is strictly positive. Since $\alpha_3\alpha_4 < 1$ by assumption, $\phi > 0$, and since $\delta < 1$, it follows that the right-hand side is strictly positive only if $\alpha_2 > 0$, which confirms the first necessary condition. Assume henceforth that indeed $\alpha_2 > 0$. Next, (A.5) is equivalent to

$$(1 - \delta)^2 + 2\phi\,(\alpha_1 - \alpha_2)\,(1 - \delta)\,\phi + (\alpha_1 + \alpha_2)^2\,\phi^2 < 0\ .$$

Since the left-hand side of this expression is a convex quadratic function of ϕ, a necessary condition for it to hold for some range of ϕ is that the roots of that quadratic function are real. One can verify that this is only the case if $(\alpha_1 - \alpha_2)^2 > (\alpha_1 + \alpha_2)^2$. Since $\alpha_2 > 0$, this is not possible if $\alpha_1 \geq 0$. Thus, we need $\alpha_1 < 0$ for $g(\omega)$ to be hump-shaped.

SECOND STEP:

To have $4t_2 + t_1(1 - t_2) < 0$, we need

$$[1 - \delta + (\alpha_1 + \alpha_2)\,\phi]\,[1 + \alpha_2\,(1 - \delta)\,\phi] < 4\alpha_2\,(1 - \delta)\,\phi\,,$$

which is in turn equivalent to the condition

$$4\alpha_2\,(1 - \delta) - (\alpha_1 + \alpha_2) - \alpha_2\,(1 - \delta)^2 > (1 - \delta)\left[\frac{1}{\phi} + (\alpha_1 + \alpha_2)\,\alpha_2\phi\right]\,. \tag{A.6}$$

Assumption 2 implies that the first of these two inequalities is reversed when $\phi = 1$ (see condition (A.4)), and therefore so is the second. Since only the right-hand side of (A.6) depends on ϕ, in order for (A.6) to hold for some $\phi < 1$ (i.e., some $\alpha_3\alpha_4 < 0$), the right-hand side of (A.6) must be increasing in ϕ over some range of $\phi < 1$. But this in turn requires $(\alpha_1 + \alpha_2) > 1/\phi^2$ for some $\phi < 1$, which is ruled out by Assumption 2 (see condition (A.4)). Hence, $\alpha_3\alpha_4 > 0$ ($\phi > 1$) is necessary for $g(\omega)$ to be hump-shaped. This completes the proof.

B. Proof Proposition 2

The eigenvalues of the system are given by the roots of $B(z^{-1}) = 0$, where $B(L) = 1 - [1 - \delta + (\alpha_1 + \alpha_2)\phi]L + \phi\alpha_2(1 - \delta)L^2$, and $\phi \equiv \frac{1}{1-\alpha_3\alpha_4}$. A sufficient condition for the system to be unstable is that $\phi\alpha_2(1-\delta) > 1$. Since $\alpha_2(1-\delta) > 1$ by Assumption 2, there must exist a $\phi^\star > 1$ for which the system is unstable, and accordingly there exists an $\alpha^\star < 1$ as stated in the proposition.

If the system has complex roots when it loses stability as ϕ increases, this will happen at the point where $\phi = \frac{1}{\alpha_2(1-\delta)}$. For the roots to be complex when $\phi = \frac{1}{\alpha_2(1-\delta)}$, it must be the case that $\left[1 - \delta + \frac{\alpha_1+\alpha_2}{\alpha_2(1-\delta)}\right]^2 < 4$. This will happen if $-\alpha_1 > \delta^2\alpha_2$ and $\alpha_2 > -\frac{\alpha_1}{(2-\delta)^2}$, that is, when α_1 is sufficiently negative and α_2 sufficiently positive.

REFERENCES

Azariadis, Costas, and Bruce Smith. 1998. "Financial Intermediation and Regime Switching in Business Cycles." *American Economic Review*, 88(3): 516–536.

Baxter, Marianne, and Robert G. King. 1999. "Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series." *The Review of Economics and Statistics*, 81(4): 575–593.

Beaudry, Paul, Dana Galizia, and Franck Portier. 2015. "Reviving the Limit Cycle View of Macroeconomic Fluctuations." National Bureau of Economic Research, Inc NBER Working Papers 21241.

Beaudry, Paul, Dana Galizia, and Franck Portier. 2016*a*. "Is the Macroeconomy Locally Unstable and Why Should We Care?" In *NBER Macroeconomics Annual*. Vol. 31, , ed. Jonathan A. Parker and Michael Woodford, 479–530. The University of Chicago Press.

Beaudry, Paul, Dana Galizia, and Franck Portier. 2016*b*. "Putting the Cycle Back into Business Cycle Analysis." National Bureau of Economic Research, Inc NBER Working Papers 22825.

Beaudry, Paul, Dana Galizia, and Franck Portier. 2018. "Reconciling Hayek's and Keynes' Views of Recessions." *Review of Economic Studies*, 85(1): 119–156.

Benhabib, Jess, and Kazuo 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.

Benhabib, Jess, and Kazuo Nishimura. 1985. "Competitive Equilibrium Cycles." *Journal of Economic Theory*, 35(2): 284–306.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist. 1999. "The financial accelerator in a quantitative business cycle framework." In *Handbook of Macroeconomics*. Vol. 1 of *Handbook of Macroeconomics*, , ed. J. B. Taylor and M. Woodford, Chapter 21, 1341–1393. Elsevier.

Boldrin, Michele, and Luigi Montrucchio. 1986. "On the Indeterminacy of Capital Accumulation Paths." *Journal of Economic Theory*, 40(1): 26–39.

Boldrin, Michele, and Michael Woodford. 1990. "Equilibrium models displaying endogenous fluctuations and chaos." *Journal of Monetary Economics*, 25(2): 189–222.

Borio, Claudio. 2014. "The financial cycle and macroeconomics: What have we learnt?" *Journal of Banking and Finance*, 45: 182 – 198.

Brave, Scott, and R. Andrew Butters. 2011. "Monitoring financial stability: a financial conditions index approach." *Economic Perspectives*, 1 Q: 22–43.

Burns, Arthur F., and Wesley C. Mitchell. 1946. *Measuring Business Cycles*. New York, NY:National Bureau of Economic Research.

Calvo, Guillermo A. 1983. "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics*, 12: 383–398.

Carlstrom, Charles T, and Timothy S Fuerst. 1997. "Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis." *American Economic Review*, 87(5): 893–910.

Christensen, Ian, and Ali Dib. 2008. "The Financial Accelerator in an Estimated New Keynesian Model." *Review of Economic Dynamics*, 11(1): 155–178.

Christiano, Lawrence J., and Robert J. Vigfusson. 2003. "Maximum likelihood in the frequency domain: the importance of time-to-plan." *Journal of Monetary Economics*, 50(4): 789–815.

Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno. 2014. "Risk Shocks." *American Economic Review*, 104(1): 27–65.

Comin, Diego, and Mark Gertler. 2006. "Medium-Term Business Cycles." *American Economic Review*, 96(3): 523–551.

Cooley, T., and E. Prescott. 1995. "Economic Growth and Business Cycle." In *Frontiers of Business Cycle Research.* , ed. T. Cooley, Chapter 1, 1–38. Princeton, New Jersey:Princeton University Press.

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

Day, Richard H. 1982. "Irregular Growth Cycles." *The American Economic Review*, 72(3): 406–414.

Day, Richard H. 1983. "The Emergence of Chaos from Classical Economic Growth." *The Quarterly Journal of Economics*, 98(2): 201–213.

Day, Richard H., and Wayne Shafer. 1987. "Ergodic Fluctuations in Deterministic Economic Models." *Journal of Economic Behavior and Organization*, 8(3): 339–361.

Fernandez-Villaverde, Jesus. 2016. "Lecture Notes for Graduate Macroeconomics." Lecture 12: RBC Models.

Frisch, Ragnar. 1933. "Propagation Problems and Impulse Problems in Dynamic Economics." In *Economic Essays in Honour of Gustav Cassel*. 171–205. London:George Allen & Unwin Ltd.

Galizia, Dana. 2018. "Saddle Cycles: Solving Rational Expectations Models Featuring Limit Cycles (or Chaos) Using Perturbation Methods." Working Paper.

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

Gourieroux, C., A. Monfort, and E. Renault. 1993. "Indirect Inference." *Journal of Applied Econometrics*, 8: S85–S118.

Grandmont, Jean-Michel. 1985. "On Endogenous Competitive Business Cycles." *Econometrica*, 53(5): 995–1045.

Grandmont, Jean-Michel. 1998. "Expectations Formation and Stability of Large Socioeconomic Systems." *Econometrica*, 66(4): 741–782.

Granger, C. W. J. 1966. "The Typical Spectral Shape of an Economic Variable." *Econometrica*, 34(1): 150–161.

Gu, Chao, Fabrizio Mattesini, Cyril Monnet, and Randall Wright. 2013. "Endogenous Credit Cycles." *Journal of Political Economy*, 121(5): 940 – 965.

Hamilton, James Douglas. 1994. *Time series analysis.* Princeton, NJ:Princeton Univ. Press.

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

Hommes, Cars. 2013. *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems. Cambridge Books*, Cambridge University Press.

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

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

Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit Cycles." *Journal of Political Economy*, 105(2): 211–48.

Kuznetsov, Yuri A. 1998. *Elements of Applied Bifurcation Theory.* Vol. 112 of *Applied Mathematical Sciences*. 2 ed., New-York:Springer.

Le Corbeiller, Ph. 1933. "Les systemes autoentretenus et les oscillations de relaxation." *Econometrica*, 1(3): 328–332.

Matsuyama, Kiminori. 1999. "Growing Through Cycles." *Econometrica*, 67(2): 335–347.

Matsuyama, Kiminori. 2007. "Credit Traps and Credit Cycles." *American Economic Review*, 97(1): 503–516.

Matsuyama, Kiminori. 2013. "The Good, the Bad, and the Ugly: An inquiry into the causes and nature of credit cycles." *Theoretical Economics*, 8(3): 623–651.

Mian, Atif R., and Amir Sufi. 2018. "Finance and Business Cycles: The Credit-Driven Household Demand Channel." National Bureau of Economic Research, Inc NBER Working Papers 24322.

Minsky, Hyman. 1986. *Stabilizing an Unstable Economy.* New Haven:Yale University Press.

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

Qu, Zhongjun, and Denis Tkachenko. 2012. "Frequency Domain Analysis of Medium Scale DSGE Models with Application to Smets and Wouters (2007)." In *DSGE Models in Macroeconomics: Estimation, Evaluation, and New Developments*. Vol. 28 of *Advances in Econometrics*, , ed. N. Balke, F. Canova, F. Milani and M. A Wynne, 319–385. Emerald Group Publishing.

Ramey, Valerie A., and Sarah Zubairy. 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data." *Journal of Political Economy*. forthcoming.

Sala, Luca. 2015. "DSGE Models in the Frequency Domain." *Journal of Applied Econometrics*, 30(2): 219–240.

Sargent, Thomas J. 1987. *Macroeconomic Theory.* . Second ed., Emerald Group Publishing Limited.

Scheinkman, José. 1990. "Nonlinearities in Economic Dynamics." *The Economic Journal*, 100(400): 33–48. Conference Papers.

Shleifer, Andre. 1986. "Implementation Cycles." *Journal of Political Economy*, 94(6): 1163–1190.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *The American Economic Review*, 97(3): 586–606.

Stock, James H., and Mark W. Watson. 1999. "Business cycle fluctuations in us macroeconomic time series." In *Handbook of Macroeconomics*. Vol. 1 of *Handbook of Macroeconomics*, , ed. J. B. Taylor and M. Woodford, Chapter 1, 3–64. Elsevier.

Wieland, V., E. Afanasyeva, M. Kuete, and J. Yoo. 2016. "New Methods for Macro-Financial Model Comparison and Policy Analysis, ." Vol. 2 of *Handbook of Macroeconomics*, 1241–1319. Elsevier.