

 Open access • Journal Article • DOI:10.1257/AER.20190789

Putting the Cycle Back into Business Cycle Analysis — [Source link](#)

[Paul Beaudry](#), [Dana Galizia](#), [Franck Portier](#)

Institutions: [Bank of Canada](#), [Carleton University](#)

Published on: 01 Jan 2020 - [The American Economic Review](#) (American Economic Association)

Topics: [Indirect Inference](#), [Limit cycle](#) and [Business cycle](#)

Related papers:

- [Putting the Cycle Back into Business Cycle Analysis](#)
- [Medium-Term Business Cycles](#)
- [Limit Cycles in Japanese Macroeconomic Data: Policy Implications from the View of Business Cycles](#)
- [Recessions and recoveries in real business cycle models](#)
- [The Asymmetric Business Cycle](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/putting-the-cycle-back-into-business-cycle-analysis-dw94xzisb9>

ADEMU WORKING PAPER SERIES

Putting the Cycle Back into Business Cycle Analysis

Paul Beaudry[†]

Dana Galizia[‡]

Franck Portier[‡]

March 2018

WP 2017/092

www.ademu-project.eu/publications/working-papers

Abstract

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 new evidence in favour of the second interpretation and, most importantly, we highlight the set of key elements that influence our answer to this question. In particular, when adopting our most preferred estimation framework, we find support for the somewhat extreme notion that business cycles may be generated by stochastic limit cycle forces; that is, we find support for the notion that business cycles may primarily reflect an endogenous propagation mechanism buffeted only by temporary shocks. The three elements that tend to favour 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. We document the sensitivity of our findings to each of these elements within the context of an extended New Keynesian model with real-financial linkages.

Jel Codes: E3, E32, E24

Keywords: Business Cycle, Limit Cycle

[†] Vancouver School of Economics, University of British Columbia and NBER

[‡] Department of Economics, Carleton University

[‡] Department of Economics, University College London and Toulouse School of Economics, Université Toulouse Capitole and CEPR

Acknowledgments

The authors thank Jess Benhabib, Kiminori Matsuyama and Morten Ravn for helpful discussions. The authors would also like to thank seminar participants at Banque de France, University of Manchester, Pompeu Fabra-Toulouse 'skiminar', NBER Summer Institute, Hydra workshop, UCL, UCLA, UCSD, Wisconsin, the University of Pennsylvania, Oxford University, EUI Florence, EIEF Roma and Northwestern University for comments. Franck Portier acknowledges financial support by the ADEMU project, 'A Dynamic Economic and Monetary Union', funded by the European Union's Horizon 2020 Program under grant agreement No 649396.

The ADEMU Working Paper Series is being supported by the European Commission Horizon 2020 European Union funding for Research & Innovation, grant agreement No 649396.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License Creative Commons Attribution 4.0 International, which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.



Introduction

Market economies repeatedly go through periods where, for sustained lengths of time, productive factors are used very intensively—with low rates of unemployment, high levels of hours worked per capita, and intensive use of productive capital—followed by periods where these utilization rates are reversed. The type of forces and mechanisms that drive those 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 such models, shocks are generally transformed due to the structure of the economy by features such as adjustment costs, capital accumulation and habit. However, it remains the case that if the persistence of the exogenous shocks were to be reduced, the prediction of such models would be that business cycles type fluctuations would essentially disappear. It is for this reason that such models are interpreted as supporting the notion that persistent exogenous shocks are central to understanding business cycles. On the other hand, there is the alternative view whereby the bulk of business cycle fluctuations are believed to be the result of forces internal to the economy 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 prolonged business cycles nevertheless arises due to the equilibrium incentives it creates.¹

Even though from a theoretical point of view both 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.

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 its auto-correlation or its variance— that is lost if one reduces the persistence of 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.

² Certain credit cycle models, such as the seminal piece by ?, 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.

Here, it is argued that the autocovariance patterns of 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 is argued, have remained elusive.

The object 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 capacity utilization, and risk premia. Using a high-resolution technique based on zero padding, we document a sizable peak in several spectral densities in US macroeconomic data at periodicities around 9 to 10 years. We also show that standard business cycle models are inconsistent with this property of the data. While the presence of such a peak does not necessarily imply strong endogenous cyclical forces³, it is an important first step in our argument, as it runs counter to the notion that the spectral properties of the data rule out such a possibility. Moreover, these spectral densities play a central role in our later structural estimation exercises.

In a second step we present a class of models that highlights key economic forces that can endogenously generate strong cyclical forces that are in turn capable of producing the spectral peaks documented in the first step. There are two main ingredients in this class of models: strategic complementarities across agents, and the accumulation of some form of capital. These two properties are found in many macroeconomic models, which makes this class of models particularly relevant. We use a generic model structure in order to illustrate how commonly used estimation methods may be biased against uncovering the full force of the endogenous cyclical mechanisms generated within such models. For example, we show that this class of models may easily give rise to limit cycles—that is, to a situation where even in the absence of shocks the economy would generate cycles—but that many standard empirical approaches would fail to uncover this feature when it is present in the data.

In the third step we present a fully specified dynamic general equilibrium model that falls within the general class of models discussed in the second step. The model contains real-financial linkages that produce complementarities, and embeds a standard New Keynesian model as a special case when the financial imperfections are shut down. In the model, unemployment risk, default risk and risk premia on borrowing are jointly determined. The model is aimed at capturing the common narrative of an accumulation-credit cycle, where booms

³One could obtain such spectral density pattern with strongly cyclical exogenous forces.

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 for this stochastic force to compete with the endogenous forces in order to evaluate their respective roles. We use this model to illustrate the different inferences one would make regarding the relative importance of exogenous versus endogenous forces in driving business cycle movements depending on how one approaches estimation. In particular, we show that if one targets the spectral densities documented in our first step, and if one adopts an estimation method that allows for a locally unstable steady state, then the estimation results suggest that business cycles are mainly driven by internal forces buffered 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 as 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 view our main contribution as illustrating 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 different estimation methods can greatly influence one’s inferences regarding the relative contribution of internal versus external mechanisms in driving business cycles. For example, we show how, when adopting an approach to estimation 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 im-

⁴It should be emphasized that there are any number of 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 as we believe they likely play a role in business cycle fluctuations (see Section 1.4), we by no means wish to rule out the possibility of other relevant sources of complementarity.

⁵The endogenous boom-bust cycles that arise in our model do not reffect multiple equilibria or indeterminacy.

⁶While 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.

portance 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 as we show how one’s inference can change from a situation where the economy appears to be driven primarily by exogenous shocks to one 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 ?, ?, ? and ?.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., ? and [?], ? and [?], ?, ?, ?; for surveys of the literature, see ? and ?). By the early 1990s, however, the interest in such models for understanding business cycle fluctuations greatly diminished and became quite removed from the mainstream research agenda.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 case nor a strong theoretical one for why they should be considered to be as or more relevant than the alternative explanations. An important contribution of this paper can therefore be seen as reviving the limit cycle view of fluctuations by offering new perspectives on these two arguments. In particular, with respect to the first argument, we directly address the criticism of the excessive regularity of limit cycles by examining instead the notion of a stochastic limit cycle, where the system is subject to exogenous shocks, but where the deterministic part of the system admits a limit cycle. Such systems have been studied little by quantitative macro-economists, but recent solution techniques now make this a tractable endeavor. We address the second argument, meanwhile, by presenting a whole class of simple models that are capable of exhibiting stochastic limit cycles, and showing that this class of models finds some support in the data.

The remaining sections of the paper are organized as follows. In Section 1 we document

⁷An earlier mention of self-sustaining cycles as a model for economic fluctuations is found in ? in the first volume of *Econometrica*.

⁸There 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 ? and ?), and a literature on endogenous credit cycles in an OLG setting (see, for example, ?, ? and [?], ? and ?). One should also mention a large literature on endogenous business cycles under bounded rationality and learning, following early ideas of ?. ? reviews this literature and the debate on endogenous business cycles versus exogenous shocks, particularly the role of stochastic limit cycles and (noisy) chaos in that framework.

the spectral properties of U.S. data on hours worked, unemployment, capacity utilization, and several indicators of financial conditions. These properties motivate our analysis and will be used later on in estimating our model. In Section 2 we present a simple dynamic set-up where agents both accumulate goods and interact strategically with one another. Following ?, these strategic interactions can be characterized either by substitutability or complementarity. We use this framework to highlight when complementarities are likely to produce deterministic cyclical behavior, while simultaneously explaining why such an outcome is more likely than outcomes associated with multiple equilibria or indeterminacy. In Section 3, we extend a standard three-equation New Keynesian model in a manner that allows for the features highlighted in model of Section 2. We present 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 stochastic limit cycles. Finally, in the last section we offer concluding comments.

1 Motivating Observations

1.1 U.S. Post-War Business Cycle

In this subsection we examine the cyclical properties of a set of quarterly U.S. macroeconomic variables covering the period 1947Q1-2015Q2.⁹ One potential way of describing the cyclical properties of stationary data is to focus on the spectral density of the series, which depicts the importance of cycles of different frequencies in explaining the data. As is well known, if the spectral density of a time series displays a substantial peak at a given frequency, this is an indication of recurrent cyclical phenomena at that frequency. The traditional view, as expressed for example in ? and ?, is that most macroeconomic time series do not exhibit peaks in their spectral densities at business cycle frequencies. This view suggests that business cycle theory should not seek to explain macroeconomic fluctuations as recurrent phenomena, where a boom may sow the seeds of a subsequent bust, but instead should focus mainly on explaining the co-movement properties of macro-variables. In this section, we aim to re-examine the validity of this consensus motivated by three key observations. First, using spectral densities to evaluate business cycle properties requires stationary data, while many macroeconomic variables are non-stationary. It has therefore been common to substantially transform (i.e., detrend) 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 compo-

⁹Sources for all data series are discussed in Appendix A.

ment. For this reason, the use of spectral methods is most easily applied to macroeconomic variables that can plausibly be argued to be stationary before any transformation.¹⁰ For example, this may be the case for labor market variables such as employment or unemployment rates. Second, the traditional definition of the business cycle focuses on movements in macroeconomics variables at periodicities between 6 and 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 38 quarters so far, having 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. Third, much of the literature using spectral methods to look at US post-war business cycles is rather old and relied on time series substantially shorter than those currently available. Thus, a new look seems in order.

We begin our re-evaluation of US business cycle properties by focusing on (the log of) U.S. non-farm business (NFB) hours worked per capita, plotted in panel (a) of Figure 1. As the Figure shows, over the sample period hours exhibited substantial fluctuations, but with limited evidence of any long-run trend. For this reason, it does not seem unreasonable to treat this series—at least as a first approximation—as stationary. Accordingly, we begin by looking directly at the spectral density of this series without any initial transformation (except de-meaning). The dark line in panel (b) of Figure 1 plots this spectral density over the range of periodicities from 4 to 80 quarters.¹¹ 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—we also plot in the Figure the spectra obtained after 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 below 50 quarters—the range of periodicities we focus on henceforth—are very

¹⁰If a series is non-stationary and it is known to be $I(1)$, then it can be reasonable to look at the spectrum of the first difference of the series. However, in such a case, the application of the first difference filter may substantially mask the properties of any cyclical component that is at somewhat low frequencies.

¹¹We 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 H in the Appendix for more details). For the sake of clarity, we do not superimpose confidence bands for the estimated spectral density. In Section H in the Appendix, we show those confidence bands for the four variables of Figures 1 and 2.

robust to whether or not one first removes very low-frequency movements from this series.

[Figure 1 about here.]

What does panel (b) of Figure 1 reveal about the cyclical properties of hours? To us, the dominant feature is the distinct peak in the spectral density at around 38 quarters, which is much more pronounced than anything found at periodicities less than 32 quarters. This suggests that a significant proportion of the fluctuations in hours is coming from some cyclical force with a periodicity of about 9-10 years. Note that this spectral peak is mainly contained in the from 32- to 50-quarter range, and therefore just slightly beyond the traditional range (6-32 quarters) usually thought of as capturing the business cycle. Note also that this peak is not capturing the medium-run phenomena emphasized in work by ?, who focus indistinctly on cycles between 32 and 200 quarters. In our opinion, this peak should be thought of as part of the business cycle, suggesting that the traditional definition of the business cycle may be slightly too narrow. To make this case more salient, in panel (c) of Figure 1 we plot the hours series after having removed fluctuations longer than 60 quarters using the high-pass filter. In the Figure, we also highlight NBER recessions in gray. As can be seen, the fluctuations in detrended hours that we observe when retaining these slightly lower frequencies match very closely the standard narrative of the business cycle, rather than reflecting, as in ?, movements unrelated to it.

To further support the notion of a peak in the spectral density of factor usage near 40 quarters, in panels (a) and (b) of Figure 2 we redo the same exercise for two other measures of work activity: total hours worked per capita (panel (a)), and the unemployment rate (panel (b)). In addition, in panel (c) we report the spectral density of a (physical capital) capacity utilization measure. In all three cases, we plot both the spectral density for the untransformed data (dark line), as well as a set of spectra obtained after first removing low-frequency movements using high-pass filters as in Figure 1 (light gray lines). We highlight in dark gray the band of periodicities from 32 to 50 quarters.¹² In all of these plots we see distinct peaks in the spectra around 36-40 quarters, regardless of whether or not we first remove very low-frequency movements. Together, panel (b) of Figure 1 and panels (a)-(c) of Figure 2 indicate that the aggregate utilization of workers and capital by firms in the U.S. exhibits important recurrent cyclical phenomena at approximately 9 to 10 year intervals.

[Figure 2 about here.]

¹²Given our stated focus on periodicities below 50 quarters, we cut the x -axis off at 60 quarters. For longer periodicities, as we have seen in panel (b) of Figure 1, high-pass filtering is likely to affect the spectral density estimates.

1.2 Spectral Implications of Standard Models

In Figure 1, with further support in Figure 2, we have shown that several cyclically sensitive macroeconomic variables exhibit substantial spectral peaks at a periodicity of around 38-40 quarters. In this subsection, we argue that this observation is a challenge for many business cycle models. To illustrate this, we first report in panel (a) of Figure 3 the spectral density for hours implied by the simplest Real Business Cycle model described in ?. In that stripped-down model, capital utilization is constant and fluctuations only come from persistent technology shocks. In panels (b) and (c) we show the spectra of hours and capital utilization after augmenting the model to include variable capital utilization as well as investment-specific technology shocks as a second source of fluctuations.¹³ Panels (d) and (e) report the same results using the rich New Keynesian model of ?, which features seven shocks and a variety of real and nominal frictions. In all cases we see a similar pattern: the spectra do not exhibit any peaks, being instead very similar to what one would obtain from a simple AR(1) process. This observations should not be too surprising: as is well known, the internal propagation of these models is rather weak and therefore the endogenous variables largely inherit the properties of the exogenous driving forces, which in these cases are mainly AR(1) processes. Although these results 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 in the data. Further, we have also checked whether the spectral shape we document could have been a spurious draw from a model in which hours are in fact driven by an AR(1) process. In particular, to explore this possibility, we first estimated an AR(1) process on our NFB hours series, then drew 1,000 samples of the same length as our data set (270 quarters) from this estimated process, and then estimated spectral densities for each of those artificial series. We found a peak in the spectral density of the size observed in the data at a periodicity below 50 quarters in less than 1 percent of the simulations. It therefore appears very unlikely that the observed peak in the data could have been generated by a model with a form close to that of an AR(1) process.

[Figure 3 about here.]

1.3 An Alternative Approach Using NBER Recession Dates

The pattern we documented for the spectrum suggests the presence of forces that favor recurrent booms and busts roughly every 9-10 years. In this section we provide further support for this idea by directly using NBER recession dates instead of spectral methods.

¹³This augmented model is taken from ?.

One of the attractive features of NBER recession dates is that they are synthesis of a number of different variables, and are clearly associated with standard ideas about business cycle activity. In Figure 4, 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 an (NBER) recession at time t . The Figure plots this probability as we vary k between 12 and 80 quarters, using all NBER recession dates from 1946:1 to 2017:2. We look in an x -quarter window around date $t+k$ since NBER recessions are rather short-lived, and show results for different window widths ranging from 1 to 10 quarters.

As can be seen in the Figure, regardless of the window width, a rather clear pattern emerges. The probability that the economy is in a recession k quarters out given that it is in recession at time t increases almost monotonically as k goes from 12 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, nicely echoes the recurrence suggested by the previous spectral results. 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. This pattern again emphasizes the fact that business cycle analysis should not be limited to fluctuations less than 8 years, but should instead include movements with periodicities of at least 10 years.

[Figure 4 about here.]

1.4 A Related Financial Cycle?

We now look at the spectral properties of a set of financial market indicators. So that the exercise is a sensible one, we again restrict attention to trend-less variables, and accordingly a discussion of the appropriate filtering method is unnecessary.¹⁴ In particular, we wish to see whether indicators of financial market conditions also exhibit peaks in their spectral densities at frequencies similar to those observed for factor utilization variables. 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 emphasize previously. To this end, we report spectral estimates for four financial market indicators in Figure 5. The measures in panels (a) and (b) are the Chicago Fed's National Financial Conditions Index (NFCI) and its Risk Subindex, respectively. As described in ?, the NFCI is computed over a long time horizon and from a large sample of financial indicators. It is a synthetic index between -1 and 1 that attempts to summarize financial conditions. The Risk Subindex captures volatility and

¹⁴We again complement the results computed from the variable in levels with results derived after first filtering the data with a high-pass filter in order to illustrate the robustness of the results.

funding risk in the financial sector. We complement these two synthetic indices with two more transparent series. In particular, in panel (c) we show the spread between the federal funds rate and BAA bonds. This is a common measure of risk discussed throughout the macro-financial linkage literature, and is often simply referred to as a measure of the market risk premium. Finally, in panel (d) we show delinquency rates on all loans to commercial banks.¹⁵ As is clear from the Figure, the spectral density peak we have identified near 40 quarters for real variables is also present for all four of these financial variables. For the two synthetic indices, the peaks are slightly more spread out than for employment, while for the BAA spread there are two peaks, with the one around 40 quarters being the largest.¹⁶ Despite these slight variations, the overall picture that emerges from Figure 5 suggests a close link between the cycle in employment and the cycle in financial market conditions.¹⁷ This suggests that, when looking to explain these patterns, it may be appropriate to turn to a model with financial-real linkages.

[Figure 5 about here.]

One of the potential drawbacks of our spectrum estimates is that they are based on only 70 years of data. In the Appendix D, we therefore examine the robustness of our results to longer historical samples.

2 Explaining Spectral Peaks: A Class of Models

In the last section we documented how several (trend-less) macroeconomic variables—most importantly employment measures—exhibit a peak in their spectral densities at a periodicity of around 36-40 quarters. We also showed that this feature of the data is, at least for employment, not well captured by standard macroeconomic models, which typically generate hours spectra that are monotonically increasing, resembling closely that of an AR(1) process. In fact, the hump-shaped spectrum observed in the data is more akin to that generated by an AR(2) process with complex roots that have modulus near one.¹⁸ For example, it is easy to verify that the AR(2) process $X_t = 1.89X_{t-1} - 0.92X_{t-2}$ has a spectral density that, at periodicities below 50 quarters, is quite similar in shape to the one we have documented for hours. While there may be many mechanisms capable of generating such a shape, it is worth

¹⁵The limitation of this latter measure is that it is available only since 1985.

¹⁶The lower peaks arrives almost exactly at half the periodicity (a harmonic) of the larger peak. Such a pattern may reflect a non-linearity in the data.

¹⁷We also find the coherence between these real and financial series to be very high.

¹⁸See ? for the precise conditions under which an AR(2) process has a hump shaped spectral density.

pointing out a class of models that will *not* produce this feature. In particular, suppose we have a variable X_t that has a solution in state space form given by

$$X_t = a_1 X_{t-1} + a_2 Z_t, \quad 0 < a_1 < 1, \quad (1)$$

$$Z_t = \rho Z_{t-1} + \epsilon_t, \quad 0 < \rho < 1, \quad (2)$$

where Z_t is an exogenous driving force with innovations ϵ_t . In this case, one may verify that X_t is an AR(2) process, but one that will never have complex roots, and therefore will generally not exhibit a peak in its spectral density. Hence, models that mainly emphasize dynamics driven by a combination of endogenous sluggishness (as captured by the parameter $a_1 \in (0, 1)$ in (1)) and an exogenous driving force that itself does not have a peak in its spectrum (e.g., the AR(1) process (2)) will not be able to explain the recurrent cyclical behavior implied by the local peak in the spectrum.

In effect, the above example suggests that, to generate hump-shaped spectral densities, one likely needs to rely on either a strong endogenous propagation mechanism capable of creating recurrent cyclical forces (e.g., a model of the form $X_t = a_1 X_{t-1} + a_3 X_{t-2} + a_2 Z_t$ such that the characteristic equation $\lambda^2 - a_1 \lambda - a_3 = 0$ has complex roots), or on an exogenous process that itself has a hump-shaped spectral density. In particular, it appears unlikely that an interaction between a non-cyclical endogenous propagation mechanism and a non-cyclical exogenous driving force in a more complex setting than that given by (1)-(2) will generate hump-shaped spectral densities for endogenous variables. For this reason, the most promising route to us appears to be to look for an internal propagation mechanism that produces recurrent cyclical behavior (as captured by complex roots); that is, to look for models where a boom tends to endogenously sow the seeds of the next bust.¹⁹ Our goal in Section 3 will be to present and explore one such model in which financial frictions play a central role. These frictions will produce strategic complementarities between agents that favor bunching of economic activity, which helps to generate recurrent cyclical patterns. However, it turns out that dynamic models with complementarities introduce challenges for model solution and estimation that are not well recognized in the literature. For this reason, before describing the microeconomic structure behind our particular model, we want to highlight several key properties of a general class of dynamic models featuring complementarities, of which our later model is a special case. We discuss these properties first as we believe they are of general interest, and because they will influence how we approach the estimation of our specific model in Section 3.

¹⁹One could alternatively look for evidence of exogenous driving forces that have the required cyclical property, but we view such a path as less promising.

2.1 Demand Complementarities as a Source of Cyclicity

At least since the work of ?, it has been recognized that models with complementarities may help explain macroeconomic behavior. In this section we build on the work of ? by analyzing a class of dynamic models featuring complementarities. In particular, we emphasize several properties and challenges that arise when these complementarities interact with a process of accumulation. First, we confirm the intuition that this simple class of models ubiquitously generates endogenous cyclical behavior—as captured by the emergence of complex roots in the endogenous propagation mechanism—as the strength of the complementarities increase. In this sense, we show why these models may be relevant for explaining the cyclical patterns we documented above for employment and other macroeconomic variables. Second, contrary to our prior expectations, this class of models does not produce indeterminacy or sunspot equilibria as the complementarities become stronger. However, it does produce a different challenge: the steady state can potentially lose local stability. In particular, if the environment is stochastic and the system loses local stability (e.g., if the system experiences a Hopf bifurcation²⁰) as the complementarities strengthen, then a stochastic limit cycle may emerge. If an econometrician is not aware of such a possibility, it can bias their estimation of the forces driving cyclical behavior. Finally, and again contrary to our prior beliefs, as this transition from local stability to limit cycles occurs in a stochastic environment, nothing qualitatively drastic occurs. Rather, the model’s statistical properties change smoothly. For this reason, we believe that the commonly analyzed locally stable models and those that exhibit stochastic limit cycles should not be viewed as conceptually distinct frameworks for thinking about the business cycle. Instead, these two classes of models should be seen as close cousins, where observable differences would be mainly quantitative.

2.1.1 The Environment

Consider an environment with a large number N of agents indexed by i , where each agent can accumulate a good X_{it} , which can be either productive capital or a durable consumption good. The accumulation equation is given by

$$X_{it+1} = (1 - \delta)X_{it} + I_{it}, \quad 0 < \delta < 1, \quad (3)$$

where I_{it} is agents i ’s investment in the good. Suppose the decision rule for agent i ’s investment is given by

$$I_{it} = \alpha_0 - \alpha_1 X_{it} + \alpha_2 I_{it-1} + \alpha_3 \mathbb{E}_{it} [I_{it+1}] + \mathbb{E}_{it} [F(I_t)] \quad (4)$$

²⁰Since our model will be formulated in discrete time, the bifurcation we consider is more appropriately referred to as a Neimark-Sacker (rather than Hopf) bifurcation. Nonetheless, we will typically follow convention in applying the term “Hopf bifurcation” to our discrete environment.

where parameters α_1 and α_3 are strictly between 0 and 1, $\alpha_2 \in [0, 1)$, and \mathbb{E}_{it} is the expectation operator summarizing agent i 's beliefs at date t . Here we do not present the primitives that give rise to this decision rule, since we do not want to take a precise stance on whether this comes from a consumer problem (in which case the accumulation would involve durable goods), or if it comes from a firm problem (where it would involve productive capital). Instead, we simply posit a decision rule that embeds, through X_{it} , I_{it-1} and $\mathbb{E}_{it}[I_{it+1}]$, three forces that are generally considered important for individual-level investment decisions. First, the presence of X_{it} reflects some underlying decreasing returns to accumulation, and accordingly is assumed to enter negatively. Second, an agent's past investment I_{it-1} is assumed to enter positively to reflect a sluggish response due, for example, to adjustment costs.²¹ Third, agent i 's expectation of their future investment I_{it+1} enters positively so as to capture typical forward-looking behavior.

In addition to allowing individual-level variables to affect investment decisions, we also allow agents to affect each other through the term $\mathbb{E}_{it}[F(I_t)]$, where $I_t \equiv \sum_j I_{jt}/N$ is the average level of investment. Thus, agents' actions depend in some way on their expectations of *other* agents' actions, as summarized by the average level of investment I_t . The nature of this dependence is captured by the function $F(\cdot)$, which we assume is thrice continuously differentiable.²² For example, if $F'(\cdot) < 0$ then $F(\cdot)$ may capture an agent's optimal response to an increase in prices caused by increased demand by others, while $F'(\cdot) > 0$ would capture some form of demand complementarity.²³ We henceforth assume that $F'(I) < 1$, so that demand complementarities, if they are present, are never strong enough to produce multiple temporary equilibria.²⁴

Since we focus in what follows on symmetric behavior, and since there are no stochastic driving forces yet,²⁵ we may henceforth drop the subscript i and the expectations operator from (4). Since the model contains forward-looking agents, we also assume the presence of a transversality condition of the form $\lim_{t \rightarrow \infty} \beta^t X_t = 0$ for some parameter $0 < \beta < 1$. Our goal is to examine how the dynamics of the system (3)-(4) are affected by the properties of the interaction effects, and in particular to make clear the relationship between $F(\cdot)$ and the emergence of cyclicity, which we define as a situation where, along almost

²¹We focus here on a system with two pre-determined variables so as to allow mathematically for the possibility of complex roots, which is our main focus.

²²When the steady state of (3)-(4) is unique, we also normalize F so that $F(I^s) = 0$, where I^s denotes the steady state level of aggregate investment.

²³In general, both of these forces (and potentially others) could be present, in which case F would capture the combined net effect.

²⁴By a temporary equilibrium we mean an equilibrium for I_{it} given values of X_{it} , I_{it-1} and $\mathbb{E}_{it}[I_{it+1}]$.

²⁵We discuss the stochastic case below in Section 2.2.

every deterministic transition,²⁶ deviations from the steady state change sign at some point. We will mainly focus on the case of “smooth” cyclicity—that is, cyclicity characterized by booms and busts that each last for more than one period (on average)—since this is likely to be the more empirically relevant case. In addition to establishing that the model can, under very general conditions, generate this smooth cyclicity, we also show that, under certain conditions, as the complementarities become sufficiently strong this smooth cyclicity transitions from a “weak” (i.e., convergent) to a “strong” form; that is, a Hopf bifurcation will occur, causing the steady state to become locally unstable and a limit cycle to emerge. Importantly, we will show that this happens *even though we restrict the strength of the demand complementarities to be too weak to create multiple equilibria or indeterminacy*.

In order to understand these dynamics, it is useful to first look at the local dynamics in the neighborhood of the steady state. As a first step, we make the following assumption to guarantee the existence of a unique steady state, so that our discussion of the local dynamics around the steady state is itself uniquely defined.²⁷

Assumption (A1). $\alpha_1 > \delta(\alpha_2 + \alpha_3)$.

Note that (A1) will be satisfied if, for example, the depreciation rate δ is sufficiently small. We then have the following proposition.

Proposition 1. *If (A1) holds, then the steady state of system (3)-(4) is unique.*

All proofs are given in Appendix G. Next, we parameterize F by $\rho < 1$ such that $F'(I^s) = \rho$. The first-order approximation of (3)-(4) around the steady state is then

$$\begin{pmatrix} X_{t+1} \\ I_{t+1} \\ I_t \end{pmatrix} = \underbrace{\begin{pmatrix} 1 - \delta & 1 & 0 \\ \frac{\alpha_1}{\alpha_3} & \frac{1-\rho}{\alpha_3} & -\frac{\alpha_2}{\alpha_3} \\ 0 & 1 & 0 \end{pmatrix}}_M \begin{pmatrix} X_t \\ I_t \\ I_{t-1} \end{pmatrix}, \quad (5)$$

where all variables are now expressed in deviations from steady state. Since M is a 3×3 matrix, it has at least one real eigenvalue. Let λ_2 denote the largest real eigenvalue of M , and $\lambda_{11}, \lambda_{12}$ the remaining two eigenvalues, with $\lambda_{11} \leq \lambda_{12}$ when they are real, and λ_{11} the eigenvalue with negative imaginary part when they are complex. When the solution to (3)-(4) is determinate, its qualitative dynamics will be dictated by two of these three eigenvalues. In particular, positive eigenvalues will be associated with monotonic dynamics, complex

²⁶By this we mean, along every deterministic transition except possibly for a measure-zero subset of them.

²⁷The case of multiple steady states could also be interesting, but to keep our analysis manageable we focus only on cases with unique steady states.

eigenvalues with smooth cyclicity, and negative eigenvalues with non-smooth cyclicity (period 2 cycles). If exactly two of the three eigenvalues are stable (i.e., inside the complex unit circle), then the solution will either exhibit monotonic convergence or weak cyclicity. If all three of the eigenvalues are stable then the solution will be indeterminate. Finally, if none or only one of the three eigenvalues are stable, then the linearized system (5) has no solution, though this does not necessarily imply that the original *non-linear* system has no solution, and in particular, as we shall see, there may in some cases be a solution featuring limit cycles.

2.1.2 Main Results

In this subsection we present our main results (contained in Propositions 2 through 5) concerning the relationship between the nature of agents' interactions (governed by F) and the qualitative dynamic behavior of the system. In particular, we suppose that, in the absence of agent interactions (i.e., with $F(I) = 0$, so that $\rho = 0$), we have determinate non-cyclical convergence (i.e., three real positive eigenvalues with exactly two of them stable); that is, without agent interactions, individual-level decisions do not exhibit cycles. This appears reasonable to us as most individual-level concave optimization problems indeed produce determinate non-cyclical dynamics. Hence, if cyclicity is to emerge, it will have to be due to some form of agent interaction. To explore this, we will mainly focus on the consequences of changing the degree of local complementarity, ρ , beginning from point $\rho = 0$ (maintaining our assumption that $\rho < 1$ at all times).

Our first main result establishes a key relationship between the degree of local complementarity and the emergence of weak cyclicity.

Proposition 2. *Suppose (A1) holds, and initially $\rho = 0$ with $0 < \lambda_{11}, \lambda_{12} < 1 < \lambda_2$.²⁸ Then:*

- (i) *As local complementarity ρ decreases (without bound), the system continues to exhibit determinate monotonic convergence.*
- (ii) *As local complementarity ρ increases, at some point $\lambda_{11}, \lambda_{12}$ become complex and stable (i.e., weak smooth cyclicity emerges), with λ_2 remaining greater than 1 (i.e., the system remains determinate).*

Proposition 2 begins by establishing that, starting from a situation with no agent interaction and no cyclicity, introducing substitution forces across agents (via $\rho < 0$) has no effect

²⁸While we focus on the case of determinate monotonicity for $\rho = 0$, all results in this section apply equally when changing ρ beginning from *any* situation featuring determinate monotonicity (i.e., regardless of the initial ρ). Further, it can be shown that such a situation always exists for ρ sufficiently small.

on the qualitative dynamics of the system. In other words, if in the absence of agent interaction the system does not cycle, then introducing substitutability will not cause cyclicity to emerge. This explains why, in general equilibrium models, agent interactions fully transmitted through price effects will not typically tend to create cyclicity, since price effects tend to produce substitutability. Thus, if cyclical behavior is to arise in our environment, it will require complementarity.

The second part of Proposition 2 establishes that, as local complementarity *increases*, cyclicity always eventually emerges, and when it emerges it does so in its weak smooth form. Thus, in a model with accumulation, having sufficiently strong complementarity in agents' actions is sufficient to generate cyclicity. Proposition 2 further establishes that the system continues to be determinate when this cyclicity emerges.²⁹

We now turn to investigating what happens as ρ continues to increase after this weak smooth cyclicity first emerges. This is given in our second key proposition of this section.

Proposition 3. *Suppose (A1) holds. Then, beginning from the situation of determinate weak smooth cyclicity noted in Proposition 2, as ρ increases all the way to 1, λ_2 remains greater than 1 (no indeterminacy), and one of the following three possibilities arises:*

- (i) $\lambda_{11}, \lambda_{12}$ remain complex and stable (weak smooth cyclicity), or*
- (ii) at some point a Hopf bifurcation occurs, at which point $\lambda_{11}, \lambda_{12}$ become complex and unstable and a limit cycle emerges (strong smooth cyclicity), or*
- (iii) at some point $\lambda_{11}, \lambda_{12}$ become real, negative and stable (weak non-smooth cyclicity).*

Proposition 3 establishes two important facts. First, since there is always at least one unstable eigenvalue (i.e., λ_2), we never have local indeterminacy—or the attendant self-fulfilling booms or busts—as ρ increases towards 1. This may seem surprising given that the literature on indeterminacy often seems to draw on complementarities to generate multiplicity. Under assumption (A1), however, the fact that $\rho < 1$ rules out this behavior; that is, if there are demand complementarities, they are too weak to generate indeterminacy. Second, as ρ increases from the point where weak smooth cyclicity first emerges, either (i) the system continues to feature weak smooth cyclicity as ρ increases all the way towards 1, (ii) at some point the system becomes locally unstable and a limit cycle featuring strong smooth cyclicity emerges, or (iii) at some point weak non-smooth cyclicity will emerge.

²⁹Note that, since this cyclicity arises without the loss of determinacy, it is quite different from that found in ? and the related literature.

Since non-smooth cyclicity does not appear to be an empirically relevant situation, it is of interest to know under what conditions we can rule out the case highlighted in Proposition 3(iii). To this end, consider the following additional assumption.

Assumption (A2). $\alpha_2 > \alpha_1 - 2(1 - \delta)^2 \alpha_3 \left[\sqrt{1 + \frac{\alpha_1}{(1 - \delta)^2 \alpha_3}} - 1 \right]$.

Assumption (A2) can be interpreted as requiring that there be sufficient sluggishness (captured by the parameter α_2) in the system. Note that if $\alpha_2 > \alpha_1$ then this condition necessarily holds. We may then obtain the following proposition.

Proposition 4. *Under (A1)-(A2) hold, $\lambda_{11}, \lambda_{12}$ cannot be real and negative.*

In order to ensure that any cyclicity is smooth, Proposition 4 states that we simply need to have enough sluggishness. This should not be too surprising: sluggishness increases agents' desire to reduce fluctuations in investment, making outcomes characterized by large fluctuations—as would be the case under non-smooth cyclicity—less likely to occur.

As a final result in this section, we wish to establish a condition under which the situation highlighted in Proposition 3(ii) (i.e., a Hopf bifurcation and the emergence of strong smooth cyclicity) occurs. Consider the following additional assumption.

Assumption (A3). $(1 - \delta)^2 \frac{\alpha_2^2}{\alpha_3} - \alpha_3 > \alpha_1 - [1 - (1 - \delta)^2] \alpha_2$.

Note that assumption (A3) can be interpreted as requiring that α_3 (the coefficient on future investment I_{t+1}) not be too large. We then have the following proposition.

Proposition 5. *Under (A1)-(A3), the Hopf bifurcation noted in Proposition 3(ii) will occur.*

Thus, Proposition 5 makes clear that, as long as there is sufficient sluggishness, and as long as the effect of future investment on current investment is not too strong, then as ρ increases toward 1, at some point the system will undergo a Hopf bifurcation. At that point, a limit cycle will appear, and thus the solution will be characterized by strong smooth cyclicity. To understand the economics for why an increase in ρ will cause the system to become unstable, consider the limiting case where $\alpha_3 \rightarrow 0$ (so that the system is fully backward-looking, and (A3) trivially holds). A high value of ρ in this case will give an individual agent a large incentive to accumulate more capital at times when other agents increase their own accumulation. This leads to a feedback effect whereby any initial individual desire to have high current investment—due to some combination of a low current capital stock and a high level of investment in the previous period—becomes amplified in equilibrium through

a multiplier-type mechanism. When this feedback effect is strong enough, it will cause small initial deviations from the steady state to grow over time, pushing the system away from the fixed point. As a result, the economy will tend to go through repeated episodes of periods of high accumulation followed by periods of low accumulation (i.e., strong cyclical), even in the absence of any exogenous shocks. Assumption (A2), meanwhile, ensures that these periods of high and low accumulation last (on average) more than one period each, so that this strong cyclical is in its smooth form. As Proposition 5 confirms, the above intuition will continue to hold as long as $\alpha_3 > 0$ satisfies (A3); that is, even if agents are forward-looking, limit cycles will appear.³⁰

For our purposes, the emergence of limit cycles is primarily of interest if they are attractive; that is, if starting away from the limit cycle the forces in the model push outcomes toward the limit cycle. Accordingly, in Appendix E.2 we discuss conditions for the limit cycle in this model to be attractive. In that appendix we also discuss issues related to the use of saddle-path stability in this context.

2.2 Stochastic Limit Cycles

As we have seen, in dynamic environments with accumulation, strategic complementarities between agents' actions can readily create limit cycles. We showed that this can arise even when individual-level behavior favors stability, in that the system would converge to the steady state in the absence of agent interactions. Moreover, in our environment the complementarities are modest, in that they imply elasticities less than one. However, if the behavior of all agents is deterministic, then the resulting cyclical dynamics are far too regular to match the patterns observed in macroeconomic data. To see this, suppose we take the limiting case of (3)-(4) as $\alpha_3 \rightarrow 0$ (i.e., the purely backward-looking model), and suppose F is given simply the simple cubic function $F(I) = \rho I - \xi I^3$, with $\xi > 0$. In Figure 6 we plot an example of the path of I_t along a limit cycle in this set-up.³¹ While the pattern is not a perfect sine wave, it is nonetheless very regular.

[Figure 6 about here.]

For limit cycles to potentially have a chance of helping to explain macroeconomic fluctuations, it is necessary to embed them within the standard paradigm that includes exogenous

³⁰As in any rational expectations model with saddle-path stability, we assume that agents coordinate on the saddle-stable path.

³¹See notes to Figure 6 for further details.

stochastic forces.³² For example, suppose we modify our agents’ decision rule (4) to

$$I_{it} = \alpha_0 - \alpha_1 X_{it} + \alpha_2 I_{it-1} + \alpha_3 \mathbb{E}_t [I_{it+1}] + F(I_t) + \mu_t \quad (6)$$

where μ_t is a stationary exogenous stochastic force.³³ How does the addition of this exogenous stochastic force affect equilibrium dynamics? In this case, when the steady state is locally unstable, it will be said to feature a *stochastic* limit cycle. It is important to understand that the stochastic term μ_t does not simply add noise around an otherwise-deterministic cycle, as for example would be the case in panel (a) of Figure 7, which is the result of adding the AR(1) exogenous process $\mu_t = \gamma\mu_{t-1} + \sigma\epsilon_t$ to the value of I_t along the deterministic cycle from Figure 6 (where ϵ_t is i.i.d. $N(0, 1)$, and we set $\gamma = 0.9$ and $\sigma = 0.15$).³⁴ In particular, the exogenous forces not only produce random “amplitude” shifts that temporarily perturb the system from the limit cycle, but also create random “phase” shifts that accelerate or delay the cycle itself. Furthermore, even though μ_t is stationary, and while the amplitude displacement caused by a shock is temporary, the phase displacement will have a permanent component. Thus, the system will eventually converge back to the limit cycle, but permanently either ahead or behind where it would have been in the absence of the shock.

[Figure 7 about here.]

To illustrate this effect, consider the stochastic model given by (3) and (6) with the same parameter values used in Figure 6 and panel (a) of Figure 7. Figure 8 shows the effect, beginning from a point on the deterministic limit cycle, of a one-time one-standard-deviation temporary shock to μ .³⁵ The light gray line shows the path that would have occurred in absence of the one-time shock, while the dark gray shows the perturbed path. As the Figure makes clear, the perturbed path eventually returns to the limit cycle, but is permanently out of sync with (behind) the unperturbed path; that is, the temporary shock induces a random-walk in the phase of the cycle. Panel (b) of Figure 7, which presents a full simulation of this system beginning from a point on the deterministic cycle, illustrates the cumulative effect of these random phase shifts: after 10 periods, the simulated path of investment (dark gray) is noticeably out of sync with the deterministic path (light gray), and after around 50 periods the phase has shifted by about half of a cycle, so that investment in the stochastic simulation is at a trough at the same time that the deterministic simulation is at its peak. Instead of observing a smooth, regular cycle as we would in the deterministic model, the stochastic

³²An alternative route, which we do not pursue here, would be to consider chaotic dynamics.

³³Note that the form of (6) reflects the fact that the stochastic driving force is common across agents, as well as our focus on symmetric equilibria.

³⁴See notes to Figure for further details.

³⁵See notes to Figure for further details.

model clearly generates data that look more like that observed for macroeconomic variables; namely, boom-and-bust cycles that have both stochastic amplitudes and durations. As we will see below, the addition of stochastic elements also changes the spectral density of a series. In a deterministic setting, the spectral density associated with a limit cycle will have extreme peaks. In contrast, including stochastic elements tends to smooth out the spectral density, making peaks less pronounced, though without generally removing them altogether.

[Figure 8 about here.]

2.2.1 Predictability in Stochastic Limit Cycle Models

A common criticism of many early models of macroeconomic fluctuations featuring limit cycles (e.g., those of ?, ?, etc.) was that they implied an unreasonably high degree of predictability.³⁶ We wish to emphasize that, once one introduces stochastic forces into the environment, this criticism is no longer so obvious. In particular, as noted above, data simulated from a stochastic limit cycle model can have properties similar to those found in actual economic data, including significantly irregular business cycles. A similar property obtains for the *predictability* of the model. In fact, the property is even more stark for predictability than it is for regularity. A deterministic cycle is both perfectly regular and perfectly predictable arbitrarily far into the future. However, while introducing a small amount of stochastic variability in the model would tend to make the cycle only slightly less regular, the degree of unpredictability (as measured, for example, by the forecast-error variance) arbitrarily far into the future will jump up discontinuously.³⁷ The reason for this discontinuity follows directly from the fact that the “phase” component of the system follows a random walk. As is well known, as long as the variance of the innovations to a random walk process is positive (even if it is arbitrarily small), the forecast-error variance of the process becomes infinite as the forecast horizon increases. In our context, this means that, as you go far enough into the future, no matter how small the variance of the shock process is, as long as it is not zero the phase of the cycle becomes completely unpredictable; that is, no matter what the current state is, many periods into the future the system is just as likely to be at the bottom of the cycle as it is at the top. In Appendix E.3 we present a quantitative example of this fact using the same model and parameters as Figures 6-8.

³⁶Another common criticism is that these dynamics would imply the existence of arbitrage forces that would tend to erase any cycles. As we discuss in the context of our structural model (see Section 3.2), this criticism does not apply to equilibrium models with rationally-optimizing agents.

³⁷Formally, let $V_t(\sigma_\mu^2) \equiv \lim_{k \rightarrow \infty} \text{Var}_t(I_{t+k})$ denote the limit, when $\text{Var}(\mu_t) = \sigma_\mu^2$, of the forecast-error variance of I , conditional on information available as of date t , as the forecast horizon k extends infinitely far into the future. Since the deterministic model is perfectly predictable we have $V_t(0) = 0$. However, it can be verified that $\lim_{\sigma_\mu^2 \rightarrow 0} V_t(\sigma_\mu^2) > 0$.

2.3 The Problem with Failing to Allow for Local Instability and Limit Cycles

What happens if the data is driven by stochastic limit cycles, but a researcher fails to recognize this possibility? Consider two possibilities. First, suppose the researcher uses a linear model. In this case, the parameter estimates she would obtain will generally indicate that the steady state of the system is locally stable, even if it is in fact locally unstable. For example, suppose the data is generated from the parameterization of (6) used in Figures 7-8, which features a stochastic limit cycle.³⁸ Suppose, however, that a researcher knows the form of (6) and of the shock process, but erroneously believes that the data is generated from a linear version of this equation (i.e., that $F(I) = \rho I$). In this case, the estimated parameters she would obtain will imply a pair of locally stable eigenvalues with modulus 0.96, even though the true process features eigenvalues (at the steady state) that are actually of modulus 1.28. Hence, while the researcher would not perceive any problem, she would in fact have obtained a significantly biased estimate of the forces governing local dynamics. Second, a similar problem arises if a researcher estimates the model by non-linear methods, but throws out all parameterization that imply local instability (as would be the case if estimation was done using a standard package like Dynare). The researcher would then infer local stability even when local instability is present. Whether this bias is large or small cannot be determined *a priori*. One needs to estimate the model allowing for limit cycles to know whether this problem is present and, if so, how substantial it is. We do this in the following section for our fully specified model.

3 A New Keynesian Model with Risky Household Borrowing and Default

In Section 1 we presented evidence of a pervasive spectral property—namely, a peak in the spectrum near 38 quarters—that standard business cycle theories do not appear to explain well. In Section 2, we showed that models featuring capital accumulation and complementarities may provide a good framework for explaining such patterns, while we also emphasized that there are important challenges in solving and estimating such models due to the possibility of local instability and (stochastic) limit cycles. In this section, we link the previous two sections by exploring a fully specified general equilibrium model with financial-real interactions that leads to equilibrium behavior similar to that captured by the reduced-form

³⁸To keep things simple, we suppose here the shock process is i.i.d. (i.e., $\gamma = 0$), but adjust σ to maintain the same unconditional variance as before (i.e., $\sigma = 0.15/\sqrt{1 - 0.9^2}$).

model of Section 2. In particular, our aim is to construct and estimate such a model to see to what extent it can capture the peaks present in the spectra of hours and of the risk premium on borrowing. We first estimate our model with standard linear methods and then, in a second stage, we estimate a non-linear version allowing for the possibility of local instability and limit cycles. The first exercise will allow us to assess whether the complementarities induced by financial frictions—which are at the core of our model—can help explain the data, while the latter exercise will allow us to check to what extent using the simpler linear methods may bias the characterization of internal dynamics.

3.1 The Model

We propose here a model of real-financial linkages that produces strategic complementarities in agents' purchasing decisions. Relative to many existing models of real-financial linkages, our model focuses primarily on the household and emphasizes how labor market risk (in the form of unemployment risk) can affect financial conditions through default risk, and how financial conditions in turn can affect consumer demand, thereby feeding back to labor market risk. The model builds on the workhorse three-equation New Keynesian model, with the most important difference being our specification of the household problem and of the functioning of the financial sector. For this reason, we begin by presenting the household setup and the determination of lending rates in the banking sector. In Appendix F, we present a full description of the model.

3.1.1 The Determination of Household Consumption Decisions and Risk Premium on Loans

Consider an environment with a continuum of identical households of mass one, each composed of a continuum of identical members (workers) of mass one and a household head. Each household, through its 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

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

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 (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. A household's holding of durable goods is denoted by X_t ,³⁹ 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

$$X_t = (1 - \delta)X_{t-1} + I_t \quad (7)$$

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 an important 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 payments. 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 is 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 $\Phi = 0$ bringing us back to a frictionless credit market. As we show in Appendix F, this financial market imperfection yields a budget constraint for household h of the form

$$D_{ht+1} = [\bar{e}_t + (1 - \bar{e}_t)\phi](1 + r_t)(D_{ht} + P_t C_{ht} + P_t^X I_{ht}^X) - (1 + i_t)Y_{ht}, \quad (8)$$

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 $[\bar{e}_t + (1 - \bar{e}_t)\phi](1 + r_t)$,

³⁹Since all households will be identical, for notational simplicity we will often drop h subscripts where no confusion will arise.

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 consumption services given by

$$U'(C_t - \gamma C_{t-1}) = \beta \frac{\xi_t}{\xi_{t-1}} [\bar{e}_t + (1 - \bar{e}_t) \phi] (1 + r_t) \mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{1 + \pi_{t+1}} \right], \quad (9)$$

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 $\phi < 1$, (9) reflects the fact that the household knows it will default on some fraction of its loans. The household will also have an Euler equation associated with the purchase of durables given by⁴⁰

$$U'(C_t - \gamma C_{t-1}) = \beta \frac{\xi_t}{\xi_{t-1}} \times \mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{(1 + \pi_{t+1}) P_t^X} \left\{ \frac{R_{t+1}^X (1 + i_{t+1})}{[\bar{e}_{t+1} + (1 - \bar{e}_{t+1}) \phi] (1 + r_{t+1})} + (1 - \delta) P_{t+1}^X \right\} \right]. \quad (10)$$

Equation (10), when combined with (9), can be interpreted as an arbitrage condition that the return to holding a durable good must satisfy.⁴¹

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 Appendix F, in a zero profit equilibrium we will have

$$1 + r_t = (1 + i_t) \frac{1 + (1 - \bar{e}_t) \phi \Phi}{\bar{e}_t + (1 - \bar{e}_t) \phi}. \quad (11)$$

We refer to $r_t^p \equiv \frac{1+r_t}{1+i_t} - 1$ as the risk premium. Using (11) to replace r_t in (9) we get

$$U'(C_t - \gamma C_{t-1}) = \beta \frac{\xi_t}{\xi_{t-1}} [1 + (1 - \bar{e}_t) \phi \Phi] (1 + i_t) \mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{1 + \pi_{t+1}} \right]. \quad (12)$$

From (9) and (11) we see how unemployment risk, financial conditions, and purchasing decisions all become interrelated due to the fact that loans to households occasionally involve default. Equation (11) indicates that as unemployment increases so does the risk premium on

⁴⁰Note that the household treats the purchase of durable goods as it would any other asset.

⁴¹In the case where $\phi = 1$ and $i_t = r_t$, equation (10) reduces to the standard asset-pricing condition

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

loans, while equation (9) indicates that a higher risk premium on loans will lead households to delay their purchases. Equation (12) 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 2.

While the household cannot completely control the employment rate of its members, it can dictate their reservation wage, and in particular household members will optimally accept all jobs for which the real wage satisfies

$$\frac{W_t}{P_t} \geq \frac{v'(1 - e_t)}{U'(C_t - \gamma C_{t-1})} \frac{[e_t + (1 - e_t)\phi](1 + r_t)}{(1 + i_t)} + \frac{1 + r_t}{1 + i_t} (1 - \phi) \left(C_t + \frac{P_t^X}{P_t} I_t^X \right). \quad (13)$$

In equilibrium, firms will offer wages that satisfy (13) with equality. Note that if $\phi = 1$ and $r_t = i_t$, then (13) implies accepting all wage offers where the real wage is higher than the marginal rate of substitution between leisure and consumption.⁴²

3.1.2 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 \tilde{C}_{kt} , from the unit mass of intermediate service firms, and combining them in a standard way according to a Dixit-Stiglitz aggregator.⁴³ 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 $F(\tilde{e}_{kt}, \theta_t)$, where \tilde{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 $\tilde{C}_{kt} = s[\tilde{X}_{kt} + F(\tilde{e}_{kt}, \theta_t)]$, $s > 0$, where \tilde{X}_{kt} is the amount of durable goods rented by firm k from households. Moreover, after using newly produced durables to produce consumption services, the firm can sell the remaining

⁴²Note that the extra term on the right-hand side of (13) reflects the fact that, by accepting a job, the household loses the possibility of being allowed to default on that worker's loan.

⁴³See Appendix F for further details.

undepreciated amount to households at the market price P_t^X .⁴⁴ 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 allow 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 to obtain 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 undepreciated portion sold to households. For rented durables, this marginal cost is simply the rental rate. Thus, we must have

$$R_t^X = \frac{W_t}{F_e(e_{kt}, \theta_t)} - \psi P_t^X \quad (14)$$

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

3.1.3 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 the form

$$1 + i_t \approx \Theta \mathbb{E}_t \left[e_{t+1}^{\phi_e} (1 + \pi_{t+1}) \right] \quad (15)$$

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

⁴⁴Note that in the market for durable goods the intermediate firms are price takers, while they are price setters in the consumption service market.

of this monetary policy rule is that it gives the equilibrium equations the bloc recursive structure.⁴⁵

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} .⁴⁶ Using the monetary policy rule (15), as shown in Appendix F 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 equations⁴⁷

$$1 + r_t^p = \frac{1 + (1 - e_t) \phi \Phi}{e_t + (1 - e_t) \phi} \quad (16)$$

$$X_{t+1} = (1 - \delta)X_t + \psi F(e_t, \theta_t), \quad (17)$$

$$U'(s(X_t + F(e_t, \theta_t)) - \gamma s(X_{t-1} + F(e_{t-1}, \theta_{t-1}))) = \beta \Theta \frac{\xi_t}{\xi_{t-1}} [e_t + (1 - e_t) \phi] (1 + r_t^p) \\ \times \mathbb{E}_t \left[U'(s(X_{t+1} + F(e_{t+1}, \theta_{t+1})) - \gamma s(X_t + F(e_t, \theta_t))) e_{t+1}^{\phi_e} \right]. \quad (18)$$

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 1.

3.1.4 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

⁴⁵The precise form of the Taylor rule we use to obtain the block recursive property is

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

This deviates slightly from (15) due to Jensen's inequality, which is why (15) is expressed with an \approx symbol.

⁴⁶The relevant equilibrium equations correspond to (9), (10), (11), (13) with equality, (14), the aggregate consumption equation $C_t = s[X_t + F(e_t, \theta_t)]$, the accumulation equation (7), the Phillips curve, and the Taylor rule.

⁴⁷Given 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.

environment with a potentially rich endogenous propagation mechanism, can capture the spectral properties of the data.

3.2 Limit Cycles and Arbitrage

As will become clear when we present our estimation results below, for certain parameterizations our model will feature limit cycles. In addition to the criticism noted in Section 2.2.1, another 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., ? and ?), it does not apply to our model. We state this formally in the following proposition.

Proposition 6. *There are no arbitrage opportunities in the equilibrium of our model, even when there are no shocks and the dynamics feature a limit cycle.*

Proposition 6 follows directly and obviously from the fact that agents in our model are rationally optimizing and forward-looking. For example, even though there may be predictable cycles 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.

3.3 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 for the purposes of estimation we normalize it to 1.⁴⁸ 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 (17)-(18) with respect to $\log(X)$, $\log(e)$, r^p , and $\mu_t \equiv -\Delta \log(\xi_t)$, we obtain equations of the form⁴⁹

$$\hat{X}_{t+1}^* = (1 - \delta) \hat{X}_t^* + \hat{e}_t, \quad (19)$$

$$\hat{e}_t = -\alpha_1 \hat{X}_t^* + \alpha_2 \hat{e}_{t-1} + \alpha_3 \hat{e}_{t+1} - \alpha_4 \hat{r}_t^p + \alpha_4 \mu_t, \quad (20)$$

$$\hat{r}_t^p = R^p(\hat{e}_t) \equiv R_1^p \hat{e}_t, \quad (21)$$

⁴⁸Allowing for deterministic growth in the model does not change any results.

⁴⁹See Section I in the Appendix for details.

where \hat{X} and \hat{e} are log-deviations from steady state, \hat{r}^p is the deviation in levels, $\hat{X}^* \equiv \hat{X}/(\alpha\delta)$, and the α_j 's are (positive) functions of the structural parameters. Note that, since the risk premium is a negative function of employment ($R_1^p < 0$), the above system has a structure identical to the model of Section 2. 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 (19) and (20). 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 (16) in its (log-)linearized form as given by (21). 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 $\psi = 0$). 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.⁵⁰ As discussed in Section 2, allowing for non-linearity in the strength of the complementarity will allow us to expand the parameter space to include situations where there is a (unique) rational-expectations solution featuring local instability and limit cycles. In particular, for this case we allow that

$$\hat{r}_t^p = R^p(\hat{e}_t) \equiv R_1^p \hat{e}_t + R_2^p \hat{e}_t^2 + R_3^p \hat{e}_t^3, \quad (22)$$

where the coefficient on the linear term, R_1^p , is the same as in the linear RP model, and we estimate the second- and third-order coefficients R_2^p and R_3^p directly.⁵¹ We take $R^p(\cdot)$ to be a cubic polynomial since, as noted in Section 2.1.2, for the model to be capable of producing (attractive) limit cycles we will typically need the third derivative of $R^p(\cdot)$ to be sufficiently positive, and a cubic polynomial is a simple way to allow for this possibility.

We estimate this model using the indirect inference method of ?, where for each parameterization the model is solved by a first-order (linear RP, canonical, and no friction models)

⁵⁰We could also allow for non linearities in both (19) and (20). However, we chose to allow for non-linearities in the risk premium as to make the analysis more transparent since it allows us to refer to results from Section 2.

⁵¹Strictly speaking, R_2^p and R_3^p are functions of the underlying structural parameters. In choosing instead to estimate R_2^p and R_3^p directly, we are allowing in the simplest possible way for sources of non-linearity beyond those explicitly modeled. It would nonetheless be straightforward to enrich the microfoundations of the model to allow for non-linearities of a form at least as general as that embodied by $R^p(\cdot)$.

or third-order (non-linear RP model) perturbation method.⁵² In the non-linear RP model, the solution and estimation is somewhat involved, as it allows for the possibility of locally unstable steady state and limit cycles in a stochastic model with forward-looking agents. To our knowledge, such an exercise is novel. 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 1).⁵³ 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).⁵⁴ 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. 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 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: ω , γ , ψ (all except canonical), ϕ_e , ϕ (linear and non-linear RP only), Φ (linear and non-linear RP only), R_2^p and R_3^p (non-linear RP only), ρ , and σ . The weighting matrix used in estimation is the identity matrix.

[Figure 9 about here.]

Figures 9-11 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 9 and 10, which show

⁵²Details of the solution and estimation are given in Section I of the Appendix.

⁵³Note that, in our model, since employed workers each work the same number of hours, hours is simply proportional to employment.

⁵⁴We 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).

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 9, 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 10 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 9, 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 importance in allowing our model to capture the salient business cycle features of the data. As shown in panel (d) of Figure 9, the canonical model, which is obtained by re-estimating the model with both the complementarity and accumulation channels shut down, tells a similar story, though even more extreme: not only does the model miss the peak near 40 quarters, it no longer exhibits any hump whatsoever. Lastly, consider panel (a) of Figures 9 and 10, which show the results for the non-linear RP model. Relative to the linear RP model, the fits of both the spectra are noticeably better, and in particular while both models generate 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 10 about here.]

[Figure 11 about here.]

The parameter estimates for the four models are presented in Table 1. 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-cyclicity 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 both increasing the degree of household interest rate counter-cyclicity and increasing the household’s consumption response to that counter-cyclicity.

[Table 1 about here.]

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 1) to the no friction model, the linear RP model, and finally to the non-linear RP model, the unconditional standard deviation of the shock process monotonically decreases. Further, moving leftward from the no friction model, the persistence and innovation standard deviation also monotonically decrease. Thus, sequentially allowing for accumulation, complementarity, and then non-linearity in the complementarity lets the model not only better fit the data (as discussed above), but do so with less reliance on exogenous stochastic forces. The upshot is that in the non-linear RP model the autocorrelation of the shock is effectively zero, so that almost all of the model’s dynamics are due to endogenous forces. These results 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 2, 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 and given respectively by 0.76 and 0.83. 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 in which deviations from steady state have endogenous half-lives of less than four quarters. In contrast, the linear RP model has a pair of complex eigenvalues, which allows it to generate spectral peaks, and these eigenvalues 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).⁵⁵ That is, when given the option, the data appear 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)—one may be significantly biasing the results of any estimation.

[Table 2 about here.]

To illustrate the deterministic mechanisms implied by the parameter estimates in the non-linear RP model, Figure 12 reports results obtained when feeding in a constant value of $\mu_t = 0$ for the exogenous process.⁵⁶ Panel (a) of Figure 12 plots a simulated 270-quarter sample⁵⁷ 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 2, which indicates that the steady state is unstable and features two complex eigenvalues. The difficulty that some earlier models had in generating cycles of quantitatively reasonable lengths may have been one of the factors leading to limited interest in using a limit cycle framework to understand business cycles. However, as this exercise demonstrates, reasonable-length endogenous cycles can be generated in our framework relatively easily, precisely because the model possesses the three key features we highlighted in the previous section: diminishing returns to capital accumulation, sluggishness, and complementarities. Second, notwithstanding the reasonable cycle length, it is clear when comparing the simulated data in panel (a) of Figure 12 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 12 plots the spectral density of hours for the deterministic model (gray line), along with the spectral density for the

⁵⁵Note 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 2) 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.

⁵⁶Note that, as it is always done when computing transitional dynamics, the model was solved (and simulated) using the estimated parameters presented in column (a) of Table 1, and in particular we did not first re-solve the model with $\sigma = 0$. Thus, agents in this deterministic simulation implicitly behave as though they live in the stochastic world. As a result, any differences between the deterministic and stochastic results are due exclusively to differences in the realized sequence of shocks, rather than differences in, say, agents' beliefs about the underlying data-generating process.

⁵⁷This is equal to the length of the sample period of the data.

data (black line) for comparison. This spectral density exhibits an extremely large peak—characteristic of a highly regular cycle—at the 38-quarter periodicity,⁵⁸ while the spectral density of the data is much flatter.

[Figure 12 about here.]

[Figure 13 about here.]

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 13 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 9), which matches the data quite well. In particular, the spectral density of the stochastic model includes a distinct peak close to 40 quarters, suggesting some degree of regularity at that periodicity, but without the exaggerated peak 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 12), 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).⁵⁹ 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 $\Phi = R_2^p = R_3^p = 0$), 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

⁵⁸The 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.

⁵⁹Note that, since we have simply fed a constant sequence $\mu_t = 0$ of shocks into our model without first resolving it under the assumption that $\sigma = 0$, this phenomenon is not due in any way to rational-expectations effects. The fact that a fall in shock volatility can lead to a rise in the volatility of endogenous variables in a limit cycle model was pointed out in ?. 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 ? for a more detailed discussion of these mechanisms in the context of an estimated reduced-form univariate equation.

amplify them, the small i.i.d. disturbances are sufficient to generate only a tiny amount of volatility in hours.

As we have already emphasized, non-linearities are crucial in order to have a steady state that is locally unstable without also having explosive dynamics, a combination which is a precondition for limit cycles to emerge. To get a sense of the estimated degree of non-linearity in our non-linear RP model, we plot in Figure 14 the expression $4[\phi_e \hat{e} + R^p(\hat{e})]$, which gives, as a function of employment \hat{e} , the approximate annualized effective real interest rate faced by households. The non-linearities coming from the non-linear response of the risk premium to economic activity are scarcely apparent in the figure. Near the steady state, a one-percent rise in hours is associated with around a 48.1-basis-point fall in the interest rate faced by the household. As we move away from the steady state, this sensitivity fades, but only mildly. For example, at the deepest trough and highest peak recorded in our data sample, a one-percent rise in hours would be associated with 46.4- and 45.6-basis-point falls in the household interest rate (3.6% and 5.2% less sensitivity than at the steady state), respectively. Such differences in the strength of the complementarities may appear small at first pass, but they are sufficient to have large quantitative impacts on the dynamics of the system. This illustrates nicely that even if non-linearities may be very minor, it may be important to allow for them in order to get an proper estimation of the internal propagation mechanism.

[Figure 14 about here.]

3.4 Sensitivity of Results to Target Frequency Range

As we noted in Section 1, 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 particular choice affects our results. In particular, since the lower end of the standard business cycle range in the literature is 6 quarters, we consider what happens if we estimate the parameters of the model restricting attention to periodicities from 6 to 50 quarters. We also consider what happens if we restrict attention further to the range beyond the traditional upper bound of 32 quarters (i.e., restricting to periodicities between 32 and 50 quarters), and we repeat these same 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 of a similar nature for the linear RP model (available upon request). As we will see, the above choices make little difference to our results. As an additional exercise, we also report results when estimating over the 2-100 quarter range, and show that in this case the results are fundamentally

changed.⁶⁰

The top row of plots in Figure 15 shows the the hours and risk premium spectral densities for period ranges of the form $(x,50)$, $x = 2, 6, 32$,⁶¹ while the second row shows results for ranges of the form $(x,60)$. The corresponding data spectral 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 well the spectral peak near 40 quarters. Further, as shown in the first six rows of Table 3, the eigenvalues associated with the solved system remain complex (indicating cyclicity) and outside the unit circle (indicating the presence of a limit cycle) for all six of these different period ranges.

[Figure 15 about here.]

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 15 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 3 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 estimate of the exogenous driving now indicate an auto-correlation above .8 with a standard error of .17, 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 3 about here.]

⁶⁰We have also explored the effect of estimating the model only on frequencies between 2-32 quarters. This tends to favor inferring—as found in much of the literature—that persistent exogenous shocks drive business cycles. For example, if we estimate the linear version of the model to match both hours worked and the risk premium (as before), the resulting shock process has an autoregressive parameter of 0.98 (with an innovation standard deviation of 0.0018). In comparison, when estimating this same model over frequencies between 2-50 quarters we have an autoregressive parameter of only 0.14.

⁶¹Note that the (2,50) results simply reproduce the information in Figures 9 and 10.

3.5 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 structural 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 do 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 12), which features booms that last 23 quarters (trough to peak) and recessions 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 16 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.⁶² 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 fragile to negative shocks (e.g., a financial disruption).

[Figure 16 about here.]

A third important way in which the non-linear and linear RP models differ is in their implications for the relationship between the variance of the shock process and the resulting variance of the endogenous variables. As we already noted above, the standard deviation

⁶²See notes to Figure for further details.

(s.d.) of hours in the non-linear RP model is in fact smaller than the s.d. of hours obtained when we feed in a constant stream of zeros for the shocks but do not re-solve the model under the assumption that the s.d. of the shocks is zero. This property in fact holds more generally. To illustrate this, in Figure 17 we plot the elasticity of the s.d. of hours with respect to σ (the s.d. of the shock innovation) for each of the linear and non-linear RP models over a range of σ 's. The horizontal axis in the Figure is $\sigma/\hat{\sigma}$, where $\hat{\sigma}$ is the estimated value of the innovation s.d. for the corresponding model (so that $\sigma/\hat{\sigma} = 1$ corresponds to the elasticity at the estimated value of σ). Note that, for the non-linear RP model, we re-solve the model for each different value of the s.d. of the shock innovation, so that the s.d. of hours is obtained from the appropriate rational expectations solution.⁶³ As the Figure shows, the elasticity for the linear RP model (light gray) is constant and equal to one, indicating that a one percent increase in σ always leads to a one percent increase in the s.d. of hours. The same is not true for the non-linear RP model (dark gray), in which the elasticity is negative over the range of standard deviations shown (and actually declining over most of that range). For example, near the estimated value of σ , a one percent increase in σ is associated with a 0.45 percent fall in the s.d. of hours. To the extent that this sort of relationship—which is typical of models featuring limit cycles—is present in the real world, it has the intriguing implication that reducing the volatility of economic shocks may not help stabilize the economy.

[Figure 17 about here.]

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 on this dimension.

We have emphasized three elements that have led us to infer that business cycles are more likely generated by strong endogenous forces than by persistent exogenous shocks. However, in concluding, we would like to emphasize the one element we view as most important for

⁶³This is unnecessary for the linear RP model, since the rational expectations solution is invariant to the s.d. of the shock in that case.

this debate, which is 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 to 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,⁶⁴ 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 a AR(1) process. Given that cyclically sensitive variables such as unemployment, capacity utilization, 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 point, then inferring that the economy has a strong internal propagation mechanism becomes much more likely. While we have made the explicit case for this inference for only one model, we conjecture that if one uses a different model and successfully explains these observed humps in the spectral density, then most likely the estimated model will require complementarities that generate a strong internal cyclical mechanism. Whether the resulting model delivers the more extreme form of endogenous propagation as implied by limit cycles, or if instead it favors more 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 one's inference regarding the role of shocks in driving business cycles is likely to be greatly diminished if one tries to explain features of the data we emphasized such as the hump shaped spectral density for hours worked and risk premia.

⁶⁴A 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 3.

Appendix to “Putting the Cycle Back into Business Cycle Analysis”

Paul Beaudry, Dana Galizia and Franck Portier

January 2018

A Data

- U.S. Population: Total Population: All Ages including Armed Forces Overseas, obtained from the FRED database (POP) from 1952Q1 to 2015Q2. Quarters from 1947Q1 to 1952Q1 are obtained from linear interpolation of the annual series of National Population obtained from U.S. Census, where the levels have been adjusted so that the two series match in 1952Q1.
- U.S. Total GDP is obtained from the Bureau of Economic Analysis National Income and Product Accounts. Real quantities are computed as nominal quantities (Table 1.1.5) over prices (Table 1.1.4.). Sample is 1947Q1-2015Q2.
- Non-Farm Business Hours, Total Hours and unemployment rate (16 years and over) are obtained from the Bureau of Labor Statistics. Sample is 1947Q1-2015Q2 (1948Q1-2015Q2 for total hours), and we do not use the observations of 2015.
- Capacity utilization: Manufacturing (SIC), Percent of Capacity, Quarterly, Seasonally Adjusted, obtained from the FRED database, (CUMFNS). Sample is 1948Q1-2015Q2.
- TFP: utilization-Adjusted quarterly-TFP series for the U.S. Business Sector, produced by John Fernald, series ID: dtfp_util, 1947Q1-2015Q2.
- Spread: Moody’s Seasoned Baa Corporate Bond Minus Federal Funds Rate, quarterly average, obtained from the FRED database, (BAAFFM). Sample is 1954Q3-2015Q2.
- Delinquency rate: Delinquency Rate on All Loans, All Commercial Banks, Percent, Quarterly, Seasonally Adjusted, obtained from the FRED database, (DRALACBS). Sample is 1985Q1-2015Q2.

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

†Department of Economics, Carleton University

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

§The authors thank Jess Benhabib, Kiminori Matsuyama and Morten Ravn for helpful discussions. The authors would also like to thank seminar participants at Banque de France, University of Manchester, Pompeu Fabra-Toulouse “skiminar”, NBER Summer Institute, Hydra workshop, UCL, UCLA, UCSD, Wisconsin, the University of Pennsylvania, Oxford University, EUI Florence, EIEF Roma and Northwestern University for comments. Franck Portier acknowledges financial support by the ADEMU project, “A Dynamic Economic and Monetary Union,” funded by the European Union’s Horizon 2020 Program under grant agreement No 649396.

- National Financial Conditions Index: Chicago Fed National Financial Conditions Index, Index, Quarterly, Not Seasonally Adjusted, obtained from the FRED database, (NFCI). Sample is 1973Q1-2015Q2.
- National Financial Conditions Risk Subindex: Chicago Fed National Financial Conditions Risk Subindex, Index, Quarterly, Not Seasonally Adjusted, obtained from the FRED database, (NFCIRISK). Sample is 1973Q1-2015Q2.
- Policy rate : computed from the FRED database as Moody’s Seasoned Baa Corporate Bond Yield (BAAFFM) minus the spread (BAAFFM). Sample is 1954Q3-2015Q2. The real policy rate is obtained by subtracting realized inflation, computed using the “Consumer Price Index for All Urban Consumers: All Items” (CPIAUCSL from the FRED database).

B Spectral Properties of Detrended GDP

Our observations of a distinct peak in the spectral density of a set of macroeconomic variables may appear somewhat at odds with conventional wisdom. In particular, it is well known, at least since ?, that several macroeconomic variables do not exhibit such peaks, and for this reason the business cycle is often defined in terms of co-movement between variables instead of reflecting somewhat regular cyclical behavior. According to us, this perspective on business cycle dynamics may be biased by the fact that it often relies on examining the spectral properties of transformed non-stationary variables, such as detrended GDP. We instead have focused on variables—which we like to call cyclically sensitive variables—where business cycle fluctuations are large in relation to slower “trend” movements. For such variables, the breakdown between low-frequency trend and cycle is potentially less problematic if the series can still be considered stationary. In contrast, if one focuses on quantity variables, for example GDP, one needs to believe that the detrending procedure used to make it stationary is allowing one to isolate the relevant cyclical properties. This is certainly questionable as the detrending procedure most often changes the spectral properties dramatically. To see this, in Figure 18 we report the same information we reported before regarding the spectral density, but in this case the series is real per capita GDP. Here we see that the spectra of the non-detrended data and of the filtered data have very little in common with each other. Since it does not make much sense to report the spectral density of non-detrended GDP (it is clearly non-stationary), in Panel (a) of Figure 19 we focus on the spectral density of GDP after removing low-frequency movements using the various high-pass filters. These spectral densities are in line with conventional wisdom: even when we have removed very low-frequency movements, we do not detect any substantial peak in the spectral density of GDP around 40 quarters. How can this be? What explains the different spectral properties of filter-output versus the level of hours worked? There are at least two possibilities. First, it could be that the filtering we implemented on GDP is simply not isolating cyclical properties. Alternatively, if one believes that such a filtering procedure is isolating cyclical properties, the answer to the puzzle may lie in the behavior of (detrended) TFP. Panel (b) of Figure 19 plots the spectral density of the (log of the) product of TFP by hours, after having removed low-frequency movements in the same way we have done for GDP and other variables. Note that the spectral density of this variable is very similar to the GDP one, as if, for periodicity below 80 quarters, GDP could be approximatively seen as being produced linearly with hours only, whose productivity would be scaled by TFP. If Hours and TFP were uncorrelated, then the spectral density of GDP (in logs) would be the sum of the ones of (log) Hours and TFP. This is approximatively what we have, as shown in panels (c) and (d). The spectral density of TFP shows a quick pick-up it just above periodicities of 40 quarters. As with GDP, we do not see any marked

peaks in the spectral density of TFP. An interesting aspect to note is that if we add the spectral density of hours worked to that of TFP, we get almost exactly that of GDP. This suggests that looking at the spectral density of GDP may be a much less informative way to understand business cycle phenomena than looking at the behavior of cyclically sensitive variables such as hours worked and capacity utilization. Instead, GDP may be capturing two distinct processes: a business cycle process associated with factor usage and a lower-frequency process associated with movement in TFP. For this reason, we believe that business cycle analysis may gain by focusing more closely on explaining the behavior of cyclically sensitive variables at business cycle frequencies.

[Figure 18 about here.]

[Figure 19 about here.]

C Non-linearity

Here we also examined whether the business cycle fluctuations which we have been focusing upon are approximately normally distributed. To this end, we perform χ^2 and J omnibus tests, which combine skewness and kurtosis into a single statistic, on our main series after using a high-pass linear filter to remove periodicities greater than 60 quarters, which allows us to retain all the variation that we have argued is relevant for the business cycle, while removing more medium- and long-run fluctuations. The null hypothesis for the test is normality. For non-farm business hours, the unemployment rate, and capacity utilization, the p -values we find are, respectively, 5%, 1% and smaller than 1% for the D’Agostino-Pearson test and 7%, 1% and close to 0% for the Jarque-Bera test. This indicates that linear-Gaussian models might not be an appropriate way to describe these data, and that one may need to allow for some type of non-linearity in order to explain these movements. Accordingly, we explore a class of explanation that allows for such non-linearities. Moreover, when we estimate our model, we use the skewness and kurtosis properties of the data to help identify parameters.

D Evidence From Longer Historical Series

While it is difficult to get comprehensive measures of hours or capacity utilization for extended periods prior to WWII, financial data is readily available. For example, the dividend yields on stocks are available since 1871, and the spreads between different bonds are available since 1919.¹ Although it is still debated, movements in the dividend yield are often argued to mainly reflect movement in risk premia. Similarly the interest rate spread between similarly-dated BAA bonds and treasuries are also generally thought to be good indicators of risk premia. To this end, in panel (a) of Figure 20 we report the spectral density of the quarterly dividend yield over the period 1871-2016, and in panel (b) we report the estimated spectral density for the spread between 10-year BAA bonds and 10-year treasuries from 1919 to 2016. In both panels we again observe a distinct peak in the spectral densities at around 9 years, suggesting recurrent cycles at this periodicity. Thus, at least for financial data, our earlier evidence appears robust to extending the sample further back in time.

As a last piece of evidence of cyclical behavior, we present in Figure 21 extended results regarding the probabilities of being in recession in period $t + k$ conditional of being in recession at

¹The dividend yield and the 10 year constant maturity treasury yield are drawn from Robert Shiller’s website, while the BAA is from FRED..

period t , using recession dates going back to 1871. It should be noted that NBER recession dates suggest that recessions were more frequent prior to WWII. While the patterns observed in Figure 21 are less distinct than those observed in Figure 4, the same general movements remain visible; that is, this measure of the conditional probability of being in a recession appears to again have a sinusoidal shape with a period around 9 years.

[Figure 20 about here.]

[Figure 21 about here.]

E Additional Results for the Reduced-Form Model

E.1 Saddle Limit Cycles

Throughout subsection 2.1, for the sake of exposition we assumed that if the system has exactly one positive and unstable eigenvalue, and either (a) a pair of unstable complex eigenvalues, or (b) one or two negative and unstable eigenvalues, then there is a solution featuring strong cyclicity. This assumption deserves some comment here. Focusing on the case where we have a pair of complex eigenvalues, as they lose stability (i.e., as a Hopf bifurcation occurs), we argue in next subsection of this appendix that as long as F''' is sufficiently negative the system will typically retain a similar saddle-path structure to the pre-bifurcation system. Specifically, a limit cycle will appear as the eigenvalues lose stability, and this limit cycle will be attractive on a two-dimensional manifold—which we refer to as the non-explosive manifold—in our three-dimensional phase space,² and repulsive elsewhere. The transversality condition will then force I_t to jump—for any given values of X_t and I_{t-1} —onto this manifold. Figure 22 illustrates this configuration, which we refer to as a saddle limit cycle. In panel (a), which shows a three-dimensional phase diagram, the non-explosive manifold (gray surface) contains a limit cycle, and beginning from any point initially on this manifold (except the steady state), the system converges to that limit cycle. Examples of such trajectories are given by the black solid and dotted lines. The dynamics beginning from any point not on the gray surface, however, are dominated by the remaining (real and positive) unstable eigenvalue, causing paths emanating from such points to explode. This is illustrated in panel (b) of the Figure, which shows the same phase space from a slightly different angle. The black dashed lines in this panel show examples of trajectories beginning away from the non-explosive manifold, which necessarily become unbounded and therefore violate the transversality condition.³

[Figure 22 about here.]

E.2 Emergence of Attractive Limit Cycles

In our reduced-form model of Section 2, we established that a Hopf bifurcation may emerge as ρ increases towards 1. We address here the question of whether the limit cycle that appears

²Formally, if A is the eigenspace associated with the two complex eigenvalues, then the non-explosive manifold is the invariant two-dimensional subspace that is tangent to A at the steady state.

³The dynamic system we are considering is three-dimensional, which allows for the possibility of such a saddle limit cycle to appear as part of a Hopf bifurcation. In a two-dimensional system, on the other hand, a Hopf bifurcation would necessarily be associated with indeterminacy, since beginning from any point in the phase space the system would converge to the limit cycle (so that the transversality condition would be satisfied). This latter case has appeared often in the macroeconomic literature. In contrast, the type of configuration implied by Figure 22 is less common.

near the bifurcation point is attractive; that is, whether the system would converge towards such an orbit given an arbitrary starting point on the non-explosive manifold. To use language from the theory of dynamical systems, a bifurcation may be either *supercritical* or *subcritical*. In a supercritical bifurcation, the limit cycle emerges on the “unstable” side of the bifurcation and attracts nearby orbits, while in a subcritical bifurcation the limit cycle emerges on the “stable” side of the bifurcation and repels nearby orbits.⁴

The emergence of a limit cycle is mainly of interest to us if it is attractive, since if it is there will be a determinate solution to the system (and this solution will be such that the system approaches the limit cycle over time), while if it is not attractive there will generally not be a solution the system. Even in environments that are two-dimensional to begin with, the conditions governing whether a Hopf bifurcation is supercritical or subcritical are somewhat complicated. In our set-up, stating these conditions would first require obtaining an analytical solution to the model (under the assumption that a solution exists) in order to reduce the number of dimensions from three to two, and then expressing the Hopf conditions in terms of the coefficients of this solution. Not surprisingly, such conditions are highly intractable, making it difficult to gain any insight from them. However, there are two special cases of the model that do yield some important insight into what properties $F(\cdot)$ must have in order for a bifurcation to be supercritical: the limiting case of $\alpha_3 \rightarrow 0$ (the purely backward-looking case), and that of $\alpha_2 = 0$ (the case with no sluggishness). Note that, in the latter case, since there is only one pre-determined variable (i.e., X), a Hopf bifurcation is not possible, and in particular our system will have two real eigenvalues: one greater than 1, and one less than 1. Thus, for this case we instead focus on a situation where, as ρ increases towards 1, the smaller eigenvalue decreases through -1 , which is referred to as a flip bifurcation.⁵ Flip bifurcations are also associated with the emergence of limit cycles (strong non-smooth cyclicity in our terminology), and as with Hopf bifurcations may be supercritical or subcritical. The following proposition establishes that, regardless of which of these two special cases we are in, the condition on F that ensures the emergence of an attractive limit cycle is the same (suggesting that a similar condition may ensure supercriticality for the general case).

Proposition E.1. *Suppose that either (a) $\alpha_3 = 0$ and as ρ increases towards 1 the system undergoes a Hopf bifurcation, or (b) $\alpha_2 = 0$ and as ρ increases towards 1 the system undergoes a flip bifurcation. The bifurcation will be supercritical as long as $F'''(I^s)$ is sufficiently negative.*

The requirement that $F'''(I^s)$ be sufficiently negative for the emergence of an attractive limit cycle can be related to economic forces. When $F'(I^s) = \rho$ is large enough so that the system is locally unstable, the demand complementarities are strong enough near the steady state that any perturbation from that point will tend, through the feedback mechanism described above, to induce outward “explosive” forces. When $F'''(I^s)$ is negative, however, as the system moves away from the steady state these demand complementarities will eventually fade out (i.e., $F'(I)$ eventually falls), so that the explosive forces that are in play near the steady state are gradually replaced with inward “stabilizing” ones. As long as $F'''(I^s)$ is sufficiently negative, these stabilizing forces will emerge quickly enough, and an attractive limit cycle will appear at the boundary between the inner explosiveness region and the outer stability region.

If instead we had $F'''(I^s) > 0$, then instead of dying out, the demand complementarities would tend to *grow* in strength as the system moves away from the steady state, so that inward stabilizing

⁴Note that the results in this subsection on the attractiveness of limit cycles in our model are not particularly novel, largely echoing existing results in the non-linear dynamics literature. We nevertheless include them here for completeness.

⁵It can be verified that such a bifurcation will occur if and only if $\alpha_3 < \alpha_1/(2 - \delta)$.

forces do not appear. In this case, when ρ is large enough for the system to become unstable, the bifurcation will be subcritical, in which case a repulsive limit cycle appears just before the system becomes unstable.⁶

The general insight we take away from Proposition E.1 is that attractive limit cycles are likely to emerge in our setting if demand complementarities are strong and create instability near the steady state, but tend to die out as one moves away from the steady state. We may refer to such a setting as one with local demand complementarities. In an economic environment, it is quite reasonable to expect that positive demand externalities are likely to die out if activity gets very large. For example, if investment demand becomes sufficient large, some resource constraints are likely to become binding, causing strategic substitutability to emerge in place of complementarities. Similarly, physical constraints, such as a non-negativity restrictions on investment and capital or Inada conditions implying that the marginal productivity of capital tends to infinity at zero are reasonable considerations in economic environments that will limit systems from diverging to zero or to negative activity. Such forces will in general favor the emergence of attractive limit cycles in the presence of demand complementarities.

E.3 (Un)predictability: An Example

Figure 23 plots, using the same parameters as in Figures 7-8, a forecast of I_{t+k} as of date t (i.e., $\mathbb{E}_t[I_{t+k}]$, dark gray solid line), as well as a 66% conditional confidence interval for I_{t+k} (dark gray dotted lines).⁷ For comparison, we also plot the unconditional forecast $\mathbb{E}[I_{t+k}]$ (light gray solid line) and an unconditional 66% confidence interval for I_{t+k} (light gray dotted lines), as well as a “deterministic” forecast as of date t (light gray dashed line) obtained by shutting down the stochastic process completely.⁸ Several important points emerge from the Figure. First, unlike the deterministic forecast which oscillates indefinitely (reflecting the existence of the limit cycle), the conditional forecast converges to the unconditional forecast as the forecast horizon k increases. Thus, in this stochastic environment even a risk-neutral investor could not earn above-average long-run returns by betting on I . Second, as the conditional confidence interval indicates, even in the shorter run when $\mathbb{E}_t[I_{t+k}]$ differs significantly from $\mathbb{E}[I_t]$, there is quite a bit of forecast error in this example. Not only does the mean conditional forecast converge to the unconditional one as k increases, the conditional confidence interval also converges to the unconditional one. To the extent that this pattern holds for all conditional moments, this would imply that the stochastic process is in fact ergodic.

[Figure 23 about here.]

⁶Note that subcriticality of a bifurcation does not necessarily imply global explosiveness on the unstable side of the bifurcation, nor does it rule out the emergence of an attractive limit cycle in that region. Rather, the results of this subsection are inherently about the local behavior of the system, where “local” in this case means “to a third-order Taylor approximation on some sufficiently small neighborhood of the steady state”. Conclusions about the *global* behavior of the system cannot in general be inferred from these local results. In particular, subcriticality only implies that if an attractive limit cycle does emerge on the unstable side, then it must involve terms higher than third order.

⁷The initial state of the system at date t is given by $X_t = -8.00$, $I_{t-1} = -1.68$, and $\mu_t = 0$.

⁸Unconditional and conditional moments obtained using 100,000 simulations. Unconditional forecasts are computed as time averages, while conditional forecasts are computed as ensemble averages.

F Steps in Deriving the Model of Section 3

The model can be seen as an extension of the canonical three-equation New Keynesian model. The extra elements we add are habit persistence, the accumulation of durable goods (or houses), and a credit market imperfection that generates a counter-cyclical risk premium.

F.1 Model Summary

The economy is populated with a continuum of households of mass one, final good and intermediate services firms, commercial banks and a central bank. Each household (or *family*) will be composed of a continuum of mass one of worker members and an atomless family head who coordinates family activities. Time is discrete, and a period of time is divided into four sub-periods.

Final good firms Final good firms buy consumption services from the set of intermediate firms, and bundle them into a final good according to a Dixit-Stiglitz aggregator.

Intermediate firms Intermediate firms produce a set of differentiated consumption services using durable goods as input. These durable goods will be composed of existing durable goods rented from the households and new ones produced by these intermediate firms using labor. After the production of consumption services occurs, the remaining undepreciated stock of new durable goods will be sold to households. We assume that firms are price takers on the market for durable goods, but that they are monopolistic competitors when selling their variety of consumption services. As standard in a New Keynesian framework, they will only be able to adjust their prices upon the arrival of price change opportunity, which occurs according to a Poisson distribution.

Commercial Banks Orders for final goods and new durable goods are made by each individual household member before going to the labor market, and these orders are financed by credit granted by banks at interest rate r . To finance their lending, banks take deposits on which they pay the risk-free nominal rate i . Because some workers will not find a job and may as a result default on their debt, the interest rate r charged by banks to borrowers will include a risk premium r^p over and above the risk-free rate. We assume free entry and perfect competition in the banking market, so that banks make no profits.

Central Bank The central bank sets the risk free nominal interest rate i according to a type of Taylor rule.

Households A family purchases consumption services and durable goods. At the beginning of each period, after the resolution of aggregate uncertainty and the clearing of the last period's debts and deposits, each worker in the family places orders for consumption services and newly produced durable goods, financed by borrowing from banks against their uncertain labor income, with the imperfect backing of their family. Because of a financial imperfection that we describe below, the cost of borrowing will be the risky rate $r = (1 + r^p)(1 + i) - 1$. 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, and will not be able to repay their debt the next period. We assume that each employed worker works a fixed number of hours normalized to one, so that e_t is also hours.

The financial friction We assume that for workers that were employed the last period, loan contracts are fully enforceable, so that their debt is always repaid. The financial friction takes the form of a costly enforcement of debt repayments by the family for the household members that do not find a job. In particular, not all non-performing loans can be recovered by going after the family. When a household member cannot pay back a loan—which will be the case when they were unable to 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 (which the bank will always choose to do), while with probability $1 - \phi$ it is too costly to pursue the household, in which case the bank is forced to accept a default. To fix ideas, we assume that, for unemployed workers who have their debt covered by the household, the household transfers to the unemployed worker the funds necessary to cover the debt.⁹

F.2 Timing

Borrowing is in one-period bonds. A family enters the period with outstanding debt and bank deposits for each of its workers, and a stock of durable goods that is managed by the family head.

In the first sub-period, interest on last-period deposits are paid, and past debts with banks are settled. Workers that were employed in the previous period pay back their loans, and return any remaining cash balances to the household.

In the second sub-period, workers first borrow from banks, and then use the proceeds to order consumption services and durable goods. Final good services firms receive demand (paid orders) and make orders to intermediate good firms, which also receive the orders for new durable goods from the households. A worker who wants to borrow can try to do it on his own, or can get backing from his family. If there is no backing by the family, a worker who borrows but does not find a job will not be able repay loans. Backing by the family will therefore reduce the risk premium on loans, and in equilibrium will always be chosen by the workers. As we have explained above, backing by the family will be imperfect: with probability ϕ the bank will pursue the household, and with probability $1 - \phi$ it is too costly pursue the household.

In the third sub-period, intermediate good firms rent durable goods from the household head and hire workers to produce any new durable goods. As noted above, there will always a measure one of workers looking for a job, but because of job indivisibility, only a fraction e_t of workers will find a job. Within a match, the firm makes a take-it-or-leave-it wage offer to the worker, the equilibrium level of which will be at a reservation wage, which will be decided by the family head.

Finally, in the fourth sub-period, production takes place, wages are paid to workers, rental payments for durable goods are made to the household head, orders are fulfilled, and all consumption services and durable goods are brought back to the household and shared equally among family members. Firms' and banks' profits, if any, are returned to the household.

F.3 Households

There is a continuum of mass one of households (families) indexed by h who purchase consumption services from the market and accumulate durable goods. The preferences of family h are given by $U(h) = \mathbb{E}_0 \sum_t \beta^t \xi_{t-1} [U(C_{ht} - \gamma C_{t-1}) + \nu(1 - e_{ht})]$, with $U'(\cdot) > 0$, $U''(\cdot) < 0$, $\nu'(\cdot) > 0$, $\nu''(\cdot) < 0$ and $0 \leq \gamma < 1 - \delta$, where δ is the depreciation rate of the durable good. C_{ht} represents the consumption services purchased by household h in period t , $C_t \equiv \int_0^1 C_{ht} dh$ denotes the average

⁹An equivalent formulation would have the household simply absorbing that workers debt, to be repaid in a future period.

level of consumption in the economy, β is the discount factor, and ξ_t denotes an exogenous shock to the discount factor at date t . Note that this preference structure assumes the presence of external habit. Each family owns a stock of durable goods X_{ht} that is rented to firms in order to produce intermediate services. This stock of durable goods is accumulated according to the equation

$$X_{ht+1} = (1 - \delta)X_{ht} + I_{ht}. \quad (\text{F.1})$$

A family is composed of a continuum of measure one of workers indexed by j . All the members of the family share the utility $U(h)$, where $C_{ht} = \int_0^1 C_{jht} dj$ is the total consumption services bought by the workers (C_{jht} for each worker) in household h , and total purchases of new durable goods is similarly given by $I_{ht} = \int_0^1 I_{jht} dj$.

We go into details of a family behavior starting with the second sub-period of period t , when family i has inherited from the past a net debt position given by D_{ht} ; that is, the family owes a bank the amount D_{ht} . It is useful to note at this stage that, because all families are identical, D_{ht} will be zero in equilibrium.

The outstanding debt of the family D_{ht} is equally shared among the unit measure of workers. Worker j is told by the family head to go to bank with its share of existing debt D_{jht} ($= D_{ht}$) to apply for a new loan $L_{jht} \geq D_{jht}$, where the first D_{jht} units of that loan are used to settle the past debt and $L_{jht} - D_{jht}$ is left and available for spending. When granting the loan, the banks open the worker a checking account (which cannot have a negative balance), where the initial amount in the deposit would be $L_{jht} - D_{jht}$. As noted above, a worker can apply for their loan either or without family backing. We only consider the case where all workers apply for loans with family backing. This is not restrictive, as the risk premia on backed and non-backed loans will be such that workers will never choose a non-backed loan.

Banks can transfer the balances in the checking account to other agents, and these deposits earn the safe (central bank-determined) interest rate i at the beginning of next period. The workers can then use this bank money to order consumption goods C_{jht} at nominal price P_t and new durable goods I_{jht} at nominal price P_t^X . When a firm receives an order, the bank money is transferred to that firm, which then uses it to similarly place orders with the intermediate good firms (in the case of a final good firm) or to pay workers and rent durables (in the case of an intermediate good firm). With these latter payments, much of the bank money will be transferred back to (employed) workers and the household head (who maintains a family bank account) at the end of the period. The remaining bank money (reflecting intermediate good firm profits) is also transferred to the household head. After paying off their own loans, employed workers are asked by the family head to transfer their remaining bank money back to the family account.

The head of the family coordinates the family activities by telling workers how much to borrow, how much to purchase and how much of the bank money to transfer to the family. On top of that, the family head manages the stock of durable goods (or houses) X_{ht} of the family. She does so by renting it on the market at rate R_t^X .

In the third sub-period, firms wish to hire e_t workers from of the unit measure pool of workers, so that $1 - e_t$ workers will be left unemployed.¹⁰ Firms make take-it-or-leave-it offers to the workers in the form of a nominal wage W_t , who accept it as long as it is not below a reservation level set by the head of the family. In equilibrium, all offered jobs will be accepted. Because there is a unit measure of families with a unit measure of workers each, we will have $e_{ht} = e_t$, so that e_t is the probability of a worker in family h being employed.

¹⁰As the supply of workers is always inelastically one, equilibrium e_t must be less than or equal to one. To simplify notation, we will also let e_t denote the probability that a worker is offered a job, whereas it technically should be written $\min\{e_t, 1\}$ to allow for off-equilibrium labor demand.

Finally, as noted above, in the fourth sub-period, production takes place, all wage and rental payments are made, profits are transferred, an consumption and new durable goods are delivered to those who ordered them, which in turn are returned to the family to be split equally among its members.

In the first sub-period of the next period, debts need to be settled after interest on bank money is paid. Workers who were employed the last period ended it with a bank account balance of $W_t + L_{jht} - D_{jht} - P_t C_{jht} - P_t^X I_{jht}$. This balance receives an interest payment $1 + i_t$ at the beginning of period $t + 1$. The worker, meanwhile, owes the bank $(1 + r_t)L_{jht}$ at this point. The bank will limit its lending such that $(1 + r_t)L_{jht} \leq (1 + i_t)W_t$ so as to ensure that employed workers can always pay back their loans. Further, since $r_t \geq i_t$ (reflecting the risk premium), it will never be optimal for workers to borrow more than they intend to spend, and thus we will have $L_{jht} = D_{jht} + P_t C_{jht} + P_t^X I_{jht}$. Thus, an amount $T_{jht}^e = (1 + i_t)W_t - (1 + r_t)(D_{jht} + P_t C_{jht} + P_t^X I_{jht}) > 0$ will be transferred back to the family account by each employed worker.

The bank account balance of an unemployed worker at the end of period t , meanwhile, is given by $L_{jht} - D_{jht} - P_t C_{jht} - P_t^X I_{jht} = 0$, which as noted above is optimally set to zero. Further, as with employed workers, unemployed workers owe an amount $(1 + r_t)L_{jht}$ to the bank at the beginning of period $t + 1$. For these workers, the bank decides whether or not to pursue the family for repayment. If it is too costly (which which happens with probability ϕ), the loan is not repaid. With probability $1 - \phi$, pursuing the household is worthwhile for the bank, in which case the family transfers $T_{jht}^u = (1 + r_t)(D_{jht} + P_t C_{jht} + P_t^X I_{jht})$ to the worker so they can repay the loan.

Therefore, family h 's net debt at the beginning of period $t + 1$ will be given by

$$\begin{aligned} D_{ht+1} &= (1 - e_t) \phi T_{jht}^u - e_t T_{jht}^e - (1 + i_t) (R_t^X X_{ht} + \Pi_t) \\ &= [e_t + (1 - e_t) \phi] (1 + r_t) (D_{ht} + P_t C_{ht} + P_t^X I_{ht}^X) - (1 + i_t) (e_t W_t + R_t^X X_{ht} + \Pi_t), \end{aligned} \quad (\text{F.2})$$

where Π_t is total profits of all firms and banks.

F.4 Banks

Banks remunerate deposits at rate i_t , receive interest r_t on the fraction $e_t + (1 - e_t)\phi$ of loans that are repaid, and incur processing costs of Φ per unit of loans to bailed-out workers (i.e., those made to the fraction $(1 - e_t)\phi$ of workers who end up unemployed and are bailed out by the household). Thus, bank profits are

$$\Pi_t^{\text{banks}} = \{[e_t + (1 - e_t) \phi] (1 + r_t) - (1 + i_t) [1 + (1 - e_t) \phi \Phi]\} L_t,$$

where L_t is the total volume of loans. Free entry into banking implies zero profits, i.e.,

$$1 + r_t = (1 + i_t) \frac{1 + (1 - e_t) \phi \Phi}{e_t + (1 - e_t) \phi}. \quad (\text{F.3})$$

Defining the risk premium as $(1 + r_t) = (1 + i_t)(1 + r_t^p)$, the above equation implies

$$1 + r_t^p = \frac{1 + (1 - e_t) \phi \Phi}{e_t + (1 - e_t) \phi} \quad (\text{F.4})$$

If there is no unemployment risk ($e_t = 1$), or no default or recovery costs ($\phi = 1$ and $\Phi = 0$), one can check that there is no risk premium ($r_t^p = 0$).

F.5 Firms

The final good sector is competitive. This sector provides consumption services to households by buying a set of differentiated intermediate services, denoted \tilde{C}_{kt} , from intermediate service firms and combining them using a Dixit-Stiglitz aggregator. We assume a measure one of intermediate service firms, indexed by k . The objective of the final good firm is thus to solve

$$\max P_t C_t - \int_0^1 \tilde{P}_{kt} \tilde{C}_{kt} dj$$

subject to

$$C_t = \left(\int_0^1 \tilde{C}_{kt}^\eta dk \right)^{\frac{1}{\eta}}$$

with $\eta \in (0, 1)$ and where \tilde{P}_{kt} is the nominal price of intermediate service k . This gives rise to demand for intermediate service k given by

$$\tilde{C}_{kt} = \left(\frac{\tilde{P}_{kt}}{P_t} \right)^{-\frac{1}{1-\eta}} C_t.$$

Details of intermediate firms are presented in Section 3.1.2.

F.6 Equilibrium Outcomes

As noted in the text, the equilibrium outcome for this model is determined from a set of nine equations in the endogenous variables. The first two equilibrium conditions are the bank zero-profit condition (F.3) and the intermediate goods optimality condition (14). We derive the three next equilibrium conditions from the household's behavior. The head of family h maximizes

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \xi_{t-1} [U(C_{ht} - \gamma C_{t-1}) + v(1 - e_{ht})]$$

subject to the intertemporal budget constraint (F.2) and the durables accumulation equation (F.1). We can use the first-order conditions of this problem and the equilibrium conditions (the bank zero profit condition (F.3), $D_{ht} = 0$ and $C_{ht} = C_t$) to obtain

$$U'(C_t - \gamma C_{t-1}) = \beta \frac{\xi_t}{\xi_{t-1}} (1 + i_t) [1 + (1 - e_t) \phi \Phi] \mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{1 + \pi_{t+1}} \right], \quad (\text{F.5})$$

$$U'(C_t - \gamma C_{t-1}) = \beta \frac{\xi_t}{\xi_{t-1}} \mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{(1 + \pi_{t+1}) P_t^X} \left\{ \frac{R_{t+1}^X}{1 + (1 - e_{t+1}) \phi \Phi} + (1 - \delta) P_{t+1}^X \right\} \right], \quad (\text{F.6})$$

with an additional wage condition given by (13). Equation (F.5) is derived from the Euler equation for debt accumulation, (F.6) is derived from the Euler equation for the durable good accumulation, and the wage condition is the reservation wage of the household, which is in equilibrium the actual wage offered by firms. We also have the aggregate law of motion for the stock of durable goods (17) and the consumption services market-clearing condition

$$C_t = s [X_t + F(e_t, \theta_t)]. \quad (\text{F.7})$$

The last two equations will be given by the nominal interest rate policy rule of the central bank (15), and the optimal pricing decision of firms under Calvo adjustment costs which, as noted in the text, we do not need to explicitly obtain.

Block Recursive Structure

Using the monetary policy rule (15),¹¹ the equilibrium equations have a block recursive structure whereby the variables C_t , e_t , X_{t+1} and r_t^p can be solved for first using equations (F.3), (F.4), (F.7) and the combination of (F.5) and (15), which is given by

$$U'(C_t - \gamma C_{t-1}) = \beta \frac{\xi_t}{\xi_{t-1}} \Theta [e_t + (1 - e_t) \phi] (1 + r_t^p) \mathbb{E}_t \left[U'(C_{t+1} - \gamma C_t) e_{t+1}^{\phi_e} \right] \quad (\text{F.8})$$

Given that the above equations determine the quantity variables, the rest of the system then simultaneously determines the remaining variables $\left\{ \frac{R_t^X}{P_t}, \frac{P_t^X}{P_t}, \frac{W_t}{P_t}, \pi_t \right\}$. In particular, as we do not consider the implications of the model for inflation, we do not need to explicitly derive the optimal pricing behavior of firms. Finally, using (F.7) to eliminate C from the system, we end up with the three equations (18)-(16)

G Proofs of Propositions

Some Lemmas

We begin by establishing some useful lemmas. Letting $\phi \equiv 1 - \rho$, the characteristic polynomial of M is given by

$$P(\lambda) \equiv \lambda^3 - \left(1 - \delta + \frac{\phi}{\alpha_3} \right) \lambda^2 + \frac{\phi(1 - \delta) + \alpha_2 - \alpha_1}{\alpha_3} \lambda - (1 - \delta) \frac{\alpha_2}{\alpha_3}.$$

It will be useful in some cases to write P as

$$P(\lambda) = (\lambda - \lambda_2) (\lambda^2 - \tau\lambda + c),$$

where the roots of the quadratic polynomial $\lambda^2 - \tau\lambda + c$ are λ_{11} and λ_{12} , and $\lambda_{11}, \lambda_{12}, \lambda_2$ are as defined in the main text. In particular, note that

$$\lambda_{11} = \frac{\tau - \sqrt{\tau^2 - 4c}}{2}, \quad \lambda_{12} = \frac{\tau + \sqrt{\tau^2 - 4c}}{2}, \quad (\text{G.9})$$

and λ_2 , τ , and c satisfy the relationships

$$\tau + \lambda_2 = 1 - \delta + \frac{\phi}{\alpha_3}, \quad (\text{G.10})$$

$$c + \tau\lambda_2 = \frac{\phi(1 - \delta) + \alpha_2 - \alpha_1}{\alpha_3}, \quad (\text{G.11})$$

$$c\lambda_2 = (1 - \delta) \frac{\alpha_2}{\alpha_3}. \quad (\text{G.12})$$

Lemma 1. *Suppose Assumption (A1) holds. Then, for any $\phi > 0$ (i.e., any $\rho < 1$), at least one root of P is real and greater than 1, and if the remaining two roots are real then they are both non-zero and of the same sign. Further, for ϕ sufficiently small (i.e., ρ sufficiently large), these two remaining roots are either negative or complex. For ϕ sufficiently large, meanwhile, these two remaining roots must both be positive and less than 1.*

¹¹Specifically, we use the version in footnote 45.

Proof. First, note that

$$P(1) = -\frac{\alpha_1 - \delta(\alpha_2 + \alpha_3) + \delta\phi}{\alpha_3},$$

which is strictly negative by Assumption (A1). Further, since $P''' > 0$, we must have $P(\lambda) > 0$ for λ sufficiently large. Thus, we must have $P(\lambda) = 0$ for some $\lambda > 1$, which confirms that P has at least one real root greater than 1. Note that this implies that $\lambda_2 > 1$. Next, note that $\lambda_{11}\lambda_{12} = c$. By (G.12) and the fact that $\lambda_2 > 0$, we must have $c > 0$, which confirms that if $\lambda_{11}, \lambda_{12}$ are real then they are non-zero and of the same sign. Next, since $c > 0$, we see from (G.9) that if $\lambda_{11}, \lambda_{12}$ are real then they are of the same sign as τ . Since $\lambda_2 > 1$ we have that $\tau + 1 < \tau + \lambda_2$, and thus (G.10) implies $\tau < \phi/\alpha_3 - \delta$. For ϕ sufficiently small, the right-hand side of this inequality is negative, and thus so is τ , which confirms that if ϕ is sufficiently small and $\lambda_{11}, \lambda_{12}$ are real, then they must be negative. Finally, note that

$$P(0) = -\frac{(1-\delta)\alpha_2}{\alpha_3} < 0, \quad (\text{G.13})$$

$$P'(0) = \frac{\phi(1-\delta) + \alpha_2 - \alpha_1}{\alpha_3}, \quad (\text{G.14})$$

$$P''(0) = -\frac{2[(1-\delta)\alpha_3 + \phi]}{\alpha_3} < 0. \quad (\text{G.15})$$

Let $Q(\lambda)$ denote the second-order Taylor approximation to P around $\lambda = 0$. Since $P'''(\lambda) > 0$, we have $P(\lambda) > Q(\lambda)$ on $\lambda > 0$. Thus, if Q has a real root in $(0, 1)$, then so does P . Since $P(0) < 0$ and $P(1) < 0$, a real root in $(0, 1)$ would necessarily in turn imply that P has two such roots, i.e., that $\lambda_{11}, \lambda_{12}$ are both positive and less than 1. Now, it can be verified that the discriminant of Q is given by

$$\frac{[\phi(1-\delta) + \alpha_2 - \alpha_1]^2 - 4[(1-\delta)\alpha_3 + \phi](1-\delta)\alpha_2}{\alpha_3^2}.$$

For ϕ sufficiently large, this discriminant is positive, in which case Q has real roots. Further, for ϕ sufficiently large $Q'(0) = P'(0) > 0$, in which case at least one of these real roots of Q must be greater than 0. It can be further verified that the product of the roots of Q is equal to $(1-\delta)\alpha_2/[(1-\delta)\alpha_3 + \phi] > 0$. For ϕ sufficiently large, this product is less than 1, and thus Q must have at least one root between 0 and 1, which confirms that P has two such roots. \square

Lemma 2. *Suppose Assumption (A1) holds, and there exist ϕ_+, ϕ_- such that $\lambda_{11}, \lambda_{12}$ are strictly positive for $\phi = \phi_+$, and strictly negative for $\phi = \phi_-$. Then there exists a range of ϕ between ϕ_+ and ϕ_- such that $\lambda_{11}, \lambda_{12}$ are complex.*

Proof. Suppose to the contrary that, as ϕ goes from ϕ_+ to ϕ_- , one or both of $\lambda_{11}, \lambda_{12}$ transition from being strictly positive to being strictly negative without ever becoming complex. This would imply that at some point at least one of $\lambda_{11}, \lambda_{12}$ are zero. But by Lemma 1, this cannot happen, which completes the proof. \square

Lemma 3. *Suppose Assumption (A1) holds and $\lambda_{11}, \lambda_{12}$ are real with $\lambda_{12} < 1$. Then $\lambda_{12} < 1 - \delta$.*

Proof. If $\lambda_{11}, \lambda_{12}$ are negative, then the conclusion follows immediately. Thus, suppose $\lambda_{11}, \lambda_{12}$ are positive (which is the only other possibility). Note that $P(1-\delta) = -(1-\delta)\alpha_1/\alpha_3 < 0$. Since $P''' > 0$, this implies that we cannot have $\lambda_{11} \leq 1-\delta \leq \lambda_{12}$. Suppose by way of contradiction, then, that $1-\delta < \lambda_{11} \leq \lambda_{12}$. Note that, since P has three distinct real roots here one of which must be

greater than 1, and since $P''' > 0$, it must be the case that $P'(\lambda_{11}) > 0 > P'(\lambda_{12})$ and $P''(\lambda_{11}) < 0$. This implies that $P'(\lambda) > 0$ and $P''(\lambda) < 0$ for $\lambda < \lambda_{11}$. Since $1 - \delta < \lambda_{11}$ by hypothesis, this in turn implies that $P'(1 - \delta) > 0$ and $P''(1 - \delta) < 0$, which require, respectively, that

$$\phi < (1 - \delta) \alpha_3 + \frac{\alpha_2 - \alpha_1}{1 - \delta}, \quad (\text{G.16})$$

$$\phi > 2(1 - \delta) \alpha_3. \quad (\text{G.17})$$

Let $Q(\lambda)$ denote the first-order Taylor approximation to P around the point $\lambda = 1 - \delta$. Since $P''(\lambda) < 0$ on $(-\infty, \lambda_{11}]$ —an interval containing $1 - \delta$ —we must have $Q(\lambda) > P(\lambda)$ on $(-\infty, \lambda_{11}]$. Thus, the root $\hat{\lambda}$ of Q must satisfy $\hat{\lambda} < \lambda_{11}$. We may obtain that

$$\hat{\lambda} - 1 = \frac{\frac{\alpha_1 - \delta \alpha_2}{1 - \delta} + \delta [\phi - (1 - \delta) \alpha_3]}{(1 - \delta) \alpha_3 + \frac{\alpha_2 - \alpha_1}{1 - \delta} - \phi}.$$

By Assumption (A1), the first term in the numerator on the right-hand side of this expression is strictly positive, and by (G.17) so is the second. Further, the denominator is strictly positive by (G.16). Thus, we have found that $1 < \hat{\lambda} < \lambda_{11} \leq \lambda_{12}$, which contradicts our initial assumption that $\lambda_{12} < 1$. Thus, if $\lambda_{12} < 1$, we must have $\lambda_{11} \leq \lambda_{12} < 1 - \delta$. \square

Lemma 4. *Suppose Assumption (A1) holds and $\lambda_{11}, \lambda_{12}$ are real with $\lambda_{11} < \lambda_{12}$. Then the following properties hold:*

- (i) *If $0 < \lambda_{11} < \lambda_{12} < 1$, then $d\lambda_{11}/d\phi < 0$ and $d\lambda_{12}/d\phi > 0$.*
- (ii) *If $\lambda_{11} < \lambda_{12} < 0$, then $d\lambda_{11}/d\phi > 0$ and $d\lambda_{12}/d\phi < 0$.*

Proof. Note that, as argued in the proof of Lemma 3, it must be the case that $P'(\lambda_{11}) > 0$ and $P'(\lambda_{12}) < 0$. Totally differentiating the expression $P(\lambda_{1j}) = 0$ with respect to ϕ , we may obtain

$$\frac{d\lambda_{1j}}{d\phi} = \frac{[\lambda_{1j} - (1 - \delta)] \lambda_{1j}}{\alpha_3 P'(\lambda_{1j})}. \quad (\text{G.18})$$

By Lemma 3, we have $\lambda_{1j} < 1 - \delta$. Thus, $d\lambda_{1j}/d\phi$ is of the same sign as $-\lambda_{1j}/P'(\lambda_{1j})$. The conclusions of the Lemma then follow immediately. \square

Lemma 5. *Suppose Assumption (A1) holds and $\lambda_{11} = \lambda_{12} = \tilde{\lambda} < 1$ for $\phi = \phi_0$. Then the following properties hold:*

- (i) *If $\tilde{\lambda} > 0$, then as ϕ increases through ϕ_0 , $\lambda_{11}, \lambda_{12}$ transition from complex to positive.*
- (ii) *If $\tilde{\lambda} < 0$, then as ϕ increases through ϕ_0 , $\lambda_{11}, \lambda_{12}$ transition from negative to complex.*

Proof. Note first that $\lambda_{11} = \lambda_{12} = \tilde{\lambda}$ implies that $\tilde{\lambda}$ is real with $P(\tilde{\lambda}) = 0$, and further that $\tilde{\lambda}$ is also a local maximizer of P , so that $P'(\tilde{\lambda}) = 0$ and $P''(\tilde{\lambda}) < 0$. If the local maximum of P is increasing (decreasing) in ϕ at $\phi = \phi_0$, then this implies that $\lambda_{11}, \lambda_{12}$ are complex (real) for $\phi < \phi_0$ and real (complex) for $\phi > \phi_0$. Letting λ^* denote the local maximizer of P , by the envelope theorem we have

$$\left. \frac{dP(\lambda^*)}{d\phi} \right|_{\lambda^*=\tilde{\lambda}} = \left. \frac{\partial P(\lambda^*)}{\partial \phi} \right|_{\lambda^*=\tilde{\lambda}} = - \frac{[\tilde{\lambda} - (1 - \delta)] \tilde{\lambda}}{\alpha_3}.$$

By Lemma 3, $\tilde{\lambda} < 1 - \delta$, so that this derivative is of the same sign as $\tilde{\lambda}$. The conclusions of the Lemma then follow immediately. \square

We now present proofs of the propositions in the text.

Proof of Proposition 1

The steady state value of investment is given by the solution to

$$\left[1 + \frac{\alpha_1 - \delta(\alpha_2 + \alpha_3)}{\delta}\right] I^s - F(I^s) = \alpha_0 \quad (\text{G.19})$$

for I^s . Since $F'(I) < 1$ by assumption, the derivative of the left-hand side of (G.19) is greater than $[\alpha_1 - \delta(\alpha_2 + \alpha_3)]/\delta$. Under Assumption (A1), this derivative is strictly positive, and thus the left-hand side of (G.19) is strictly monotonic, and therefore the steady state is unique. \square

Proof of Proposition 2

To see part (a), note that as ρ decreases (i.e., ϕ increases) beginning from a situation with $0 < \lambda_{11} \leq \lambda_{12} < 1$, Lemmas 4 and 5 imply that $\lambda_{11}, \lambda_{12}$ initially remain real with λ_{11} decreasing and λ_{12} increasing, continuing to do so as long as we have $\lambda_{11} > 0$ and $\lambda_{12} < 1$ (since λ_{11} and λ_{12} are moving apart, they cannot become complex without one of these two inequalities first becoming violated). Suppose one of these two inequalities becomes violated at some point. By Lemma 1, λ_{11} and λ_{12} are of the same sign, and thus the first violation cannot involve λ_{11} becoming negative. Further, by Lemma 3 we have $\lambda_{12} < 1 - \delta$, and thus the first violation cannot involve λ_{12} becoming greater than 1. Thus, a violation cannot happen, which confirms that we continue to have $0 < \lambda_{11} \leq \lambda_{12} < 1$ as ρ decreases without bound. By Lemma 1, λ_2 must remain greater than 1, which confirms the first part of the proposition.

To see part (b), note that as ρ increases (i.e., ϕ decreases), Lemma 4 implies that as long as $\lambda_{11} < \lambda_{12}$, λ_{11} will increase and λ_{12} will decrease, with both remaining stable. By Lemma 1, at some point $\lambda_{11}, \lambda_{12}$ must become negative or complex for ϕ sufficiently small, and by Lemma 2 this necessarily implies that the first time this happens they become complex. Thus, at some point as ϕ continues to fall we must have $\lambda_{11} = \lambda_{12} < 1$, and at that point by Lemma 5 they must transition to complex and stable. Further, again by Lemma 1, λ_2 must remain greater than 1, which completes the proof. \square

Proof of Proposition 3

That λ_2 remains greater than 1 as ρ increases all the way to 1 (i.e., as ϕ decreases all the way to zero) follows immediately from Lemma 1. To see the remaining part of the proposition, note that the only case ruled out is one in which at some point $\lambda_{11}, \lambda_{12}$ return to being real, positive and stable. By Lemma 2, if this were to happen it must involve a transition from complex to real, positive and stable, but Lemma 5 rules out such a transition as ϕ falls, which completes the proof. \square

Proof of Proposition 4

Note that Assumption (A2) can be rephrased as the requirement that if $\alpha_1 > \alpha_2$ then

$$(\alpha_1 - \alpha_2)^2 < 4\alpha_2(1 - \delta)^2\alpha_3. \quad (\text{G.20})$$

Note from (G.13)-(G.15) that if $\phi \geq (\alpha_1 - \alpha_2)/(1 - \delta)$, then P is negative, weakly increasing, and concave at $\lambda = 0$. Since $P''' > 0$, this implies that P must be strictly increasing on $\lambda < 0$, and since $P(0) < 0$, it is not possible for P to have a negative real root in this case. Thus, suppose

instead that $\phi < (\alpha_1 - \alpha_2)/(1 - \delta)$, which can only happen if $\alpha_1 > \alpha_2$. In this case we see from (G.14) that $P'(0) < 0$. Let $Q(\lambda)$ be the second-order Taylor approximation to P around $\lambda = 0$ and note that if Q has any real roots then they must be negative. Note also that, since $P''' > 0$, we have $Q(\lambda) > P(\lambda)$ on $\lambda < 0$. Thus, if P has a negative root then so must Q . A sufficient condition to rule out negative roots of P is thus to rule out the possibility of Q having real roots. Now, the discriminant of Q is given by

$$\Delta_Q = \frac{[\phi(1 - \delta) - \alpha_2 - \alpha_1]^2 - 4\alpha_2\alpha_1 - 4\alpha_2(1 - \delta)^2\alpha_3}{\alpha_3^2}.$$

We have

$$\frac{\partial \Delta_Q}{\partial \phi} = \frac{2(1 - \delta)[\phi(1 - \delta) - \alpha_2 - \alpha_1]}{\alpha_3^2} < -\frac{4(1 - \delta)\alpha_2}{\alpha_3^2} < 0$$

where the first inequality follows from the fact that $\phi < (\alpha_1 - \alpha_2)/(1 - \delta)$ by hypothesis. Thus, Δ_Q is maximized when $\phi \rightarrow 0$. Now,

$$\lim_{\phi \rightarrow 0} \Delta_Q = \frac{(\alpha_1 - \alpha_2)^2 - 4\alpha_2(1 - \delta)^2\alpha_3}{\alpha_3^2}.$$

Since we have $\alpha_1 > \alpha_2$ in this case, under Assumption (A2) inequality (G.20) holds, which implies that $\lim_{\phi \rightarrow 0} \Delta_Q < 0$. Thus, $\Delta_Q < 0$ for $\phi < (\alpha_1 - \alpha_2)/(1 - \delta)$, and thus neither Q nor P can have negative real roots in this case either. This completes the proof. \square

Proof of Proposition 5

Note first that, since we maintain Assumption (A2), Proposition 4 rules out case (c) of Proposition 3. Thus, establishing that a Hopf bifurcation will occur is equivalent to establishing that there is some $\phi > 0$ for which $\lambda_{11}\lambda_{12} = 1$. Since $\lambda_{11}\lambda_{12} = c$, substituting $c = 1$ into (G.10)-(G.12) and solving for ϕ in terms of parameters, we may obtain

$$\phi = \frac{(1 - \delta)^2 \frac{\alpha_2^2}{\alpha_3} - \alpha_3 - \alpha_1 + [1 - (1 - \delta)^2] \alpha_2}{(1 - \delta) \left(\frac{\alpha_2}{\alpha_3} - 1 \right)}. \quad (\text{G.21})$$

A Hopf bifurcation will occur if this solution is strictly positive. Under Assumption (A3), the numerator of the right-hand side of (G.21) is strictly positive. Assumption (A3) can also be rearranged to produce $\alpha_2 - \alpha_3 > \alpha_1/[1 + (1 - \delta)^2\alpha_2/\alpha_3] > 0$, which ensures that the denominator on the right-hand side of (G.21) is positive as well, and thus there exists a positive value of ϕ for which $c = 1$, which completes the proof. \square

Proof of Proposition E.1

Case (a): $\alpha_3 = 0$ (Hopf bifurcation):

For this case, we make use of ? theorem and of the formulation given by ? (see ? for a comprehensive exposition of bifurcation theory). In this case, we may write our non-linear system (in deviations from the steady state) as

$$\begin{pmatrix} I_t - F(I_t) \\ X_t \end{pmatrix} = \begin{pmatrix} \alpha_2 - \alpha_1 & -\alpha_1(1 - \delta) \\ 1 & 1 - \delta \end{pmatrix} \begin{pmatrix} I_{t-1} \\ X_{t-1} \end{pmatrix}. \quad (\text{G.22})$$

To study the stability of the limit cycle when this system undergoes a Hopf bifurcation, we need to first transform the system into the “standard form” at the bifurcation point

$$\begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \end{pmatrix} + \begin{pmatrix} f(y_{1t-1}, y_{2t-1}) \\ g(y_{1t-1}, y_{2t-1}) \end{pmatrix}, \quad (\text{G.23})$$

where $(y_1, y_2)'$ is a smooth invertible function of $(I, X)'$, and f and g are $O(\|(y_1, y_2)'\|^2)$. Define $H(I_t) \equiv I_t - F(I_t)$. Under our restriction $F'(\cdot) < 1$, H is a strictly increasing function, and is therefore invertible. Let $G(\cdot) \equiv H^{-1}(\cdot)$, and note that $G'(0) = 1/(1 - \rho)$. Letting

$$m(I_{t-1}, X_{t-1}) \equiv G((\alpha_2 - \alpha_1)I_{t-1} - \alpha_1(1 - \delta)X_{t-1}) - \frac{\alpha_2 - \alpha_1}{1 - \rho}I_{t-1} + \frac{\alpha_1(1 - \delta)}{1 - \rho}X_{t-1},$$

we may re-write system (G.22) as

$$\begin{pmatrix} I_t \\ X_t \end{pmatrix} = \underbrace{\begin{pmatrix} \frac{\alpha_2 - \alpha_1}{1 - \rho} & -\frac{\alpha_1(1 - \delta)}{1 - \rho} \\ 1 & 1 - \delta \end{pmatrix}}_M \begin{pmatrix} I_{t-1} \\ X_{t-1} \end{pmatrix} + \begin{pmatrix} m(I_{t-1}, X_{t-1}) \\ 0 \end{pmatrix}, \quad (\text{G.24})$$

Note that m does not contain any constant or first-order terms (i.e., m is $O(\|(I, X)'\|^2)$).

Near the Hopf bifurcation, the eigenvalues of M are complex, and can be denoted $\lambda, \bar{\lambda}$ (which implicitly depend on ρ). At the Hopf bifurcation these eigenvalues have modulus 1, and we can write $\lambda = \lambda_0 \equiv \cos \theta + i \sin \theta$ for some $\theta \in (0, \pi)$. Let Λ and C be the matrices

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

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

Note that, by construction, λ_0 and $\bar{\lambda}_0$ are the eigenvalues of C . We introduce matrices

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

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

whose columns are eigenvectors of C and M , respectively. Suppose we are at the Hopf bifurcation point, so that $\lambda = \lambda_0$. Then $C = V_C \Lambda V_C^{-1} = V_C V_M^{-1} M V_M V_C^{-1} = B M B^{-1}$, where

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

Letting $(y_{1t}, y_{2t})' = B(I_t, X_t)'$, we obtain a system in the “standard form” (G.23), with

$$f(y_{2t-1}, y_{2t-1}) = 0,$$

$$g(y_{1t-1}, y_{2t-1}) = -G \left(\frac{\gamma_1}{\sin \theta} y_{1t-1} - \gamma_2 y_{2t-1} \right) + \frac{1}{1 - \rho} \left(\frac{\gamma_1}{\sin \theta} y_{1t-1} - \gamma_2 y_{2t-1} \right),$$

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

Let $\rho_0 < 1$ denote the value of ρ (the ‘‘bifurcation parameter’’) at which the Hopf bifurcation occurs. Define

$$b \equiv \left. \frac{d|\lambda|}{d\rho} \right|_{\rho=\rho_0}$$

and

$$a \equiv -\operatorname{Re} \left(\frac{(1-2\lambda_0)\bar{\lambda}_0^2}{1-\lambda_0} \xi_{11}\xi_{20} \right) - \frac{1}{2} |\xi_{11}|^2 - |\xi_{02}|^2 + \operatorname{Re}(\bar{\lambda}_0 \xi_{21}),$$

where

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

and the $f_{(\cdot)}$'s and $g_{(\cdot)}$'s indicate the relevant second- and third-order partial derivatives of f and g , respectively, evaluated at the steady state. According to ?, the Hopf bifurcation is supercritical if $b > 0$ and $a < 0$. To verify the first condition, note that $|\lambda| = |M| = \frac{\alpha_2(1-\delta)}{1-\rho}$, so that

$$b = \frac{(1-\rho) + \alpha_2(1-\delta)}{(1-\rho)^2} > 0.$$

Consider next the expression for a . Since $G(I)$ is the inverse of $H(I) = I - F(I)$, we have

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

Thus, when F''' becomes large (in absolute terms) and negative, so does G''' . In the expression for a , the first three terms do not depend on F''' , while the last term is given by

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

The coefficient on G''' in this expression is strictly positive. Thus, for G''' sufficiently negative, $\operatorname{Re}(\bar{\lambda}\xi_{21})$ will in turn be sufficiently negative so that $a < 0$. Thus, we have verified that if F''' is sufficiently negative, by ? theorem the Hopf bifurcation will be supercritical.

Case (b): $\alpha_2 = 0$ (Flip bifurcation):

In this case our non-linear system is given (in deviations from the steady state) by

$$\begin{aligned} X_{t+1} &= (1-\delta)X_t + I_t, \\ I_{t+1} &= \frac{\alpha_1}{\alpha_3}X_t + \frac{1}{\alpha_3}H(I_t), \end{aligned}$$

where as above $H(I_t) \equiv I_t - F(I_t)$. We proceed in two steps. First, we obtain a candidate solution of the form

$$I_t = \Phi(X_t), \tag{G.25}$$

$$X_{t+1} = \chi(X_t) \equiv (1-\delta)X_t + \Phi(X_t), \tag{G.26}$$

where Φ and χ are some functions. Equation (G.26) here captures the dynamics of this solution as a one-dimensional system. The second step will then be to establish that if this solved system undergoes a flip bifurcation as ρ increases toward 1, then for F''' sufficiently negative that bifurcation will be supercritical.

To a first-order approximation, we may write our system as

$$\begin{pmatrix} X_{t+1} \\ I_{t+1} \end{pmatrix} = \underbrace{\begin{pmatrix} 1-\delta & 1 \\ \frac{\alpha_1}{\alpha_3} & \frac{1-\rho}{\alpha_3} \end{pmatrix}}_M \begin{pmatrix} X_t \\ I_t \end{pmatrix}.$$

The eigenvalues of the matrix M are given by

$$\lambda_1 = \frac{1}{2} \left[1 - \delta + \frac{1-\rho}{\alpha_3} - \sqrt{\left(1 - \delta - \frac{1-\rho}{\alpha_3}\right)^2 + 4\frac{\alpha_1}{\alpha_3}} \right],$$

$$\lambda_2 = \frac{1}{2} \left[1 - \delta + \frac{1-\rho}{\alpha_3} + \sqrt{\left(1 - \delta - \frac{1-\rho}{\alpha_3}\right)^2 + 4\frac{\alpha_1}{\alpha_3}} \right].$$

It can be verified that $\lambda_2 > 1$ (so that we cannot have indeterminacy), and that if $\alpha_3 < \alpha_1/(2-\delta)$, then there is a $\rho_0 < 1$ such that if $\rho = \rho_0$ then $\lambda_1 = -1$ (i.e., the system undergoes a flip bifurcation at ρ_0). Assume henceforth that this is the case.

Next, using (G.25)-(G.26) in our system, we may obtain the relationships

$$\chi(X_t) = (1-\delta)X_t + \Phi(X_t), \quad (\text{G.27})$$

$$\Phi((1-\delta)X_t + \Phi(X_t)) = \frac{\alpha_1}{\alpha_3}X_t + \frac{1}{\alpha_3}H(\Phi(X_t)). \quad (\text{G.28})$$

Sequentially differentiating (G.28) three times and evaluating at the steady state $X = 0$, we obtain the equations

$$\Phi'^2 + \left(1 - \delta - \frac{1-\rho}{\alpha_3}\right)\Phi' - \frac{\alpha_1}{\alpha_3} = 0, \quad (\text{G.29})$$

$$\Phi'' = \frac{H''\Phi'^2}{\alpha_3 \left[(1-\delta + \Phi')^2 + \Phi' \right] - (1-\rho)}, \quad (\text{G.30})$$

$$\Phi''' = \frac{H''' \Phi'^3 + 3\Phi'' [H''\Phi' - \alpha_3(1-\delta + \Phi')\Phi'']}{\alpha_3 \left[(1-\delta + \Phi')^3 + \Phi' \right] - (1-\rho)}, \quad (\text{G.31})$$

which implicitly give the coefficients of the first three terms in the Taylor series representation of Φ . It can be verified that the two values of Φ' satisfying (G.29) are given by $\lambda_j - (1-\delta)$, $j = 1, 2$. It can be further verified that if there is a solution satisfying the transversality condition, then it cannot be the one associated with $j = 2$ (since we would then have $\chi' > 1$). Thus, take $\Phi' = \lambda_1 - (1-\delta)$, so that $\chi' = \lambda_1$. In this case, as noted above, system (G.26) will undergo a flip bifurcation at some point $\rho = \rho_0$.

Substituting $\Phi' = \lambda_1 - (1-\delta)$ into (G.30) one can obtain Φ'' , and from this one can obtain Φ''' from (G.31). Since $\chi' = \lambda_1$, $\chi'' = \Phi''$, and $\chi''' = \Phi'''$, we can then write

$$\chi(X) = \lambda_1 X + \frac{1}{2}\Phi'' X^2 + \frac{1}{6}\Phi''' X^3 + O(|X|^4).$$

As noted in ?,¹² since $\partial\chi/\partial\rho|_{X=0} = 0$, the flip bifurcation that this system undergoes at $\rho = \rho_0$ will be supercritical if

$$\frac{\partial\chi(X)}{\partial X\partial\rho} = \frac{\partial\lambda_1}{\partial\rho} \neq 0$$

and

$$a \equiv \left[\frac{1}{2}\Phi''^2 + \frac{1}{3}\Phi''' \right]_{\rho=\rho_0} > 0.$$

It is straightforward to verify that the first condition holds using the expression above for λ_1 . To check the second condition, note that when $\rho = \rho_0$ we have $\Phi' = -(2 - \delta) < 0$, and thus from (G.31) we have $\partial\Phi'''/\partial H''' > 0$. Thus, since $H''' = -F'''$, if F''' is sufficiently negative we will have $a > 0$, in which case the flip bifurcation will be supercritical. This completes the proof. \square

Proof of Proposition 6

The proof follows obviously from the fact that, if an arbitrage opportunity existed, the rationally optimizing agents in our model would attempt to take advantage of it, and thus the model could not have been in equilibrium to begin with.

H Spectral Density Estimation

H.1 Schuster's Periodogram

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

$$X_k = X \left(e^{i\frac{2\pi}{T}k} \right) = \sum_{t=0}^{T-1} x_t e^{-i\frac{2\pi}{T}kt}, \quad (\text{H.32})$$

for $k = 1, \dots, T - 1$. We then can compute samples of the Sample Spectral Density (SSD) S_k from samples of Schuster's periodogram I_k ¹³ according to

$$S_k = I_k = \frac{1}{T} |X_k|^2 \quad (\text{H.33})$$

Taking advantage of the fact that X is even, this amounts to evaluating the spectral density at T frequencies equally spaced between 0 and π .¹⁴

¹²See Theorem 3.5.1 on p.158 and the subsequent discussion.

¹³Another approach for obtaining the spectral density is to take a Fourier transform of the sequence of autocovariances of x . We show below that this method gives essentially the same result when applied to our hours series.

¹⁴See ? for a detailed exposition of spectral analysis, ? for practical implementation and ? for a quick introduction.

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

As we have computed only T samples of the DTFT $X(e^{i\omega})$, we might not have a detailed enough picture of the shape of the underlying function $X(e^{i\omega})$, and therefore of the spectral density $|X(e^{i\omega})|^2$. This problem is particularly acute if one is interested in the behavior of the spectrum at longer periodicities (i.e., lower frequencies). Specifically, since we uniformly sample frequencies, and since the periodicity p corresponding to frequency ω is given by $p = \frac{2\pi}{\omega}$, the spectrum is sparser at longer periodicities (and denser at shorter ones). While the degree of accuracy with which the samples of X_k describe the shape of $X(e^{i\omega})$ is dictated and limited by the length T of the data set, we can nonetheless increase the number of points at which we sample the DTFT in order to increase the graphic resolution of the spectrum. One common (and numerically efficient) way to do this is to add a number of zeros to the end of the sequence x_t before computing the DFT. This operation is referred to as zero-padding. As an example, suppose that we add exactly T zeros to the end of the length- T sequence $\{x_t\}$. One can easily check that this has no effect on the DFT computed at the original T sampled frequencies, instead simply adding another set of T sampled frequencies at the midpoints between each successive pair of original frequencies.¹⁵

If one is interested in the behavior of the spectral density at long enough periodicities, zero-padding in this way is useful. We will denote by N the number of samples at which the DTFT (and thus the SSD) is sampled, meaning that $T' = N - T$ zeros will be added to the sequence $\{x_t\}$ before computing the DFT. In the main text, we have set $N = 1,024$.¹⁶

H.3 Smoothed Periodogram Estimates

We obtain the raw spectrum estimate of a series non-parametrically as the squared modulus of the DFT of the (zero-padded) data sequence, divided by the length of the data set.¹⁷ This estimate is called Schuster's periodogram, or simply the periodogram. It turns out that the periodogram is asymptotically unbiased, but is not a consistent estimate of the spectrum, and in particular the estimate of the spectrum at a given frequency ω_k is generally quite unstable (i.e., it has a high standard error). Notwithstanding this fact, the over all pattern of the spectrum is much more stable, in the sense that the average value of the estimated spectrum within a given frequency band *surrounding* ω_k is in fact consistent. In order to obtain a stable and consistent estimate of the spectrum, we exploit this fact by performing frequency-averaged smoothing. In particular, we obtain our estimate of the SSD $S(\omega)$ by kernel-smoothing the periodogram $I(\omega)$ over an interval of frequencies centered at ω . Since the errors in adjacent values of $I(\omega)$ are uncorrelated in large samples, this operation reduces the standard errors of the estimates without adding too much bias. In our estimations, we use a Hamming window of length $W = 13$ as the smoothing kernel.¹⁸

¹⁵This is true when the number of zeros added to the end of the sample is an integer multiple of T . When instead a non-integer multiple is added, the set of frequencies at which the padded DFT is computed no longer contains the original set of points, so that the two cannot be directly compared in this way. Nonetheless, the over all pattern of the sampled spectrum is in general unaffected by zero-padding.

¹⁶As is well known, standard numerical routines for computing the DFT (i.e., those based on the Fast Fourier Transform algorithm) are computationally more efficient when N is a power of 2, which is why we set $N = 1,024$ rather than, say, $N = 1,000$.

¹⁷Note that we divide by the original length of the series (i.e., T), rather than by the length of the zero-padded series (i.e., N).

¹⁸Using alternative kernel functions makes little difference to the results.

H.4 Confidence Intervals

For the estimated smoothed SSD $\widehat{S}(\omega)$, we have asymptotically

$$\widehat{S}(\omega) \sim S(\omega) \frac{\chi_{2L}^2}{2L},$$

where L is the length of the smoothing window. The point estimates and 95% confidence bands for the spectral density of Non-Farm Business Hours, Total Hours per capita, Unemployment rate and Capacity utilization rate (all in levels) are displayed on Figure 24.

[Figure 24 about here.]

H.5 Smoothing and Zero-Padding with a Multi-Peaked Spectral Density

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

Each of the x_i follows an AR(2) process

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

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

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

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

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

We set $\bar{\omega}_i$ equal to 20, 40, and 100 quarters, respectively, for the three processes, and ρ_{i2} equal to -0.9, -0.95, and -0.95. The corresponding values for ρ_{i1} are 1.802, 1.9247, and 1.9449. We set σ_i equal to 6, 2, and 1. We then construct $x_t = x_{1t} + x_{2t} + x_{3t}$. The theoretical spectral density of x is shown in Figure 25. As in the factor utilisation series we are using in the main text, the spectral density shows long-run fluctuations, but the bulk of the business cycle movements is explained by movements at the 40-quarter periodicity, although we observe another peak at periodicity 20 quarters.

[Figure 25 about here.]

We simulate this process 1,000,000 times, with $T = 270$ for each simulation, which is the length of our observed macroeconomic series. We estimate the spectral density for various values of N (zero-padding) and W (length of the Hamming window). Higher N corresponds to higher resolution, and higher W to more smoothing. On each panel of Figure 26, we report the mean of the estimated spectrum over the 1,000,000 simulations (solid grey line), the mean \pm one standard

deviation (dashed lines), and the theoretical spectrum (solid black line). As we can see moving down the figure (i.e., for increasing W), more smoothing tends to reduce the error variance, but at the cost of increasing bias. Effectively, the additional smoothing “blurs out” the humps in the true spectrum. For example, with no zero-padding ($N = 270$), the peak in the spectral density at 40 quarters is (on average) hardly detected once we have any smoothing at all. Meanwhile, moving rightward across the figure (i.e., for increasing N), we see that more zero-padding tends to reduce the bias (and in particular, allows for the humps surrounding the peaks to be better picked up on average), but typically increases the error variance. As these properties suggest, by appropriately choosing the combination of zero-padding and smoothing, one can minimize the error variance while maintaining the ability to pick up the key features of the true spectrum (e.g., the peaks at 20 and 40 quarters).

[Figure 26 about here.]

H.6 Smoothing and Zero-Padding with Non-Farm Business Hours per Capita

Figure 27 presents estimates of the spectral density of U.S. Non Farm Business Hours per Capita for different choices of the zero-padding parameter (N) and different lengths of the Hamming window (W). The results indicate that, as long as the amount of zero-padding is not too small (i.e., N larger), we systematically observe the peak at around 40 quarters in the spectral density. In fact, it is only with minimal zero-padding (N low) and a wide smoothing window (W high) that the peak is entirely washed out. We take this as evidence of the robustness of that peak.

[Figure 27 about here.]

H.7 Detrending with a Polynomial Trend

In this section, we check that detrending our hours series with a polynomial trend of degrees 1 to 5 does not affect our main finding; namely, the existence of a peak in the spectrum at a periodicity around 40 quarters. Plots confirming that our finding is robust to polynomial detrending are shown in Figure 28.

[Figure 28 about here.]

H.8 Alternative Estimators

As another robustness test, we estimate the spectrum using the `SPECTRAN` package (for `MATLAB`), which is described in ?. The spectrum is computed in this case as the Fourier transform of the covariogram (rather the periodogram as we have done thus far). Smoothing is achieved by applying a window function of length M to the covariogram before taking its Fourier transform.¹⁹ Three different window shapes are proposed: the Blackman-Tukey window, Parzen window, and Tukey-Hanning window. The width of the window used in estimation is set as a function of the number of samples of the spectrum. In the case where no zero-padding is done ($N = 270$), these “optimal”

¹⁹Specifically, the k -th-order sample autocovariance is first multiplied by $w(|k|)$, where the window function w is an even function with the property that $\max_k w(k) = w(0) = 1$, and the window length $M > 0$ is such that $w(|k|) \neq 0$ for $|k| = M - 1$ and $w(|k|) = 0$ for $|k| \geq M$.

widths correspond to lengths of, respectively, $M = 68, 89,$ and 89 quarters for the three methods.²⁰ Figure 29 shows the estimated spectrum of Non Farm Business hours for the three windows and with or without zero-padding ($N = 270, 512,$ or 1024). Results again confirm the existence of a peak at a periodicity around 40 quarters, as long as there is enough zero-padding.

[Figure 29 about here.]

²⁰Note that, in contrast to the kernel-smoothing case, in this case a wider window corresponds to less smoothing.

I Solving and Estimating the Model

I.1 Deriving the Estimated Equations

Substituting the functional forms $U(c) = (c^{1-\omega} - 1)/(\omega - 1)$ and $F(e, \theta) = \theta e^\alpha$, as well as the normalizations $s = \theta = 1$ and the definition $X_t^* \equiv X_t^{1/(\alpha\delta)}$ into equations (18) and (19) from the main text, equations (20) and (21) can then be obtained by linearizing these two equations around the non-stochastic steady state with respect to $\log(e)$, $\log X^*$, r_t^p , and $\mu_t \equiv -\Delta \log(\xi_t)$. In particular, the coefficients in (21) are given by

$$\begin{aligned}\alpha_1 &\equiv \frac{\psi\alpha\delta^2(1-\delta-\gamma)}{(1-\delta)\kappa}, \\ \alpha_2 &\equiv \frac{\gamma\alpha\delta(1-\delta-\psi)}{(1-\delta)\kappa}, \\ \alpha_3 &\equiv \frac{\alpha\delta - \tau\phi_e}{\kappa}, \\ \alpha_4 &\equiv \frac{\tau}{\kappa},\end{aligned}$$

where $\tau \equiv (1-\gamma)(\psi+\delta)/\omega$, $\kappa \equiv (1+\gamma-\psi)\alpha\delta + \tau\Xi$, $\Xi \equiv (1-\phi)e^s/[e^s + (1-e^s)\phi]$, and e^s is the steady state employment rate. Further, for the linear and non-linear RP models, we may obtain the (log-)linearized version of equation (17) as $\hat{r}_t^p = R_1^p \hat{e}_t$, where $R_1^p = -\Xi\{1 + \phi\Phi/[(1-\phi)R^p]\}$ and $R^p \equiv [1 + (1-e^s)\phi\Phi]/[e^s + (1-e^s)\phi]$ is the steady state gross risk premium.

I.2 Solution Method

For the three linear versions of the model (linear RP, no friction, and canonical), we use standard (linear) perturbation methods to solve the model (see, e.g., ?). For the non-linear RP model, since we would like to allow for the possibility of local instability and limit cycles, standard non-linear perturbation methods are not appropriate. In particular, as a first step standard methods require obtaining a rational-expectations solution to the linearized (i.e., linear RP, in this case) model, and if the model is locally unstable (as is necessarily the case if it features limit cycles) such a solution cannot exist. We therefore use instead the perturbation method discussed in ?. The method involves finding a two-dimensional manifold such that, when the system is projected onto that manifold by choice of the jump variable (e_t in our case), the system remains bounded (in expectation), while any other choice of the jump variable would cause the system to explode.²¹ When the system has a typical saddle-path stable structure (i.e., one unstable and two stable eigenvalues), this manifold is tangent to the two-dimensional real linear subspace spanned by the two “stable” eigenvectors.²² When the system instead has a “saddle limit cycle” (e.g., one of the form shown in Figure 9), it will have one positive unstable eigenvalue and either a complex pair of unstable eigenvalues or at least one negative unstable eigenvalue (with the other strictly less than 1). The manifold we use (if it exists) is again tangent to a two-dimensional real linear subspace spanned by a pair of eigenvectors, but in this case the eigenvectors in question are the ones associated with the two eigenvalues that

²¹See Figure 9 in the text for a graphical illustration.

²²Note that, for this case, standard perturbation methods and the method we use coincide.

are not positive and unstable.²³ See ? for further details on the method.

In practice, we solve the non-linear RP model in this manner to a third-order approximation. Given such a parameterization and an associated solution in state-space form, we need to verify that the system indeed remains bounded in expectation. It is not possible to confirm this analytically, so we employ numerical methods instead. In particular, to minimize computational burden, we simply check whether, for a given initial (non-zero) state of the system,²⁴ the deterministic path for hours (i.e., the one obtained by feeding in a constant stream of zeros for the shock) explodes, where we define an explosion as a situation where, within the first 270 simulated quarters (the length of our data set), the absolute value of hours exceeds $10 \times \bar{L}$, where \bar{L} is the maximum absolute value of hours (in log-deviations from the mean) observed in our data set ($\bar{L} = 14.586$).

I.3 Estimation Procedure

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

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

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

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

I.3.1 The Parameter Space

We estimate the parameters of the model imposing several restrictions on the parameter space Θ . First, we require that $0 \leq \gamma, \psi < 1$. Second, we require that the policy rate reacts positively to expected hours, but not so strongly as to cause current hours to fall in response to an increase in expected hours, i.e., $0 < \phi_e < \alpha\delta/\tau$, where τ is defined above. Third, we impose that the degree of complementarity is never so strong as to generate temporary multiple equilibria.²⁵ This latter property is ensured if the function $\hat{e}_t + \frac{\tau}{\kappa} R^p(\hat{e}_t)$ (see equation (22) in the main text) is strictly increasing in \hat{e}_t (so that it is invertible). This in turn requires $1 + \frac{\tau}{\kappa} R_1^p > 0$ and $R_3^p \geq R_2^{p^2}/[3(\kappa/\tau + R_1^p)]$. Fourth, we impose that the shock process is stationary, i.e., $|\rho| < 1$. Finally, we require that the parameters be such that a solution to the model exists and is unique (see Appendix

²³If the system has more than one positive unstable eigenvalue, we assume it has no rational expectations solution, and discard the associated parameterization. We also discard any parameterization for which there are no positive unstable eigenvalues, as such configurations will typically be associated with indeterminacy. Note that we do not find any evidence that this constraint on the parameter space—i.e., that there not be indeterminacy—is binding for either the linear or non-linear RP models.

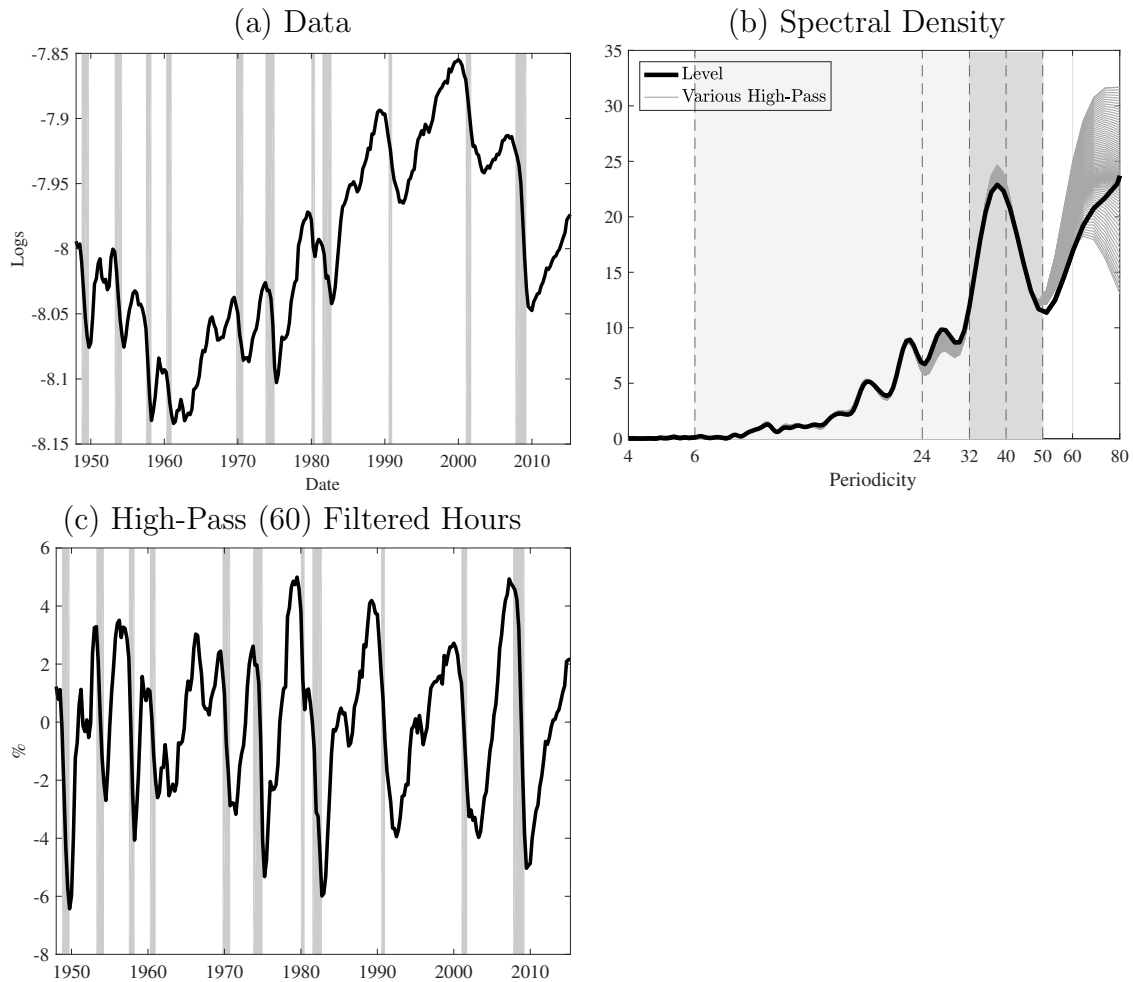
²⁴We use the same initial point for every parameterization.

²⁵By temporary multiple equilibria, we mean a situation where, for a given X_t , e_{t-1} and expectation about e_{t+1} , there are multiple values of e_t consistent with the dynamic equilibrium condition.

I.2). Note that none of the estimated parameters in either the linear or non-linear RP models is on the boundary of the set of constraints we have imposed.

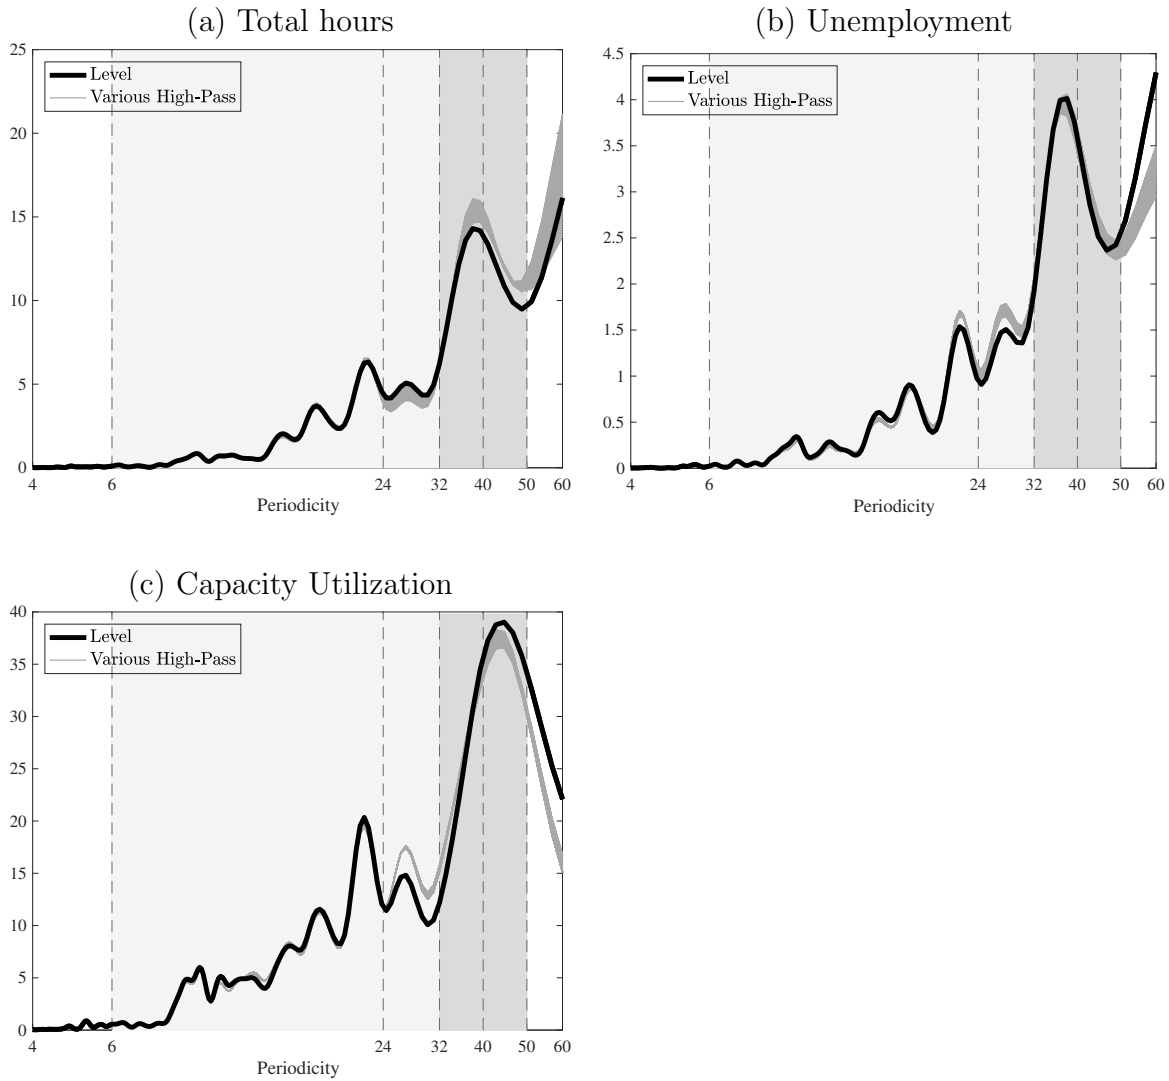
Figures

Figure 1: Properties of Hours Worked per Capita



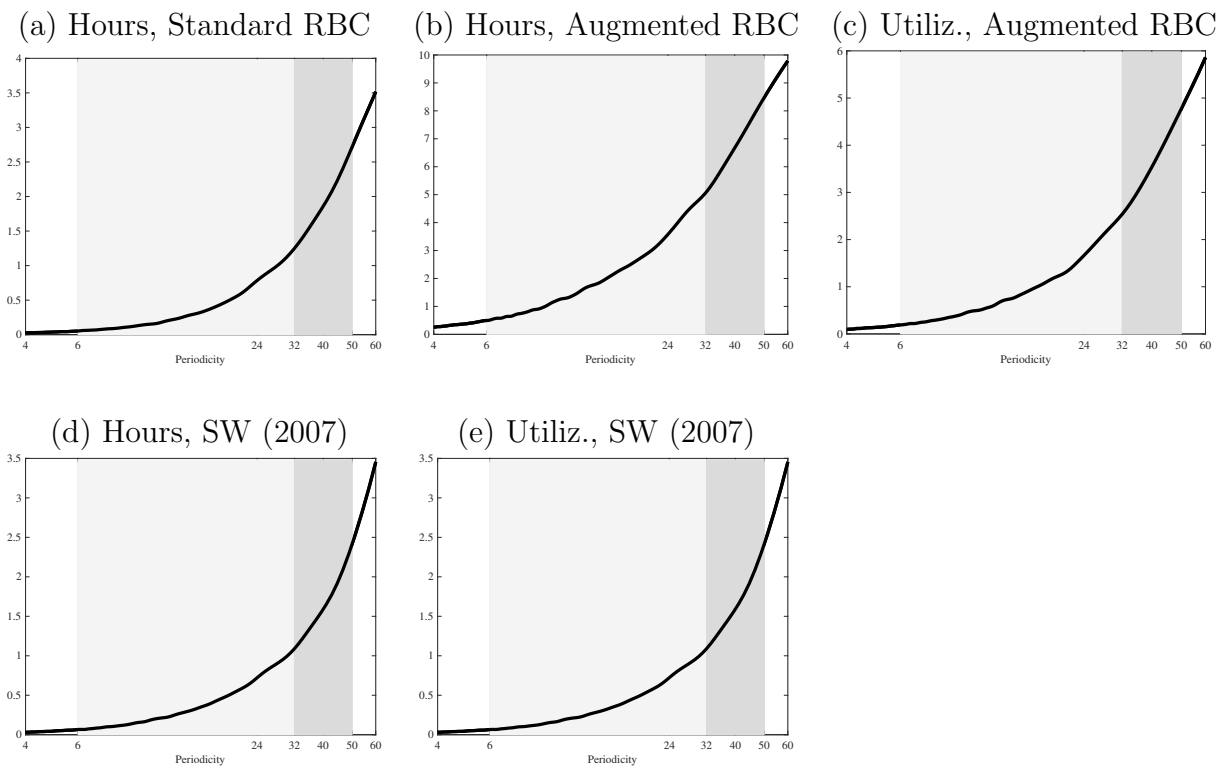
Notes: Panel (a) plots the log of Non-Farm Business Hours divided by Total Population. Panel (b) is an estimate of the spectral density of hours in levels (black line) and for 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (gray lines). A high-pass (P) filter removes all fluctuations of period greater than P . Panel (c) displays high-pass (60) filtered hours.

Figure 2: Spectral Densities For Other Cyclical Variables



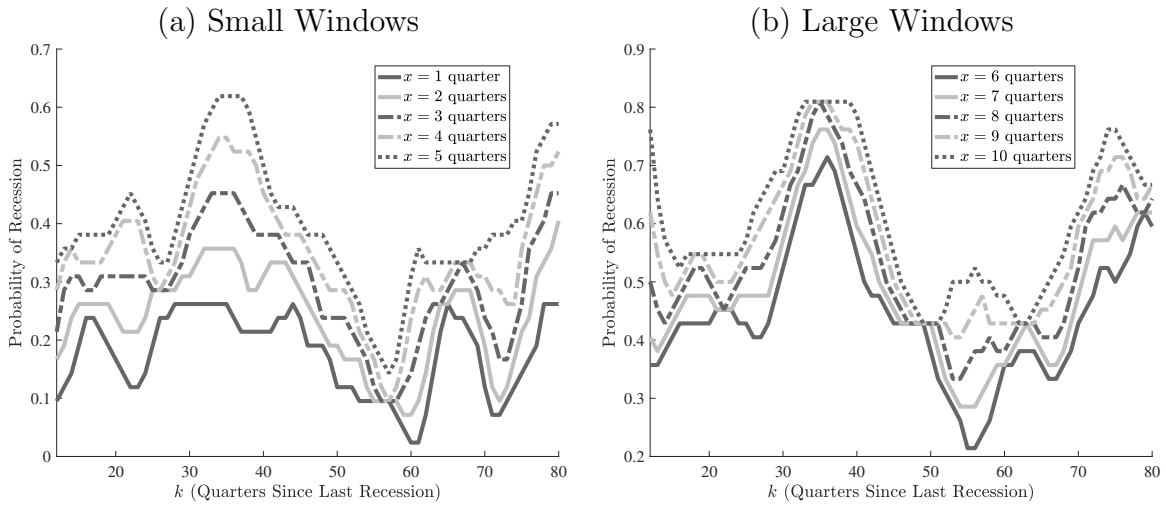
Notes: Panels (a) to (c) show estimates of the spectral density of Total Hours per capita, Unemployment rate and Capacity utilization rate in levels (black line) and for 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (gray lines). A high-pass (P) filter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

Figure 3: Spectral Densities of Hours and Capital Utilization in Standard Models



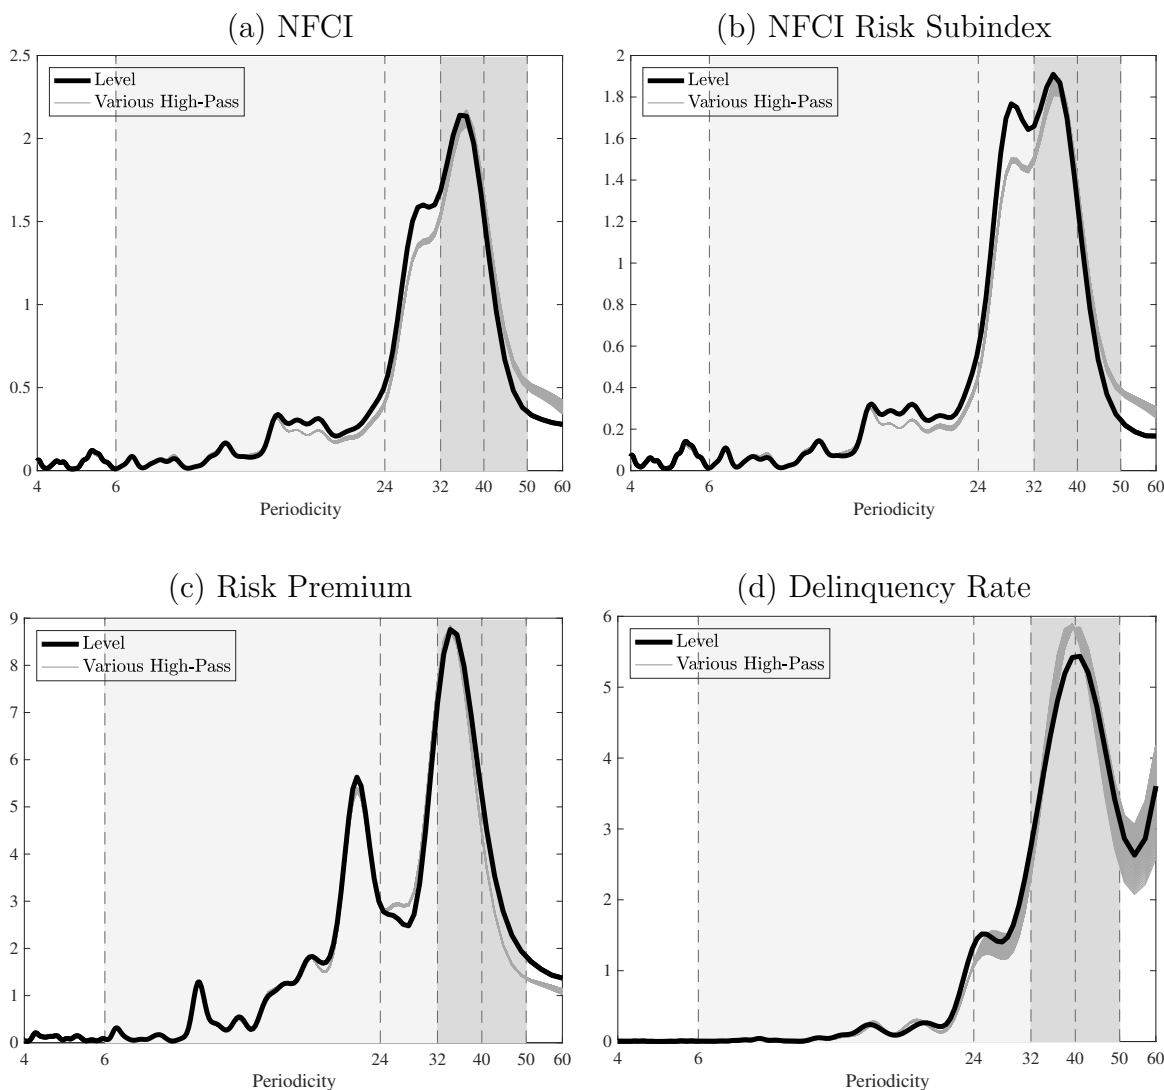
Notes: Each panel display the mean spectral density of hours or capital utilization over 1000 simulations of length 270. Models and parameters values are ? for the standard RBC model, ? for the augmented RBC (with variable capital utilization and investment specific technology shocks) and ?.

Figure 4: Conditional Probability of Being in a Recession




Notes: Each panel 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 1 and 5 quarters in panel (a) and between 6 and 10 quarters in panel (b). Figure was constructed using NBER recession dates over the period 1946:1-2017:2.

Figure 5: Spectral Density of Some Financial Series




Notes: 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 . Series are (a) Chicago Fed National Financial Conditions Index, (b) Chicago Fed National Financial Conditions Risk Subindex, (c) the spread between the federal funds rate and BBA bonds and (d) Delinquency Rate on All Loans, All Commercial Banks. All series end in 2015Q2, and start in 1973Q1 for (a) and (b), 1954Q3 for (c) and 1985Q1 for (d).

Figure 6: Deterministic Simulation of I_t




Notes: Figure shows deterministic simulation of (3)-(4) when $\alpha_3 \rightarrow 0$, $F(I) = \rho I - \xi I^3$, $\alpha_1 = 0.04$, $\alpha_2 = 0.6$, $\delta = 0.05$, $\rho = 0.65$, and $\xi = 0.1$.

Figure 7: Stochastic Limit Cycles



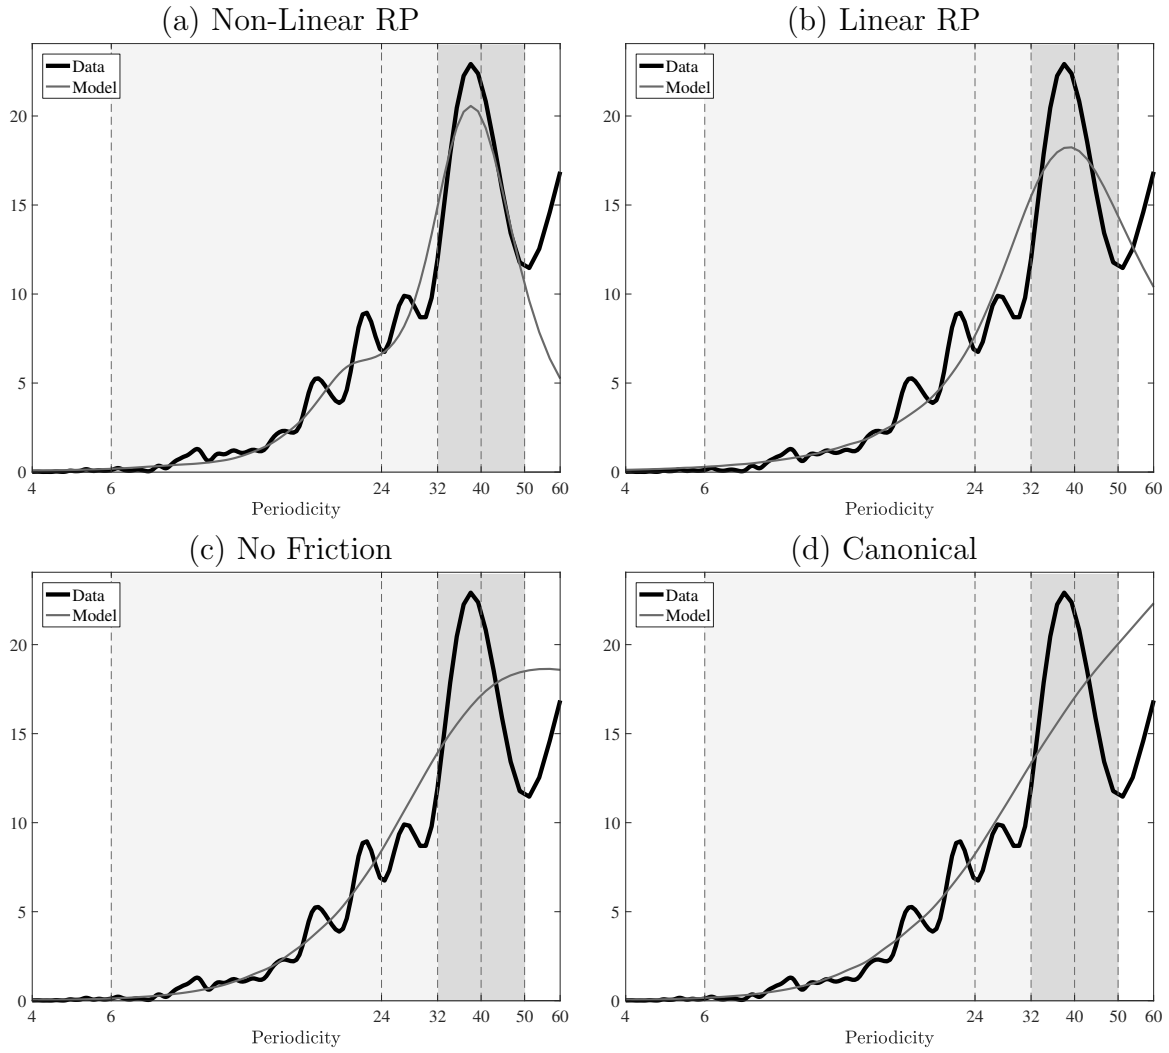
Notes: Light gray lines in both panels of the Figure plot the deterministic limit cycle from Figure 6. The dark gray line in panel (a) is obtained by adding a simulation of the AR(1) stochastic process $\mu_t = \gamma\mu_{t-1} + \sigma\epsilon_t$ to the deterministic limit cycle (i.e., to the light gray line). The dark gray line in panel (b) is obtain by simulating (3) and (6) with the same parameters as in Figure 6 and the same values of the shock as panel (a). The parameters of the AR(1) process are $\gamma = 0.9$ and $\sigma = 0.15$, where ϵ_t is i.i.d. $N(0, 1)$.

Figure 8: Permanent Phase Shift



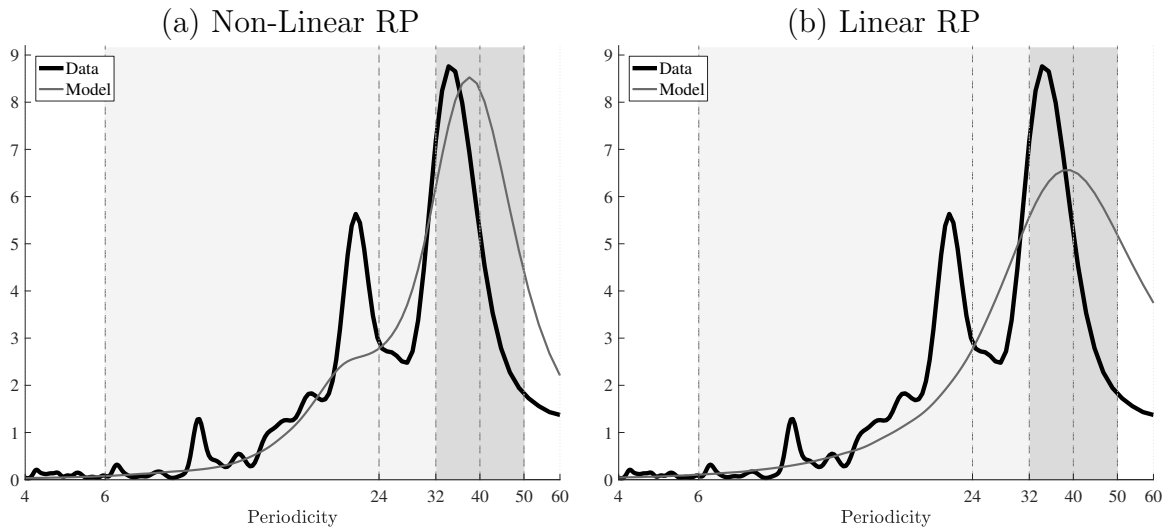
Notes: Figure shows impact, beginning from a point on the deterministic cycle, of a one-time one-standard-deviation shock to μ at date $t+1$ (i.e., we set $\epsilon_{t+1} = \sigma/\sqrt{1-\text{gamma}^2}$, and $\epsilon_\tau = 0$ for all $\tau \neq t+1$). Light gray line shows the path of I in the unperturbed system, while the dark gray line shows the path of the perturbed system. Parameters are the same as those reported in the notes to Figure 7.

Figure 9: Fit of Hours Spectral Density




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

Figure 10: Fit of Risk Premium Spectral Density




Notes: 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 11: Fit of Other Moments (Non-Linear RP)




Notes: 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.

Figure 12: Non-Linear RP Model When Shocks Are Turned Off




Notes: 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 model (gray line).

Figure 13: Sample Path of Hours in the Non-Linear RP Model with Shocks



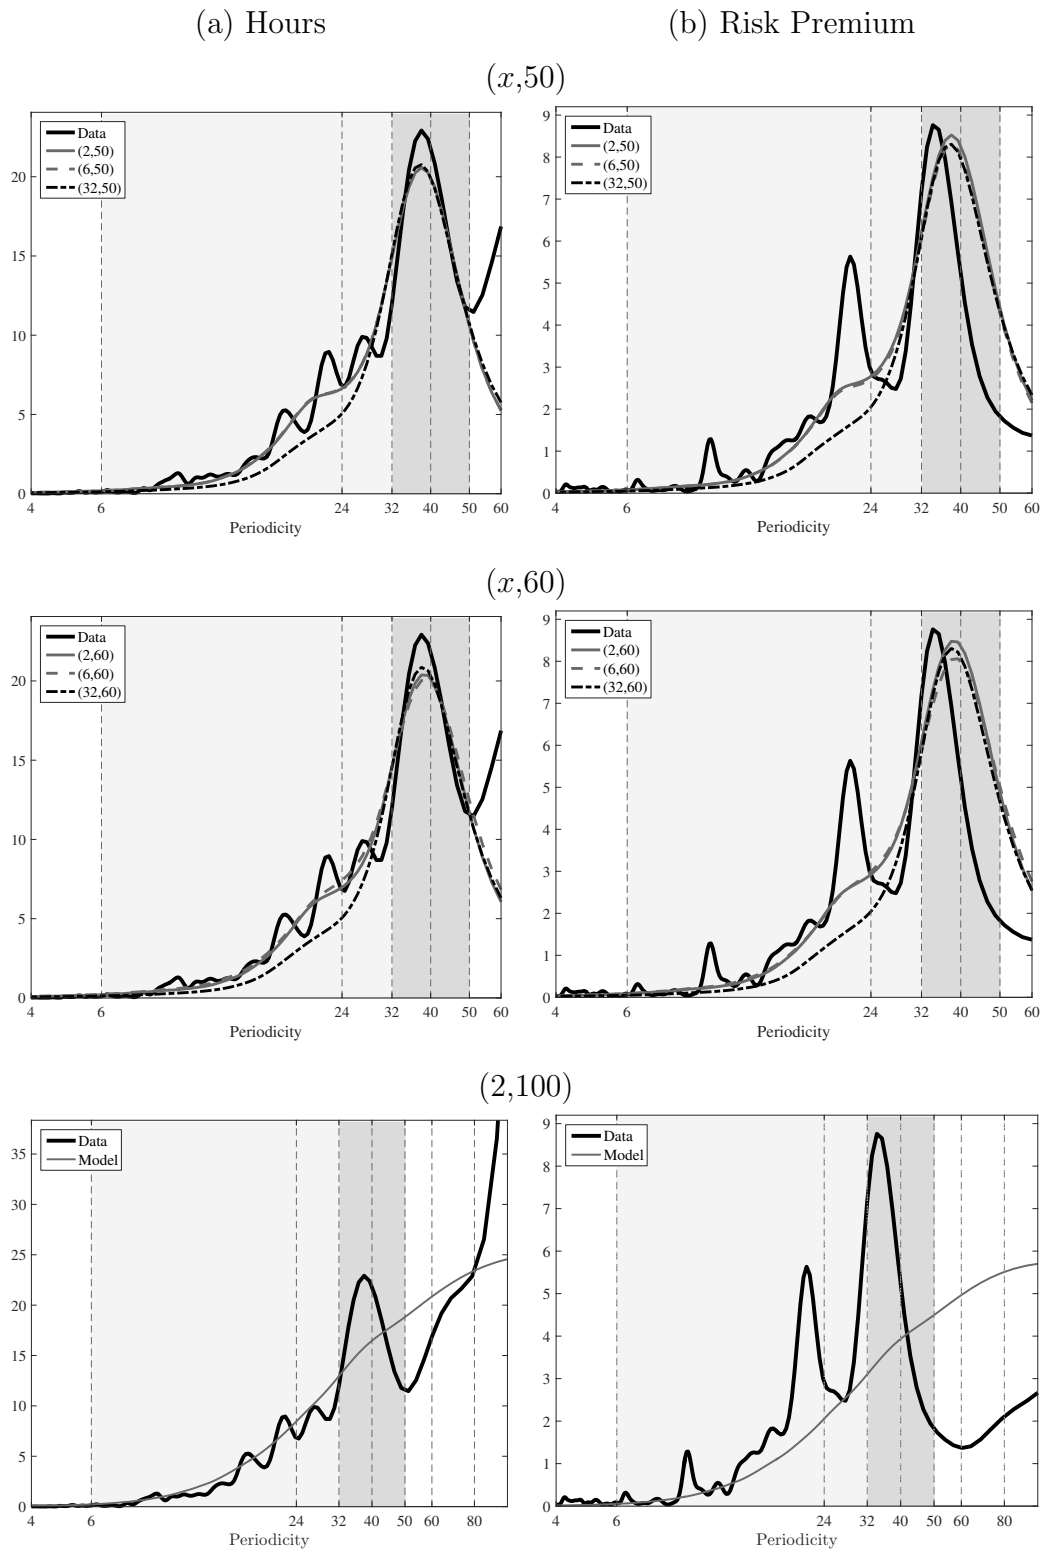
Notes: This figure shows one particular 270-quarter simulation of the estimated non-linear RP model.

Figure 14: Sensitivity of Household Interest Rate to Economic Activity (Non-Linear RP)



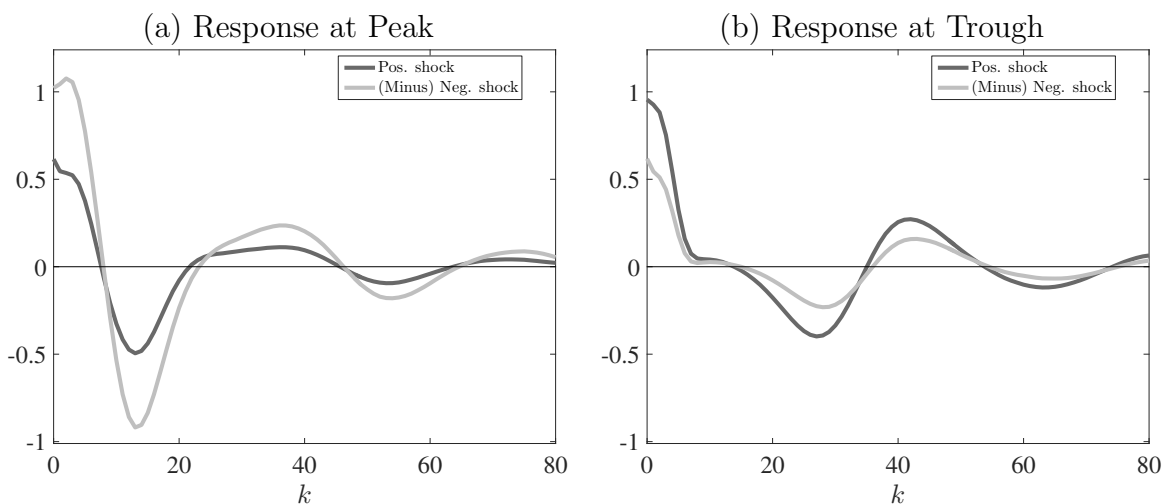
Notes: This figure plots the annualized real interest rate function $4[\phi_e \hat{e} + R^p(\hat{e})]$ for the non-linear RP model, with estimated coefficients and over the range of hours deviations that we observe in the data.

Figure 15: Estimated and Data Spectral Densities for Various Estimations




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

Figure 16: Hours Impulse Response at Cycle Peak (Non-Linear RP)



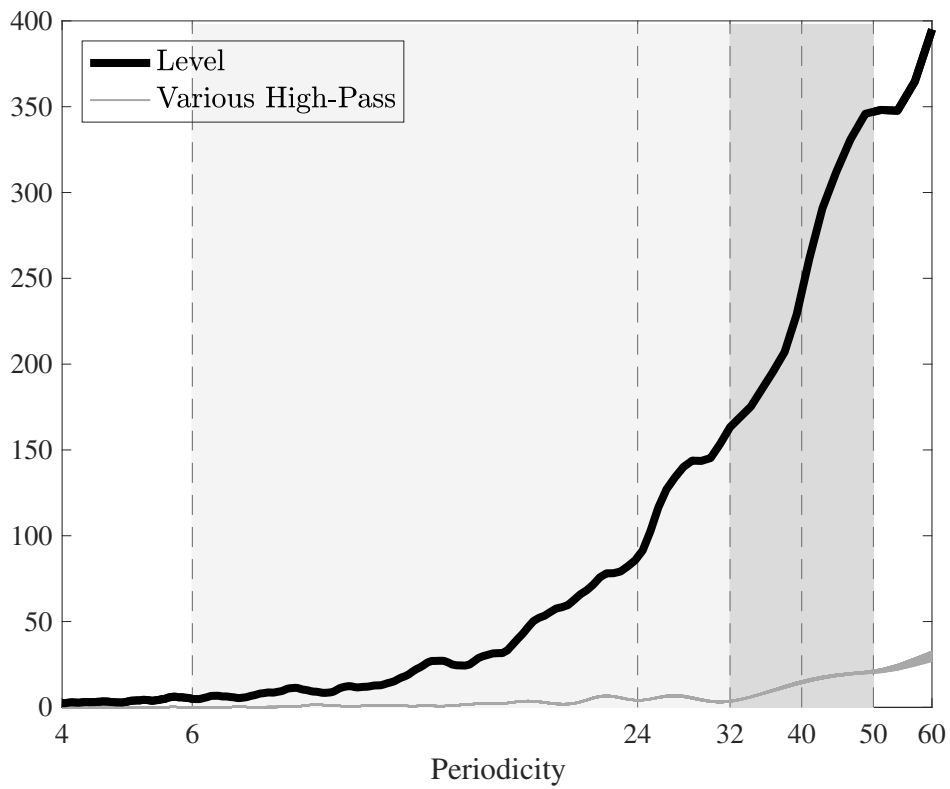
Notes: Figure shows the average response of hours in the non-linear RP model to a one-standard-deviation shock to ϵ_t as at date $t + k$, for $k = 0, \dots, 80$, conditional on initially being at a peak (panel (a)) or trough (panel (b)) of the business cycle. Dark gray line shows response to a positive 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.

Figure 17: Elasticity of s.d.(e) w.r.t. σ (Linear and Non-Linear RP)



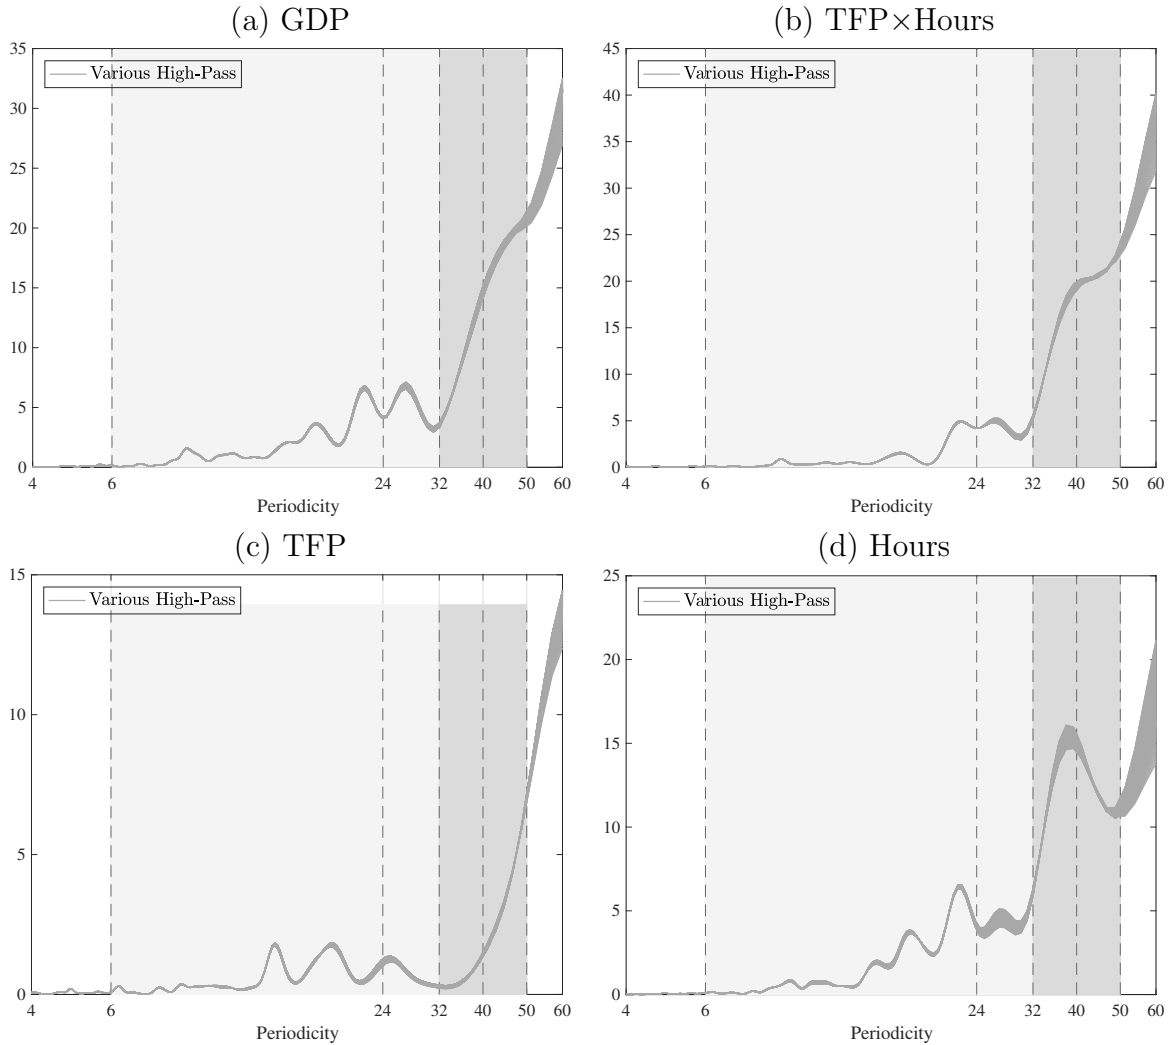
Notes: Figure plots the elasticity of the s.d. of hours with respect to σ for the linear RP (light gray line) and non-linear RP (dark gray line) models. The elasticity for the linear RP model can be shown analytically to equal 1 at all times. To compute the elasticity for the non-linear RP model, for each value of σ we re-solved the model, and then computed the sample s.d. of hours from 10,000,000 simulated periods of data.

Figure 18: Spectral Density of U.S. GDP per Capita, Levels and Various High-Pass Filters



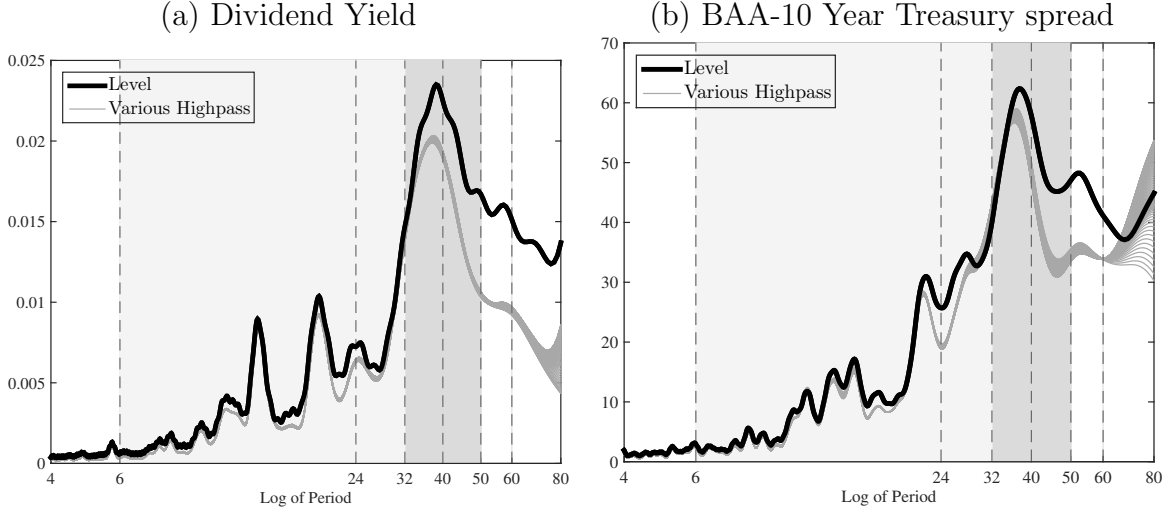
Notes: This figure shows an estimate of the spectral density of U.S. GDP per capita in levels (black line) and for 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (grey lines). A high-pass (P) filter removes all fluctuations of period greater than P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

Figure 19: Decomposing the Spectral Density of GDP



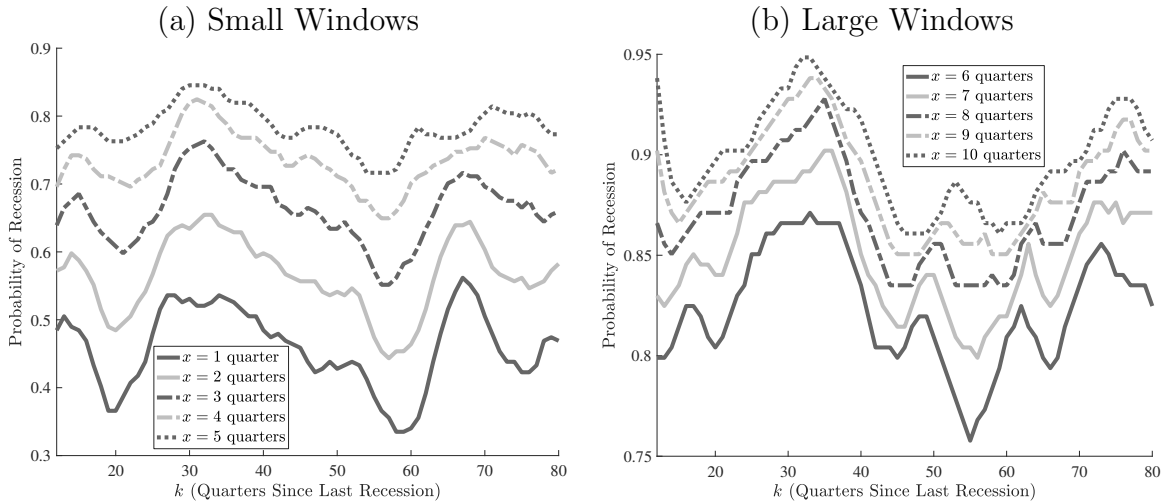
Notes: This figure shows an estimate of the spectral density of U.S. GDP per capita (panel (a)) and Total Factor Productivity (panel (b)) for 101 series that are high-pass (P) filtered version of the levels series, with P between 100 and 200 (grey lines). A high-pass (P) filter removes all fluctuations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

Figure 20: Spectral Densities for Historical Measures of Risk Premia



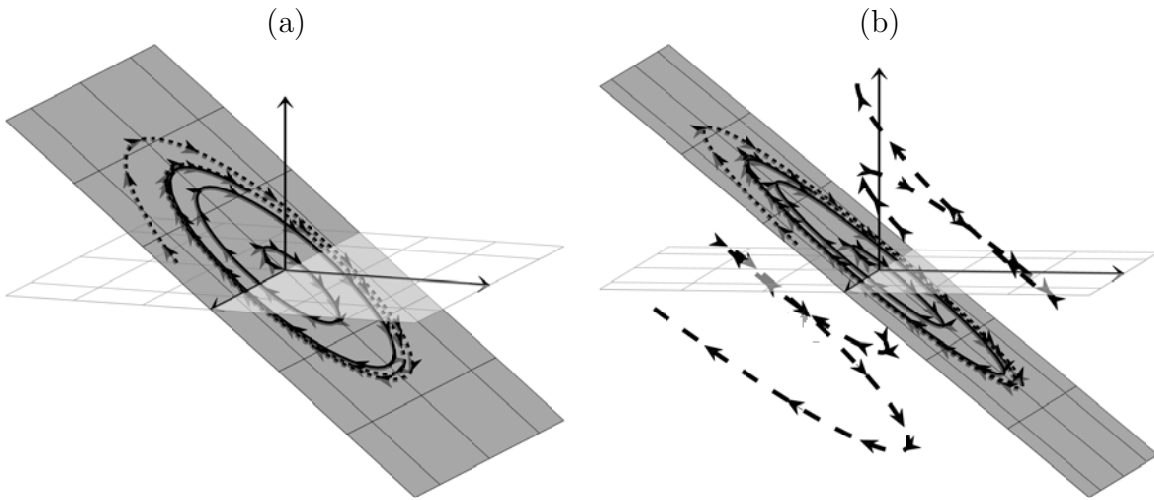
Notes: Panel (a) reports the estimated spectral density of the quarterly dividend yield over the period 1871 to 2016. Panel (b) reports the estimated spectral density for the (quarterly) spread between the interest rate on 10-year BAA bonds and the interest rate on 10-year treasuries over the period 1919 to 2016. Data taken from www.macrohistory.net.

Figure 21: Conditional Probability of Being in a Recession, 1871-2016




Notes: Each panel 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 1 and 5 quarters in panel (a) and between 6 and 10 quarters in panel (b). Figure was constructed using NBER recession dates over the period 1871:1-2017:2.

Figure 22: A Saddle Limit Cycle



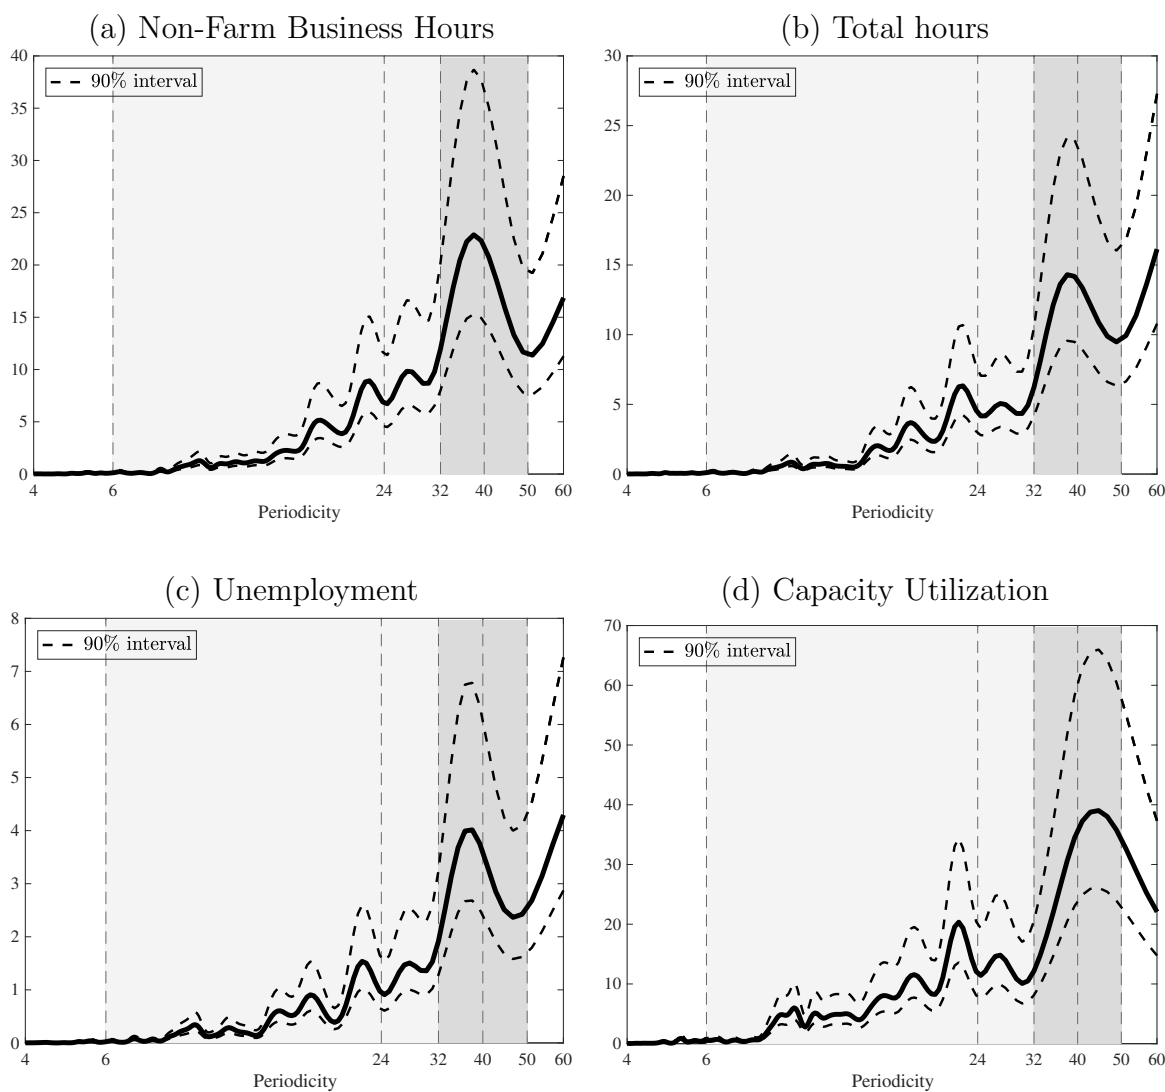
Notes: The gray surface is the non-explosive manifold (which we show here as a plane, but which will in general be non-linear). In panel (a), the black solid and dotted lines are two paths that converge to the limit cycle (which is located on the non-explosive manifold), one from the inside and one from the outside. In panel (b), which shows the same phase space (from a slightly different angle), the dashed black lines are two paths for which the jump variable has not placed the system onto the non-explosive manifold, and which therefore violate the transversality condition.

Figure 23: Unpredictability




Notes: Figure shows unconditional (light gray line and dotted lines) and date- t -conditional (dark gray line and dotted lines) forecasts and 66% confidence intervals for I_{t+k} . Parameters are the same as those reported in the notes to Figure 7. Unconditional and conditional moments obtained from 100,000 simulations of the model. Unconditional forecasts are computed as time averages, while conditional forecasts are computed as ensemble averages.

Figure 24: Spectral Densities For Cyclical Variables



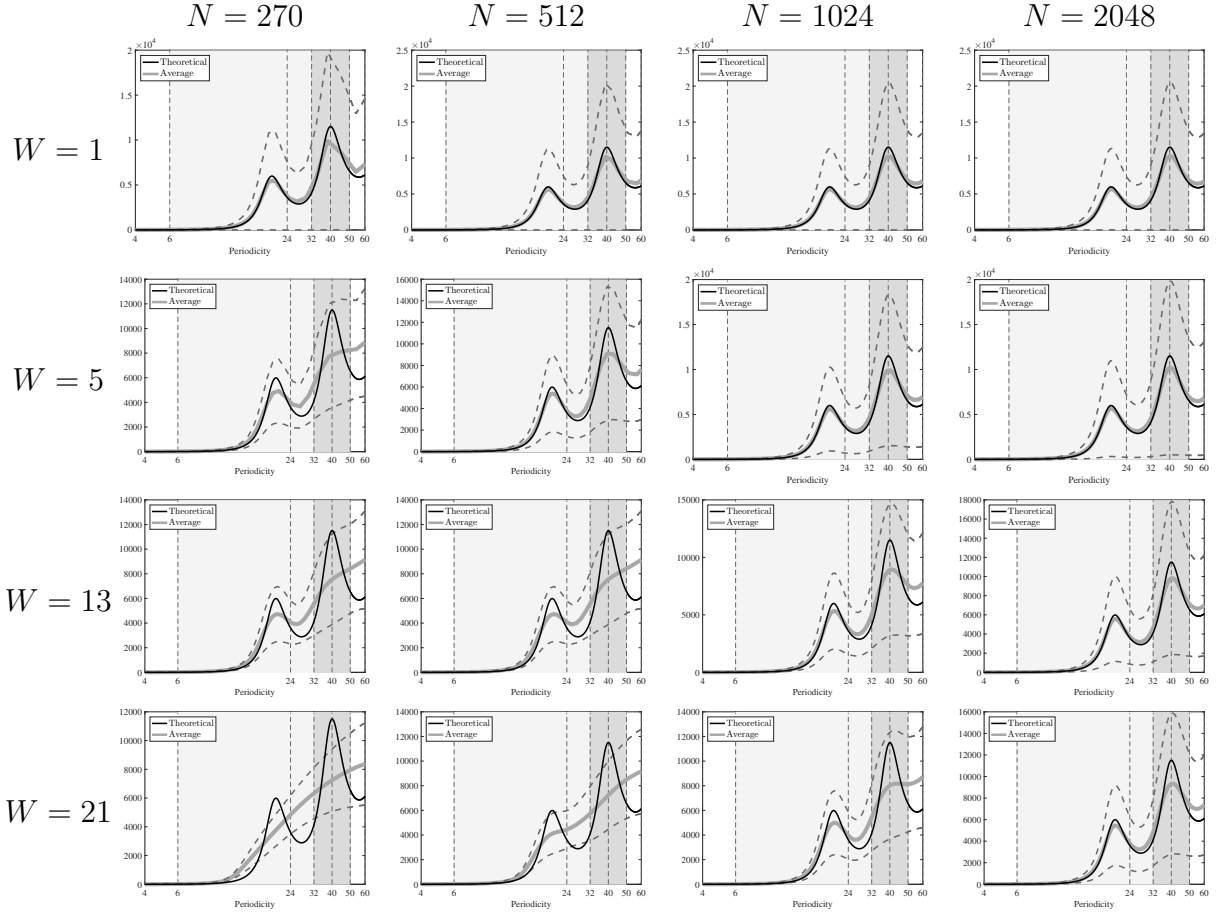
Notes: Panels (a) to (d) show estimates of the spectral density of Non-Farm Business Hours, Total Hours per capita, Unemployment rate and Capacity utilization rate in levels (black line), together with the 5% and 95% asymptotic bounds.

Figure 25: Theoretical Spectral Density (Sum of Three AR(2))



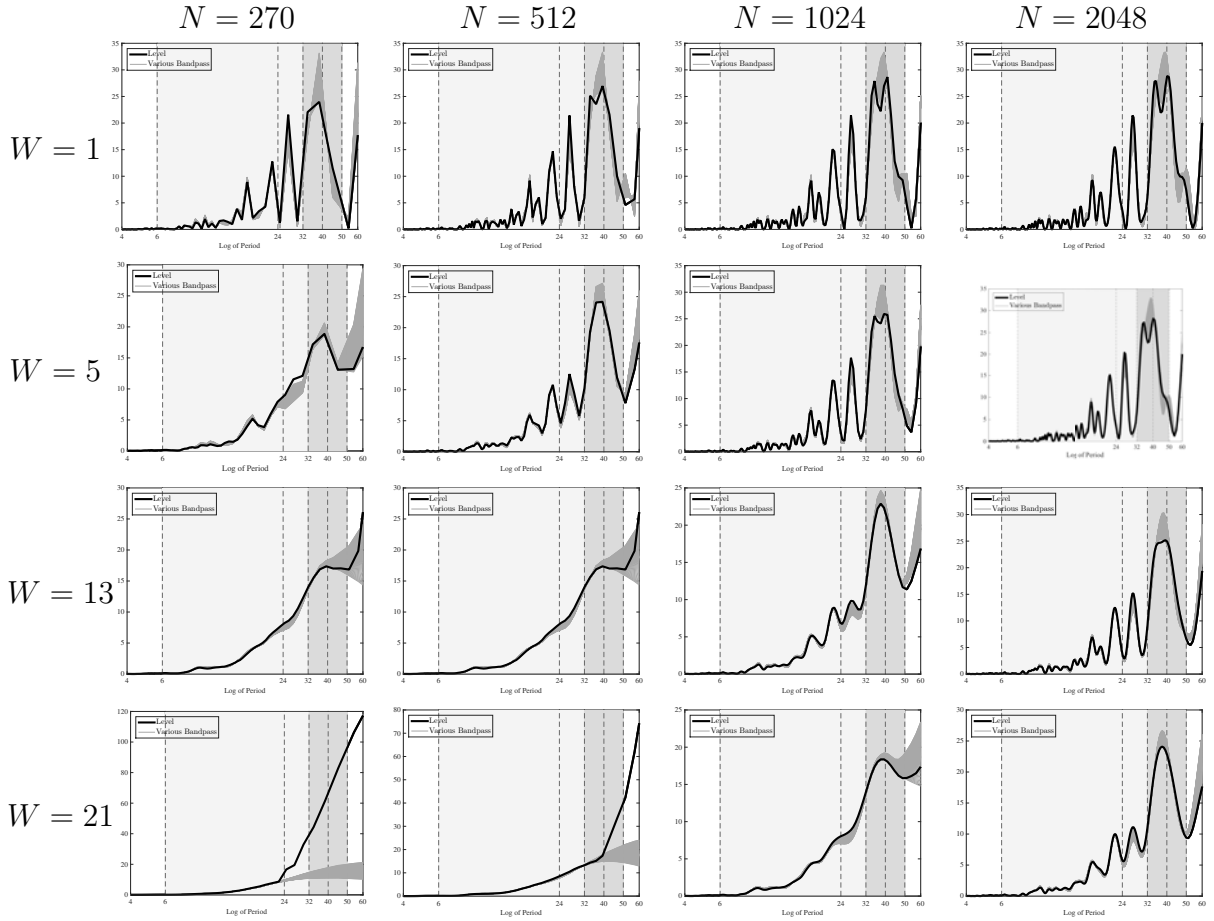
Notes: Figure shows the theoretical spectral density of the sum of three independent AR(2) processes, which have peaks in their spectral densities at, respectively, 20, 40 and 100 quarters.

Figure 26: Effects of Smoothing and Zero-Padding (Sum of Three AR(2))



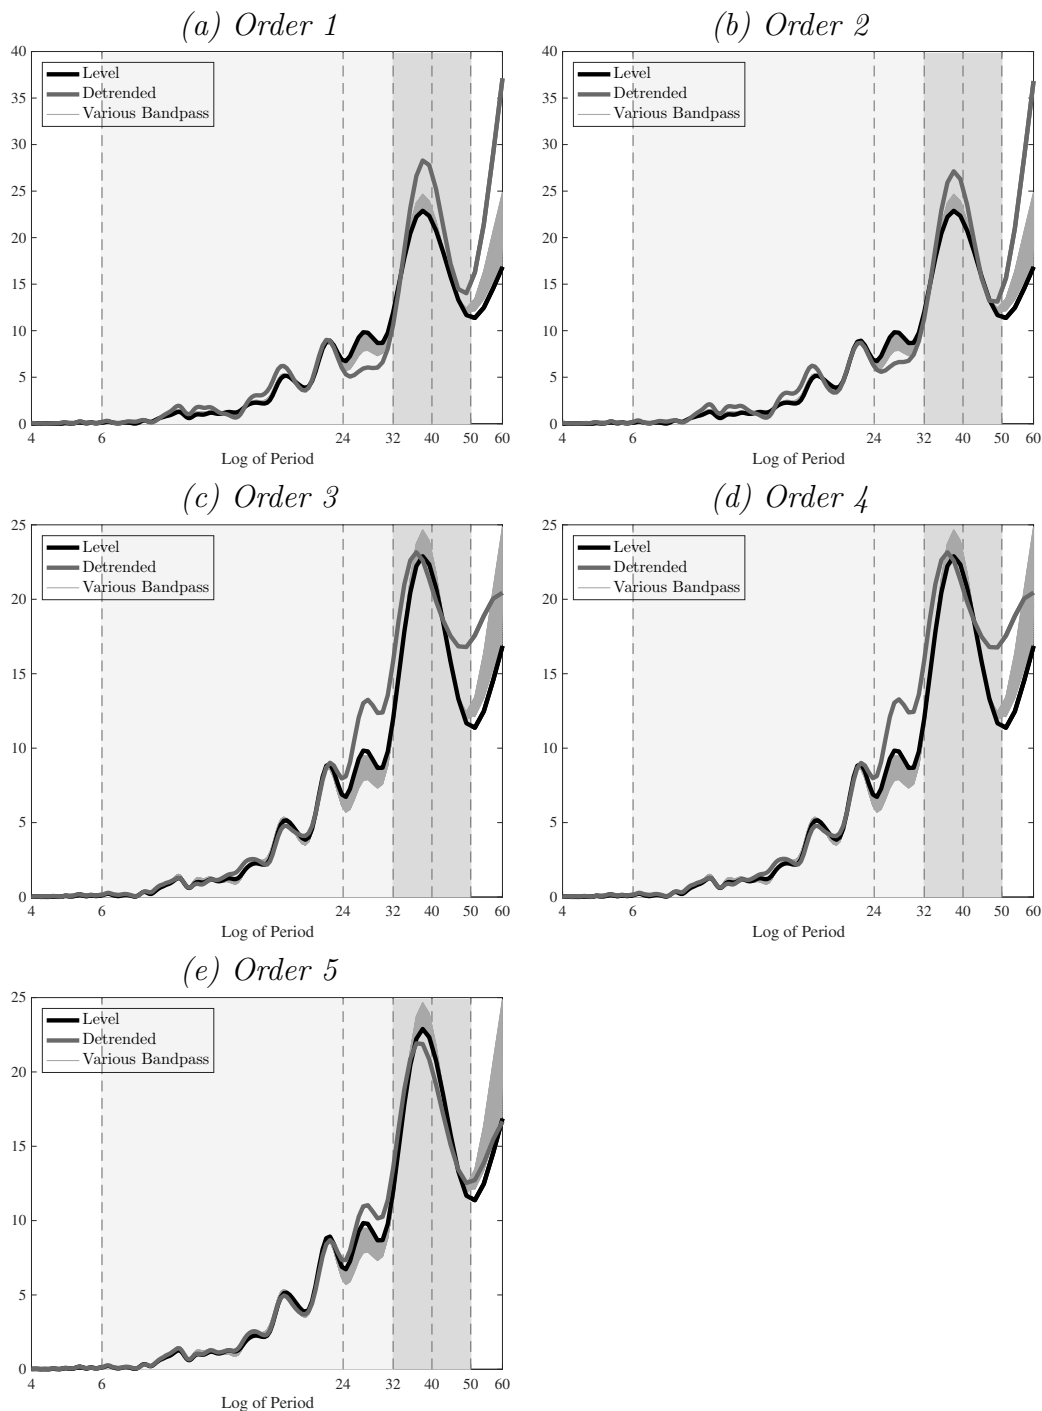
Notes: Figure shows estimates of the spectral density using simulations of the sum of three independent AR(2), which have peaks in their spectral densities at, respectively, 20, 40 and 100 quarters. The black line is the theoretical spectrum, the solid grey line is the average estimated spectrum over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

Figure 27: Changing Smoothing and Zero-Padding



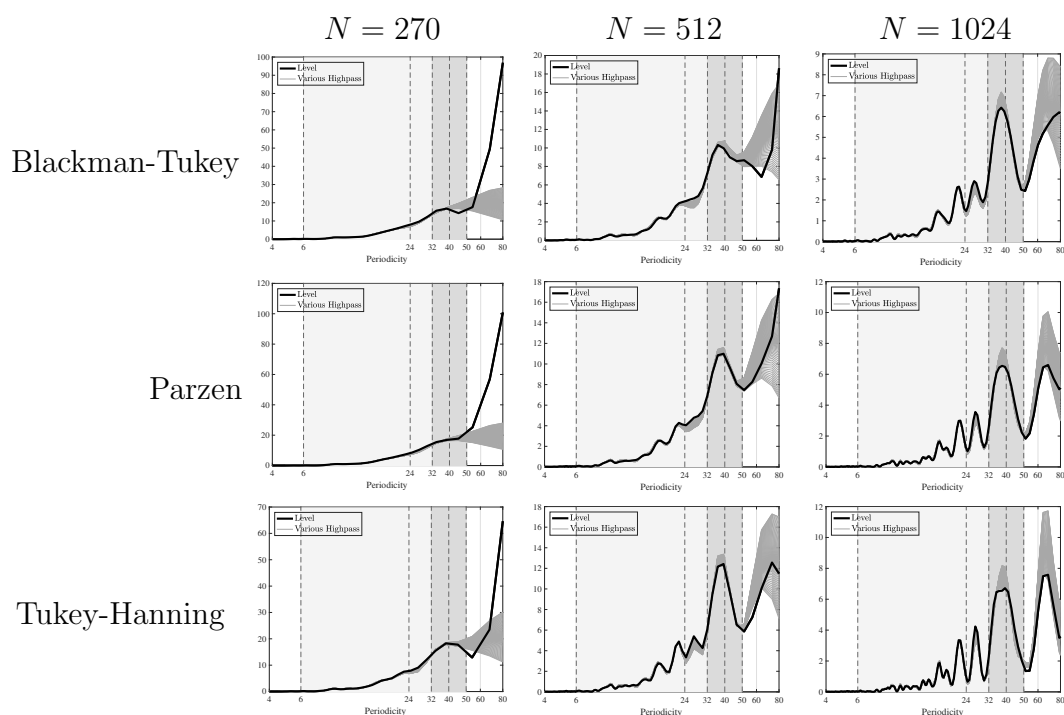
Notes: Figure shows estimates of the spectral density of U.S. Non-Farm Business Hours per Capita over the sample 1947Q1-2015Q2. The different lines correspond to estimates of the spectral density of hours in levels (black line) and of 101 series that are high-pass (P) filtered version of the levels series,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

Figure 28: Using a Polynomial Trend of Various Orders for Benchmark Smoothing ($W = 13$) and Zero-Padding ($N = 1024$)



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

Figure 29: Non-Farm Business Hours with Various Windows and Estimation Using the Covariogram



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

Tables

Table 1: Estimated Parameter Values

	(a)	(b)	(c)	(d)
	Non-Linear RP	Linear RP	No Friction	Canonical
ω CRRA parameter	0.2997	0.2408	0.2408*	0.2408*
γ Habit	0.5335	0.5876	0.8346	0.7635
ψ One minus initial dep.	0.4000	0.2994	0.1154	0*
ϕ_e Taylor rule	0.0421	0.0467	0.0000	0.0000
ϕ Debt backing	0.8668	0.8827	1*	1*
Φ Recovery cost	0.0441	0.0474	0*	0*
R_2^p Risk prem. (2nd order)	0.00017	—	—	—
R_3^p Risk prem. (3rd order)	0.00006	—	—	—
ρ Autocorrelation	-0.0000	0.1387	0.8346	0.7636
σ Innovation s.d.	0.0157	0.0272	0.0335	0.1252
$s.d.(\mu)$ Implied uncond. s.d.	0.0157	0.0275	0.0609	0.1940

Notes: Table displays the estimated parameters of the model for each of the four estimation scenarios. Asterisks indicate calibrated values. Note that, in the no friction and canonical models, the CRRA parameter ω is not separately identified from the Taylor rule parameter ϕ_e and the shock innovation standard deviation σ . To facilitate comparisons, we have fixed ω in these cases to be equal to its estimated value from the linear RP model.

Table 2: Eigenvalues at the Steady State

	Non-Linear RP	Linear RP	No Friction	Canonical
$\lambda_{11}, \lambda_{12}$	$1.1032 \pm$ $0.2164i$	$0.9258 \pm$ $0.1372i$	0.8346, 0.8346	0.7635
$ \lambda_{1j} $	1.1242	0.9359	0.8346	0.7635

Notes: 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.

Table 3: Eigenvalues at the Steady State for Various Estimations

Est. Range	$\lambda_{11}, \lambda_{12}$	$ \lambda_{1j} $
(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

Notes: 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.