Can a Limit-Cycle Model Explain Business Cycle Fluctuations?

Dana Galizia*†

December 31, 2014

Abstract

In conventional models of the business cycle, all fluctuations are ultimately caused by the arrival of random shocks. As a result, individual booms and busts are largely unrelated phenomena. An alternative to this viewpoint is that booms and busts are inherently related, which suggests that fluctuations are at least in part driven by deterministic cyclical forces.

This paper shows (1) how a purely deterministic general-equilibrium model can give rise to limit-cycle fluctuations, and (2) that this model can replicate business cycle features once it includes a small amount of exogenous variation. The deterministic limit cycle arises through a simple micro-founded mechanism in a rational-expectations environment, and does not rely on the existence of multiple equilibria or dynamic indeterminacy.

Since these cycles would indefinitely repeat themselves in the absence of shocks, a TFP shock is introduced in order to create irregularities. The model is estimated to match US hours data, and is shown to be able to match the data closely. The TFP shock in the model is of a reasonable persistence and relatively small size, accounting for only around a fifth of the standard deviation of hours in the model. This highlights that models capable of generating deterministic fluctuations do not require the addition of large, persistent shocks in order to match patterns in the data, which is a common criticism of conventional models.

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

*PhD Candidate, Vancouver School of Economics, The University of British Columbia; dana.galizia@gmail.com.

†I wish to thank Paul Beaudry especially for his invaluable advice and guidance. I would also like to thank Franck Portier, Henry Siu, Yaniv Yedid-Levi, Amartya Lahiri, Francesco Trebbi, Matilde Bombardini, and Jess Benhabib for their helpful comments. All errors are mine alone.

1 Introduction

In conventional models of the business cycle, all fluctuations are ultimately caused by the arrival of random shocks. One implication of this is that a boom is usually caused by one collection of random events, and any subsequent recession is caused by another collection of random events. In this sense, in conventional models individual booms and busts are largely unrelated phenomena.

An alternative to this viewpoint is that booms and busts are inherently related, and that, for example, a protracted boom can sow the seeds for a subsequent recession, which can in turn lead to the next boom, and so on. According to this view, fluctuations are at least in part driven by deterministic cyclical forces that do not fit neatly into the usual shocks-and-propagation-mechanisms characterization of conventional models. While an earlier literature made an attempt to formalize the forces that can produce these deterministic fluctuations,[1] in recent decades this research area has gone dormant.

This paper revisits the idea of deterministic fluctuations, with the aim of showing (1) how a purely deterministic general-equilibrium model featuring strategic complementarity[2] near the steady state can give rise to a stable limit cycle,[3] and (2) that this model can replicate business cycle features once it is augmented to include a small amount of variation from exogenous sources. The limit cycle in the model arises through a simple micro-founded mechanism in a rational-expectations environment. In contrast to most conventional models, the model I present does not require shocks in order to generate fluctuations, nor does it rely on the existence of multiple equilibria or dynamic indeterminacy.[4] Cycles emerge endogenously and would indefinitely repeat themselves in the absence of shocks. Shocks are only introduced into the model to create irregularities in the cycles.

The key deterministic mechanism in the model, which is based closely on Beaudry et al. (2014), will center around the accumulation of a stock of capital, interpreted here as a stock

[1] To fix terminology, I will say a dynamic system exhibits "deterministic fluctuations" if, in the absence of stochastic shocks, its state vector neither diverges to infinity nor converges to a single point. See Appendix C for a more formal definition.

[2] Two agents' decisions are strategic complements if a rise in any agent's choice variable results in a rise in the marginal value of the other agent's choice variable. Their decisions are strategic substitutes if the reverse is true, i.e., if a rise in any agent's choice variable leads to a *fall* in the marginal value of the other agent's choice variable.

[3] A "limit cycle" is a deterministic fluctuation that exactly repeats itself every k periods. A "stable limit cycle" is a limit cycle that acts as an attractor for the system, so that nearby points converge to it over time. See Appendix C for further detail.

[4] Even though they may not require shocks to fundamentals to generate fluctuations, models featuring multiple equilibria or indeterminacy still require some time-varying equilibrium-selection device, and this device is usually taken to be some exogenous shock that affects agents' beliefs in a coordinated way (see, e.g., Farmer and Guo (1994), Jaimovich (2007), Kaplan and Menzio (2014)). Thus, these models can still be seen as requiring shocks in order to generate fluctuations.

of durable goods and/or housing. When this stock is high following a boom period, agents reduce their demand for new capital goods with the goal of running down the stock of capital, which leads to a bust. Because of a demand externality in the model, this bust is excessively large, and can lead new purchases of capital goods to become sufficiently depressed that the stock of capital overshoots its steady-state level. Once this happens, the capital stock is too low and the reverse mechanisms come into play, producing a boom and a subsequent capital stock that is too high, and so on.[5]

The demand externality, meanwhile, arises because of two key imperfections in the model. First, there is a matching friction in the labor market in the spirit of Diamond-Mortensen-Pissarides, which creates the possibility that a household may not find employment when looking for a job. Second, households are unable to perfectly insure against this idiosyncratic unemployment risk. The combination of these two imperfections causes agents to react to a fall in the unemployment rate by increasing their demand for new goods, which in turn reduces the unemployment rate further. If sufficiently strong, this strategic complementarity—where an initial increase in one agent's purchases leads all other agents to increase their own purchases—causes the economy to be very sensitive to small changes in the stock of capital goods. In turn, this causes the unique steady state of the model to be locally unstable, so that even in the absence of shocks the economy will in general fail to converge to the (unique) steady state. When the economy reaches full employment, however, the above demand externality is no longer operative. Instead, a rise in one household's purchases causes the price of goods to increase, which tends to favor a *decrease* in expenditure by other households. As a result of this strategic *substitutability*, when far enough away from the steady state the economy exhibits locally the types of stabilizing forces that are present *globally* in conventional models. This prevents the system from exploding. The combination of this non-explosiveness property with an unstable steady state yields conditions under which a stable limit cycle may appear.

Because the model does not rely exclusively on shocks to drive fluctuations, it is capable of addressing one of the common criticisms of conventional business-cycle models, which is their frequent reliance on poorly motivated and/or implausibly large shocks. For example, the widely cited model of Smets and Wouters (2007) features two "mark-up" shocks that are highly persistent[6] and together account for over half of the 40-quarter forecast-error variance

[5] Because this mechanism operates through a demand channel, to make this channel as clear as possible the stock of capital enters directly into the household utility function in the model without affecting the production side of the economy. It is for this reason that I interpret the stock of capital here as including only durables and housing, and not productive capital. Nonetheless, subject to several technical restrictions the important elements of the model also extend to a productive-capital environment (see Beaudry et al. (2014)).

[6] At the reported median posterior parameter estimates, the price mark-up process has persistence 0.90, while the wage mark-up process has persistence 0.97.

(FEV) of output and hours worked.[7] Outside of their apparent usefulness in helping the model fit moments of the data, however, little empirical justification for the size and persistence of these exogenous shocks is offered.

Mark-up shocks are just one example drawn from a large number of shocks proposed in the literature that may help a particular model fit the data, but that do not necessarily have any clear empirical justification. For example, investment shocks, liquidity shocks, news/noise shocks, risk shocks, financial shocks, government spending shocks, confidence shocks, intertemporal preference shocks, and ambiguity shocks have all been argued in recent years to be quantitatively important drivers of fluctuations, based primarily on their ability to help a particular model fit the data. Without a direct compelling argument for the empirical relevance of these shocks, though, the case can be made that, while the conventional class of models may have had some success in fitting the data, it has perhaps been less successful at *explaining* the data.

The second goal of this paper is therefore to show that the model can match fluctuations in US data with a minimal amount of exogenously-caused variation. With this in mind, I add a single shock—a simple TFP process—to the model, then estimate the model parameters to match as closely as possible the spectrum of hours worked in US data over the period 1960-2012. I focus on the hours spectrum here for two reasons. First, as discussed further in section 2, one of the main criticisms of earlier deterministic-fluctuations models was that they produced cycles that were far too regular. This regularity shows up clearly as a large spike in the spectrum. The ability for the model to match the much flatter spectrum found in the data will thus be an important test of its ability to generate realistic data. Second, as compared with other data series, hours is arguably less likely to be directly impacted by various exogenous shocks. For example, while GDP is directly affected by things like shocks to total factor productivity, over the business cycle we may expect hours to largely respond only indirectly to exogenous shocks. To the extent that this is true, variation in hours is more likely to be caused by the endogenous mechanisms that are the primary focus of this paper. Nonetheless, after estimating the model, I evaluate its ability to fit several other data series, including GDP.

The main quantitative results are as follows. First, I show that the purely deterministic version of the model (i.e., with the TFP shock shut down) is capable of generating cycles in hours of an empirically reasonable length. For example, at the baseline parameterization, the model generates cycles with a roughly 30-quarter period. This contrasts with the earlier

[7] Smets and Wouters (2007) are not the only ones to find that mark-up shocks explain a large portion of variation in a model. For example, Schmitt-Grohé and Uribe (2012) report that a wage mark-up shock explains nearly 70% of the variance of hours growth in their model.

literature on deterministic-fluctuations models, where cycles tended to be either far too short or far too long.

Second, these deterministic cycles are highly regular, as evidenced both by a clear repeating pattern in the simulated path of hours as well as by a large spike in its spectrum. Including the TFP process, however, easily rectifies this problem, resulting in simulated data that appears realistically irregular and a spectrum that closely matches the one estimated from the data.

Third, the estimated TFP process is close to that estimated directly from TFP data, with an unconditional variance that is, if anything, slightly smaller. In many conventional business-cycle models, this TFP process on its own would generate fluctuations that are far smaller than those found in the data. In the model presented here, however, the bulk of the fluctuations come from the deterministic forces, which account for 79% of the standard deviation of hours. Instead, the TFP shock serves primarily to accelerate and decelerate the deterministic cycles at random, causing significant fluctuations in their length while only minimally affecting their amplitude. This result again highlights one of the more general insights of the paper: conventional models usually require large persistent shocks in order to produce the types of large persistent fluctuations found in the data. This is the primary apparent motivation for a large number of the shocks now found throughout the literature. Models capable of generating deterministic fluctuations, however, are able to produce large persistent fluctuations without any shocks whatsoever. Shocks are only required in order to help match the irregularity of the fluctuations found in the data, and this can be accomplished with shocks that are of a much more plausible type and size.

The remainder of the paper proceeds as follows. Section 2 briefly reviews an older deterministic-fluctuations literature and highlights some of its important failures. Section 3 discusses the key features of the data that the model will attempt to match. Section 4 discusses intuitively the basic properties an economic model must have in order to produce a limit cycle. Section 5 introduces the model and discusses its key properties, and includes several theoretical results establishing conditions under which a limit cycle may appear. Section 6 presents the main quantitative results of the paper, while section 7 concludes.

2 Literature: Deterministic fluctuations

Since at least as far back as Kaldor (1940), economists have considered models that are capable of producing deterministic fluctuations. In the 1970s and 1980s in particular, a large literature emerged that examined the conditions under which qualitatively and quantitatively reasonable economic fluctuations might occur in a purely deterministic setting (see, e.g., Benhabib and Nishimura (1979, 1985), Day (1982, 1983), Grandmont (1985), Boldrin and Montrucchio (1986), Day and Shafer (1987); for surveys of the literature, see Boldrin and Woodford (1990) and Scheinkman (1990)). By the early 1990s, however, this literature seemed to have largely gone dormant.

There appear to be several key reasons why interest in deterministic fluctuations may have waned, each of which are addressed in the present paper. First, the earlier literature on deterministic fluctuations can be broadly sub-divided into two categories: models with and without fully-microfounded, forward-looking agents.[8] The latter category, which were generally more capable of producing reasonable deterministic fluctuations than the former, likely fell out of favor as macro in general moved toward more microfounded models.

Second, in the category of models featuring forward-looking agents, the primary focus was on models with a neoclassical, competitive-equilibrium structure.[9] Such models were often found to require relatively extreme parameter values in order to generate deterministic fluctuations. For example, the Turnpike Theorem of Scheinkman (1976) establishes that, under certain basic conditions met by these models, for a sufficiently high discount factor—i.e., for agents that are "forward-looking" enough—the steady state of the model is globally attractive, so that persistent deterministic fluctuations cannot appear.[10] While in principle this does not rule out deterministic fluctuations completely, in practice the size of the discount factor needed to generate them was often implausibly low. For example, in a survey of deterministic-flucutuations models by Boldrin and Woodford (1990), discount factors for several of the models they discuss were on the order of 0.3 or less.[11] As the present paper illustrates, however, if one departs from the assumptions of a neoclassical, competitive-equilibrium environment—for example, if there is a demand externality as in the model presented in section 5—then a discount factor arbitrarily close to one can relatively easily support deterministic fluctuations in equilibrium.

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

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

[10] See the discussion in section 5.3 for further details.

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

Third, as suggested above, models producing periodic cycles—that is, cycles which exactly repeat themselves every k periods—are clearly at odds with the data, where such consistenly regular cycles cannot be found. This can be observed by looking at the spectrum of data generated by such a model, which will generally feature one or more large spikes at frequencies associated with k-period cycles. Spectra estimated on actual data generally lack such spikes,[12] which suggests less regularity in real-world cycles. To address this issue, papers from the earlier literature largely sought to establish conditions under which such irregular cycles could emerge in a purely deterministic setting (i.e., via chaotic dynamics[13]). While in a number of cases this was found to be possible, the conditions appear to have been significantly more restrictive even than those required to generate simple periodic cycles. In constrast, rather than restricting attention to a purely deterministic setting, this paper embeds deterministic (but highly regular) cyclical mechanisms into a stochastic environment for which irregularity emerges naturally.

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

3 Data: Hours worked

As noted above, the focus of this paper is on attempting to explain patterns in hours worked. In this section, I discuss the key properties of the hours data series used as a focal point for both the qualitative and quantitative discussion that follows.

The hours series I use is the quarterly index of nonfarm business hours worked from the US Bureau of Labor Statistics, divided by civilian noninstitutional population obtained from the FRED database. The full sample period is from 1948Q1 to 2014Q1, though I focus on the 1960Q1-2012Q4 subsample.[14] The series was transformed by taking logs and then running the result through a band-pass (BP) filter[15] in order to remove long-run trend components,

[12]See panel (b) of Figure 1 for an example.

[13]Informally, chaotic fluctuations are deterministic fluctuations (see footnote 1) that do not converge to periodic cycles and for which the paths emanating from two different initial points cannot be made arbitrarily close by choosing those initial points sufficiently close together. See, e.g., Glendinning (1994) for a formal definition.

[14]As alluded to below, prior to 1960 the business cycle appears to have been more irregular. The analysis considered here will be concerned with the more regular post-1960 period.

[15]See Baxter and King (1999) and Christiano and Fitzgerald (2003). Code to implement Christiano and Fitzger-

defined here as cyclical components with periods greater than 80 quarters (20 years).[16,17] The filter was applied to the full 1948Q1-2014Q1 sample, after which observations outside of the 1960Q1-2012Q4 sub-sample were discarded.

Panel (a) of Figure 1 plots the resulting series, with NBER-dated recessions indicated by shaded areas. Two things should be noted from the figure. First, it confirms that the BP-filtered hours series exhibits fluctuations that correspond closely to conventional business cycle definitions, as evidenced by the large downward movements during NBER recessions. Second, with the exception of the 1971-1983 period where fluctuations were somewhat more frequent, over the past half-century a full cycle in hours appears to have taken anywhere from 8 to 11 years (32 to 44 quarters) to complete. This pattern, which suggests some degree of regularity, can be confirmed by looking at the spectrum, which is obtained by first orthogonally decomposing the BP-filtered data series into sinusoidal components of different period lengths, then computing the variance of each such component. Panel (b) of Figure 1 plots the result, with the period of the component on the horizontal axis and the associated variance on the vertical one. From the figure, we see that the bulk of the variation in hours occurs at periodicities greater than 24 quarters, with a peak at around 40 quarters. It should be noted that cycles of this length are outside the typical range conventionally associated with business cycles. For example, in their Handbook article, Stock and Watson (1999) define business cycles as fluctuations between 6 and 32 quarters in length. While this range appropriately reflected the length of the business cycle over the broad historical time frame considered in that paper,[18] Figure 1 suggests that over the more recent time frame considered here the business cycle has become longer and, apparently, somewhat more regular.

4 Conditions for a limit cycle

The basic conditions under which a limit cycle may appear are most easily understood by way of example. The example presented here is for a continuous-time bivariate system, but the basic intuition extends to discrete-time systems and to systems with an arbitrary number

ald's (2003) filter in several environments is available at `https://www.frbatlanta.org/cqer/researchcq/bpf/`.

[16] Understanding the mechanisms underyling these low-frequency fluctuations, such as cultural and demographic factors, are beyond the scope of this paper.

[17] Note that none of the high-frequency fluctuations were removed, i.e., a BP(2,80) filter was used. Note also that the choice of this particular filter is not crucial. In section 6, I verify that the key results are robust to a number of alternative filtering choices.

[18] Stock and Watson (1999) report that of the 30 full cycles (peak to peak) identified by the NBER over the period 1858 to 1996, 90% were 32 quarters in length or shorter, with the shortest being six quarters. This observation formed the basis for their definition of the business cycle.

of state variables.[19]

Suppose we have a bivariate system whose evolution is characterized by the differential equations $\dot{z}(t) = g(z(t))$ and the boundary condition $z(0) = z_0$, where (suppressing explicit dependence on t when no confusion will arise) $z \equiv (x, y)'$, $g : \mathbb{R}^2 \to \mathbb{R}^2$ is some function, and $\dot{z}$ indicates the time-derivative of z. It turns out to be more convenient for our purposes to re-cast this system into polar coordinates, expressing the state as $(\varphi, r)'$, where φ is the angle between the vector z and the positive x-axis, and $r \geq 0$ is the magnitude of z. Since it will not play an important role in the basic conditions for a limit cycle, assume for simplicity that φ evolves as

$$\dot{\varphi} = \bar{\theta}$$

for some constant rate of rotation $\bar{\theta} \neq 0$. Next, suppose we have

$$\dot{r} = -r(ar + b)$$

for some parameters a and b, with $ab \neq 0$.[20] Note that, regardless of the values of a and b, this system always has a unique steady state at $r = 0$, i.e., at the point $x = y = 0$. Thus, assuming $r > 0$, the vector z rotates around the origin (in (x, y)-space) at a constant rate, but with a potentially fluctuating length. Depending on the signs of a and b, however, the system will exhibit different qualitative properties.

Conventional economic models generally feature a globally stable steady state. This corresponds here to the case where $\dot{r} < 0$ for all $r > 0$, so that the length of z is always shrinking. This property holds when $a, b \geq 0$. Panel (a) of Figure 2 shows a phase plot of such a case. The individual arrows indicate the forces acting on the length r at that point in the state space. Also plotted are paths tracking the evolution of the system beginning from two different starting points: one close to the steady state and the other far away from it. We can see in this case that there are forces everywhere that tend to push the system inward towards the steady state. As a result, starting from any point the system converges to the steady state.

The case of a globally *unstable* steady state (i.e., $\dot{r} > 0$ for all $r > 0$), meanwhile, corresponds to the case where $a, b \leq 0$, an example of which is plotted in panel (b) of Figure 2. Here we see that forces exist everywhere which push the system *outwards*, away from the

[19] There is a technical issue here that I will sidestep throughout this paper. Certain types of deterministic fluctuations share many of the basic qualitative features of a limit cycle, but never exactly repeat themselves. For example, in a bivariate discrete-time system characterized by rotation around the unit circle by θ radians per period, if θ/π is irrational then the system will never return to the same point twice. I will ignore these uninteresting technicalities and apply the term "limit cycle" loosely, with the understanding that the term may not always strictly apply.

[20] N.B.: The fact that $\dot{\varphi}$ does not depend on r, and vice versa, will not generally hold for an arbitrary g.

steady state, so that if the system begins anywhere but at the steady state it will diverge to infinity.

The conditions under which a limit cycle may appear correspond to the remaining cases, i.e., where $ab < 0$. When this is true there now exists, in addition to the steady-state value $r = 0$, a second (non-negative) solution to $\dot{r} = 0$, given by

$$r = \hat{r} \equiv -\frac{b}{a} > 0$$

When the system begins with $r = \hat{r}$, it will rotate in the plane at the constant rate $\bar{\theta}$, but move neither toward nor away from the steady state. The set of points z such that $\|z\| = \hat{r}$ is the limit cycle of this system.

Of particular interest is the case where the limit cycle is stable, i.e., where neighboring points will tend to converge to the limit cycle over time. This will occur if (1) when the system begins *inside* the limit cycle, forces tend to push it *outwards*; and (2) when the system begins *outside* the limit cycle, forces tend to push it *inwards*. This corresponds to the case where $a > 0$ and $b < 0$, so that when r is small $\dot{r} \approx -br > 0$, while for r large $\dot{r} \approx -ar^2 < 0$. An example of this is shown in panel (c) of Figure 2. Here we see that, regardless of where the system begins (as long as it is not exactly at the steady state) it will converge to the limit cycle.

For the sake of completeness, the final case where $a < 0$ and $b > 0$ is shown in panel (d) of Figure 2. This corresponds to an unstable limit cycle: if the system begins on the limit cycle (shown as the dashed circle in the figure), it will remain there forever, but if it begins off the limit cycle, it will either converge to the steady state (if it begins inside the limit cycle) or diverge to infinity (if it begins outside the limit cycle).

The preceding analysis highlights the two key properties needed to obtain a stable limit cycle in a general setting: (1) when the system is close to the steady state, there are forces which tend to push it *away* from the steady state (i.e., the steady state is unstable); and (2) when the system is far away from the steady state, there are forces which tend to push it *towards* the steady state (i.e., the system is non-explosive). A limit cycle then emerges as the set of points where these outward and inward forces precisely balance. A natural question to ask, then, is whether there are reasonable conditions under which a dynamic economy may exhibit these two key properties. It turns out that one possible set of such economic conditions is as follows. First, suppose that in a neighborhood of the steady state, individual agents' actions are *strategic complements*. As is well known, strategic complementarity often leads to situations where small changes in state variables can lead to relatively large changes in equilibrium outcomes, which is precisely the state of affairs needed for a system to be (locally) unstable. Second, suppose that when far enough away from the steady state, individual agents' actions act as *strategic substitutes*. When this is the case, small changes in state variables tend to produce *small* changes in equilibrium outcomes, and as a result the system tends to drift back towards the steady state. In the following section, an economic model is presented which possesses these two features and which, as a result, will have the potential to generate limit cycles.

5 The unemployment-risk model

In this section I present a simple economic model that is capable of generating limit cycles. In the model, which is based closely on Beaudry et al. (2014), households begin each period with a stock of durable goods and must decide how many additional goods to purchase in the goods market. Abstracting for the moment from forward-looking behavior, there are two key static factors that affect this decision. First, household demand is decreasing in the size of the current stock of durables: when their existing stock of durables is low, households want to purchase more, and vice versa. Second, because of a self-insurance motive, household demand is decreasing in the unemployment rate. There are two imperfections in the model that cause this self-insurance behavior to emerge. First, there is a matching friction in the spirit of Diamond-Mortensen-Pissarides, which creates the possibility that a household may not find employment when looking for a job. Second, households are unable to perfectly insure against this idiosyncratic unemployment risk. The upshot is that an increase in the unemployment rate causes them to reduce their demand for new goods.

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

5.1 Static version

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

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

Preferences for the first sub-period are represented by

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

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

There is a large set of potential consumption goods firms in the economy who can decide to search for workers in view of supplying goods to the market. Each firm can hire one worker and has access to a decreasing-returns-to-scale production function $F(\ell)$, where ℓ is the number of hours worked for the firm.[21] Production also requires a fixed cost k in terms of the output good, so that the net production of a firm hiring ℓ hours of labor is $F(\ell) - k$. Firms search for workers and, upon finding a worker, they jointly decide on the number of hours worked and on the wage to be paid. The fixed cost k is paid before firms can look for workers. Upon a match, the determination of the wage and hours worked within a firm is done efficiently though a competitive bargaining process,[22] so that in equilibrium $pF'(\ell) = w$, where w is the wage, expressed in terms of household services.[23]

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

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

where the left-hand side is total purchases of consumption goods and the right-hand side is the total available supply after subtracting firms' fixed costs. Firms enter the market up to the point where expected profits are zero. This condition can be written as[24]

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

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

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

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

[23] As discussed in Beaudry et al. (2014), the assumption of a competitive bargaining process is for simplicity. The main mechanisms are robust to alternative bargaining protocols.

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

V, is given by[25]

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

where $v, \tau > 0$ are parameters. This function is piecewise linear and concave, with a kink at $a = 0$.[26] Here, the marginal value of assets is given by v when assets are positive and $(1+\tau)\, v$ when assets are negative. Since buyers in general face unemployment risk when making their purchase decisions, the wedge between the marginal value of assets when in deficit and that when in surplus generates self-insurance behavior, whereby a fall in the employment rate causes buyers to reduce their purchases out of increased concern that they will end up in the costly unemployment state. This mechanism is central to the strategic complementarity that emerges in the model, which in turn is what will allow the dynamic version of the model to generate limit-cycle behavior. The strength of this mechanism, meanwhile, is governed by the parameter τ. Given the above value function V, the buyer's problem is to choose e to maximize

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

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

5.1.1 Equilibrium

Letting e_j denote purchases by household j and e the average level of purchases in the economy, one may show that household j's optimal consumption-choice decision is characterized by[27]

$$U'(X + e_j) = p(e)\, v\left[1 + \tau - \tau\phi(e)\right] \tag{1}$$

where $p(\cdot)$ and $\phi(\cdot)$ are the price of consumption goods and the employment rate, respectively, expressed as functions of aggregate purchases. The left-hand side of (1) is simply household j's marginal utility of consumption. The right-hand side, meanwhile, captures buyer j's expected marginal-utility cost of funds. When the economy is at full employment ($\phi(e) = 1$), this is simply equal to the price $p(e)$ of consumption goods in terms of household services, times the marginal value v of those services when assets are non-negative. When there is unemployment, however, the buyer faces some positive probability of ending up in the negative-asset state, which is associated with a higher marginal value of assets (i.e., $(1+\tau)\, v$). As a result, the

[25] See Section 2.2 in Beaudry et al. (2014) for a discussion of how to microfound such a value function.

[26] As noted in Beaudry et al. (2014), what matters here is that the marginal value of assets be smaller in surplus than in deficit. The piecewise linearity property is assumed only for tractability.

[27] See Beaudry et al. (2014).

expected marginal-utility cost of funds is higher and, all else equal, household j would choose a lower level of purchases.

An equilibrium for this economy is given by a solution to (1) with the additional restriction that $e_j = e$. To understand how the equilibrium is affected by shifts in X, note the following properties of the equilibrium functions $p(\cdot)$ and $\phi(\cdot)$. First, one may show that $\phi(e) = \min\{e/e^*, 1\}$, where e^* is the output (net of fixed costs) produced per firm when there is a positive level of unemployment.[28] Second, one may show that $p(\cdot)$ is a continuous function of e, with $p'(e) = 0$ for $e < e^*$, and $p'(e) > 0$ for $e > e^*$.[29] The consequences of these two properties for the marginal-utility cost of funds (i.e., the right-hand side of (1)) are illustrated by the curve labelled "cost of funds" in panel (a) of Figure 3. For e sufficiently small, the curve is downward-sloping: as e rises, output is increased along the extensive labor margin, lowering the unemployment rate and making purchases feel less expensive to households. Once e reaches the full-employment level e^*, however, additional increases in output come via the *intensive* labor margin, which is associated with a rising price and thus an increased cost of funds.

The two regimes—unemployment and full employment—are associated with different equilibrium responses to a rise in the endowment X.[30] Panel (a) of Figure 3 shows the case for the unemployment regime. The economy is initially in equilibrium at the level e_1 of purchases, which occurs at the intersection of the cost of funds curve and the solid marginal-utility function $U'(X + e)$. A rise in the endowment by ΔX then shifts this marginal-utility function to the left by ΔX units, as represented by the dashed curve in the figure. We see that the equilibrium level of purchases falls as a result of the rise in X, and furthermore that it falls by *more* than ΔX (so that total consumption $c = X + e$ falls). This amplified response is due to the strategic complementarity that exists in the unemployment regime: a rise in the endowment causes households to reduce their demand for new goods which, via an extensive labor margin adjustment, lowers the employment rate ϕ, which in turn raises the cost of funds, causing households to reduce purchases further, further lowering the employment rate, etc.

In contrast, panel (b) of Figure 3 shows the same experiment but beginning from the

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

[29] Combining the household's labor supply condition and the firm's labor demand condition, one may obtain $p = \nu'(\ell) / [vF'(\ell)]$. As pointed out in footnote 28, when $e < e^*$ we have $\ell = \ell^*$, so that $p = p^* \equiv \nu'(\ell^*) / [vF'(\ell^*)]$. Further, once the economy achieves full employment, a rise in output must come through the intensive margin of labor (i.e., through a rise in ℓ), which causes $p(\cdot)$ to be increasing in e on $e > e^*$.

[30] As shown in Beaudry et al. (2014), if τ is sufficiently large there may be more than one equilibrium. While this is an interesting theoretical possibility, the evidence obtained from the quantitative exercise of section 6, though not conclusive, gives no indication that multiple equilibria are of concern. I therefore restrict attention throughout this paper to the case where the equilibrium is unique, i.e., where τ is not too large.

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

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

5.2 Baseline dynamic model

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

$$X_{t+1} = (1-\delta)\left(X_t + \gamma e_t\right) \tag{2}$$

where X_t is the stock of durables brought into period t and e_t is quantity of consumption-goods purchases in period t. For simplicity, I assume that a constant fraction $\gamma \in (0,1]$ of these purchases are durable.[31] $\delta \in (0,1]$ is the depreciation rate.

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

$$\sum_{t=0}^{\infty} \beta^t \left\{ U\left(c_t\right) + \phi_t \left[-\nu\left(\ell_t\right) + v\left(w_t \ell_t - p_t e_t\right)\right] - \left(1-\phi_t\right)(1+\tau)\, v p_t e_t \right\} \tag{3}$$

[31] In the quantitative exercise below, I will interpret "durables" as including both conventional durable goods as well as residential investment, which is conceptually similar.

subject to $c_t = X_t + e_t$ and the accumulation equation (2), and taking ℓ_t as given.[32]

I assume that a steady state for this economy exists and is unique. As is the case in the static model, it can be verified that this is true as long as τ is not too large.[33] I further assume that this steady state satisfies $\ell < \bar{\ell}$, so that the household's time constraint is not binding at the steady state.

5.3 Limit cycles in the dynamic model

5.3.1 The myopic case

Conditions under which limit cycles may appear in this model can be understood most easily in the myopic case where $\beta = 0$. In this case, we simply have a repeated sequence of the static economy discussed in section 5.1, with the only linkage between them being the inherited stock of durable goods. We may characterize the equilibrium evolution of the stock of durables over time as

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

where $e\left(X_t\right)$ expresses the equilibrium level of purchases at date t as a function of the only state variable, X_t. This equilibrium is determined entirely as it was in Figure 3, with the unemployment regime characterized by strategic complementarity and the full-employment regime by strategic substitutability.

Recall the two basic conditions discussed in section 4 which are required to generate a stable limit cycle: (1) a locally unstable steady state, and (2) global non-explosiveness. Letting $\bar{X}$ denote the steady state level of durables, these two conditions correspond mathematically to (1) $\left|g'\left(\bar{X}\right)\right| > 1$, and (2) $\left|g'\left(X\right)\right| < 1$ for $\left|X - \bar{X}\right|$ sufficiently large, where

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

It is straightforward to verify that the second condition necessarily holds here, as follows. Suppose X is sufficiently small so that the economy is in the full-employment regime. As was shown earlier, a rise in X in this regime is associated with a fall in e, but by less than the rise

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

[33] See Beaudry et al. (2014). See also the comments in footnote 30, which apply equally here.

in X, i.e., $-1 < e'(X) < 0$. Thus, $(1-\delta)(1-\gamma) < g'(X) < 1-\delta$, and therefore $|g'(X)| < 1$ clearly holds.[34] Suppose instead that X is very large. In this case it can be verified that the non-negativity constraint $e \geq 0$ binds, so that $e'(X) = 0$ and therefore $g'(X) = 1-\delta$, and thus again $|g'(X)| < 1$ holds.

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

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

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

5.3.2 The general case

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

In particular, consider a general deterministic dynamic economy with date-t state vector $z_t \in \mathbb{R}^n$. Let $\mathcal{W}(z_t, z_{t+1})$ denote the period-t return function when the current state is z_t and the subsequent period's state is z_{t+1}.[35] The following theorem characterizes the solution to the problem of maximizing lifetime utility $\sum \beta^t \mathcal{W}(z_t, z_{t+1})$, where β is the discount factor.

Turnpike Theorem. (Scheinkman (1976)) *If $\mathcal{W}$ is concave, then there exists a $\bar{\beta} < 1$ such that if $\bar{\beta} \leq \beta \leq 1$ then the steady state is unique and globally stable.*[36]

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

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

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

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

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

Proposition 1. *Given any initial endowment of durables X_0, $\limsup_{t\to\infty} |X_t| < \infty$.*

Proof. All proofs in Appendix A. □

Proposition 1 ensures that in the limit the system either exhibits deterministic fluctuations (such as a limit cycle) or converges to a fixed point. The following proposition establishes that, in contrast to models for which the Turnpike Theorem applies, local instability is possible in this model for an arbitrarily high discount factor.

Proposition 2. *There exists parameter values and functional forms such that, for some $\bar{\beta} < 1$, if $\bar{\beta} \leq \beta < 1$ then the (unique) steady state is locally unstable.*

In combination with Proposition 1, Proposition 2 confirms that there are parameter values and functional forms for which the model will generate deterministic fluctuations even if β is arbitrarily close to one. The reasons for the failure of the Turnpike Theorem to hold for this model can be clarified as follows. Suppose the steady state of the model is in the unemployment regime, and let $\mathcal{W}(X_t, X_{t+1})$ be a period-t return function such that the solution to the problem

$$\max_{\{X_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \mathcal{W}(X_t, X_{t+1}) \tag{4}$$

implements the equilibrium of the model in a neighborhood of this steady state.[37] If it turns out that $\mathcal{W}$ is concave, then the Turnpike Theorem implies that the model cannot generate limit-cycle dynamics. The following proposition establishes that in fact $\mathcal{W}$ may not be concave.

[37] An example of such a $\mathcal{W}$ is found in the proof of Proposition 3.

Proposition 3. *There exists parameter values and functional forms such that, in the neighborhood of an unemployment-regime steady state, $\mathcal{W}$ is not concave.*

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

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

Proposition 4. *For τ sufficiently close to zero, the steady state is stable.*

Proposition 4 thus confirms that, if τ is not sufficiently large, the degree of strategic complementarity is too small to produce an unstable steady state.

6 Quantitative Exercise

This section presents the main quantitative results of the paper. I estimate a version of the dynamic model discussed above, with the primary goal of establishing that it is capable of matching the key quantitative features of the hours data discussed in section 3.

The baseline dynamic model presented in section 5 was constructed with an eye toward analytical tractability. As a result, that model lacks many of the features which are known to be helpful in quantitatively matching the data, and includes several others which, while not central to the key mechanisms, turn out to be restrictive in a quantitative setting. Since the main purpose of the exercise in this section is quantitative in nature, I make several adjustments to the model designed to help it in that regard.

First, as is well known, dynamic systems with a single state variable have considerable difficulty in producing deterministic fluctuations with the basic qualitative properties that we observe in macroeconomic aggregates. In particular, deterministic fluctuations in such models tend to be erratic, with the system often jumping back and forth from one side of the steady state to the other every few periods or less. Thus, if the unemployment-risk model is to have any chance of successfully replicating key features of the data, it will require the addition of at least one other state variable. To this end, and following much of the quantitative business cycle literature, I now assume that the household exhibits internal habit-formation in consumption,[38] so that its period utility for consumption is now given by

$$U\left(c_t - hc_{t-1}\right)$$

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

Second, the relatively simple structure of the baseline model produces a stark dichotomy, whereby in the unemployment regime all output adjustments occur along the extensive labor margin, while in the full-employment regime all adjustments occur along the intensive margin. In order to relax this stark dichotomy, in the quantitative version of the model I allow firms to be heterogeneous in terms of their fixed costs. That is, rather than assuming that all firms have fixed cost k, I assume that the n-th firm has fixed cost $k(n) \geq 0$, where $k(\cdot)$ is a non-decreasing function. This will allow for the possibility of there being regions where both extensive and intensive labor margin adjustments may occur.[39,40]

Third, as discussed earlier and in contrast to what is observed in the data, purely deterministic models of economic fluctuations tend to yield cycles of a constant length. This can be observed either as a very regular pattern in a plot of time series data generated from the model, or as one or more large spikes in the spectrum estimated from that data.[41] One of the key contributions of this paper is to show that by introducing a relatively small amount of randomness into a limit-cycle model it becomes possible to produce realistically irregular fluctuations. To this end, I also include in the model an exogenous random TFP process, $\tilde{\theta}_t$.[42]

---

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

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

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

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

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

6.1 Functional forms, calibration and estimation

Production is assumed to be of the Cobb-Douglas form

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

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

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

while disutility of labor is taken to be of the form

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

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

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

where n_0, η and $\bar{k}$ are parameters. This function is piecewise linear with three regimes: low-n firms have fixed cost zero, high-n firms have fixed cost $\bar{k}$, and over the intermediate range k rises linearly from zero to $\bar{k}$.[43] Finally, I assume the TFP process is given by

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

Several of the model parameters were directly calibrated. In particular, I set the labor share at a standard value of $\alpha = 2/3$. The inverse Frisch elasticity was calibrated at the widely used level $\omega = 1$. I set the depreciation rate and discount factor at standard values of $\delta = 0.025$ and $\beta = 0.99$, respectively, and normalize the maximum fixed cost at $\bar{k} = 1$. Finally, the fraction of purchases entering the durables stock was calibrated at $\gamma = 0.192$, which is the average ratio of durables to total consumption in the National Income and Product Accounts data.[44] The remaining parameters were estimated.

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

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

Solving the model for a particular parameterization was done using the parameterized expectations (PE) approach.[45] Given this solution, a large data set ($T = 100,000$ periods in length) was simulated and, after taking logs of the resulting hours series and detrending it with the same BP filter as used for the data, the spectrum of log-hours was estimated. The non-calibrated parameters were then estimated so as to minimize the average squared difference between the model spectrum and the spectrum estimated from the data (see panel (b) of Figure 1). Further details of the solution and estimation procedure are presented in Appendix B.

Estimated parameter values are reported in Table 1. Several things should be noted. First, the TFP process is close to the process that would be estimated directly from productivity data. For example, using John Fernald's (2014) measure of business-sector labor productivity growth over the sample period (1960Q1-2012Q4),[46] after cumulating, linearly detrending, and fitting an AR(1) process, one obtains a persistence estimate of 0.974 and an innovation standard deviation of 0.713%, yielding an unconditional productivity standard deviation of 3.16%.[47] The corresponding parameters estimated for the unemployment-risk model, meanwhile, are $\rho = 0.969$ and $\sigma = 0.570$, respectively, which yields an unconditional standard deviation of 2.30%. The fact that the model only features a single shock, and that the variance of that shock in the model is, if anything, smaller than its data counterpart highlights the more general observation that models featuring deterministic fluctuations may not require the presence of large amounts of exogenous variation in order to generate empirically reasonable business cycles.

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

The remaining parameters in Table 1 are composed mainly of uninteresting scale parameters, and parameters for which few if any precedents exist. The parameter τ, which captures the strength of the household's desire to reduce spending in response to a rise in unemployment risk, falls into the latter category. Given its central role in the model, however, it deserves some comment. If interpreted narrowly as a one-period financial premium on debt vis-à-vis saving, the estimate of $\tau = 0.27$, or 27%, clearly exceeds typical borrowing-lending

[45] See, for example, den Haan and Marcet (1990) and Marcet and Marshall (1994). Details can be found in Appendix B.

[46] Available at http://www.frbsf.org/economic-research/economists/jfernald/quarterly_tfp.xls.

[47] Similar values are obtained when using Fernald's TFP or utilization-adjusted TFP measures instead of labor productivity.

spreads as reported in the literature. However, there are several reasons to think this view of τ may be overly restrictive. First, in order to avoid significantly complicating the model, conditional on the employment rate an individual worker's probability of being employed is assumed to be independent from quarter to quarter. If the actual employment state of an individual exhibits persistence, then considering only one-period financial costs may understate households' desire to reduce spending in response to an increase in unemployment. Second, borrowing-lending spreads that reflect average borrowing rates faced by all households may not accurately reflect rates faced by unemployed individuals, which are likely to be higher. Third, many unemployed individuals may in fact be unable to access financial markets at all, instead being forced to rely on costly asset liquidations and/or reduced consumption levels in order to meet their obligations, the potential for either of which may cause households to strongly reduce their desired spending. To the extent that any or all of these factors should be subsumed into τ, the value estimated here may not be unreasonable.

6.2 Main Results

To illustrate the deterministic mechanisms, I first report results obtained when shutting down the TFP shock (i.e., setting $\sigma = 0$).[48] Panel (a) of Figure 4 plots a simulated 212-quarter sample[49] of log-hours generated from this deterministic model. Two key properties should be noted. First, the model is clearly capable of generating cycles of a reasonable length, which in this case is approximately 30 quarters. As noted in section 2, the apparent inability of models of deterministic flutcutations to generate cycles of quantitatively reasonable lengths appears to have been one of the factors leading to the abandonment of this literature. As this exercise demonstrates, however, unreasonable cycle lengths are by no means an unavoidable property of these models. Second, notwithstanding the reasonable cycle length, it is clear when comparing the simulated data in Figure 4 to the actual data in Figure 1 that the fluctuations in the deterministic unemployment-risk model are far too regular,[50] a shortcoming shared by many earlier models of deterministic fluctuations.

These properties of the deterministic model—i.e., a highly regular 30-quarter cycle—can

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

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

[50]Note that the cycles clearly do not exactly repeat themselves. As alluded to in footnote 19, this property is due to the discrete-time formulation of the model. In a continuous-time version of the model, the cycles would necessarily repeat themselves, a direct consequence of the Poincaré-Bendixson Theorem (see, e.g., Guckenheimer and Holmes (2002), p. 44).

also be seen clearly in the frequency domain. Panel (b) of Figure 4 plots the spectrum for the deterministic model (dashed line), along with the spectrum for the data (solid line) for comparison.[51] Consistent with the pattern in the time domain, the spectrum exhibits a peak at around 30 quarters. Further, the regularity of the cycle is manifested as a large spike in the spectrum. In contrast, the spectrum estimated from the data is much flatter.

Re-introducing the TFP shock into the model, we see a markedly different picture in both the time and frequency domains. Panel (a) of Figure 5 plots a 212-quarter sample of log-hours generated from the stochastic model. While clear cyclical patterns are evident in the figure, it is immediately obvious that the inclusion of the TFP shock results in fluctuations that are significantly less regular than those generated in the deterministic model, appearing qualitatively quite similar to the fluctuations found in Figure 1 for actual data. This is confirmed by the spectrum, which is plotted in panel (b) of Figure 5 alongside the data spectrum. Also plotted is a pointwise 90% simulated confidence interval from the model for data sets of the same length as the data (i.e., 212 quarters).[52] The stochastic model clearly matches the data quite well in this dimension, including possessing a peak near 40 quarters and, as compared to the deterministic model, lacking any large spike. The good fit of the model can also be seen by looking at the autocovariance function (ACF) of hours, i.e., $Cov\left(L_t, L_{t-k}\right)$, where k is the lag (in quarters). Panel (a) of Figure 6 plots the result for the first 40 lags for both the data and model, along with pointwise 90% confidence intervals. As the figure shows, the curves lie nearly on top of one another, indicating that the model matches the data very well in this dimension also.[53]

To verify that the good fit of the spectrum is not driven by the choice of filter, Figure 7 plots the data and model spectra for hours under four alternative filtering choices.[54] Panels (a)-(c) present results for three alternative band-pass filters with different upper bounds (100, 60, and 40 quarters, respectively), while panel (d) plots spectra using a Hodrick-Prescott filter with parameter 1600. As the figure shows, the model fits the data very well in all cases.

Next, it should be emphasized that the exogenous shock process in this model primarily accelerates and decelerates the endogenous cyclical dynamics, causing significant random fluctuations in the length of the cycle while only modestly affecting its amplitude. For example, in the deterministic version of the model the standard deviation of log-hours is 0.026, while

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

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

[53]Note that the ACF is simply the inverse Fourier transform of the spectrum. Since the spectrum of the model and data are similar, we would expect the ACF to be similar as well, a property clearly verified in Figure 6.

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

in the stochastic model it is 0.033, implying that 79% of the standard deviation of hours is due to deterministic mechanisms. In contrast, if this TFP process were the only shock process operating in the widely-cited model of Smets and Wouters (2007), for example, it would generate a standard deviation of log-hours of only 0.005. This again suggests the more general point that, if one is willing to consider the class of models capable of generating deterministic fluctuations, then a very parsimonious set of shocks that are small in magnitude can potentially yield qualitatively and quantitatively reasonable fluctuations.

As a final exercise in this section, it is worth briefly further comparing the above results to those of Smets and Wouters (2007). Their model has received much attention in the literature for its ability to fit well a number of key macroeconomic data series. Panel (a) of Figure 8 shows the spectrum for hours worked as generated by the Smets and Wouters (2007) model at the reported median posterior parameter values. As suggested by the relatively close fit, their model also matches patterns in the hours data reasonably well, though not quite as well as the unemployment-risk model.[55]

More insight into the drivers of fluctuations in the Smets and Wouters (2007) model can be obtained by looking at a spectral variance decomposition; that is, by decomposing the total variance at each individual periodicity into the portions that are attributable to each of the shocks in that model. Panel (b) of Figure 8 presents such a decomposition. It is clear from the figure that, in the range of periodicities reponsible for the bulk of the variance of hours, the two mark-up shocks (price and wage) in the Smets and Wouters (2007) model account for by far the largest portion. In fact, the proportion of the total hours variance that is explained by the mark-up shocks rises monotonically with periodicity, explaining around a third of the variance of hours by the 24-quarter periodicity and over half by the 36-quarter periodicity.[56] In contrast, the unemployment-risk model presented here is equally capable of matching the spectrum in hours, but does so with only a reasonably-sized TFP shock and without relying on poorly motivated mark-up shocks.

6.3 Additional Results

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

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

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

Panel (a) of Figure 9 compares the spectrum of output for the data and the stochastic model.[57] As shown in the figure, the model spectrum matches the data reasonably well, though it is somewhat too large (indicating too much output variance in the model), and the average periodicity is somewhat too low. The second observation should not be too surprising, as the model does not include capital as a factor of production. Since productive capital tends to exhibit lower-frequency fluctuations than labor (the other factor of production), all else equal its omission from the model will cause the average periodicity of output to be too small. Panel (d) of Figure 6, meanwhile, plots the ACF for output, which confirms the first observation: the variance of output in the model (i.e., the autocovariance at lag $k = 0$) is slightly larger than in the data. Notwithstanding this, however, the spectrum and ACF for output in the data lies well within a 90% confidence interval for the model, suggesting a relatively good overall fit.

Next, panel (b) of Figure 9 plots the coherence between hours and output for the data and for the stochastic model.[58] Coherence is analogous to a regression R^2, giving the proportion of the variance of hours that can be linearly predicted by output at a given periodicity. A coherence of one would thus indicate that hours and output are perfectly correlated at that periodicity, while a coherence of zero would indicate that hours and output are orthogonal. In the data (solid line in the figure), we see that at the lowest periodicities hours and output are modestly correlated, with coherence around 0.4-0.5. As the periodicity rises, the coherence initially increases relatively rapidly, reaching a peak of 0.87 at around 13 quarters. Over this range, as indicated by the dashed line in the figure the model coherence matches the data very well. Beyond the 13-quarter periodicity, however, the data and model begin to diverge somewhat. The data coherence largely flattens out, with a gradual downward slope, reaching 0.82 at the 80-quarter periodicity. The model coherence, meanwhile, rises somewhat over this range. As with the spectrum of output, the discrepancy between the data and model coherences at higher periodicities can be explained by the lack of productive capital in the model.[59] Notwithstanding this discrepancy, however, the basic qualitative properties of the relationship between hours and output in the data—namely, moderate correlation at higher

[57] Data series for output is the log of nominal GDP, deflated by population and the GDP deflator, then detrended using a BP(2,80) filter using the same procedure as with hours worked (see section 3). Output in the model is the sum of wage earnings and firm profits, which is equal to total production net of fixed costs, i.e., $\tilde{\theta}_t \left[\phi_t F(\ell_t) - \int_0^{n_t} k(x)\,dx\right]$, where n_t is the number of firm entrants at date t.

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

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

frequencies but significant correlation at medium-to-low frequencies (including the range of frequencies in which the bulk of variation occurs)—are well-captured by the model.

While coherence measures the strength of the relationship between two series at a given periodicity, it provides no information about the sign of this relationship or whether one series tends to lead the other. To address how well the model fits in these dimensions, panels (b) and (c) of Figure 6 plot the cross-covariance function (CCF) for hours and output. Two things should be noted from these plots. First, hours and output are positively correlated in both the model and data. Second, in the model hours and output are in phase (i.e., the peak of the CCF occurs at a lag of $k = 0$), while in the data the peak occurs at the point where output leads hours by one quarter. Nonetheless, the CCF is close to flat in the data between its peak and $k = 0$,[60] suggesting that any lead of output is weak at best. Further, as suggested by the reported 90% confidence intervals, over all the cross-covariance between output and hours is well-captured by the model.

Finally, while we have established that the model does a good job of matching patterns in total hours, consider the model's implications for its two component parts, the employment rate, ϕ_t, and hours-per-worker, ℓ_t. Panel (c) of Figure 9 shows spectra for the data and stochastic model for these two series.[61] From the figure, we see that the spectrum of the employment rate from the model matches fairly well the one from the data, and in particular the employment rate exhibits an overall level of volatility that is close to the volatility in the data. Thus, this model addresses one of the frequent criticisms of many models of unemployment in the literature, which is that they generate too little employment volatility.[62]

On the other hand, the model does a relatively poor job of matching behavior in hours-per-worker. In particular, while the basic pattern of the model spectrum is close to that in the data, the model spectrum is in most places too small, especially beyond the lowest periodicities. This suggests that the model features too little in the way of movements along the intensive labor margin.[63] To understand why, recall that when the economy moves into a region where the fixed-cost function $k(\cdot)$ is increasing, forces come into play which cause output fluctuations to occur on both intensive and extensive labor margins. Recall also that the former are associated with strategic substitutability (through changes in the price of

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

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

[62] See for example Shimer (2005).

[63] As Figure 9 shows, extensive-margin fluctuations are an order of magnitude larger than intensive-margin fluctuations in both the model and the data. As a result, even though the model does not capture well the intensive-margin fluctuations, this has little impact on the fit of total hours, which is driven primarily by extensive-margin fluctuations.

goods), while the latter are associated with strategic complementarity (through changes in unemployment risk). If a given change in output occurs too much along the intensive margin (as is the case for this parameterization of the model), the associated strategic substitutability tends to push the economy back towards the steady state quickly, so that any change in hours-per-worker is relatively small and short-lived.[64]

6.4 Multiple Equilibria and Indeterminacy

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

In particular, define

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

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

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

Proposition 5. *The steady state of the unemployment-risk model is unique if $\tau < \bar{\tau}$. The period-t static equilibrium is unique if $\tau < \tilde{\theta}_t^2 \tau^*$.*

Ex ante, it is not clear whether imposing the constraints on τ from Proposition 5 is restrictive in practice. The results from the estimation reported above, however, give no indication that these constraints are binding. In particular, at the parameter values reported in Table 1, we have $\bar{\tau} = 2.61$ and $\tau^* = 0.82$, both well above the value of $\tau = 0.27$. Clearly, the constraint ensuring a unique steady state is not binding at the optimal parameter values. The constraint ensuring a static equilibrium, meanwhile, depends on the level of productivity $\tilde{\theta}_t$, which can in principle be arbitrarily small, and thus the constraint may be violated with

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

strictly positive probability. Nonetheless, given the size of the estimated TFP shock, this probability is negligible in practice. For example, in 100,000 simulated periods, the smallest value of $\tilde{\theta}_t^2 \tau^*$ that occurred was 0.69, still more than twice the value of τ.

While there are relatively simple analytical conditions that can be obtained to ensure uniqueness of the steady state and of static equilibrium, verifying dynamic determinacy—that is, the presence of a unique path converging to the limit cycle for a given initial state—is more challenging, since no analytical results are available in general. Nonetheless, in numerical simulations I was unable to find any evidence of indeterminacy. In particular, given initial values for the state variables and arbitrary initial values for the jump variables (chosen in practice from some neighborhood of the PE solution), one may simulate a non-stochastic version of the model forward.[65] If, for a given initial state, the system were to converge to the limit cycle for multiple combinations of initial jump variables, this would indicate the presence of indeterminacy. However, performing this experiment many times beginning from different initial conditions, in all cases the system eventually exploded, which suggests to me that indeterminacy is not likely to be an issue here.

7 Conclusion

Conventional models of the business cycle usually feature fluctuations driven exclusively by exogenous shocks around a unique stable steady state. In these models, booms and subsequent recessions are typically unrelated, each being driven by different independent realizations of the underlying shocks. A contrasting view is that booms and busts are inherently related, with a boom sowing the seeds of a subsequent bust, which then sets the stage for the next boom. This second view received some attention in an older literature, but formal attempts to model underlying mechanisms appear to have been largely abandoned due to the perception that implausible assumptions or parameter values were necessary to generate quantitatively reasonable fluctuations.

In this paper, I present a purely deterministic general-equilibrium model featuring strategic complementarity near the steady state and show that it can give rise to a stable limit cycle. The limit cycle arises through a simple micro-founded mechanism in a rational-expectations environment. Cycles emerge endogenously, and thus the model does not require shocks in order to generate fluctuations, nor does it rely on the existence of multiple equilibria or dynamic indeterminacy.

Since cycles would indefinitely repeat themselves in the absence of shocks, a TFP shock

[65] See Appendix D for details.

is introduced into the model in order to create irregularities. The model is then estimated to match the spectrum of US hours. In contrast to results suggested in the earlier literature, I find that the model is able to match this spectrum quite closely. The TFP shock in the model is also shown to be of a reasonable persistence and relatively small size, accounting for around a fifth of the standard deviation of hours in the model. This result highlights the important insight that models capable of generating deterministic fluctuations do not require the addition of large, persistent, poorly-motivated shocks in order to match the patterns in the data, which is a common criticism of conventional models.

References

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

Beaudry, Paul, Dana Galizia, and Franck Portier (2014). "Reconciling Hayek's and Keynes' views of recessions." Working Paper.

Benhabib, Jess and Kazuo Nishimura (1979). "The Hopf Bifurcation and the Existence and Stability of Closed Orbits in Multisector Models of Optimal Economic Growth." *Journal of Economic Theory*, 21(3):421–444.

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

Blanchard, Olivier J. and Charles M. Kahn (1980). "The Solution of Linear Difference Models under Rational Expectations." *Econometrica*, 48(5):1305–1311.

Boldrin, Michele, Nobuhiro Kiyotaki, and Randall Wright (1993). "A dynamic equilibrium model of search, production, and exchange." *Journal of Economic Dynamics and Control*, 17(5):723–758.

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

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

Christiano, Lawrence J. and Terry J. Fitzgerald (2003). "The Band Pass Filter." *International Economic Review*, 44(2):435–465.

Coles, Melvyn G. and Randall Wright (1998). "A Dynamic Equilibrium Model of Search, Bargaining, and Money." *Journal of Economic Theory*, 78(1):32–54.

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

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

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

den Haan, Wouter J. and Albert Marcet (1990). "Solving the Stochastic Growth Model by Parameterizing Expectations." *Journal of Business & Economic Statistics*, 8(1):31–34.

Diamond, Peter and Drew Fudenberg (1989). "Rational expectations business cycles in search equilibrium." *Journal of Political Economy*, 97(3):606–619.

Farmer, Roger E. A. and Jang-Ting Guo (1994). "Real Business Cycles and the Animal Spirits Hypothesis." *Journal of Economic Theory*, 63(1):42–72.

Fernald, John (2014). "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Federal Reserve Bank of San Francisco Working Paper 2012-19.

Glendinning, Paul (1994). *Stability, Instability and Chaos: an introduction to the theory of nonlinear differential equations.* New York, NY: Cambridge University Press.

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

Guckenheimer, John and Philip Holmes (2002). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields.* New York: Springer.

Hammour, Mohamad Lutfi (1989). *Endogenous Business Cycles: Some Theory and Evidence.* PhD Thesis, Massachusetts Institute of Technology.

Jaimovich, Nir (2007). "Firm dynamics and markup variations: Implications for sunspot equilibria and endogenous economic fluctuations." *Journal of Economic Theory*, 137(1):300–325.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2010). "Investment shocks and business cycles." *Journal of Monetary Economics*, 57(2):132–145.

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

Kaplan, Greg and Guido Menzio (2014). "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations." Working Paper.

Lagos, Ricardo and Randall Wright (2005). "A Unified Framework for Monetary Theory and Policy Analysis." *Journal of Political Economy*, 113(3):463–484.

Marcet, Albert and David A. Marshall (1994). "Solving Nonlinear Rational Expectations Models by Parameterized Expectations: Convergence to Stationary Solutions." Federal Reserve Bank of Minneapolis Discussion Paper 91.

Rocheteau, Guillaume and Randall Wright (2005). "Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium." *Econometrica*, 73(1):175–202.

Scheinkman, José A. (1976). "On Optimal Steady States of n-Sector Growth Models when Utility is Discounted." *Journal of Economic Theory*, 12(1):11–30.

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

Schmitt-Grohé, Stephanie and Martín Uribe (2012). "What's News in Business Cycles." *Econometrica*, 80(6):2733–2764.

Shimer, Robert (2005). "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *The American Economic Review*, 95(1):25–49.

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

Stock, James H. and Mark W. Watson (1999). "Business Cycle Fluctuations in US Macroeconomic Time Series." In John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, vol. 1A. Amsterdam: Elsevier.

A Proofs

A.1 Proof of Proposition 1

Recall that

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

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

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

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

□

A.2 Proof of Proposition 2

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

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

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

$$X_{t+1} = (1-\delta)(X_t + e_t) \tag{A.2}$$

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

$$b > vp^* \frac{\tau}{e^*} \equiv b_0 \tag{A.3}$$

where e^* is output per firm (net of fixed costs) when the economy is in the unemployment regime and p^* is the price in the unemployment regime, as described in section 5.1 for the static model (see footnote 29 regarding p^*). I henceforth assume that (A.3) holds.

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

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

with the steady-state level of X then given by

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

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

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

which is clearly implied by (A.3).

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

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

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

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

Thus, if $\beta = 1$ the steady state is in the unemployment regime by construction, and by continuity of $\bar{e}$ in β the steady state is also necessarily in the unemployment regime for β sufficiently close to one. This implies the existence of a $\underline{\beta} < 1$ such that the steady state is in the unemployment regime when $\beta > \underline{\beta}$. Assume henceforth that $\beta \in \left(\underline{\beta}, 1\right)$ and note that this implies that $p'(\bar{e}) = 0$ and $\phi'(\bar{e}) = 1/e^*$.

Next, linearizing equations (A.2)-(A.3) around this steady state and solving, we may obtain in matrix form

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

Thus, the steady state is locally stable if and only if at least one of the two eigenvalues of A

lies inside the complex unit circle. These eigenvalues are given by

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

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

$$b = \left[1 + q\left(1-\delta\right)^2\right] b_0 \tag{A.4}$$

for some $q > 0$, and note that as long as $\delta < 1$, which I henceforth assume, such a value of b satisfies (A.3). One may then show that the eigenvalues are complex as long as

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

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

□

A.3 Proof of Proposition 3

Let

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

where e^* is output per firm (net of fixed costs) when the economy is in the unemployment regime and p^* is the price in the unemployment regime, as described in section 5.1 for the static model (see footnote 29 regarding p^*). It can be verified that maximizing

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

subject to (2) implements the de-centralized equilibrium outcome in the neighborhood of an unemployment-regime steady state. Thus, using

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

in problem (4) satisfies the desired properties. Next, we may obtain

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

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

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

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

□

A.4 Proof of Proposition 4

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

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

$$U'\left(X_t+e_t\right)-vp\left(e_t\right)=\beta\left(1-\delta\right)\left(1-\gamma\right)U'\left(X_{t+1}+e_{t+1}\right)-\beta\left(1-\delta\right)vp\left(e_{t+1}\right)$$

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

Assume the steady state is in the unemployment regime, so that $p\left(e_t\right)=p^*$ in a neighborhood of the steady state.[66] Linearizing around this steady state, assuming $\beta\left(1-\delta\right)\left(1-\gamma\right)>0$ we may obtain in matrix form

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

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

□

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

B Solution and estimation

B.1 Solution

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

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

$$\mu_t = \mathbb{E}_t\left\{\beta h\left[a - b\left(X_{t+1} + \tilde{\theta}_{t+1}\tilde{e}_{t+1} - hc_t\right)\right]\right\} \tag{B.2}$$

$$\lambda_t = \mathbb{E}_t\left\{\beta\left[a - b\left(X_{t+1} + \tilde{\theta}_{t+1}\tilde{e}_{t+1} - hc_t\right) + (1-\delta)\,\lambda_{t+1} - \mu_{t+1}\right]\right\} \tag{B.3}$$

$$c_t = X_t + \tilde{\theta}_t\tilde{e}_t \tag{B.4}$$

$$X_{t+1} = (1-\delta)\left(X_t + \gamma\tilde{\theta}_t\tilde{e}_t\right) \tag{B.5}$$

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

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

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

where $e^* \equiv \frac{\alpha}{1-\alpha}\bar{k}$ and $\bar{e} \equiv (n_0 + \eta)\, e^*$. Meanwhile, μ_t and λ_t are the Lagrange multipliers on the definition of consumption and the durables accumulation equations ((B.4) and (B.5)), respectively.

Conditional on the state variables X_t, c_{t-1} and θ_t, and on values of the Lagrange multipliers μ_t and λ_t, equation (B.1) can be solved for $\tilde{e}_t$. To obtain values of μ_t and λ_t, I employ the method of parameterized expectations as follows. Let $Y_t \equiv \left(X_t - \bar{X}, c_{t-1} - \bar{c}, \theta_t\right)'$ denote the vector of state variables (expressed as deviations from steady state). The expectations in equations (B.2) and (B.3) are assumed to be functions only of Y_t, i.e.,

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

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

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

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

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

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

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

These values and values for the state variables can be plugged into (B.1) to yield a solution for $\tilde{e}_t$, which can then be replaced in (B.4) and (B.5) to obtain values for the subsequent period's state. In practice, I use $N = 2$.[67]

To obtain Θ_μ and Θ_λ, I proceed iteratively as follows. Begin with some initial guesses $\Theta_{\mu,0}$ and $\Theta_{\lambda,0}$,[68] and generate a sample of length $T = 100,000$ of the exogenous process θ_t. Next, given $\Theta_{\mu,i}$ and $\Theta_{\lambda,i}$, assume that $g_j (Y_t) = \Theta_{j,i}' Y_t^{(N)}$ and simulate the path of the economy for T periods. Given this simulated path, let $\mathbf{Y}^{(N)}$ denote the matrix whose t-th row is given by $Y_t^{(N)\prime}$, and construct T-vectors $\tilde{g}_\mu$ and $\tilde{g}_\lambda$, the t-th elements of which are given respectively by

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

and

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

i.e., the terms inside the conditional-expectation operators in equations (B.2) and (B.3). Then

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

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

update the guesses of Θ_j via

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

and iterate until convergence.

B.2 Estimation

As discussed in section 6.1, estimation was done by searching for parameters to minimize $\overline{S^2}$, the average squared difference between the model spectrum and the spectrum estimated from the data.

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

C Definitions

A deterministic dynamic system characterizing the evolution of a state vector $z(t)$ over time can be expressed as a function $G : \mathcal{T} \times \mathcal{Z} \to \mathcal{Z}$. Here, $\mathcal{T}$ is the set of dates at which the system is defined (e.g., $\mathcal{T} = \mathbb{R}$ in a continuous-time formulation, and $\mathcal{T} = \mathbb{Z}$ in a discrete-time formulation), while $\mathcal{Z} \subset \mathbb{R}^n$ is the n-dimensional state space. The function G takes a date $t \in \mathcal{T}$ and a date-0 state $z(0) = z_0$ as inputs and returns $z(t) = G(t, z_0)$.[69] We focus here on time-invariant dynamic systems, i.e., systems for which $G(t + \Delta t, z_0) = G(\Delta t, G(t, z_0))$. We have the following definition.

Definition 1. *G exhibits **deterministic fluctuations** if, for some z_0, the following hold.*

(*a*) $\limsup_{t\to\infty} \|G(t, z_0)\| < \infty$.

(*b*) $\lim_{t\to\infty} G(t, z_0)$ *does not exist.*

In words, the system exhibits deterministic fluctuations if it neither diverges to infinity nor converges to a single point. Of particular interest for us will be one type of deterministic fluctuation, the limit cycle, defined as follows.

[69]Note that by definition $G(0, z_0) = z_0$.

Definition 2. *A subset $L \subset \mathcal{Z}$ is a **limit cycle** of G with prime period $k > 0$ if the following hold.*

- (a) *For any $z \in L$ and $\Delta t \in (0, k)$, we have $G(k, z) = z$ and $G(\Delta t, z) \neq z$.*
- (b) *For any $z, z' \in L$ there is some $\Delta t \geq 0$ such that $G(\Delta t, z) = z'$.*
- (c) *Let $\Gamma(t, z)$ denote the distance between $G(t, z)$ and the closest point in L.*[70] *Then there exists some $z \notin L$ such that either $\lim_{t\to\infty} \Gamma(t, z) = 0$ or $\lim_{t\to-\infty} \Gamma(t, z) = 0$.*

The first property here says that, if the system starts at a point in L, then it will return to that point $k > 0$ periods later (and no sooner). The second property says that as the system evolves beginning from any point in L it will visit every other point in L at some subsequent date. Finally, the third property says that there is some point *not* in the limit cycle such that, beginning from that point, the system will eventually converge to the limit cycle as it moves either forward or backward through time.

D Solving the model forward

In the non-stochastic case, we may re-arrange equation (B.1) to yield

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

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

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

where the conditions in Proposition 5 ensure that H is an invertible function. This value of e_t then gives c_t and X_{t+1} via equations (B.4) and (B.5), respectively. From equations (B.2) and (B.3) we can then solve for μ_{t+1} and λ_{t+1} as

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

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

[70]Properties (a) and (b) ensure that L is necessarily a closed set, so that such a closest point always exists.

Table 1: Parameter Values

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

Figure 1 – Hours Worked Data (1960-2012)

(a) Data Series

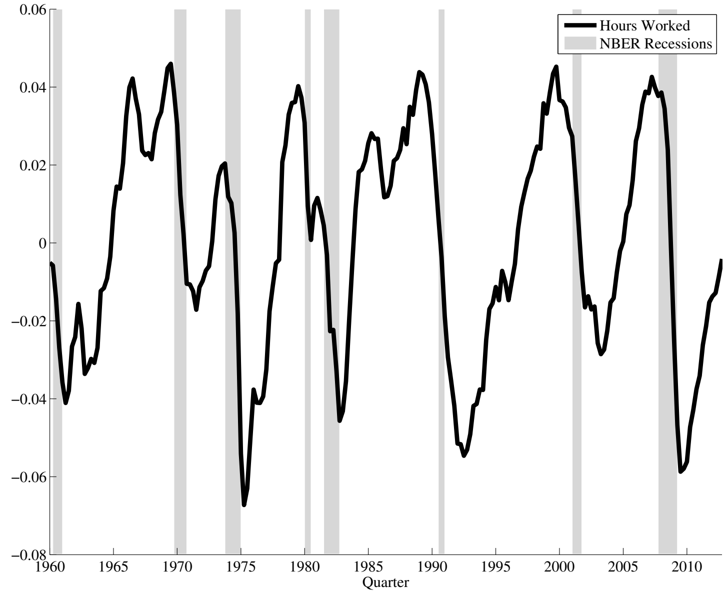

(b) Spectrum

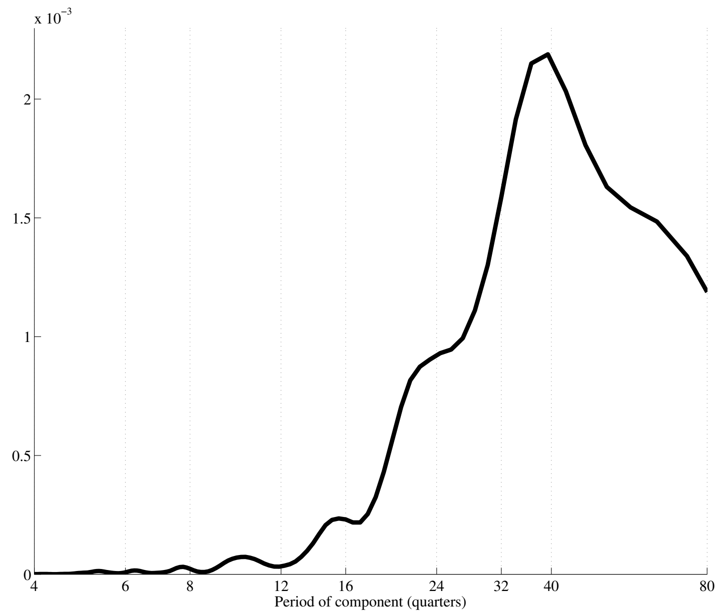

Notes: Hours Worked series is the log of BLS nonfarm hours worked divided by population, detrended with a BP filter to remove fluctuations with periods greater than 80 quarters. In panel (a), shaded areas are NBER-dated recessions. For panel (b), raw spectrum in obtained as the squared modulus of the discrete Fourier transform of the data series (scaled so that the integral with respect to angular frequency over the interval $[-\pi, \pi]$ equals the variance of the series). Spectrum in figure is kernel-smoothed raw spectrum. Kernel is a Hamming window with bandwidth parameter 11.

Figure 2 – Conditions for a Limit Cycle

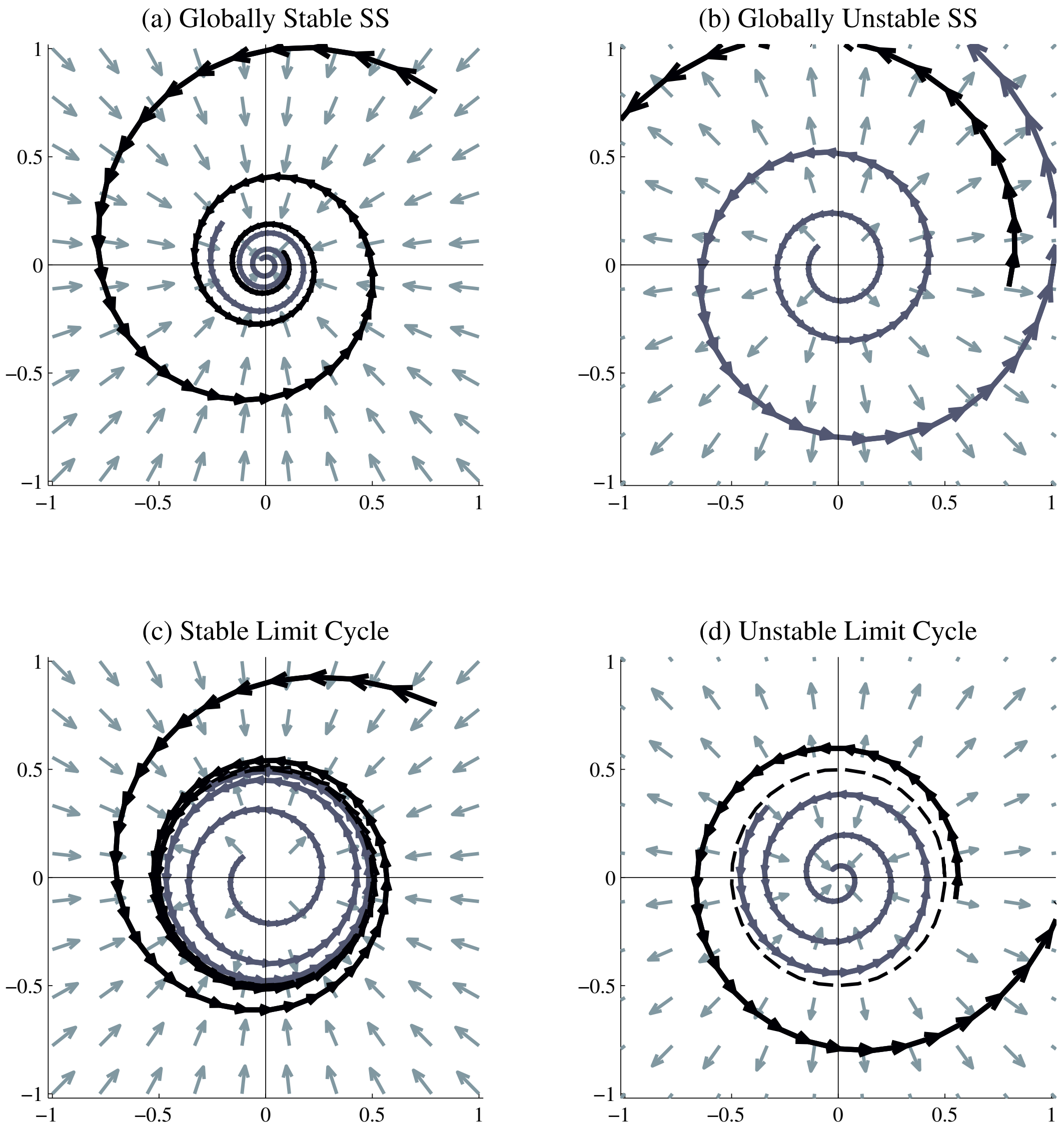

Panel (a): $a = 0.1$, $b = 0.2$. Panel (b): $a = -0.1$, $b = -0.2$. Panel (c): $a = 1$, $b = -0.5$. Panel (d): $a = -1$, $b = 0.5$. In all cases, $\bar{\theta} = 0.6\pi$.

Figure 3 – Static Equilibrium Determination

(a) Unemployment Regime

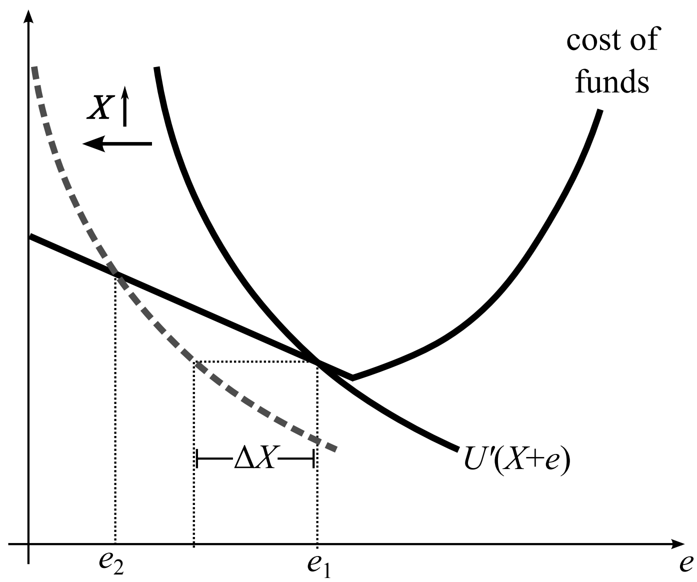

(b) Full-employment Regime

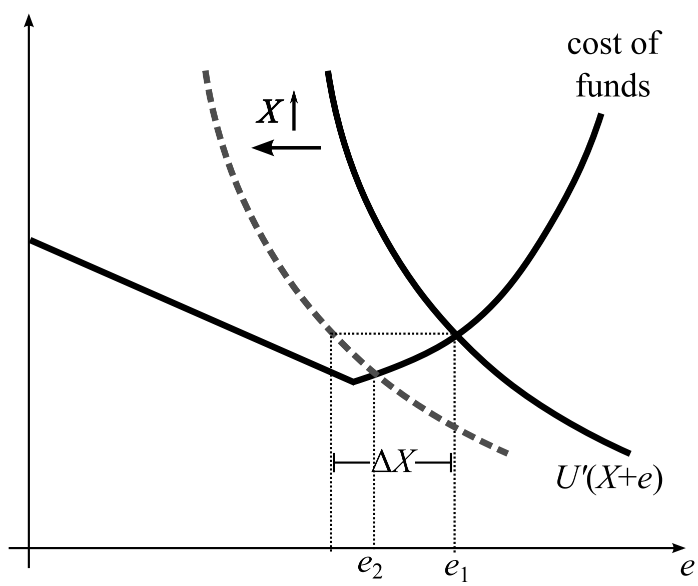

Figure 4 − Deterministic Model

(a) Sample of Simulated Hours Worked

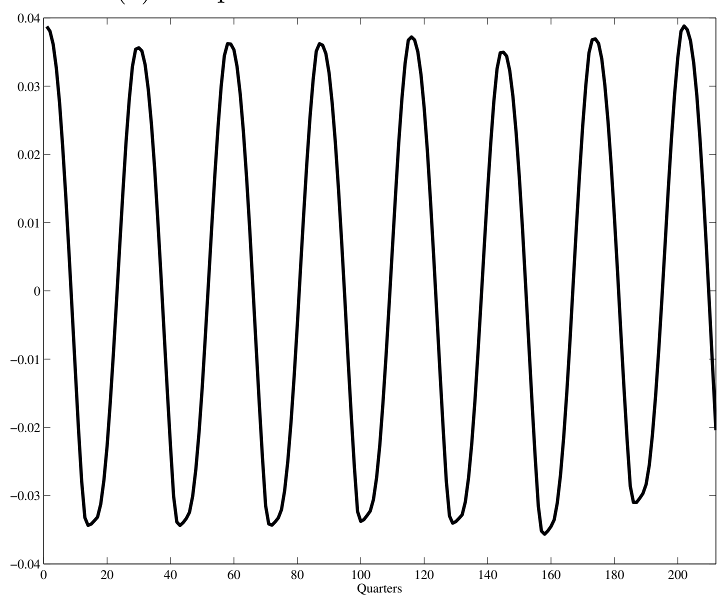

(b) Spectrum of Hours Worked

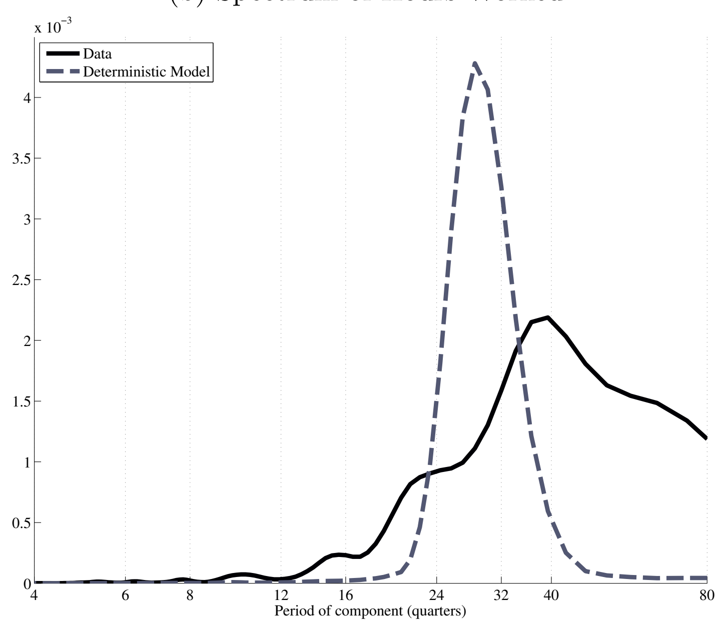

Note: Panel (a) shows 212-quarter simulated sample (same size as data set) of BP-filtered log(hours worked) ($\phi_t \ell_t$) generated from the deterministic model. Initial simulated series was 252 quarters long, with first and last 20 quarters discarded after BP-filtering. Details for computation of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

Figure 5 – Stochastic Model

(a) Sample of Simulated Hours Worked

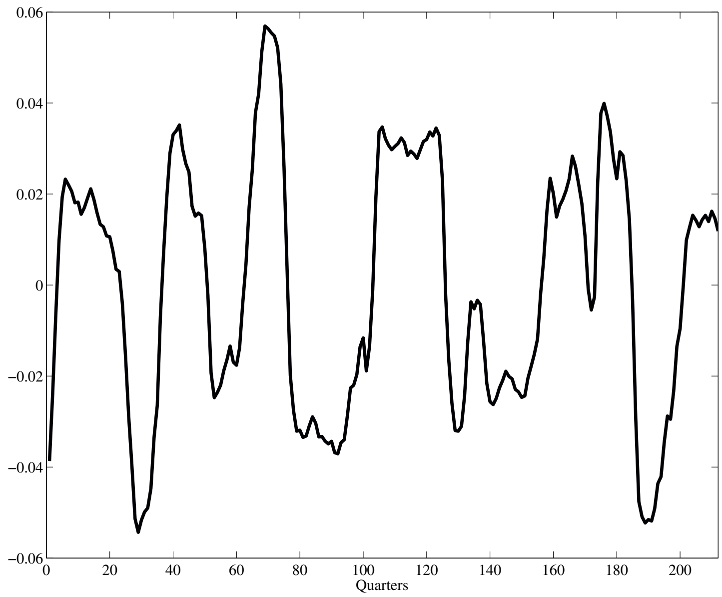

(b) Spectrum of Hours Worked

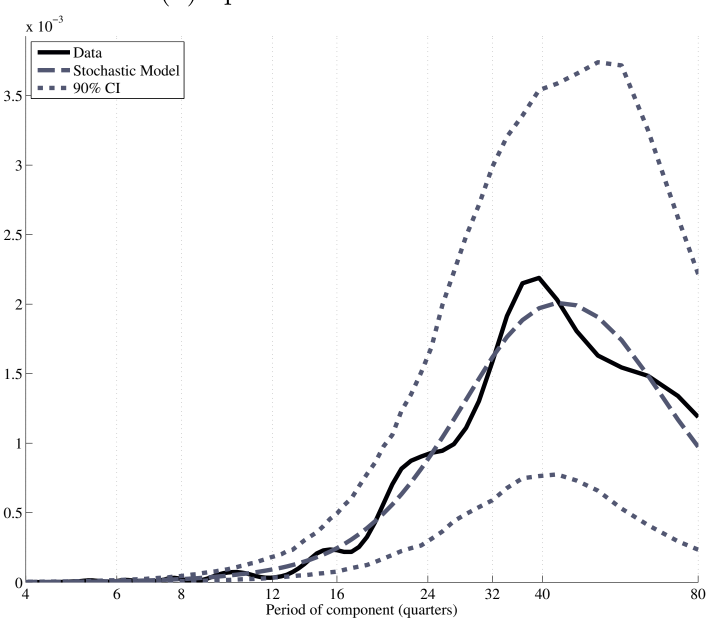

Note: Panel (a) shows 212-quarter simulated sample (same size as data set) of BP-filtered log(hours worked) ($\phi_t \ell_t$) generated from the stochastic model. Initial simulated series was 252 quarters long, with first and last 20 quarters discarded after BP-filtering. Details for computation of model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

Figure 6 – Autocovariance: Hours Worked (L) and Output (y)

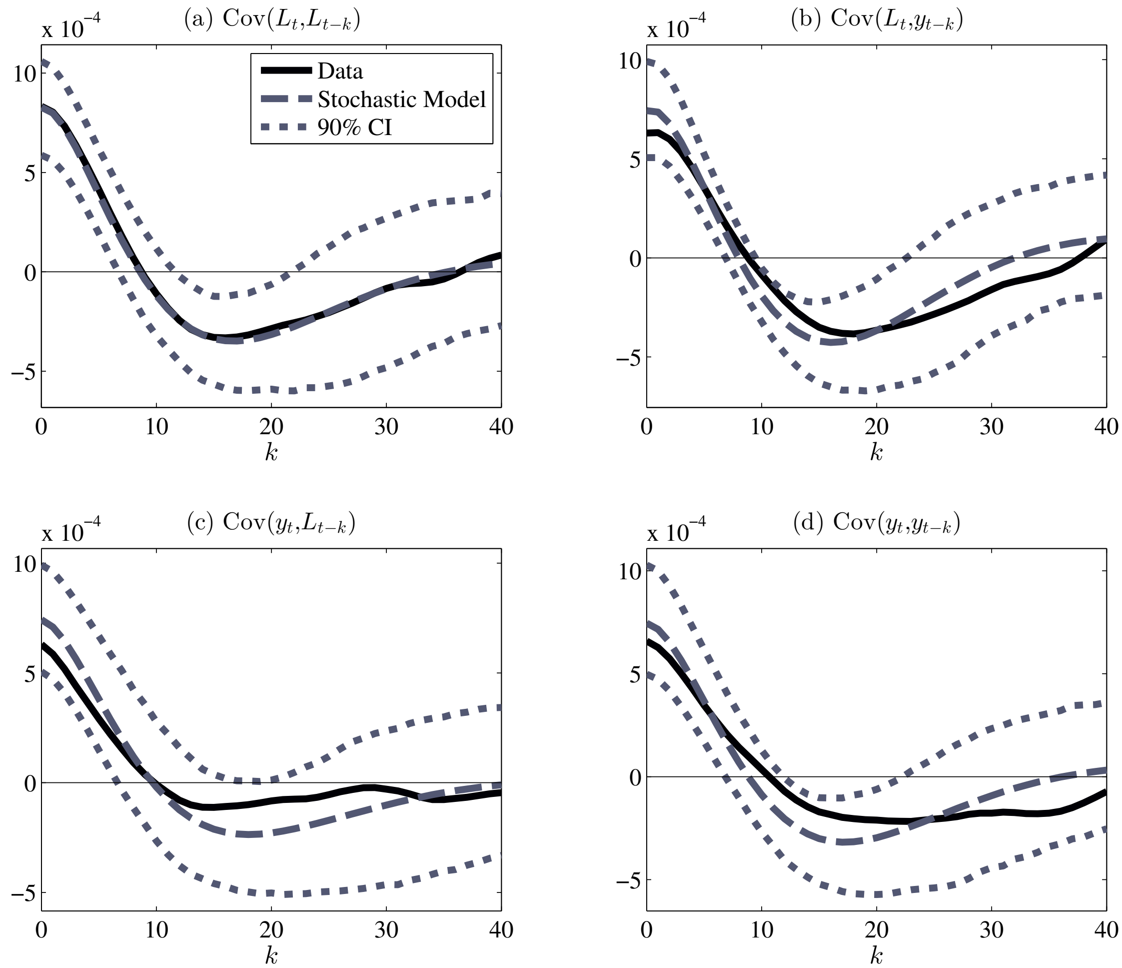

Note: Figure shows autocovariances of BP(2,80)-filtered hours and output in the data and stochastic model. k is the lag in quarters. Data series for output is the log of nominal GDP, deflated by population and the GDP deflator. Output in the model is the sum of wage earnings and firm profits, which is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

Figure 7 – Spectrum: Hours Worked (Alternative Filters)

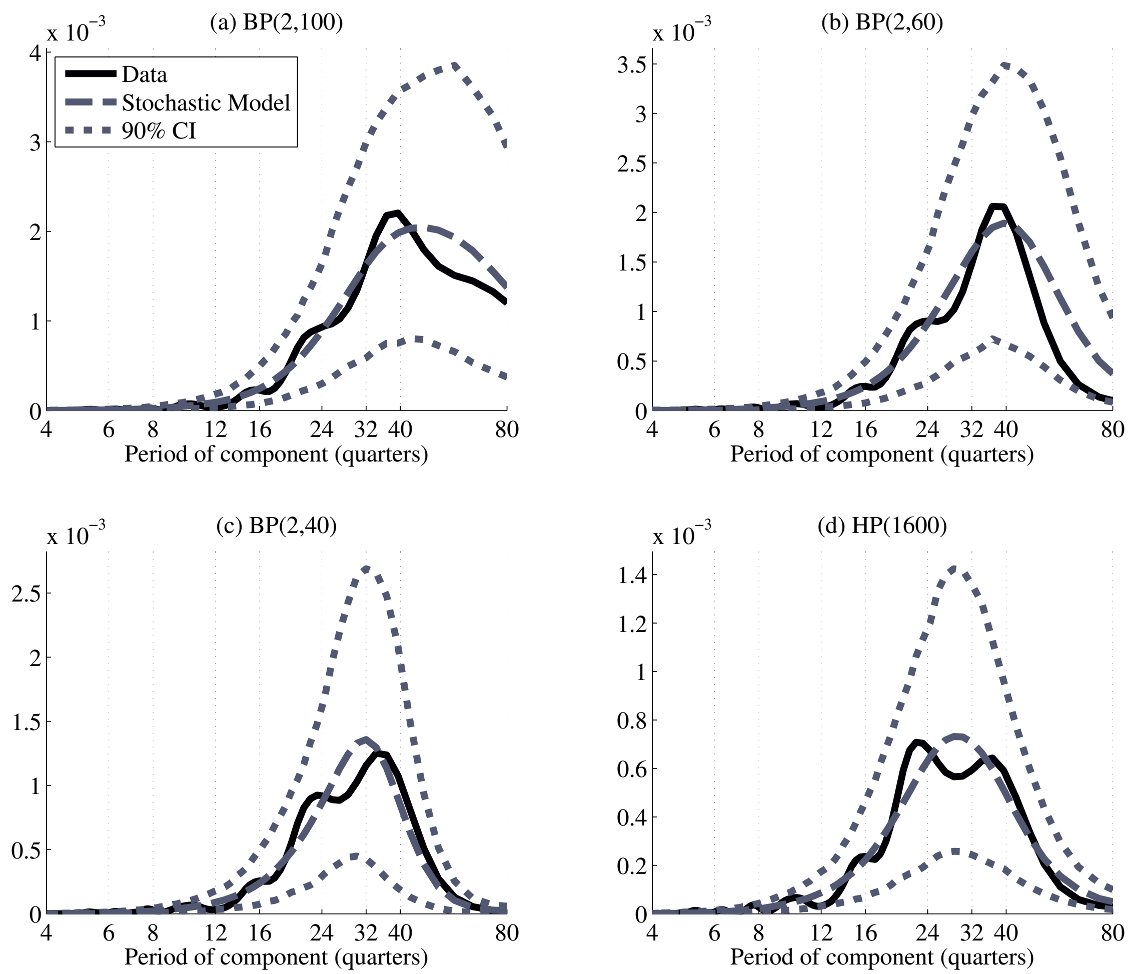

Note: Each panel plots corresponding data (solid) and model (dashed) spectrum using the reported filter instead of the baseline BP(2,80) filter. Dotted lines show pointwise 90% confidence intervals for the spectrum that would be estimated from a model-generated data set of the same length as the actual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

Figure 8 – Hours Worked in Smets-Wouters

(a) Data and Smets-Wouters Spectrum

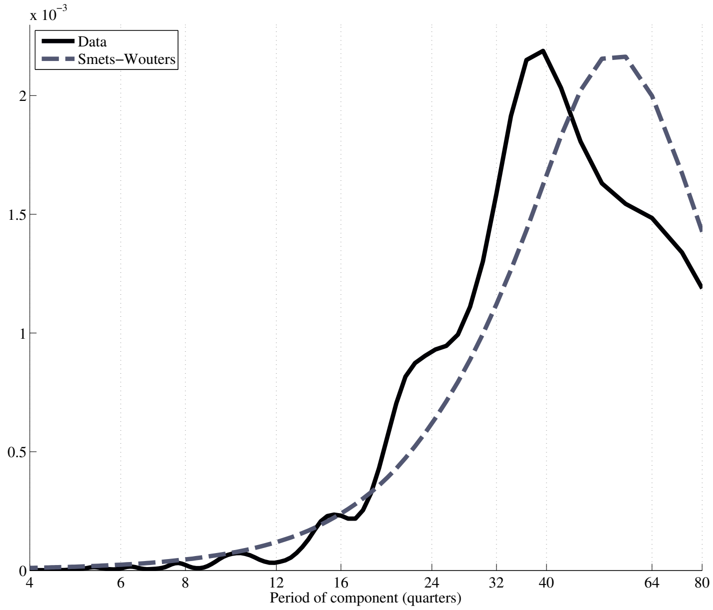

(b) Spectral Variance Decomposition (Smets-Wouters)

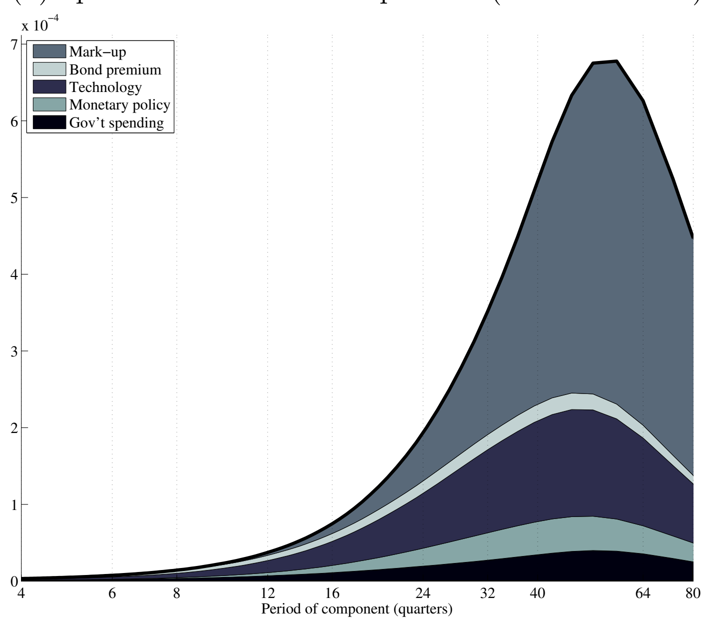

Notes: Data spectrum is as in Figure 1. Spectrum for Smets-Wouters (SW) obtained by simulating 10,000 data sets of the same size as the actual data series. For each simulation, the data was detrended and the spectrum estimated using the same procedures as for the actual data. A point-wise average was taken across all simulated spectra. Because the hours series used by SW for their estimation differs somewhat from the series used here, for purposes of comparability, in panel (a) the SW spectrum was scaled by a constant so that the total variance is the same as in the data. Panel (b) shows portion of variance at each periodicity attributable to each of the following shock groupings: "Mark-up" – price and wage mark-up shocks; "Bond Premium" – bond premium shock; "Technology" – TFP and investment-specific technology shocks; "Monetary policy" – monetary policy shock; "Gov't spending" – government spending shock.

Figure 9 – Spectrum: Output (Data and Stochastic Model)

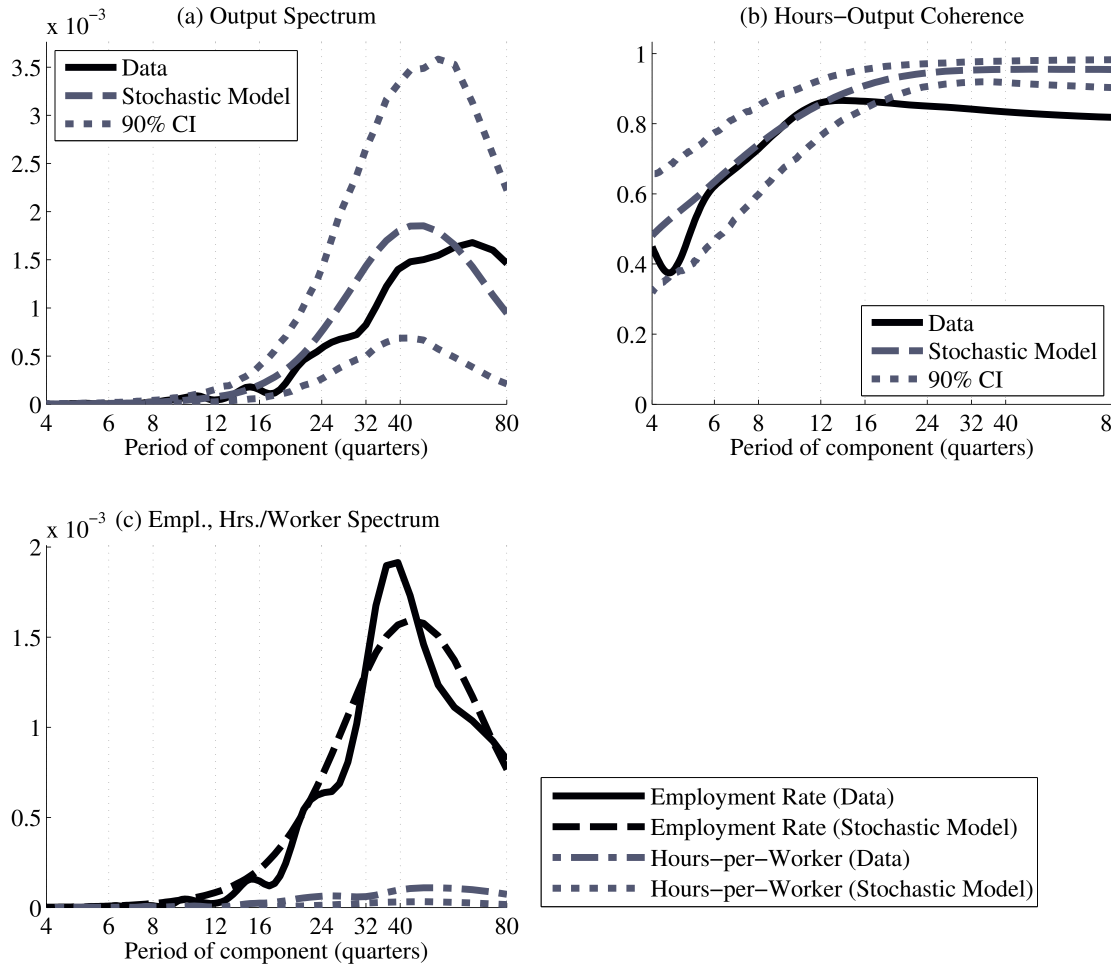

Notes: Data series for output is the log of nominal GDP, deflated by population and the GDP deflator. Data series for the employment rate is the log of the BLS's index of nonfarm business employment divided by population. Data series for hours-per-worker is the log of nonfarm business hours divided by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