Stagnant productivity and low unemployment: stuck in a Keynesian equilibrium

Wendy Carlin* and David Soskice**

Abstract: A major challenge is to build simple intuitive macroeconomic models for policy-makers and professional economists as well as students. A specific contemporary challenge is to account for the prolonged slow growth and stagnant productivity that has followed the post-financial crisis recession, along with low inflation despite low unemployment (notably in the UK). We set out a simple three-equation model, which extends the core model in our two recent books (Carlin and Soskice, 2006, 2015) to one with two equilibria and two associated macroeconomic policy regimes. One is the standard inflation-targeting policy regime with equilibrium associated with central bank inflation targeting through monetary policy. It is joined by a second, Keynesian policy regime and equilibrium, with a zero lower bound (ZLB) in the nominal interest rate and a ZLB in inflation in which only fiscal policy is effective (Ragot, 2015). Our approach is related to the Benigno and Fornaro (2016) Keynesian–Wicksellian model of growth with business cycles. It diverges from New Keynesian models because although we attribute model-consistent expectations to the policy-maker, we do not assume that these are the basis for inflation and growth expectations of workers and firms. We compare our approach to Ravn and Sterk's related multiple equilibrium New Keynesian model (Ravn and Sterk, 2016).
Keywords: stagnation, liquidity trap, zero lower bound, strategic complementarity, multiple equilibria, inflation-targeting
JEL classification: E32, E43,E52, O42

* University College London, CEPR, and Centre for Macroeconomics, e-mail: w.carlin@ucl.ac.uk
* London School of Economics, e-mail: d.w.soskice@lse.ac.uk
We are grateful to David Vines for prompting us to work on this topic. We have had very useful feedback and advice from Gianluca Benigno, Olivier Blanchard, Samuel Bowles, Steve Nickell, Bob Rowthorn, participants in the *Oxford Review of Economic Policy* seminars, and an anonymous referee. We thank Simon Wren-Lewis and Xavier Ragot for key insights. We are also grateful to Davide Melcangi for his assistance.

I. Introduction

In this paper we set out a three-equation model with two alternate parts—one summarized by a vertical and the other by a horizontal long-run Phillips curve. The former is appropriate to 'normal' macroeconomic conditions and the latter to a world following a financial crisis that occurred in a low-inflation environment. In doing so we borrow from some of the most interesting recent developments in macroeconomics that seek to analyse the long drawn-out post-financial crisis period of low growth and very low inflation by introducing heterogeneous agents, incomplete markets, and precautionary savings.

However, in contrast to many New Keynesian models, we pay attention to investment. In particular, we model investment as a multiple equilibrium strategic complementarities game in which uncertainty is the key to the choice of equilibrium. The simple two-part three-equation model with mechanisms that help account for recent experience builds on our earlier work on small tractable macroeconomic models targeted at policy-makers and students, and motivated by the need to understand contemporary macroeconomic problems.[1]

The aim is to highlight mechanisms that could explain the very slow recovery from the financial crisis, which has specific and unusual features. Since the crisis, the UK, along with other high-income countries, has been characterized by very low growth of *per capita* GDP, low core inflation, and a real interest rate close to zero. The slow recovery has lasted longer than the overhang of the financial effects of the crisis, which suggests other mechanisms are at play in addition to the deflationary deleveraging implied by private- and public-sector debt overhang effects and the impairment of the financial system emphasized by, among others, Reinhart and Rogoff (2009, 2010). A further puzzling aspect of current macroeconomic performance is the persistence of low growth and low and stable inflation while unemployment has fallen to low levels, and in the UK and some other countries, the employment rate has reached historically high levels.

We set out the 'Wicksellian–Keynesian' model in section III. The model we develop has two equilibria and two associated macroeconomic policy regimes. The Wicksellian equilibrium is the familiar one in a model of an inflation-targeting central bank: inflation is at target, unemployment is at the equilibrium of the labour market, and the real interest rate is positive and at its 'natural' or 'Wicksellian' level. Of particular interest for current developments is the other—Keynesian—equilibrium. In this equilibrium, inflation is constant, but at zero. The nominal interest rate is at the zero lower bound (ZLB), also zero. The implication is that the economy is at a labour market equilibrium with the real interest rate equal to zero. The economy becomes trapped in a Keynesian equilibrium because a sufficiently large fall in aggregate demand in a low inflation environment takes it to the ZLB and renders stabilizing monetary policy ineffective. A floor on wage inflation at zero prevents a downward deflationary spiral and explains why the economy remains in the Keynesian equilibrium.

We introduce investment as a multiple equilibrium game of strategic complementarity in which investment is 'trapped' at a low level in the Keynesian equilibrium, and together with a model of precautionary saving accounts for persistent low aggregate demand. Consistent with models of endogenous growth, we model productivity growth as a function of investment (productivity increases are embodied in the capital stock), so that productivity growth is also at its ZLB in a Keynesian equilibrium with very low net investment.

[1] In the first book (Carlin and Soskice, 1990), we introduced wage- and price-setting agents with the result that involuntary unemployment characterized the equilibrium of the labour market. The model was well-suited to analysing the trends and cross-country comparisons in unemployment across high-income countries that first emerged in the 1970s. In the second (Carlin and Soskice, 2006), the model was extended to introduce an inflation-targeting central bank and used to analyse policy-making by contemporary central banks. In the third (Carlin and Soskice, 2015; hereafter, CS2015), inflation-targeting was extended to the open economy. We also showed how the policy-maker's intended equilibrium at target inflation and equilibrium unemployment could be subject to a financial cycle fuelled by a financial accelerator process based on the use of housing as collateral and on a value at risk model of investment bank behaviour.

When the economy is in the Wicksellian regime, fluctuations in aggregate demand and productivity shocks are stabilized by conventional inflation-targeting monetary policy.[2] In the Keynesian regime, monetary policy is ineffective. We discuss circumstances under which fiscal policy can release the economy from the stagnation trap.

The paper is organized as follows. In section II, we summarize stylized facts of the period since the global financial crisis, focusing on the UK but presenting data for the US, Germany, and Japan as well. The basic two-part Keynesian–Wicksellian three-equation model is set out in section III. In section IV, we explain the persistence of low investment and stagnant productivity in the Keynesian equilibrium. In section V we investigate the puzzle of a Keynesian equilibrium at high employment. Section VI pulls together the microeconomic assumptions underlying the model so as to provide a comparison with two other small models—the standard New Keynesian model of Clarida *et al*. (1999) and a multiple equilibrium New Keynesian model proposed by Ravn and Sterk (2016).

II. Economic performance since the global financial crisis

The motivation for developing the simple model in this paper is the unusual combination of characteristics of economic performance experienced by major economies following the global financial crisis. The features are particularly pronounced in the UK, but as the following charts illustrate, are present elsewhere. The charts show data for the UK, US, Germany, and Japan. We date the post-crisis period from 2010 and it is shaded in each figure.

Figure 1 shows long-run hourly productivity trends for the economy as a whole from 1973. In all four countries, productivity growth has fallen below 1 per cent per annum, with that in the UK close to zero. Figure 2 charts the downward trend in the real interest rate (on 10-year sovereign bonds) for the G-7 countries (weighted by GDP). The real interest rate appears to have shifted down to a rate close to zero, after being between 2 and 4 per cent for the previous two decades.

Figure 3 presents data for the UK, US, Germany, and Japan on the growth in capital services per hour worked over the period since 1985. There is a common, dramatic fall in the rate of capital deepening across countries in the 2010–15 period. This applies to Germany and Japan, which did not experience domestic leverage-based financial crises in their economies in 2008–9, as well as to the US and UK.

From Figure 4 we see that, with the exception of Germany, real wage growth (in terms of the consumer price index, CPI) has more or less ceased in the post-crisis period.

[2] In order to concentrate in this paper on the Keynesian regime, we do not integrate into the two-equilibrium model the destabilizing financial cycle that can take the economy away from the Wicksellian equilibrium as developed in ch. 6 of CS2015.

Figure 5 shows that in the post-crisis period, developments in the labour markets of these economies are also notable: unemployment rates have fallen to levels below 6 per cent and employment rates are, with the exception of the US, historically high.

Finally, in spite of labour market tightness, nominal wage growth as shown in Figure 6 has been between –0.5 and 3 per cent per annum in all four countries in the post-crisis period. Core CPI inflation remains low, but with no deflationary spiral following the financial crisis.

To summarize, we need a model that will help to explain how an economy with an inflation-targeting central bank can exhibit the following features:

- virtual stagnation of productivity, capital services per hour and real wages;
- a real interest rate close to zero;
- low unemployment and high employment rates;
- low growth of nominal wages.

Figure 1: Trend hourly productivity growth (1973–2015)

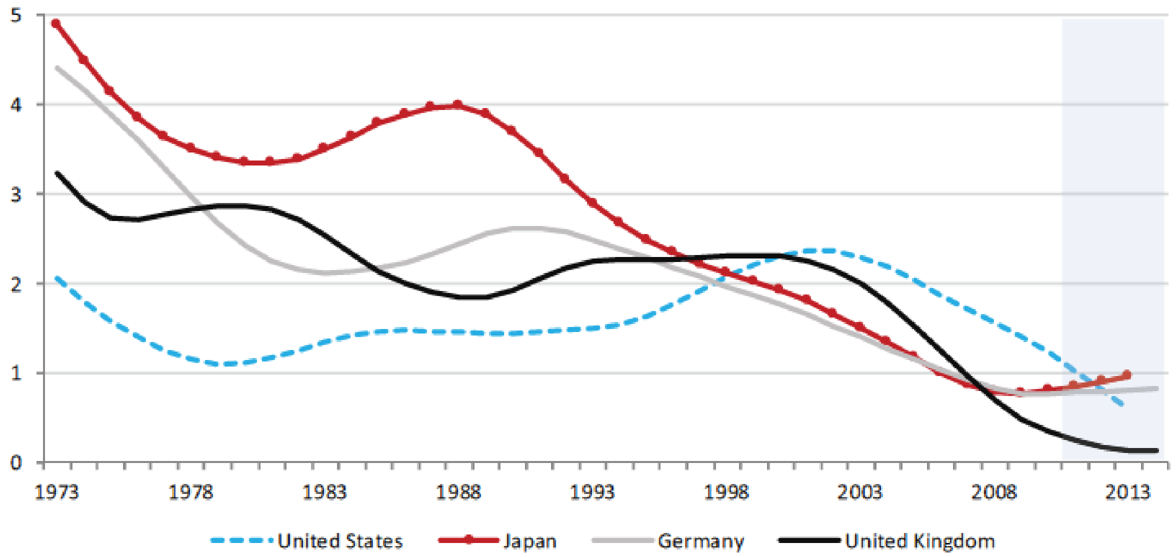

Notes: Labour productivity, total economy, is defined as GDP per hour worked and its growth rate is calculated as first natural-log difference. Trends are estimated using the Hodrick-Prescott filter.
Source: OECD Productivity Statistics (database), April 2017.

Figure 2: G7 real sovereign bond yield (1985–2016)

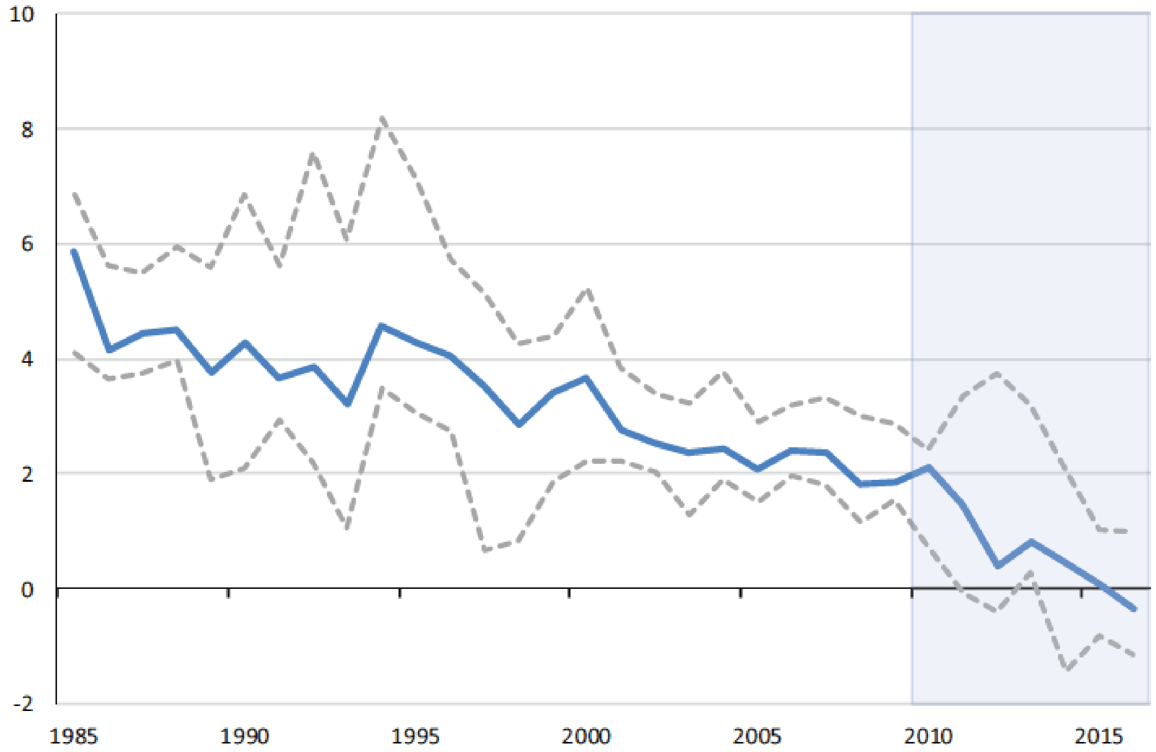

Notes: Real 10-year government bond yield, G7 weighted by GDP. Dashed lines show the minimum and the maximum.
Source: OECD.

Figure 3: Capital deepening: growth in capital services per hour worked (1985–2015)

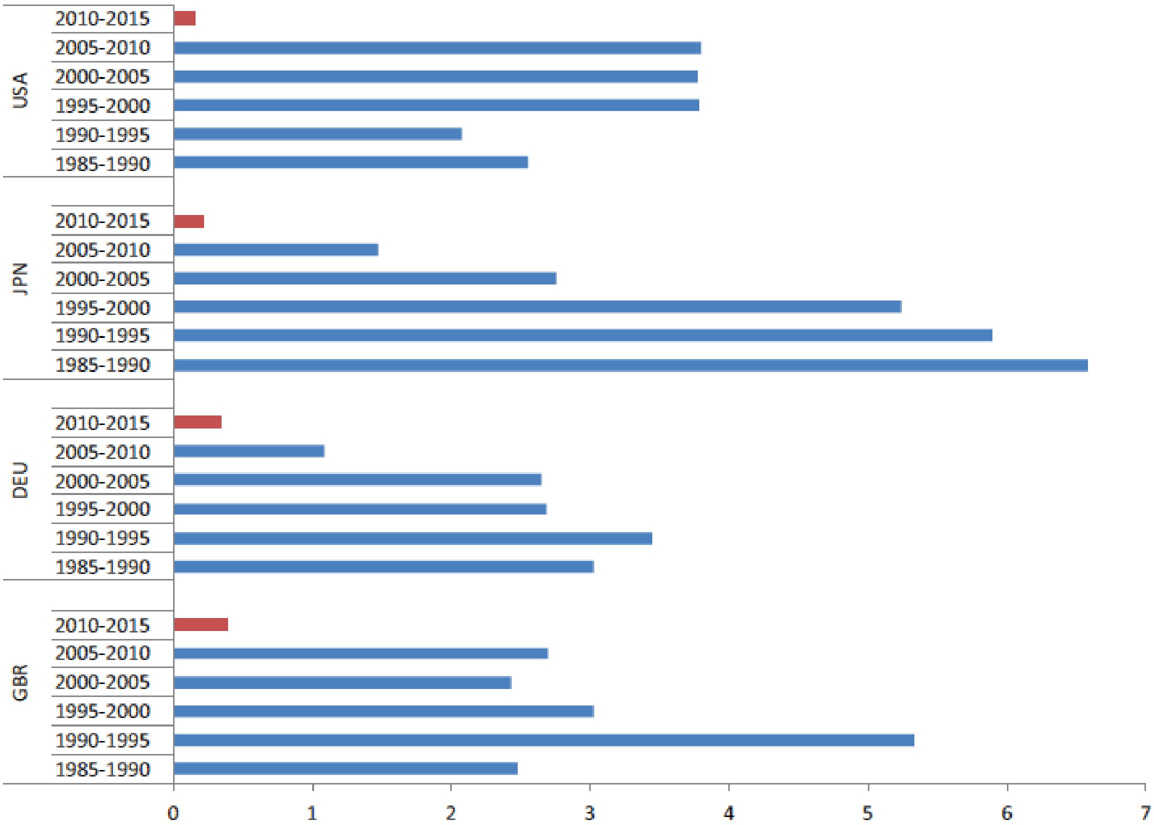

Notes: Total economy, percentage change at annual rate and percentage of GDP.
Source: OECD Productivity Statistics (database), April 2017.

Figure 4: Real compensation per hour worked, deflated by CPI (1995–2016)

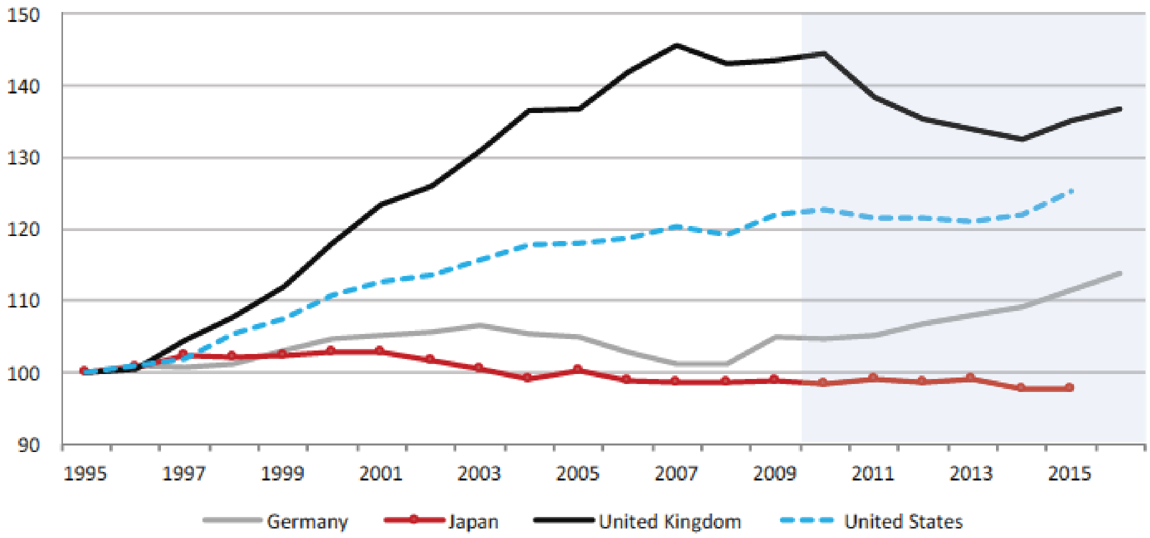

Dataset: LFS.
Source: OECD.stat, 2017.

Figure 5: Unemployment and employment rates, %, 3-year moving average (1986–2015)

(a) Unemployment rate

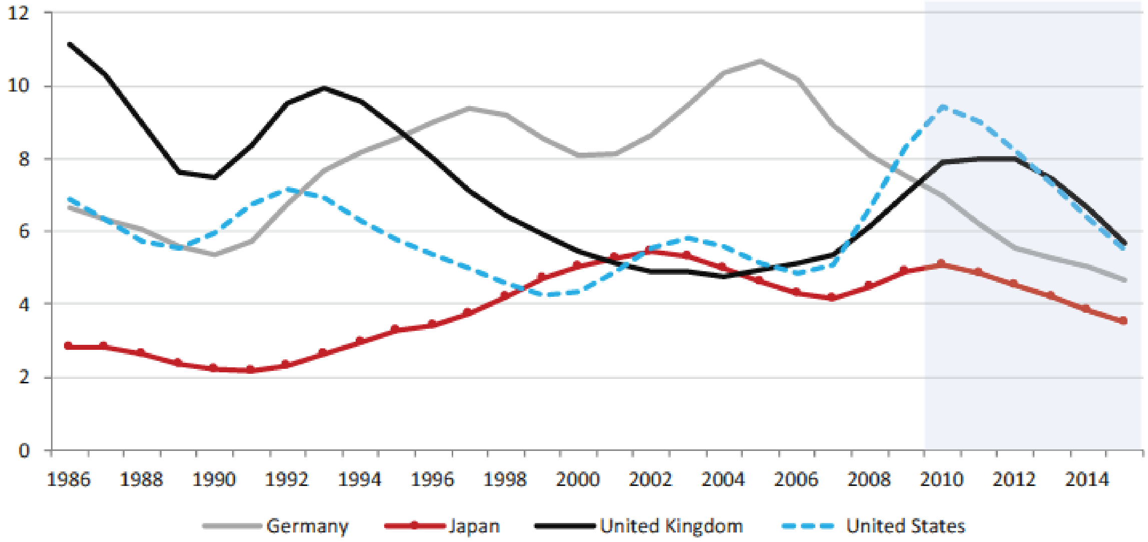

(b) Employment rate

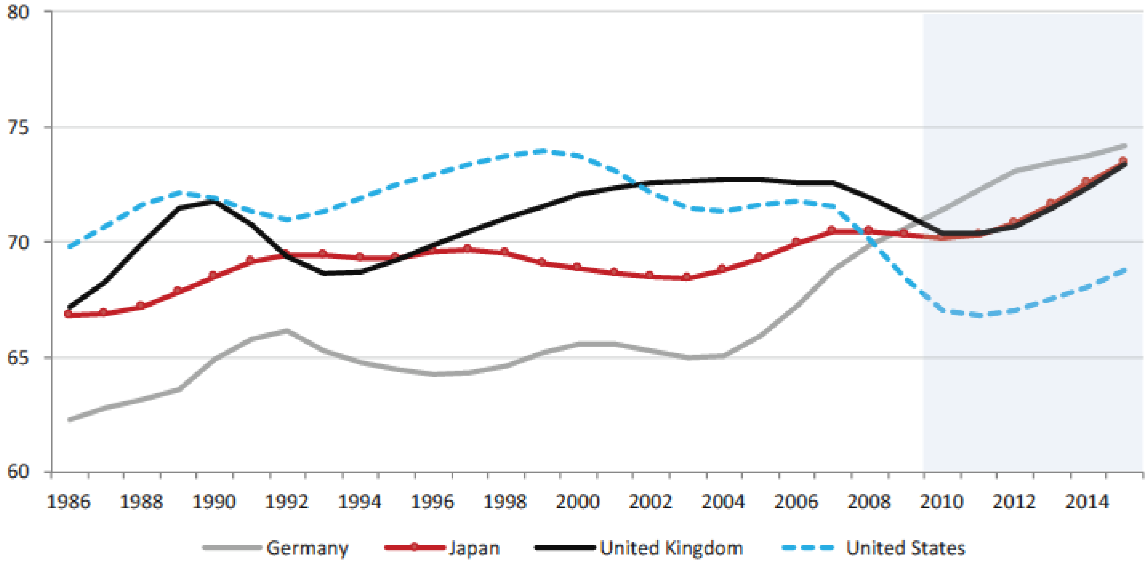

Dataset: LFS.
Source: OECD.stat, 2017.

Figure 6: Nominal compensation per hour worked (1997–2015), total economy, % p.a.

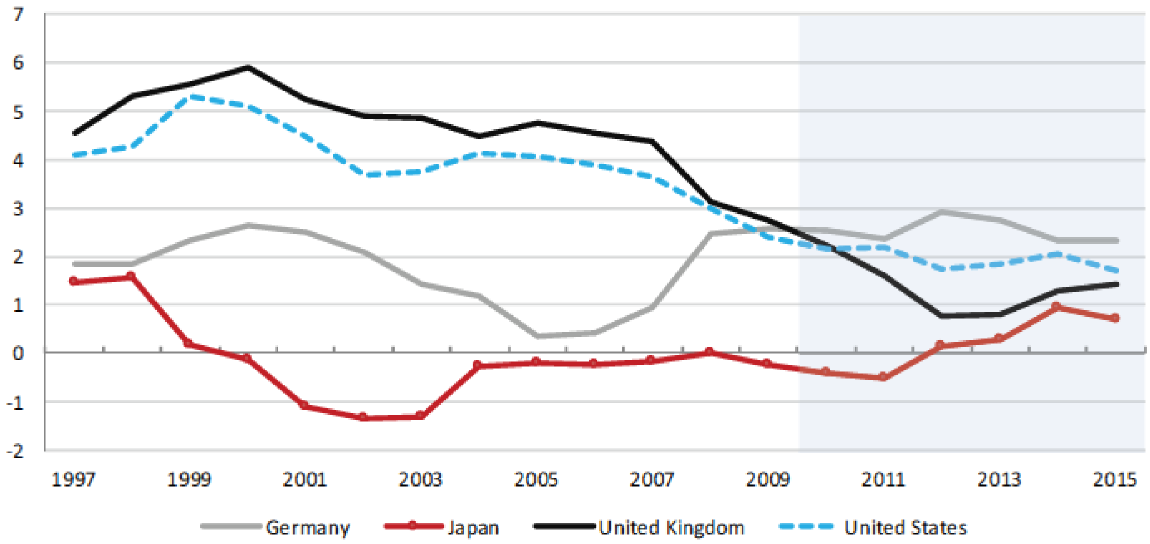

Source: LFS, OECD.stat 2017.

Figure 7: Core CPI inflation (1985–2016)

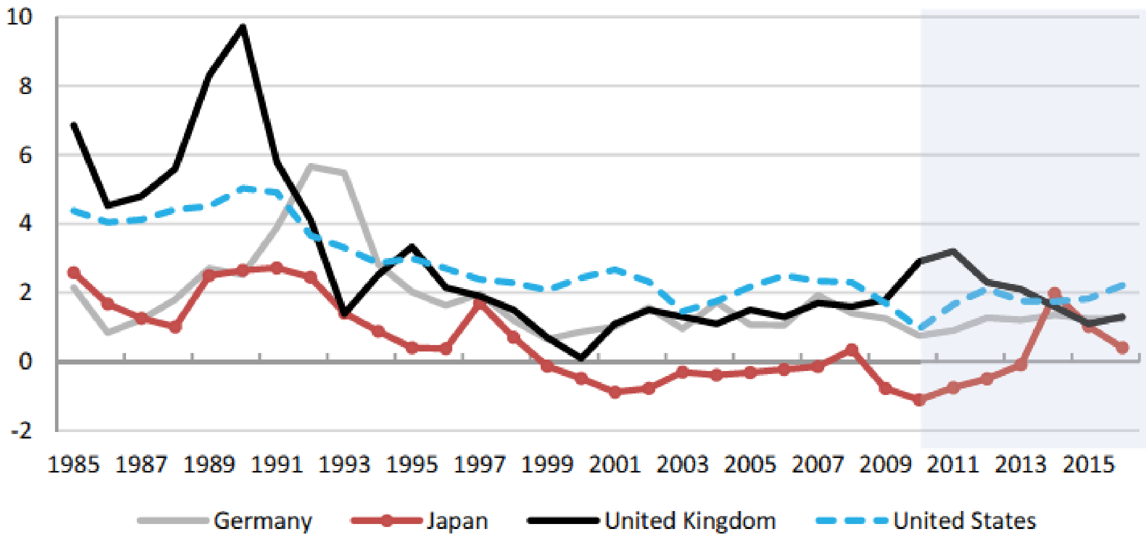

Source: OECD.stat 2017.

III. The two-part two equilibrium model

To provide a framework for interpreting the stylized facts of the period since the financial crisis presented in section II, we set out a small macro model with two equilibria—a Keynesian one and a Wicksellian one. By including a second type of equilibrium, this model extends the closed economy three-equation model (Carlin and Soskice, 2005, 2006, 2015). We begin with an intuitive explanation of the two equilibria and then show how the model is represented graphically.

The standard components of the model are the *IS* curve, where aggregate demand is a function of the real interest rate, the Phillips curve derived from the interaction of aggregate demand and wage and price-setting (*PC*), and the central bank's monetary policy reaction function (*MR*). Underlying the Phillips curve are the price-setting real wage curve, which is the real wage implied by profit-maximizing price setting, and the wage-setting real wage curve (resulting from wage-setting in a labour market with incomplete contracts). Equilibrium employment requires that the price-setting and wage-setting real wages are equal. At equilibrium, inflation is constant and there is involuntary unemployment. The central bank has an inflation target and implements its policy by affecting the real interest rate and hence aggregate demand. Finally, the equilibrium is characterized by a real interest rate determined by the *IS* relation at equilibrium employment.

A minimal amount of information about the microeconomic foundations of the model is required in order to understand the way it works. This is developed further in section VI. Firms produce differentiated products and set prices. In the standard case (Dixit and Stiglitz, 1977), consumers maximize CES utility functions implying a constant elasticity demand function, ε, and constant returns with productivity λ, hence the price-setting real wage is $w^{PS} = \lambda \frac{\varepsilon - 1}{\varepsilon}$. This holds in both the Keynesian and Wicksellian parts of the model.

The wage-setting real wage differs in the two parts. In the Wicksellian case the wage-setting real wage w^{WS} can be seen as an explicit or implicit bargain or agreement in which w^{WS} is higher the higher the rate of employment, *n*, and the higher is productivity λ, scaled by α as a measure of the bargaining effectiveness of the labour force. Thus $w^{WS} = \alpha\lambda n$.

Hence, equilibrium employment in the Wicksellian case is given by

$$w^{W,e} = w^{PS} = \lambda \frac{\varepsilon - 1}{\varepsilon} = w^{WS} = \alpha\lambda n^{W,e}$$

$$\rightarrow n^{W,e} = \frac{1}{\alpha\left(\frac{\varepsilon}{\varepsilon - 1}\right)}$$

(where the interpretation is that equilibrium employment is reduced by low product market competition, and stronger labour market bargaining power α).

In the Keynesian case the key assumption is that nominal wage inflation is bounded by a floor around zero, which captures the ability of employees to coordinate on opposition to nominal wage cuts. For simplicity, we assume the floor is at zero. The desire of wage-setters to avoid the negotiating costs involved in a rational expectations-based Phillips curve explains the widespread use of compensation for previous inflation in wage-setting, and we make that assumption, referring to the 'compensation' Phillips curve.

Equilibrium employment in the Keynesian state, absent nominal wage cuts, is that real wages remain at $w^{WS} = w^{PS} = \lambda \frac{\varepsilon - 1}{\varepsilon}$ whenever $n < n^{K,e}$. Hence $n^{K,e} = n^d$, where n^d is the employment rate determined by the level of aggregate demand.

Thus there are two states of the economy with employment at equilibrium and two associated macroeconomic policy regimes.

In one of the equilibria, the inflation rate is equal to the central bank's inflation target. This is the policy-maker's intended equilibrium. The real interest rate at the intended equilibrium is variously referred to as the stabilizing, natural, or Wicksellian interest rate. For this reason, we call this the Wicksellian equilibrium of the model.

In the second equilibrium, unemployment is higher than at the Wicksellian equilibrium. The real interest rate and inflation are both zero. Whereas employment in the Wicksellian equilibrium is determined in the labour market by the unique intersection of the wage- and price-setting curves, in the second equilibrium, it is determined by the level of aggregate demand. For this reason, it is called the Keynesian equilibrium of the model. Although determined by aggregate demand, inflation is constant at the equilibrium so it must also be a labour market equilibrium where wage- and price-setting real wages coincide.

Before setting out the model graphically, it is useful to contrast the role of aggregate demand in defining the Keynesian equilibrium with the way fluctuations in aggregate demand play out in the neighbourhood around the Wicksellian equilibrium. Beginning at a Wicksellian equilibrium with inflation at target, a negative shock to aggregate demand leads to a fall in employment and output. Given lags in policy-making, inflation falls and pulls forecast inflation below the target. The central bank takes the new compensation-based expectations-augmented Phillips curve as its constraint and optimizes by setting the interest rate (lower) using its monetary rule to boost aggregate demand and output. Inflation rises and the economy returns to equilibrium. If the shock to demand is permanent, the new equilibrium will have a lower stabilizing real interest rate, but all other characteristics of the equilibrium, including the unemployment rate, are unchanged.[3]

Now consider a large negative demand shock. Output, employment, and inflation fall, and the central bank seeks to stabilize as above. However, it is prevented by the ZLB on the nominal interest rate from creating the positive output gap required to get the economy on a path back to its intended equilibrium. Given the persistence of a negative output gap, inflation would be

[3] This is set out in detail in Carlin and Soskice (2005) and in CS2015 ch. 3.

expected to fall, taking the economy further away from equilibrium. Falling inflation pushes up the real interest rate, which depresses aggregate demand further. The dynamics of the model suggest an unstable path with growing deflation and rising unemployment.[4] But this does not match the data well. Indeed, there appears to be a floor to inflation close to zero. In Japan, inflation has fluctuated around zero for almost 20 years.

If we assume a ZLB not only to the nominal interest rate but also to inflation, then instead of a deflationary spiral, the economy will settle at equilibrium with the real and nominal interest rates, and inflation all equal to zero. Maximum employment will be determined by aggregate demand when the real interest rate is equal to zero. Once the economy is at the ZLB for wage inflation, then there is no negative bargaining gap (where the wage on the wage-setting curve is below that on the price-setting curve) because inflation cannot fall. This explains why there is a *labour market equilibrium* associated with low aggregate demand when the economy is at the ZLB for the nominal interest rate and inflation. As in the Wicksellian equilibrium, there is involuntary unemployment. Unlike the standard business cycle fluctuation away from the Wicksellian equilibrium discussed above that produces a rise in *cyclical* unemployment to which the central bank can successfully respond, the additional unemployment in the Keynesian equilibrium is '*structural'*.

We now set out the model in graphical form. There is more detail on the assumptions about the microeconomic mechanisms in section VI. In the lower panel of Figure 8 is the labour market diagram, with the wage- and price-setting curves. The intersection of the two solid curves determines the equilibrium employment level in the Wicksellian equilibrium at $n^{W,e}$, labelled I. This implies a vertical 'long-run' Phillips curve in the middle panel, indicating that there is no long-run trade-off between inflation and unemployment. The downward-sloping MR_1 curve maps the tangencies between ellipsoid central bank indifference curves centred on the equilibrium labelled I and positively sloped 'short-run' expectations Phillips curves (not shown). It indicates the central bank's best response to any deviation in inflation from target: if inflation rises above target, the central bank uses the MR_1 at that rate of inflation to identify the negative output gap with employment below $n^{W,e}$ that will put it on the path back to target inflation. The central bank uses the IS curve to calculate the change in the policy interest rate that will produce the required output gap. In equilibrium, the economy is at the intersection of the MR_1 and PC curves. The real interest rate in equilibrium depends on the position of the IS curve: with IS_1, the stabilizing real interest rate is r_1^W.[5]

The regime in which there is a Keynesian rather than a Wicksellian equilibrium is defined by the ZLB in the nominal interest rate and in inflation. In Figure 8, the Keynesian equilibrium is labelled II with employment at $n^{K,e}$ below the Wicksellian level. Note that IS_2 intersects the horizontal axis to the left of $\underline{IS}$ and hence at employment below the minimum Wicksellian equilibrium employment level, which is defined by a real interest rate of zero.

[4] See CS2015, pp. 104–7.
[5] This is set out in detail in Carlin and Soskice (2005) and in CS2015 chs 3 and 12.

The contrast between the two regimes is shown in Figure 8:

- In the top panel, the IS_2 curve is to the left of the grey $\underline{IS}$. The IS curve labelled $\underline{IS}$ marks the boundary between the Wicksellian and Keynesian regimes.
- In the middle panel, the ('long-run') PC is horizontal at the ZLB of inflation, $ZLB_{\pi=0}$ and vertical in the Wicksellian regime.
- In the bottom panel, the wage-setting curve, $w_{II}^{WS} = w_1$ is horizontal when $ZLB_\pi = 0$. In the Wicksellian regime, the labour market equilibrium is at the intersection of the upward-sloping WS curve and the horizontal PS curve.

Figure 8: Wicksellian and Keynesian equilibria

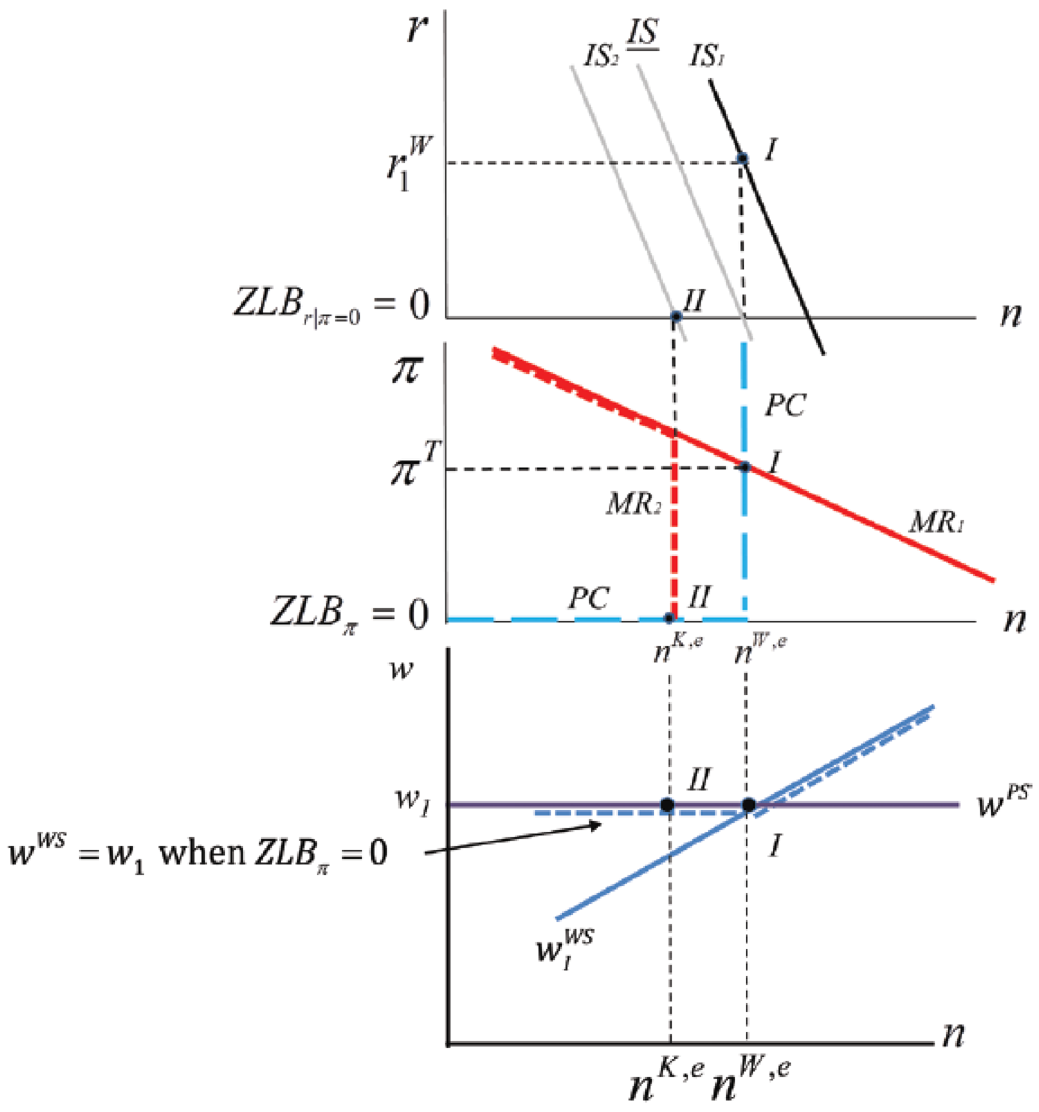

To show graphically how an economy gets stuck in a Keynesian equilibrium we look at the way the central bank's ability to respond to a large negative output gap is constrained by the ZLB on the nominal interest rate.

In Figure 9, we show how the economy adjusts from an initial Wicksellian equilibrium, I, to the Keynesian equilibrium II. Initially, the negative demand shock leads to a fall in output and employment (A to B). From the lower panel, we can see that the ('short-run') Phillips curve for expected inflation at the target of 2 per cent implies deflation at the lower employment level (point C). However, by assumption, inflation cannot fall below zero, shown by the dashed part of the Phillips curve beneath the horizontal axis. Given expected inflation of zero, the central bank's constraint is the Phillips curve indexed by $\pi^E = 0$: the central bank wishes to shift the economy next period to point D on the MR_1 curve. But this is not feasible because of the ZLB on the real interest rate: the economy will settle at point E with inflation, nominal interest rate, and real interest rate equal to zero. The economy is in a Keynesian equilibrium.

Figure 9: A large negative demand shock results in the economy settling at the Keynesian equilibrium

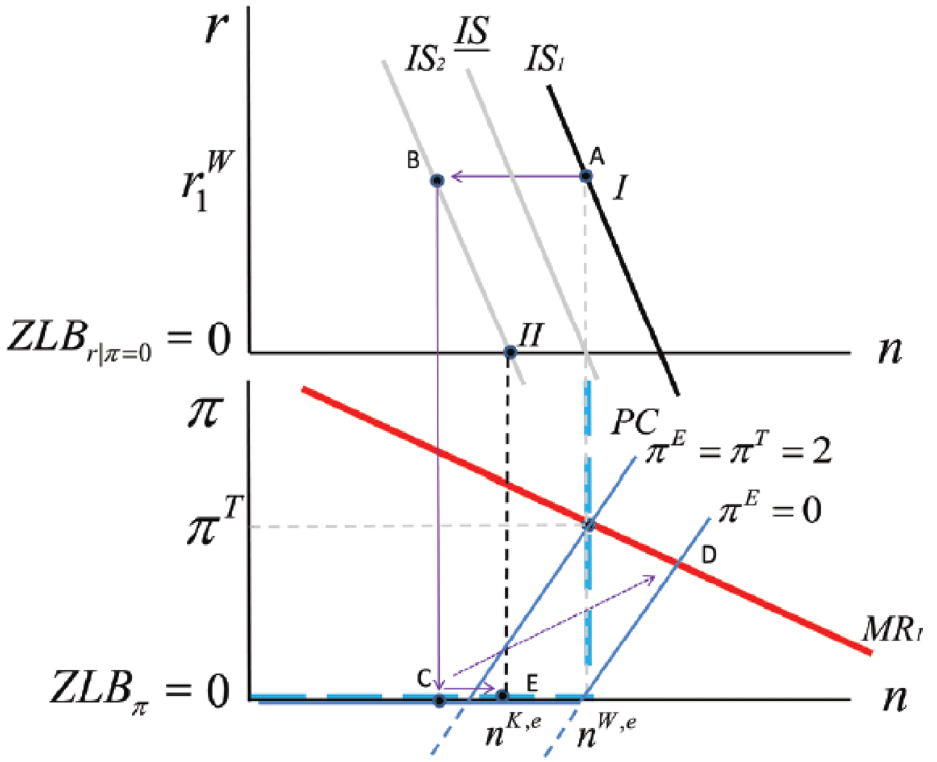

To conclude this section, we draw out the implications for macroeconomic policy in the two regimes. In the neighbourhood of the Wicksellian equilibrium, when the economy is disturbed by demand, supply, or inflation shocks, the central bank uses changes in the interest rate to alter aggregate demand so as to keep the economy close to the inflation target at equilibrium unemployment. Given the policy-maker's objectives, monetary policy is both necessary and sufficient as stabilizer. The equilibrium fiscal multiplier is zero because equilibrium output and employment are determined by the supply side. Higher government spending will *ceteris paribus* result in a higher stabilizing real interest rate; otherwise, the equilibrium is unaffected.

Monetary policy is ineffective in the Keynesian equilibrium. An expansionary fiscal policy sufficient to raise inflation above the target (point C in Figure 10) will take the economy out of the Keynesian equilibrium and restore the role of monetary policy. With inflation above target, the central bank will raise the nominal interest rate to get the economy on to the MR_1 curve to the left of the Wicksellian equilibrium (at point D). The usual adjustment process of a falling interest rate and rising employment will take the economy back to the Wicksellian equilibrium along the MR_1 curve at $n^{W,e}$.

Figure 10: Fiscal policy can be effective in a Keynesian equilibrium

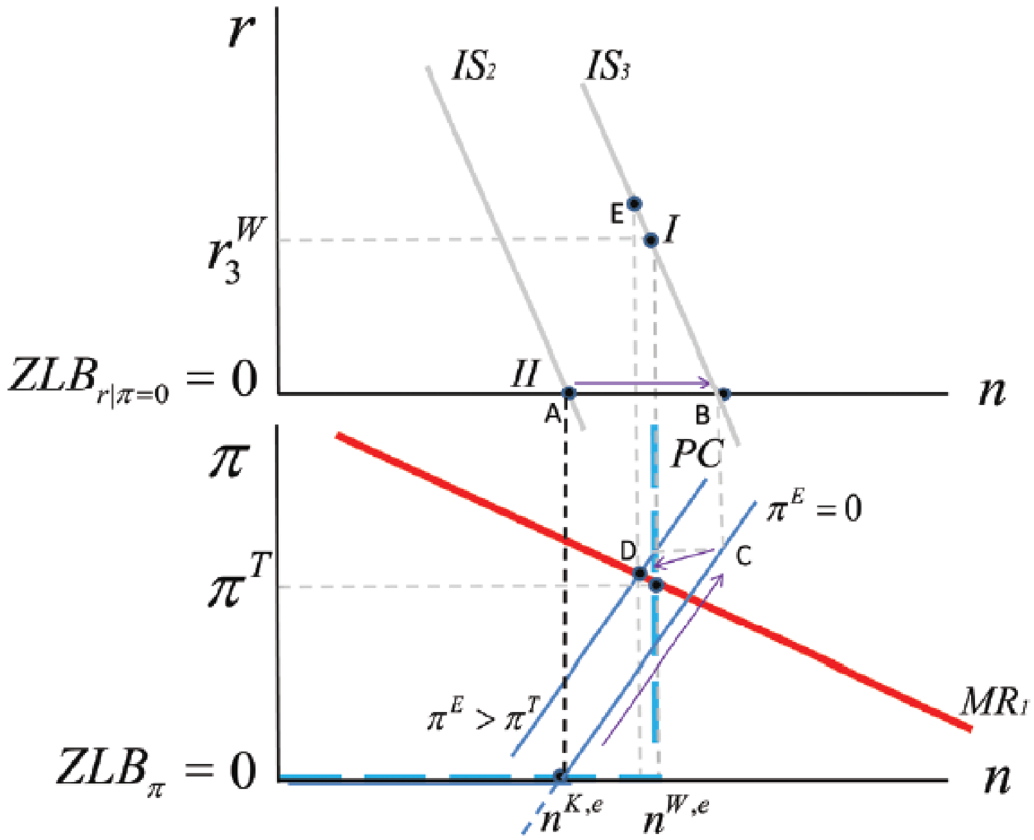

But as we will see in the next two sections, a revival of aggregate demand—either via fiscal policy or via consumption—may not be sufficient to release the economy from a Keynesian equilibrium, even if it is able to deliver low unemployment. To explain this we need to introduce productivity growth and its relationship to investment (in section IV) and in section V, to show how a Keynesian equilibrium can persist at low unemployment.

IV. Low investment and stagnant productivity in the Keynesian equilibrium

In this section we look at the role of aggregate demand, and in particular at investment and productivity growth. In our model of two equilibria, investment plays a major role. We come back to consumption at the end of the section.

As the data in section II make clear, a specific characteristic of the post-financial crisis period is that productivity growth has virtually ceased in the UK and is unusually low in all four countries. It is also the case that the growth of capital services per worker is exceptionally low. Our hypothesis is that low productivity growth is a consequence of low investment (which entails a model in which productivity growth is endogenous), and that persistent low investment is the outcome of coordinated beliefs of firms around pessimistic expectations of market growth.

Caballero documents the lumpiness of investment in firm-level data, which is consistent with its sunk-cost character (Caballero, 1999; Bloom *et al*., 2007; Bloom, 2014). Because of the limited second-hand markets in capital goods, net investment decisions are largely irreversible and are made when management has strong and positive views about the future. In a Keynesian equilibrium, in the absence of strongly held optimistic expectations, firms invest at a low level because the potential loss if expectations are mistaken is low. The *i*th firm's investment will depend on its degree of confidence in the expected growth of its market. Bloom (2014) argues that uncertainty is strongly positively related to recession.

The return to investment to each firm depends on aggregate output. Thus, firms are influenced by the actions of other firms and a natural way to model the problem is as a game of strategic complementarities among firms where there are two stable equilibria and an unstable one (Vives, 2005). The stable equilibria are a low investment level characterized by high uncertainty and pessimistic expectations, and a high investment level associated with high certainty and optimism.

In modelling expectations, we refer to Tuckett's psychological behaviouralist approach where formally independent decision-makers are either in a (radically) uncertain state or coordinate on an optimistic state, which is believed in with high probability (Tuckett, 2012; Tuckett and Nikolic, 2017; King, 2016). We refer to this as the *Tuckett hypothesis.* Without specifying exactly how optimistic expectations come about, individual decision-makers will adopt optimistic expectations simply based on shared optimism (perhaps airing a shared argument or narrative—'this time is different', for example). According to Tuckett, it is not that individual decision-makers look to see the weakness or lack of rationality of the shared optimism, so much as their concern to ensure the optimism is indeed shared. This then has the effect *ex post* that decision-makers do take the 'right' decisions jointly that enable the high investment equilibrium.

Bringing together the assumption that firm *i*'s expectations of its market growth depend on aggregate demand in the economy (consumption plus the investment of all other firms) with the Tuckett hypothesis about the role of uncertainty, provides a model of two stable investment equilibria in which the switch from one to the other occurs because of changes in beliefs in a setting of radical uncertainty. By radical uncertainty, we simply mean to signal the difference between a situation where probabilistic decisions are possible and one in which they are not (King, 2016).

This model is represented graphically in Figure 11. On the horizontal axis is aggregate investment per worker for all but the *i*th firm; on the vertical is investment by the *i*th firm. A symmetric equilibrium exists along the 45-degree line. The S-shaped function indicates that there are two stable equilibria (at low and high investment). At the low one, high uncertainty and pessimism keep investment in the rest of the economy—and therefore in the *i*th firm—low.

At the high one, certainty and optimism generate high investment in all the *j* firms—to which the *i*th firm's best response is to invest at a high level.

Figure 11: Investment behaviour in a game of strategic complementarity, and with radical uncertainty

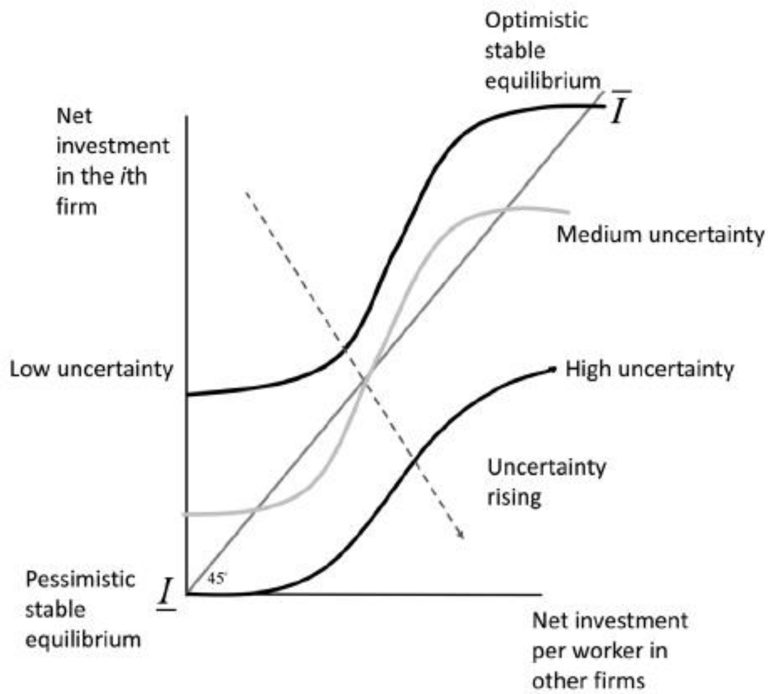

From the geometry in Figure 11, when the S-shaped function cuts the 45-degree line from above, the equilibrium is stable: for example, the equilibria ($\underline{I}, \overline{I}$) are both stable. When the function cuts the 45-degree line from below, the equilibrium is unstable: following a disturbance, investment will converge to either a low or a high equilibrium.

Vives's strategic complementarities game captures the slope and existence conditions for the two stable and one unstable equilibria noted above, and nicely characterizes the difference between stable Tuckett 'pessimistic' low expectations at low levels of investment, $\underline{I}$, and stable 'optimistic' high expectations with high investment, $\overline{I}$, with rapidly changing expectations in between as the *i*th and the other firms make major upwards expectational adjustments.

Using the diagram we can provide an account of how the economy initially trapped in the lower stable equilibrium, jumps to a high investment equilibrium. As uncertainty falls, the S-shaped function shifts up. If uncertainty continues to fall, the S-shaped function will eventually be tangential to the 45-degree line—a further fall will see the lower equilibrium disappear and investment will rise until the optimistic stable equilibrium is reached with investment at $\overline{I}$.

This picture is obviously greatly simplified. But it enables us to see how a move from a high-level coordinating investment-equilibrium can occur in 'one jump' to a low level equilibrium. This might plausibly be as a result of a financial crisis, either directly or indirectly. An historically significant event such as the financial crisis is a commonly observed signal to which many can respond with changing beliefs about macroeconomic uncertainty. And given the initial conditions of low inflation and the small gap between the prevailing interest rate and the ZLB, this in turn moves the economy from a Wicksellian to a Keynesian equilibrium.

Since it is typically the case that much new technology and hence productivity growth is embedded in new investment, a move to a low investment equilibrium is a likely explanation for a period of very low productivity growth. Productivity remains stagnant in the Keynesian equilibrium, with the lower bound of productivity growth at zero, because investment is at a very low level.

The market growth expectations of firms are affected not only by the investment decisions of other firms but also by consumption spending. Because of incomplete insurance and credit markets, workers in this model self-insure through precautionary savings to cover periods of unemployment rather than buying insurance or being able to borrow when unemployed. This reflects asymmetric information by banks (and insurance companies) in the face of idiosyncratic risk, and moral hazard in the face of aggregate unemployment, which is characteristic of the Wicksellian as well as of the Keynesian equilibrium.

Instead of saving primarily for 'smoothing' motivations (as in the standard Euler equation approach based on complete credit markets and insurable unemployment risk), many households are 'hand to mouth' consumers who spend what they earn. Rather than spending on the basis of future income expectations, spending is based on past income. In a recessionary environment household expenditure is based on (low) past income.

Moreover, if households are nervous about the possibility of future unemployment, financial institutions are unlikely to be prepared to lend to those who have become unemployed. Instead as Challe and Ragot argue, households have to self-insure rather than borrow; thus households save in response to recessionary fears (Challe and Ragot, 2016).

What our model shows is that there may be many Keynesian employment equilibria, below the unique Wicksellian equilibrium. As an example, take a simple aggregate demand equation, with 'hand to mouth' consumption, and profits saved, and with a low investment equilibrium given by $I = \underline{I}_{Low}$, and no other component of aggregate demand. Consumption is equal to $\lambda \frac{\varepsilon - 1}{\varepsilon} n$, the price-setting real wage multiplied by *n*. Hence $\lambda n_d = \lambda \frac{\varepsilon - 1}{\varepsilon} n_d + \underline{I}_{Low} \rightarrow n^{K,e} = n_d = \frac{\varepsilon}{\lambda} \underline{I}_{Low}$ is a Keynesian equilibrium. If $\underline{I}_{Low}$ changes because of coordinated changes in expectations, then a new Keynesian equilibrium emerges.

In a much more sophisticated model, Benigno and Fornaro's paper on stagnation traps (2016) also sets out a dual Keynesian–Wicksellian framework with broadly similar assumptions. Their concern is to analyse growth and cycles together within an explicit Aghion–Howitt type endogenous growth framework. The major difference in the results is that in Benigno and Fornaro's model both the Wicksellian and the Keynesian employment equilibria are unique, whereas we emphasize the range of possible Keynesian equilibria once the economy is in the stagnation trap.

In terms of our simplified approach, in the Keynesian part of the model, a Keynesian equilibrium is given by $n^{K,e} = \frac{\varepsilon}{\lambda}\underline{I}_{Low}$ and there may be many such equilibria reflecting different values of $\underline{I}_{Low}$.

By contrast, in Benigno and Fornaro's model, the low investment level is uniquely pinned down with the level of employment in the Keynesian equilibrium: $\underline{I}_{Low} = \underline{I}_{Low}(n^{E})$. In this case the expected value of the employment rate, n^{E}, will be equal—at least eventually—to the actual value of the employment rate. Hence

$$n^{K,e} = \frac{\varepsilon}{\lambda}\underline{I}_{Low}(n^{K,e})$$

so that there is a unique Keynesian equilibrium.

In both models, the Wicksellian equilibrium is unique, $n^{W,e}$, so that $n^{W,e} = \frac{\varepsilon}{\lambda}\overline{I}_{high}(n^{W,e}, r^{W})$. Thus the major difference in the characterization of the stagnation trap between the Benigno and Fornaro model and the approach set out here is the uniqueness or non-uniqueness of the Keynesian equilibrium.

In both cases if productivity growth is embedded in investment, it is plausible to assume that (very) slow productivity growth is associated with the Keynesian (stagnant) equilibrium. And a jump up to a Wicksellian equilibrium, again with productivity growth embedded in investment, now implies a jump to a higher level of productivity growth and in consequence growth in real wages.

While our approach is less sophisticated than that of Benigno and Fornaro, and while the latter usefully integrates stagnation traps with endogenous growth theory, the broadly similar result is a contrast between perhaps long-lasting periods of unemployment with very slow productivity growth and moves (which may be quite rapid) to Wicksellian equilibria with more normal productivity growth and relatively high employment rates. Wicksellian equilibria may also be long-lasting (although subject to destabilizing financial cycles).

V. A Keynesian equilibrium at high employment

The two-equilibrium model can account for a lengthy period during which monetary policy is paralysed and unemployment is high. Yet the data in section II suggest that although several features of the Keynesian equilibrium continue to persist in the UK and elsewhere (very low real interest rates, low inflation, and ineffective monetary policy), unemployment rates are now low and employment rates are at historically high levels.

One hypothesis for the UK's performance consistent with the model is that the weakness of investment persists, which is reflected in Figures 1 (productivity) and 3 (capital deepening), but that aggregate demand has revived through consumption spending. Although real wages have not been rising, households have had access to cheap credit, especially for cars and consumer durables, in the low interest rate environment. This has shifted the *IS* curve to the right.

Increased aggregate demand in the context of zero productivity growth (a consequence of the continued weakness of investment) implies higher employment. According to Figure 5, the employment rate is higher than it was prior to the financial crisis (in three of the four countries). We need to explain how employment rose above the pre-crisis Wicksellian equilibrium level without triggering inflation—and, therefore, to explain how a Keynesian equilibrium with the real interest rate and inflation close to zero can co-exist with low unemployment.

In graphical terms, this requires a downward shift in the wage-setting curve, which in the context of the nominal interest rate at the ZLB and wage inflation at zero, implies the extension of the flat part of the wage-setting curve to the level of employment consistent with a new lower unemployment rate at a Wicksellian equilibrium. Figure 12 illustrates. The new Wicksellian equilibrium is marked by *I* and the new Keynesian equilibrium at low unemployment where the economy is trapped is labelled *II* .

Figure 12: A Keynesian equilibrium at low unemployment

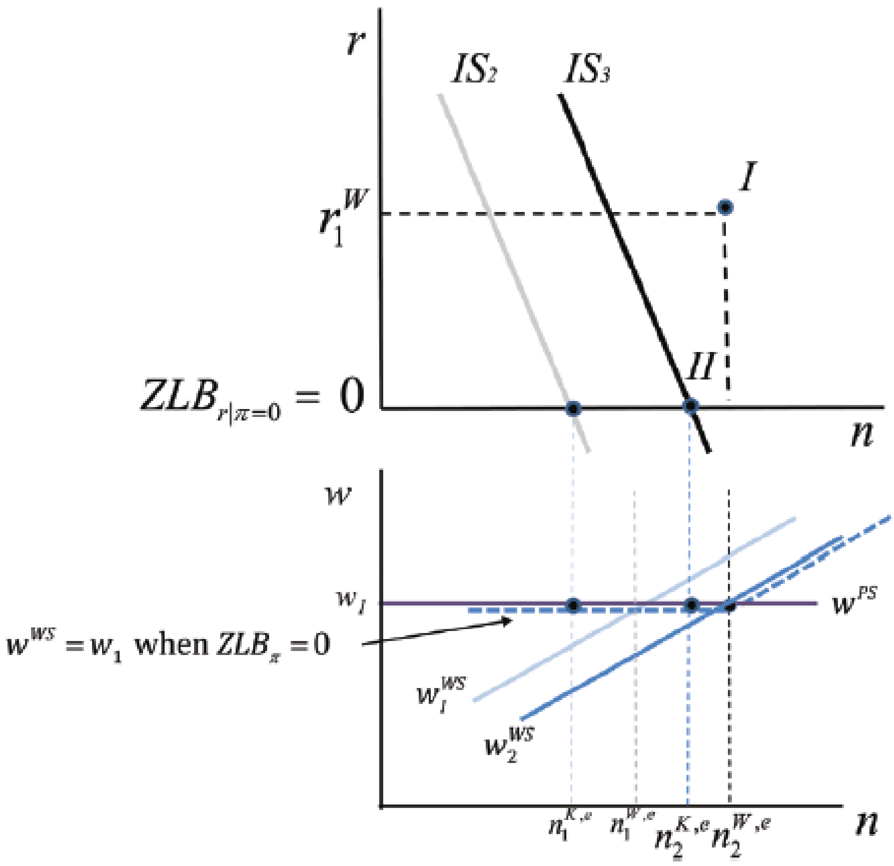

The graphical illustration of the high employment Keynesian equilibrium begs the question of why the wage-setting curve should have shifted downwards. Anything that weakens the reservation position of employees shifts the curve down: lower unemployment benefits, less legal protection for workers, weaker unions, and increased labour supply (for example, as a result of migration) can all increase the expected cost of job loss. New forms of contracts, including the so-called zero hours contracts in the UK, have increased job insecurity. All of these developments, some of which are associated with the continued shift of employment out of industry and with new forms of work organization facilitated by new technology, reduce the wage the employer needs to pay in order to secure adequate worker effort at any given unemployment rate.

VI. Microeconomic mechanisms underlying the two-equilibrium model: comparison with standard New Keynesian and heterogeneous agent (HANK) models

In this section we contrast the mechanisms sustaining the Keynesian and Wicksellian equilibria in the small model we have set out with those in two other models. The first is the widely cited three-equation representative agent New Keynesian model of Clarida *et al*. (1999) and the second is a new small model motivated by the empirical issues of the post-crisis stagnation (section II), which is called a heterogeneous agent New Keynesian model, HANK (Ravn and Sterk, 2016). All three models can be presented graphically using IS, PC, and MR-type curves, although as we shall see, the underpinnings of the curves reflect quite different assumptions about the behaviour of the economic actors. This comparison will help to clarify that although ours is a heterogeneous agent model with multiple equilibria like Ravn and Sterk's model, it is not 'New Keynesian'. We call it HAWK for short—heterogeneous agent, Wicksellian-Keynesian model. Neither the Clarida *et al*. nor the Ravn and Sterk model includes investment.

Unlike real business cycle and New Keynesian models, our model does not impose model-consistent expectations on all agents. Model-consistency of expectations implies that it is common knowledge that all relevant agents have full knowledge of the whole model and know that everyone else does. In particular, we part company with the New Keynesian modelling of the Phillips curve. The requirement of model-consistency of expectations combined with sticky prices implies that with a New Keynesian Phillips curve new information leads to jumps in inflation, which are counter-factual.[6] Some agents—such as those in the foreign exchange market and Central Banks—do form model-consistent expectations *given the expectations formed by others*.[7]

[6] As an example of the concern of policy-makers, the brief for the Bank of England's 2016 review of its in-house New Keynesian model says: 'More generally on expectations, the rational expectations assumption within COMPASS produces responses that too often appear implausible (large and front-loaded).' They call for 'an alternative expectations scheme that was well tested, well understood and gave rise to plausible impulse responses'.

[7] See ch. 16 of CS2015 for a detailed comparison of the CS model and the New Keynesian dynamic stochastic general equilibrium model.

Since the actors in our model set wages, prices, output, and interest rates, we require that a relation in the model describing how *x* takes decisions such as setting wages, or prices, or output, or the interest rate should be comprehensible to the actor in question. Few price-setters would understand Calvo price-setting, let alone see this as approximating how prices are in fact set. The requirement in NK DSGE (New Keynesian dynamic stochastic general equilibrium) models for wage and price variables to jump on to a perfect foresight path runs counter to our requirement for comprehensibility of actions to actors. Whereas in a price-taking model, comprehensibility is not relevant, it is important in a model where the strategies of actors are central.

As in our previous work, the microfoundations of the modelling in our two-equilibrium HAWK model are post-Walrasian in the sense of Stiglitz (1993), Bowles and Gintis (2000), and Howitt (2006). Principal–agent relationships arising from asymmetric information problems characterize interactions in the labour market between workers and employers, and in the credit market between borrowers and lenders. A macroeconomic model defined by principal–agent problems is by definition a heterogeneous agent model and, because information problems make it impossible to write complete labour or credit market contracts, there is no frictionless benchmark. Unemployment and credit rationing are present in any equilibrium.

Before drawing the contrasts with the two New Keynesian models, we summarize the assumptions of our model.

(i) Key assumptions

The model is of a closed economy. The ten key features of our HAWK model (some shared by Clarida *et al*. and Ravn and Sterk as explained below) are that:

i. in the Blanchard–Kiyotaki (1987) tradition, there are Keynesian microfoundations with firms setting prices and producing what is demanded at that price based on Dixit–Stiglitz (1977) monopolistic competition; this applies to the Wicksellian as well as the Keynesian part of the model (as in Clarida *et al*. and Ravn and Sterk);

ii. it is a heterogeneous agent model, with employed and unemployed workers; and there is a distinct class of equity owners of firms, so that the effect of income inequality can be understood; we assume all profits are saved along Kalecki–Pasinetti lines; (as in Ravn and Sterk);

iii. with the exception of the Central Bank, we assume agents do not form model-consistent expectations about the model as a whole; this is on empirical grounds and distinguishes the model from New Keynesian ones;

iv. asymmetric information and moral hazard imply (*a*) incomplete loan and insurance markets—generating credit constraints and precautionary savings (as in Ravn and Sterk and Challe and Ragot (2016)); and (*b*) involuntary unemployment at any labour market equilibrium;

v. there are no constraints on the timing of price-setting, but incomplete contracts and 'influence costs' (Milgrom and Roberts, 1990) concerning employees' productivity contributions make periodic wage-setting efficient; this is moreover overwhelmingly the case in practice;

vi. because model-consistent expectations are ruled out, the desire to avoid the negotiating costs involved in an expectations-based Phillips curve explains the widespread use of compensation for previous inflation in wage-setting, and we make that assumption, referring to the compensation Phillips curve;

vii. in line with much empirical evidence there is *de facto* (something like) a ZLB in nominal wage inflation in wage-setting; a simple analytic explanation may be that nominal wage-cutting is often highly visible across groups of employees: it thus constitutes a focal point in a coordination game in which employees can choose whether or not they feel well-disposed towards the employer in their effort decision; in any case we assume that wage-inflation has a ZLB;

viii. we assume that investment decisions by firms are actions in a multiple equilibrium dynamic game of strategic complementarities with stable optimistic and pessimistic choices (Vives, 2005; Caballero, 1999);

ix. while oversimplified, we assume productivity is largely embodied in the capital stock so that low investment implies static productivity and, given the profit mark-up, static real wages; there is no exogenous technical progress;

x. given the presence of an inflation-targeting central bank (which imposes a ceiling on inflation), we assume a ZLB on the nominal rate of interest; with a ZLB on nominal wage and price inflation, there is a non-negative lower bound on the real rate of interest.

The two parts of the model can be summarized in the following equations where the superscript *ps* refers to precautionary savings:

(1) The Keynesian part:

$$w = \frac{\lambda(\varepsilon - 1)}{\varepsilon} = w^{K,e}$$

$$\pi \approx 0$$

$$r \approx 0$$

$$c = nw(1 - s^{ps}(n))$$

$$I = \underline{I}_{Low}$$

$$y = c + I + g$$

$$n = y / \lambda$$

$$\dot{\lambda} \approx 0$$

$$n^{K,e} = n < \frac{\varepsilon - 1}{\alpha\varepsilon}$$

Note that, ignoring precautionary savings:

$$
\begin{aligned}
& y = wn_d + \underline{I}_{Low} + g \\
& \rightarrow \lambda n_d = \lambda \frac{\varepsilon - 1}{\varepsilon} n_d + \underline{I}_{Low} + g \\
& \rightarrow n^{K,e} = n_d = \left(\frac{\varepsilon}{\lambda}\right)\left(\underline{I}_{Low} + g\right) \\
& \rightarrow y = \lambda n_d = \varepsilon\left(\underline{I}_{Low} + g\right) \\
& \rightarrow \Delta y = \varepsilon \Delta\left(\underline{I}_{Low} + g\right)
\end{aligned}
$$

which takes us from one Keynesian equilibrium to another.

(2) The Wicksellian part:

$$
\begin{aligned}
& w = \frac{\lambda(\varepsilon - 1)}{\varepsilon} = w^{W,e} \\
& \pi = \pi_{-1} + \alpha(n - n^{W,e}) \geq 0 \\
& r = r^W + \tau(\pi - \pi^T) \geq 0 \\
& c = wn \\
& y = c + \overline{I}_{High}(r) + g \\
& n = y / \lambda \\
& n^{W,e} = \frac{\varepsilon - 1}{\alpha\varepsilon}
\end{aligned}
$$

The equilibrium multiplier is zero since $y^{W,e} = ln^{W,e}$.

(ii) Comparison with the New Keynesian models

The representative agent New Keynesian model (Clarida *et al*., 1999) has a transparent modelling of central bank behaviour that relates to how central banks describe their deliberations and actions.[8] However, when subjected to shocks, the model incorporates the unappealing 'jumping' behaviour of inflation characteristic of the New Keynesian Phillips curve.[9] In the Clarida *et al*. model, there is a single equilibrium (type I) at the policy-maker's inflation target so it is not a useful model for addressing the problems of a low growth, low inflation equilibrium. By assuming a single representative agent, it omits key features of the world that we are seeking to understand.

[8] For an example, see Mark Carney's speech January 2017 at the LSE entitled 'Lambda' (Carney, 2017).
[9] The characteristics of the New Keynesian Phillips curve are explained in ch. 16 of CS2015.

In response to these challenges, in the last few years models with heterogeneous agents have been developed (e.g. Kaplan *et al*. (2016) and the many papers cited there, and by Ravn and Sterk, 2016, 2017). In many such models consumption behaviour varies according to the employment status of the agent. This is often tied in with an extension to the search and matching labour market model (SAM) in which the assumption that there are complete insurance markets for unemployment risk is dropped, with the result that unemployment becomes a major concern.

In these models, new mechanisms can lead to large fluctuations in aggregate demand and employment. They emerge from the joint presence of a precautionary savings motive, unemployment risk, and nominal rigidities:

> intuitively, a wave of pessimism among households about their employment prospects could be self-fulfilling as the increased desire to build precautionary savings reduces aggregate demand, causing firms to hire fewer workers when prices are sticky and stabilization policy is insufficiently responsive. (Ravn and Sterk, 2016, p. 2).

When these features are brought into a typical quantitative NK DSGE model, it can only be solved numerically. From the perspective of a policy-maker, this is problematic because the modeller is not able to make the causal mechanisms and policy transmission channels transparent, and the presence and nature of multiple equilibria cannot be rigorously analysed (Ravn and Sterk, 2016).

To deal with these difficulties, Ravn and Sterk (2016) set up a small-scale analytical version of a HANK–SAM model with incomplete markets, search and matching frictions, and nominal rigidities. It is a model with three agents/classes in equilibrium: unemployed, employed, and owners/entrepreneurs. Their paper delivers a model that can be represented graphically and the qualitative predictions compared with the model of Clarida *et al*. (1999). Neither model includes investment.

In Ravn and Sterk's model there are three equilibria—the central bank's intended equilibrium, a second, liquidity trap equilibrium at the ZLB, and a third 'unemployment trap' equilibrium with low aggregate demand, high unemployment, and low inflation, but at an interest rate above the ZLB. Positive feedbacks between unemployment and precautionary savings reinforce the unemployment trap.

We simplify the Ravn–Sterk model by concentrating on the HANK element, abstracting from the details of the search and matching model of the labour market. We explain how a simplified version of the Ravn–Sterk model works and then compare it with the two-equilibrium HAWK model presented earlier in this paper. This helps to bring out important features of the two models as they relate to the modelling of consumption, the Phillips curve, and the monetary policy rule.

The Ravn–Sterk model is a major departure from the benchmark representative agent New Keynesian model with its unique equilibrium at target inflation and labour market clearing.

Ravn and Sterk's model has an inflation-targeting central bank, but there is involuntary unemployment at the intended equilibrium, and there are 'belief-driven' fluctuations so the economy can become trapped in equilibria characterized by very low inflation and high unemployment.

The motivation for Ravn–Sterk's model is very similar to that for the two-equilibrium model presented above: to match the post-crisis reality where low inflation and high unemployment persisted for a number of years. And, as we shall see, the interaction between the risk of unemployment and the motive for precautionary saving plays a role in both models (there is no investment in their model). In their words:

> The unemployment trap that can arise . . . offers an alternative perspective of secular stagnation which ties together low real interest rates, high unemployment and low activity. Moreover, because of the low job-finding rate, there is a strong incentive for precautionary savings which drives down the real interest rate. Intriguingly, the unemployment trap can occur in our model purely because of expectations and thus does not rely on sudden changes in population growth, technological progress or financial tightening. (Ravn and Sterk, 2016, p.17)

The Ravn and Sterk model is shown in the figure below, where a measure of labour market tightness is on the horizontal axis and the inflation rate is on the vertical. The diagram captures the three equations in the model (the 'IS', the 'PC', and the 'MR'). Ravn and Sterk combine the IS and monetary rule into a single equation called EE.

Figure 13: Ravn–Sterk model with incomplete markets and sticky prices; three equilibria

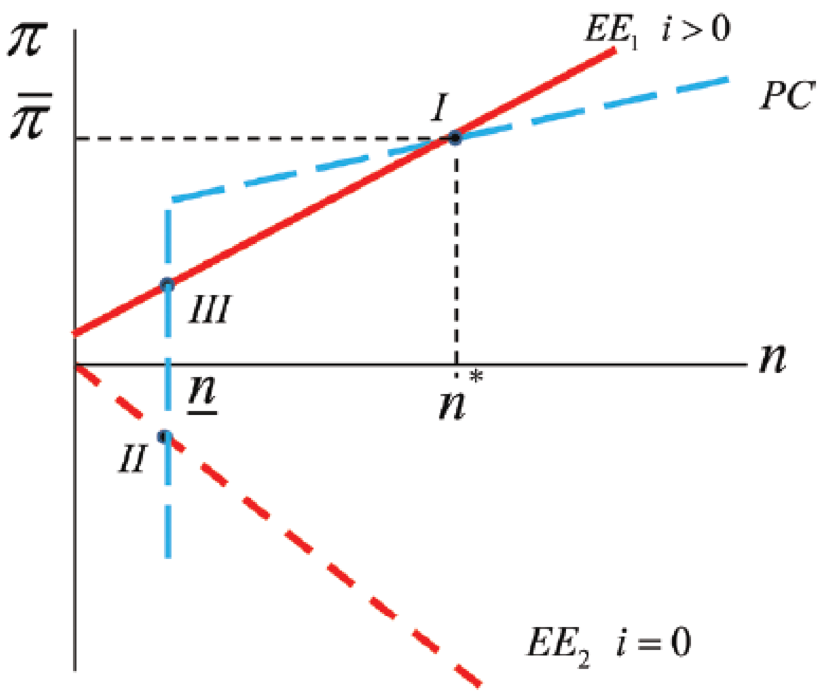

We compare Figure 13 with Figure 8, which shows our two-equilibrium Wicksellian–Keynesian model.

The first observation is that there are multiple equilibria in both models where the PC and EE (or MR) curves intersect. In each model, the equilibrium intended by the policy-maker is labelled I. The Clarida *et al*. model has a single equilibrium, which is the one intended by the policy-maker. But in both Ravn and Sterk and our model, there is a second equilibrium labelled II at which monetary policy is ineffective when the interest rate is at the ZLB. The Ravn and Sterk model has a third equilibrium, III, which is called the unemployment trap, where inflation is low but the economy is not constrained by the ZLB on the nominal interest rate.

Looking at the geometry of Figures 8 and 13, it is evident that the structural relationships summarized in the Phillips curve and the EE (Monetary policy / IS) curve that produce the two equilibria (I and II) are very different. When comparing the two figures, the striking differences are in the slopes of the PC and EE curves. Notice first that the Phillips Curve in the Ravn–Sterk model is very flat—it must be flatter than the EE curve to produce the third unemployment trap equilibrium (III), and second, that the EE curve is positively sloped in Ravn–Sterk and negatively sloped in our HAWK model. In Ravn and Sterk's model, the EE curve is negatively sloped only in the special case where the economy is at the ZLB.

The models have similar objectives (tractability, transparent mechanisms, relevance to contemporary conditions) and superficially similar results: an equilibrium at the policy-maker's intended outcome and a ZLB equilibrium. Both incorporate an inflation-targeting policy-maker and heterogeneous agents so as to include precautionary savings, uninsured unemployment risk, and feedback at high unemployment between precautionary savings and unemployment. Yet the nature of the key relationships tying together the supply side (PC), the demand side (IS), and the policy-maker (MR/EE) are quite different as signalled by the contrasting geometry in the figures.

We next provide some intuition behind the differently sloped Phillips and EE/MR curves in the two models. In the online Appendix, we set out the equations of a simplified version of the Ravn and Sterk model in order to explain in more detail the sources of the differences.

Three key modelling choices lie behind the differences between the Phillips curve and EE/MR relations in the Ravn and Sterk and our HAWK model illustrated in Figures 8 and 13.

Ravn and Sterk assume:

- high price adjustment costs in equilibrium, which is necessary to flatten the 'long run' Phillips Curve and permit the third equilibrium to arise;
- there is no instrument for precautionary saving (in spite of the presence of the precautionary motive for saving), which makes the EE curve positively sloped; and
- a model of central bank behaviour without a specific inflation target. If there is an inflation target, then the third 'unemployment trap' equilibrium disappears.

The relatively flat equilibrium Phillips curve

The reason for the upward-sloping PC in the Ravn and Sterk model is their assumption that more rapid adjustment of prices is costly. As we show in the Appendix, dropping the terms in the Rotemberg price adjustment cost from the Ravn–Sterk first order condition produces a wage-setting real wage (as a function of *n*) equal to the price-setting real wage at a unique degree of labour market tightness, just as in our HAWK model. This results in a vertical PC relation (adding inflation adjustment costs to the HAWK model, similarly produces an upward-sloping PC).

In the steady-state equilibrium, the rise in inflation permits the real wage to increase and the exactly corresponding fall in the real profit mark-up enables the competing claims on output of wages and profits to be met and thus preserves constant inflation equilibrium in the conditions of a tighter labour market. It is not clear what justifies the assumption of price adjustment costs following wage changes, and certainly why they should be so high as to flatten the Phillips curve to the degree necessary.

The positively sloped EE curve
This curve incorporates the IS relation along with the central bank's monetary rule. Recall that in our model, the equivalent to this curve is downward sloping: when inflation is judged too high by the central bank, it will tighten monetary policy by raising the interest rate according to the IS relation in order to create a negative output gap so as to reduce inflation. In the Ravn and Sterk model, the curve is upward-sloping. This means that although higher inflation makes the policy-maker want to reduce the output gap in order to dampen inflation (by raising the interest rate in the normal way), the expectation of this policy response will only be consistent with the consumption behaviour of employed workers if employment and the wage are in fact higher.

To explain this, we need to look more closely at an assumption that appears to be made for technical reasons. We find it has strong implications (the argument is set out using equations in the Appendix). The assumption is that although employed workers have a precautionary savings motive, they cannot implement it because there is no savings instrument. Ruling out saving implies that those in employment must be so-called 'hand to mouth', which means that their consumption each period is equal to their wage. And yet, employed workers are assumed also to satisfy the Euler equation, which means they are making their optimal consumption choice. To bring their desired consumption dictated by the Euler condition into line with the wage (hand to mouth consumption) requires that when inflation is higher and the central bank is expected to raise the interest rate, employment and the wage must be higher. This produces combinations of higher inflation and higher employment that satisfy both the *IS* relation and the monetary rule. This is not an intuitively appealing mechanism.

The third 'unemployment trap' equilibrium
In the simplified version of the Ravn and Sterk model set out in the Appendix, we find that when central bank behaviour is adjusted to bring it into line with the more standard Clarida *et al*. approach with an inflation target chosen by the policy-maker, the specific equilibrium that their paper prioritizes (the unemployment trap, III) disappears because the EE curve becomes downward-sloping as it is in our HAWK model.

VII. Conclusion

We set out a small analytical model that has recognizable central-bank and private-sector behaviour and includes heterogeneous agents and models of consumption and investment that are consistent with two equilibria: the intended one and the one at the ZLB. The ZLB equilibrium is characterized by a lower bound of zero on wage inflation and productivity growth as well as on the nominal interest rate.

This small model provides a way of accounting for the persistence of the long post-crisis stagnation—with neither recovery to the neighbourhood of the normal Wicksellian equilibrium nor a destabilizing deflationary process. Getting stuck in the stable Keynesian regime is due to the combination of low net investment (keeping productivity growth close to the ZLB), low aggregate demand (rendering monetary policy impotent because of the ZLB on the nominal interest rate), and the functioning of the labour market that sets a lower bound at zero on wage inflation, and prevents a deflationary spiral.

We attach particular importance to investment as a key driver, with investment modelled as a game of strategic complementarities and low and high coordinating equilibria. Once in a low equilibrium, the economy can get caught in a prolonged period of stagnation. In addition, productivity growth is embedded at least in part in investment: hence investment-induced stagnation can tie down productivity growth to very low levels.

In the context of this model, the more recent phenomenon, most starkly present in the UK, of the combination of low and stable inflation, a very low real interest rate, and productivity growth close to zero with low unemployment can be accounted for by developments on both the supply and demand side. In particular, it is consistent with, on the one hand, a weakening of the bargaining power of labour (due to policy, organizational and technological changes), and on the other, to a revival of consumer demand based on the extension of consumer credit in the conditions of low interest rates. Weaker worker bargaining power implies lower unemployment at a new Wicksellian equilibrium and extends the zone of the Keynesian ZLB trap to higher levels of employment, without upward pressure on inflation. More buoyant consumer demand under the ZLB conditions expands employment. However, without a shift to the optimistic scenario about future growth, investment and productivity growth do not revive.

In the Wicksellian regime, monetary policy is effective in stabilizing supply, demand, and inflation shocks—additional policies are required to prevent a financial cycle-induced crisis. In the Keynesian regime, monetary policy is ineffective in restoring employment to the Wicksellian equilibrium. Fiscal policy may be effective in enabling the economy to escape the stagnation trap, but for a return to normal conditions of positive productivity growth, a reduction of uncertainty about expected market growth is necessary to revive investment and, with it, productivity and real wage growth.

References

Bank of England (2016), 'The COMPASS and Suite Transition (CAST) Review Brief', Simon Hayes, cited by Muellbauer (2016).

Benigno, G., and Fornaro, L. (2016), 'Stagnation Traps', CFM DP 2016-06.

Blanchard, O. J., and Kiyotaki, N. (1987), 'Monopolistic Competition and the Effects of Aggregate Demand', *American Economic Review*, **77**(4), 647–66.

Bloom, N. (2014), 'Fluctuations in Uncertainty', *Journal of Economic Perspectives*, **28**(2), 153–75.

— Bond, S., and van Reenen, J. (2007), 'Uncertainty and Investment Dynamics', *Review of Economic Studies*, **74**(2), 391–415.

Bowles, S., and Gintis, H. (2000), 'Walrasian Economics in Retrospect', *Quarterly Journal of Economics*, **115**(4), 1411–39.

Caballero, R. J. (1999), 'Aggregate Investment', *Handbook of Macroeconomics*, **1**, 813–62.

Carlin, W., and Soskice, D. W. (1990), *Macroeconomics and the Wage Bargain*, Oxford, Oxford University Press.

— — (2005), 'The 3-equation New Keynesian Model—A Graphical Exposition', *The B.E. Journal of Macroeconomics*, **5**(1), 1–38.

— — (2006), *Macroeconomics: Imperfections, Institutions and Policies*, Oxford, Oxford University Press.

— — (2015), *Macroeconomics: Institutions, Instability and the Financial System*, Oxford, Oxford University Press.

Carney, M. (2017), 'Lambda', speech given by Mark Carney, Governor of the Bank of England, at the London School of Economics.

Challe, E., and Ragot, X. (2016), 'Precautionary Saving Over the Business Cycle', *The Economic Journal*, **126**(590), 135–64

Clarida, R., Galí, J., and Gertler, M. (1999), 'The Science of Monetary Policy: A New Keynesian Perspective', *Journal of Economic Literature*, **37**(4), 1661–707.

Dixit, A. K., and Stiglitz, J. E. (1977), 'Monopolistic Competition and Optimum Product Diversity', *American Economic Review*, **67**(3), 297–308.

Howitt, P. (2006), 'Monetary Policy and the Limitations of Economic Knowledge', in D. Colander (ed.), *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*, Cambridge, Cambridge University Press.

Kaplan, G., Moll, B., and Violante, G. (2016), 'Monetary Policy According to HANK', NBER Working Paper No. 21897.

King, M. (2016), *The End of Alchemy: Banking, the Global Economy and the Future of Money,* Little Brown.

Milgrom, P., and Roberts, J. (1990), 'Bargaining Costs, Influence Costs, and the Organization of Economic Activity', ch. 3 in J. E. Alt and K. A. Shepsle (eds), *Perspectives on Positive Political Economy,* Cambridge, Cambridge University Press.

Muellbauer, J. (2016), 'The COMPASS and Suite Transition (CAST) Review', Nuffield College, Oxford.

Ragot, X. (2015), 'Le Retour de l'Economie Keynesienne', *Revue de Economie Financiere,* **121**, 173–85.

Ravn, M., and Sterk, V. (2016), 'Macroeconomic Fluctuations with HANK & SAM: An Analytical Approach', CFM DP2016-33.

— — (2017), 'Job Uncertainty and Deep Recessions', *Journal of Monetary Economics,* **90**(October), 125–41.

Reinhart, C. M., and Rogoff, K. (2009), 'The Aftermath of Financial Crises', *American Economic Review,* **99**(2), 466–72.

— — (2010), 'Growth in a Time of Debt', *American Economic Review,* **100**(2), 573–8.

Stiglitz, J. E. (1993), 'Post Walrasian and Post Marxian Economics', *Journal of Economic Perspectives*, **7**(1), 109–14.

Tuckett, D. (2012), 'Financial Markets are Markets in Stories: Some Possible Advantages of Using Interviews to Supplement Existing Economic Data Sources', *Journal of Economic Dynamics and Control,* **36**, 1077–87

— Nikolic, M. (2017), 'The Role of Conviction and Narrative in Decision-making Under Radical Uncertainty', *Theory and Psychology*, **27**(4), 501–23.

Vives, X. (2005), 'Complementarities and Games: New Developments', *Journal of Economic Literature*, **43**(June), 437–79

Stagnant productivity and low unemployment: stuck in a Keynesian equilibrium

Wendy Carlin* and David Soskice**

Appendix

Comparison of the Ravn–Sterk HANK model with the two-equilibrium HAWK model[10]

We derive a simplified version of the Ravn–Sterk HANK model to facilitate comparison of the key relationships with our HAWK model. This parallels the simple comparison in the text of section VI of the paper.

(i) The Phillips Curve relation

In the Ravn and Sterk model in Figure 13, the 'long run' PC is positively sloped and very flat. A positively sloped PC means that there is a range of *equilibrium* combinations of the inflation rate and the unemployment rate at which inflation is constant.

In the Ravn and Sterk model in Figure 13, the positively sloped Phillips Curve is integral to the predictions of the model. It is an equilibrium relationship that results from assuming there are costs of price adjustment. We need to explain both why the long-run relation between the job-finding rate and inflation is positive and why it is flatter, the higher are price adjustment costs.

Intuitively, the argument is as follows: for a given adjustment cost parameter, a higher inflation rate is associated with higher costs. Hence, firms are less inclined to mark up cost increases. The equilibrium real wage therefore lies closer to the wage-setting curve. Since the real wage is higher, the economy can operate at a higher job-finding rate with constant yet higher inflation. Similarly, as price adjustment costs increase, the real wage will lie closer to the wage-setting curve for any given job-finding rate, flattening the equilibrium relationship between inflation and the job-finding rate.

Note that the PC in the Ravn and Sterk model is vertical when the economy is at the ZLB. This arises because $\bar{n}$ is defined as the minimum possible job-finding rate: if n cannot be reduced below $\bar{n}$, then the relative price cannot be increased, which implies that the rate of inflation of the relative price is zero and of the absolute price level is indeterminate. For a job-finding rate above $\bar{n}$ the equilibrium moves to equilibrium I. If $n = \bar{n}$ we remain on the vertical part of *PC*. This is the liquidity trap equilibrium in the Ravn and Sterk model.

The long-run trade-off in the Ravn and Sterk model is reminiscent of the original Phillips Curve relation. The long-run trade-off was replaced in macro models by the vertical Friedman–Phelps-accelerationist Phillips curve from the late 1960s. The HAWK model in Figure 8 shows the standard vertical accelerationist Phillips curve. In the HAWK model, wage- and price-setters do not solve the forward-looking model. Prices respond rapidly to cost shocks and wage-setters cannot bring forward wage adjustments. Nominal stickiness arises from wage- not from

[10] We are very grateful for Bob Rowthorn's help with the Appendix.

price-setting as reflected in the compensation Phillips curve described in sections III and VI. In equilibrium firms are on the price-setting curve achieving the mark-up consistent with competitive conditions. Hence the Phillips curve is vertical at the equilibrium rate of unemployment.

The equations used to derive the Phillips curve in a simplified version of the Ravn–Sterk model are as follows. The economy consists of a large number of identical 'small' firms. The demand for the output of firm *j* is given by:

$$y_j = \left(\frac{P_j}{P}\right)^{-\gamma} y$$

where P_j and y_j are, respectively, the price and output of this firm in the current period, and P and y are the price and output of the average firm. Employment in firm *j* is given by

$$e_j = \frac{y_j}{A}$$

where A is labour productivity. The real wage rate is the same for all firms and is given by:

$$w_j = w(n)$$

where *n* is a measure of the job finding rate and $w'(n) > 0$.

Suppose that the cost of filling vacancies is zero. Then equation (10) of Ravn and Sterk implies that firm *j* chooses P_j so as to maximize:

$$V_j = \left(\frac{P_j}{P}\right) y_j - w_j e_j - \frac{\phi}{2}\left(\frac{P_j - P_{j,-1}}{P_{j,-1}}\right)^2 y - \Lambda \frac{\phi}{2}\left(\frac{P_{j,+1} - P_j}{P_j}\right)^2 y_{+1}$$

where $P_{j,-1}$ and $P_{j,+1}$ are the price of good *j* in the previous and next periods respectively, y_{+1} is average output in the next period, and Λ is a discount factor. The above equation can be expressed as follows:

$$V_j = \left(\frac{P_j}{P}\right)^{1-\gamma} y - \frac{w(n)}{A}\left(\frac{P_j}{P}\right)^{-\gamma} y - \frac{\phi}{2}\left(\frac{P_j - P_{j,-1}}{P_{j,-1}}\right)^2 y - \Lambda \frac{\phi}{2}\left(\frac{P_{j,+1} - P_j}{P_j}\right)^2 y_{+1}$$

The first order condition for a maximum is:

$$\frac{P}{y}\frac{\partial V_j}{\partial P_j} = (1-\gamma)\left(\frac{P_j}{P}\right)^{-\gamma} - \frac{w(n)}{A}\left(\frac{P_j}{P}\right)^{-(1+\gamma)} - \phi\left(\frac{P}{P_{j,-1}}\right)\left(\frac{P_j - P_{j,-1}}{P_{j,-1}}\right) + \Lambda\phi\left(\frac{PP_{j,+1}}{P_j^2}\right)\left(\frac{P_{j,+1} - P_j}{P_j}\right)\frac{y_{+1}}{y}$$
$$= 0$$

In equilibrium, $P_j = P,\ P_{j,-1} = P_{-1},\ P_{j,+1} = P_{+1}$, and the above condition can be expressed thus:

$$(1-\gamma)-\gamma\frac{w(n)}{A}-\phi(\pi+1)\pi+\Lambda\phi(\pi_{+1}+1)\pi_{+1}\frac{y_{+1}}{y}=0$$

where $\pi=\frac{P-P_{-1}}{P_{-1}}$ and $\pi_{+1}=\frac{P_{+1}-P}{P}$ are the rates of inflation in the current and next periods.

Rearranging:

$$\Lambda\phi(\pi_{+1}+1)\pi_{+1}=\frac{y_{+1}}{y}\phi(\pi+1)\pi-(1-\gamma)+\gamma\frac{w(n)}{A}$$

which can be readily compared with the Phillips curve in our HAWK model:

$$\pi=\pi_{-1}+\alpha(y-y^e) \qquad \text{PC (HAWK model)}$$

.

In the steady state $\pi_{-1}=\pi$ and the HAWK PC is vertical since $y=y^e$. In the steady state, it is also the case that $\pi_{+1}=\pi$ and $y_{+1}=y$, and hence the Ravn and Sterk HANK PC is upward-sloping:

$$\phi(1-\Lambda)\pi(\pi+1)=\left(1-\gamma+\gamma w(n)/A\right) \qquad \text{PC (HANK model)}$$

As noted in the text, the reason for the upward-sloping PC in the Ravn and Sterk model is their assumption that faster price increases are costly. Dropping the terms in the Rotemberg price adjustment cost, ϕ, from the RS first order condition produces a wage-setting real wage (as a function of n) equal to the price-setting real wage at a unique degree of labour market tightness $(1-\gamma+\gamma w(n)/A=0)$ just as in our HAWK model. This results in a vertical PC relation.

Similarly, if we add inflation adjustment costs, $C(\pi)$, to our HAWK model, an upward-sloping PC is the outcome.

In the steady-state equilibrium, the rise in inflation permits the real wage—and hence n—to increase. As noted above, the exactly corresponding fall in the real profit mark-up enables the competing claims on output of wages and profits to be met and thus preserves constant inflation equilibrium in the conditions of a tighter labour market.

Returning to the Ravn–Sterk Phillips Curve, we can derive the steady-state Phillips curve relation in a form that highlights the source of the positive slope and its magnitude. The positive root to the above HANK quadratic equation is:

$$\pi=\frac{-1+\sqrt{1+\frac{4\left(1-\gamma+\gamma w(n)/A\right)}{\phi(1-\Lambda)}}}{2} \qquad \text{PC (HANK model)}$$

$$\approx\frac{1-\gamma+\gamma w(n)/A}{\phi(1-\Lambda)}$$

The larger is ϕ, the cost of changing prices, the flatter is the PC.

(ii) The EE curve (aggregate demand and the monetary policy rule)

It is useful to set out the simplified equations to amplify the explanation given graphically and verbally in section III. We begin by writing an IS curve and a TR, a Taylor-type monetary policy rule used by Ravn and Sterk. (Note, however that in the Taylor equation, the central bank's target is the steady state inflation rate, $\bar{\pi}$, rather than a specific target inflation rate, π^T. We return to this issue in (iii) below.)

$$IS: n = A - \alpha(i - \pi)$$
$$TR: i - \pi = [i-\pi]^S + \tau^{\pi}(\pi - \bar{\pi}) + \tau^{n}(n - \bar{n}) \,.$$

By substituting the Taylor rule into the IS, we get the relation called EE:

$$n = A - \alpha\left([i-\pi]^S + \tau^{\pi}(\pi - \bar{\pi}) + \tau^{\eta}(n - \bar{n})\right)$$
$$\rightarrow n(1+\alpha\tau^{\eta}) = A - \alpha\left([i-\pi]^S + \tau^{\pi}(\pi - \bar{\pi}) - \tau^{n}\bar{n}\right)$$
$$n = \frac{A - \alpha([i-\pi]^S - \tau^{\pi}\bar{\pi} - \tau^{n}\bar{n})}{(1+\alpha\tau^{n})} - \frac{\alpha\tau^{\pi}}{(1+\alpha\tau^{n})}\pi = \bar{A} - \tau\pi \qquad \text{EE relation}$$

This produces a downward-sloping EE relation. The intuition follows directly from the IS and the Taylor rule: a rise in inflation leads via the central bank's MR to an increase in the real interest rate about the 'normal' stabilizing or Wicksellian real interest rate ($= r^W = [i-\pi]^S$), which, in turn, reduces aggregate demand and output via the IS relation. A negative output gap is required to bring inflation down to the steady-state rate. And the more inflation-averse is the central bank, the flatter is the EE curve.

But the EE relation in Ravn and Sterk's model is *upward-sloping* outside the ZLB case.

It is simplest to show the origins of the upward-sloping relation if we use a log utility function for consumption. Consider an employed worker with money wage W in the current period and expected money income EZ_{+1} in the next period. This worker chooses c to maximize their expected utility $\log c + \Lambda E \log c_{+1}$ subject to the budget constraint $P_{+1}Ec_{+1} = P_{+1}Ez_{+1} + (1+i)(wP - cP)$. In steady state the (constant) rate of inflation is π and hence $P_{+1} = (1+\pi)P$. Assuming that $E\log c_{+1} \approx \log Ec_{+1}$ the Lagrangian for this problem is:

$$L = \log c + \Lambda \log Ec_{+1} + \mu\left[(1+\pi)PEz_{+1} + (1+i)W - Ec_{+1}(1+\pi)P - (1+i)cP\right]$$

This yields the solution

$$c = \left(\frac{1+\pi}{1+i}\right)\frac{1}{\Lambda}Ec_{+1}$$

$$\approx \frac{(1+\pi-i)}{\Lambda}Ec_{+1} \qquad \text{Euler equation for employed worker in the HANK model}$$

Suppose that at the end of the production period a fraction ω of employees randomly lose their jobs, of whom a fraction n are immediately re-employed in another firm. The risk of an employed worker being unemployed in the next period is equal to $\omega(1-n)$ and hence

$$Ec_{e,+1} = \omega(1-n)c_u + (1-\omega(1-n))c_e$$

The Euler equation is therefore

$$c_e = \frac{(1+\pi-i)}{\Lambda}\left[\omega(1-n)c_u + (1-\omega(1-n))c_e\right]$$

For employed workers in the Ravn–Sterk model, consumption, c_e, is equal to the real wage, *w*, reflecting 'hand-to-mouth' behaviour, and for unemployed workers, consumption is equal to *v*, the amount of home production. Thus:

$$w = \frac{(1+\pi-i)}{\Lambda}\left[\omega(1-n)v + (1-\omega(1-n))w\right].$$

Rearranging yields

$$\frac{1}{(1+\pi-i)} = \frac{1}{\Lambda}\left[\omega(1-n)\frac{v}{w} + 1 - \omega(1-n)\right]$$

which can be approximated as follows

$$i - \pi \approx \frac{1}{\Lambda}\left[\omega(1-n)\frac{v}{w} + 1 - \omega(1-n)\right] - 1.$$

Moreover, the real wage is also a function of the employment rate. Thus,

$$i - \pi \approx \frac{1}{\Lambda}\left[\omega(1-n)\frac{v}{w(n)} + 1 - \omega(1-n)\right] - 1.$$

This can be combined in a monetary policy rule of the following type

$$i-\pi=[i-\pi]^{S}+\tau^{\pi}(\pi-\bar{\pi})+\tau^{n}(n-\bar{n})$$

If the weight on the measure of labour market tightness is zero ($\tau^{n}=0$), the above takes the following form:

$$i-\pi=[i-\pi]^{S}+\tau^{\pi}(\pi-\bar{\pi})$$

Combining this with the Euler equation yields:

$$\tau^{\pi}(\pi-\bar{\pi})\approx\frac{1}{\Lambda}\left[\omega(1-n)\frac{v}{w(n)}+1-\omega(1-n)\right]-1-[i-\pi]^{S}\quad \text{EE relation HANK model}$$

Differentiating:

$$\frac{\tau^{\pi}\Lambda}{\omega}\frac{\partial\pi}{\partial n}=\left(1-\frac{v}{w(n)}\right)-\frac{(1-n)v}{\left(w(n)\right)^{2}}w'(n)$$

The sign of this differential can be positive or negative, depending on the function $w(.)$ and the size of n. For any plausible $w(.)$ the differential will be positive for n sufficiently close to 1. More generally, suppose that $w(n)=a(1-n)^{-\theta}$, where $\theta>0$ and $v>(1+\theta)a$. In this case, the differential is positive for all $n\in[0,1]$.

The positive slope of the EE curve (away from the ZLB), arises because higher inflation arises because of an increase in the job-finding rate (labour market tightness). This is a direct consequence of the assumption that although employed workers have a precautionary savings motive, they cannot implement it because there is no savings instrument. Ruling out saving implies that those in employment are 'hand to mouth', consuming their wage. Imposing this assumption dictates the upward-sloping EE relation. The only way that workers, given their precautionary savings motive, can be on their Euler equation and wish not to save when inflation is higher and the central bank is expected to respond with a higher interest rate, is if employment and the wage are higher.

(iii) The intended equilibrium (I) in the Ravn–Sterk model

Here we examine the role of target inflation in the Ravn–Sterk model and look at the implications of substituting the Taylor-type rule incorporated into their EE relation by a best-response monetary policy rule, MR (as in Clarida *et al*. and in our model).

In the Ravn–Sterk model, inflation in the steady state at the intended equilibrium (I) is defined by whatever is the value of inflation in the steady state. This contrasts with other small models where there is an inflation target chosen by the policy-maker. Let us now assume that in steady-state equilibrium, there is an additional constraint: inflation must be at the inflation target, i.e.

$\bar{\pi} = \pi^T$.

As before, the EE relation in the Ravn–Sterk model is:

$$\tau^{\pi}(\pi - \bar{\pi}) \approx \frac{1}{\Lambda}\left[\omega(1-n)\frac{v}{w(n)} + 1 - \omega(1-n)\right] - 1 - [i - \pi]^S .$$

In the steady-state equilibrium with inflation at the target rate, this can be written as

$$0 = \frac{1}{\Lambda}\left[\omega(1-n)\frac{v}{w(n)} + 1 - \omega(1-n)\right] - 1 - [i - \pi]^S .$$

This implies a specific value of market tightness, n^*, when inflation is on target in the steady state. In turn, via the Euler equation, this pins down the level of aggregate demand (consumption) in the steady state (the vertical dotted line) in Figure 13 consistent with n^*.

Figure A1: A puzzle in the Ravn–Sterk model

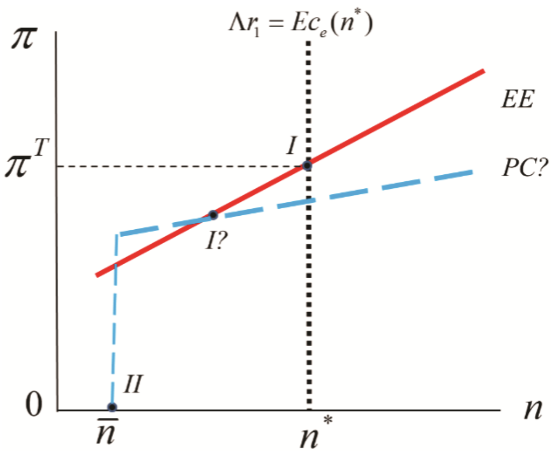

This poses a puzzle: the equilibrium (I) in the RS model is at the intersection of the PC and the EE, yet once we introduce an inflation target, the equilibrium must be on the EE at the combination (n^*, π^T). In terms of the diagram, the PC must move to intersect the EE at this point, suggesting contrary to RS, that the PC and EE are not independent, and not independent of the inflation target.