BANK OF ENGLAND

Staff Working Paper No. 627

Deflation probability and the scope for monetary loosening in the United Kingdom

Alex Haberis, Riccardo M Masolo and Kate Reinold

November 2016

BANK OF ENGLAND

Staff Working Paper No. 627

Deflation probability and the scope for monetary loosening in the United Kingdom

Alex Haberis,[1] Riccardo M Masolo[2] and Kate Reinold[3]

Abstract

In this paper, we use an estimated DSGE model of the UK economy to investigate perceptions of the effectiveness of monetary policy since the onset of the 2007–08 financial crisis in a number of measures of deflation probability — the Survey of Economic Forecasts, financial-market option prices, and the Bank of England's Monetary Policy Committee's (MPC) forecasts. To do so, we use stochastic simulations of the model to generate measures of deflation probability in which the effectiveness of monetary policy to offset deflationary shocks is affected by different assumptions about the existence and level of a lower bound on policy rates. We find that measures of deflation probability are consistent with the perception that the MPC was not particularly constrained in its ability to offset deflation shocks in the post-crisis period.

Key words: Deflation, forecasting and simulation, models and applications, interest rates, monetary policy.

JEL classification: E31, E37, E43, E47, E52.

(1) Bank of England and Centre for Macroeconomics. Email: alex.haberis@bankofengland.co.uk
(2) Bank of England and Centre for Macroeconomics. Email: riccardo.masolo@bankofengland.co.uk
(3) Bank of England. Email: kate.reinold@bankofengland.co.uk

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

In the aftermath of the 2007-08 financial crisis the UK economy fell into its deepest recession for decades, and, although headline CPI inflation increased in response to the large depreciation of sterling, underlying measures of inflation fell back. Consistent with this fall in inflationary pressure, the probability of *de*flation increased across a variety of measures—the Survey of Economic Forecasts,[1] financial-markets measures derived from option prices, and the Bank of England's Monetary Policy Committee's (MPC) own inflation projections. At the same time, the MPC lowered the policy rate in the UK (Bank Rate) to a record-low level of 0.5%. In 2009, the MPC judged that further cuts in Bank Rate would not deliver the desired monetary stimulus and instead turned to the unconventional measure of quantitative easing (QE). As an instrument of monetary policy, QE had not been used in the UK before and there was some uncertainty about its effectiveness. Moreover, there was also uncertainty about the efficacy of conventional monetary policy, in the form of interest-rate cuts, at such low levels of Bank Rate. Therefore, the post-crisis period had been characterised by higher than usual uncertainty about the overall effectiveness of monetary policy. If there were perceptions that monetary policy was constrained in its ability to provide stimulus, that could have contributed to the MPC missing its inflation target in the medium term and fed into weaker real economic outturns (via lower inflation expectations and higher expected real interest rates).

Economic agents' perceptions of any constraints on monetary policy are, by their nature, not observable. Therefore, in this paper, we study the three measures of deflation probability described above to infer what they might imply for perceptions of monetary-policy constraints since the onset of the crisis. We define a deflationary episode as any quarter in which the annual inflation rate is negative. Our approach involves comparing measures of deflation probability derived from stochastic simulations of a DSGE model estimated on UK data with the observed measures of deflation probability. Using the model allows us to produce simulated probabilities of deflation under alternative assumptions about the ability of monetary to offset the effects of deflationary shocks. This means that we are able to provide an economic interpretation to differences in the deflation-probability measures from the simulations and, in turn, use that to make inferences about possible drivers of the observed measures of deflation probability that we consider.

In the simulations, we allow for a lower bound on nominal interest rates to proxy for the effectiveness of monetary policy in offsetting the effects of deflationary shocks. Setting the lower bound to be 0.5%, in line with the MPC's 2009 assessment of the "effective" lower bound (ELB), is consistent with further cuts to Bank Rate not being possible and unconventional policy being entirely ineffective. By contrast, considering simulations in the absence of any form of lower bound is tantamount to assuming that unconventional policies are equivalent to further cuts in short rates, including to well-below zero. In other words, the "shadow short rate"—which translates unconventional policy into its equivalent in short-rate space, as discussed by Krippner (2013)—is unconstrained.

We find that the observed deflation-probability measures do not appear to embody a perception that monetary policy was particularly constrained in its ability to offset the effects of deflationary shocks during the crisis. To be more specific, our simulations identify three episodes of particularly elevated deflation probability since the crisis: in

[1] While the Survey of Economic Forecasts only included an option for deflation from 2009, the probability of inflation of less than 1% fell markedly relative to pre-crisis.

the immediate aftermath of the crisis in 2009, in 2012-3, and uring 2015. The causes of the heightened deflation probability appear to differ across the episodes. In the 2015 and 2009 episodes, the weakness of the outlook for inflation is looks to be the primary factor behind the elevated deflation probability. In the case of 2012-13, by contrast, inflation was above target, and it was lack of scope for monetary policy to loosen that generated a high probability of deflation. In the observed deflation-probability measures, the episodes of heightened deflation probability that our simulation results suggest are primarily due to a weak inflation outlook—in 2009 and in 2015—are present. But the 2012-13 episode, which is driven by a limited scope for policy easing in our simulation results, is not. As such, our simulation results are consistent with the observed deflation-probability measures reflecting a relative absence of perceptions that monetary policy was particularly constrained over this period.

To produce the simulations, we consider a standard piecewise-linear framework. This assumes that there are two regimes characterised by linearity—one in which the ELB binds and one in which it is slack—but that these are separated by the non-linearity introduced by the ELB. To implement the lower bound in our simulations, we use the shadow-price shocks approach of Holden and Paetz (2012). There are parallels between our approach and the one outlined in Guerrieri and Iacoviello (2015). Neither approach captures any precautionary motives associated with the possibility that the ELB will bind in the future. That said, our approach applies more readily to a model of this size compared to alternatives, such as that proposed by Fernández-Villaverde et al. (2015). Our analysis follows that of Coenen and Warne (2014), who conduct a similar exercise for the euro area.

To add greater realism to our simulations, we base them around the MPC's forecasts as published in the *Inflation Report*. This is consistent with the assumptions in Coenen and Warne (2014), who also condition their projections. The MPC's projections are for inflation and GDP growth and are conditioned on the assumption that Bank Rate follows the path expected by financial markets—also referred to as the "market path." They are published in the form of density forecasts, covering 90% of the perceived possible outcomes for the two variables—the so-called "fan charts." As such, we can compare the deflation probability implied by our model-based approach with those in the MPC's fans.

Other work our analysis is related to includes that of Swanson and Williams (2013, 2014). These authors investigate the perceived constraints on nominal interest rates in the US and the UK by studying how long rates respond to data surprises (with a fall in sensitivity indicating monetary policy constraints). Their results suggest that, both in the US and in the UK, perceptions of the degree to which policymakers were constrained by the lower bound varied in the post-2009 period.

The rest of the paper proceeds as follows. Section 2 discusses observed measures of deflation probability; Section 3 discusses our model and methodology for the stochastic simulations, including the imposition of the ELB; Section 4 describes our results, beginning with the unconditional probability of deflation in the model as well as the evolution and drivers of deflation probability since the crisis; Section 5 concludes.

2 Quantifying the probability of deflation: observed measures

In this section we present different measures of the probability of deflation. First, the MPC's forecasts as published in its quarterly *Inflation Report*; second, distributions for inflation derived from financial-market options; third, the Survey of Economic Forecasters.

2.1 MPC fan charts

The *Inflation Report* projections are produced as forecast densities—widely known as the "fan charts"—so we can compute the probability of deflation over time from these distributions. The projections look ahead over the next 13 quarters, and are published on a quarterly basis (in February, May, August and November). The forecasts we use in this section are for annual CPI inflation and are produced under the assumption that monetary policy follows the path expected by financial markets in the weeks preceding the date of their publication.

The fan charts are based on a judgemental forecasting exercise. Although the MPC uses a central model, known as COMPASS, in its forecast process—see Burgess et al. (2013) for details on the model—stochastic simulations are not used to produce the forecast distributions. Rather, the fan charts are produced by translating Committee judgement into the parameters of a two-part normal distribution for each quarter of the forecast horizon. In particular, the width of the fan chart is scaled judgementally, informed by past forecast errors. The MPC also incorporate skews into the distributions if they think that the risks to the outlook are weighted in one direction.[2] Hence, we use the two-part normal distribution and the published moments to compute the cumulative probability of inflation falling below zero at each point in the forecast horizon and for each forecast vintage.[3]

The two panels in Figure 1 show the probability of deflation across forecast vintages and over different horizons under alternative assumptions about skews in the forecast distributions. Here, the probability of deflation is defined as a fall in the Consumer Price Index on an annual basis in that quarter. In the left-hand panel, we plot the probability of deflation assuming the distributions have the skews that the MPC published. In the right-hand panel, we plot the deflation probability under the assumption that the first and second moments of the distributions match those of the MPC's projections, but that distributions are symmetric. For each forecast horizon and each forecast vintage, the figures plot the share of simulated paths for which inflation is below zero (green corresponding to a low share of simulated deflation episodes and red to high). The horizontal axis gives the origin of each forecast, and the vertical axis the quarter being forecasted. Since the MPC produces forecasts spanning three years, the width of the band is constant, but each quarter, the first quarter of the forecast rolls on one (hence the upward slope). So reading vertically up the chart shows the probability of deflation at different horizons, based on a forecast produced at the date on the x-axis. Reading

[2] See the box on pages 48-49 of the May 2002 *Inflation Report* (Monetary Policy Committee, 2002) for a fuller description of the MPC's fan charts and what they represent.

[3] The MPC's fans only cover 90% of the distribution of outcomes. The distribution outside the fans is not necessarily assumed to follow the two-part normal in the MPC's projections. However, for our purposes, we are assuming that the remaining 10% is distributed in the same way as the central 90%.

Figure 1: Probability of deflation at different forecast horizons across forecast vintages from MPC fan charts

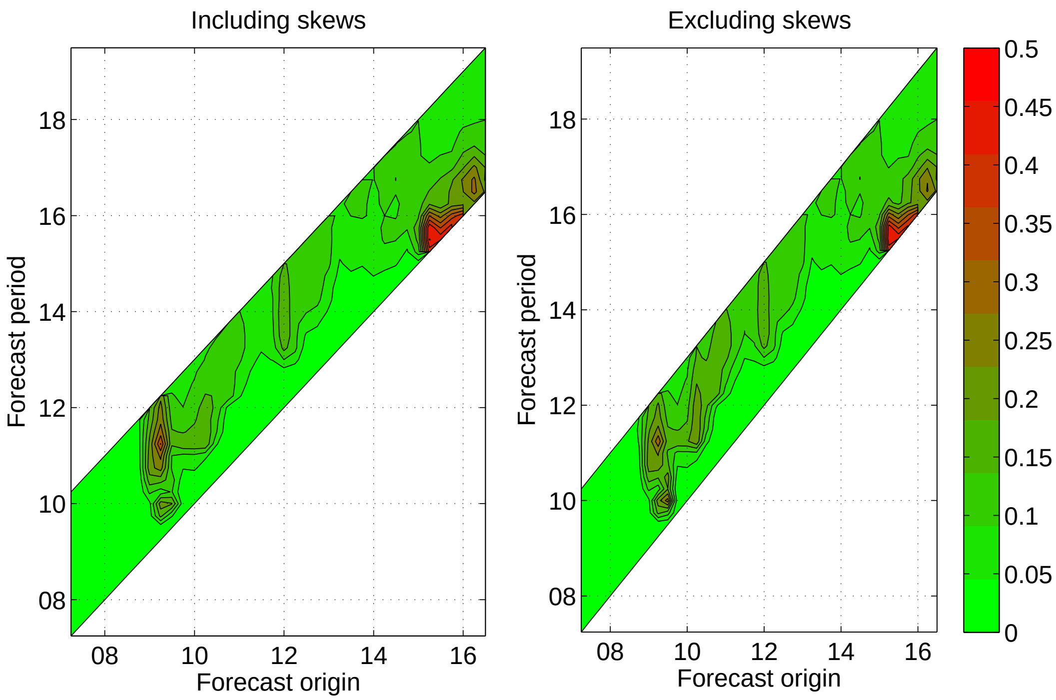

Notes: For forecasts made at each forecast origin (horizontal axis) and for each forecast horizon (vertical axis), colours refer to the probability of deflation in that quarter. Red corresponds to a high probability and green to a low probability.

horizontally shows how the probability of deflation in a particular quarter evolved over successive forecasts.

In the MPC's projections, including the published skews, there are episodes of heightened probility of deflation soon after the crisis, around 2009, and in the most recent period. Between 2010 and late 2014, the MPC's projections included relatively low probabilities of deflation. This is also the case when we assume consider the projections in the absence of the published skews. Comparing the two panels shows the importance of the skews in the MPC's fan charts. In practice, the difference is fairly small, which probably reflects two factors. First, deflation tends to be in the tails of the inflation distribution and so it would take large skews to materially increase the probability of deflation occurring. A second factor is that the MPC explicitly excluded the worst of the risks associated with the euro-area crisis from their fan charts. This exercise, therefore, is likely to understate the true probability of deflation over this period.[4]

2.2 Financial-market option-implied deflation probabilities

An alternative way to gauge the probability of deflation at different points in time is to consider the probability density functions (PDFs) implied by options on UK inflation traded in financial markets.[5] Investors trade these options to take a position on what

[4]See the box on page 38 of the August *Inflation Report* (Monetary Policy Committee, 2011) for further details on the impact the impact of euro-area developments on the UK were treated in the MPC's projections.

[5]Strictly, the inflation-linked instruments used for the inflation PDFs are caps and floors (see Smith (2012)) .

future inflation outturns might be. The prices can, therefore, be used to infer a probability distribution for inflation at different horizons (Smith, 2012). As well as computing the implicit distribution of inflation at different horizons, we can also extract the probability of particular thresholds (including the deflation probability).

There are two important assumptions underpinning the calculations in Figure 2. First, the probability of deflation implied by these options is only equivalent to the true probability of deflation under the assumption that options prices do not build in a risk premium. Were investors actually to be risk averse, and were the dislike of deflation to rise, it would look like a rise in the probability.[6] Second, an adjustment has been made to reflect the fact that the options are traded on the Retail Price Index (rather than the Consumer Price Index on which the simulations were based). These indices differ according to the statistical approach used to aggregate prices (formula effect) and the coverage (and in particular the coverage of mortage interest payments in the RPI). On average, RPI inflation is around 1pp higher than the CPI index and so we compute the deflation probability as RPI inflation of less than 1%. As a result, we correct the data by that amount, at the risk of missing out potential changes in the wedge or in the markets' expecations of the wedge.

Figure 2: Probability of deflation from UK option-implied PDFs

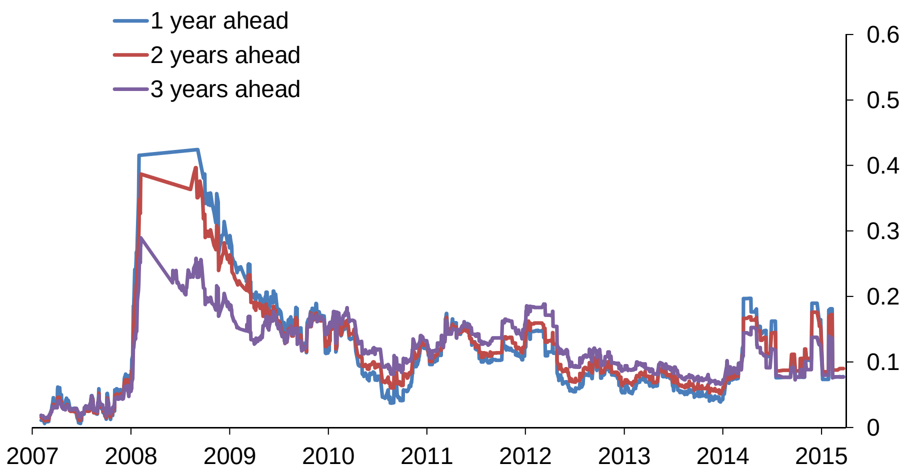

Notes: Bloomberg and Bank calculations. Probability of RPI inflation below 1% (equivalent to deflation in CPI given an average 1pp wedge) at one, two and three year horizons, as implied by the inflation distribution imputed from options prices. Straight lines indicate an absence of data (particularly prominent in 2008 and early 2009) and reflects the volatility of inflation over that period (see Smith (2012)).

Figure 2 plots the probability of deflation from option-implied PDFs at the one, two and three year horizon. The option-implied PDFs imply a heightened deflation probability in the early part of the crisis, which subsequently fell back and troughed in around 2011 (at all horizons). There is a small pickup in the deflation probability in 2012 and 2013, and then again in 2015. A noteworthy feature of these option-implied probabilities is the relative ranking of deflation probability at the horizons of one, two and three years. For most of the sample, the probability is lower at the one-year horizon,

[6]There is a parallel between this assumption and our use of a linearised model for our model-based simulations. By linearising the DSGE model, we are imposing certainty equivalence and hence a form of risk-insensitive behaviour on the behalf of agents.

the exceptions occurring in the immediate aftermath of the crisis and in 2014 and 2015, in line with the results we illustrate in the previous section for MPC forecasts.[7]

2.3 Survey of Economic Forecasters

A third metric for the probability of deflation can be gleaned from the economic projections of external forecasters collected by the Bank of England each quarter (the Survey of Economic Forecasters, SEF). Respondents are asked to assign probabilities for CPI inflation one, two and three years ahead lying in particular buckets, which from 2009 included a "less than 0%" category. These are typically presented as mean values.[8]

The mean estimate of the probability of deflation for one, two, and three years ahead is plotted in Figure 3. As with the MPC's forecasts and measures derived from options, the SEF points to a heightened probability of deflation in 2009. This then subsided and remained relatively low until 2015. From 2015, on the one-year ahead measure, the probability of deflation picked up again. Furthermore, consistent with option-implied measures, the deflation probability at the one-year horizon is the lowest for a majority of the sample. The ranking is reversed in the aftermath of the crisis and since 2015.

One issue to note is that the SEF points to lower probabilities of deflation compared to either the MPC's forecasts or the options-implied measure. In part, this could reflect biases in the level of survey-based expectations. Typically, the mean projection for inflation does not match up with the average of inflation in the data. Therefore, we focus on the movements in deflation probabilities according to this measure—that it started high, before falling back, and then rising again somewhat—rather than the reported level.

Figure 3: Probability assigned to deflation by Survey of Economic Forecasters

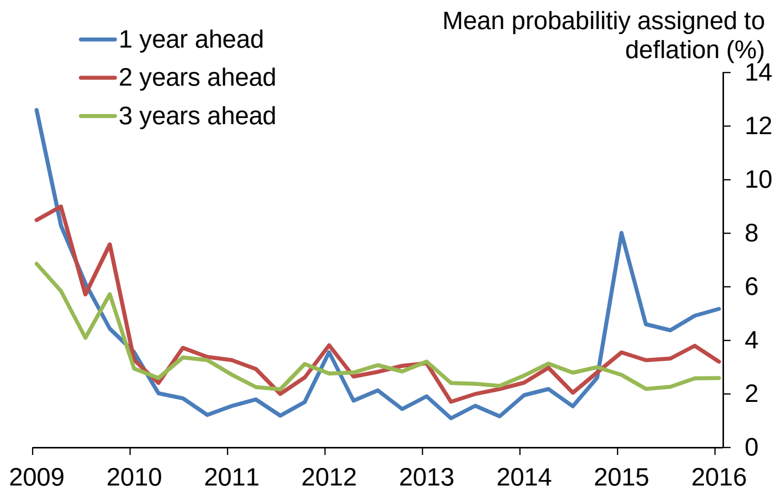

Notes: Bank of England Survey of Ecomomic Forecasters. Probability that a sample of external forecasters assigned to CPI inflation falling below 0% one, two and three years ahead.

[7] Some caution needs to be taken when interpreting the precise levels of the probabilities in the first three years. While the sum of the PDFs for the first three years is pinned down by actual prices, how that sum is distributed across the years is sensitive to how the options prices are interpolated.

[8] The median is very similar.

2.4 Stylised facts

While there are differences in the absolute levels of the probability of deflation across the three measures we consider, two robust findings emerge:

(i) the probability of deflation was relatively high during the crisis, largely subsided between 2010 and 2014 and increased again afterwards;

(ii) since 2015, the probability of deflation has been more pronounced at shorter horizons.

The remainder of the paper will investigate how our model matches up to these findings and will use the model to study unobserved characteristics of the economy, in particular, perceptions of the scope for monetary easing—as proxied by existence of an effective lower bound on interest rates—that best reconciles this evidence with our model simulations.

3 Methodology

3.1 Model

For our simulations, we use a medium-sized New Keynesian DSGE of the UK economy. This model is based on that described in Burgess et al. (2013). It is in the class of models popularised by Christiano et al. (2005) and Smets and Wouters (2003), which are now used widely by central banks around the world. The model of Burgess et al. (2013) departs from those of Christiano et al. (2005) and Smets and Wouters (2003) insofar as it includes a small open economy extension. Compared to Burgess et al. (2013), the model we use has three main differences. First, the estimation sample is extended to cover 1993 Q1 through 2012 Q4. Second, the pricing scheme allows for firms that price their output according to a rule-of-thumb (RoT). Third, the rest-of-the-world block is enriched to include a world interest rate.[9]

Incorporating post-crisis data into our estimate should allow us to better capture economic relationships over the period in our stochastic simulations (see Campbell et al., 2016). But it does introduce additional challenges around dealing with the effects of unconventional policy measures when estimating parameters. In extending the sample to 2012Q4 we take into accounts the effects of quantitative easing along the lines of Fawcett et al. (2015).[10] Moreover, we allow for a productivity trend-break over the crisis.[11]

Turning to our specification for price setting, we assume that optimising firms set their prices according to a Calvo scheme.[12] We also allow some firms to use a simple price updating rule as opposed to solving the optimal price-setting problem (for example, see Galí and Gertler, 1999; Steinsson, 2003; Dorich et al., 2013). Survey data supports the idea that a large share of firms set their prices in a rule-of-thumb fashion based on backward looking information (Greenslade and Parker, 2012). We use this survey information to calibrate the share of firms that do not optimise prices but rather follow

[9]More details of the model available upon request.

[10]As described on page 5 of that paper, we use a 'shadow' measure of the policy rate that aguments Bank Rate to include an in-house estimate of the effect of QE.

[11]Between 1993 and 2007 we calibrate trend productivity to deliver consumption growth at its average over that period, while after 2008 we calibrate the model to deliver trend consumption growth at its *historical average*, computed between 1955 and 2014.

[12]As opposed to the Rotemberg (1992) pricing scheme adopted in Burgess et al. (2013).

a mechanical updating rule to 20 percent.[13] The overall effect is a weakening of the link between current economic conditions (measured by the marginal costs in the different sectors) and inflation. In particular, the Phillips curve for each of the sectors in our DSGE model reads:

$$\begin{aligned}\pi_t^i &= \hat{\mu}_t^i + \frac{(1-\phi_i)\left(1-\phi_i\beta\Gamma^H\right)(1-\omega_i)}{\phi_i\left(1+\xi_i\beta\Gamma^H\right)+(1-\phi_{it})\,\omega_i} mc_t^i + \frac{\xi_i\left(\phi_i+(1-\phi_i)\,\omega_i\right)}{\phi_i\left(1+\xi_i\beta\Gamma^H\right)+(1-\phi_i)\,\omega_i}\pi_{t-1}^i \\ &+ \frac{\phi_i\beta\Gamma^H}{\phi_i\left(1+\xi_i\beta\Gamma^H\right)+(1-\phi_i)\,\omega_i}\mathbb{E}_t\pi_{t+1}^i\end{aligned}$$

where π_t^i is (quarterly) inflation for sector i in period t, mc_t^i is real marginal cost, and $\hat{\mu}_t^i$ is a mark-up shock. In addition, ϕ_i is the probability of not being able to change prices, ω_i is the share of rule-of-thumb price setters, ξ is the indexation parameter, β is the discount factor and Γ^H is the growth rate of hours. The index i indicates the sector to which the pricing refers (since the form of the Phillips curve is the same for the five pricing decisions in the model). The standard specification is nested, corresponding to $\omega_i = 0$.

In this case, it is instructive to express current inflation as a function of all future marginal costs:

$$\begin{aligned}\pi_t^i &= \xi_i\pi_{t-1}^i + \frac{(1-\phi_i)\left(1-\phi_i\beta\Gamma^H\right)(1-\omega_i)}{\phi_i+(1-\phi_i)\,\omega_i}\sum_{j=0}^{\infty}\left(\frac{\phi_i\beta\Gamma^H}{\phi_i+(1-\phi_i)\,\omega_i}\right)^j\mathbb{E}_t mc_{t+j}^i \\ &+ \frac{\phi_i+(1-\phi_i)\,\omega_i+\beta\Gamma^H\phi_i\xi_i}{\phi_i+(1-\phi_i)\,\omega_i-\beta\Gamma^H\phi_i\rho^{\mu^i}}\hat{\mu}_t^i\end{aligned}$$

The effect of rule-of-thumb pricer setters is, therefore, to lower the relative forward lookingness of price setting as ω_i shows up in the denominator of the expression determining the weights on future expected marginal cost.

Finally, we use an enriched specification of the rest-of-the-world block. The main purpose of this development to have a measure of the world interest rate that can vary over the cycle. In particular, the AR(1) processes in Burgess et al. (2013) are replaced with a VAR(2) on world GDP, world CPI inflation and a measure of the world interest rate, where sign restrictions—applied via the priors in the Bayesian estimation—are used to identify a world demand, a world supply and a world monetary shock. The key resulting change is in the UIP condition for the real exchange rate q_t, which features a measure of the deviation of world interest rate from its steady state value $(r_t^F - \mathbb{E}_t\pi_{t+1}^F)$:

$$q_t = \mathbb{E}_t q_{t+1} + \left(r_t - \mathbb{E}_t\pi_{t+1}^Z\right) - \left(r_t^F - \mathbb{E}_t\pi_{t+1}^F\right) - \hat{\varepsilon}_t^{B^F}$$

For a small open economy like the UK, movements in the exchange rate and import prices are a key factor influencing inflation, which is the motivation for enriching the world block for our exercise. Table 1 reveals that world shocks drive up to around a sixth to one fifth of the forecast error variance of inflation (depending on the firm pricing scheme), so can play a significant role in generating deflation in our simulations.[14]

[13] We do not include these parameters in our Bayesian estimation because of challenges around identifying multiple parameters associated with price setting (Calvo parameters, rule-of-thumb shares, indexation parameters.

[14] Including exchange rate risk premium, world demand, world supply, world monetary policy, world trade and world export price shocks

Table 1: Share of world shocks in forecast error variance decomposition for inflation

Horizon:	1st quarter	5th quarter	9th quarter	13th quarter
World shocks	4%	12%	14%	14%
World shocks (no RoT pricing)	8%	20%	21%	21%

3.2 Stochastic simulations

To address the question of how the probability of deflation has evolved, we conduct a stochastic-simulation exercise. The stochastic simulations involve, for each forecast date we consider (2007 Q1 to 2016 Q2), producing three-year projections using the model, a conditioning assumption (described below), and a large number of randomly drawn sequences of shocks. This gives us, for each particular forecast date, a forecast density for each of the model's endogenous variables. As such, we are able to characterise the distributions around the outlook for these variables at different points since the financial crisis.

In light of our interest in exploring the effect on deflation probability of possible constraints on monetary policy's ability to loosen, we produce the simulations under alternative assumptions about a lower bound on nominal interest rates. Once we have these distributions, we are able to interrogate their properties—e.g., computing the probability of deflation in any particular quarter of the projection. In the stochastic simulations, we impose the ELB using the shadow-price shocks approach of Holden and Paetz (2012). The overall procedure in which we embed this step has similarities to the one outlined in Guerrieri and Iacoviello (2015). For further details on the implementation see Appendix A.

Our exercise is a conditional one. For each forecast we consider, we base our simulations around the projections from the MPC's *Inflation Report* (which are conditioned on the path for Bank Rate expected by financial markets at the time). This approach adds greater realism to the exercise because it incorporates information known at the time that the model misses (for example expectations for oil prices). Compared to an unconditional forecast, the conditional forecasts revert to steady state more slowly, and in a way that is more consistent with the way the data subsequently evolved.[15] Conditional forecasts are, therefore, likely to give a more realistic picture of the persistence of the headwinds facing the economy (as perceived by policymakers) over this period.

The MPC forecasts that we use are for annual CPI inflation and GDP growth in the UK over the next 13 quarters. We gather the forecasts published between February 2007 and May 2016. These are published as density forecasts, but for the purposes of our stochastic simulations, we focus on the modal forecasts (or "central projections").[16] The first step is to use our model to compute the shocks and underlying states that would be consistent with those modal forecasts (our "base forecast"). We apply an inversion on all of the model's shocks (except for the monetary policy shock) under the assumption that they are anticipated by agents in the model, to match the profiles for quarterly CPI inflation, quarterly GDP growth, and the level of Bank Rate with the modal path for

[15] For example see a comparison of a sample of the forecasts on which we condition and unconditional forecasts from the model for inflation, GDP growth and the interest rate in Figure B.1. The conditional forecasts for the interest rate return to long run averages relatively more slowly, particularly after 2012. There is also more persistence in some of the deviations of inflation from target, particularly around forecasts made in 2009. Forecasts for annualised unconditional GDP growth, however, tend to display more overshooting.

[16] We explore the information contained in the rest of the distributions in Section 2.1

each forecast vintage (or the mean market path in the case of Bank Rate).[17] Note that for all of the baseline forecasts, we use the same version of the model and the latest vintage of back data.

Finally, it is important to note that while excluding the monetary policy shock from the inversion is not necessarily an insignificant assumption in general, for the purposes of our results it is relatively innocuous. Exclusion of the monetary-policy shock means that the monetary-policy rule matches the market path exactly. In turn, that implies that the set of underlying non-policy shocks that we compute in the inversion needs to be consistent with the projections for inflation and GDP growth, conditional on the path for policy rates prescribed by the rule following the market path. If we included the monetary policy shock in the inversion, the set of non-policy shocks would likely be different. This would mean that the projections for the endogenous variables that are not conditioned (that is, all variables other than inflation, GDP growth, and the policy rate) would be different (given they would be responding endogenously to the different set of policy and non-policy shocks). In practice, however, the two variables that matter most for deflation probability are inflation and the nominal interest rate, which are conditioned in our baseline projections. As such, these variables will not be affected by the different set of shocks that underpin our baseline projection.

4 Simulation results

In this section we discuss the evolution of deflation probability, as calculated in our model-based simulations and defined as the chance of a negative annual inflation rate in that quarter, since the onset of crisis. Before turning to those results, we present two experiments to build intuition. The first presents the deflation probability in unconditional simulations, which gives a sense of the average deflation probability across the estimation sample of the model. The second demonstrates the effects that imposing an ELB has on the distribution of simulated paths for the economy at a single point in time.

4.1 Unconditional probability of deflation

Table 2: Unconditional probability of deflation

	a) Baseline model	b) Lower steady state interest rate	c) As b) & higher shock standard deviation
Deflation probability (ELB=0.5%)	0.16	0.18	0.23
Deflation probability (No ELB)	0.16	0.16	0.21
Probability of ELB of 0.5% binding	≈0	0.09	0.12

In this section we consider the model-implied probability of deflation from simulations that are initialised at the model's steady state. This exercise is intended to give a sense of the average probability of deflation implied by the model's properties. As we will come on to describe, it also hints at the importance of the conditioning assumption in our exercise in Section 4.3. In the absence of a lower bound on nominal interest rates, it would be

[17] The MPC's forecasts typically report annual inflation and annual GDP growth. However, given the quarterly nature of the model we prefer to condition on the quarterly rates for these variables.

possible to compute the implied probability of deflation by considering the asymptotic distribution for inflation in the model itself. When there is a lower bound in play, we compute the probability of deflation using stochastic simulations with an occasionally binding ELB, starting from its steady state. In effect, in these simulations, the deflation probability is a property of the ergodic distribution for inflation for the model subject to a lower bound on the policy rate.

In the model, two features affect how likely an ELB is to bind, and hence the effect of the ELB on the deflation probability. First, the distance between the steady-state nominal interest rate and the ELB—the bigger this gap, the less likely the ELB is to bind. Second, the standard deviation of the shocks. When the shocks are relatively small, the probability of nominal rates hitting the ELB is lower. Therefore, for this section, we also consider the sensitivity of the probability of inflation to changes to these assumptions.

Table 2 shows the unconditional probability of at least one quarter of deflation at any point over a thirteen quarter forecast horizon. For illustrative purposes, we consider the probability of deflation under the alternative assumptions of an ELB of 0.5% and of no ELB, and for alternative variants of the model. Beginning with the baseline model, the probability of deflation is around 16% irrespective of the assumption about whether the ELB is a binding constraint. That reflects the fact that from the model's perspective, a sufficiently adverse selection of shocks to drive the policy rule to the ELB is very unlikely, given we are initialising the simulations at the model's steady state (bottom row).

The second column considers an alternative model where the steady-state nominal interest rate has been lowered to 2% (from 4.3% in the baseline model, which is the sample average from 1993Q1 to 2012Q4). By shrinking the gap between the average policy rate and the ELB, the probability of a binding ELB rises from essentially zero to 9%. Where the ELB is assumed to exist, the deflation probability is also higher than under the assumption of no ELB, as a binding ELB starts to constrain the extent to which policy can be loosened in the face of deflationary shocks.

As an additional exercise to examine the sensitivity of the model-implied probability of deflation, the final column considers a version of the model in which the steady-state nominal interest rate has been lowered to 2%, *and* the standard deviation of the model's shocks has been increased by 10%.[18] The additional effect of raising the shock standard deviations is similar to that of lowering the steady-state interest rate, but for different reasons. First, it will increase the volatility of the variables in the model, making it more likely that absent any ELB inflation would fall into negative territory. In addition, it makes shocks that are sufficiently bad to make an ELB bind more likely, so raising the probability of deflation further.

These unconditional simulations illustrate that, in a mechanical sense, the determinants of the probability that the ELB binds are the gap between the starting level of the interest rate and the level of the ELB and the average size of the shocks. As we will come on to show, in an analogous way, the gap between the starting level of inflation and zero is a key determinant of the deflation probability in our analysis of deflation probility since the onset of the crisis (Section 4.3). In our conditional stochastic simulations, we always use the estimated shock standard deviations from the model. However, by conditioning on the starting level for inflation and the policy rate for different forecast vintages, and

[18]This assumption is modest. Mumtaz (2011) allows for time-varying volatility in the structural shocks of VARs for the US and finds that over the past fifty years the variation in the standard deviations has been considerably greater than 10%.

imposing different ELBs, we vary the size of the gaps between inflation and zero and the interest rate and the ELB.

4.2 A case-study: the November 2015 forecast

Imposing an effective lower bound on the nominal interest rate in our simulations introduces downside skews to the forecast distributions for inflation and GDP growth and upside skews to the distributions for nominal and real interest rates. By constrast, in the absence of an ELB, the assumption that the shocks are Gaussian implies that the forecast distributions for these variables are symmetric. Figure 4 shows the forecast distributions—plotted as "fan charts"—for annual CPI inflation, annual real GDP growth, the policy interest rate ("Bank Rate"), and the ex-ante real interest rate around the November 2015 MPC forecast under the assumptions that there is no ELB, that there is an ELB of zero, and that there is an ELB of 0.5%.[19] The just-described skewed pattern of responses is evident in this figure. Furthermore, it is clear that the probability of deflation is greater under the assumption that there is an ELB of 0.5%—a larger proportion of the forecast distribution for inflation lies below zero in this case.

The reason for this pattern of skews and the greater probability of deflation when there is an ELB of 0.5% (and, to a lesser extent, 0%) is straight-forward. The ELB places a limit on monetary policy's ability to loosen in the face of deflationary shocks, generating an upward skew in the distribution of the nominal policy rate relative to the symmetric unconstrained distribution. Given sticky prices, this leads to higher real interest rates than would be the case if the policymaker was unconstrained. Higher real interest rates, in turn, generate a downside skew in the the distribution of GDP growth in the near term and, ultimately, a downward skew in inflation.

4.3 Model-based assessment of deflation probability since the crisis

We can now turn to our analysis on the evolution of the probability of deflation since the onset of the 2007-08 financial crisis, as implied by our model simulations. Figure 5, which can be read in the same way as Figure 1, shows the evolution of the probability of deflation around forecasts made since 2007 assuming that there is an ELB of 0.5% and that there is no lower bound.

Three episodes of heightened deflation probability are noticeable when we impose an ELB of 0.5% (first panel, Figure 5). The first occurs during mid-to-late 2009. Forecasts made during this period include deflation probability in some quarters of up to around 30-40%. The forecasts include heightened risk at the start of the projection, which then wanes, before increasing again towards the end of the forecast horizon. The second episode of heightened deflation probability occurs in 2012-13. The probability of deflation in projections during this episode is higher than the in the immediate aftermath of the crisis, rising to up to around 50% in some quarters. Furthermore, the projections are for more protracted periods of deflation probability, rather than isolated periods at the start and end of the forecast horizon: there is heightened deflation probability over much

[19] The fan charts described here are the result of our stochastic simulation, not the fan charts published by the MPC. Rather than taking a stochastic simulation approach, they use information on past forecast errors and additional judgements about the size and balance of risks to the economy at different points in time. The differences between the two approaches discussed in more detail in Section 2.1.

Figure 4: Fan charts for simulations centred on November 2015 *IR* for an ELB of 0.5%

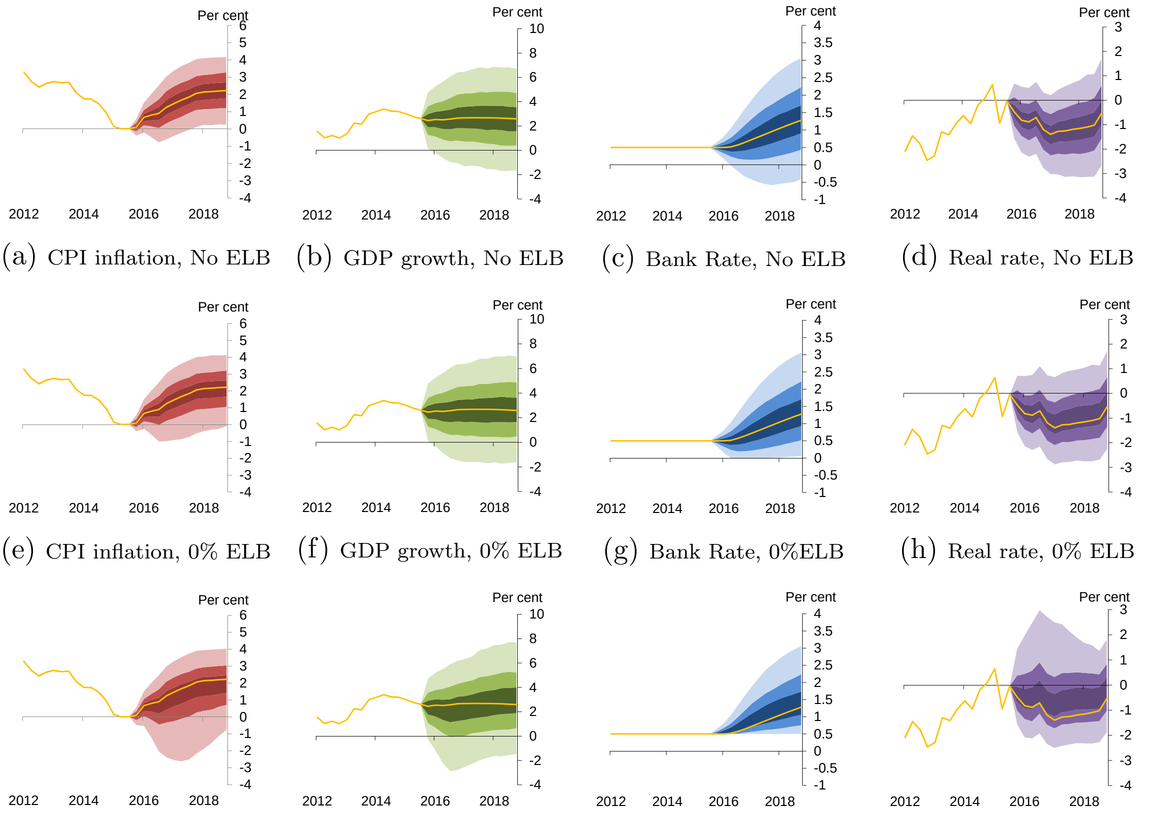

(a) CPI inflation, No ELB (b) GDP growth, No ELB (c) Bank Rate, No ELB (d) Real rate, No ELB

(e) CPI inflation, 0% ELB (f) GDP growth, 0% ELB (g) Bank Rate, 0%ELB (h) Real rate, 0% ELB

(i) CPI infl., 0.5% ELB (j) GDP grth., 0.5% ELB (k) Bank Rate, 0.5% ELB (l) Real rate, 0.5% ELB

Notes: Each pair of shaded bands represents 30% of simulated paths. The yellow line plots the central projection from the November 2015 *Inflation Report* forecast (or in the case of real rates, the path which those forecasts imply when imposed in the model), on which the simulations are centered.

of the forecast horizon in this episode. Finally, the third episode of heightened deflation probability occurs in the most recent period, in 2015-16. Here, the probability of deflation is around 35-45% in some quarters. The heightened deflation probability in 2015 is particularly marked in the near term. For the episode in late 2015 and early 2016, the probability of deflation is fairly consistent across the horizon of the forecast.

When we consider deflation probability in the absence of a lower bound on nominal rates (second panel, Figure 5), there continues to be some heightened probability of deflation in 2009 and 2015-16, but not in 2012-13. The pattern of deflation probability in the 2009 episode is broadly consistent across the two assumptions about the effective lower bound. In the latest episode, the heightened deflation probability in projections at the start of 2015 is also apparent when we assume that there is no lower bound on nominal rates. However, the heightened deflation probability in late 2015 and early 2016 is largely absent when we do not impose an ELB.

Another way to assess deflation probability is to consider the probability that annual inflation will turn negative in *at least one* of the thirteen quarters in each forecast. This is shown in Figure 6, which plots this probability under three assumptions about the existence of an ELB: that there is an ELB of 0.5%, 0%, and no ELB at all. The measure of deflation probability in this figure is distinct from that depicted in Figure 5, which is the probability of inflation in the particular quarter in question. The probability that inflation will be negative in at least one of the forecast quarters will, typically, be higher

Figure 5: Probability of deflation at different forecast horizons across forecast vintages

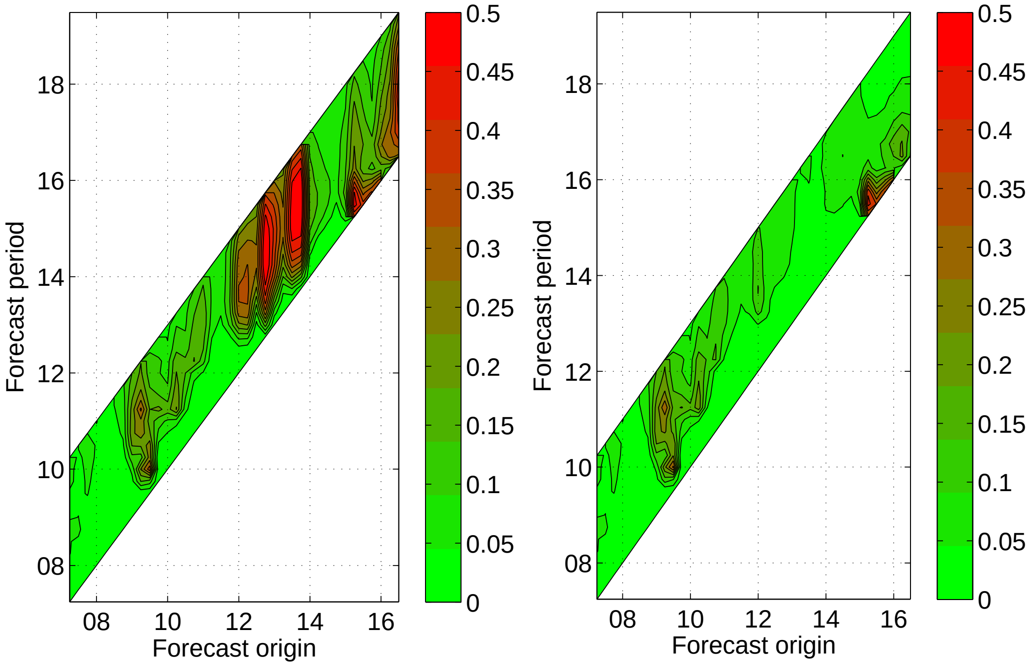

Notes: Simulated deflation probabilites for ELB=0.5% (left) and no ELB (right). For forecasts made at each forecast origin (horizontal axis) and for each forecast horizon (vertical axis), colours refer to the probability of deflation in that quarter. Red corresponds to a high probability and green to a low probability.

Figure 6: Probability of deflation in any quarter of forecast horizon across forecast vintages

Notes: Simulated deflation probabilites for ELB=0.5% (blue lines), 0% (green line) and no ELB (red line). The probability of deflation is calculated as the probability that inflation will be negatve in at least one quarter of the course of the forecast made at the dates along the x-axis.

than the probability of deflation in any particular quarter.

Under the assumption of an ELB of 0.5%, the three episodes of heightened deflation probability are also evident in Figure 6. The figure also reinforces the conclusion that the

ELB assumption is particularly important for the 2012-3 episode. When we assume there is no ELB, then the probability falls back to a little below 20% during 2012-13 (similar to the unconditional probability discussed in Section 4.1).

4.4 Drivers of deflation probability in our simulations

The model-based measures of deflation probability reported in Section 4.3 suggest that the existence of a lower bound on nominal interest rates plays an important role in the probability of deflation. But the results also suggest that this is only part of the picture: in both 2009 and in 2015-16, model-implied deflation probability is elevated even in the absence of an ELB. However, a common feature of the economic outlook in both these episodes was a relatively low starting point for inflation (in both cases, below the MPC's 2% target) and a corresponding weakness in the MPC's inflation projection. Therefore, there are two obvious candidates to explain differences in the probability across our alternative scenarios:

(i) the level of inflation over the forecast horizon, especially at the start;

(ii) the degree to which monetary policy is constrained, as reflected by the gap between the market path for interest rates and the level of an effective lower bound on policy rates.

These two factors play different roles in each of the episodes. The higher deflation probability associated with forecasts made in 2009 reflected the fact that the expectation at the time was that inflation would remain low over the next three years (Table 3).[20] As such, the shocks needed to drive inflation into deflationary territory would not have been very large. However, the market path was fairly steep which meant that any ELB had little bearing on the probability of deflation.

In the 2015 episode, the key determinant was also the low level of inflation, but primarily in the early part of the forecast (it was expected to return to 2% fairly quickly). As a consequence, stochastic simulations around these central projections generate deflationary outcomes primarily in the early part of the forecast window. The market path was also very flat at this point, but it does little to raise the deflation probability in the early part of the projection; even if the policy stance can be loosened, it will not be very effective at preventing deflation at very short horizons. And since the forecasts on which the simulations are centred rise quickly back towards target and the shocks from the stochastic simulation quickly dissipate, imposing an ELB only raises the probability of deflation at longer horizons a little.

The 2012-13 episode, in contrast, corresponded to a period when inflation was expected to be at or above target (second column of Table 3), but while the market path was extremely flat and close to 0.5%. As a consequence, the probability of deflation did not materialise in the near-term horizons in the forecast because it would taken several quarters worth of deflationary shocks to drive inflation into negative territory. At the same time, the scope for policy loosening was limited (when the 0.5% ELB is considered), hence deflationary shocks, the effects of which would be normally counteracted by a policy loosening ,led to persistent and severe deflationary episodes in years two and three of

[20]This highlights the importance of using conditional forecasts for our exercise. Given the faster speed with which unconditional forecasts return to target, the deflation probabilities absent conditioning would be around zero at this time (Figure B.2).

Table 3: Inflation and market path slopes across high deflation probability episodes

	Immediate post crisis episode	2012-3 episode	2015 episode
Modal inflation forecast (%)			
Average over first year	1.3	2.7	0.3
Average over second year	1.2	2.1	1.6
Average over third year	1.2	1.9	2.1
Market path statistics			
Slope	1.4	0.1	0.4
Average quarterly change (pp)	0.15	0.03	0.07
Memo: Episode dates	*2009Q1-Q3*	*2012Q1-2013Q4*	*2015Q1-Q4*

Notes: Slope of market path defined as average level of market path minus level of yield curve in first quarter

the forecast horizon. This explains why in the 2012-13 episode the probability of deflation materialises primarily towards the end of the forecast horizon and why relaxing the ELB constraint substantially reduces the possibility that our simulations might generate deflation.

Figure 7 illustrates the role of the slope of the market path at different horizons. For each forecast vintage, we plot the slope of the market path (defined as the average short-term rate over the forecast minus the level in the first quarter) against the deflation probability one quarter ahead, two years and three years ahead. In practice, both the slope *and* the starting level of the market path are important.[21] By excluding observations pre-2009 Q2 we can abstract from the role of the starting level since the market path for all forecasts after date have the same starting point of 0.5%. The blue dots show that for the majority of forecast vintages, there is a very low probability of deflation in the first quarter (with the exception of forecasts made in 2015). Moreover, it bears very little relation to the slope of the market path. However, for the two- and three-year ahead deflation probabilities, there is a negative slope: flatter market paths are associated with higher deflation probabilities.

4.5 Discussion: model-based and alternative measures of deflation probability

The discussion of our model-based measures of deflation probability and their drivers in Sections 4.3 and 4.4 suggests that the measures of deflation reported in Section 2 are consistent with a perception that the MPC has been relatively unconstrained in its ability to ease the stance of policy since the onset of the crisis.

The first thing to note is that the probability of deflation implied by the MPC's fan charts (Figure 1) looks very similar to the probability of deflation in our simulations where we do not impose an ELB (the right hand panel of Figure 5). In particular, it identifies very similar probabilities of deflation in the 2009 and 2015 episodes. Moreover, the MPC fan charts do not identify an episode of elevated deflation probability in 2012-3, consistent with our simulations when we do not impose an ELB. This is in contrast with our simulations under an ELB of 0.5%, in which the probability of deflation is

[21] This is equivalent to the long rate minus the short rate up to the term premia.

Figure 7: Slope of the market path and deflation probability at different horizons

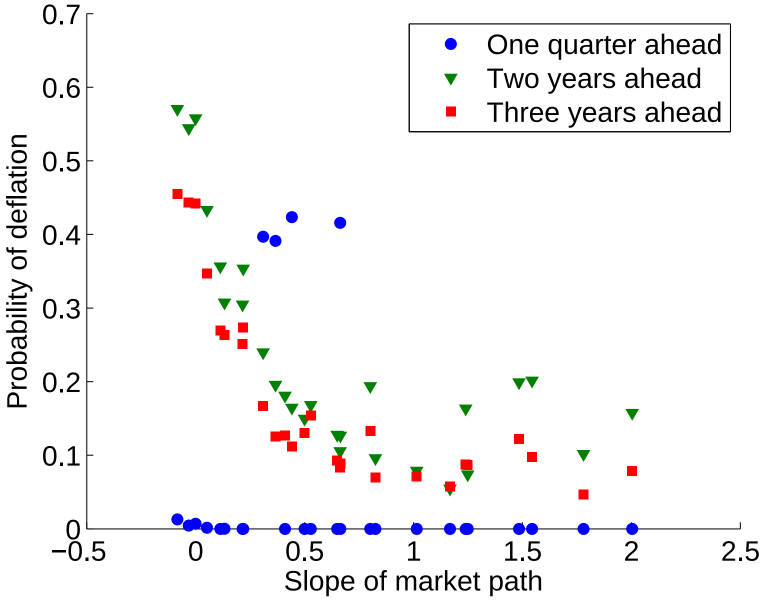

Notes: Forecast vintages before 2009Q2 are excluded from the chart (since the slope of the market path becomes most relevant once Bank Rate has reached 0.5%). The slope of the market path is defined as the mean level of the yield curve over the forecast horizon minus the level of the market path in the first quarter.

heightened in 2012-13. In other words, the MPC's fan charts are consistent with having been produced under the (implicit) assumption that monetary policy was unconstrained since the onset of the crisis. Furthermore, it is consistent with the MPC's clarification in 2015 that, were deflationary shocks to occur, there would be scope to cut Bank Rate towards zero from 0.5%, which had hitherto been regarded as an ELB, as well as engage in further rounds of QE, if necessary (see Carney, 2015).

To provide a more comprehensive picture, Figure 8 reports all the measures of deflation probability that we are considering in our analysis.[22] The evolution of the different measures bears significant similarities.[23] All measures identified a heightened probability of deflation in 2009, which fell back in the years to 2011. All our measures also show a remarkable degree of comovement at the one-year horizon over the last couple of years, when probabilities of deflation upwards of 30 percent have been recorded across the board.

The 2012-3 period, by contrast, shows quite marked differences between our data measures and simulations under the most restrictive assumption that 0.5% consitutes a lower bound for policy rates. Comparing the alternative measures to the swathe of deflation probabilities from our simulations, it is clear that the observed measures do not exhibit the spike in deflation probability over this time. Judgemental and market-based forecasts are, instead, much closer to simulations in which the ELB is either lower than 0.5% or non-existent.

In sum, our simulations, despite the linearity of the model, do a good job of capturing the dynamic evolution of three different data-based measures of the probability of deflation when we allow for policy rates to fall below the 0.5%. We find this sugges-

[22]The chart for the three year horizon is found in Figure B.3.
[23]The survey-based measure is lower in absolute level.

Figure 8: Summary of alternative measures of deflation probability

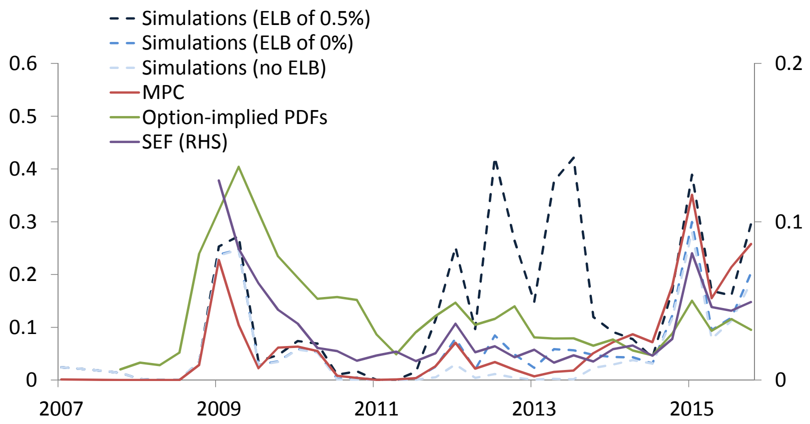

(a) One year ahead

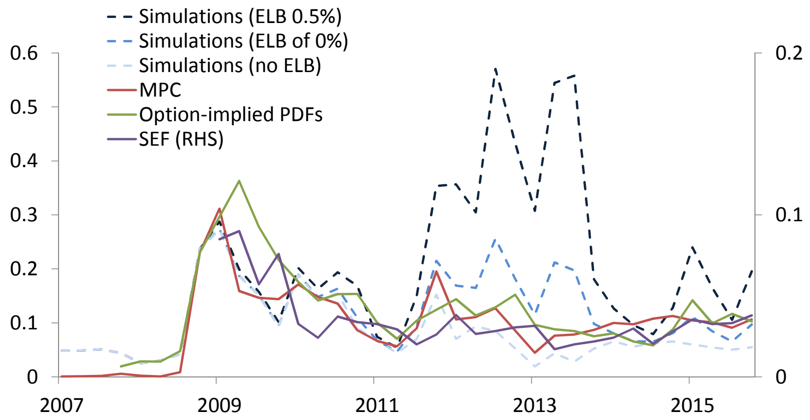

(b) Two years ahead

Notes: Option-implied PDF measures have been converted into quarterly averages to be consistent with other measures of deflation probability.

tive that economic forecasters, market participants and policymakers alike expected that rates could be cut below the 0.5% level if economic conditions warranted the move. Or alternatively, that unconventional policy measures such as QE could substitute for rate cuts.

5 Conclusions

In this paper we use a quantitative DSGE model of the UK economy to uncover what observed measures of deflation probability imply for perceptions about the degree to which monetary policy has been constrained in its ability to offset the effects of deflationary shocks since the onset of the financial crisis in 2007-08. These observed measures of deflation probability point to heightened probabilities of deflation around 2009 and in 2015. In measures of deflation probability generated from stochastic simulations of our DSGE model, these episodes of heightened deflation probability are also evident. However, when we assume that monetary policy is constrained in its ability to offset the effects of deflationary shocks by imposing a lower bound on policy rates rates of 0.5%, we also identify an episode of heightened deflation probability in 2012-13. That this episode is not evident in the observed measures of deflation suggests that monetary policy was not perceived to have been constrained at that point.

Bibliography

Burgess, S., Fernandez-Corugedo, E., Groth, C., Harrison, R., Monti, F., Theodoridis, K., and Waldron, M. (2013). The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models. Bank of England working papers 471, Bank of England.

Campbell, J. R., Fisher, J. D. M., Justiniano, A., and Melosi, L. (2016). Forward Guidance and Macroeconomic Outcomes Since the Financial Crisis. In *NBER Macroeconomics Annual 2016, Volume 31*, NBER Chapters. National Bureau of Economic Research, Inc.

Carney, M. (2015). Open letter to the Chancellor.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, 113(1):1–45.

Coenen, G. and Warne, A. (2014). Risks to Price Stability, the Zero Lower Bound, and Forward Guidance: A Real-Time Assessment. *International Journal of Central Banking*, 10(2):7–54.

Dorich, J., Johnston, M., Mendes, R., Murchison, S., and Zhang, Y. (2013). ToTEM II: An updated version of the Bank of Canada's Quarterly Projection Model. *Bank of Canada Technical Report, no 100.*

Fawcett, N., Koerber, L., Masolo, R., and Waldron, M. (2015). Evaluating UK point and density forecasts from an estimated DSGE model: the role of off-model information over the financial crisis. Bank of England working papers 538, Bank of England.

Fernández-Villaverde, J., Gordon, G., Guerrón-Quintana, P., and Rubio-Ramírez, J. F. (2015). Nonlinear adventures at the zero lower bound. *Journal of Economic Dynamics and Control*, 57(C):182–204.

Galí, J. and Gertler, M. (1999). Inflation Dynamics: A Structural Economic Analysis. CEPR Discussion Papers 2246, C.E.P.R. Discussion Papers.

Greenslade, J. V. and Parker, M. (2012). New Insights into Price Setting Behaviour in the UK: Introduction and Survey Results. *Economic Journal*, 122(558):F1–F15.

Guerrieri, L. and Iacoviello, M. (2015). OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics*, 70(C):22–38.

Holden, T. and Paetz, M. (2012). Efficient simulation of DSGE models with inequality constraints. School of Economics Discussion Papers 1612, School of Economics, University of Surrey.

Krippner, L. (2013). Measuring the stance of monetary policy in zero lower bound environments. *Economics Letters*, 118(1):135–138.

Monetary Policy Committee (2002). Inflation Report, May.

Monetary Policy Committee (2011). Inflation Report, August.

Mumtaz, H. (2011). Estimating the impact of the volatility of shocks: a structural VAR approach. Bank of England working papers 437, Bank of England.

Smets, F. and Wouters, R. (2003). An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area. *Journal of the European Economic Association*, 1(5):1123–1175.

Smith, T. (2012). Option-implied probability distributions for future inflation. *Bank of England Quarterly Bulletin*, 52(3):224–234.

Steinsson, J. (2003). Optimal monetary policy in an economy with inflation persistence. *Journal of Monetary Economics*, 50(7):1425–1456.

Swanson, E. T. and Williams, J. C. (2013). Measuring the effect of the zero lower bound on yields and exchange rates. Technical report.

Swanson, E. T. and Williams, J. C. (2014). Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates. *American Economic Review*, 104(10):3154–85.

A Stochastic simulation subject to the ELB

A.1 Without an ELB

When monetary policy is not assumed to be constrained by an ELB, producing stochastic simulations is straightforward. It involves projecting forward for H quarters (in our case the 13 of the MPC's projections) from an initial condition provided by the data, using the model's rational expectations solution and an H-quarter-long sequence of unanticipated shocks (except the monetary policy shocks). These are drawn from a normal distribution with mean zero and the shocks' estimated standard deviations. This process is repeated a large number of times (5,000) to give a distribution of forecasts. This results in symmetric fan charts centred on the baseline forecasts, where the width of the fan is determined by the standard deviation of the shocks in the model. To illustrate this point, fan charts without an ELB imposed can be found in the left hand column of Figure **??** at the back of the paper.

A.2 With an ELB

In producing stochastic simulations with a binding ELB, we impose the constraint using the approach of Holden and Paetz (2012). The starting point is a linear rational-expectations model, for which the structural equations for the endogenous variables, x_t, can be written as:

$$H^F \mathbb{E}_t x_{t+1} + H^C x_t + H^B x_{t-1} = \Psi z_t$$

To separate out the monetary policy rule for the nominal interest rate r_t from the others variables (denoted by $\tilde{x}_t$), the model can be partitioned as follows:

$$\begin{aligned} \begin{bmatrix} H^F_{\tilde{x}\tilde{x}} & H^F_{\tilde{x}\tilde{r}} \\ H^F_{r\tilde{x}} & H^F_{rr} \end{bmatrix} \mathbb{E}_t \begin{bmatrix} \tilde{x}_{t+1} \\ r_{t+1} \end{bmatrix} &+ \begin{bmatrix} H^C_{\tilde{x}\tilde{x}} & H^C_{\tilde{x}\tilde{r}} \\ H^C_{r\tilde{x}} & H^C_{rr} \end{bmatrix} \begin{bmatrix} \tilde{x}_t \\ r_t \end{bmatrix} + \begin{bmatrix} H^B_{\tilde{x}\tilde{x}} & H^B_{\tilde{x}\tilde{r}} \\ H^B_{r\tilde{x}} & H^B_{rr} \end{bmatrix} \begin{bmatrix} \tilde{x}_t \\ r_t \end{bmatrix} \\ &= \begin{bmatrix} \Psi_{\tilde{x}\tilde{x}} & 0 \\ 0 & \Psi_{rr} \end{bmatrix} \begin{bmatrix} z^{\tilde{x}}_t \\ z^r_t \end{bmatrix} \end{aligned}$$

where z^r_t is the monetary policy shock and $z^{\tilde{x}}_t$ are the other shocks in the model. When the nominal interest rate cannot go below a certain value, b (an ELB), the policy rule is a non-linear equation which has the following form:

$$r_t = \max\{b, -H^F_{rr}\mathbb{E}_t x_{t+1} - H^C_{rx} x_t - H^B_{rr} x_{t-1}\} \tag{1}$$

For the purposes of the implementation, this will be replaced with:

$$r_t = -H^F_{rr}\mathbb{E}_t x_{t+1} - H^C_{rx} x_t - H^B_{rr} x_{t-1} + \Psi_{rr} \sum_{s=0}^{H} z^r_{t+s} \tag{2}$$

where z^r_{t+s} are the shadow-price shocks that ensure that the constraint b is respected (and which are anticipated by agents in the model). The solution minimises the size of the anticipated shocks needed to enforce the ELB, subject to additional criteria: the shadow-price shocks can only be positive (they can only increase the interest rate) and the shadow-price shock at a particular horizon can only be non-zero if the bound binds at that horizon. That is, the shocks are the optimum of a quadratic programming problem.

The approach is incorporated into the stochastic simulation on a period-by-period basis, which ensures that the shocks used to impose the ELB are anticipated only once the (unanticipated) shocks that have driven the policy rate to the ELB have hit. A stochastic simulation for H periods subject to the ELB, therefore, can be constructed by combining H perfect foresight projections from an initial condition that is updated sequentially with the unanticipated shocks, to move from forecast period 0 to forecast period $H-1$. The steps proceed as follows:

i. In the first period, a projection absent the ELB is produced from a random draw of the shocks in that period around the base forecast from an initial condition (period 0).

   a. If the interest rates implied by the policy rule are below the ELB, solve for the (potentially large) sequence of anticipated monetary policy shocks that enforce the ELB.

   b. If the ELB is not binding, the vector of anticipated monetary policy shocks would be equal to zero.

ii. Update the projection with this sequence of anticipated monetary policy shocks and retain the endogenous variables in first period.

iii. Repeat these steps in the second period using the endogenous variables from the previous period as the initial condition, and continue to produce an H-quarter projection subject to the ELB.

B Figures

Figure B.1: Conditional and unconditional forecasts for inflation, GDP growth and interest rates

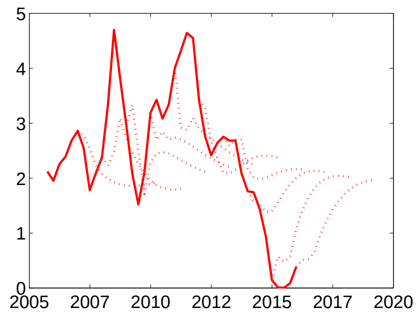

(a) Inflation: Unconditional

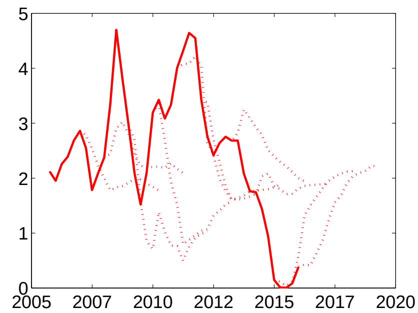

(b) Inflation: Conditional

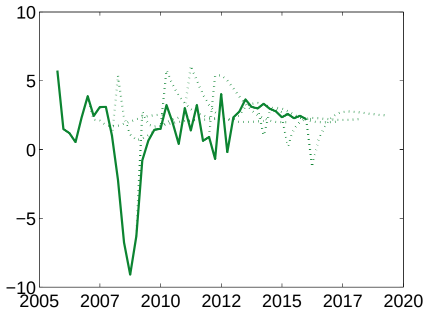

(c) GDP: Unconditional

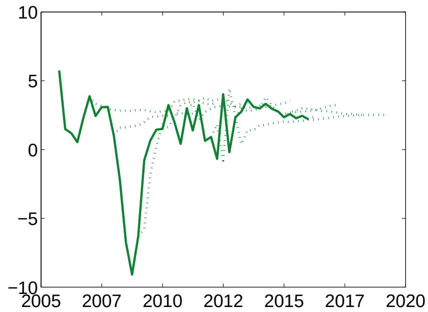

(d) GDP: Conditional

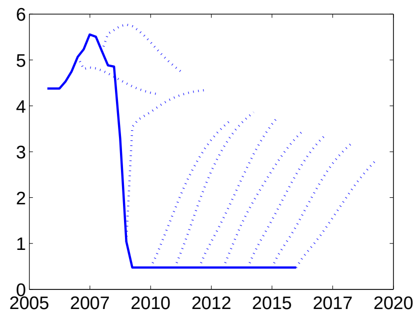

(e) Interest rate: Unconditional

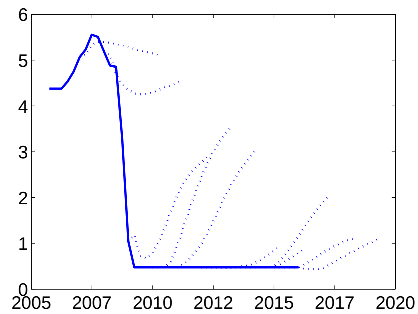

(f) Interest rate: Conditional

Notes: Solid lines report data. Dashed lines report forecasts for annual inflation, annual GDP growth and the interest rate produced in Q1 of 2006 to 2016. Left-hand column reports model forecasts. The right-hand column reports forecasts conditioned to match the MPC's *Inflation Report* forecasts.

Figure B.2: Probability of deflation across forecast horizons based on unconditional model forecasts

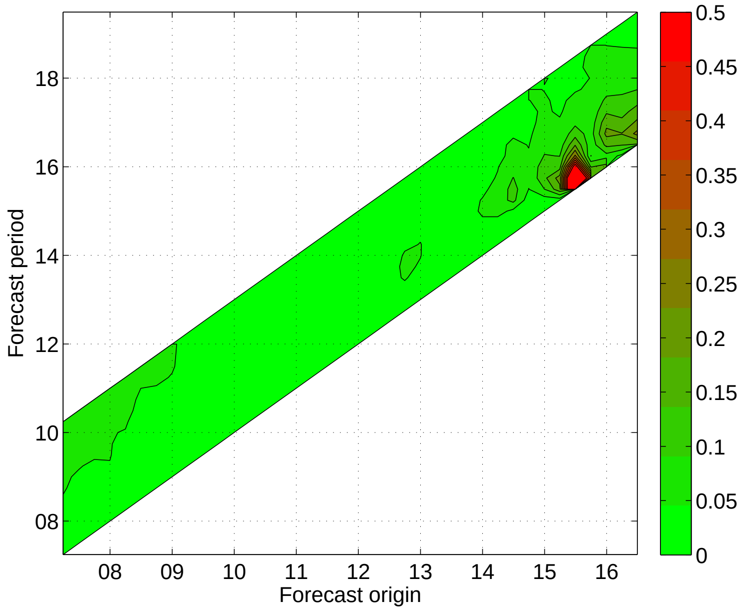

(a) ELB of 0.5%

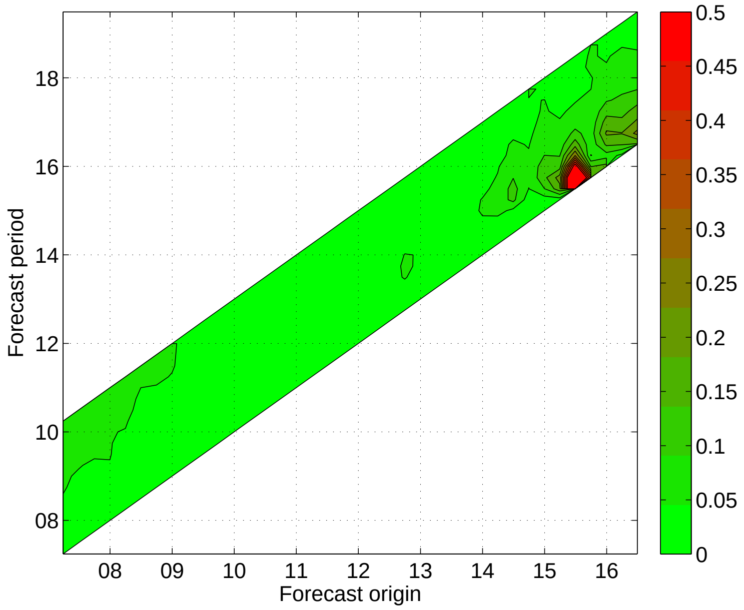

(b) No ELB

Notes: Simulated deflation probabilities for ELB=0.5% (top) and no ELB (bottom) based on unconditional model forecasts. For further notes on how to interpret the chart see Figure 5.

Figure B.3: Summary of alternative measures of deflation probability three years ahead

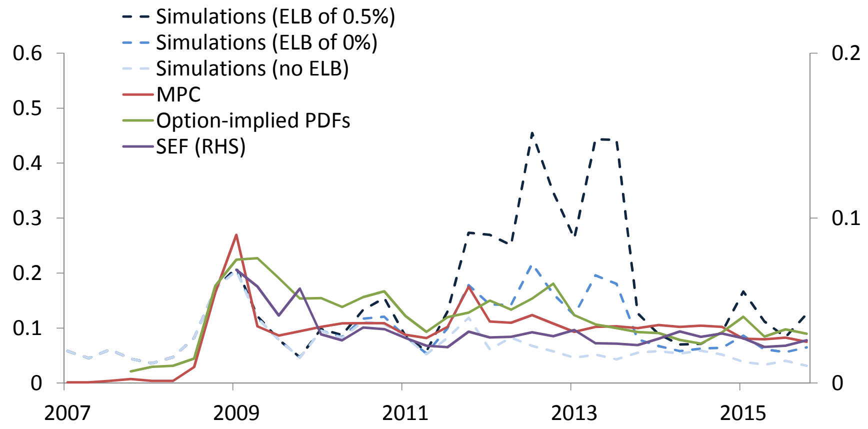

Notes: Option-implied PDF measures have been converted into quarterly averages to be consistent with other measures of deflation probability.