

Risk shocks and monetary policy in the new normal[☆]

Martin Seneca^a

April 2019

^a*Bank of England, Threadneedle Street, London EC2R 8AH, United Kingdom
martin.seneca@bankofengland.co.uk*

Abstract

Risk shocks give rise to a trade-off for monetary policy between inflation and output stabilisation in the canonical New Keynesian model if they are large relative to the distance between the nominal interest rate and its lower bound. The trade-off inducing effects operate through expectational responses to the interaction between the perceived volatility of conventional level shocks and the available monetary policy space. At the same time, a given monetary policy stance becomes less effective. Optimal time-consistent monetary policy therefore calls for potentially sharp cuts in interest rates when risk is perceived to be elevated, even if this risk does not materialise in any actual disturbances to the economy. The new normal for monetary policy may be one in which policymakers should both constantly lean against a tendency for inflation expectations to anchor below target—operating the economy above potential in the absence of disturbances—whilst accepting that inflation will settle potentially materially below target, and respond nimbly to changes in public perceptions of economic risk.

Keywords: risk shocks, uncertainty, zero lower bound, monetary policy.

JEL Codes: E52, E58.

[☆]The views expressed are solely those of the author and do not necessarily reflect those of the Bank of England or its committees. A previous version of this paper has appeared as Bank of England Staff Working Paper No. 606 (2016). I thank, without implicating, seminar participants at the University of Sheffield, the Bank of Canada, Danmarks Nationalbank, ISCEF 2016 in Paris, CEF 2016 in Bordeaux, SNDE 2017 in Paris, SED 2017 in Edinburgh, EEA 2017 in Lisbon, Workshop on Time-Varying Uncertainty in Macro 2017 in St Andrews, RES 2018 in Brighton, as well as Michal Brzoza-Brzezina, Menzie Chinn, Ambrogio Cesa-Bianchi, Richard Dennis, Jesús Fernández-Villaverde, Andrea Ferrero, Jordi Galí, Richard Harrison, Keith Kuester, Luisa Lambertini, Roland Meeks, Valerie A. Ramey, Christoph Grosse Steffen, Mathias Trabandt, Ricardo Reis, Garry Young, and Matthew Waldron for helpful conversations.

1. Introduction

During the Great Moderation, there was a general consensus that spells at the zero lower bound (ZLB) would be rare and short (see e.g. Coenen et al., 2004, Reifschneider and Williams, 2000, and Schmitt-Grohé and Uribe, 2010). Over the past decade, this pre-crisis consensus has been revised in light of the incoming data (e.g. Chung et al., 2012, Kiley and Roberts, 2017, and Williams, 2014). Equilibrium real interest rates are now widely expected to recover only to levels that fall short of historical averages—reducing the scope for future cuts in policy interest rates—and disturbances are expected to be larger—increasing the occasional need for such cuts. In the future, policy rates are deemed likely to hit their lower bounds more frequently than previously assumed. Nevertheless, an optimistic view holds that unconventional monetary policies such as quantitative easing and Odyssean forward guidance can be relied upon as substitutes for conventional reductions in policy rates (e.g. Bernanke, 2017, Harrison, 2017, Kiley, 2018, and Reifschneider, 2016). Whilst the ZLB may bind from time to time, monetary policy’s extended toolkit will rarely be constrained according to this view. But suppose the public are not fully convinced by such assurances. Suppose the economy recovers to a ‘new normal’ (in the terminology of El-Erian, 2010) in which people occasionally find reason to worry that policymakers may not be able to respond to future adverse disturbances with sufficient monetary stimulus. How should monetary policy be conducted in such an environment? Are the prescriptions for good monetary policy in ‘normal times’ developed during the Great Moderation sufficient guideposts for determining the appropriate stance of policy?

This paper points to two key differences that set monetary policy in the ‘new normal’ apart from that of the greatly moderated pre-crisis economy. First, when risk is high relative to the available monetary policy space, policymakers should operate

the economy above its efficient potential in normal times. This stimulatory bias leans against a tendency for inflation expectations to re-anchor too far below target, but it does allow inflation to settle potentially materially below target in the absence of disturbances. Welfare may be improved by appointing an independent central banker with a slightly higher inflation target than the social optimum. Second, because of constraints on monetary policy alone, changes in the public's perception of risk affect the appropriate stance of monetary policy through time-variation in the appropriate trade-off between inflation and real stability. A spike in uncertainty, for example, has a negative cost-push effect and makes a given stance of monetary policy less effective. Potentially sizeable changes in interest rates may be warranted even if risk does not materialise in any actual disturbance to the economy. This is in sharp contrast to conventional guidelines that monetary policy should respond only to disturbances that have actually occurred (or are fully anticipated), or at least that risk should only affect the stance of monetary policy to the extent that it affects precautionary behaviour by households and firms.

I derive these results in a simple version of the canonical New Keynesian model. The monetary policy design problem for this model served as the 'science of monetary policy' for the Great Moderation (Clarida et al., 1999), and it remains the theoretical foundation for the kind of flexible inflation targeting effectively practised by major central banks today (Svensson, 2010). Throughout, I focus on responses under optimal time-consistent monetary policy. The period-by-period nature of decision-making under discretion makes it a realistic description of the actual conduct of monetary policy in a flexible inflation targeting regime (see e.g. Bean, 2013). In particular, policymakers set interest rates policy meeting by policy meeting to achieve good outcomes given their operational targets. Neither do they follow an instrument rule mechanically, nor do they commit both themselves and future incumbents to a policy plan that will later

turn out to be undesirable. In line with conventional wisdom, policymakers cannot bootstrap the economy out of a ZLB episode by promising a future economic boom as advocated by Eggertsson and Woodford (2003).¹ As noted by Kiley (2018), it is doubtful that any central bank has attempted such purely Odyssean forward guidance in response to binding lower bounds on short-term policy interest rates.

Specifically, I solve for the risky steady state and I study optimal responses to risk shocks around that steady state in a quasi-linear version of the New Keynesian model augmented with a ZLB. In line with the definition in Coeurdacier et al. (2011), the risky steady state is the point at which the economy settles when disturbance have abated but agents are aware that innovations may occur in the future. The *risky* state state differs from the perfect-foresight or *deterministic* steady state, in which agents do not consider the possibility of future shocks. By risk shocks I mean changes in the standard deviations of conventional level shocks in the model. I trace out responses to such changes along the zero-shock path, i.e. trajectory of the economy over time in which innovations to level shocks do not actually occur. At it were, nothing actually happens in this paper. But crucially, agents remain aware that level shocks could hit the economy at any time in the future and take the economy off the zero-shock path. For analytical tractability, I maintain the assumption that agents form expectations at any given point in time in the belief that current risk levels will persists. With these assumptions, the effects of risk shocks can be thought of as the economy's responses to changes in a broad notion of the public's perception of risk.

I refer to these changes in second moments as risk shocks following the traditional Knightian distinction between risk and uncertainty.² Agents in the model economy

¹Adam and Billi (2006) show that risk increases welfare gains from time-inconsistent policy plans if policymakers could find ways to credibly commit to them.

²While the macroeconomic literature often follow Bloom (2009) in referring to second moment shocks as uncertainty shocks, this term is also used to describe a range of related phenomena, see

have well-defined probability distributions of economic shocks in mind when making decisions, and at any given point in time these distributions coincide with the actual ones. Probabilities are only subjective in the limited sense that agents do not foresee future changes in the levels of standard deviations. We can think of changes in these probabilities as changes in perceived risk only because there is an observational equivalence between the effects of changes in actual and subjective probabilities along the zero-shock path. Elevated risk in the model may well stand in for Knightian uncertainty in reality, but I make no attempt to model such uncertainty about probability distributions directly as, for example, in the work by Ilut and Schneider (2014) and Masolo and Monti (2017).

It is important to note that the macroeconomic risk shocks considered here are very different from the cross-sectional risk shocks analysed by Christiano et al. (2014).³ In their paper, a ‘risk shock’ refers to a disturbance to the *ex post* realisation of the dispersion of the quality of capital acquired by entrepreneurs. When this dispersion widens, the agency problem associated with financial intermediation becomes more severe. As credit spreads increase, entrepreneurs demand less capital and aggregate demand contracts for a given stance of policy. In the simple New Keynesian model I consider, such a scenario would correspond to a negative *level* shock to the efficient-equilibrium real rate of interest. To keep aggregate demand in line with the economy’s supply potential, monetary policy would have to counter a cross-sectional risk shocks with a looser policy stance.

The simple quasi-linear structure I consider comes with the benefit that risk affects the economy exclusively through its interaction with the ZLB. Schmitt-Grohé

e.g. Kozeniauskas et al. (2018) for a recent discussion. Fernández-Villaverde et al. (2011) also refer to second moments shocks as changes in risk, and LeRoy and Singell (1987) discuss the Knightian terminology.

³Similarly, it is different from the volatility considered by Arellano et al. (2018).

and Uribe (2004) show that risk does not affect decision rules to a first-order approximation, and only the constant term up to second order, for a general class of models *without* inequality constraints. This class of models includes the textbook New Keynesian model without a ZLB. In my analysis of the New Keynesian model, risk affects decision rules to a first-order because of the additional inequality constraint on policy rates—and for that reason only. Second moment shocks have real effects in the model through expectations if and only if the ZLB is binding in some conceivable states of the world. This allows me to focus on the defining feature of the new environment—an expected regular recurrence of a binding ZLB—and its implications for monetary policy in normal times, without conflating them with effects of risk not stemming directly from a potentially binding ZLB.

I solve the model following the stochastic extended path approach in Evans et al. (2015). On the one hand, the approach follows Adjemian and Juillard (2013) in relaxing the assumption of certainty equivalence maintained in the extended path procedure first proposed by Fair and Taylor (1983). On the other hand, a number of assumptions are imposed to simplify the analysis and focus on the expectational effects of risk directly caused by the ZLB.⁴ Compared to a full global solution procedure, the approach relies on (i) a first-order approximation of behavioural relations with the ZLB imposed as the only non-linearity (through which risks have first-order effects); (ii) the absence of intrinsic persistence (which could occur, for example, because of habits in consumption, indexation in pricing, or capital accumulation); (iii) optimal monetary policy under discretion (so that policymakers do not seek to affect expectations through commitments); and (iv) the approximation of continuous autoregressive shock

⁴These simplifications avoid the curse of dimensionality that forces Adjemian and Juillard (2013) to prune the tree of forward histories when calculating expectations. See Maliar et al. (2015) for a discussion as well as a generalised procedure.

processes as discrete-space Markov chains. As these assumptions combine to eliminate state variables, the model can be solved in a simple recursive procedure, providing a mapping between risk and macroeconomic outcomes. The assumption that agents do not foresee future changes in risk further allows me to use the approach to illustrate how those macroeconomic outcomes change—and how monetary policy optimally responds—as risk changes over time.

The paper is organised as follows. First, Section 2 discusses the relation to the literature. Section 3 then describes the model, Section 4 the solution method, and Section 5 the calibration. Section 6 presents the risky state state, and Section 7 shows responses to risk shocks. Section 8 presents normalisation scenarios after a ZLB episode. Finally, Section 9 concludes.

2. Relation to the literature

The analysis follows previous studies of the implications for discretionary monetary policy of deviations from certainty equivalence in New Keynesian models with a ZLB. Adam and Billi (2007) and Nakov (2008) first showed how the interaction of risk and the ZLB may give rise to a negative bias in expectations in the New Keynesian model's stochastic equilibrium. Both papers illustrate how this bias amplifies the economy's responses to negative shocks to the level of the equilibrium real rate of interest and leads to trade-offs when monetary policy is driven close to the ZLB. They show how gains from commitment are significantly larger as a result when risk is taken into account. In a recent application, Nakata and Schmidt (2014) further suggest that the skew in expectations provides a justification for a weight-conservative central banker. Similarly, Evans et al. (2015) emphasise that higher risk generally calls for looser monetary policy when the ZLB may bind. They find that lift-off from a ZLB episode should be delayed when agents are concerned about the risk of future episodes.

Compared to these seminal papers, my contribution is, first, to characterise the economy’s risky steady state when monetary policy’s room for manoeuvre may be deemed inadequate in normal times. Writing during the ‘old normal’ when ZLB episodes were expected to be rare, Adam and Billi (2007) and Nakov (2008) did not pay close attention to steady-state outcomes or distinguish clearly between different notions of the steady state and average outcomes. Instead, they emphasised the amplification of effects following a large persistent fall in the equilibrium real rate of interest from its normal level.⁵ Second, I trace out dynamics around the risky state state following changes in risk, illustrating how a risk shock propagates through dynamics in expectations in a way that induces time-varying trade-offs for monetary policy. I show that this time-variation calls for monetary policy responses even in the absence of actual disturbances (such as level shocks to the equilibrium real rate), and I show that these responses do not depend on the source of risk (i.e. the particular shock process for which risk has changed). Moreover, I show that impulse responses to positive and negative risk shocks are asymmetric around the risky steady state because of a non-linearity in the mapping between risk and economic outcomes.

The simple quasi-linear framework sets my analysis apart from the recent related work by Basu and Bundick (2015, 2017). In these papers, second moment shocks have real effects through precautionary savings behaviour by households. The authors show that these higher-order effects are greatly amplified at the ZLB as monetary policy fails to respond to them.⁶ Moreover, a feedback mechanism sets in: monetary policy’s inability to respond to what amounts to a further fall in the equilibrium real rate re-

⁵Also, Evans et al. (2015) operate in a part of the parameter space for which the solution to the New Keynesian model with a ZLB is explosive, pinned down only by the assumed expectational horizon.

⁶Basu and Bundick (2015, 2017) compute third-order approximations to behavioural relations to capture effects operating through precautionary behaviour when monetary policy is unconstrained, while they turn to global solution methods when imposing the ZLB.

duces the expected mean of outcomes, in turn inducing further precautionary saving.⁷ In my analysis, by contrast, risk shocks affect the economy solely because of changes in expected mean paths for macroeconomic variables for a given equilibrium real rate of interest. The quasi-linearity thus serves both to illustrate that risk shocks may have significant real effects through the ZLB without precautionary behaviour and to isolate such direct effects from those effectively operating through shifts in the equilibrium real rate.⁸ There are two other significant differences. First, Basu and Bundick (2015, 2017) either impose a simple instrument rule for monetary policy or allow for commitment policies that essentially assume away the effect of the ZLB on mean expectations. Instead, I consider policy prescriptions under the more realistic assumption that optimising policymakers do not have access to such commitment devices. Second, Basu and Bundick (2015, 2017) are concerned with the contributions of second-moment shocks to macroeconomic volatility over the past. Specifically, they explain how elevated macroeconomic risk could have contributed significantly to the Great Recession, while changes in second moments amounted to inconsequential background noise before the financial crisis. Hence, they distinguish sharply between periods when the ZLB is both non-binding and a negligible risk, and periods when the ZLB is in fact binding. By contrast, I derive practical normative prescriptions for monetary policy in an environment in which the ZLB is not necessarily binding but nevertheless constantly looming as a cloud on the horizon.

My paper is also closely related to the contemporaneous and independent work by Hills et al. (2016). In their paper, the authors also compare the risky and deterministic

⁷See also Bloom (2009) and Nakata (2017) for discussions of how other non-linearities may give rise to real effects from movements in risk that may be amplified by the ZLB. Fernández-Villaverde et al. (2015) and Johannsen (2014) suggest that uncertainty about fiscal policy in particular has larger implications when monetary policy is constrained.

⁸See Paoli and Zabczyk (2013) for an analysis of the effect of precautionary saving on the level of the effective equilibrium real rate of interest.

steady states in a stylised New Keynesian model, finding results fully consistent with mine. But they then go on to quantify the difference between these steady states in a richer model calibrated to match key features of the US data over the past decades. I instead focus on the risky steady state in a hypothetical new normal, and I proceed to consider the implications of potential time-variation in the underlying level of macroeconomic risk. Moreover, while Hills et al. (2016) consider outcomes under a simple monetary policy rule taken to be representative of monetary policy’s reaction pattern in the past, I derive normative prescriptions by solving for optimal monetary policy responses for the future. Finally, while they solve their model in non-linear form using global methods, I keep my analysis within a quasi-linear framework to facilitate a direct comparison of these prescriptions with those of the influential ‘science of monetary policy’ (Clarida et al., 1999).⁹ I thus take the two papers to be complementary, each providing a different perspective on the interactions between risk and the ZLB.

More broadly, my paper relates to a growing literature on the effects of risk and uncertainty. Following the work of Bloom (2009), there has been a surge of interest in this issue. Whilst the empirical literature has struggled to identify structural risk shocks from the volatility and forecast disagreement measures that are usually taken to be proxies for risk and uncertainty, theoretical work has provided clear channels through which risk shocks may affect the economy, see e.g. the survey by Bloom (2014). The expectational mechanism at the core of my analysis is a ‘bad news channel’ in the terminology of Bernanke (1983). It arises because monetary policy is sometimes unable to provide sufficient stimulus in response to large adverse disturbances while it can always act to contract the economy when needed. The model’s prediction that the effect of a given risk shock is larger, the closer the economy is to the ZLB, is in

⁹See e.g. Carney (2017) for an application to monetary policy in practice.

line with the empirical evidence provided by Caggiano et al. (2017) and Plante et al. (2018). Similarly, the finding that monetary policy is less effective when risk is high is consistent with the evidence in Aastveit et al. (2017). Finally, Caggiano et al. (2018) provide recent evidence that monetary in the United States has indeed responded to changes in aggregate risk.

3. The model

The model is the canonical New Keynesian model, expressed in log-deviations from its deterministic steady state, extended with a ZLB on interest rates. In addition to a specification of monetary policy, it consists of the equations:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + u_t \tag{1}$$

$$x_t = E_t x_{t+1} - \frac{1}{\varsigma} (i_t - E_t \pi_{t+1} - r_t^*) \tag{2}$$

$$i_t + i^* \geq 0 \tag{3}$$

where E_t is the expectations operator, π_t is inflation at time t in deviation from its target π^* , x_t is the output gap defined as output in deviation from its efficient level, and i_t is the nominal interest rate in deviation from its normal deterministic steady-state value i^* . The first equation is the New Keynesian Phillips curve, the second is the forward-looking IS curve, and the third imposes the ZLB. The model is derived from its microfoundations by Galí (2008) and Woodford (2003) among others.

There are two shock processes in the model. The term u_t is a cost-push process, and r_t^* is the efficient equilibrium real interest rate in deviation from its steady state level $r^* \approx i^* - \pi^*$. I assume that the latter is the sum of a deterministic but potentially time-varying component ρ_t and a stochastic process ϵ_t so that $r_t^* = \rho_t + \epsilon_t$. Both the stochastic component of the equilibrium real interest rate and the cost-push shock are

given as first-order autoregressive processes with zero-mean Gaussian innovations:

$$\epsilon_t = \mu_\epsilon \epsilon_{t-1} + \nu_{\epsilon,t} \quad (4)$$

$$u_t = \mu_u u_{t-1} + \nu_{u,t} \quad (5)$$

where $\nu_{\epsilon,t} \sim N(0, \sigma_{\epsilon,t}^2)$ and $\nu_{u,t} \sim N(0, \sigma_{u,t}^2)$. I allow the standard deviations of the innovations to vary over time as indicated by the time subscripts in $\sigma_{\epsilon,t}$ and $\sigma_{u,t}$.

I define a risk shock as a change in one or both of these standard deviations. The standard deviations are independent of each other and can also change independently. As a baseline, however, I consider the special case in which $\varsigma^{-1} \sigma_{\epsilon,t} = \sigma_{u,t} = \sigma_t$ with

$$\sigma_t = \sigma + \mu_\sigma (\sigma_{t-1} - \sigma) + \nu_{\sigma,t} \quad (6)$$

where $\nu_{\sigma,t}$ is an innovation to risk, and σ is an underlying level of risk in the absence of risk disturbances.¹⁰ In what follows, agents observe the current level of risk, but they do not expect further changes in this level to occur.

Under optimal policy under discretion, a policymaker, hypothetically unconstrained by the ZLB in (3), minimises the period loss function

$$L \propto \pi_t^2 + \lambda x_t^2 \quad (7)$$

each period subject to the Phillips curve in (1) while taking expectations as given. The loss function can be derived as a quadratic approximation of the utility of the representative household in the full New Keynesian model (again, see e.g. Galí, 2008, and

¹⁰I only consider realisations of risk that are strictly larger than zero, and I specify innovations to risk directly. In applications with stochastic volatility, the risk shock could be specified in logs to rule out non-zero realisations.

Woodford, 2003). The optimality condition takes the form of a conventional targeting rule

$$\pi_t = -\frac{\lambda}{\kappa}x_t \quad (8)$$

stating the optimal policy trade-off between inflation and the output gap. Following trade-off inducing disturbances, monetary policy seeks to keep deviations of inflation from target and of output from potential of opposite signs, letting inflation absorb more—and output less—of the adjustment the higher the weight on the output gap in the loss function and the flatter the Phillips curve (i.e. the higher the sacrifice ratio). The interest rate consistent with this optimal allocation can now be found from the IS curve in (2). Since the policymaker is, in fact, constrained by (3), the targeting rule in (8) is replaced by the Kuhn-Tucker conditions

$$(i_t + i^*) (\lambda x_t + \kappa \pi_t) = 0 \quad (9)$$

$$i_t + i^* \geq 0 \quad (10)$$

$$\lambda x_t + \kappa \pi_t \leq 0 \quad (11)$$

As shown formally in Appendix A, these conditions imply that the interest rate will be set to the maximum of the level consistent with satisfying the targeting rule in (8) and zero.

4. Solution method

I solve the quasi-linear version of the canonical model following the approach in Evans et al. (2015). I approximate the shock processes by independent Markov processes using the Rouwenhorst (1995) algorithm provided by Galindev and Lkhagvasuren (2010). I then solve the model backwards from a distant future period T , beyond which

there is no risk and all shocks are zero so that $E_t\pi_{t+1} = E_tx_{t+1} = 0$ for all $t > T$. In each step, I take expectations as given and calculate the unconstrained outcome under each policy regime for a state grid of values for the shock processes, (ϵ, u) . I then check if this outcome is consistent with the ZLB in (3) for each node in the grid. If so, I take the unconstrained outcome as the solution for this particular node. If not, I calculate the outcome from the model equations with $i_t = -i^*$ imposed. I then update the *ex ante* expectations of inflation and the output gap using the Markov transition matrices before progressing to the previous period. The solution consists of the $n_\epsilon \times n_\epsilon$ matrices for inflation, the output gap and the interest rate, to which this algorithm converges in the initial period $t = 0$. See Appendix B for further details.

The values for the nodes $(\epsilon = 0, u = 0)$ represent outcomes in the event that no non-zero shock has actually materialised. This converged zero-shock solution at $t = 0$ represents the risky steady state of the model as defined by Coeurdacier et al. (2011). This is the resting point to which the economy returns when all disturbances have dissipated. Non-zero realisations of shocks will of course continuously drive the economy away from this point, and the risky steady state potentially differs from the deterministic one exactly because it accounts for agents' expectations that such deviations will occur. Unconditional expectations are averages weighted by unconditional probabilities over outcomes across the state space. Given the Markov structure, these probabilities can easily be derived from the eigenvectors associated with the unitary eigenvalues of the transition matrices, see e.g. Ljungqvist and Sargent (2000). Appendix C provides further details.

I find responses to a risk shock by running a double loop. The outer loop moves forward from period $t = 0$, while the inner loop solves the model backwards from period T to the period of the current iteration of the outer loop. For each iteration of the outer loop, I reduce the value of σ_ϵ and/or σ_u from an initial spike according

to the assumed process. The economy’s responses to the risk shock is the sequence of zero-shock solutions found in the outer loop.

In line with the broad definition in Koop et al. (1996), this impulse response function measures the effects of the shock on macroeconomic variables over time. The conceptual experiment is a comparison of the profile for the economy when a risk shock hits at time t with a profile where risk stays at its baseline level, keeping all other shocks (in this case the level shocks) dormant. It is a traditional impulse response function with the caveat that agents have myopic rather than rational expectations about the risk shock process itself (while otherwise continuing to form rational expectations). The advantage of taking this approach is that risk does not become a state variable. It is also a natural starting point for considering occasional changes in risk perceptions: each period, economic agents simply assign a number to the level of risk they think is present in the economy. But it is a limitation that agents always expect risk to stay constant at a given point in time. Alternative specifications where agents are allowed to see a stochastic autoregressive profile for risk are, however, likely to generate similar qualitative responses in this quasi-linear model, and deriving them would add significant computational complexity to the analysis.¹¹

5. Calibration

I calibrate the model to fit a hypothetical ‘new normal’ distribution for desired policy rates—the short-term interest rates a policymaker would choose to implement if the ZLB were not a constraint—shown in dashed-dotted black lines in both panels in Figure 1. To motivate the choice of this distribution, the figure compares it with recent pre-crisis historical experience in the United States (left panel) and the United King-

¹¹Specifically, time-variation in transition probabilities makes the model non-stationary. See Maliar et al. (2015) for a rigorous analysis of such models.

dom (right panel). The normal probability density functions shown in solid blue lines share the means and standard deviations with the observed Federal Funds Rate and Bank Rate, respectively, from 1968 through 1992. Blue dotted lines show corresponding kernel density estimates. An observer looking back at these distributions around the time when 2% inflation targets were emerging in the early 1990s (see e.g. Svensson, 2010), would not have found much reason to worry about the ZLB. Policy rates had been very volatile over the past quarter of a century (standard deviations were 3.2pp and 2.9pp, respectively), but they had also been high (with means of 8.1% and 10.6%). The probability that interest rates should be negative would have seemed negligible. Over the subsequent 15 years, the distributions of policy rates shifted sharply to the left as inflation targeting became established (formally or informally), and inflation expectations anchored at lower levels, see the dashed red lines for normal approximations and dotted red lines for kernel density estimates. In isolation, the lower means (4.0% and 5.4%) would have increased probability mass below zero. But as volatility fell substantially at the same time, the risk of a binding ZLB did not appear to have increased significantly over the period.

Any Gaussian model informed by interest rate data over the two periods would have produced a probability of negative negative interest rates very close to zero, including both the semi-structural and the DSGE model that were increasingly used for such purposes, see e.g. Coenen et al. (2004), Schmitt-Grohé and Uribe (2010) and Chung et al. (2012). Moreover, assessments of the fit of models estimated using full-information techniques, for example by comparing fitted normals to kernel distributions (a tougher test than usually applied), would have pointed to a satisfactory if not perfect fit.¹² It is therefore not surprising that a consensus emerged in the pre-crisis period

¹²The sample periods in Figure 1 are such that estimated distributions are close to normal without detrending. In particular, Bank Rate is very close to normal in each of the two subsamples. The

that the ZLB should be of no great practical concern. As it turned out, however, the estimated models were too sanguine about prospects of a binding ZLB.

Instead of focusing on past distributions, I therefore consider a hypothetical scenario for the normal for monetary policy to which the economy may recover after the financial crisis. Compared to the historical distributions over the 1993-2008 period, it is defined by a further shift to the left and an increase in the spread for desired policy rates. The shift captures an assumed decline in the trend component of the efficient equilibrium real rate of interest. The higher spread reflects an end to the Great Moderation so that larger disturbances call for stronger policy responses than in the pre-crisis inflation targeting period. I set the average desired policy rate to 3%—in line with the forecasts emphasised by Reifschneider (2016), and the recent estimates in Kiley (2015), Laubach and Williams (2016) and Del Negro et al. (2017)—and I fix the standard deviation at 2.2, in between the values for the two pre-crisis subsamples for the US and the UK. Hence, macroeconomic risk is higher than during the Great Moderation, but not as high as in the preceding decades. With the lower average level of the policy rate, the moderate increase in the desired spread increases probability mass below zero to about 9%. In this ‘new normal’ scenario, therefore, the probability that policymakers would set policy rates to negative values if they could is clearly non-negligible.¹³

In the baseline calibration, I parameterise the model so that it features a distribution for the unconstrained optimal policy rate corresponding to the new normal distribution in Figure 1. The (annualised) inflation target is assumed to be $\pi^* = 2\%$ and the

picture is a bit less clear for the United States, where the data point as much to a gradual decline as to a single structural shift in the level around which the Federal Funds Rate fluctuates. The sharp response to the 2001 recession followed by rapid normalisation also results in a near-bimodal shape of the estimated kernel for the Federal Funds Rate over the 1993-2008 period.

¹³The assumption that desired policy rates are normally distributed is not restrictive. What matters for the qualitative results in the following is that probability mass below zero is non-negligible more than the specific shape of the distribution.

deterministic steady-state level of the real interest rate $r^* = 1\%$. The normal nominal interest rate is then approximately $i^* = 3\%$ with a discount factor $\beta = 0.9975$. The deterministic component of r_t^* is set to $\rho_t = 0$ for all t , and the inverse of the elasticity of intertemporal substitution is set to $\zeta = 1$ as is common in the literature for this poorly identified parameter (see e.g. Galí, 2008). The slope of the Phillips curve is assumed to be $\kappa = 0.02$. This value is at the lower end of the 0.02-0.05 range of empirical estimates collected by Woodford (2005), in line with the hypothesis that Phillips curves have (if anything) flattened in recent decades (e.g. Blanchard et al., 2015). The weight on the output gap in the loss function is similarly set to $\lambda = 0.02$. This is larger than the $\lambda = \kappa/\zeta$ imposed when the loss function is derived from household utility in the basic New Keynesian model, where $\zeta > 1$ is the elasticity of substitution between product varieties under monopolistic competition. But for conventional values of ζ around 6 and empirically plausible values of κ , the weight on output would be much smaller than actual mandates for monetary policy seem to imply. The assumed value of λ corresponds to a weight on output stabilisation of about a third in annualised terms, a reasonable interpretation of the degree of flexibility in inflation targeting in practice (see e.g. Carney, 2017, English et al., 2015, and Svensson, 2010). Moreover, the basic New Keynesian model is likely to underestimate the appropriate weight on the output gap, see e.g. Debortoli et al. (2016) and Walsh (2014). With an implied targeting rule for monetary policy with a slope of minus one, policymakers seek to let quarterly inflation and the output gap share the burden of adjustments to tradeoff-inducing disturbances equally. Finally, level shocks are assumed to be moderately persistent with $\mu_u = 0.25$ and $\mu_\epsilon = 0.75$. With these parameter values, an underlying level of risk given by $\sigma = 0.0027$ delivers a standard deviation of the unconstrained nominal policy rate of 2.2 when $\zeta^{-1}\sigma_{\epsilon,t} = \sigma_{u,t} = \sigma_t$ so that the desired policy rate is negative with probability 9%.

For comparison, I also consider a low-risk scenario with $\sigma = 0.0012$. With this low level of underlying risk, the dispersion of desired interest rates is kept low while the mean shifts down to 3%. Specifically, the standard deviation of desired interest rate is 1pp, similar to the level observed for the United Kingdom between 1993 and 2008. The probability that interest rate should be negative remains negligible in this case despite a low level of r^* . In Section 7, I show impulse responses to risk shocks in this low-risk scenario as well as in the baseline new normal.

Table 1 compares the ‘new normal’ baseline and the low-risk scenarios with alternative calibrations in which σ is set to fit the distribution of the unconstrained policy rate to the historical distributions in Figure 1. In the historical episodes, I match the mean by assuming that $r^* = 3.75$ in the 1968-1992 period, and that $\pi^* = 2\%$ between 1993 and 2016. In all of the alternative calibrations, the probability of negative interest rates are close to zero, rounding to 1% only for the United States. Even if the model is too stylised to capture the covariance structure of the data more broadly, the table shows that it gives a fairly good fit to observed inflation volatility across the historical periods. Hence, it does not seem unreasonable a priori that it can shine some light on inflation outcomes in hypothetical scenarios. Nevertheless, I make no claim to provide empirical accuracy in this simple framework. The numerical results should be taken as qualitative indication of the sign and order of outcomes rather than quantitative estimates.

In what follows, I assume that the lower bound is exactly zero as specified in (3). The unconstrained model is hardly a good guide for monetary policy when the desired level of interest rates is negative with a non-negligible probability as in the hypothetical new normal. Absent substantial reform to the payment system, it is improbable that policymakers can persistently drive interest rates into negative territory (e.g. Rogoff, 2015). It is more likely that unconventional policies such as quantitative easing may

act as substitutes for negative short-term interest rates (e.g. Haldane et al., 2016). The interest rate in the model may best be thought of as a *shadow* rate implicitly incorporating the effects of unconventional policy tools (Black, 1995). The shadow rate may be negative if these tools are operational whenever the ZLB binds. But if either the availability or the effectiveness of such tools is limited, the shadow rate will also be bounded from below at some level less than zero. With the ZLB imposed, I effectively assume that cash has not been phased out to allow for negative interest rates, and that unconventional policies cannot act as perfect substitutes for negative interest rates. The analysis becomes irrelevant if either of these assumptions is fully reversed so that interest rates can go negative without difficulty or side-effects, or unconventional policy tools can be relied upon as perfect substitutes for changes in policy rates. But the conclusions hold for any constellation of the distribution of the desired shadow rate and its effective lower bound such that monetary policy is constrained in its ability to stimulate the economy with a non-negligible probability. I take such a constellation to be the defining feature of the new normal.

6. Risky steady state

The first row in Table 2 shows risky steady state outcomes for the interest rate (i), inflation (π) and the output gap (x) in the full model with the ZLB under the baseline calibration. In the new normal, public perceptions of risk are high enough that, because of the ZLB, policymakers are not expected to be able to respond sufficiently to some of the large negative disturbances that are deemed likely to hit the economy in the foreseeable future. By contrast, the public know that monetary policy can always be tightened appropriately, also in response to large inflationary shocks. This asymmetry introduces a negative skew in expectations as first emphasised by Adam and Billi (2007) and Nakov (2008). Under optimal discretionary policy, policymakers lean against the

tendency for inflation expectations to anchor below target by operating the economy above potential in normal times through a stimulatory bias in policy rates. Such a long-run effect on real output is feasible because the monetary policy stance interacts with inflation expectations to determine real interest rates in the risky steady state. In this sense, monetary policy is no longer neutral in the new normal’s long-run. But there are limits to policymakers’ willingness to overheat the real economy, and subdued inflation expectations are allowed to weigh on the price-setting of firms. Consequently, inflation settles about 20 basis points below target.

Hence, the point at which the economy comes to rest when shocks have faded away does not coincide with the deterministic steady state, in which inflation is on target, the output gap is closed, and the interest rate is at its normal level. Moreover, as Table 2 also shows, unconditional expectations deviate from steady-state values. In expectation, output falls short of potential because of spells at the ZLB. These episodes drive the average interest rate above its risky steady-state level, and expected inflation falls somewhat further below target. Notice also that, because of monetary policy’s inability to deliver the desired stimulus, the frequency of ZLB episodes is higher than the probability that interest rates should be negative in the unconstrained model. In the new normal, a 9% probability that desired policy rates are negative translates into a 14% probability that the ZLB is binding.

The remaining rows in Table 2 illustrate the sensitivity of these statistics to assumptions about risk and the available monetary policy space. In the baseline, I assume that $\varsigma^{-1}\sigma_{\epsilon,t} = \sigma_{u,t} = \sigma_t$ for convenience. But in general, the levels of risk for the two shock processes may not be related by a simple multiple. The second row shows the effect of completely removing the risk of cost-push shocks, while keeping the risk of r^* shocks at the baseline value. The third row shows the opposite case without a perceived risk of r^* shocks. In both cases, the risky steady state deviates from the deterministic

one with inflation settling below target. The marginal contributions of the two shocks are similar, but the deviations are much smaller with inflation rates of 1.93 and 1.98, respectively. In the new normal, agents are particularly concerned about the inability of policymakers to respond when large adverse disturbances to the cost-push process and the equilibrium real rate coincide.

Higher risk for individual shocks may, however, result in similar biases as in the benchmark. Increasing σ_ϵ to about 0.0032 when the risk of cost-push shocks is absent (row 4), or σ_u to about 0.0043 when the risk of r^* shocks is negligible (row 5), leads to similar biases in inflation and the output gap as under the baseline calibration. Inflation may fall short of target in the new normal regardless of the source of risk. Notice also that, in all of these cases, policymakers have substantial room for manoeuvre. Inflation falls short of target in normal times only because private agents worry that the ZLB may bind in the future. The re-anchoring of inflation expectations occurs whenever risk is perceived to be high relative to the available monetary policy space.

For a given level of risk, the effect on expectations therefore also depends on the normal distance to the ZLB as illustrated in the remaining cases in Table 2. The closer the economy operates to the ZLB, the larger are the effects of risk on outcomes. If the distance to the ZLB is reduced by about 25 basis points, either because the equilibrium real rate of interest is lower (row 6), or because monetary policy targets a lower inflation rate (row 7), risky steady-state inflation falls below target by a further 15 basis points. When inflation in the deterministic steady state itself is lower, inflation settles almost 40 basis points lower than in the baseline. By contrast, if i^* increases to about 3.25%, the negative bias in inflation is reduced by 7-8bp. With a higher inflation target, the component of i^* that can be chosen by policymakers, inflation settles around 2.1% (row 9). To fully eliminate the negative bias in inflation, however, policymakers will have to target a rate of inflation above 4% (row 10).

Since inflation is closer to 2% and the output gap is closer to zero when monetary policy targets an inflation rate of 2.25% (as in row 9 in Table 2), welfare losses are unambiguously lower in the risky steady state when evaluated using a loss function (7) that is centred around a 2% optimal rate of inflation. In the new normal, welfare may be improved by appointing an independent central banker with a slightly higher operational inflation target than the social optimum. But as shown in Figure 2, setting the operational target too high, for example at the level which eliminates the negative bias in expectations, comes with substantial welfare costs if the socially optimal rate of inflation is 2%. Specifically, only operational targets in the open interval (2.00%, 2.30%) result in smaller welfare losses both in the risky steady state and in expectation, where expected period welfare losses are calculated as the probability weighted sum of losses across the state space. Provided that monetary policy responds optimally, the simple analysis does not by itself provide a case for increasing the inflation target to, say, 3% or 4% as suggested as a potential response to low equilibrium interest rates e.g. by Ball et al. (2016), Blanchard et al. (2010), Krugman (2014) and Williams (2009).

Finally on the risky steady state, Figure 3 illustrates that the solution does not depend on the horizon in the extended path procedure, set to $T = 1000$, or on its initialisation with zero one-period ahead conditional expectations of inflation deviations from target and output gaps. The figure shows the zero-shock paths in the iterative solution from period $t = T$ to $t = 0$, i.e. conditional on the particular realisation of the two shock processes shown in the top-left panel in which level shocks do not actually occur. For the new normal initialised with zero conditional gap expectations at time T (solid blue lines), the algorithm converges with significantly fewer iterations than 1000. Hence, results do not depend on the choice of T as long as it is sufficient large. If the procedure is re-initialised with the true conditional expectations, convergence will be immediate (dashed-dotted lines). If the initial guesses are further away from the

true conditional expectations, convergence will be slower, as illustrated for an arbitrary initial conditional inflation expectation of 2.5% (dotted blue lines). But the procedure converges to the same solution. For the low-risk case initiated with zero one-period ahead expectations, convergence at the deterministic steady state is immediate (dotted red lines).

7. Impulse responses to risk shocks

Now suppose that risk may vary over time. How does the economy adjust to changes in the perception of risk? To build intuition, I first present impulse responses to the baseline risk shock in (6) starting from a low-risk steady state. The low-risk case differs from the new normal only in that $\sigma = 0.0012$. With this low level of underlying risk, the probability that interest rates should be negative remains negligible despite a low level of r^* (see the second row in Table 1). The risky steady state therefore practically coincides with the deterministic steady state. I then consider a risk shock for each of the two level shocks in turn, before turning to a baseline risk shock in the risky steady state implied by the new normal scenario.

7.1. *The case of low underlying risk*

Solid blue lines in Figure 4 are impulse responses to the baseline risk shock along the zero-shock path starting from the low-risk steady state. The risk shock represents a scenario in which risk is temporarily elevated so that agents expect shocks to be drawn from distributions with higher spreads for some time in the future. But the economy is not actually hit by any level shocks along this adjustment path; it is only the perception of risk that changes. When risk spikes up, agents begin to worry about the monetary policymaker's inability to respond to large adverse shocks as a consequence of the ZLB. Therefore, inflation expectations fall short of the inflation target, and

output expectations of potential. By (1), the risk shock has a negative cost-push effect: for any given level of the output gap, inflation falls in response to lower inflation expectations. This effect induces a trade-off for the policymaker as reflected in the targeting rule in (8). Under optimal discretion, the policymaker loosens policy enough to bring output above its efficient potential. The expansion in the economy works to limit the fall in inflation and appropriately balance deviations from target with real economic outcomes. As risk falls back, the ZLB becomes less of a concern and the economy gradually returns to the low-risk steady state.

The dynamics induced by the risk shock are similar to those following a level cost-push shock in the New Keynesian model. But with a risk shock, the interest rate has to be reduced more to achieve the optimal balance between inflation and the output gap. There are two reasons for this. First, lower inflation expectations raises the real interest rate for a given level of the nominal rate. And second, since output expectations have also been adversely affected by the risk shock, policy needs to bring about a lower real interest rate to boost aggregate demand through (2). In this sense, the increase in risk has made monetary policy less effective.

Importantly, a trade-off arises in uncertain times even if shocks do not actually happen. The only prerequisite is that the risk shock is large enough that the ZLB becomes a concern. Small increases and reductions in risk around the low-risk steady state leave economic outcomes unaffected. Of course, the closer the economy operates to the ZLB, the more risk shocks become ‘large’ in this sense. Reversely, if underlying risk is high, the trade-off for monetary policy becomes a permanent feature of the economy as in the new normal described above.

7.2. On the sources of risk

Figure 4 also shows the effects of a positive risk shock around a low-risk steady state for each of the two shocks in turn. Again, the standard deviations for the two shock processes moves together in the baseline risk shock mainly for convenience. There is no reason to rule out a priori that risk cannot move independently for the two types of shocks. Qualitatively, however, the economy responds in the same way to the two individual shocks. Spikes in risk lead to cost-push effects both when risk is elevated for r^* only (dashed red lines) and for the cost-push process only (dashed-dotted black lines). It is simply the numerical increases in risk required to induce similar quantitative dynamics that are different (top-left panel). In both cases, responses are driven by an increase in the likelihood that policymakers cannot provide sufficient stimulus. But the sources of the potential adverse shocks are immaterial. For cost-push shocks, a negative bias in inflation expectations occurs because monetary policy cannot always engineer a sufficient boom in the economy to prevent inflation from falling too much after large negative shocks. In the case of r^* shocks, sufficiently negative realisations make it impossible for monetary policy to provide enough support for aggregate demand to keep up with supply.¹⁴ A trade-off arises for monetary policy as the prospect of such demand-driven recessions feed into inflation expectations when risk is elevated.¹⁵

Notice, however, that shocks to r_t^* are not necessarily demand shocks in the traditional sense. In the canonical New Keynesian model, fluctuations in the efficient equilibrium real rate of interest are driven by changes in the expected growth rate of total factor productivity in addition to changes in preferences and exogenous spending,

¹⁴If monetary policy were unrestricted by the ZLB, shocks to r_t^* could always be perfectly offset by an appropriate stance of policy. In this case, the output gap would remain closed, and inflation would be on target by the divine coincidence (Blanchard and Galí, 2007).

¹⁵As illustrated by Adam and Billi (2007), a trade-off arises for persistent negative *level* shocks to r^* of an intermediate size for a similar reason: when the economy moves closer to the ZLB, more future shocks can potentially cause a recession for a given level of risk.

see e.g. Galí (2008). Heightened uncertainty about the future growth potential of the economy is therefore an example of a risk shock to r_t^* . A scenario in which such an increase in perceived risk is associated with a fall in expected future growth rates would correspond to a combination of a positive risk shock and a negative level shock to r^* in this framework.

Notice also that the analysis ignores the effect of risk on precautionary saving. Similarly, the simple New Keynesian model does not allow for negative demand effects from the option value associated with postponing irreversible investments when risk is high (Bernanke, 1983). The adverse effects from risk shocks arise solely because of adjustments to the expected mean paths for output and inflation when monetary policy is constrained. The advantage of this assumption is that the effects stemming from the constraints on policy are clearly separated from higher-order behavioural effects. But in reality, a risk shock of any kind is likely to be accompanied by what would be a negative level shock to r_t^* in this framework as households seek to build a buffer stock of savings while firms put investment projects on hold.¹⁶

7.3. Risk shocks in the new normal

Starting from the new normal steady state, variation in risk have both positive and negative cost-push effects as shown in Figure 5. Responses to a positive shock (solid blue lines) are as before, except that the economy reverts to the risky steady state with a negative bias in inflation. An increase in risk increases the bias in expectations, worsening the trade-off for monetary policy. But a negative risk shock (dashed red lines) now has a positive cost-push effect. As risk falls, agents stop worrying about the ZLB, and inflation expectations realign with the inflation target. Policymakers

¹⁶The second-moment shock considered by Basu and Bundick (2015, 2017) would correspond to such a combination of a level shock and a risk shock to r^* in a version of the quasi-linear model with a simple monetary policy rule. See also Paoli and Zabczyk (2013).

increase interest rates in response, while the output gap closes. Gradually, as risk returns to its underlying level, the economy reverts to the high-risk steady state. As the responses show, optimal monetary policy in the new normal responds nimbly when risk perceptions change in both directions.

The asymmetry in the responses to positive and negative risk shocks around the risky steady state reflect a non-linearity in the effect of risk on economic outcomes as illustrated in Figure 6. With low levels of risk, the economy operates in the deterministic steady state in the absence of level shocks. As risk increases, the ZLB eventually becomes binding in some states of the world. For small increases, the effects are small. But as risk increases further, the frequency of ZLB episodes increases and effects begin to accelerate. Beyond a certain critical point (around $\sigma = 0.0032$ under the baseline calibration), interest rates are driven to the ZLB by the bias in expectations itself. In this unpleasant scenario, negative expectations—caused by a concern about the policymaker’s inability to respond to adverse shocks—become self-fulfilling as the policymaker is, in fact, unable to respond sufficiently to these expectations because of the ZLB. As a result, the economy enters a downward spiral with hyperdeflation and a collapse of output.¹⁷

8. Normalisation scenarios

To illustrate the implications of a binding ZLB for the propagation of risk shocks, Figure 7 shows a normalisation scenario in which the economy is gradually recovering from a ZLB episode caused some time in the past by a large and persistent negative shock to the level of the equilibrium real interest rate. The nature of this initial shock, say a financial crisis, is well understood by agents in the economy by now. Specifically,

¹⁷The possibility of explosive dynamics corresponds to the potential non-existence of equilibria analysed by Mendes (2011) and Nakata and Schmidt (2014).

the deterministic component driven by ρ_t is known to follow the path shown in the top-right panel of Figure 7 (dotted green line) so that the equilibrium nominal interest rate gradually returns to a new normal level of 3%. Uncertainty surrounding this recovery is perceived to be low ($\sigma = 0.0012$).

At around period $t = 4$, the efficient nominal interest rate turns positive and the policymaker, who operates under optimal discretion, is preparing to lift interest rates off the ZLB. In the absence of risk, the policymaker would simply follow the equilibrium interest rate on its trajectory back towards normal levels once it exceeds the ZLB.¹⁸ But as long as the equilibrium interest rate is this close to the ZLB, even small shocks are ‘large’, and the possibility that a shock may drive the economy back to the ZLB in the future is sufficient to optimally delay lift-off even when risk is low.¹⁹

Now suppose that agents suddenly become more uncertain about economic prospects, perhaps reflected in turmoil across financial markets. Specifically, suppose the economy is hit by a baseline risk shock corresponding to the one shown in in Figure 4 at time $t = 5$, just as lift-off was supposed to take place in the absence of any disturbances to the economy. Now that the economy is close to the ZLB, the impact effect of the risk shock on expectations is larger than before as the monetary policymaker is constrained by the ZLB in its response to the shock. As shown in Figure 7 (dashed red lines), inflation falls more as a consequence, and lift-off from the ZLB is further delayed. Now because of the binding ZLB, output also falls further below potential. Only as risk abates will the optimal interest rate path catch up with the equilibrium rate. The longer risk stays elevated, i.e. the more persistent the risk shock, the longer lift-off is

¹⁸This corresponds to the perfect foresight case analysed by Adam and Billi (2007) and Guerrieri and Iacoviello (2015).

¹⁹This is the argument made in Evans et al. (2015). But in Figure 7 the ZLB binds because of an initial level shock to the equilibrium real rate of interest and not, as in their analysis, because of an explosively high risk level that may keep the economy at the ZLB for an arbitrary length of time depending on the expectational horizon.

optimally delayed even if the economy is not actually exposed to any shocks during the recovery.

Following this temporary temporary risk shock, the economy eventually returns to a low-risk steady state with inflation on target. If the shock instead takes the form of a permanent increase in underlying risk to the level associated with the new normal, the economy instead gradually settles in the risky steady state as shown in dashed dotted black lines. In this normalisation scenario, optimal policy lifts off from the ZLB late and continues to lean against low inflation expectations. The optimal trade-off, however, requires the policymaker to accept that that inflation settles below below target as the economy recovers to its new normal.

9. Conclusion

In the canonical New Keynesian model, expectations are negatively skewed when risk is high relative to the available monetary policy space. Inflation settles materially below target in the absence of disturbances under optimal discretionary policy. Changes in the perception of risk give rise to cost-push effects regardless of the source of risk. The model is too simple to assign any great significance to quantitative results, and dynamic responses are likely to be too immediate in the purely forward-looking framework. But the results are indicative of the direction of the effects of risk in actual economies operating in an environment in which agents have reason to worry that monetary policy may be constrained in the foreseeable future. The new normal may be one in which monetary policy should lean against a re-anchoring of inflation expectations below target by operating the economy above potential in normal times. The analysis further points towards a monetary policy strategy aiming for an inflation rate *above*, but close to, 2%. The results also suggest that monetary policymakers should respond nimbly to changes in the perception of risk even as the economy escapes the ZLB.

References

- Aastveit, K. A., G. J. Natvik, and S. Sola (2017). Economic uncertainty and the effectiveness of monetary policy. *Journal of International Money and Finance* 76, 50–67.
- Adam, K. and R. M. Billi (2006). Optimal monetary policy under commitment with a zero bound on nominal interest rates. *Journal of Money, Credit and Banking* 38(7), 1877–1905.
- Adam, K. and R. M. Billi (2007). Discretionary monetary policy and the zero lower bound on nominal interest rates. *Journal of Monetary Economics* 54(3), 728–752.
- Adjemian, S. and M. Juillard (2013). Stochastic extended path approach. Unpublished manuscript, CEPREMAP, Ecole normale supérieure.
- Arellano, C., Y. Bai, and P. J. Kehoe (2018). Financial frictions and fluctuations in volatility. *Journal of Political Economy*, forthcoming.
- Ball, L., J. Gagnon, P. Honohan, and S. Krogstrup (2016). *What else can central banks do?* CEPR Press: Geneva Reports on the World Economy.
- Basu, S. and B. Bundick (2015). Endogenous volatility at the zero lower bound: Implications for stabilization policy. Research Working Paper RWP 15-1, Federal Reserve Bank of Kansas City.
- Basu, S. and B. Bundick (2017). Uncertainty shocks in a model of effective demand. *Econometrica* 85, 937–958.
- Bean, C. (2013). Global aspects of unconventional monetary policies. Remarks at the Federal Reserve Bank of Kansas City Economic Policy Symposium, Jackson Hole, Wyoming, 24 August.
- Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics* 98(1), 85–106.

- Bernanke, B. S. (2017). Monetary policy in a new era. Peterson Institute for International Economics.
- Black, F. (1995). Interest rates as options. *The Journal of Finance* 50(5), 1371–1376.
- Blanchard, O., E. Cerutti, and L. H. Summers (2015). Inflation and activity: Two explorations and their monetary policy implications. Working Paper Series WP15-19, Peterson Institute for International Economics.
- Blanchard, O., G. Dell’ariccia, and P. Mauro (2010). Rethinking macroeconomic policy. *Journal of Money, Credit and Banking* 42(s1), 199–215.
- Blanchard, O. and J. Galí (2007). Real wage rigidities and the New Keynesian model. *Journal of Money, Credit and Banking* 39(s1), 35–65.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica* 77(3), 623–685.
- Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives* 28(2), 153–76.
- Caggiano, G., E. Castelnuovo, and G. Nodari (2018). Risk management-driven policy rate gap. *Economics Letters* 171(C), 235–238.
- Caggiano, G., E. Castelnuovo, and G. Pellegrino (2017). Estimating the real effects of uncertainty shocks at the zero lower bound. *European Economic Review* 100(C), 257–272.
- Carney, M. (2017). Lambda. Speech, London School of Economics, London, 16 January.
- Christiano, L. J., R. Motto, and M. Rostagno (2014). Risk shocks. *American Economic Review* 104(1), 27–65.
- Chung, H., J.-P. Laforte, D. Reifschneider, and J. C. Williams (2012). Have we underestimated the probability of hitting the zero lower bound? *Journal of Money, Credit and Banking* 44(2), 47–48.

- Clarida, R., J. Galí, and M. Gertler (1999). The science of monetary policy: A New Keynesian perspective. *Journal of Economic Literature* 37(4), 1661–1707.
- Coenen, G., A. Orphanides, and V. Wieland (2004). Price stability and monetary policy effectiveness when nominal interest rates are bounded at zero. *Advances in Macroeconomics* 4(1).
- Coeurdacier, N., H. Rey, and P. Winant (2011). The risky steady state. *American Economic Review* 101(3), 398–401.
- Debortoli, D., J. Kim, J. Linde, and R. Nunes (2016). Designing a simple loss function for the Fed: Does the dual mandate make sense? Discussion Paper Series 1601, Institute of Economic Research, Korea University.
- Del Negro, M., D. Giannone, M. Giannoni, and A. Tambalotti (2017). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity Spring*.
- Eggertsson, G. B. and M. Woodford (2003). The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity* 34(1), 139–235.
- El-Erian, M. A. (2010). *Navigating the new normal in industrial countries*. Washington, D.C.: Per Jacobsson Foundation.
- English, W. B., J. D. López-Salido, and R. J. Tetlow (2015). The Federal Reserve’s framework for monetary policy: Recent changes and new questions. *IMF Economic Review* 63(1), 22–70.
- Evans, C., J. Fisher, F. Gourio, and S. Krane (2015). Risk management for monetary policy near the zero lower bound. *Brookings Papers on Economic Activity Spring*.
- Fair, R. C. and J. B. Taylor (1983). Solution and maximum likelihood estimation of dynamic nonlinear rational expectations models. *Econometrica* 51(4), 1169–1185.

- Fernández-Villaverde, J., P. Guerron-Quintana, K. Kuester, and J. Rubio-Ramirez (2015). Fiscal volatility shocks and economic activity. *American Economic Review* 105(11).
- Fernández-Villaverde, J., P. Guerrón-Quintana, J. Rubio-Ramirez, and M. Uribe (2011). Risk matters: The real effects of volatility shocks. *American Economic Review* 101(6), 2530–2561.
- Galí, J. (2008). *Monetary Policy, Inflation, and the Business Cycle. An Introduction to the New Keynesian Framework*. Princeton, New Jersey: Princeton University Press.
- Galindev, R. and D. Lkhagvasuren (2010). Discretization of highly persistent correlated AR(1) shocks. *Journal of Economic Dynamics and Control* 34(7), 1260 – 1276.
- Guerrieri, L. and M. Iacoviello (2015). OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics* 70, 22 – 38.
- Haldane, A. G., M. Roberts-Sklar, T. Wieladek, and C. Young (2016). Qe: the story so far. Staff Working Paper 624, Bank of England.
- Harrison, R. (2017). Optimal quantitative easing. Bank of England working papers 678, Bank of England.
- Hills, T. S., T. Nakata, and S. Schmidt (2016). The risky steady state and the interest rate lower bound. Finance and Economics Discussion Series 2016-009, Board of Governors of the Federal Reserve System.
- Ilut, C. L. and M. Schneider (2014). Ambiguous business cycles. *American Economic Review* 104(1), 2368–2399.
- Johannsen, B. K. (2014). When are the Effects of Fiscal Policy Uncertainty Large? Finance and Economics Discussion Series 2014-40, Board of Governors of the Federal Reserve System.

- Kiley, M. T. (2015). What can the data tell us about the equilibrium real interest rate? Finance and Economics Discussion Series 2015-77, Board of Governors of the Federal Reserve System (U.S.).
- Kiley, M. T. (2018). Quantitative easing and the “new normal” in monetary policy. Finance and Economics Discussion Series 2018-004, Board of Governors of the Federal Reserve System (U.S.).
- Kiley, M. T. and J. M. Roberts (2017). Monetary policy in a low interest rate world. *Brookings Papers on Economic Activity Spring*.
- Koop, G., M. H. Pesaran, and S. M. Potter (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics* 74, 119–147.
- Kozeniauskas, N., A. Orlik, and L. Veldkamp (2018). What are uncertainty shocks? *Journal of Monetary Economics* 100, 1–15.
- Krugman, P. (2014). Inflation targets reconsidered. Draft paper for ECB Sintra conference, May 2014.
- Laubach, T. and J. C. Williams (2016). Measuring the natural rate of interest redux. *Business Economics* 51(2), 57–67.
- LeRoy, S. F. and L. D. Singell (1987). Knight on risk and uncertainty. *Journal of Political Economy* 95(2), 394–406.
- Ljungqvist, L. and T. J. Sargent (2000). *Recursive macroeconomic theory*. Cambridge, Massachusetts: MIT Press.
- Maliar, L., S. Maliar, J. Taylor, and I. Tsener (2015). A tractable framework for analyzing a class of nonstationary markov models. NBER Working Papers 21155, National Bureau of Economic Research.

- Masolo, R. and F. Monti (2017). Ambiguity, monetary policy and trend inflation. Discussion Paper 2017-09, Centre for Macroeconomics.
- Mendes, R. R. (2011). Uncertainty and the zero lower bound: A theoretical analysis. MPRA Paper 59218, University Library of Munich, Germany.
- Nakata, T. (2017). Uncertainty at the zero lower bound. *American Economic Journal: Macroeconomics* 9(3), 186–221.
- Nakata, T. and S. Schmidt (2014). Conservatism and liquidity traps. Finance and Economics Discussion Series 2014-105, Board of Governors of the Federal Reserve System.
- Nakov, A. (2008). Optimal and simple monetary policy rules with zero floor on the nominal interest rate. *International Journal of Central Banking* 4(2), 73–127.
- Paoli, B. D. and P. Zabczyk (2013). Cyclical risk aversion, precautionary saving, and monetary policy. *Journal of Money, Credit and Banking* 45(1), 1–36.
- Plante, M., A. W. Richter, and N. A. Throckmorton (2018). The zero lower bound and endogenous uncertainty. *Economic Journal* 128(611), 1730–1757.
- Reifschneider, D. (2016). Gauging the ability of the FOMC to respond to future recessions. Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series 2016-068.
- Reifschneider, D. and J. C. Williams (2000). Three lessons for monetary policy in a low inflation era. *Journal of Monetary, Credit and Banking* 32(4), 936–966.
- Rogoff, K. (2015). Costs and benefits to phasing out paper currency. *NBER Macroeconomics Annual* 29(1), 445–456.
- Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. In T. Cooley (Ed.), *Frontiers of Business Cycle Research*, pp. 294–330. Princeton, NJ: Princeton University Press.

- Schmitt-Grohé, S. and M. Uribe (2004). Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control* 28(4), 755–775.
- Schmitt-Grohé, S. and M. Uribe (2010). The optimal rate of inflation. Volume 3 of *Handbook of Monetary Economics*, pp. 653 – 722. Elsevier.
- Svensson, L. E. (2010). Inflation targeting. In B. M. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics*, Volume 3, pp. 1237–1302. Amsterdam: Elsevier B.V.
- Walsh, C. (2014). Multiple objectives and central bank tradeoffs under flexible inflation targeting. CESifo Working Paper Series 5097, CESifo Group Munich.
- Williams, J. (2009). Heeding Daedalus: Optimal inflation and the zero lower bound. *Brookings Papers on Economic Activity*.
- Williams, J. C. (2014). Monetary policy at the zero lower bound: Putting theory into practice. Hutchins Center on Fiscal & Monetary Policy at Brookings.
- Woodford, M. (2003). *Interest and Prices. Foundations of a Theory of Monetary Policy*. Princeton, New Jersey: Princeton University Press.
- Woodford, M. (2005). Firm-specific capital and the New Keynesian Phillips curve. *International Journal of Central Banking* 1(2).

Tables

Table 1: Alternative calibrations

Episode	Data				Unconstrained model					
	$E(i)$	$\sigma(i)$	$E(\pi)$	$\sigma(\pi)$	$E(i)$	$\sigma(i)$	$E(\pi)$	$\sigma(\pi)$	100σ	P_0
New normal	–	–	–	–	3.02	2.20	2.00	2.48	0.27	0.09
Low risk	–	–	–	–	3.02	1.00	2.00	2.11	0.12	0.00
US 1968-1992	8.07	3.16	5.96	3.73	8.07	3.16	4.16	4.66	0.39	0.01
US 1993-2008	3.97	1.74	2.55	3.59	3.97	1.74	2.00	2.31	0.22	0.01
UK 1968-1992	10.56	2.86	8.77	6.83	10.56	2.86	6.56	6.83	0.35	0.00
UK 1993-2008	5.36	1.03	1.93	2.09	5.36	1.03	2.00	2.11	0.13	0.00

Note: $E(\cdot)$ and $\sigma(\cdot)$ denote means and standard deviations, respectively, of the nominal interest rate and inflation. Interest rates are measured by the FFR and BR for the US and UK, respectively. Inflation is annualised quarterly CPI inflation (source: Datastream). P_0 denotes the probability of negative interest rates in the unconstrained model as well as in normal distributions with means and standard deviations as in the data.

Table 2: New Keynesian model with ZLB under optimal discretion

	Scenario	Interest rate			Inflation			Output gap			P_{ZLB}
		i^*	i	$E(i)$	π^*	π	$E(\pi)$	x^*	x	$E(x)$	
1)	New normal	3.02	2.73	2.81	2.00	1.80	1.79	0.00	0.05	-0.01	0.14
2)	Normal r^* risk only	3.02	2.94	2.94	2.00	1.93	1.92	0.00	0.02	-0.00	0.07
3)	Normal u risk only	3.02	2.98	2.99	2.00	1.98	1.97	0.00	0.01	-0.00	0.04
4)	High r^* risk only	3.02	2.77	2.81	2.00	1.80	1.79	0.00	0.05	-0.01	0.13
5)	High u risk only	3.02	2.63	2.81	2.00	1.80	1.79	0.00	0.05	-0.01	0.17
6)	Lower r^*	2.76	2.25	2.40	2.00	1.66	1.64	0.00	0.09	-0.01	0.20
7)	Lower π^*	2.77	2.27	2.43	1.75	1.43	1.41	0.00	0.08	-0.01	0.20
8)	Higher r^*	3.27	3.10	3.14	2.00	1.88	1.87	0.00	0.03	-0.01	0.10
9)	Higher π^*	3.27	3.09	3.14	2.25	2.12	2.11	0.00	0.03	-0.00	0.10
10)	Very high π^*	5.04	5.03	5.03	4.00	3.99	3.99	0.00	0.00	-0.00	0.01

Note: i , π and x are risky steady state values, stars denote deterministic steady state values, $E(\cdot)$ and $\sigma(\cdot)$ denote means and standard deviations, respectively, and P_{ZLB} denotes the frequency of a binding ZLB. Deterministic steady state values for interest rates and inflation satisfy $1 + i^*\% = (1 + r^*\%)(1 + \pi^*\%)$.

Figures




Figure 1: Normal probability density functions with means and standard deviations set equal to sample means and standard deviations for the Federal Funds Rate (left panel) and Bank Rate (right panel) for the sample periods 1968-1992 (solid blue lines) and 1993-2008 (dashed red lines) as well as for a hypothetical new normal (dashed-dotted black lines). Dotted lines are estimated Kernel distributions with an optimised bandwidth for the normal kernel function for the 1968-1992 (in blue) and 1993-2008 (in red) subsamples.




Figure 2: Period welfare losses with an optimal annual inflation rate of 2% as a function of an operational inflation target denoted by $\tilde{\pi}^*$. Left panel shows losses evaluated in the risky steady state, and right panel unconditionally expected losses.

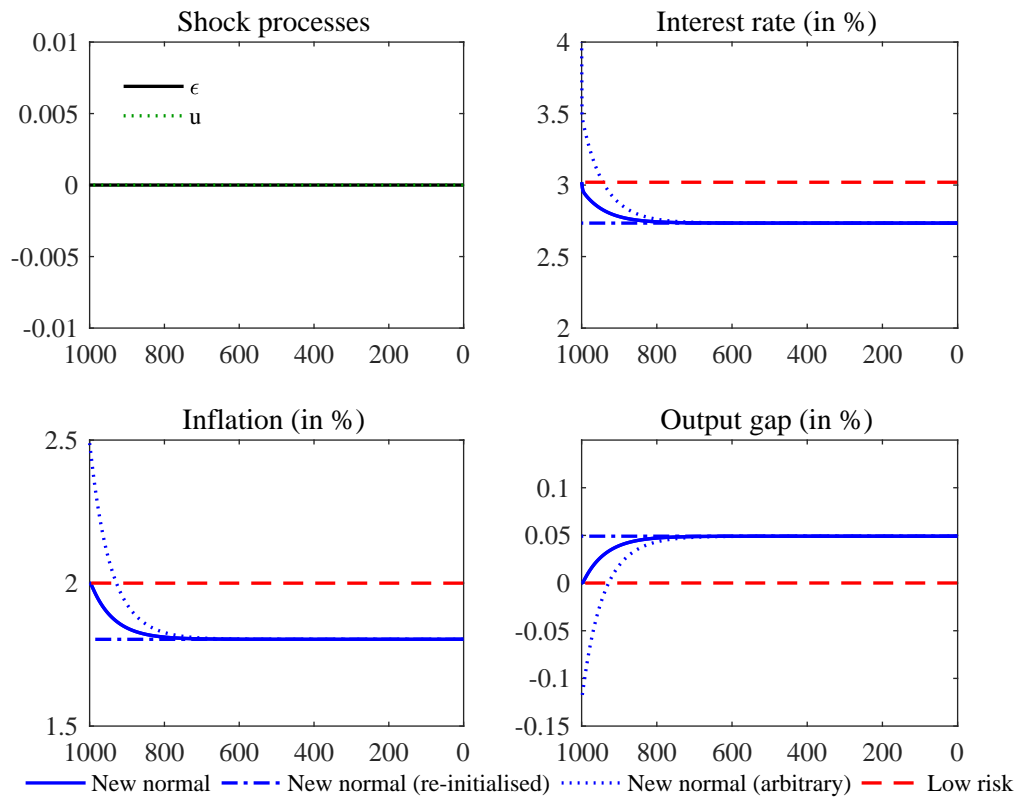


Figure 3: Zero-shock path of the solution to the canonical New Keynesian model with a ZLB on interest rates under optimal discretion with baseline underlying risk ($\sigma = 0.0027$) initiated at zero one-period ahead conditional gap expectations (blue solid lines), re-initiated at the true expectation (dashed-dotted blue lines), and initiated at an arbitrary inflation expectation of 2.5%, as well as the solution path for the model with low underlying risk ($\sigma = 0.0012$) initiated with zero one-period ahead conditional gap expectations (dashed red lines).




Figure 4: Impulse responses to a baseline risk shock (solid blue lines), a shock to r^* risk only (dashed red lines), and to cost-push risk only (dashed-dotted black lines), around a low-risk steady state ($\sigma = 0.0012$) in the canonical New Keynesian model with a ZLB on interest rates under optimal discretionary monetary policy.




Figure 5: Impulse responses to a positive (solid blue lines) and a negative (dashed red lines) baseline risk shock around the new normal risky steady state ($\sigma = 0.0027$) in the canonical New Keynesian model with a ZLB on interest rates under optimal discretionary monetary policy.

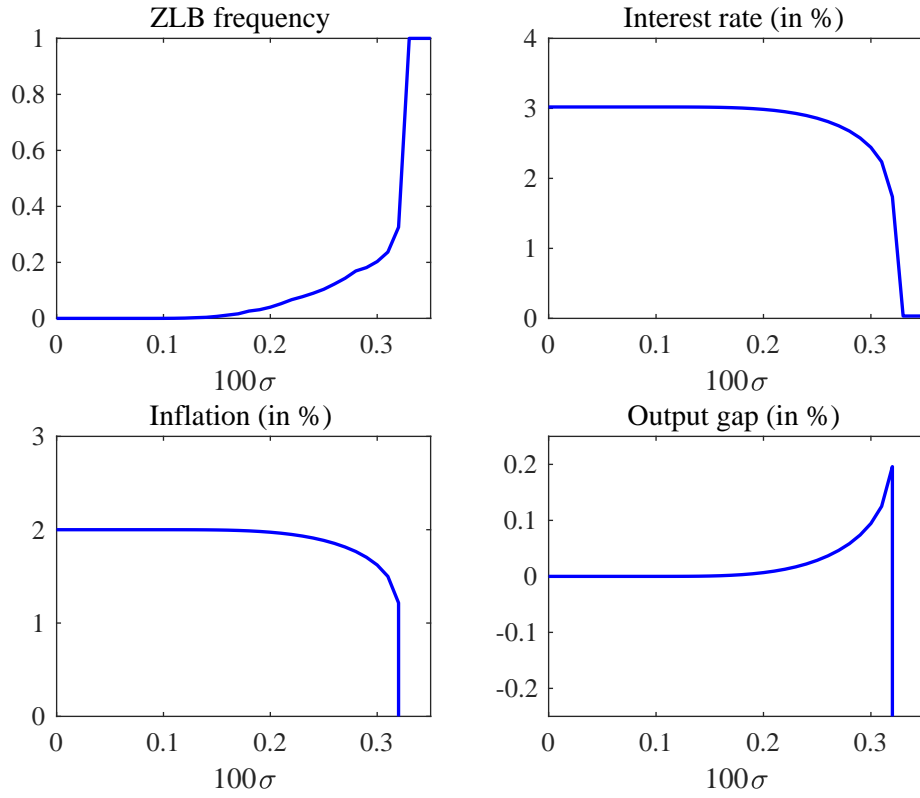


Figure 6: Economic outcomes as a function of risk in the canonical New Keynesian model with a ZLB on interest rate under optimal discretionary policy




Figure 7: Recovery from a ZLB episode to as reflected in the path of the equilibrium interest rate (dotted green lines in the top-right panel) under optimal discretionary policy in a low risk scenario (solid blue lines), in a low risk scenario with a baseline risk shock (dashed red lines), and with a shift in risk (dashed-dotted black lines).

Appendix A. Optimality conditions

The policymaker minimises (7) subject to (1), (2) and (3), taking expectations as given so that $E_t x_{t+1} = \bar{x}_{t,t+1}^e$ and $E_t \pi_{t+1} = \bar{\pi}_{t,t+1}^e$. To solve the optimal monetary policy problem, form the Lagrangian

$$\begin{aligned} \mathcal{L} = & \frac{1}{2} (\pi_t^2 + \kappa x_t^2) + \mu_{\pi,t} [\pi_t - \beta \bar{\pi}_{t,t+1}^e - \kappa x_t - u_t] \\ & + \mu_{x,t} \left[x_t - \bar{x}_{t,t+1}^e + \frac{1}{\varsigma} (i_t - E_t \pi_{t+1} - r_t^*) \right] - \mu_{i,t} (i_t + i^*) \end{aligned}$$

where μ_{π} , μ_x and μ_i are the multipliers on the three constraints. In addition to (1), (2) and (3), the Kuhn-Tucker conditions are

$$\begin{aligned} \pi_t + \mu_{\pi,t} &= 0 \\ \lambda x_t - \mu_{\pi,t} \kappa + \mu_{x,t} &= 0 \\ \mu_{x,t} \frac{1}{\varsigma} + \mu_{i,t} &= 0 \\ (i_t + i^*) \mu_{i,t} &= 0 \\ \mu_{i,t} &\geq 0 \end{aligned}$$

By substitution, these conditions can be reduced to

$$(i_t + i^*) (\lambda x_t + \kappa \pi_t) = 0 \tag{A.1}$$

$$\lambda x_t + \kappa \pi_t \leq 0 \tag{A.2}$$

There are two cases for condition (A.1) to consider.

Case 1. If $\lambda x_t + \kappa \pi_t = 0$, (A.2) also holds, and the solution is determined by (1), (2) and (8). The Kuhn-Tucker conditions are all satisfied in this case if and only if $i_t + i^* \geq 0$. If it happens that $i_t + i^* = 0$, the ZLB is just binding.

Case 2. If $i_t + i^* = 0$, condition (3) holds, and dynamics are determined by (1) and (2)

with $i_t = -i^*$. Now, we must have $\lambda x_t + \kappa \pi_t \leq 0$ for the Kuhn-Tucker conditions to be satisfied. If it happens that $\lambda x_t + \kappa \pi_t = 0$, the ZLB is just binding.

If there exists an $i_t = i_1 > -i^*$ such that (1), (2) and (8) holds for some values $x_t = x_1$ and $\pi_t = \pi_1$, then setting $i_t + i^* = 0$ would imply that $x_t > x_1$ and $\pi_t > \pi_1$ by (1) and (2) so that $\kappa \pi_t + \lambda x_t > 0$. Hence, if the Kuhn-Tucker conditions are satisfied in Case 1 given realisations of $\bar{x}_{t,t+1}^e$, $\bar{\pi}_{t,t+1}^e$, u_t and r_t^* , Case 2 will not be a candidate for a solution.

In sum, the policymaker chooses the unconstrained optimal policy allocation characterised by (1), (2) and (8) whenever it is feasible, and sets $i_t + i^* = 0$ when it is not. In the latter case, monetary policy will be insufficiently stimulatory in the sense that (A.2) holds with inequality rather than equality as in (8).

Appendix B. Solution

In each state (ϵ, u) in the $n_\epsilon \times n_u$ state space in period t , expectations are taken as given so that $E_t x_{t+1} = \bar{x}_{t,t+1}^e(\epsilon, u)$ and $E_t \pi_{t+1} = \bar{\pi}_{t,t+1}^e(\epsilon, u)$. Combining (1) and (8) in the form

$$\begin{aligned}\pi_t(\epsilon, u) &= \beta \bar{\pi}_{t,t+1}^e(\epsilon, u) + \kappa x_t(\epsilon, u) + u_t(\epsilon, u) \\ \pi_t(\epsilon, u) &= -\frac{\lambda}{\kappa} x_t(\epsilon, u)\end{aligned}$$

gives the unconstrained optimal allocation

$$\begin{aligned}\pi_t^{opt}(\epsilon, u) &= \frac{\lambda}{\lambda + \kappa^2} [\beta \bar{\pi}_{t,t+1}^e(\epsilon, u) + u_t(\epsilon, u)] \\ x_t^{opt}(\epsilon, u) &= -\frac{\kappa}{\lambda + \kappa^2} [\beta \bar{\pi}_{t,t+1}^e(\epsilon, u) + u_t(\epsilon, u)]\end{aligned}$$

The interest rate consistent with this allocation follows from (2):

$$i_t^{opt}(\epsilon, u) = \bar{\pi}_{t,t+1}^e(\epsilon, u) + r_t^*(\epsilon, u) - \sigma [x_t^{opt}(\epsilon, u) - \bar{x}_{t,t+1}^e(\epsilon, u)]$$

If $i_t^{opt}(\epsilon, u) \geq -i^*$, $\{x_t^{opt}(\epsilon, u), \pi_t^{opt}(\epsilon, u), i_t^{opt}(\epsilon, u)\}$ is the conditional solution for state (ϵ, u) in period t . If the ZLB is binding so that $i_t^{opt}(\epsilon, u) < -i^*$, the interest rate is set to $i_t^{zlb}(\epsilon, u) = -i^*$. Now from (2) and (1):

$$\begin{aligned} x_t^{zlb}(\epsilon, u) &= \bar{x}_{t,t+1}^e(\epsilon, u) - \frac{1}{\sigma} [-i^* - \bar{\pi}_{t,t+1}^e(\epsilon, u) - r_t^*(\epsilon, u)] \\ \pi_t^{zlb}(\epsilon, u) &= \beta \bar{\pi}_{t,t+1}^e(\epsilon, u) + \kappa x_t^{zlb}(\epsilon, u) + u_t(\epsilon, u) \end{aligned}$$

Hence, the conditional solution for $y_t(\epsilon, u) \in \{x_t(\epsilon, u), \pi_t(\epsilon, u), i_t(\epsilon, u)\}$ for all nodes (ϵ, u) in the state grid is

$$y_t^{sol}(\epsilon, u) = \begin{cases} y_t^{opt}(\epsilon, u) & \text{if } i_t^{opt}(\epsilon, u) \geq -i^* \\ y_t^{zlb}(\epsilon, u) & \text{if } i_t^{opt}(\epsilon, u) < -i^* \end{cases}$$

Ex ante conditional expectations across the state grid can now be found as

$$\begin{aligned} \bar{x}_{t-1,t}^e &= \mathbf{P}_\epsilon \mathbf{x}_t^{sol} \mathbf{P}'_u \\ \bar{\pi}_{t-1,t}^e &= \mathbf{P}_\epsilon \pi_t^{sol} \mathbf{P}'_u \end{aligned}$$

where \mathbf{P}_ϵ and \mathbf{P}_u are Markov transition matrices of dimensions $n_\epsilon \times n_\epsilon$ respectively $n_u \times n_u$.

The solution algorithm iterates the conditional solution backwards from some period $t = T \gg 0$ to $t = 0$, initiated with $\bar{x}_{T,T+1}^e = \bar{\pi}_{T,T+1}^e = \mathbf{0}$. The state-contingent model solution is then $\{x_0^{sol}(\epsilon, u), \pi_0^{sol}(\epsilon, u), i_0^{sol}(\epsilon, u)\}$.

Appendix C. Calculation of model statistics

The stochastic steady-state of model variable $y_t \in \{i_t, \pi_t, x_t\}$ is simply

$$y \equiv y_0^{sol}(\epsilon = 0, u = 0)$$

To calculate unconditional expectations, note that for each shock process $z \in \{\epsilon, u\}$ with Markov transition matrix \mathbf{P}_z , the stationary unconditional distribution, \mathbf{p}_z , satisfies $\mathbf{p}'_z =$

$\mathbf{p}'_z \mathbf{P}_z$, or equivalently $(\mathbf{I} - \mathbf{P}'_z) \mathbf{p}_z = 0$. Hence, the unconditional distribution can be found as the normalised eigenvector associated with the unitary eigenvalue of \mathbf{P}_z . The unconditional probability distribution over the state space (ϵ, u) is then the $n_\epsilon \times n_u$ vector product $\mathbf{P} = \mathbf{p}_\epsilon \mathbf{p}'_u$. The unconditional expectation of model variable y_t is found as

$$E(y) \equiv \sum_{\epsilon} \sum_u P(\epsilon, u) y_0^{sol}(\epsilon, u)$$

The unconditional variance is

$$\sigma^2(y) \equiv \sum_{\epsilon} \sum_u P(\epsilon, u) [y_0(\epsilon, u) - E(y)]^2$$