Risk shocks and monetary policy in the new normal*

Martin Seneca

Bank of England

First draft: 7 September 2015
Current draft: 4 September 2017

Abstract

Risk shocks give rise to cost-push effects in the canonical New Keynesian model if they are large relative to the distance between the nominal interest rate and its zero lower bound (ZLB). Therefore, stochastic volatility introduces occasional trade-offs for monetary policy between inflation and output gap stabilisation. The trade-off inducing effects operate through expectational responses to the interaction between perceived shock volatility and the ZLB. At the same time, a given monetary policy stance becomes less effective when risk is high. Optimal monetary policy calls for potentially sharp reductions in the interest rate when risk is elevated, even if this risk never materialises. If the underlying level of risk is high, inflation will settle potentially materially below target in a risky steady state under optimal monetary policy.

JEL codes: E52, E58

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

*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, 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, as well as Menzie Chinn, Ambrogio Cesa-Bianchi, Richard Dennis, Jesús Fernández-Villaverde, Andrea Ferrero, Jordi Galí, Richard Harrison, Valerie A. Ramey, Mathias Trabandt, Ricardo Reis 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). After nearly a decade of ultralow interest rates across major advanced economies, this pre-crisis assessment has been revised in light of the incoming data (e.g. Chung et al., 2012, 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 rates—and disturbances are expected to be larger—increasing the occasional need for such cuts. In the foreseeable future, therefore, policy rates are deemed likely to hit their lower bound more frequently and for longer periods of time than previously assumed. Nevertheless, an optimistic view holds that unconventional monetary policy—such as quantitative easing and forward guidance—can be relied upon as substitutes for conventional reductions in policy rates (e.g. Reifschneider, 2016). Hence, whilst the ZLB may bind from time to time, monetary policy will (with its extended toolkit) rarely be constrained in its ability to stabilise inflation and the real economy 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 always be able to respond to future adverse disturbances with sufficient monetary stimulus. Perhaps people have been psychologically scarred by the Great Recession in ways that give them an exaggerated sense of the risk of bad shocks (Haldane, 2015). Or perhaps they simply have less confidence in the feasibility and effectiveness of unconventional monetary policy than the policy optimists. 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?

In this paper, I point to two key differences that set monetary policy in this hypothetical '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 (slightly) above its efficient potential in normal times. This stimulatory bias leans against a downside deanchoring of inflation expectations, but it allows inflation to settle potentially materially below target in the absence of disturbances. Second, because of the constraints on monetary policy alone, changes in the public's perception of risk affect the appropriate stance of policy. A spike in uncertainty, for example, calls for potentially sizeable reductions in interest rates even if risk does not materialise in any actual disturbances to the economy. This is in sharp contrast to the conventional guidelines that monetary policy should only respond to disturbances that have actually occurred (or are fully anticipated), or at least that risk should only affect the stance of policy to the extent that it affects precautionary behaviour by households and firms.[1]

I derive these results in a simple version of the canonical New Keynesian model. The monetary design problem in 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 practiced by major central banks (Svensson, 2010). I focus on optimal policy under discretion. This policy assumption neither requires policymakers to follow an instrument rule mechanically, nor does it require them to commit to a time-inconsistent plan. Arguably, the period-by-period nature of policymaking under discretion makes it a more realistic description of actual monetary policy (see e.g. Bean, 2013).

Specifically, I study optimal discretionary monetary policy responses to risk shocks in a quasi-linear version of the New Keynesian model augmented with a ZLB. By risk shocks I mean changes in the second moments of conventional level shocks that affect behavioural relations in model. Throughout, I show the economy's impulse responses to such changes along the zero-shock path, i.e. the trajectory of the economy through the state space along which innovations to the level shocks do not actually occur. As risk never materialises along this path—as it were, nothing actually happens in this paper—the effects can be thought of as responses to changes in the *perceptions* of risk, where risk is defined as the set of variances of the independent distributions from which innovations to level shock processes are drawn.[2] I find that these risk shocks give rise to negative tradeoff-inducing cost-push effects if they are large relative to the distance between the nominal interest rate and its zero lower bound (ZLB). At the same time, a given monetary policy stance becomes less effective when risk is high. Optimal monetary policy under discretion therefore calls for potentially sharp movements in the interest rate when the level of perceived risk changes. If the underlying level of risk is high, inflation settles below target in a risky steady state (as defined by Coeurdacier et al., 2011) under optimal policy.

[1]For a general class of models without inequality constraints, Schmitt-Grohe and Uribe (2004) show that risk does not affect decision rules to a first-order approximation, and only the constant term up to a second order.

[2]Even so, to remain loyal to the Knightian distinction between risk and uncertainty, I refer to the changes in second moments as *risk* rather than *uncertainty* shocks; within the context of the model, agents actually assign a number to the risk that they perceive, though it may well be that elevated risk in the model stands in for Knightian uncertainty in reality.

The analysis follows previous studies of the implications for optimal discretionary monetary policy of the *presence* of risk in models with a ZLB. The main mechanism driving results was first identified by Adam and Billi (2007) and Nakov (2008). By deviating from the usual assumption of perfect foresight, these authors showed how the ZLB may give rise to a negative skew in expectations in the New Keynesian model's stochastic equilibrium. Recently, this mechanism has featured in a number of applications. For example, Nakata and Schmidt (2014) study the implications for the optimal degree of conservatism in monetary policy, while Evans et al. (2015) investigate the optimal timing of lift-off from a ZLB episode. I adopt the recursive method to solving the New Keynesian model implemented by Evans et al. (2015). But further to their analysis, since I am interested in the implications for monetary policy of changes in risk in 'normal times' rather than in a specific ZLB episode, I characterise the model's stochastic steady state and I trace out dynamics around it. I show how a risk shock propagates through dynamics in the skew in expectations in the stochastic equilibrium.

The emphasis on dynamics in normal times also sets the analysis apart from a recent related study by Basu and Bundick (2015), who also derive impulse responses to risk shocks, but who focus on dynamics at the ZLB. There are other important differences. First, they specify an instrument rule for monetary policy, confirming the results in Adam and Billi (2007) that commitment is a powerful strategy to alleviate the negative effects of risk when current policy is constrained. And second, they consider a fully non-linear version of the New Keynesian model in which the effects of the ZLB interacts with other non-linearities. The simple quasi-linear model that I consider comes with the benefit—beyond being easier to solve—that effects stemming from the constraints on policy are clearly separated from higher-order behavioural effects.[3] My results on the stochastic steady state, however, are fully consistent with the contemporaneous analysis by Hills et al. (2016), who study the stochastic steady state in a non-linear New Keynesian model with a ZLB. In an empirical application to the United States, the authors estimate that the level of inflation falls short of target by about 25 basis point in the risky steady state.

More broadly, the 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. While the empirical literature has struggled to identify structural risk shocks from the volatility measures that are usually taken to be proxies for risk and uncertainty, the theo-

[3]For how higher-order effects can be amplified by the ZLB, see e.g. Nakata (2013) and Basu and Bundick (2014). Johannsen (2014) and Fernandez-Villaverde et al. (2015) suggest that uncertainty about fiscal policy in particular has larger implications for the economy when monetary policy is constrained.

retical literature provides clear channels through which risk shocks may affect the economy as discussed in the recent survey by Bloom (2014). The mechanism that I emphasise is a further 'bad news channel' (in the terminology of Bernanke, 1983) arising from the inability of monetary policy to respond to large adverse shocks with sufficient stimulus, but never with contractionary action 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 line with the empirical evidence provided by Plante et al. (2014) and Castelnuovo et al. (2015). Similarly, the finding that monetary policy is less effective when risk is high is consistent with the evidence in Aastveit et al. (2013).

The paper is organised as follows. Section 2 describes the model and its solution. Section 3 presents the quantitative analysis. It describes the parameterisation and the numerical solution of the stochastic steady state, and it presents impulse responses to persistent shocks to risk both away from the ZLB and when the ZLB is binding. Section 5 concludes.

2 The model

The model is the canonical forward-looking New Keynesian model extended with a ZLB on interest rates. In addition to a specification of monetary policy, it consists of the equations:

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

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

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

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 canonical model is derived from its microfoundations by Woodford (2003) and Galí (2008) among others.

There are two shock processes in the model. u_t is a cost-push process, and r_t^* is the efficient equilibrium real interest rate in deviation from its steady state level $r^* = i^* - \pi^*$. 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:

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

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

where $\nu_{\epsilon,t} \sim N(0, \sigma^2_{\epsilon,t})$ and $\nu_{u,t} \sim N(0, \sigma^2_{u,t})$. Importantly, 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. Specifically, I let a baseline risk shock be such that $\varsigma^{-1}\sigma_{\epsilon,t} = \sigma_{u,t} = \sigma_t$ with

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

where $\nu_{\sigma,t}$ is the innovation to risk, and σ is the underlying level of risk in the absence of risk disturbances. Stochastic volatility generally refers to a scenario where these innovations are drawn from a fixed distribution each period.[4]

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

$$L = \pi_t^2 + \lambda x_t^2 \tag{7}$$

each period subject to the Phillips curve in (1) while taking expectations as given. This gives rise to the targeting rule

$$\pi_t = -\frac{\lambda}{\kappa} x_t \tag{8}$$

stating the optimal policy trade-off between inflation and the output gap. 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 interest rate will be set to the maximum of this optimal level and zero. For comparison, under an alternative regime the policymaker follows the simple instrument rule

$$i_t = \max\{-i^*, \phi_\pi \pi_t + \phi_x x_t\} \tag{9}$$

I solve this quasi-linear version of the canonical model following the approach in the recent analysis of this model by Evans et al. (2015). I approximate the shock processes by independent Markov processes using the Rouwenhorst (1995) method. I then solve the model

[4] I only consider realisations of risk that are strictly larger than zero. In applications emphasising stochastic volatility, the risk shock could be specified in logs to rule out non-zero realisations of risk.

backwards from a distant future period T, beyond which there is no risk and all shocks are zero so that $E_t\pi_{t+1} = E_t x_{t+1} = 0$ 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. 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. See the Annex for details. The solution consists of the values for inflation, the output gap and the interest rate, to which this algorithm converges in the initial period $t = 0$. The zero-shock solution can be found as the node in the grid with $(\epsilon, u) = (0, 0)$.

I find impulse 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 $\sigma_\epsilon = \sigma_u$ from an initial spike according to (6). The impulse response can now be found as sequence of zero-shock solutions found in the outer loop. The advantage of taking this approach is that risk does not become a state variable. It is also a natural staring point for considering variation in risk perceptions: each period, economic agents assign a number to the level of risk 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 the autoregressive profile for risk are likely to generate similar qualitative results in this quasi-linear model.

3 Calibration

To guide the parameterisation of the model, Figure 1 provides a simple illustration of how the hypothetical new normal differs from recent pre-crisis historical experience in the United State (left panel) and the United Kingdom (right panel). The normal probability density functions in solid blue share the means and standard deviations with the observed Federal Funds Rate (FFR) and Bank Rate (BR) from 1968 through 1992. An observer looking back at these distributions in the early 1990s would not have worried about the ZLB. Policy interest rates had been very volatile (standard deviations were 3.2 and 2.9, respectively), but means had also been high enough (8.1% and 10.6%, respectively) that the probability that interest rates should be negative seemed negligible for both countries. Over the subsequent 15 years, the distributions of both the FFR and BR shifted sharply to the left as shown

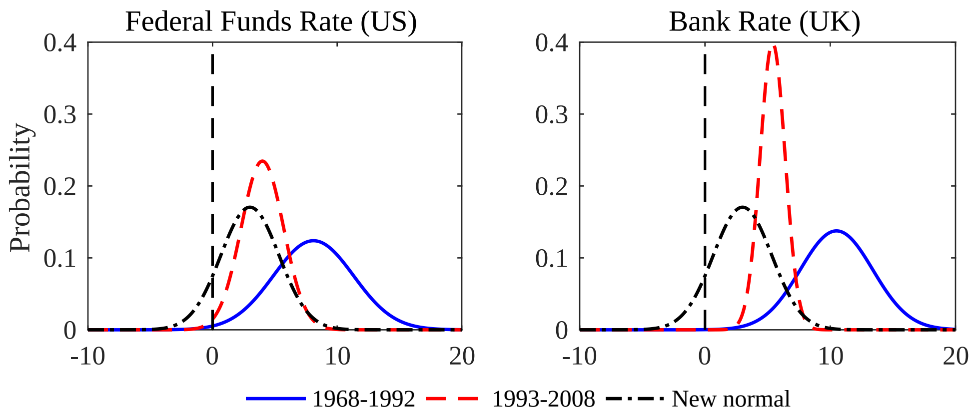

Figure 1: Normal probability density functions with means and standard deviations set equal to sample means and standard deviations for the Federal Funds Rate (left panel) and Bank Rate (right panel) for the sample periods 1968-1992 (solid blue lines) and 1993-2016 (dashed red lines) as well as for a hypothetical new normal (dashed-dotted black lines).

in dashed red lines. In isolation, lower means (4.0% and 5.4%, respectively) would increase probability mass below zero. But as volatility fell substantially at the same time—perhaps because policymakers faced smaller disturbances—the risk of a binding ZLB did not appear to have increased: the probability that interest rates should be negative seemed to remain comfortably close to zero.[5] From this perspective, it is not surprising that a consensus emerged in the pre-crisis period that the ZLB should be of no great practical concern.

Looking forward after close to a decade with policy interest rates close to zero, I define the hypothetical new normal by a further shift to the left of the distribution for desired policy rates, combined with an increase in the spread compared to the 1993-2008 period. The shift captures an assumed decline in the trend component of the efficient equilibrium real rate of interest, and the fatter tails reflect an end to the Great Moderation so that larger disturbances call for stronger policy responses then in the pre-crisis inflation targeting period. This combination of a reduction in mean and an increase in spread increases probability mass below zero. In the new normal shown in dashed-dotted black lines in Figure 1, I let the probability that desired interest rates should fall below zero be 9%. Whilst somewhat arbitrary, this probability results from setting the average desired policy rate to 3%—in line with the forecasts of the longer-run level of the FFR emphasised by Reifschneider (2016), and the estimates of long-run real rates in Laubach and Williams (2016) and Del Negro et al. (2017)—and fixing the standard deviation at 2.2, in between the values for the two pre-crisis subsamples for the US and the UK.

[5]For the UK, a naive extrapolation from the 1993-2008 distribution for BR could lead to the conclusion that the ZLB should be expected to bind for as little as 20 out of the remaining 2.5 billion years of potentially habitable conditions on Earth (the central estimate in Rushby et al., 2013).

Table 1: Alternative calibrations

	Data				Unconstrained model					
Episode	$E(i)$	$\sigma(i)$	$E(\pi)$	$\sigma(\pi)$	$E(i)$	$\sigma(i)$	$E(\pi)$	$\sigma(\pi)$	100σ	$P(i<0)$
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	10.59	2.86	8.77	6.83	10.59	2.86	6.59	6.86	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: $E(.)$ and $\sigma(.)$ 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). $P(i<0)$ 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.

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 inflation target is assumed to be $\pi^* = 2\%$ and the deterministic steady-state level of the efficient equilibrium real rate of interest is assumed to be constant at $r^* = 1\%$. The normal level for the nominal interest rate is then approximately $i^* = 3\%$ with a discount factor $\beta = 0.9975$. The slope of the Phillips curve is assumed to be $\kappa = 0.02$ and the inverse of the elasticity of intertemporal substitution $\varsigma = 1$ in line with values often used in the literature. Similarly, the weight on the output gap in the monetary policy loss function is set to be fairly low at $\lambda = 0.02$ in keeping with derivations of the loss function from household preferences (see e.g. Galí, 2008). The shock processes 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 so that, when the ZLB is not imposed, the policy rate is negative with probability 9%.

For comparison, I also consider a hypothetical scenario in which the dispersion of desired interest rates is kept low (at a level similar to that observed in the UK between 1993 and 2008) while the mean shifts down. Underlying risk is low enough in this case that the probability that interest rate should be negative remains negligible despite a low level of r^*. The standard deviation of interest rates is set to 1, which, with a mean of 3%, is obtained by letting $\sigma = 0.0012$ in the unconstrained model.

Table 1 compares the new normal and the alternative low-risk scenario with parameterisations of i^* and σ that allow the unconstrained model to match the historical distributions in Figure 1. Here, I assume that $r^* = 3.75$ in the 1968-1992 period, and $\pi^* = 2\%$ between 1993 and 2016. Whilst the model is too simple to capture the covariance structure of the data more broadly, the table shows that it does reasonably well in capturing inflation volatility in the four historical periods. Hence, it does not seem unreasonable a priori that it can shine some light on inflation outcomes in the hypothetical scenarios.

Of course, 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. Absent substantial reform to the payment system, it is improbable that policymakers can drive interest rates deep into negative territory for anything other than very short periods of time (see e.g. Rogoff, 2015). It is more likely, perhaps, that unconventional policies such as forward guidance and quantitative easing may act as substitutes for negative short-term interest rates. The interest rate in the model may best be thought of as a *shadow* interest rate implicitly incorporating the effects of unconventional policy tools as defined by Black (1995), and the shadow rate may be negative if these tools are operational whenever short rates are constrained. But if either the availability or the effectiveness of such tools is unclear, the shadow rate will also be bounded from below.

In what follows, I shall assume that the lower bound is exactly zero as specified in (3). 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 in the new normal. The analysis becomes irrelevant if either of these assumptions is fully reversed so that interest rates can go negative without difficulty, or unconventional policy tools can be relied upon as perfect substitutes for changes in policy rates in a way that is perfectly understood by the public. But the conclusions hold for any constellation of the desired 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.

4 Stochastic steady state

Turning to the full model with a ZLB, Figure 2 shows zero-shock paths of the model solution under optimal discretionary policy by backward induction from period $t = T$ to period $t = 0$ for the baseline new normal as well as the low-risk alternative. As a cross-check, the figure also shows the solution for the new normal when monetary policy follows a Taylor rule with $\phi_\pi = 1.5$ and $\phi_x = 0.5/4$. The zero-shock paths are conditional on the particular realisation

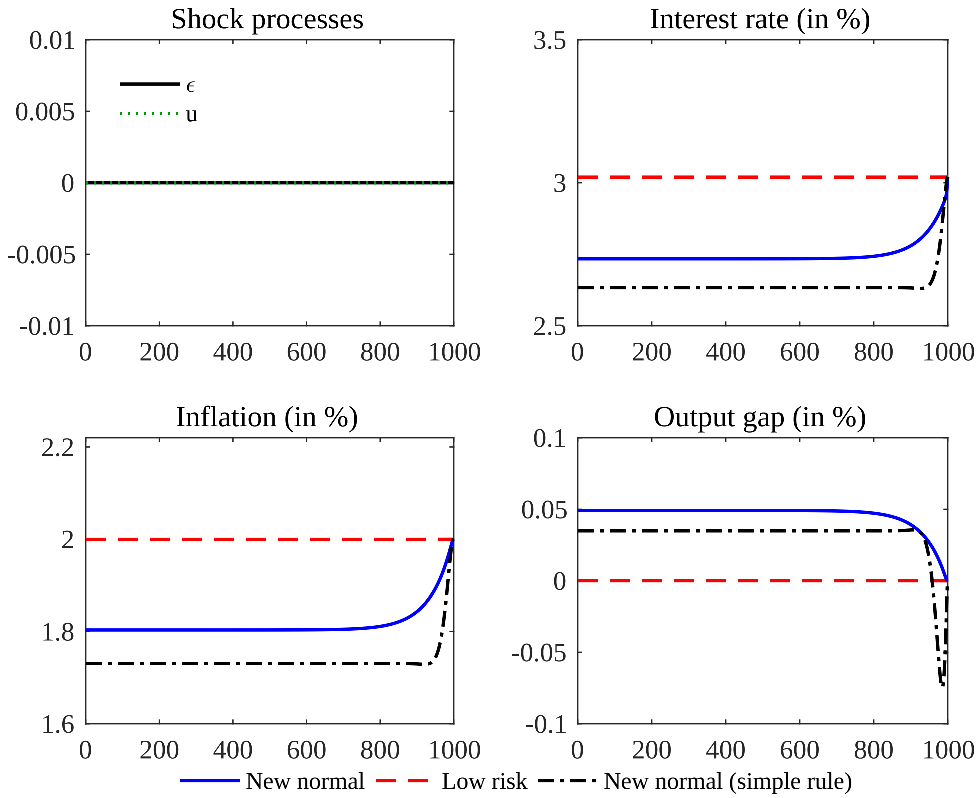

Figure 2: Zero-shock path of the solution to the canonical New Keynesian model with a ZLB on interest rates by backward induction in the new normal under optimal discretionary policy ($\sigma = 0.0029$; solid blue lines), with low risk ($\sigma = 0.0012$; dashed red lines), and in the new normal with a simple monetary policy rule ($\sigma = 0.0029$; dashed-dotted black lines).

of the two shock processes shown in the top-left panel in which no non-zero shocks actually occur. In this sense, the converged zero-shock paths at time $t = 0$ represent stochastic or *risky* steady states of the model as defined by Coeurdacier et al. (2011). The steady states represent resting points for the economy when all disturbances have dissipated. Non-zero realisations of shocks will of course continuously drive the economy away from any of these steady states.[6] The stochastic steady state potentially differs from the deterministic one because it accounts for agents' expectations that such deviations from steady state will occur.

In the new normal ($\sigma = 0.0027$), risk is perceived to be high enough that the monetary policymaker cannot respond sufficiently to some negative disturbances. Expectations are negatively skewed even in the absence of any actual shocks to the economy through the mechanism emphasised by Adam and Billi (2007) and Nakov (2008). Low inflation expectations will weigh on the price-setting of firms and actual inflation will be below target. This deflationary effect from the ZLB results in a loosing bias in optimal discretionary monetary

[6]For each iteration t, the full solution is given by an $n_\epsilon \times n_\epsilon$ matrix as described in the Annex. The zero-shock path is the set of values for the node ($\epsilon = 0, u = 0$) of this matrix for $t = 0, 1, 2, ..., T$.

Table 2: New Keynesian model with ZLB under optimal discretion

		Interest rate			Inflation			Output gap			
	Episode	i^*	i	$E(i)$	π^*	π	$E(\pi)$	x^*	x	$E(x)$	$P(ZLB)$
1)	New normal	3.00	2.73	2.81	2.00	1.80	1.79	0.00	0.05	−0.01	0.14
2)	r^* shocks only	3.00	2.94	2.94	2.00	1.93	1.92	0.00	0.02	0.00	0.07
3)	u shocks only	3.00	2.98	2.99	2.00	1.98	1.97	0.00	0.01	0.00	0.04
4)	Large r^* shocks	3.00	2.78	2.82	2.00	1.81	1.80	0.00	0.05	−0.01	0.13
5)	Large u shocks	3.00	2.64	2.82	2.00	1.81	1.80	0.00	0.05	−0.01	0.15
6)	Lower r^*	2.75	2.24	2.41	2.00	1.66	1.64	0.00	0.09	−0.01	0.20
7)	Lower π^*	2.75	2.27	2.43	1.75	1.43	1.41	0.00	0.08	−0.01	0.20
8)	Higher r^*	3.25	3.10	3.14	2.00	1.88	1.87	0.00	0.03	−0.01	0.10
9)	Higher π^*	3.25	3.09	3.14	2.25	2.12	2.11	0.00	0.03	0.00	0.10
10)	High π^*	5.00	5.03	5.03	4.00	3.99	3.99	0.00	0.00	0.00	0.01

Note: $E(.)$ and $\sigma(.)$ denote means and standard deviations, respectively. $P(ZLB)$ denotes the frequency of a binding ZLB.

policy. The new normal steady state is therefore one in which policymakers optimally lean against the deanchoring of inflation expectations by operating the economy (slightly) above potential with low interest rates. Under optimal discretionary policy, shown in solid blue lines, inflation settles about 20 basis points below target. Despite a somewhat stronger loosening bias, inflation is lower still under the simple rule with standard parameters, shown in dashed-dotted black lines. Notice that, in both cases, the algorithm converges to the steady state at expectational horizons much lower than the assumed $T = 1000$. The dashed red lines show the solution path in the alternative low-risk scenario ($\sigma = 0.0012$). In this case, despite a low i^*, the ZLB is never a concern, and expectations never deviate from zero in the absence of disturbances. As a consequence, inflation is on target, the output gap is zero, and the interest rate is at its normal level independently of the monetary policy regime. The stochastic steady state coincides with the deterministic one in this case.

The first row in Table 2 shows these steady-state outcomes for interest rates, inflation and the output gap in the new normal under optimal policy along with their values in the deterministic steady state and their unconditional means. As just described, the stochastic steady state level of inflation falls short of the target attained in the deterministic steady state, whilst the economy operates above potential as a consequence of an easing bias in monetary policy. On average, however, output falls short of potential because of spells at the ZLB, which drive the average interest rate above its stochastic steady-state level. As monetary policy is less stimulatory on average than in the absence of disturbances, average inflation settles somewhat further below target. Because of of monetary policy's inability to provide sufficient stimulus, the frequency of ZLB episodes is higher than the probability that desired rates should be negative in the unconstrained model. In the new normal, the ZLB binds with probability 14% as shown in the last column.

The remaining rows in 2 illustrate the sensitivity of these statistics to assumptions about risk and the available monetary policy space. The second row shows the effect of keeping the risk of cost-push shocks low so that risk is only high for shocks to the efficient equilibrium real interest rate. The third row shows the opposite case with low risk of r^* shocks and high risk of shocks to the Phillips curve. In both cases the stochastic steady state deviates from the deterministic one with inflation settling below target. Also, the marginal constributions of the two shocks to the baseline are similar if somewhat larger for r^* shocks. But the deviations are much smaller with stochastic steady-state inflation rates of 1.93 and 1.98, respectively. These results suggest that agents are particularly concerned about the inability of policymakers to respond when large adverse disturbances to the cost-push process and the equilibrium rate coincide. Of course, higher risk for individual shocks results in larger biases. Increasing σ_ϵ to about 0.0056 when the risk of cost-push shocks is low (row 4) or σ_u to about 0.0036 when the risk of r^* shocks is low (row 5) leads to similar biases in inflation and the output gap as under the baseline parameterisation.

Notice that, in all of these cases, inflation falls below target in the stochastic steady state only because agents worry that the ZLB may bind in future. With a policy rate above three per cent, monetary policy has a substantial distance between it and the ZLB. But when risk is perceived to be high, agents worry about policy's inability to respond to large adverse shocks in the future even when the policy rate is currently well above the ZLB.

For a given level of risk, the effect on expectations depends on the available monetary policy space. As illustrated by the remaining cases in Table 2, the closer the economy operates to the ZLB, the larger are the effects of risk on outcomes in the stochastic steady state. If the distance to the ZLB is reduced by about 25 basis points in the deterministic steady state, either because the equilibrium real rate of interest is lower (row 6) or because monetary policy targets a lower inflation rate (row 7), the stochastic steady-state inflation rate falls below target by a further 15 basis points or so. In the case in row 7, where inflation in the deterministic steady state itself is lower, this implies that inflation is almost 40 basis points lower than in the baseline. By contrast, if i^* increases to about 4%, the negative bias in inflation is reduced by 6-7 basis points. With a higher inflation target, the component of i^* that can actually be chosen by policymakers, inflation settles around 2.1%. To fully

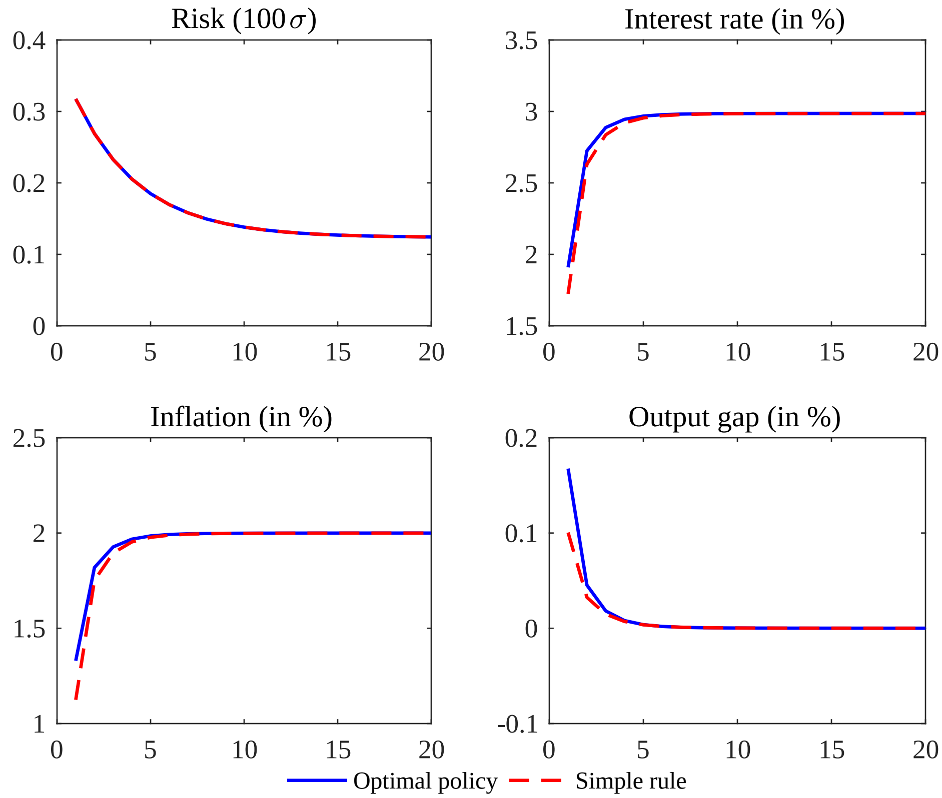

Figure 3: Impulse responses along the zero-shock path to a risk shock around a low-risk steady state ($\sigma = 0.012\%$) in the canonical New Keynesian model with a ZLB on interest rates under optimal discretionary monetary policy (solid blue lines) and with a simple interest rate rule (dashed red lines).

eliminate the negative bias in inflation, however, policymakers will have to set a target for inflation above 5% (Case 10) when risk is high.

4.1 Impulse responses to a baseline risk shock

Now suppose that risk varies over time. Specifically, to build intuition, consider a baseline risk shock to an economy operating in a low-risk steady state so that the standard deviations of the innovations to the two shock processes jump from 0.12% to 0.32% with a gradual return to 0.12% according to the process in (6). The risk shock represents a scenario in which risk is temporarily elevated so that agents expect innovations to level shocks to be drawn from a distribution with fatter tails for some time in the future.

Figure 3 shows the impulse responses along the zero-shock path. That is, the economy is not actually hit by any shocks along this adjustment path; it is only the perception of risk that changes. When risk is high, agents 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),

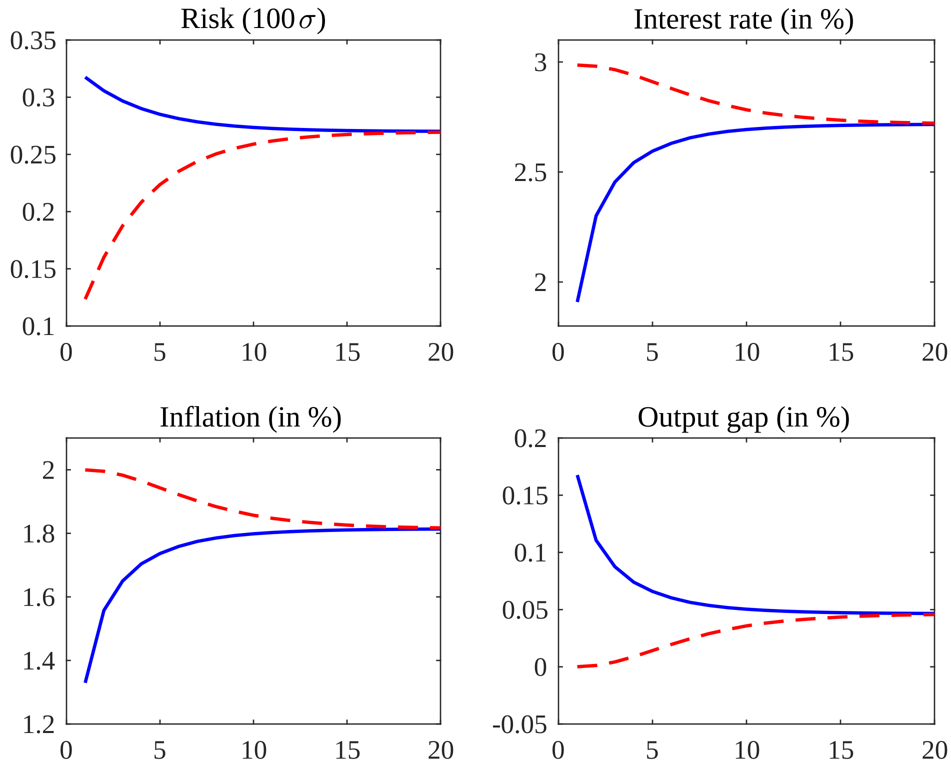

Figure 4: Impulse responses along the zero-shock path to a positive (solid blue lines) and a negative (dashed red lines) risk shock in a hypothetical new normal ($\sigma = 0.27\%$) in the canonical New Keynesian model with a ZLB on interest rates under optimal discretionary monetary policy.

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 (solid blue lines), the policymaker loosens policy enough to bring output above its efficient potential. But compared to a conventional cost-push shock with the same impact effect through (1), 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). As risk falls back, the ZLB becomes less of a concern and the economy gradually returns to the low-risk steady state. Similarly with the simple policy rule (dashed red lines), the policymaker temporarily stimulates the economy in response to falling inflation.

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

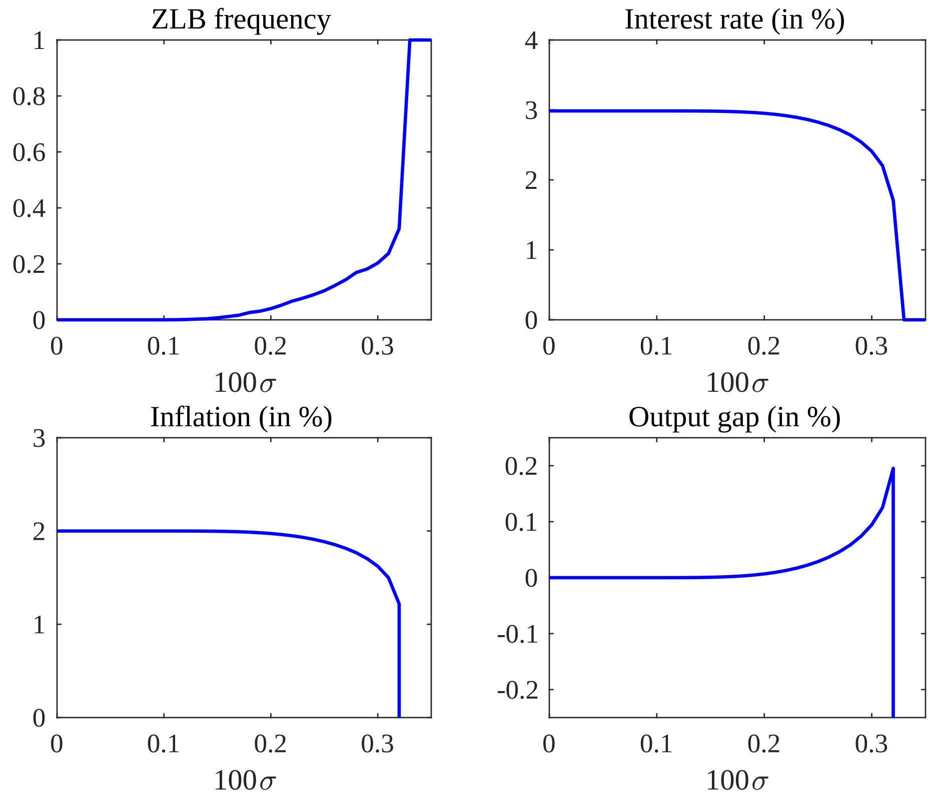

Figure 5: Economic outcomes as a function of risk in the canonical New Keynesian model with a ZLB on interest rate under optimal discretionary policy

outcomes unaffected. Of course, if the economy operates close to the ZLB, more risk shocks become 'large' in this sense. Notice also that both inflation and the output gap would fall after a large increase in risk if monetary policy did not accommodate the risk shock with lower interest rates.

In the new normal, underlying risk is high enough ($\sigma = 0.0027$) relative to the available monetary policy space that expectations are negatively skewed in steady state. In this case, variation in risk have both positive and negative cost-push effects as illustrated in Figure 4. Responses to a positive shock (solid blue lines) are as before, except that the economy reverts to the new-normal steady state with a negative bias in inflation. But a negative risk shock (dashed red lines) now has a positive cost-push effect. As risk falls to low levels, agents stop worrying about the ZLB, and inflation expectations realign with the inflation target. Policymakers increase interest rates in response, while the output gap closes. Gradually, as risk returns, the economy reverts to the high-risk steady state.

The asymmetry in the responses to positive and negative risk shocks around the high-risk steady state reflect a non-linearity in the effect of risk on economic outcomes as illustrated in Figure 5. 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 conceivable states of the world. For small increases, the effects are small. But as risk increases further, the frequency of ZLB episodes increases, and the effects begin to accelerate. For levels of risk beyond a certain point (around $\sigma = 0.0032$), dynamics become explosive with hyperdeflation and a collapse of output. 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.

4.2 On the sources of risk

The baseline risk shock considered so far affects the standard deviations of both shock processes in the model. Figure 6 shows the effects of a positive risk shock around a low-risk steady state for each of the two shocks in turn. Qualitatively, the economy is seen to respond in the same way to the two risk shocks along the zero-shock paths. Spikes in risk lead to cost-push effects both when risk is elevated for the shock to the equilibrium real rate of interest only (solid blue lines) and for the cost-push shock only (dashed red lines). In both cases, optimal monetary policy responds by stimulating the economy enough to push output above potential. It is only the numerical increases in risk required to induce similar quantitative dynamics that are different (top-left panel). The responses are driven by an increase in the likelihood that monetary policy cannot provide sufficient stimulus in response to adverse disturbances when risk is elevated. The sources of the potential adverse shocks are immaterial.

Specifically, the trade-off for monetary policy does not rely on potential cost-push effects from the level shocks themselves. It arises also following spikes in risk to the efficient equilibrium real rate of interest alone. If 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). But with a binding ZLB, monetary policy cannot fully offset large negative shocks to r_t^*. With insufficient monetary stimulus, demand cannot keep up with potential output, and inflation falls below target. A trade-off arises for monetary policy

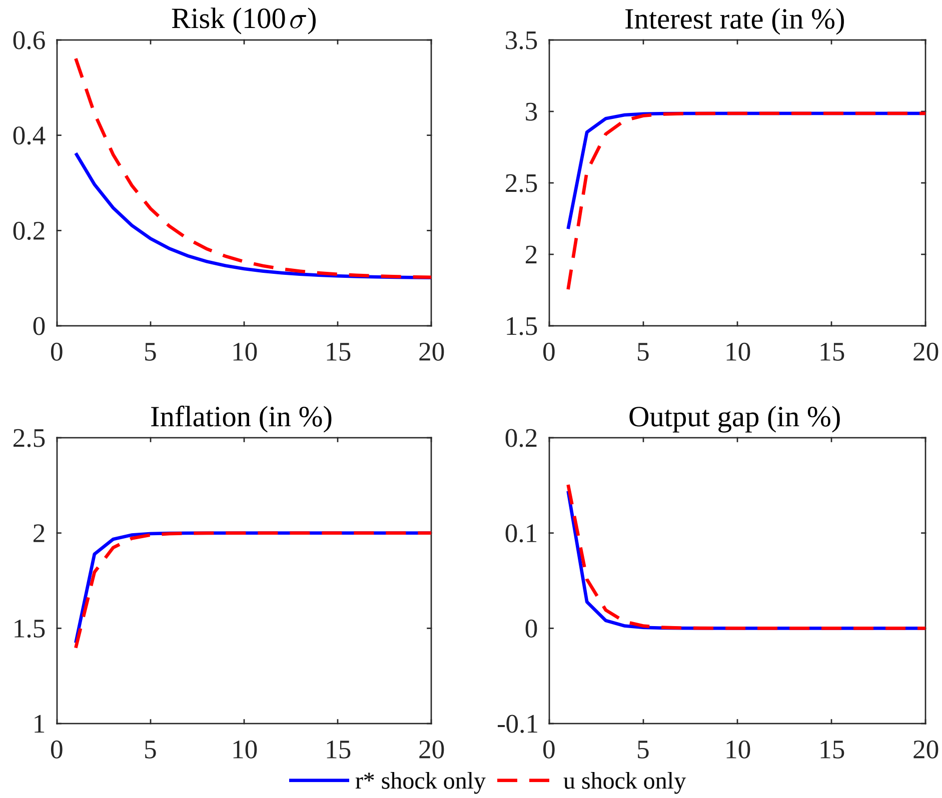

Figure 6: Impulse responses along the zero-shock path to a negative risk shock to the equilibrium real rate of interest only (solid blue lines) and to the cost-push shock only (dashed red lines)around a low-risk steady state ($\sigma = 0.01\%$) in the canonical New Keynesian model with a ZLB on interest rates under optimal discretionary monetary policy.

as the prospect of such demand-driven recessions feed into inflation expectations when risk is elevated.[7] [8]

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, 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.

[7] As 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.

[8] For cost-push shocks, a negative bias in inflation expectations occurs because monetary policy cannot always achieve the appropriate balance between inflation and output following large negative cost-push shocks, while it can always achieve such a balance after positive shocks, see e.g. Evans et al. (2015).

Notice also that, in contrast to Basu and Bundick (2015), 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 simplification 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.[9]

Finally, I remark that the risk shocks considered here are very different from the cross-sectional shocks analysed by Christiano et al. (2014). 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. Within the simple New Keynesian model, such a scenario corresponds to a negative shock to the level of r_t^*.

4.3 Risk shocks at the ZLB

Around a stochastic steady state, optimal discretionary monetary policy responds to risk shocks to mitigate their effects on the economy through expectations. Along the zero-shock path, monetary policy is not actually constrained by the ZLB – except following extreme spikes in risk that lead to hyperdeflations. Effectively, policymakers act now because they may be constrained in future. By contrast, if policy is constrained by the ZLB when a risk shock hits, policymakers are unable to provide further stimulus.

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, the deterministic component is known to follow the path shown in the top-right panel of Figure 7 (dashed green line) so

[9]See Paoli and Zabczyk (2013) for an analysis of the effect of precautionary saving on the equilibrium real rate of interest.

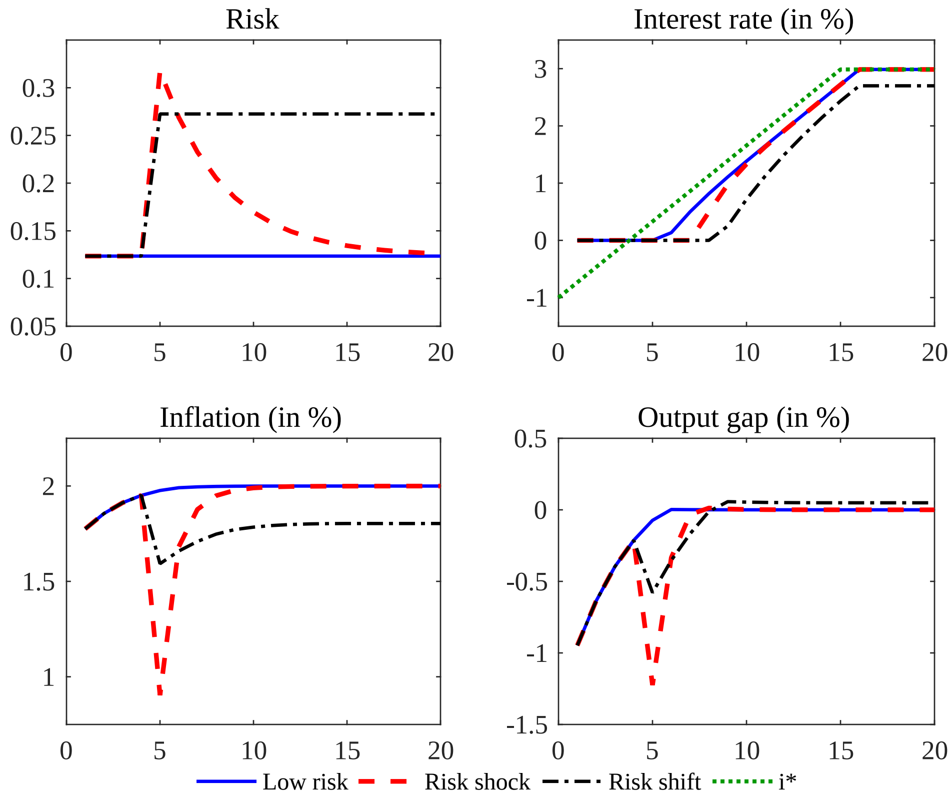

Figure 7: Recovery from a ZLB episode as reflected in the path of the equilibrium interest rate (dashed-dotted black lines in the top-left panel) under optimal discretionary policy in a low risk scenario (solid blue lines) versus a scenario with an unexpected risk shock at time $t = 5$ (dashed red lines).

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.[10] 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 drives the economy back to the ZLB in future is sufficient to optimally delay lift-off even when risk is low.[11]

[10]This corresponds to the perfect foresight case analysed by Adam and Billi (2007) and Guerrieri and Iacoviello (2015)

[11]This 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. Evans et al. (2015) set higher values for the persistence parameters and the weight on output in the loss function, which makes their solution explosive and their results highly sensitive to the choice of T. I can reproduce their results with an uncertainty horizon of $T = 51$. This possibility of explosive

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 3 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 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 stochastic steady state of the new normal as shown in dashed dotted black lines. In this normalisation scenario, optimal policy lifts off from the ZLB late and continues to provide support to economy to lean against low inflation expectations. The optimal trade-off, however, requires the policymaker to accept that that inflation settles below below as the economy recovers to the new normal.

5 Conclusion

Because the ZLB impairs monetary policy's ability to respond to large adverse shocks with sufficiently stimulatory policy action—but never its ability to respond with contractionary action when needed—expectations of inflation and output will be negatively skewed when risk is high given the available monetary policy space. In uncertain times, inflation may settle materially below the policymaker's target in the absence of disturbances even when interest rates are well above the ZLB under optimal discretionary policy.

By implication, variation in risk has potentially large effects on economic outcomes. Even if nothing actually happens, changes in the perception of risk affect the economy through expectations. In the canonical New Keynesian model, risk shocks that are large relative to policymakers' room for manoeuvre give rise to cost-push effects regardless of the source of

dynamics corresponds to the potential non-existence of equilibria analysed by Mendes (2011) under a simple instrument rule, and by Nakata and Schmidt (2014) for the case with optimal discretion.

risk. Around a low-risk steady state, stochastic volatility introduces occasional trade-offs for monetary policy between nominal and real stability, and optimal discretionary monetary policy calls for potentially sharp reductions in the interest rate when risk is elevated. When the underlying risk is high and the economy evolves around a high-risk steady state, variation in risk has both negative and positive cost-push effects. If policy is initially constrained by the ZLB, risk shocks have larger effects on the economy and lift-off is optimally delayed as long as risk is elevated.

The analysis is informative for monetary policy deliberations in inflation targeting countries faced with an effective lower bound on interest rates. While responses are likely to be too immediate in the highly stylised and purely forward-looking model presented here, they are indicative of the direction of the propagation of variations in risk in actual economies operating in an environment in which agents have reason to worry that monetary may be constrained in future. The new normal for monetary policy may be one in which policymakers should respond to changes in the perception of risk, even as the economy escapes the lower bound.

A Solution for each step

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

$$\pi_t(\epsilon, u) = \beta \bar{\pi}^e_{t,t+1}(\epsilon, u) + \kappa x_t(\epsilon, u) + u_t(\epsilon, u)$$
$$\pi_t(\epsilon, u) = -\frac{\lambda}{\kappa} x_t(\epsilon, u)$$

gives the unconstrained optimal allocation

$$\pi^{opt}_t(\epsilon, u) = \frac{\gamma}{\gamma + \kappa^2} \left[\beta \bar{\pi}^e_{t,t+1}(\epsilon, u) + u_t(\epsilon, u) \right]$$
$$x^{opt}_t(\epsilon, u) = -\frac{\kappa}{\gamma + \kappa^2} \left[\beta \bar{\pi}^e_{t,t+1}(\epsilon, u) + u_t(\epsilon, u) \right]$$

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

$$i^{opt}_t(\epsilon, u) = \bar{\pi}^e_{t,t+1}(\epsilon, u) + r^*_t(\epsilon, u) - \sigma \left[x^{opt}_t(\epsilon, u) - \bar{x}^e_{t,t+1}(\epsilon, u) \right]$$

If $i^{opt}_t(\epsilon, u) \geq -i^*$, $\left\{ x^{opt}_t(\epsilon, u), \pi^{opt}_t(\epsilon, u) \right\}$ is the solution for state (ϵ, u). If the ZLB is binding so that $i^{opt}_t(\epsilon, u) < -i^*$, the interest rate is set to $i^{zlb}_t(\epsilon, u) = -i^*$. Now from (2) and (1):

$$x^{zlb}_t(\epsilon, u) = \bar{x}^e_{t,t+1}(\epsilon, u) - \frac{1}{\sigma} \left[-i^* - \bar{\pi}^e_{t,t+1}(\epsilon, u) - r^*_t(\epsilon, u) \right]$$
$$\pi^{zlb}_t(\epsilon, u) = \beta \bar{\pi}^e_{t,t+1}(\epsilon, u) + \kappa x^{zlb}_t(\epsilon, u) + u_t(\epsilon, u)$$

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

$$\left\{ x^{sol}_t(\epsilon, u), \pi^{sol}_t(\epsilon, u) \right\} = \begin{cases} \left\{ x^{opt}_t(\epsilon, u), \pi^{opt}_t(\epsilon, u) \right\} & \text{if } i^{opt}_t(\epsilon, u) \geq -i^* \\ \left\{ x^{zlb}_t(\epsilon, u), \pi^{zlb}_t(\epsilon, u) \right\} & \text{if } i^{opt}_t(\epsilon, u) < -i^* \end{cases}$$

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

$$\bar{\mathbf{x}}^e_{t-1,t} = \mathbf{P}_\epsilon \mathbf{x}^{sol}_t \mathbf{P}'_u$$
$$\bar{\pi}^e_{t-1,t} = \mathbf{P}_\epsilon \pi^{sol}_t \mathbf{P}'_u$$

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

References

Aastveit, K. A., G. J. Natvik, and S. Sola (2013). Economic uncertainty and the effectiveness of monetary policy. Working Paper 2013/17, Norges Bank.

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.

Basu, S. and B. Bundick (2014). Uncertainty shocks in a model of effective demand. Research Working Paper RWP 14-15, Federal Reserve Bank of Kansas City.

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.

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.

Black, F. (1995). Interest rates as options. *The Journal of Finance 50*(5), 1371–1376.

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.

Castelnuovo, E., G. Caggiano, and G. Pellegrino (2015). Estimating the real effects of uncertainty shocks at the zero lower bound. "Marco Fanno" Working Paper 200, Universita Degli Studi di Padova.

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.

Del Negro, M., D. Giannone, M. Giannoni, and A. Tambalotti (2017). Safety, liquidity, and the natural rate of interest. Staff Reports 812, Federal Reserve Bank of New York.

El-Erian, M. A. (2010). *Navigating the new normal in industrial countries.* Per Jacobsson Foundation.

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* (BPEA Conference Draft, March 19-20, 2015).

Fernandez-Villaverde, J., P. Guerron-Quintana, K. Kuester, and J. Rubio-Ramirez (2015). Fiscal volatility shocks and economic activity. *American Economic Review* (forthcoming).

Galí, J. (2008). *Monetary Policy, Inflation, and the Business Cycle. An Introduction to the New Keynesian Framework.* Princeton, New Jersey: Princeton University Press.

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. (2015). Stuck. *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.

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.

Laubach, T. and J. C. Williams (2016). Measuring the natural rate of interest redux. *Business Economics 51*(2), 57–67.

Mendes, R. R. (2011). Uncertainty and the zero lower bound: A theoretical analysis. MPRA Paper 59218, University Library of Munich, Germany.

Nakata, T. (2013). Uncertainty at the zero lower bound. Finance and Economics Discussion Series 2013-09, Board of Governors of the Federal Reserve System.

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, 02). Cyclical risk aversion, precautionary saving, and monetary policy. *Journal of Money, Credit and Banking 45*(1), 1–36.

Plante, M. D., A. Richter, and N. Throckmorton (2014). The zero lower bound and endogenous uncertainty. Working Papers 1405, Federal Reserve Bank of Dallas.

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.

Rushby, A. J., M. W. Claire, H. Osborn, and A. J. Watson (2013). Habitable zone lifetimes of exoplanets around main sequence stars. *Astrobiology 13*(9), 833–849.

Schmitt-Grohe, 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). Chapter 13 - 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.

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.