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

Martin Seneca[a]

July 2018

[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 downside deanchoring of inflation expectations—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, 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 downside deanchoring of inflation expectations, but it allows 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 then 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.[1]

[1]For a general class of models without inequality constraints, Schmitt-Grohe and Uribe (2004)

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).[2] As noted by Kiley (2018), no 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. By the risky steady state I mean the point at which the economy settles when disturbances have abated but agents are aware that innovations may occur in the future (in line with the definition in Coeurdacier et al., 2011). By risk shocks I mean changes in the standard deviations of conventional level shocks in the model. I trace out the 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 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 public's *perception* of risk, where risk is defined as the set of variances of the independent distributions from which people believe innovations to level shock processes are drawn.[3] The simple quasi-linear structure comes with the benefit that risk affects the economy exclusively through its interaction with the ZLB. This allows me to zoom in on the defining feature of the new environment—the 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.[4]

show that risk does not affect decision rules to a first-order approximation, and only the constant term up to a second order.

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

[3]I write *risk* rather than *uncertainty* shocks to remain loyal to the old Knightian distinction; agents in the model actually assign a number to the level of risk, though it may well be that elevated risk in the model stands in for Knightian uncertainty in reality.

[4]For analyses of how higher-order behavioural effects of risk (such as precautionary savings by

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 ZLB gives rise to a negative bias in expectations in the New Keynesian model's stochastic equilibrium. In a recent application, Nakata and Schmidt (2014) suggest that the skew 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. My contribution is to characterise the economy's risky steady state and trace out dynamics around it when risk changes to facilitate a consideration of monetary policy when the ZLB is not currently binding but nevertheless constantly looming as a cloud on the economic horizon. I show how a risk shock propagates through dynamics in the skew in expectations in this environment in a way that induces trade-offs for monetary policy regardless of the source of risk.[5]

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. 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 further '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 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, whilst Section 3 describes the calibration. Section 4 presents the risky state state, and Section 5 impulse responses to risk shocks. Section 6 contains normalisation scenarios after a ZLB episode. Finally, Section 7 concludes.

---

households and cautious investment decisions by firms) may be amplied when the ZLB is binding, see e.g. Basu and Bundick (2014, 2015) and Nakata (2017). Fernandez-Villaverde et al. (2015) and Johannsen (2014) suggest that uncertainty about fiscal policy in particular has larger implications when monetary policy is constrained.

[5]My results on the risky steady state are fully consistent with the independent and contemporaneous analysis by Hills et al. (2016). The authors estimate an inflation target shortfall of 25bp in a risky steady state for the US.

2. 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:

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

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

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

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 model is derived from its microfoundations by Galí (2008) and Woodford (2003) 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 $\nu_{\epsilon,t} \sim N(0, \sigma_{\epsilon,t}^2)$ 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. Specifically, I let a baseline risk shock be such that $\varsigma^{-1}\sigma_{\epsilon,t} = \sigma_{u,t} = \sigma_t$ with

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

where $\nu_{\sigma,t}$ is the innovation to risk, and σ is the underlying level of risk in the absence of risk disturbances.[6] In what follows, agents observe the current level of risk, but they do not expect further changes in this level to occur (as I discuss further below).

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

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

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

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

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

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 interest rate will be set to the maximum of this optimal level and zero.[7]

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 $E_t\pi_{t+1} = E_t x_{t+1} = 0$ 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, (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 the Appendix 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).

[7] This can be confirmed formally by deriving the Kuhn-Tucker conditions, see e.g. Nakov (2008).

I can now 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 function is the 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. Effectively, I deviate from certainty equivalence only for level shocks. It is also a natural starting point for considering variation 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 would, however, generate similar qualitative responses in this quasi-linear model.

3. Calibration

Figure 1 provides a simple illustration of how the hypothetical new normal differs from recent pre-crisis historical experience in the United States (left panel) and the United Kingdom (right panel). The normal probability density functions in 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 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 both the FFR and BR shifted sharply to the left as inflation targeting became established and inflation expectations anchored at lower levels, see the dashed red lines. 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. Any Gaussian model fitted to interest rate data over the period would have produced a probability of negative interest rates close to zero. 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, I define the 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. The higher spread reflects 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 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

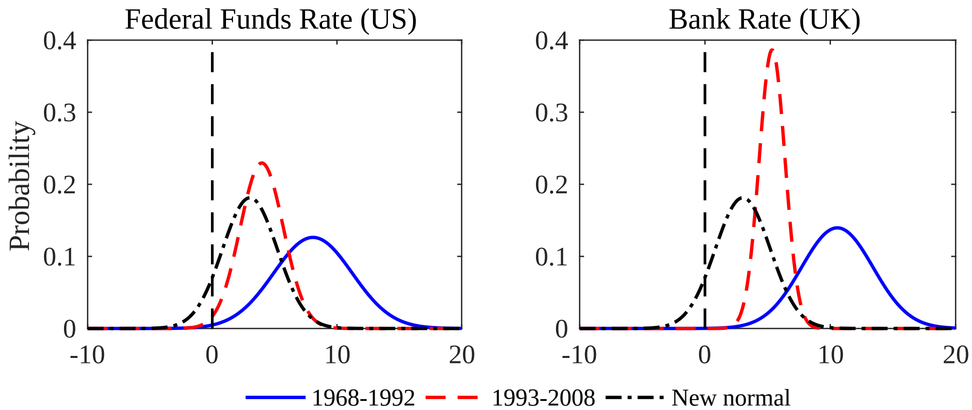

Figure 1: Normal probability density functions with means and standard deviations set equal to sample means and standard deviations for the Federal Funds Rate (left panel) and Bank Rate (right panel) for the sample periods 1968-1992 (solid blue lines) and 1993-2016 (dashed red lines) as well as for a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

arbitrary, this probability results from setting the average desired policy rate to 3%—in line with the forecasts emphasised by Reifschneider (2016), and the estimates in Kiley (2015), 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.

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 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 to $\varsigma = 1$ 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 interpretion of the degree of flexibility in inflation targeting in practice (see e.g. Carney, 2017, English, López-Salido, and Tetlow, English et al., 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

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_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: $E(.)$ 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_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.

(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 so that the desired policy rate is negative with probability 9%.

Table 1 compares the new normal with an alternative low-risk scenario as well as with calibrations that allow the unconstrained model to match the historical distributions in Figure 1. In the low risk scenario, 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 set to 1pp (similar to the level observed for the UK between 1993 and 2008), which requires $\sigma = 0.0012$. The probability that interest rate should be negative remains negligible in this case despite a low level of r^*. In Section 5, I show impulse responses to risk shocks in this low-risk scenario as well as in the baseline new normal. For the historical episodes, I assume that $r^* = 3.75$ in the 1968-1992 period, and $\pi^* = 2\%$ between 1993 and 2016. These examples serve simply to show that, even if the model is too simple to capture the covariance structure of the data more broadly, it gives a fairly good fit to observed inflation volatility. 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 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, perhaps, that unconventional policies such as 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* 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.

In what follows, I 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. 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. 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. Risky steady state

Turning to the full model with a ZLB, the first row in Table 2 shows risky steady state outcomes for the interest rate (i), inflation (π) and the output gap (x) with 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 deanchoring of inflation expectations by operating the economy above potential in normal times through a stimulatory bias in policy rates. Such an outcome is feasible because the monetary policy stance interacts with inflation expectations to determine the real interest rate. Monetary policy is non-neutral in the new 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

Table 2: New Keynesian model with ZLB under optimal discretion

		Interest rate			Inflation			Output gap			
	Scenario	i^*	i	$E(i)$	π^*	π	$E(\pi)$	x^*	x	$E(x)$	P_{ZLB}
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, $E(.)$ and $\sigma(.)$ denote means and standard deviations, respectively, and P_{ZLB} denotes the frequency of a binding ZLB.

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. The second row shows the effect of 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 deanchoring 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

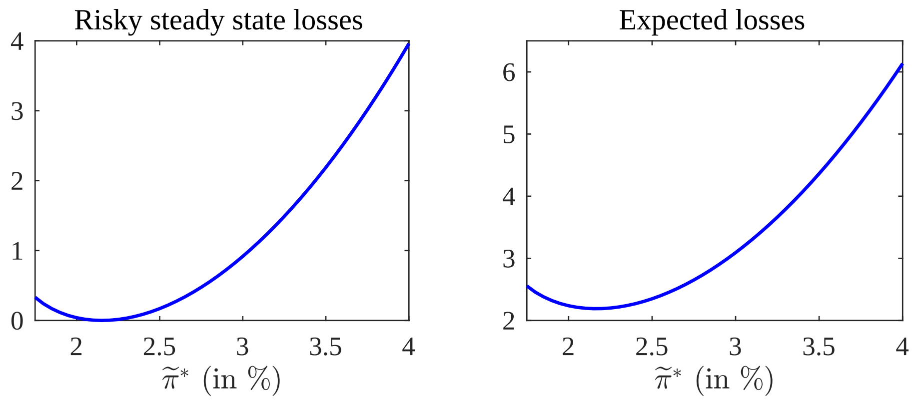

Figure 2: Period welfare losses with an optimal annual inflation rate of 2% as a function of an operational inflation target denoted by $\widetilde{\pi}^*$. Left panel shows losses evaluated in the risky steady state, and right panel unconditionally expected losses.

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).

5. 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.

5.1. The case of low underlying risk

Solid blue lines in Figure 3 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 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.

5.2. On the sources of risk

Figure 3 also 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 responds in the same way to the two individual shocks. Spikes in risk lead to cost-push effects both

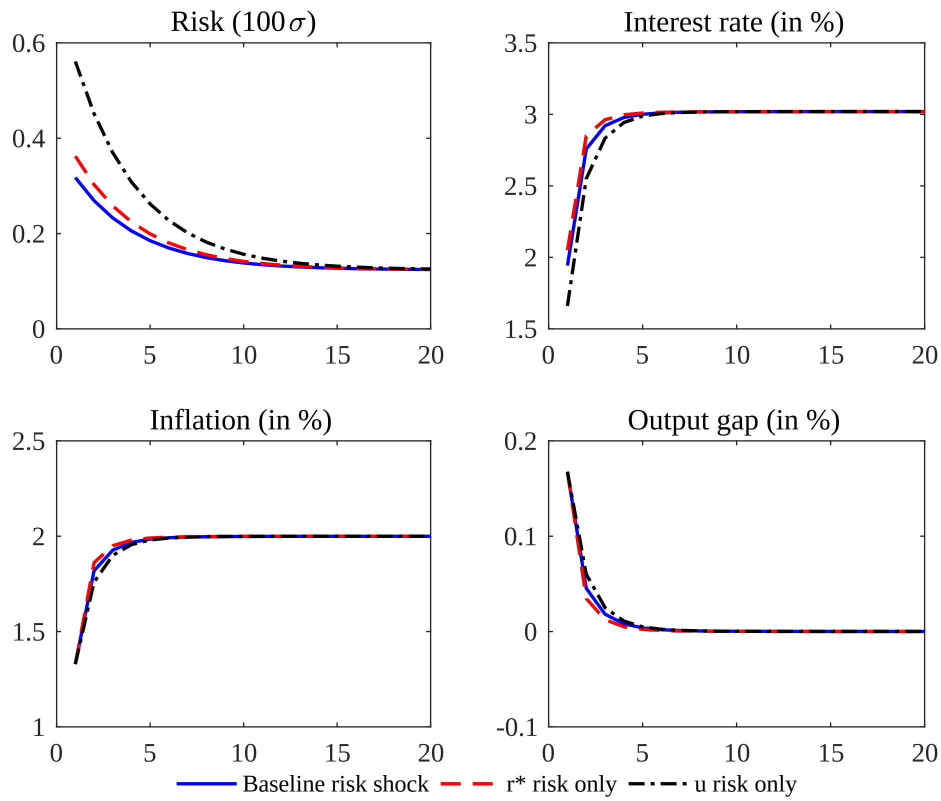

Figure 3: Impulse responses to a baseline risk shock (solid blue lines), a shock to r^* risk only (dashed red lines), and to cost-push risk only (dashed-dotted black lines), around a low-risk steady state ($\sigma = 0.0012$) in the canonical New Keynesian model with a ZLB on interest rates under optimal discretionary monetary policy.

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.[8] A trade-off arises for monetary policy as the prospect of such demand-driven recessions feed into inflation expectations when

[8]If monetary policy were unrestricted by the ZLB, shocks to r_t^* could of course 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).

risk is elevated.[9]

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.

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.[10]

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^*.

5.3. 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 4. 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

[9] 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.

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

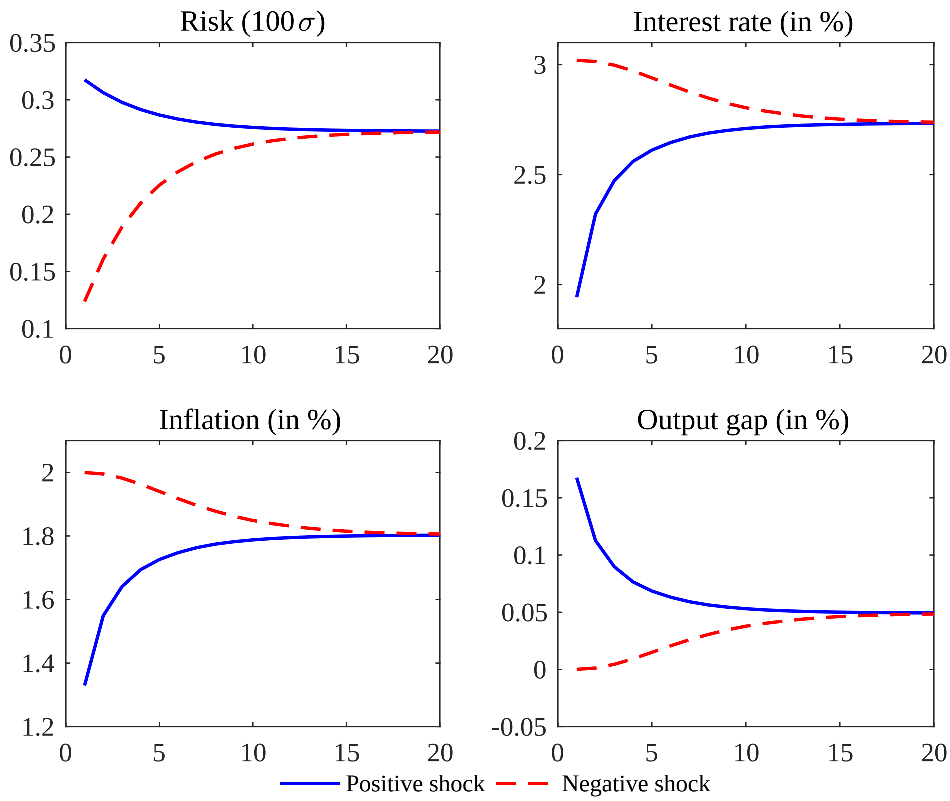

Figure 4: Impulse responses to a positive (solid blue lines) and a negative (dashed red lines) baseline risk shock around the new normal risky steady state ($\sigma = 0.0027$) in the canonical New Keynesian model with a ZLB on interest rates under optimal discretionary monetary policy.

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 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 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 states of the word. For small increases, the effects are small. But as risk increases further, the frequency of ZLB episodes increases. Effects begin 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

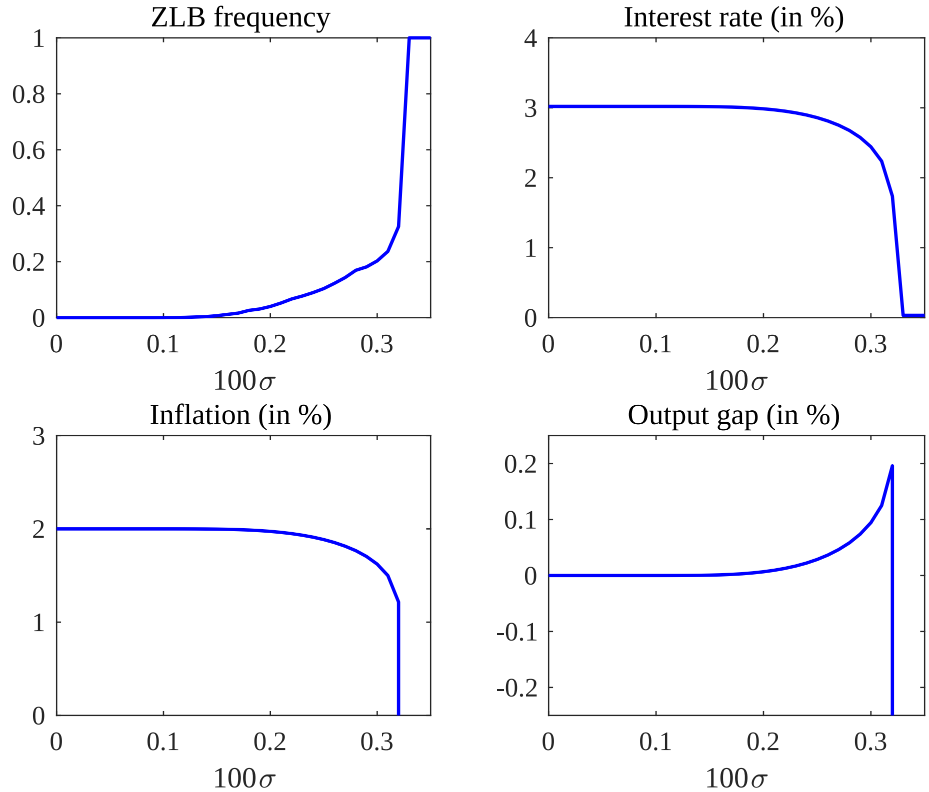

Figure 5: Economic outcomes as a function of risk in the canonical New Keynesian model with a ZLB on interest rate under optimal discretionary policy

a collapse of output.[11]

6. Normalisation scenarios

To illustrate the implications of a binding ZLB for the propagation of risk shocks, Figure 6 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 driven by ρ_t is known to follow the path shown in the top-right panel of Figure 6 (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

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

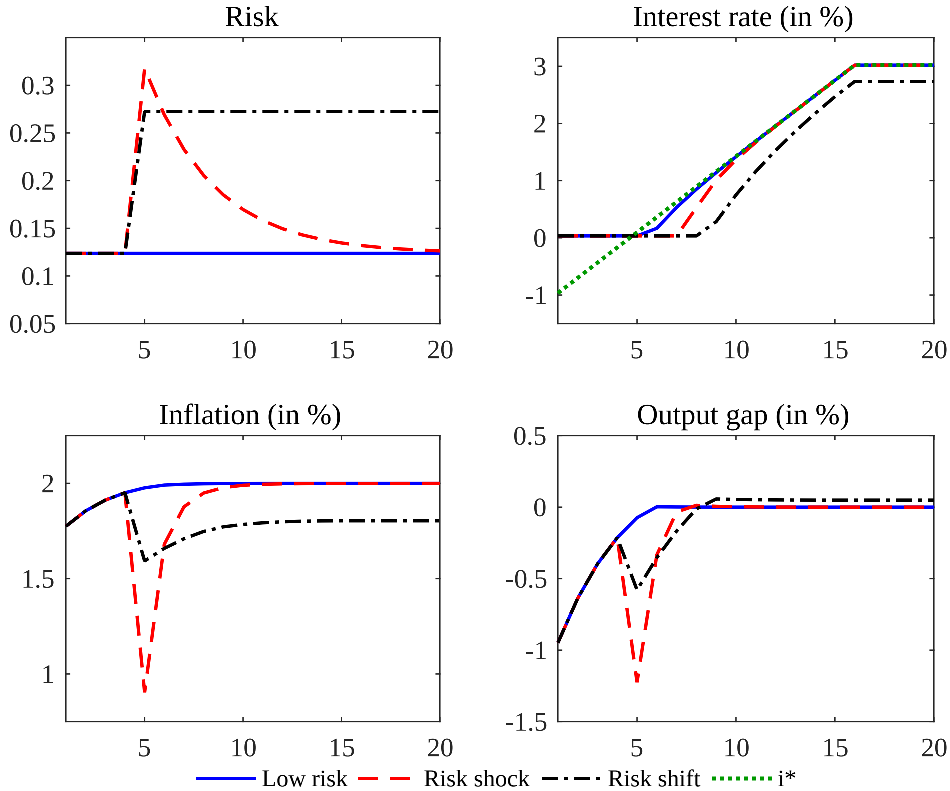

Figure 6: Recovery from a ZLB episode to as reflected in the path of the equilibrium interest rate (dotted green lines in the top-right panel) under optimal discretionary policy in a low risk scenario (solid blue lines), in a low risk scenario with a baseline risk shock (dashed red lines), and with a shift in risk (dashed-dotted black lines).

interest rate on its trajectory back towards normal levels once it exceeds the ZLB.[12] 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.[13]

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 6 (dashed red lines),

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

[13]This is the argument made in Evans et al. (2015). But in Figure 6 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.

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 risky steady state as shown in dashed dotted black lines. In this normalisation scenario, optimal policy lifts off from the ZLB late and continues 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.

7. 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. While responses are likely to be too immediate in the highly stylised and purely forward-looking model, they 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 future. The new normal may be one in which monetary policy should lean against a downside deanchoring of inflation expectations by operating the economy above potential in normal times, and respond to changes in the perception of risk even as the economy escapes the ZLB.

Appendix 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

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

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 (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.

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 (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.

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, November). 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.

Carney, M. (2017). Lambda. Speech, London School of Economics, London, 16 January.

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.

Debortoli, D., J. Kim, J. Linde, and R. Nunes (2016). Designing a simple loss function for the Fed: Does the dual mandate make sense? Technical report.

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. The Federal Reserve's framework for monetary policy: Recent changes and new questions.

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*.

Fernandez-Villaverde, J., P. Guerron-Quintana, K. Kuester, and J. Rubio-Ramirez (2015). Fiscal volatility shocks and economic activity. *American Economic Review 105*(11).

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.

Harrison, R. (2017, September). 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.

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*.

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.

Ljungqvist, L. and T. J. Sargent (2000). *Recursive macroeconomic theory*. Cambridge, Massachusetts: MIT Press.

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

Nakata, T. (2017, July). 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, 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.

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). 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).