Optimal monetary policy, least squares learning, and the zero bound to interest rates

February 18, 2019

Petr Harasimovič
Loughborough University

Anthony Yates
City University of London and Centre for Macroeconomics

Abstract

We revisit the question of optimal monetary policy at the zero lower bound by modifying the workhorse assumption of rational expectations used in most previous analysis to instead assume that expectations are formed using a constant gain least squares learning algorithm. The main result is that, unlike in the case of RE, there is no tendency for optimal policy at the zero bound to imply an overshooting of the inflation target. Under RE, expected future overshooting provides an additional stimulus by lowering the real rate. In our model, overshooting risks destabilising belief coefficients in future periods and so the policy maker foregoes this.

Keywords: optimal monetary policy, New Keynesian, learning.
JEL classification: E52, E61, E66, C14, C18

1 Introduction

This paper computes optimal interest rate policy in a widely studied New Keynesian sticky price model of the business cycle, but where the

---

Email: P.Harasimovic@lboro.ac.uk
Email: anthony_yates@yahoo.com
This work was carried out using the computational facilities of the Advanced Computing Research Centre, University of Bristol.

more usual assumption that expectations are model-consistent (often dubbed 'rational') is replaced with the assumption that agents form expectations using constant-gain, least-squares learning algorithms. In addition, optimal policy is studied when the central bank policy-maker faces the constraint that interest rates cannot be reduced below their natural floor of zero. Since several economies in the developed world have hit the ZLB recently (e.g. Japan in 1999, UK and US in 2008, or later the Eurozone, Sweden, and Norway), studying the design of optimal monetary policy in a liquidity trap has become of paramount importance for policy makers and not just an academic matter. Although the US and the UK have begun a lift-off from the zero lower bound, the question is still germane, since the length of the zero-lower-bound experience made the central banks more aware of the risks of hitting the bound gain. Our analysis thus contributes to the literature on this subject that has flourished recently.

The paper thus hopes to make useful comparisons with two separate strands in the previous literature on monetary policy design. The first is the literature on optimal monetary policy in the face of the ZLB but when expectations are rational. The classic reference here is Eggertsson and Woodford (2003) (EW). In this paper, the authors compare the dynamics of the central bank instrument and inflation under increasingly large shocks to the natural rate of interest that threaten to take rates to the ZLB, and, at some point, cause the ZLB to bind. They observed that as the size of the shock increased, the response of rates under optimal policy became one where rates were lower, for longer. Consequently, in response to shocks in which the ZLB was in play, the central bank allowed a corresponding overshoot of inflation from target.

'Lower for longer' meant that interest rates would be lower, proportionately, than in the case where the ZLB did not bind, and, likewise, 'for longer' referred to the fact that rates would be away from steady-state for longer. This result became a reference point for central banks facing the ZLB during the recent crisis, who engaged in forms of what they often described as 'forward guidance' over future interest rates. Typically, those central banks kept some degree of scepticism and resisted encouraging the expectation that they would tolerate a future overshoot of their inflation targets, but nevertheless, the EW result was influential in pushing central banks to explain why they thought that EW's prescriptions should not be followed to the letter. Our findings suggest that the caution on the side of the central banks was, indeed, justified if one did not have enough trust that the public had fully rational expectations.

The crux of EW's advice was that, in the face of being unable to reduce current interest rates further, due to the ZLB, a central bank acting credibly could commit to lowering future interest rates instead. With such a promise factored into the expectations of rational agents, this would drive today's expected inflation rate up, lower the real rate, and increase demand and the output gap. Through the Phillips Curve then, today's inflation rate would also increase. Even though such a promise came at the cost of a future inflation overshoot, this, up to a point, would be worth it for the sake of curtailing the undershoot in inflation and the output gap experienced today.

Following EW, Adam and Billi (2006, 2007) extend the analysis to a model with uncertainty and both with and without commitment. Their model is the closest rational expectations counterpart to our model. Werning (2011) then analyzes optimal policy at the zero lower bound in conjunction with fiscal policy in a model cast in continuous time.

Rather obviously, the efficacy of EW's policy advice depends on the extent to which agents can be taken to behave as though they form rational expectations. If agents do not form their expectations by computing model-consistent forecasts, including therefore correct assumptions about the element of the model that corresponds to policymakers' behaviour, then vocal commitments to future lower interest rates will be in vain. We use the model of least squares learning popularized by the likes of Bray, Sargent and Marcet, and Evans and Honkapohja.[1] Our contribution, therefore, is to see how optimal policy in the face of the ZLB compares when the assumption of rational expectations is replaced with one in which agents form expectations using least-squares learning algorithms and which is generally considered a small departure from rational expectations. The precise form we use nests the rational expectations equilibrium (REE) and, without the ZLB, we observe a convergence to it.

Our main contribution is to show that with this mild departure from rational expectations there is little to no tendency for inflation to overshoot, in fact, the amount of inflation that the optimal policy generates is lower for longer episodes of binding ZLB, while in EW the opposite is the case. Nevertheless, somewhat surprisingly, the above-

1. See e.g. Bray (1982), Marcet and Sargent (1989), and also Evans and Honkapohja (2001) for a review of the early literature.

mentioned lower-for-longer result of EW still holds in some ways: in response to negative shocks, and in the presence of the ZLB, the interest rate is lowered more aggressively, it reaches the ZLB sooner, it is kept at the bound for longer, and it returns to pre-shock levels slower than in the absence of the ZLB. The central bank is not willing to tolerate excessive inflation though. Thus, one could perhaps say that the interest rates are lower for longer but that they are not very much lower and for not very much longer.

Our results also address and contribute to the prior literature on the design of monetary policy that has made use of the sticky-price business cycle model with expectations formed using learning. First, we are the first in this literature, so far as we can tell, to add the ZLB into the analysis. Second, even leaving the zero bound aside, our work complements and contributes to prior work in a number of ways. Notable prior papers – none of which consider the ZLB – are Gaspar, Smets, and Vestin (2006) and Molnár and Santoro (2014). Prior to these two papers, there were several papers that studied sparsely parameterised, simple Taylor-like policy rules in a New Keynesian model with expectations formed using least squares learning (e.g. Orphanides and Williams 2008). Relative to that literature, our paper offers computations of fully optimal policy.

Molnár and Santoro (MS) generate analytical results on optimal monetary policy in the presence of 'cost push' shocks, but at the price of simplifying the learning algorithm somewhat. In their model, there is a recursion over agents forecasts (so called *steady state learning*), whereas we adopt the framework of Evans and Honkapohja (2001) in which there is a recursion over agents forecasting functions. Compared to MS, we find that, outside the rational expectations equilibrium, the demand shock introduces a policy trade-off making it a relevant policy consideration. Such a trade-off does not occur in models with rational expectations and, as far as we know, we are the first to observe it in a model with learning.

Gaspar, Smets, and Vestin (2006), like us, derive optimal monetary policy using numerical methods but they simplify the model in two respects relative to our work. First, they characterise monetary policy as a choice over the output gap (in the absence of the zero lower bound, this is without any loss of usefulness or generality, in fact). Second, their agents form expectations using an evolving AR(1) relating inflation to its lag, whereas our agents use recursions over two functions that link future inflation and future output to the shock, the form that nests rational expectations.

Our work is also related to studies of deflationary traps which may occur due to a binding effective lower bound in models where the policy maker uses a simple Taylor-like rule both with rational expectations and with least squares learning.[2] In comparison, our policy maker uses fully optimal policy. Although, we do not focus specifically on the question of the existence of a deflationary trap in our setup, we emphasize the parallel between the optimal policy in our model and the aggressive monetary policy rule of Evans, Guse, and Honkapohja (2008).

Besides adaptive learning, our work is, more generally, related to the literature which analyzes government policies under the assumption of bounded rationality. One such strand of literature, which is gaining popularity, uses level-k thinking. For recent examples see Farhi and Werning (2017) and Iovino and Sergeyev (2017), who analyze the effectiveness of forward guidance and quantitative easing, respectively. Both adaptive learning and level-k thinking converge to the rational expectations equilibrium, under some conditions, and thus, from all the ”wilderness” of bounded rationality approaches, these are reasonable attempts at relaxing the strict rational expectations assumption.[3] Compared to level-k thinking, econometric learning is purely backward looking and, hence, we deal with the extreme case of complete ineffectiveness of policy announcements.

To summarize our main findings, we observe that, in contrast to EW, the policy maker does not allow as much inflation overshooting and the optimal policy becomes more restrictive the longer the period of binding ZLB. Nevertheless, similarly to EW, the constrained central bank responds to negative demand shocks disproportionately more than its unconstrained counterpart, it reaches the ZLB sooner, it keeps the interest rate at the bound for longer, and it allows slower convergence back to pre-shock levels. At times when the ZLB is not binding, the central bank follows an over-aggressive policy to create counter-cyclical expectations as a means to alleviate the cost of a potential liquidity trap.

Without the effective lower bound in place, we find that the economy converges to the rational expectations equilibrium, despite the fact that our agents use constant gain learning, and in response to

2. See Benhabib, Schmitt-Grohé, and Uribe (2001), Evans and Honkapohja (2005), and Evans, Guse, and Honkapohja (2008).

3. The term ”wilderness” comes from Sims (1980).

demand shocks, when not in the REE, the central bank faces a policy trade-off between inflation and output. The latter property contrasts with models with rational expectations and has previously not been recognized in models with learning.

The rest of the paper is organized as follows. In section 2 we present the model, in section 3 we discuss calibration and give a brief overview of the computational procedure, leaving the details for the appendix. Sections 4 and 5 discuss the results for the benchmark and alternative calibrations, respectively, and section 6 contains concluding remarks.

2 The Model

As the model framework we use the, now stylized, New Keynesian 3-equation sticky-price model in its linearized form. The model consists of two structural equations dubbed Phillips curve (PC) and IS curve (IS) and a monetary policy rule equation (MPR). Using notation standard in the literature the IS and PC equations can be stated as

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + \varepsilon_t^{\pi} \tag{1}$$

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1}) + \varepsilon_t^{x}. \tag{2}$$

Equation (1), the Phillips curve, relates the deviation of quarterly inflation from its steady state value, π_t, to discounted expected future inflation, and the output gap, x_t. The gap concept here, in line with the New Keynesian literature, denotes the difference between output that would obtain under flexible prices and actual output. This equation emerges as an approximation to optimal price-setting by firms in the presence of price-stickiness in the tradition of Calvo (1983) and Rotemberg (1982). The equation includes an inflation shock, ε_t^{π}, which has the interpretation of a shock to desired mark-ups (cost push shock) in the original nonlinear model to which this equation is an approximation.

Equation (2), the IS curve, is an approximation to the aggregate consumption-Euler equation and it draws the link between the central bank's interest rate, i_t, and the output gap. The IS curve also includes a shock, ε_t^{x}, which is to be interpreted as a shock to the natural rate of interest, or, equivalently, a technology shock.

The shocks are both assumed to follow a stationary $AR(1)$ stochastic process of the form[4]

$$\varepsilon_t^\pi = \rho_\pi \varepsilon_{t-1}^\pi + \eta_t^\pi, \quad \eta_t^\pi \sim iid\, N(0, \sigma_\pi^2) \tag{3}$$

$$\varepsilon_t^x = \rho_x \varepsilon_{t-1}^x + \eta_t^x, \quad \eta_t^x \sim iid\, N(0, \sigma_x^2) \tag{4}$$

where η_t^π and η_t^x are the innovations with variances σ_π^2 and σ_x^2, respectively.

Conventional applications posit rational expectations but we assume that agents form their expectations using a least-squares learning recursion. The expected values of $Y \equiv [\pi, x]$ are thus given by

$$E_t[Y_{t+1}] \equiv E_t \begin{bmatrix} \pi_{t+1} \\ x_{t+1} \end{bmatrix} = A_t \begin{bmatrix} \varepsilon_t^\pi \\ \varepsilon_t^x \end{bmatrix}, \tag{5}$$

where A is a matrix of expectations coefficients.

If we substitute agents' beliefs (eq. 5) in equations (1) and (2) we can write the IS and PC curves compactly as

$$Y_t = B^{-1} D i_t + B^{-1}(CA_t + F)\varepsilon_t, \tag{6}$$

with

$$B = \begin{bmatrix} 1 & -k \\ 0 & 1 \end{bmatrix}, C = \begin{bmatrix} \beta & 0 \\ \sigma & 1 \end{bmatrix}, D = \begin{bmatrix} 0 \\ -\sigma \end{bmatrix}, F = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \varepsilon = \begin{bmatrix} \varepsilon_\pi \\ \varepsilon_x \end{bmatrix}. \tag{7}$$

The expectations matrix, A_t, evolves over time according to the following learning recursion

$$A_{t+1} = A_t + \gamma R_{t+1}^{-1} \varepsilon_t (Y_t^{\mathrm{T}} - \varepsilon_t{}^{\mathrm{T}} A_t{}^{\mathrm{T}}) \tag{8}$$

$$R_{t+1} = R_t + \gamma(\varepsilon_t \varepsilon_t{}^{\mathrm{T}} - R_t). \tag{9}$$

In this system there are two exogenous state variables, stacked in the vector ε_t, and the endogenous states comprise the four elements of matrix A_t and the four elements of matrix R_t, the moment matrix. Since the two-shock version of the model is nearly intractable once the zero lower bound is assumed, we disregard the inflation shock. In

4. Often in the literature the shocks are *iid*, however, for an analysis of the zero lower bound we need the model to exhibit some form of persistence. Using $AR(1)$ shocks is a cheap way of achieving this. Two alternatives, often found in the literature, would be using inflation indexation and habit formation. Both would however increase the number of state variables and hence increase the computational cost.

this case the system comprises of two elements of A_t, one element of R_t, and one shock, the technology shock ε_t^x.[5]

The motivation for using the learning recursion as described by equations (8) and (9) as a small departure from rational expectations comes from the fact that if we set $\gamma = \frac{1}{t}$, an assumption known as decreasing gain, agents will, in some circumstances, learn the rational expectations functions. Our assumption of constant gain means that agents are forced to inappropriately downweight data early in the sample (the observations are weighted geometrically) thus they can never learn the rational expectations equilibrium. Such an assumption resembles the idea of a rolling window regression and the motivation is that either agents suspect (incorrectly) that there might be structural change in the economy, and thus older data are less relevant, or that they have short memories, or insufficient processing facilities to handle long time series.

To close the model we need to specify a monetary policy rule (MPR), which is the remaining equation in the 3-equation model. Typically the MPR takes the form of a simple, Taylor-like, rule but here we focus on the problem of choosing the interest rate optimally. Therefore the interest rate is given by the central bank's policy function, $h(A_t, R_t, \varepsilon_t)$, which determines the interest rate as a function of current period state variables – the beliefs and moment matrices, A_t and R_t, and the shocks, ε_t. This function is derived from the central bank's preferences represented by the loss function $L(\pi_t, x_t) \equiv \pi_t^2 + \lambda x_t^2$, where λ is the weight the policy maker puts on stabilizing output relative to inflation.

This loss function is not purely ad hoc as it has micro-foundations in the New Keynesian model with rational expectations. As observed by Woodford (2003), under rational expectations, the use of the weight $\lambda = \frac{\kappa}{\theta}$ (i.e. proportional to the strength with which changes in the output gap induce changes in inflation) renders this criterion function one that maximises the welfare of the representative agent in this economy.[6] This will not be precisely true in our economy under learning, but we use this assumption nonetheless presuming it to be a reasonable approximation and to make the problem tractable. It

5. There are techniques for solving medium and large scale models, e.g. Smolyak interpolation or the Generalized Stochastic Simulation Algorithm of Judd, Maliar and Maliar, but those are based on low-order polynomial approximation, which generally does not work very well when the policy function contains a kink.

6. See Woodford (2003), section 6.2.

also allows us to make our results comparable to prior work in the RE and non-RE traditions, which have used criterion functions for policy makers of the same kind.

When setting the interest rate the policy maker is constrained by the presence of the zero lower bound on the nominal interest rate given by the rate of return on money. Due to the fact that economic agents always have the option to hold their wealth in money, which yields zero return, the interest rate cannot be lowered below this level. Since the model is expressed in terms of log deviations from the (zero inflation) steady state and the steady state interest rate is positive, the zero lower bound is not zero but to be interpreted as the log deviation from the steady state interest rate that would bring the gross nominal rate of interest to its lower bound of one. Thus the value of the zero lower bound in terms of this log deviation is given by $-\log(\frac{1}{\beta}+\pi^*)$, where π^* is the central bank's inflation target (hence the steady state rate of inflation).[7]

Our central bank policy maker then seeks to minimise the loss function taking the New Keynesian model equations, the learning recursions, and the zero lower bound as constraints. We assume that unlike private agents the central bank has rational expectations. This assumption is relatively common in the literature and can be motivated by the apparent informational advantage of central banks over the private sector coming from the large amount of resources, namely human capital, central banks typically dedicate to understanding the economic environment. The policymaker's problem is formally stated as

$$\min_{\{i_t\}_{t=0}^{\infty}} \left\{ E_0 \left[\sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \lambda x_t^2) \right] \right\} \tag{10}$$

subject to (1), (2), (4), (8), (9), the zero lower bound $i_t \geq i^{ZLB}$ and the initial conditions A_0, R_0, and ε_0 (we disregard the inflation shock as mentioned earlier).

It can be represented in its recursive form by the following Bellman equation (where we have collapsed the state variables in the vector $z = (a_1, a_2, r, \varepsilon)$ and a_1, a_2 are elements of A in equation (5))

7. We use the approximation to Fisher's equation of $1+i = 1+r+\pi$ when deriving the value of ZLB. Also, see the calibration section for the discussion of the apparent contradiction of using positive inflation target while linearizing the model around the zero-inflation steady state.

$$
\begin{aligned}
V(z) &= \min_i \left\{ L(z,i) + \beta E\left[V(z')|\varepsilon \right] \right\}, \text{s.t.} && (11)\\
L(\cdot) &= \pi^2 + \lambda x^2 && (12)\\
\pi &= -\sigma\kappa i + [(\kappa + (\beta + \sigma\kappa)a_1 + \kappa a_2]\varepsilon && (13)\\
x &= -\sigma i + [1 + \sigma a_1 + a_2]\varepsilon && (14)\\
a_1' &= a_1 + \frac{\gamma}{r_1'}\varepsilon[\pi(z,i) - a_1\varepsilon] && (15)\\
a_2' &= a_2 + \frac{\gamma}{r_1'}\varepsilon[x(z,i) - a_2\varepsilon] && (16)\\
r' &= r + \gamma(\varepsilon^2 - r) && (17)\\
\varepsilon' &= \rho\varepsilon + \eta, \quad \eta \sim \, iid\, N(0, \sigma_\varepsilon^2) && (18)
\end{aligned}
$$

where primes denote next period variables, $\pi(z,i)$ and $x(z,i)$ are shorthands for (13) and (14) and ε is the shock to the natural interest rate as in equation (2). Equations (15) – (18) define the transition function $g(z,i)$.

The solution of this problem is the central bank's policy function $h(z)$ minimizing the social welfare loss. The central bank bases its policy on the shock as well as the current state of agents' expectations and it is aware of the effects of its decisions on future expectations. At the same time, due to the backward looking nature of agents' expectations, it can only affect future expectations through the choice of its current instrument, the short-term interest rate, and it does not have access to tools like forward guidance which would affect future expectations directly and which a central bank facing agents with rational expectations could use.

3 Computation and Calibration

In this section we briefly describe the computational procedure and discuss calibration. The interested reader can find details of the computational procedure in the appendix.

3.1 Computational Procedure

The model does not have an analytical solution and, although it is based on a linearized version of the New Keynesian model, the learning recursions and the zero lower bound constraint make the model highly non-linear. We therefore have to resort to non-linear numerical methods. This section contains a brief description of the numerical procedure we use and we defer the presentation of the algorithm and the discussion of the details of individual steps to the appendix.

We approach the central bank's problem with the traditional tools of value function iteration. We discretize the state space, which involves approximating the continuous AR(1) stochastic process for the shock by a Markov chain, and for each grid point we minimize the right-hand-side of the Bellman equation (11) using a numerical optimization procedure (we use the Golden Section Search method, which is somewhat slow but robust). We then iterate the Bellman equation until convergence. To approximate the unknown functions – the value function $V(\cdot)$ and the policy function $h(\cdot)$ – between the grid points we use cubic spline interpolation. With these choices we subject ourselves to the 'curse of dimensionality' problem, however, they are dictated by the presence of the zero lower bound constraint, which causes a kink in the policy function. Using low-order polynomials (e.g. Chebyshev polynomials) would be an attractive alternative, especially when coupled with a sparse grid interpolation scheme like Smolyak interpolation.[8] Unfortunately the low order polynomials are known to handle kinks in the decision rules rather poorly and therefore we decided to use the slower but more flexible splines.[9]

To assess the optimal response of the central bank to an unanticipated disturbance to the natural interest rate once the model is solved and the optimal decision rule $h(\cdot)$ obtained, we do not use the standard tool of a linear impulse response function employed in linear models since non-linear models produce impulse responses which are history- and shock-dependent as discussed in Potter (2000). Instead we follow the approach of Gallant, Rossi, and Tauchen (1993) who develop the notion of a *non-linear impulse response function* (NIR).[10]

8. See e.g. Krueger and Kubler (2004) and Judd et al. (2014)

9. This we experienced ourselves when we experimented with Chebyshev polynomials based Smolyak interpolation in the 2-shock version of the model.

10. See also Potter (2000) and Koop, Pesaran, and Potter (1996), who generalize the concept of a non-linear impulse response and call it a *generalized impulse response* (GIR), or more recently Gourieroux and Jasiak (2005). We, however, use the term GIR in the sense of Gallant, Rossi, and Tauchen (1993), who use it to describe some sort of aggregate over the distribution of the NIRs. In particular they suggest to average the NIRs over initial conditions. We use the term GIR to refer to any aggregator we may use (most often the median response).

The procedure for calculating a non-linear impulse response is as follows. We use a large number of simulations to simulate the economy for a large number of periods and for each of these simulations we store the state variables from the last period.[11] This gives an ergodic set for this economy, which will be used as the set of initial conditions. Then for each element in the ergodic set we run two simulations using the same series of random shocks except in the first period the shock in one of the two simulations differs by an impulse δ. For each initial condition we repeat this step many times using different sequences of shocks and calculate the median across the multiple runs.[12] Following Gallant, Rossi, and Tauchen (1993) the NIR is then the net effect of the impulse and it is defined as the difference of the model variables between the two simulated (median) trajectories. Clearly the impulse responses depend on the state of the economy at the time of the impulse, in addition to the future shocks, reflecting the non-linearities in the system, therefore we report the (cross-sectional) average as well as the median responses and several other quantiles of interest (typically the lower and upper quartiles).

To solve the central bank's problem and to calculate the impulse responses we use Fortran with OpenMP parallelization and we run the codes on a machine with sixteen 2.6 GHz Intel SandyBridge processors and 64 GB of RAM. On this machine, using Intel Fortran compiler with the -O3 flag, it takes about 6-10 hours to solve the CB problem depending on parametrization. For some experiments, typically involving either high β, high ρ or low σ, solving the CB problem can take over 20 hours. Calculating the ergodic set and a long simulation takes about two minutes each and calculating the distribution of impulse responses to one impulse takes about 15 minutes.[13] Overall, solving the CB problem and calculating the complete set of experiments for the benchmark parameterization takes approximately 30 hours. Occasionally, when experimenting with the best settings, we had to use a high memory machine with 256 GB of RAM.

11. By a large number we mean dozens of thousands; the appendix contains more details on the computational procedure including the parameter choices.

12. Gallant, Rossi, and Tauchen (1993) use the average, however, we found that, due to the ZLB, the average provides very distorted view of the dynamics and does not capture well the non-linearities occurring around the zero lower bound.

13. This crucially depends on the parameter R (see appendix).

3.2 Calibration

Before we are able to solve the model numerically we need to set the model parameters. Our chosen values are relatively standard and are summarized in table 1.

Table 1: Calibration

Parameter	Value	Description
β	0.99	subjective discount rate
σ	1	IS slope
α	0.66	Calvo pricing parameter
θ	7.66	price elasticity of demand for intermediate goods
ω	0.47	elasticity of marginal cost
κ	0.057	PC slope
λ	0.007	welfare weight
π^*	0.004693	inflation target
γ	0.03	constant gain
ρ	0.9	shock persistence
σ_η	0.004	std of shock innovations

As the model period we choose one quarter. The value of β is chosen to match the average annual real rate of interest of 4%. The value of σ is common in the recent literature on the zero lower bound and it comes from the real business cycle tradition. Yet a few comments are in order. Adam and Billi (2006), following Rotemberg and Woodford (1998), set $\sigma = 6.25$ using the argument that a higher value of σ may capture some unmodelled interest rate sensitivity of investment demand. We, however, find that with the value of σ so high the ZLB ceases to bind and becomes irrelevant for the conduct of monetary stabilization. Eggertsson and Woodford (2003) on the other hand use $\sigma = 0.5$ in order not to exaggerate the output losses when the ZLB becomes binding if σ were too high. Very low values, however, restrict the ability of the central bank to avoid the ZLB as larger increases in the interest rate are required when facing a negative shock of a given size thus effectively bringing the ZLB closer.

The slope of the Phillips curve, κ, captures the details of the firms' price setting process. We assume Calvo pricing, under which κ becomes[14]

$$\kappa = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\frac{\sigma^{-1}+\omega}{1+\omega\theta}, \tag{19}$$

where α is the proportion of firms who cannot change their prices in the current period, ω is the output elasticity of the firm's marginal cost, and θ is the price elasticity of demand for the firm's goods. The values of α, θ, and ω are taken from Rotemberg and Woodford (1998) and the same values are used in Adam and Billi (2006).

The welfare weight in the central bank's loss function is given by $\lambda = \frac{\kappa}{\theta}$ as discussed in section 2 and the calibrated value follows from our choice of κ and θ. We will, however, explore the sensitivity of the results to alternative weights to reflect on some of the more practical central bank research which has tended to use non-microfounded welfare weights.

The shock persistence parameter, ρ, we take from Fernández-Villaverde et al. (2015) but, compared to them, we set the standard deviation of the innovations, σ_ε, to a higher value. With their value of 0.0025 the ZLB rarely binds in our simulations, which makes it an almost irrelevant policy consideration. We therefore need the stochastic process to exhibit more volatility, which could be achieved by choosing either higher σ_ε or higher ρ. We decided to increase the volatility of the innovations and explore the effect of higher persistence in the robustness section. The chosen value of σ_ε still falls in the range typical in the RBC literature and, perhaps even more importantly, it results in a fraction of periods spent at the ZLB, in our simulations, that we consider reasonable. There is no clear consensus in the literature on how often the ZLB should bind, nevertheless, our value of 5.9% of periods spent at the ZLB is comparable to a number of recent papers, e.g. Fernández-Villaverde et al. (2015) and Gavin et al. (2015), who report the ZLB binding in over 5% of the periods.

The π^* parameter affects the allowable maximum negative deviation of central bank interest rates from the steady state. It is chosen to match an inflation target of 2 percent annually. We choose a positive inflation target despite the fact that the model equations (1) and (2) have been linearized around the zero-inflation steady state. Linearizing the model around a steady state exhibiting positive trend inflation would require taking account of the effect of price dispersion as Ascari and Sbordone (2014) demonstrate. This would alter the

14. See Woodford (2003), chapter 3 for details.

Phillips curve equation (eq. 1) and introduce an additional equation describing the law of motion for price dispersion.

We would argue that steady state inflation of 2% is still sufficiently small to be treated as an acceptable approximation. There are two advantages of this approach: it allows a comparison of our results to the existing literature and it increases the tractability of the problem. From the point of view of our numerical solution procedure using the properly microfounded Phillips curve as in Ascari and Sbordone (2014) would require adding two more state variables to our model – a measure of price dispersion and the expectations coefficient associated with it. This would increase the number of state variables from the current four to six, imposing a significant penalty in terms of the number of grid points required to solve the policy maker's problem.[15] Although still tractable, the curse of dimensionality problem would severly impact the performance of our algorithm. We leave the analysis of optimal policy with trend inflation for future research.

The gain parameter, γ, is arguably the key parameter in our model. There is some uncertainty over the estimates for this key parameter in the literature, see e.g. Orphanides and Williams (2008) for a brief discussion. Our value is slightly higher than theirs but it still falls in the range of typical estimates. Values much higher than this tend to lead to the learning recursions exploding frequently and values much below this tend to eliminate fluctuations in expectations formation. We nevertheless explore the sensitivity of our results to the choice of alternative values of this parameter below.

4 Results

In this section we present the results of our calculations. We start with a discussion of the long-run properties of optimal policy and then we focus our attention on the optimal response to shocks.

4.1 Long-run properties of optimal policy

We proceed in two steps. More as a test of our computer code, we present the results for an economy that has never been exposed to

15. If we were similarly conservative in the choice of the number of grid points in each dimension as we were for the present version of the model, then the number of grid points would increase by two orders of magnitude to the total of about 15 million points

the zero lower bound, then we discuss the long-run effects of the zero lower bound and contrast them with the previous case.

4.1.1 Without ZLB

As is well known, in the model with rational expectations without the zero lower bound the demand shock does not introduce a trade-off between inflation and output.[16] Therefore the central bank can minimize both gaps at the same time by following the changes in the natural interest rate one-to-one and thus offsetting the shock perfectly.

In a model with learning, however, expectations are not anchored by the rational expectations equilibrium, therefore, in addition to the demand shock, the central bank needs to respond to changes in agents' expectations. This can make the central bank's task of stabilizing the economy more difficult and can potentially make the one-to-one response known from the rational expectations world sub-optimal. Nevertheless, in a model with learning, similar to ours, Molnár and Santoro (2014) observe that in the case of a demand shock the central bank is able to stabilize the economy perfectly under all circumstances.[17] Similarly to them, we find that in our model the central bank can stabilize the economy in the long run but, in contrast, we find that if expectations deviate from the long-run equilibrium then the central bank's optimal policy deviates from the one-to-one response.[18]

The long-run outcome is demonstrated in figure 1, which shows the probability density functions estimated from the simulated ergodic set (see appendix A.4 for details).[19] The densities reveal that, without the zero lower bound, also in our model the central bank can stabilize the economy very well in the long run. Since the agents expect the inflation and output gaps to be zero, which can be seen by the expectation coefficients densities being degenerate,[20] the central bank

16. See e.g. Clarida, Gali, and Gertler (1999).

17. Technically, in their model the demand shock does not affect the first-order conditions.

18. We call *long-run equilibrium* a set of allocations with $x = \pi = a_1 = a_2 = 0$. Such allocations are supported by $i = \frac{\varepsilon}{\sigma}$

19. The densities were estimated by a Kernel Density Estimator with Gaussian kernel using the implementation from the python SciPy library (see https://docs.scipy.org/doc/).

20. In fact, due to computer arithmetics, the coefficients are not exactly zero but of the order of $\pm 1e^{-10}$ and lower.

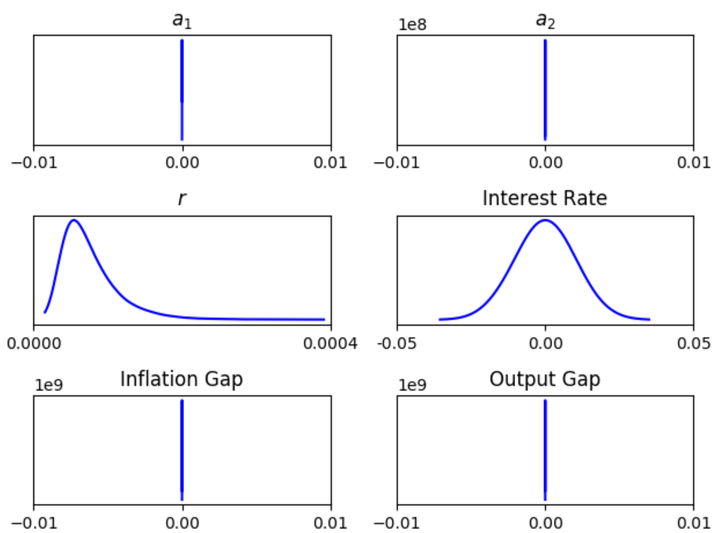

Figure 1: Ergodic distributions for the nonZLB economy

The densities were calculated from a sample of 25000 observations using Gausian kernel density estimator.

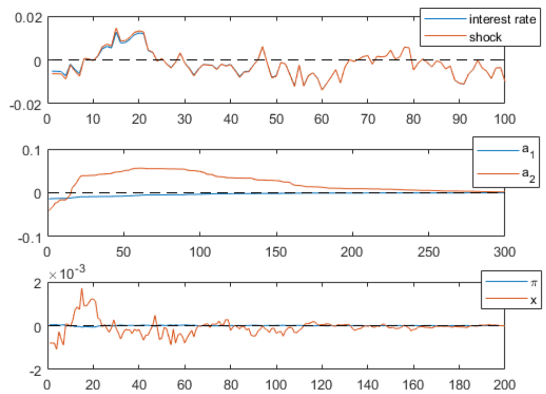

Figure 2: Simulated convergence from an arbitrary state towards the long-run equilibrium in an unconstrained economy. The initial condition is the median state from the ergodic set of the constrained economy, see table 2 for the exact values.

responds one-to-one to the demand shock[21] and this allows it to keep both gaps at zero. And since the CB keeps the gaps closed the agents learn that monetary policy does stabilize the economy and, therefore, expect both gaps to be zero. The expectations coefficients are thus zero too. As a result, once expectations converge to zero the interest rate follows the pattern of the stochastic process of the demand shock, in particular, the estimated densities of the interest rate and the demand shock coincide (the latter not shown).

If the expectations deviate from the long-run equilibrium, however, the one-to-one response is no longer optimal. Figure 2 sheds some light on the role of expectations for optimal policy. It shows a sample simulation starting from the median state of the constrained economy, which is not in the ergodic set of the unconstrained one, and illustrates convergence towards the long-run equilibrium.

Since the expectations coefficients are initially negative, we would expect the central bank to change the nominal interest rate less than

21. More precisely, due to the way we calibrate the model, the central bank changes the interest rate by a $\frac{1}{\sigma}$ multiple of the change in the shock. For the benchmark calibration this means a one-to-one response.

one-to-one with the shock (with positive coefficients we would expect a more-than-proportional response) and over time, as expectations adjust, we would expect the optimal policy to converge to the one-to-one type of policy prevalent in the long run. This is exactly what we see in the top panel of figure 2. With negative expectations coefficients, a negative demand shock creates expectations of positive inflation and output gaps, hence, to neutralize the effect of the shock the central bank does not have to lower the interest rate as much as would be necessary if the agents were expecting the central bank to stabilize the economy fully.

However, the bottom panel of figure 2 shows that optimal policy does not actually offset the shock completely, allowing inflation and output to deviate from their steady-state values. In fact the plot suggests that the central bank may be facing a trade-off between inflation and output. To verify these observations, in figure 3 we present the impulse responses to a large shock. The plot shows that if the shock occurs outside the long-run equilibrium the central bank does not respond to the shock one-to-one, moreover, it does not fully stabilize the economy due to a trade-off between inflation and output, which is apparent in panels 2 and 3 even though the effect seems to be rather small. The existence of such a trade-off appears to be robust to both the size and the sign of the shock as well as a wide range of initial conditions.[22]

This observation is new in the literature, as far as we are aware. Adam and Billi (2007) observe a policy trade-off in the case of a constrained central bank which has not reached the lower bound yet. In their model, this is explained by the combination of the lower bound distortions and the presence of a mark-up shock. In our model, however, there is no mark-up shock and this trade-off occurs in the case of an unconstrained central bank rather than the constrained one. The trade-off is caused by the departure from the assumption of rational expectations and manifests itself when the economy has deviated from the long-run equilibrium. Interestingly, we do not observe the occurence of this trade-off in the case of a constrained central bank which has not reached the lower bound yet (see section 4.2.1).

We now turn to the discussion of the effects of the zero lower bound.

22. We tried shocks of various sizes ranging from $0.1\sigma_\varepsilon$ to $4\sigma_\varepsilon$ and the trade-off is apparent across the whole ergodic set of the constrained economy (see also Figures 10 and 7).

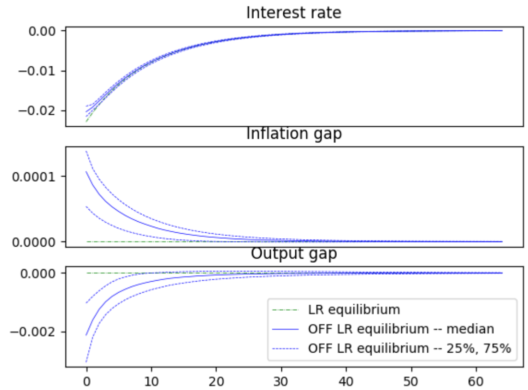

Figure 3: Generalized impulse response without the zero lower bound.
The impulse has the size of $-2.5\sigma_\varepsilon$ and for the off LR equilibrium paths the initial condition is the ergodic set of the contrained economy.

4.1.2 With ZLB

When the central bank is constrained by the zero lower bound it can no longer stabilize the economy in the long run. This can be seen by the fact that the estimated densities shown in figure 4 are no longer degenerate and that there is large dispersion of both the gaps and the expectations coefficients.

The failure of the central bank to stabilize the economy is no surprise, of course. In times when the zero lower bound becomes binding the central bank cannot prevent negative inflation and output gaps, which, in turn, leads to agents expecting the central bank to fail to stabilize in the future too. Consequently, the expectations coefficients start increasing and eventually, if the liquidity trap situation persists, become positive. As discussed in the previous section, when expectations deviate from the long-run equilibrium, an optimally behaving central bank opts for imperfect stabilization even in times when the ZLB is not binding. Therefore, the change in expectations induced by the binding ZLB magnifies its destabilizing effect. Since the ZLB becomes binding relatively often, this causes substantial deviations from the long-run equilibrium.

Figure 5 gives an illustration of the dynamics, the destabilizing effect of prolonged periods spent at the ZLB is especially apparent

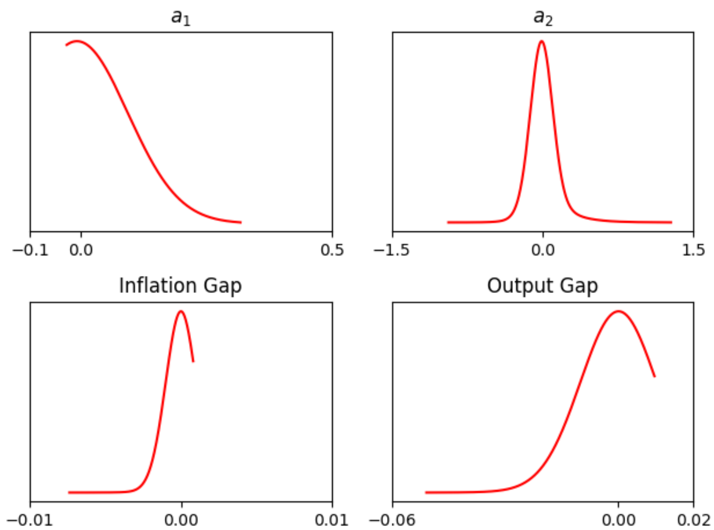

Figure 4: Ergodic distributions for the ZLB economy
The interest rate is shown in Figure 6 and r is independent of the ZLB regime (thus its density function is identical to the one in Figure 1).

as is the increased volatility of the interest rate that follows – to alleviate the welfare consequences of the deterioration of expectations the central bank's policy becomes a lot more aggressive. Periods of liquidity trap as long as those in Figure 5 are very rare events though (see table 2).

An interesting consequence of optimal policy is that the expectation coefficients are often negative. This can be seen both from the densities and from the sample simulation. In fact, for inflation the coefficient is negative most of the time (compare the mean and median values in table 2). This is a somewhat unintuitive result, for negative expectation coefficients mean that after a negative shock agents expect a boom and above-average inflation. Typically, we have negative demand shocks associated with deflation and recessions but here the agents tend to be optimistic about the future instead. The long-run distribution of interest rates shown in Figure 6 helps to shed some light on this. As can be seen from the figure, in the presence of the zero lower bound, the central bank tends to follow a more aggressive policy. The closer to the zero lower bound, the lower the interest rate is relatively to the unconstrained benchmark. Interestingly, on the

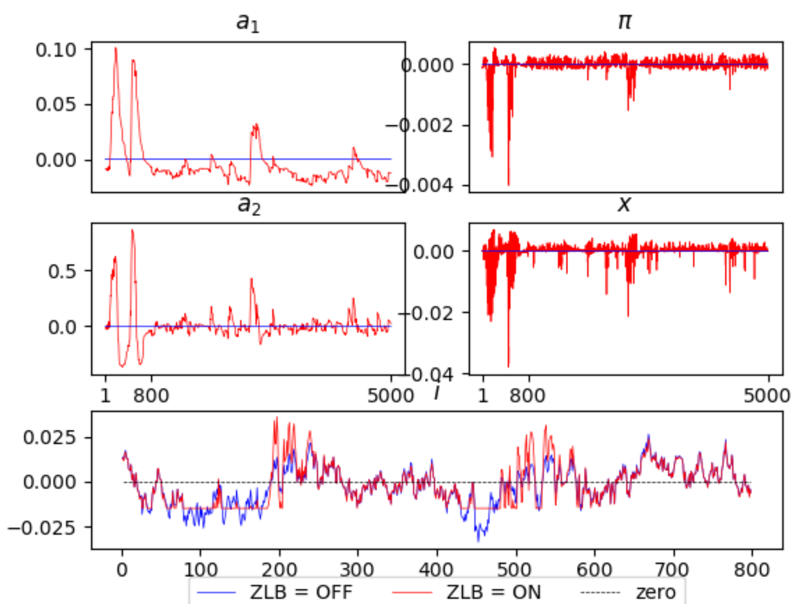

Figure 5: Sample simulation

The bottom panel shows a zoom of the simulated interest rate, the blue line is the optimal policy from the unconstrained economy and it coincides with the shock as discussed in section 4.1.1. The shocks are identical in simulations with and without the zero lower bound. The two prolonged periods of binding ZLB are very rare events (see table 2).

other side of the spectrum we observe a similar pattern even though we might perhaps expect a convergence to the unconstrained policy since the ZLB does not constrain the policy maker when setting high interest rates. Nevertheless, we observe that when interest rates are above the steady-state value, they tend to be higher in the constrained economy compared to the unconstrained one, although the difference seems to be rather small.

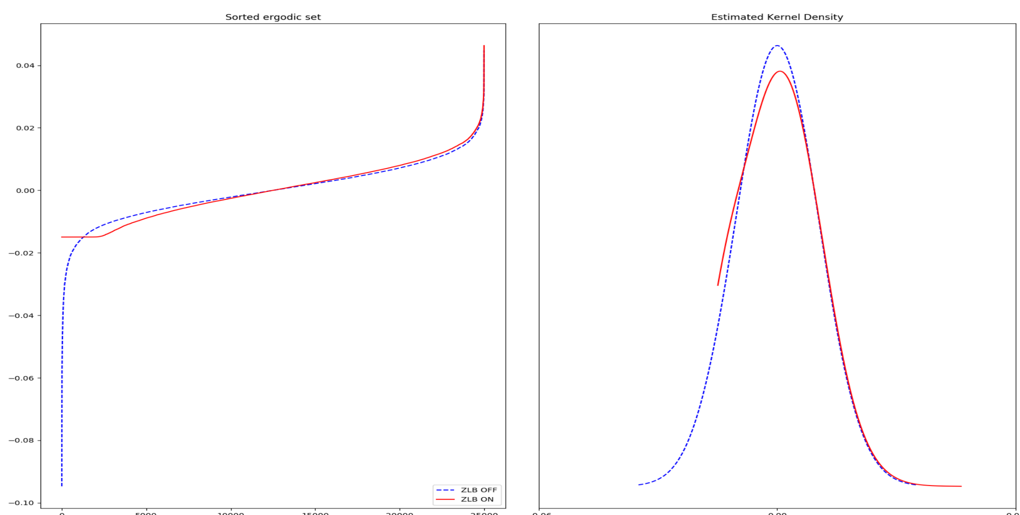

Figure 6: Ergodic distribution of interest rates
Left panel shows the ergodic set of interest rates sorted from lowest to highest. Right panel shows the estimated density functions.

We think that there are two reasons for this. First, when approaching the ZLB the central bank aims to overstimulate the economy to create inflation expectations. These increased expectations then alleviate the liquidity trap should the central bank hit the ZLB soon. This way the central bank trades off smaller welfare losses of future liquidity traps for larger losses of imperfect stabilization today. Due to the convexity of the central bank's loss function, however, such a trade-off is welfare improving. Since the closer to the ZLB, the higher the risk of the bound becoming binding, the urge to trade off the future welfare losses for the current ones becomes stronger and, hence, the deviation of optimal policy from the unconstrained benchmark is larger. A similar argument may apply when setting the interest rate above its steady-state value. This would typically be associated with a positive shock, thus encouraging counter-cyclical expectations in this situation requires more restrictive policy.

Second, upon occurence of a liquidity trap the expectations coefficients may become positive, as the agents learn to expect negative inflation and output gaps in times of depressed demand, and become larger with the duration of the liquidity trap episode. To contain this undesirable rise in expectations the central bank has to engage in more aggressive policy. Overall, the over-aggressive policy then creates the somewhat unintuitive counter-cyclical expectations. In the next section we turn to the impulse response functions to explore these mechanisms further.

Table 2: Long simulation statistics

	mean		median	std	
	non-ZLB	ZLB	ZLB	non-ZLB	ZLB
a_1	0.00	-0.006080	-0.010149	0.00	0.016600
a_2	0.00	0.004590	-0.006720	0.00	0.115000
r	0.000085	0.000085	0.000075	0.000043	0.000043
ε	-0.000128	-0.000128	-0.000183	0.009210	0.009210
i	-0.000128	-0.000182	-0.000193	0.009210	0.009300
π	0.00	-0.000011	-0.000004	0.00	0.000231
x	0.00	-0.000060	0.000078	0.00	0.002350
π^e	0.00	-0.000008	-0.000002	0.00	0.000170
x^e	0.00	-0.000105	0.000012	0.00	0.001280
i^{ZLB}		-0.01495			
ZLB %	5.29	5.88			
ZLB 1q %		60.9			
ZLB $<$ 5q %		91.9			
ZLB $<$ 9q %		97.6			
ZLB $>$ 15q %		0.46			

Note: Simulation 500000 periods. Values rounded to six decimal places. π^e and x^e are the expected inflation and output gaps as in eq (5). ZLB % shows the percentage of periods in which the ZLB is binding and the last four rows show the percentage of liquidity trap periods of certain duration (in quarters).

4.2 Optimal response to shocks

We now turn our attention to the main purpose of the paper, the question of how the central bank responds to shocks that (threaten to) take it to the zero lower bound. In a model with rational expectations, Eggertsson and Woodford (2003) find that after reaching the zero lower bound it is optimal for the central bank to keep the interest rate at the bound even after the natural interest rate ceases to be negative. This is then accompanied by a corresponding overshoot of the inflation target. They also find that the central bank should respond disproportionately more to negative shocks of a larger size and, as a corollary, it should lower the interest rate to the ZLB before the natural interest rate becomes negative. In this section, we examine these predictions by analysing impulse responses to such shocks. If these predictions carried over to our model then we would observe the constrained central bank pursuing a more expansionary policy, unless obstructed by the ZLB, and thus its impulse response lying below the one of the unconstrained central bank.

Since in non-linear models an impulse response depends on both the future realizations of the shocks and the initial condition, as discussed in section 3.1, we calculate a large number of responses, each for a different initial condition, and report various percentiles from the resulting distribution. The initial condition for each response is taken from the ergodic set of the constrained economy. Where we compare responses for constrained and unconstrained central banks the two sets of impulse responses use the same initial conditions and the same realizations of the stochastic process. We can think about this intuitively as comparing responses of two central banks which have both been subject to the zero lower bound constraint but at the beginning of period zero one of them unexpectedly discovered a way to set negative interest rates. We refer to such a central bank as unconstrained even though it has been previously constrained by the ZLB. Occasionally we also make a comparison to a central bank that has never been exposed to the zero lower bound constraint and, therefore, the agents expect it to stabilize the economy perfectly at all times. We refer to such a central bank as being in a long-run equilibrium (as defined in footnote 18).

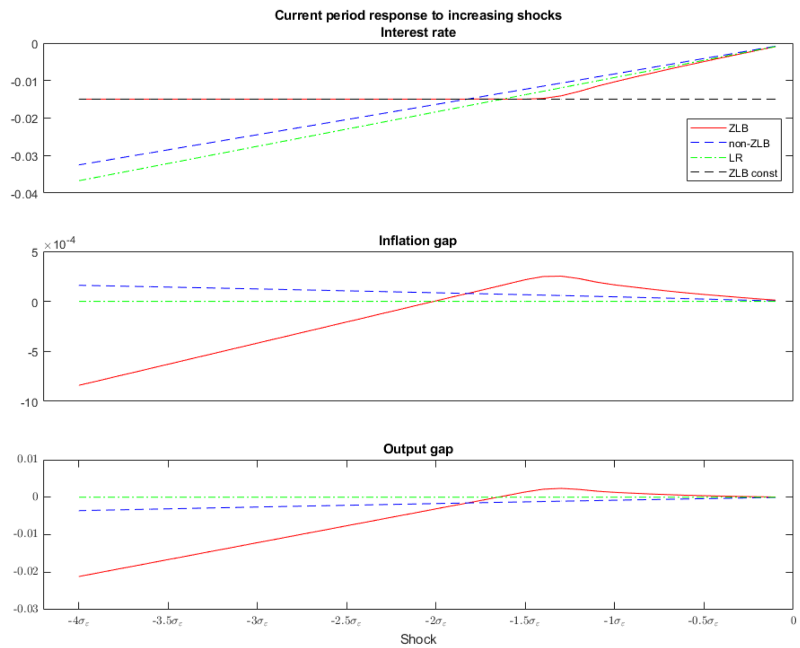

Figure 7: A current period response to shocks of varying sizes.

The initial condition is the median state from the ergodic set of the constrained economy.

4.2.1 Response to shocks in normal times

We start with a single, representative initial condition, which represents a situation when the central bank is far from the effective lower bound and it has not been obstructed by the bound for a long time. We call this normal times. A current period response as a function of the shock, when the other variables are set to the median values they attain in the constrained economy, is presented in Figure 7.[23] The most notable feature of the chart is that, due to the presence of the zero lower bound, optimal policy is strongly distorted towards lower interest rates in response to shocks that threaten to take the economy to the effective lower bound. Also these distortions are larger for larger shocks. As a corollary, the central bank reaches the zero lower bound sooner than its unconstrained counterpart.

The closer the shock brings the central bank to the lower bound the larger the risk that the bound will be reached in the near future, which would be accompanied by substantial welfare losses. The central bank attempts to mitigate these losses in anticipation of the liquidity trap occuring by overstimulating the economy in the vicinity of the lower bound in order to create (or reinforce) inflationary expectations. As Krugman (1998) points out, and what has become a stylized argument, the real interest rate can be reduced by creating expectations of future inflation, even when the nominal interest rate cannot be lowered any further. This has become the crux of the, now popular, advice of Eggertsson and Woodford (2003) to create inflation expectations by committing to overly expansionary policy once the liquidity trap has ended. In an economy with learning, however, influencing inflation expectations once the liquidity trap has occurred is not possible and, thus, this advice fails.

The central bank, therefore, has to create such expectations in advance of the liquidity trap by reacting to negative shocks with excessive monetary stimulus, which an unconstrained central bank would not find desirable.[24] This overly loose monetary policy creates welfare losses in the current period, which are traded off for lower losses in times of a liquidity trap, should one occur.

23. Using a single initial condition allows us to make a direct comparison to the zero lower bound. Using the 25-th or 75-th percentiles, as opposed to the median, does not change the conclusions.

24. Note that normal times are characterized by counter-cyclical expectations and thus negative expectations coefficients, a_1 and a_2.

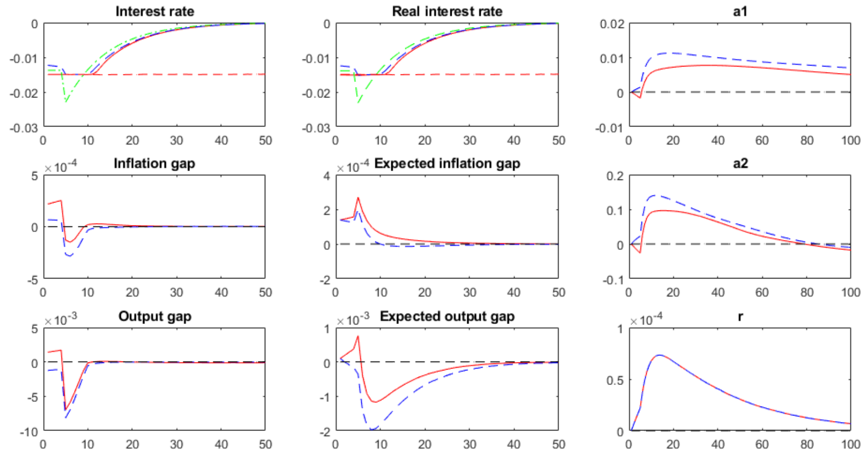

Figure 8: Comparing optimal and naive policy.

Red lines indicate the optimal constrained policy, blue a naive policy, and green the long-run (one-to-one) policy. The naive policy is the unconstrained policy truncated to the lower bound. The initial condition is the median from the ergodic set of the constrained economy.

To investigate this channel further we run the following experiment. The economy is hit by a medium negative shock, $1.5\sigma_\varepsilon$ in size, which does not cause a liquidity trap but is large enough to make a constrained central bank lower the interest rate to the effective lower bound. This shock is kept constant for four periods and then it is replaced by a larger negative shock, $2.5\sigma_\varepsilon$ in size, which then decays according to the AR(1) specification as usual. This latter shock does cause a liquidity trap. Figure 8 shows the associated dynamics and compares the optimal policy to a naive policy which is identical to the unconstrained optimal policy whenever it is above the effective lower bound. When it reaches the bound, though, it becomes stuck.

The impulse responses show that, by being more aggressive in stimulating the economy in a run-up to the lower bound, the optimally behaving central bank achieves lower losses during the liquidity trap period. The crucial difference between these two policies is manifested in the behaviour of expectations. Since the naive central bank, being unaware of the impossibility of negative nominal rates, aims to enforce the convergence of expectations to the long-run equilibrium, it attempts to lower expectations regardless of the vicinity of the effective lower bound. This is, however, a bad policy, resulting in larger welfare losses in the case of the interest rate getting stuck at the bound.

In contrast, the optimally behaving central bank stimulates expectations in anticipation of the effective lower bound becoming binding in the near future. This can be seen both on the expected gaps and the forecasting functions plots. Both banks then hit the lower bound, in this experiment, which forces expectations to decline for the duration of the trap. Since the optimally behaving central bank was better prepared for the episode and because it is more expansionary during the recovery, expectations are faster to converge to pre-shock levels.

It is also instructive to see how the naive central bank which is more restrictive than a central bank in a long-run equilibrium before the liquidity trap occurs becomes more expansionary after the trap disappears.[25] This is solely due to the distortion of expectations the unexpected presence of the zero lower bound created.

Potentially, such a naive policy could, in conjuction with the effective lower bound, allow a deflationary trap to develop. In Evans, Guse, and Honkapohja (2008) it is emphasized that monetary policy, if not aggressive enough, can, indeed, take the economy to a deflationary trap, if one exists. The authors demonstrate, by the means of an example, that such a deflationary trap may be avoided if monetary policy is sufficiently aggressive (although this may not suffice in all cases). They design a simple threshold rule, which they call an *aggressive monetary policy rule* (AMPR) by which the central bank, otherwise using a standard Taylor rule, lowers the interest rate discountinuously to the effective lower bound whenever inflation drops below a certain level. Whether such a deflationary trap exists in our model is an important question. Nevertheless, answering this question fully goes beyond the scope of this paper.[26]

We would, however, like to emphasize the similarity of our optimal policy to the AMPR. Both policies incorporate the idea that

25. Note that, outside the ZLB, the naive policy is identical to the optimal unconstrained policy but the episode of binding zero lower bound changes the state of expectations making the policy more expansionary (compare to Figure 9).

26. All we can say here is that we did not observe a deflationary trap developing anywhere in our simulations. See footnote 34 for the description of some of the experiments that were specifically targetted at discovering a deflationary trap.

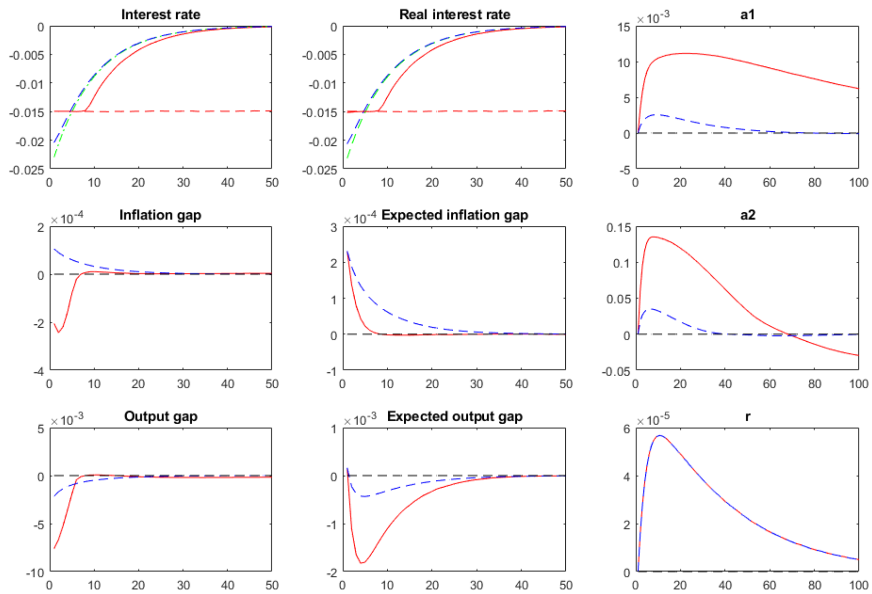

Figure 9: Impulse responses to a large shock ($-2.5\sigma_\varepsilon$)
Red lines are for the ZLB policy, blue without the ZLB, and green the long-run equilibrium. The initial condition is the median from the ergodic set of the constrained economy.

inflation expectations have to be stimulated by aggressively lowering the interest rate before a shock takes the economy to the effective lower bound where deflationary pressures start to develop. The optimal policy does it in a smoother manner, in order to spread the cost of this pre-emptive monetary easing over a longer time horizon, and without the need for discontinuous jumps in the policy instrument (clearly, such jumps are not optimal). We would argue that the requirement to prevent a deflationary trap is embedded in the notion of optimality.

We make one more observation about Figure 7. The precautionary behaviour in the vicinity of the zero lower bound is reminiscent of the behaviour of optimal policy in models with rational expectations. We notice one difference, though. In the bottom two panels we, yet again, notice the inflation-output trade-off of the unconstrained central bank, which has been discussed earlier, but, perhaps somewhat unexpectedly, we do not see the same trade-off arising in the case of a constrained central bank that has not yet reached the zero lower bound. In our model optimal policy is so aggressive that the central bank accepts both gaps increasing above their pre-shock values. We do not see this in models with rational expectations.[27]

When the shock is large enough to cause a liquidity trap, see Figure 9, the unconstrained central bank sets a negative nominal interest rate but it is not able to offset the shock fully, which is due to the existence of the policy trade-off. The failure to offset the shock then causes a necessary adjustment in agents' forecasting functions. The policy is also little more restrictive than the long-run one because agents have counter-cyclical expectations initially. The constrained central bank, on the other hand, becomes stuck at the effective lower bound and witnesses a sharp reduction in both inflation and output, which, in turn, results in substantial deterioration of expectations, further reinforcing the need for monetary stimulus. Such a stimulus can only come once the liquidity trap has disappeared and we can see that, then, the constrained policy becomes more expansionary than the unconstrained one.

This pattern is consistent with the prediction of Eggertsson and Woodford (2003), and we shall investigate it in more detail, however, it is crucially reliant on the impulse responses having been constructed from the conditional median profiles as opposed to the conditional mean ones (see footnote 12 and appendix A.4). Since the dynamics depends on the realizations of the future shocks the way one treats these shocks may affect the resulting shape of the impulse response. Figure 19 in appendix B shows the alternative impulse responses when conditional mean profiles are used. We would argue that using the mean profiles does not adequately capture the non-linearities of the zero lower bound. It can even change the interpretation of the results dramatically, in fact, the constrained central bank would appear to be more restrictive than its unconstrained counterpart while we are going to argue the opposite is the case.

4.2.2 Generalized impulse response

So far we have only considered a single initial condition, albeit one that occurs relatively frequently in simulations. To investigate to what extent our conclusions to this point are the reflections of the general

27. See Eggertsson and Woodford (2003) and Adam and Billi (2007).

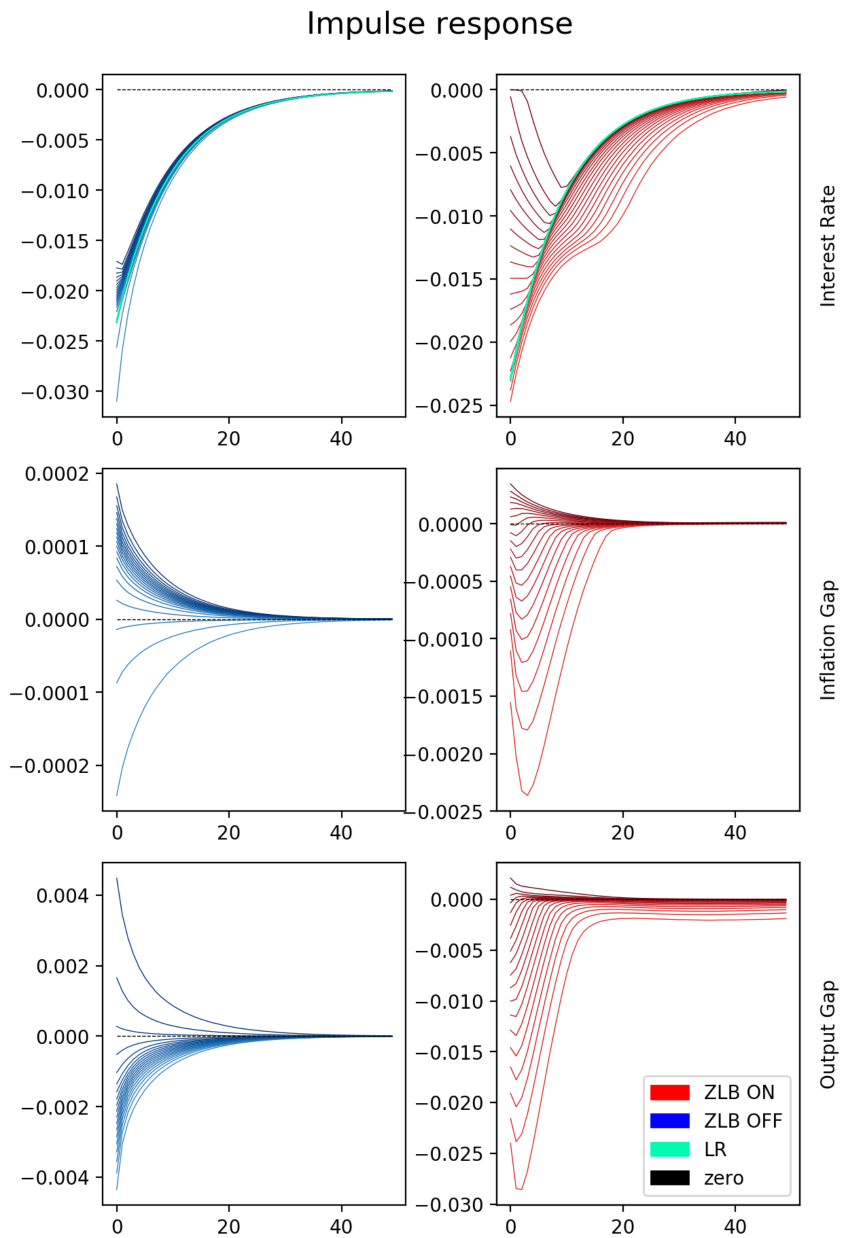

Figure 10: The distribution of impulse responses to a large shock $(-2.5\sigma_\varepsilon)$
The figure displays every fifth percentile from 5^{th} to 95^{th}, the initial condition for all responses is the ergodic set of the constrained economy (except LR, for which expectations are zero). The horizontal axis shows time.

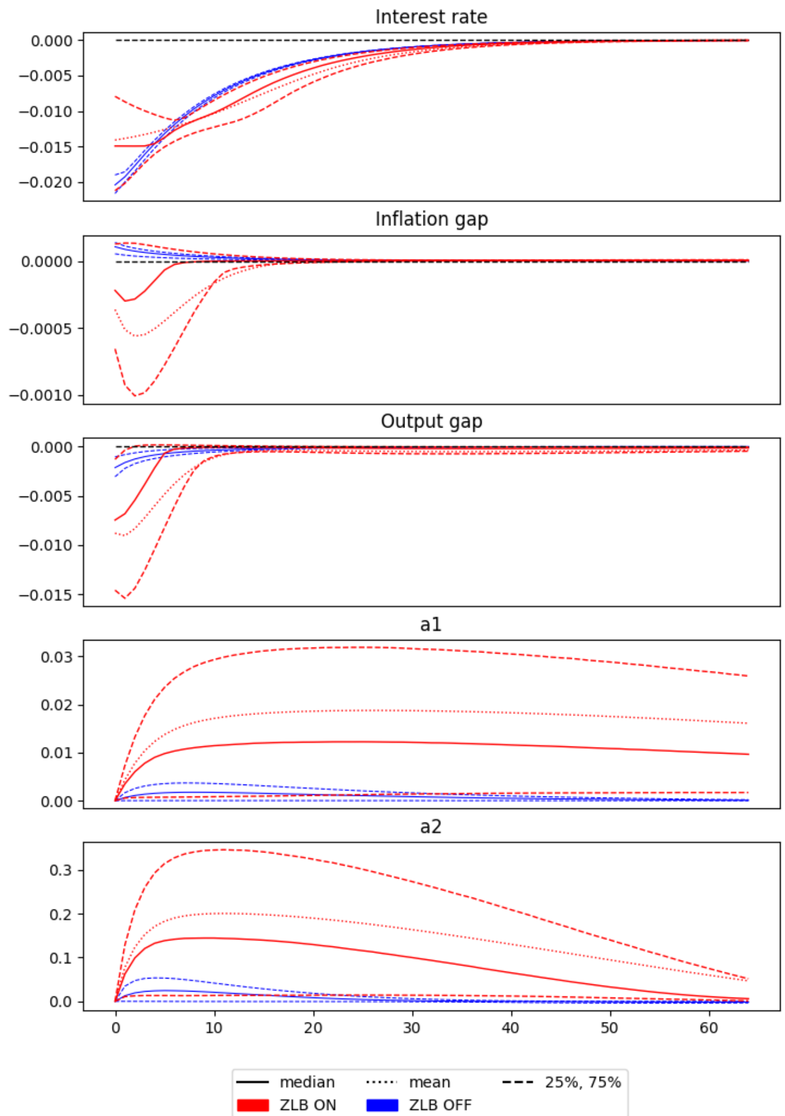

Figure 11: Impulse responses to a large shock ($-2.5\sigma_\varepsilon$)
red lines are for the ZLB policy and blue without the ZLB, the initial condition for all responses is the ergodic set of the constrained economy; for the unconstrained economy only the median response is shown (see figure 3 for comparison).

behaviour of the model, we will now explore the whole distribution of the impulse responses across all initial conditions. We consider a negative shock of size 2.5 standard deviations of the AR(1) shock, which is large enough to bring the economy to the zero lower bound for most initial conditions. A view of the full distribution of impulse responses to this shock is shown in Figure 10 but we will focus our attention on a selected set of responses from this distribution as shown in Figure 11.

Each impulse response represents a deviation of the respective variable from a counter-factual scenario that would prevail if there had been no shock in period zero and it is indexed by the probability that this or lower deviation is observed. Since these are not deviations from a steady state, or from any given initial condition for that matter, it is not possible to represent the ZLB constraint on the same graph.[28] In the context of a non-linear impulse response, the constraint needs to be interpreted as a maximum allowable reduction of the interest rate in a given state of the economy. The state of the economy, however, is not shown on the graph, in fact, it is not possible to associate any of the impulse responses with a specific state since multiple initial states of the world may have resulted in a response of the same size. The best we can do is to calculate the probability that a given impulse response is associated with the central bank hitting the ZLB (such probabilities are shown in Figure 12). We can, therefore, only make an indirect inference about the role of the lower bound in these responses.

So what do the impulse responses tell us? The unconstrained impulse responses show very little dispersion and thus the previous conclusions seem to have general validity. We still see the policy trade-off and, due to this trade-off, we see the forecasting functions going through a period of mild adjustment. On the other hand, the dynamics expressed by the constrained impulse responses can vary a lot. In some states of the world the central bank is close to the effective lower bound before the shock arrives and in some states it is so far from it that it barely hits the bound, or not at all. The initial distance to the lower bound determines the duration and the severity of the resulting liquidity trap. The longer and more severe the trap is the larger the losses are and the more expectations are disturbed. And, as we have seen earlier, expectations have an important influence

28. In the previous section we could do that because we controlled the initial condition.

on the resulting dynamics.

The solid curve represents the median response and it shows a typical picture in a liquidity trap episode. The interest rate declines as much as it can and the inflation and output gaps as well as the expectations all deteriorate as we saw earlier. For comparison, we also show the mean dynamics, represented by the dotted curve. As we would expect, the top panel suggests that the mean does not capture the non-linearities caused by the zero lower bound very well, therefore, we do not pay much attention to the mean in the subsequent discussion.[29]

The two dashed curves in each panel represent the lower and upper quartile responses. Unfortunately, it is not obvious which of the two curves in any given panel is associated with which of the dashed curves in the other panels. Our argument about which curves in the different panels are associated with each other comes from our indirect inference about the zero lower bound. The bottom dashed curves in panels two and three and the top dashed curves in panels four and five clearly show symptoms of an economy in a liquidity trap – substantial decline in output and inflation and large deterioration of expectations – and, hence, it must be associated with the top dashed curve from panel one.

It may not be easy to see that the top dashed curve in panel one actually shows that the central bank reached the zero lower bound. In fact, it does not look like the interest rate hit the lower bound at all, instead, it makes the impression that the central bank is rather cautious lowering the interest rate initially and that monetary policy is only loosened slowly over time. This, however, would not be the right way of interpreting the response. It is more plausible that the central bank started in an unfavourable initial condition and that it was initially much closer to the ZLB so that it hit it quickly. Since such poor initial conditions are relatively uncommon (the probability that the CB would start farther away from the ZLB is 75%) the counterfactual is likely to involve increasing interest rates, which makes the response appear more restrictive initially and more relaxed later while, in fact, the interest rate is at the ZLB all the time.[30]

29. This is related to the issue of averaging over future shocks discussed earlier.

30. As explained previously, the counterfactual depends on the realizations of future shocks. To eliminate this dependence we aggregate over the future shocks by calculating the median profile and the impulse response is then calculated as the difference between two such profiles, with and without the initial impulse. Therefore, by construction, even

To shed some light on this behaviour, Figure 13 presents the after-shock interest rates, from the impulse response simulations.[31] It reveals that, indeed, 85% of the simulated interest rates reach the zero lower bound after the shock. Therefore, we infer that responses smaller in absolute value are associated with reaching the zero lower bound, larger drops in output and inflation, and larger deterioration of expectations (i.e. larger increases in the expectations coefficients).

Coming back to Figure 11, the bottom dashed line in the top panel, then, shows the response of a central bank that was lucky enough to experience the shock when it was far away from the ZLB so that it had enough room to lower the interest rate and it only barely hit the lower bound. This curve is associated with the top dashed curve in the second and third panels and with the bottom dashed curve in the last two panels. It would be tempting to say that this is because the central bank is, in this case, virtually unconstrained and thus the responses are expected to be similar to the unconstrained ones – we even observe the same trade-off between inflation and output after all.

However, such an interpretation would be misleading, the central bank is still potentially constrained and it has to take into account the risk of the lower bound becoming binding in the future. Similar logic applies here as it did in section 4.2.1. Since there are strong non-linearities close to the zero lower bound, which make the central bank more expansionary in an attempt to avoid hitting the bound in the future, the constrained and unconstrained policies may still differ substantially even when the former is not immediately constrained. In fact, the emergence of the said trade-off in panel two and three is a sign that the central bank actually did hit the lower bound as otherwise it would have been encouraged to stimulate both output and inflation at the same time to alleviate the cost of a liquidity trap potentially occuring in the future.

To summarize, in response to a large shock the central bank lowers the interest rate to the zero lower bound, in most cases, which lowers the inflation and output gaps and, consequently, the expectations coefficients start to increase as the agents learn that the central bank is not able to carry out sufficiently large monetary expansion to offset the shock. Where the interest rate does not hit the bound,

---

if the interest rate was at the lower bound at the time of the shock, the impulse response would appear to be decreasing over time while the interest rate would be at the ZLB for many periods (this is the case of the top curve in Figure 10).

31. The figure was constructed from a sample of 50 million interest rates in each period.

or where it reaches it pre-emptively, we may or may not observe the same trade-off between inflation and output as we observe in the case of an unconstrained central bank.

4.2.3 The Lower for Longer Property

The discussion in the previous sections suggests that, perhaps, the first prediction of Eggertsson and Woodford (2003), the lower for longer property of optimal policy after a period of liquidity trap, holds in our model too. Probably the most convincing so far is the comparison of the distributions of constrained and unconstrained impulse responses in Figure 10, where we also show the impulse response of a central bank in a long-run equilibrium. The comparison to the long-run policy provides a convenient separation between the unconstrained and constrained policies. While the former is more restrictive most of the time, the latter is more expansionary except in the cases where its policy has clearly become obstructed by the zero lower bound.

Due to the impossibility to obtain a direct comparison to the effective lower bound, though, such a conclusion might be little premature. To gather more evidence, we, therefore, calculate the probability that the interest rate reaches the zero lower bound for each impulse response. We construct a number of bins, each 5 percentiles wide and centered at every 5$^{\text{th}}$ percentile (we disregard the bottom and top 2.5 percent). Then, for all impulse responses which fall in a given bin we calculate the proportion of interest rates which hit the lower bound in the underlying simulations.[32]

These probabilities are reported in Figure 13 and they confirm that the constrained central bank stays at the lower bound more often than its unconstrained counterpart. There is a small number of cases where this is not the case but those are driven by the fact that due to computer arithmetic and numerical/approximation errors the interest rate will never be exactly equal to the lower bound. When we consider all interest rates which agree with the ZLB constraint on four decimal places as having hit the lower bound then the probabilities for the constrained central bank are, in all cases, larger than for the unconstrained one.

32. Given we use 50 thousand initial conditions and 2 thousand simulations per initial condition, the total number of simulations for all experiments involving the whole distribution of impulse responses is 50 million. This gives a sample of 2.5 million interest rates in each bin (per period).

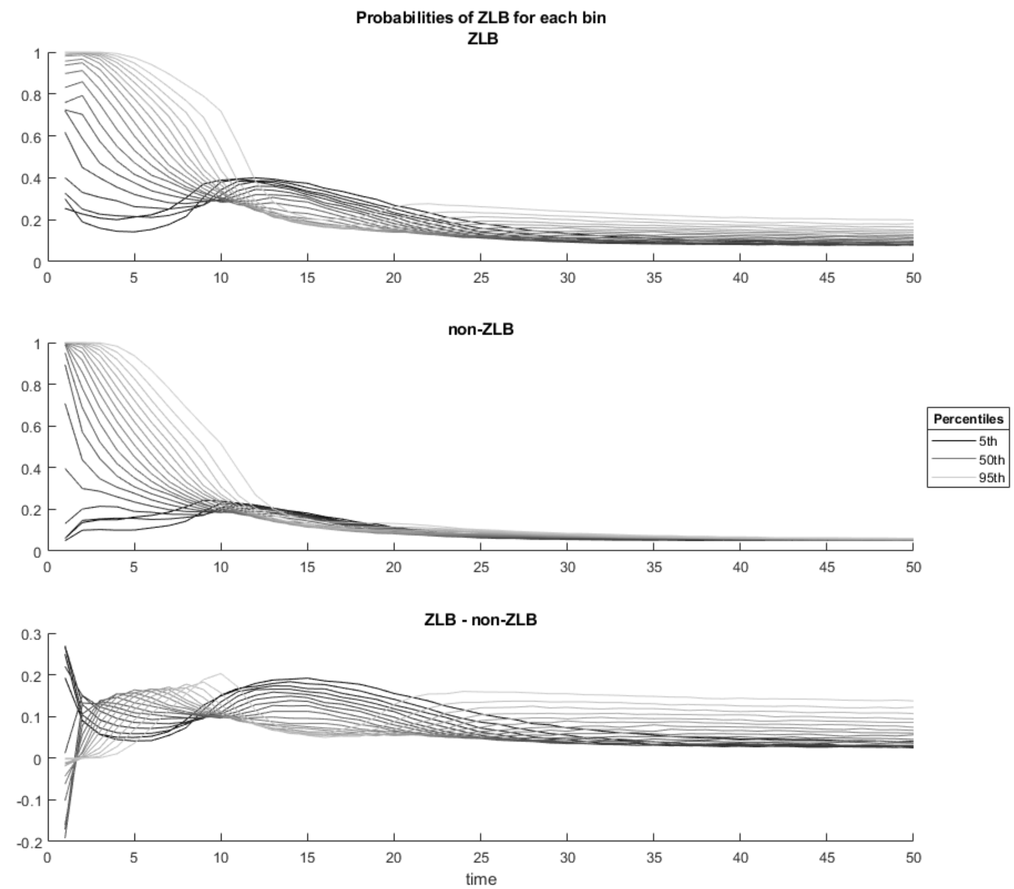

Figure 12: Distribution of the probability of staying on or below the ZLB (defined as $i - ZLB^{value} < ZLB^{tol}$) n periods after impulse.

Bins correspond to the distribution of impulse responses for the interest rate as in Figure 10. Bottom panel shows the difference between the top two panels, a positive number means that the constrained central bank has higher probability of hitting the ZLB than the unconstrained one. In all panels $ZLB^{tol} = 10^{-5}$ (for $ZLB^{tol} = 10^{-4}$ there would be no negative values in panel 3).

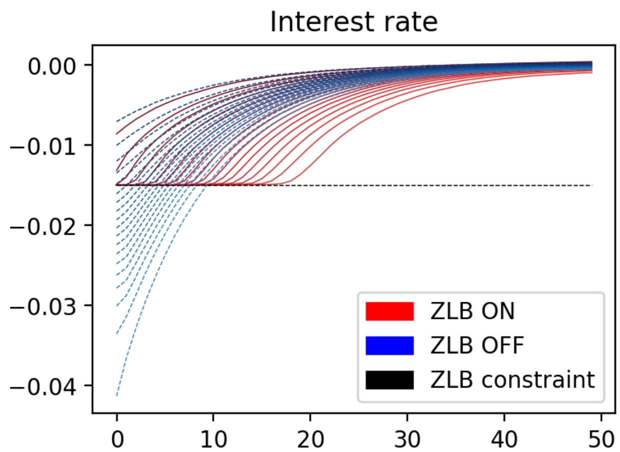

Figure 13: Distribution of interest rates after a large shock ($-2.5\sigma_\varepsilon$).
The figure displays every fifth percentile from 5^{th} to 95^{th} for the raw interest rates (not the impulse responses), the initial condition is the ergodic set of the constrained economy. The horizontal axis shows time.

We also compare the plain after-shock interest rates (as opposed to the impulse responses) of both central banks as in Figure 13. This comparison shows that after a large shock hits the economy the constrained central bank lowers the interest rate faster, keeps the interest rate at the bound for longer, and returns the interest rate to the pre-shock levels slower than its unconstrained counterpart.

Finally, we calculate the probability that the constrained interest rate is below the unconstrained interest rate whenever the latter is above the lower bound. This probability is displayed in Figure 14 and it shows that for the first approximately 25 periods after the shock this is always the case as long as the shock is still negative. All this gives solid evidence of the lower for longer property of optimal policy in our setting.

When the shock turns positive, however, it is the exact opposite. The reason for this property depending on the sign of the shock is that the ZLB destabilizes agents' expectations through the central bank's failure to control the gaps once the interest rate is lowered to its effective lower bound. Since, at that moment, the CB loses the ability to stimulate the economy, the inflation and output gaps start

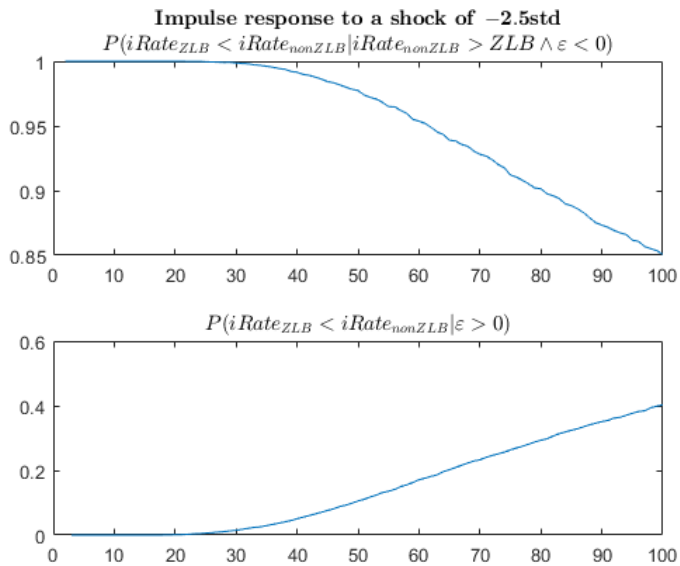

Figure 14: Probability of lower for longer

Horizontal axis shows time.

to decrease and so do agents' expectations. The, initially inflationary, expectations can even turn into deflationary ones if the liquidity trap episode lasts for long enough. While counter-cyclical expectations, which prevailed before the shock, are of a stabilizing nature, the pro-cyclical expectations, resulting from a period of liquidity trap, become a source of volatility when the shock changes its sign.

The mechanics of this are just the opposite from what we have discussed previously. Under pro-cyclical expectations, when the shock becomes expansionary, agents become expecting above average inflation and economic activity, and, hence, responding one-to-one to the shock would make the central bank overshoot its targets. The zero lower bound can thus have significant distortionary effects even when policy is not constrained by it. It can take several dozen quarters for the central bank to contain these expectations, and for the constrained and unconstrained policies to converge, as can be seen for instance from the sample simulation shown in Figure 5. In that chart we can also see how the central bank becomes very restrictive in response to positive shocks that occur after prolonged periods spent at the zero lower bound.

What we, however, also notice is that the optimal policy, although exhibiting the lower for longer property, produces only a small amount of inflation overshooting (as in Figure 11). We especially note the fact that the amount of overshooting decreases with the duration of the liquidity trap, which is in stark contrast to models with rational expectations under commitment. Most of the overshooting comes from the pre-emptive monetary expansion where the central bank tries to avoid hitting the zero lower bound.

The most likely explanation is that the episode of liquidity trap creates deflationary expectations which, moreover, get worse as the liquidity trap progresses. Therefore, even though monetary policy is more relaxed, and stimulates inflation and output, most of the extra monetary stimulus is absorbed by the increasing deflationary expectations. This is in contrast to rational expectations models where, under the optimal policy, the liquidity trap episode creates expectations of above average inflation in the future.

To investigate this channel in more detail we run the following experiment. We choose the median state from the ergodic set of the constrained economy as the initial condition, simulate the economy for several hundred periods with the shock innovations set to zero, and then we introduce a shock of $-2.5\sigma_\varepsilon$, which we keep constant for a given number of periods.[33] After the last period the shock becomes $-0.5\sigma_\varepsilon$, we set the shock innovations to zero again, and simulate the economy for a several dozen more periods. The purpose of this experiment is to compare the response to a small shock that would normally not take the economy to the zero lower bound when the economy has spent a different number of periods in the liquidity trap. The only difference between the responses for different durations is the effect of the ZLB on expectations.

The results are shown in Figure 15 and they highlight the importance of expectations on central bank's precautionary behaviour. Although the shock is too small to take the economy to the zero lower bound under normal circumstances, the graph shows that the central bank may react very strongly even to small shocks if its objectives have been disturbed by the zero lower bound previously. Interestingly, having spent even just a single period at the zero lower bound

33. In principle, we could do this for the whole ergodic set and calculate the median response. Focusing on a single initial condition gives us better control over the expectations channel.

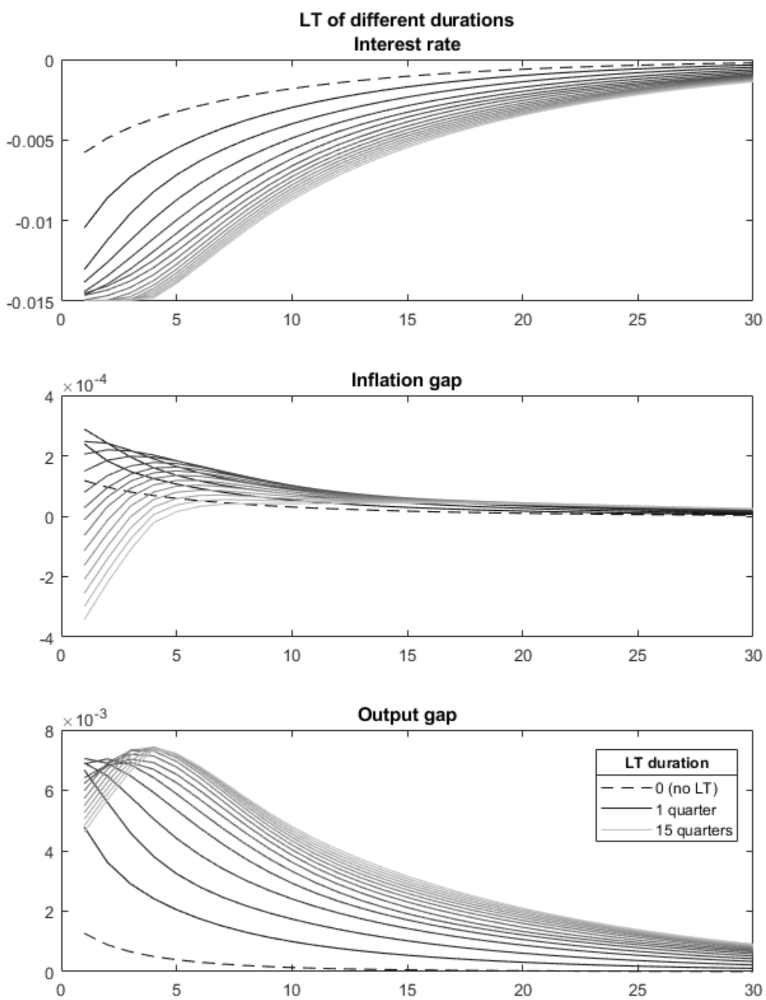

Figure 15: Response to a small shock after a liquidity trap of various durations.

Small shock of $-0.5\sigma_\varepsilon$ takes place in period 1. The economy has been subject to a shock of $-2.5\sigma_\varepsilon$ for n periods starting in period $1-n$ (not shown). Horizontal axis shows time.

produces a substantial distortion of optimal policy. The size of the response doubled purely through the effect a single period at the ZLB had on expectations. This experiment gives another confirmation of the lower-for-longer property of optimal policy, the central bank that has spent a longer period of time at the zero lower bound sets lower interest rates and even keeps the interest rate at the bound for longer still.[34]

In models with rational expectations, the lower-for-longer property is the result of binding announcements of future overshooting of the inflation target, dubbed forward guidance. Such overshooting is then delivered by an excessive monetary loosening when the period of liquidity trap is over. Such announcements have no effect with learning since agents adjust their expectations based on experience rather than on promises. Therefore, it is interesting to observe that even in a model with learning the optimal policy mimics the one from a rational expectations world with commitment.

Nevertheless, the mechanism is very different and it works through two distinct channels. On the one hand, lower interest rates on exit from the liquidity trap are driven by the proximity of the zero lower bound whereby the central bank puts a large weight on the risk of hitting the bound again and, hence, engages in precautionary monetary expansion. On the other hand, having spent several periods at the effective lower bound disturbs agents' expectations, which may even turn from counter-cyclical to pro-cyclical. When the liquidity trap period is over, the low expectations drag both inflation and economic activity downwards and the central bank responds to that by additional monetary loosening. This effect is not present in models where monetary policy is based on the forward guidance principle. We can, colloquially, summarize the difference by saying that the lower for longer in our model is not driven by delivering on a promise of high inflation but rather by dealing with the disappointment from not delivering on agents' past expectations. In this sense optimal policy with learning shares the history-dependence feature of forward guidance, although the nature of this history dependence differs.

34. We have run experiments with even longer periods of liquidity traps, up to 200, and starting with extreme values of expectations (close to upper grid boundaries, those values are outside the ergodic set) to see whether this would lead to a deflationary trap developing. It would not.

4.3 Accuracy

Since our calculations are inevitably affected by numerical and approximation errors we need to test the accuracy of our solution. The usual practice is to test the solution by calculating the Euler equation errors as is done e.g. in Fernández-Villaverde et al. (2015). We were not able to derive the Euler equation in terms of economic primitives, however, the solution of the central bank's problem must satisfy the Bellman equation, too.[35] Hence, we use the BE errors instead.

Denoting the left-hand and right-hand sides of equation (20) as $\widehat{LHS}$ and $\widehat{RHS}$, respectively, the absolute BE errors are calculated as $abs(\widehat{LHS} - \widehat{RHS})$ and the relative errors are calculated as $abs(1 - \frac{\widehat{RHS}}{\widehat{LHS}})$.

We present both absolute and relative errors in log_{10} units, which shall be interpreted as the digit after the decimal point on which the first non-zero value occurs. The relative errors are in some cases unrealistically high due to dividing by a number close to zero. This occurs for the unconstrained central bank since in the long-run it stabilizes the economy so well that it achieves near zero welfare losses. Overall, we judge the errors acceptable. Although not directly comparable, the relative errors are of similar magnitude to Fernández-Villaverde et al. (2015).

5 Robustness

In this section we discuss the effects of alternative parameterizations. There are three potential sources of sensitivity of our results: the process of expectations formation as captured by the learning rate, γ, the preferences of the policy maker as captured by the welfare weight, λ, and the degree to which the effective lower bound inhibits stabilization. The latter is captured by the frequency with which the natural interest rate becomes negative and it is affected mainly by the parameters β, π^*, σ, ρ, and σ_ε.[36]

The individual rate of time preference, β, and the inflation target, π^*, affect the value of the zero lower bound directly. Since a

35. In our model the envelope theorem does not eliminate the unknown derivatives of the value function.

36. The natural interest rate is the same as the long-run one-to-one policy, Figure 26 in the appendix shows the proportions of periods in which it is below the zero lower bound.

Table 3: Bellman equation errors

	non-ZLB				ZLB			
	max	min	mean	median	max	min	mean	median
random grid	-5.05	-18.67	-6.84	-7.72	-3.92	-12.32	-5.88	-7.08
	(-1.86)	(-10.17)	(-3.40)	(-3.61)	(0.64)	(-8.56)	(-2.57)	(-3.42)
long simulation	-22.10	-28.44	-23.44	-23.61	-5.07	-14.71	-8.49	-10.18
	(-0.98)	(-7.33)	(-2.34)	(-2.50)	(-1.72)	(-9.57)	(-3.96)	(-5.03)
ergodic set	-21.90	-29.13	-23.45	-23.61	-4.92	-15.64	-8.57	-10.18
	(-0.79)	(-8.04)	(-2.34)	(-2.50)	(-1.68)	(-10.44)	(-3.97)	(-5.03)
BE grid	-9.53	-27.33	-11.13	-12.70	-10.51	-10.51	-10.51	-10.51
	(-1.15)	(-8.12)	(-3.44)	(-7.49)	(-4.38)	(-8.37)	(-5.73)	(-6.18)

Note: The table shows the absolute Bellman equation errors in $\log_{10}$ units (relative errors in parentheses). BE grid means the grid on which the Bellman equation was solved, by design these errors should be small. For the random grid 100K random points were used. The relative errors for the unconstrained economy are in many cases affected by the division-by-zero problem.

lower π^* and a higher β reduce the steady state interest rate they effectively bring the lower bound closer. The interest elasticity of aggregate demand, σ, does not affect the value of the effective lower bound directly but it has a similar effect as just described. A lower value of σ reduces the effectiveness of monetary policy in influencing aggregate demand and, in response to a negative shock of given size, forces the central bank to lower the interest rate by more. This gives the central bank less room to maneuver before the ZLB is reached, effectively bringing the lower bound closer. The last most important factor is the volatility of the demand shock process. Both a higher variance of the innovations, σ_ε, and a higher persistence of the AR(1) component, ρ, make the process more volatile, which results in the natural rate becoming negative more often. The central bank then lowers its instrument to the effective lower bound more often too. All these factors then have a similar effect in that they result in the central bank reaching the zero lower bound more frequently. To simplify the exposition we will, therefore, only illustrate this effect with the shock persistence parameter, ρ.

We have solved the model with different values of γ, λ, and ρ, changing only one at a turn, and we observed our main result un-

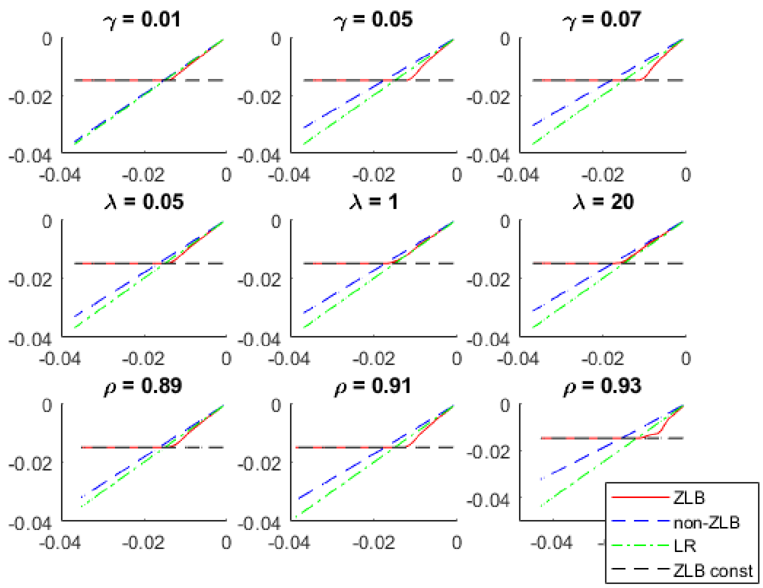

Figure 16: Impact response

The initial condition is the median of the constrained economy.

changed in all experiments. This is illustrated by the generalized impulse response charts in the appendix (see Figures 27 – 29), which show the lack of inflation overshooting discussed earlier. Nevertheless, some interesting patterns emerge.

Generally, increasing volatility, either endogenously by faster expectations updating by the agents or exogeneously through the nature of the stochastic process (or, more generally, in a way orthogonal to expectations formation), the central bank tends to respond stronger to negative shocks. This can be seen in the top and bottom rows of Figure 16. By responding stronger the central bank encourages higher expectations, i.e. lower expectations coefficients, which stimulates inflation expectations. These then serve as a cushion against the ZLB becoming binding in the near future. Interestingly, higher volatility causes the unconstrained central bank to be more restrictive relatively to the long-run one-to-one policy.

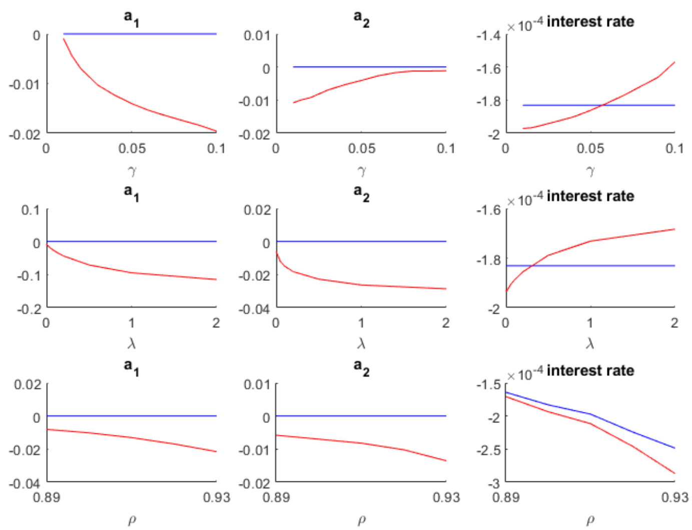

Figure 17: Long run values as functions of parameters

The curves show the median value of the given variable in a long run simulation when γ, λ, and ρ are changing.

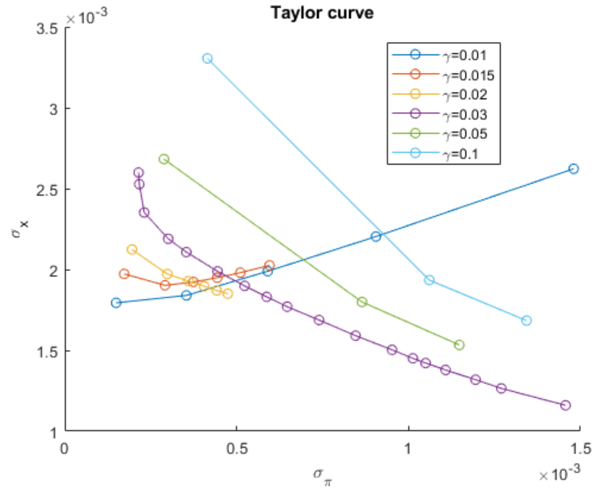

Figure 18: Taylor curves

λ ranges from its benchmark value (0.007) to 0.4 for $\gamma = 0.01$, to 0.5 for $\gamma < 0.03$, to 1.0 for $\gamma > 0.03$ and 10.0 for $\gamma = 0.03$.

5.1 Taylor curve

Taylor (1979) presented the idea that, under rational expectations, there is a trade-off between inflation volatility and output volatility and this trade-off is driven by the central bank's preference for inflation vs. output stabilization (parameter λ in the model). In our model with learning we also observe this trade-off. The combinations of lowest achievable volatility of inflation and output represent an efficiency frontier and they are depicted by, so called, Taylor curve, which we show in Figure 18 for different values of the learning gain parameter, γ.[37] We can see that lower values of the learning gain cause the Taylor curve to shift inwards and pivot. The inward shift means that the central bank is more successful in stabilizing the economy when agents update their expectations more slowly, hence, without sudden shifts.

For very low values of γ, however, the Taylor curve becomes posi-

37. For high values of λ and/or low values of γ the computations of the Taylor curve become rather costly, therefore, we economized on the number of points representing the Taylor curves. For the benchmark (purple) we calculated the Taylor curve for λ up to 10, which is a reciprocal value to 0.1 weight on output.

tively sloped, meaning that higher welfare weight attached to output stabilization causes more volatility of both output and inflation and, thus, such an excessive output stabilization is inefficient. By putting less weight on output stabilization the policy maker could achieve lower volatility of both the inflation and the output gaps. Overall, we conclude that the central bank's policy is more efficient in stabilizing economic fluctuations when agents change their expectations more smoothly.

This is because episodes of binding zero lower bound are the result of random shocks, which are not under the central bank's control. When the central bank fails to stabilize the economy during these random episodes, the agents learn to expect such failures to occur in the future again. However, with lower learning gain parameter the expectations are affected to a lesser extent. Once the period of the binding zero lower bound is over the central bank restores the control over agents' expectations and, with a suitable, now unconstrained, policy, it can prevent the episode to have a prolonged effect on the economy. Our results thus suggest that the welfare costs of the zero lower bound are magnified by more volatile expectations formation.

6 Conclusion

This paper set out to study the design of optimal monetary policy when central banks are constrained by the zero lower bound (ZLB) and agents form expectations using constant gain, least-squares learning. Compared to models with rational expectations in which the central bank can commit to a future policy path, in our setting the central bank has no direct influence on agents' future expectations and the indirect influence, which it has through its current period policy instrument, is lost when that instrument reaches the effective lower bound. In this setting, therefore, the policy advice of Eggertsson and Woodford (2003), often referred to as 'forward guidance', which is to announce such a policy path that makes the policy maker overshoot their inflation target at some point in the future when the ZLB no longer binds, is ineffective for it cannot alter agents' expectations.

This ineffectiveness of forward guidance then carries over to our results not showing the inflation overshooting, which is the key factor shaping optimal policy under rational expectations. In our model, the amount of inflation that the central bank optimally allows is even decreasing with the duration of the liquidity trap episode. This is in stark contrast to EW and it is our main point. Nevertheless, optimal monetary policy with learning still shows some similarities to optimal policy derived under rational expectations with commitment: in response to smaller shocks that threaten to take it to the effective lower bound the central bank responds disproportionately more to shocks of larger size, and hence it reaches the ZLB sooner; in response to larger shocks, which cause a liquidity trap, it keeps the interest rate at the bound for longer, and it allows the interest rate to stay away from the pre-shock levels for longer too.

When responding to smaller shocks, there is a clear precautionary motive in pursuing this type of monetary expansion before the ZLB is reached. In fact, the central theme of optimal policy under the ZLB, and this is a distinctive feature of our model, is that by reacting more aggressively to both negative and positive shocks the central bank manipulates the agents' expectations in such a way that they become counter-cyclical – in times of a recession agents expect expansion and positive inflation. This then lowers the cost of the recession should the ZLB become binding. In this sense, a rational policy maker can mitigate the cost of the effective lower bound even in the absence of a direct lever over agents' expectations.

The response to large shocks, then, is driven partly by the same precautionary motive in face of the risk of the liquidity trap re-occuring and partly by deflationary pressure, which develops during the period of binding ZLB and which calls for additional monetary stimulus. At the same time, and in contrast to EW, this stimulus is not large enough to result in much inflation overshooting, following a period of binding ZLB. The optimal policy even becomes more restrictive the longer the duration of the liquidity trap. This is due to the episode of binding ZLB destabilizing the belief coefficients, which the policy maker does not want to allow to carry over to future periods. In principle, there is also the potential for a deflationary trap to develop but we do not observe it in our simulations. We leave a detailed investigation of the existence of deflationary traps in our setup for future research.

In addition to the above, we obtain two new results when the ZLB is not allowed to bind. First, notwithstanding the presence of constant gain learning, the economy converges to the rational expectations equilibrium (REE). Second, when not in the REE, the demand shock causes an inflation-output trade-off, which does not occur in models with rational expectations. These findings are new in the literature, as far as we are aware.

Appendices

A Details of the numerical procedure

This appendix sketches the algorithm we used to solve the central bank's optimization problem and discusses some issues in more detail.

Let $\mathscr{Z} \subset \mathbb{R}^4$ be the state space and $z = (a_1, a_2, r, \varepsilon)$, $z \in \mathscr{Z}$ be a vector of state variables. We choose a bounded subset of the state space, $\overline{\mathscr{Z}} \subseteq \mathscr{Z}$, and define a fixed grid of points on it giving the finite set $\mathrm{Z} = \{z_i : z_i \in \overline{\mathscr{Z}}, i = 1, \ldots, n_z\}$. Denoting the values of the value function on this grid as an n-tuple $V^{\mathrm{Z}} = \{(v_1, v_2, \ldots, v_{n_z}) : v_i = V(z_i), z_i \in \mathrm{Z}, i = 1, \ldots, n_z\}$, we guess the values v_i and construct an approximation $\widehat{V}(z), z \in \overline{\mathscr{Z}}$ such that $\widehat{V}(\mathrm{Z}) = V^{\mathrm{Z}}$.[38] To solve the central bank's problem we iterate on the Bellman equation (11) in the following form

$$\widehat{V}_{j+1}(z) = \min_i \left\{ L(z,i) + \beta E\left[\widehat{V}_j(z')|\varepsilon\right]\right\}, \tag{20}$$

where j denotes the iteration step. The solution to the problem are functions $\widehat{V}(\cdot)$ and $\widehat{h}(\cdot)$ such that $\widehat{V}(z) = L(z, h(z)) + \beta E\left[\widehat{V}(g(z, h(z)))|\varepsilon\right]$ up to the required tolerance level.

The algorithm is as follows:

1. Choose a grid for the state variables, Z
   This amounts to choosing both $\overline{\mathscr{Z}} \subseteq \mathscr{Z}$ and $\mathrm{Z} \subset \overline{\mathscr{Z}}$.
2. Choose a functional form for the approximation of the value function, $\widehat{V}(\cdot)$ [39]
   Note that $\widehat{V}(\cdot)$ is parameterized by a parameters vector b, which is to be determined in step 4.1. Thus $\widehat{V}(\cdot) \equiv \widehat{V}(\cdot; b)$, where there is no danger of confusion we make the dependence on b implicit.
3. Choose the initial guess for the value function, V_0^{Z}
4. Iterate on eq. (20), i.e. for $j = 0, \ldots, j_{max}$
   4.1 Using Z and V_j^{Z} construct the interpolant $\widehat{V}_j(\cdot) \equiv \widehat{V}(\cdot; b_j)$

38. In other words $\widehat{V}(\cdot)$ is an interpolant

39. Actually, there is more to it than just choosing the functional form. As we do an interpolation, besides the functional form we need to choose an interpolation scheme.

4.2 Precalculate the expected value for the approximation of the value function
To do this construct the interpolant $\widehat{V}_j^e(z), z \in \overline{\mathscr{Z}}$ with the property $\widehat{V}_j^e(z_i) = E_{\varepsilon'}\left[\widehat{V}_j(a_1, a_2, r, \varepsilon')|\varepsilon\right], \forall a_1, a_2, r, \varepsilon :$ $z_i = (a_1, a_2, r, \varepsilon) \in \mathrm{Z}$. This means that $\widehat{V}_j^e(\cdot)$ agrees with $E_{\varepsilon'}\left[\widehat{V}_j(\cdot)|\varepsilon\right]$ on all grid points. This step is unnecessary but saves computational time in step 4.3.

4.3 For each grid point $z_i \in \mathrm{Z}$ use a numerical optimization procedure to solve the minimization problem on the RHS of (20) (replace $E[\widehat{V}_j(z')|\varepsilon]$ with $\widehat{V}_j^e(z)$), the optimal value on the RHS gives V_{j+1}^{Z} and the optimal interest rate gives H_j^{Z} (with H^{Z} defined analogously to V^{Z}).

4.4 Verify the convergence of V^{Z}
If $\|V_{j+1}^{\mathrm{Z}} - V_j^{\mathrm{Z}}\| > V_{tol}$ then if $j = j_{max}$ declare failure to converge and quit, if $j < j_{max}$ set $j = j + 1$ and go back to step 4.1. Otherwise declare convergence, set $\widehat{V}(\cdot) = \widehat{V}_j(\cdot)$, and proceed to step 5.

5. Upon convergence recover the policy function
Using the same functional form as in step 2 use H_j^{Z} obtained in step 4.3 together with Z to construct an approximation $\widehat{h}(\cdot)$ for the policy function $h(\cdot)$.

In the following subsections we discuss some of the steps of the algorithm in more detail.

A.1 Choice of the grid

Although we present steps 1 and 2 separately, they are not totally independent of each other. The choice of the grid and the approximation/interpolation scheme have to be made with the view of each other. On one hand the choice of the grid may be dictated by the chosen interpolation scheme, an example would be the choice of Chebyshev polynomials interpolation, which desires the grid points to be placed at the roots of the Chebyshev polynomial. Another example can be Smolyak interpolation, which uses a carefully constructed sparse grid to provide an alternative to the usual tensor product grid interpolation. On the other hand the chosen grid may require a specific interpolation scheme that is flexible enough to accomodate such a grid. One example can be Spline interpolation, which is typically used in problems that exhibit kinks in the policy function (such as ours) due to the need to place some grid points arbitrarily in the region where the kink occurs. Another example would be Barycentric interpolation, which is defined on triangular grids.

Since we use cubic spline interpolation the grid is required to be rectangular. We therefore choose an independent grid for each state variable, the details of which are discussed below, and construct their tensor product. Denoting the unidirectional grids as $\mathrm{G}^j, j = \{a_1, a_2, r, \varepsilon\}$ the four-dimensional grid is given by $\mathrm{Z} = \mathrm{G}^{a_1} \otimes \mathrm{G}^{a_2} \otimes \mathrm{G}^r \otimes \mathrm{G}^\varepsilon$. The choices discussed below give us 103935 grid points.

Due to the presence of the zero lower bound constraint in the present problem the choice of the grid for individual state variables requires some care. The ZLB constraint causes a kink in the policy function therefore to capture the kink well and not to lose too much accuracy in its neighbourhood more grid points need to be placed around the kink. The obvious complication is that it is not known in advance where the kink occurs moreover the location of the kink can differ depending on the values of all four state variables and also the kink can move substantially when we experiment with different model parameters. Additionally, due to the curse of dimensionality problem we have to be rather conservative in the choice of the number of grid points. The grid boundaries are chosen such that they are never hit in a five hundred thousand periods simulation.

In simulations the inflation expectations coefficient, a_1, tends to be positive and when negative the values are relatively small in absolute value. Therefore we choose an asymmetric range of $[-0.15, +1.5]$ and make the grid quadratic around zero. We use 13 points in this direction.

The output gap expectations coefficient, a_2, tends to be very volatile since the ZLB prevents the central bank frequently from stabilizing the output gap, which also carries very small weight in the central bank's objective function. Therefore we use a wider range of $[-5, +5]$ and a larger number of points, 15. For some parameter values the ZLB becomes binding more frequently, which makes the a_2 coefficient more volatile. In those cases we enlarge the range up to $[-10, +10]$ and use more grid points, up to 21 as needed. As in the previous case we make the grid quadratic around zero.

As is apparent from equation (9) the dynamics of the r coefficient does not depend on central banks policy and is fully determined by the stochastic process for ε. In our simulations the median value of r is slightly below the shock variance therefore we center the grid around σ_ε^2. Moreover equation (8) reveals that the learning recursion can be sensitive to small values of r and in fact our results confirm that the policy function has a lot of curvature close to the lower grid boundary. Therefore we place double the points below σ_ε^2 compared to the number of points above and we make the grid linear below σ_ε^2 while making it quadratic above. The grid boundaries are chosen as [0.00001,0.0015], which is one order of magnitude below and above the median of r. We choose 13 grid points in this direction.

Since we approximate the continuous AR(1) process for ε by a Markov chain the grid should be determined by the chosen approximation method. At the same time, however, since the kink in the central bank's policy function is especially apparent along the shock dimension we require the approximating method to allow some flexibility in choosing the grid points. This restricts available choices of the approximation method.[40]

While Kopecky and Suen (2010) advocate the use of Rouwenhorst method arguing that it delivers an approximation which represents the stochastic properties of the underlying AR(1) process most accurately (for a given number of grid points), especially for highly persistent processes. In fact the main advantage of the Rouwenhorst method is that it delivers comparable accuracy as other methods using a smaller number of nodes. We however need a large number of points to capture the curvature of the value function and the kink of the policy function accurately, hence, this advantage of the Rouwenhorst method would not be of much benefit to us. Furthermore, and more importantly, in Rouwenhorst method the grid is created as evenly spaced and the grid boundaries are proportional to the standard deviation of the AR(1) process being approximated while the factor of proportion is positively related to the number of grid points. This means that for a large number of grid points the grid would be unnecessarily wide (thus sparse) and we would not be able to choose the location of the nodes either, leaving us unable to place them where we expect the kink to occur.

Therefore we approximate the shock's stochastic process using the Tauchen method (Tauchen 1986). This method also distributes the

40. This is also a reason to favour Markov chain approximation over using Gauss-Hermitte quadrature.

grid points evenly but there is no restriction that would prevent us from choosing the points arbitrarily, in fact the author himself mentions that the choice of equidistant grid points is mere convenience and that there is likely to be more efficient placement of the nodes (p. 179). We therefore choose the grid points arbitrarily and use Tauchen method to determine the transition probabilities. The rationale for this choice is that for a given number of nodes we insert an additional node in the region of interest, which should not harm the quality of the approximation. Since we are using a rather large number of points their placement should not be much of a concern.

We choose the grid boundaries to be $\pm 5\sigma_\varepsilon$. This value seems extreme, however, even with a rather generous choice of $\pm 3.5\sigma_\varepsilon$ (a typical value in practice is $\pm 2.5\sigma_\varepsilon$) we observed about 0.5% of the shocks in our simulations to hit the grid boundary. Since we focus on the effect of the zero lower bound on central bank's response to shocks and since the ZLB is associated with especially large shocks we better adopted a wide grid. Another point worth mentioning is that when drawing a large shock when calculating the impulse responses we still need some room for drawing more bad shocks after the initial impulse and the grid has to allow for this. For calculating the impulse responses we use shocks as large as $4\sigma_\varepsilon$.

With as wide grid as we use we need a large number of grid points. We use 41 points in total and place two thirds of the points in the negative domain while making the grid linear and one third of the points we place in the positive domain and make the grid quadratic. The rationale for this choice is that the kink occurs for negative values of the shock and given our parameterization it may not be located too close to zero thus quadratic grid in this region would be inadequate. For some combinations of parameters, however, the ZLB becomes binding much closer to zero and in such cases we make the grid quadratic.

For some model parameters configurations (typically high β, high ρ and low σ) the ZLB binds so frequently and leads to so high volatility that the standard parameterization described above is inadequate. Either the grid boundaries are hit in the simulations or we may not obtain convergence. In such cases it is necessary to expand the grid boundaries and increase the number of grid points in some dimensions. We increase the upper bound on a_1 up to 3 and enlarge the grid boundaries for a_2 up to ± 10, at the same time we increase the number of grid points in these dimension to 17 and 21, respectively, which gives us 190281 points in total.

A.2 The Interpolation Scheme

To gain more flexibility in approximating the value and policy functions around the kink we use cubic splines in the B-spline form, which is especially convenient for numerical computations (especially in higher dimensions). Following de Boor (1978) we give a brief (and somewhat sketchy) description of the ideas underlying B-spline interpolation.[41]

Let $t = \{t_i\}_{i=1}^{n+k}$ be a non-decreasing sequence of knots and let $S_{k,t}$ and $B_{i,k,t}$ be a Spline and an i-th B-spline, respectively, of order k given the sequence of knots t. Making the dependence on k and t implicit we may denote them as S and B_i where there is no danger of confusion. A spline function is a piecewise polynomial, which satisfies certain continuity conditions captured by the knot sequence t. Typically, splines of order k are C^{k-2} functions (e.g. standard cubic splines have continuous second derivative). An i-th B-spline of order k is defined by the recursive relationship[42]

$$B_{i,k}(x) = \frac{x - t_i}{t_{i+k-1} - t_i} B_{i,k-1}(x) + \frac{t_{i+k} - x}{t_{i+k} - t_{i+1}} B_{i+1,k-1}(x), \qquad (21)$$

$\forall x \in \mathbb{R}$, with $B_{i,1} = 1, x \in [t_i, t_{i+1})$, and 0 otherwise. Among the properties of B_i belong the following, B_i is positive on and zero outside the interval $[t_i, t_{i+k}]$, only k neighbouring B-splines can be non-zero on any particular interval $[t_j, t_{j+1}]$, namely $B_{j-k+1}, \dots, B_j$, and $\sum_i B_i = 1$. It can then be shown that the sequence of B-splines, $\{B_i\}_{i=1}^n$, forms a basis for the linear space of splines, i.e. $\mathbb{S}_{k,t} = \{\sum_i \alpha_i B_{i,k,t} : \alpha_i \in \mathbb{R}, \forall i\}$. $S_{k,t} \in \mathbb{S}_{k,t}$ are considered functions on $[t_k, t_{n+1}]$.

Given a set of Lagrange data, $\{(y_i, x_i)\}_{i=1}^{n_x}$ with x_i strictly increasing, we can construct a spline interpolant, S, by a suitable choice of t (among other requirements $t_k \leq x_1$ and $t_{n+1} \geq x_{n_x}$) and by solving the system of the following interpolation conditions for α_i, $\sum_i \alpha_i B_i(x_j) = y_j, j = 1, \dots, n_x$.

In multiple dimensions then the spline $S \in \mathbb{S}^m$ can be constructed as the tensor product of the univariate B-splines. Taking the two-dimensional case as an example, $\mathbb{S}^2 = \mathbb{S}_{k,t} \otimes \mathbb{S}_{h,s} = \{\sum_{i,j} \gamma_{i,j} B_{i,k,t} \widetilde{B}_{j,h,s} : \gamma_{i,j} \in \mathbb{R}, \forall i,j\}$. For a set of Lagrange data $\{\{(y_{i,j}, x_{i,j})\}_{i=1}^{n_x^1}\}_{j=1}^{n_x^2}, x_{i,j} \in \{x_i^1\}_i \otimes \{x_j^2\}_j$, $\{x_i^1\}_i$ and $\{x_j^2\}_j$ strictly increasing, the parameters $\gamma_{i,j}$

41. The most relevant are chapters IX, XIII, and XVII

42. An alternative definition involves divided differences of the function $(t-x)_+^{k-1}$ given x, see chapter IX

of $S \in \mathbb{S}^2$ are determined by the system of interpolation conditions $\sum_{i,j} \gamma_{i,j} B_i(x_l^1)\widetilde{B}_j(x_r^2) = y_{l,r}, \forall l, r$. Note that it is not necessary to solve this system as the parameters $\gamma_{i,j}$ can be calculated efficiently knowing $B_i(x_l^1), \widetilde{B}_j(x_r^2)$ and the associated parameters vectors α and $\widetilde{\alpha}$ that define S and $\widetilde{S}$.

The above then maps in our algorithm as follows, in step 2 for $j = \{a_1, a_2, r, \varepsilon\}$ we set $k^j = 4$ and we construct the knot sequences t^j, thus this is done only once. In step 4.1 we calculate the coefficients α^j and γ and we evaluate the basis functions B^j_{i,k^j,t^j} on each grid point in dimension j. Hence this is done only once per iteration. Finally, in steps 4.2 and 4.3 we evaluate the interpolant S^4 at different points as needed. This is done many times in each iteration and it is the most computationally expensive part of the algorithm due to the need to bracket the point z^j on the grid G^j.[43]

A.3 Details of the Value Function Iteration Algorithm

For the initial guess in step 3 we use the zero function, which we found to work well, although there are other possibilities, some of which might even result in smaller number of iterations needed for convergence. For the version with zero lower bound we make the use of the non-zero lower bound solution to construct a better initial guess. We start with the non-ZLB value function and iterate the non-ZLB policy function (with the ZLB imposed on it) skipping the optimization step 4.3 until convergence. We use the resulting value function as the initial guess and redo step 4.

When evaluating the right-hand-side of the Bellman equation, (20), one needs to calculate the expected value term. The conventional approach is the straightforward calculation of the expectations "on-the-fly" whenever the RHS needs to be evaluated (i.e. skipping step 4.2). This however results in unnecessary and often very costly overhead. We therefore precalculate the expected value function before the optimization problem is solved hence avoiding unnecessary calculations of the integral that defines the expected value.[44] Dur-

43. To carry out the required calculations we use a modernised version of the original Fortran library from de Boor (1978) available at https://github.com/jacobwilliams/bspline-fortran

44. In our case, due to the Markov chain approximation, the integral reduces to a sum.

ing the optimization step, 4.3, we then evaluate an approximation of the expected value term avoiding unnecessary repeated calculations. This approach is advocated for example in Judd et al. (2017) and most recently used in Druedahl and Jorgensen (2017).

To solve the optimization problem in step 4.3 we use the *Golden Section Search* method (see Press et al. 1986). This is a derivative free method and is very robust, however, it has linear rate of convergence, which makes it somewhat slow. Alternatively we could use for instance the Newton-Raphson method, which has a quadratic rate of convergence, it however requires the first two derivatives of the value function and a good initial guess otherwise it can be unstable. We decided to favour robustness of the Golden section search method. As a technical side note, before solving the minimization problem we tranformed it to the one of maximization. This is for mere convenience to be able to reuse existing libraries. Observing that $\min U = -\max\{-U\}$ we transformed eq. (10) in

$$V(a_1, a_2, r, \varepsilon) = -\max_i \left\{-(\pi_t^2 + \lambda x_t^2) - \beta E\left[V(a_1', a_2', r', \varepsilon'|\varepsilon)\right]\right\} \quad (22)$$

and in each iteration step we applied the numerical optimization routine on that.

A.4 Generalized Impulse Responses

Here we present an overview of the key ideas from Gallant, Rossi, and Tauchen (1993).

Let $\{y_t\}_{t=-\infty}^{\infty}, y \in \mathbb{R}^M$ be a strictly stationary process with a conditional density function that depends upon at most L lags. Denote the L lags of y_{t+1} by $x_t = (y_{t-L+1}', \dots, y_t') \in \mathbb{R}^{ML}$ and write $f(y|x)$ for the (one-step ahead) conditional density, which is time invariant (i.e. does not depend on t) due to the strict stationarity assumption.

Define the conditional mean profile $\{\hat{y}_j(x)\}_{j=0}^{\infty}$ corresponding to initial condition x by $\hat{y}_j(x) = E[y_{t+j}|x_t = x] = \int y f^j(y|x)dy$, where $f^j(y|x)$ denotes the j-step ahead conditional density

$$f^j(y|x) = \int \dots \int \left[\prod_{i=0}^{j-1} f(y_{i+1}|y_{i-L+1}, \dots, y_i)\right] dy_0 \dots dy_{j-1}, \quad (23)$$

with $x = (y_{-L+1}', \dots, y_0')'$. (If a dummy variable of integration coincides with an element of x then that integration is omitted.)

For reasons explained in section 4.2.1 we use a conditional median profile instead. Let us denote it $\{\tilde{y}_j(x)\}_{j=0}^{\infty}$ to distinguish from the mean profile and define it as $\tilde{y}_j(x) | \int_{y_{min}}^{\tilde{y}_j(x)} f^j(y|x)dy \geq 0.5 \wedge \int_{\tilde{y}_j(x)}^{y_{max}} f^j(y|x)dy \geq 0.5$ with $f^j(y|x)$ defined as above. In what follows $\hat{y}_j$ would be replaced with $\tilde{y}_j$.

Let δy^+ represent a small perturbation to the contemporaneous y_0 and set $x^+ = (y'_{-L+1}, \dots, y'_0)' + (0, 0, \dots, \delta y^{+'})'$ and $x^0 = (y'_{-L+1}, \dots, y'_0)'$. Thus x^+ is an initial condition corresponding to an impulse added to contemporaneous y_0 and x^0 represents the base case with no impulse. Then define $\hat{y}_j^+ \equiv \hat{y}_j(x^+)$ and $\hat{y}_j^0 \equiv \hat{y}_j(x^0)$.

The non-linear impulse response is then naturally defined as the net effect of the impulse δy^+, i.e. $\hat{y}_j^+ - \hat{y}_j^0$. The integrals in the conditional moment profile are then calculated by Monte Carlo integration. Let $\{y_j^r\}_{j=1}^{\infty}, r = 1, 2, \dots, R$ denote R simulated realizations of the process starting from $x_0 = x$. In other words, y_1^r is a random draw from $f(y|x)$ with $x = (y'_{-L+1}, \dots, y'_0)'$, y_2^r is a random draw from $f(y|x)$ with $x = (y'_{-L+2}, \dots, y'_0, y_1^{r'})'$, etc. Then $\hat{y}_j = \int \dots \int y_j \left[\prod_{i=0}^{j-1} f(y_{i+1}|y_{i-L+1}, \dots, y_i)\right] dy_0 \dots dy_j \doteq \frac{1}{R}\sum_{r=1}^{R} y_j^r$ (analogously for $\tilde{y}_j$).

To obtain some sort of representative impulse response, let's call it *Generalized Impulse Response*, one can either choose a representative initial condition x^0 or calculate the set of non-linear impulse responses for a large number of different initial conditions and average the responses over the initial conditions. The former strategy is simpler but perhaps less representative of the dynamics of the economy. In our model, however, it is useful for a direct comparison to the zero lower bound, which would otherwise not be possible.

With the view of the preceding discussion we approach the computations as follows. After we obtain the policy function in step 5, we first calculate the ergodic set of the economy running m independent simulations for $n^b + 1$ periods and discard the first n^b of them. This gives us a sample of m states of the economy, which we call the *Ergodic Set* and which we use as initial conditions for calculating the non-linear impulse responses. We use $m = 25000$ and $n^b = 50000$, we also experimented with values as low as $m = 5000$ and $n^b = 2000$ and the results were not noticeably different.

To calculate the generalized impulse response we calculate the non-linear impulse response for each element in the ergodic set using $R = 2000$ and, in each period, taking the median among these R

simulations.[45] Then, for the whole set of impulse responses, we sort the observations in each period using a quicksort algorithm[46] and report the desired quantiles and the mean. To generate the shocks in the simulations we draw random values for the innovations of the AR(1) process (see eq. 18) and interpolate the policy function for values of ε that do not lie on the grid. For generating the random numbers we use the Fast Mersenne Twister generator (commonly referred to as SFMT19937).[47]

45. When using the mean profile instead of the median, we used $R = 100$. We also tried $R = 500$ and the differences were very small, almost unnoticeable. When using the median profile, though, a higher R is necessary to avoid jumps in the impulse response. When calculating an NIR for only a single initial condition (as opposed to calculating the whole distribution across which we would aggregate) we used $R = 100000$.

46. See e.g. Press et al. (1986)

47. We use the implementation from Intel Math Kernel Library

B Additional Figures

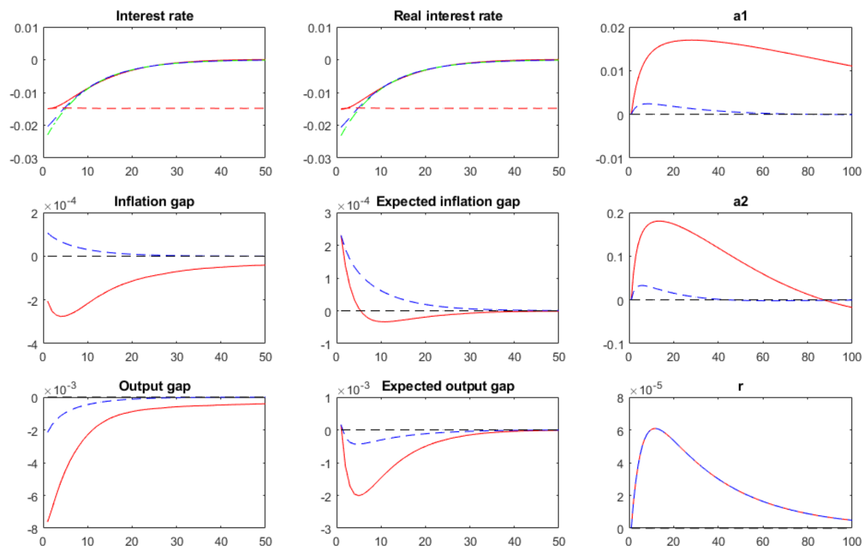

Figure 19: Impulse responses to a large shock ($-2.5\sigma_\varepsilon$)

The same as Figure 9 except the impulse responses are constructed using the **mean** profiles (as opposed to median). Red lines are for the ZLB policy, blue without the ZLB, and green the long-run equilibrium. The initial condition is the median from the ergodic set of the constrained economy.

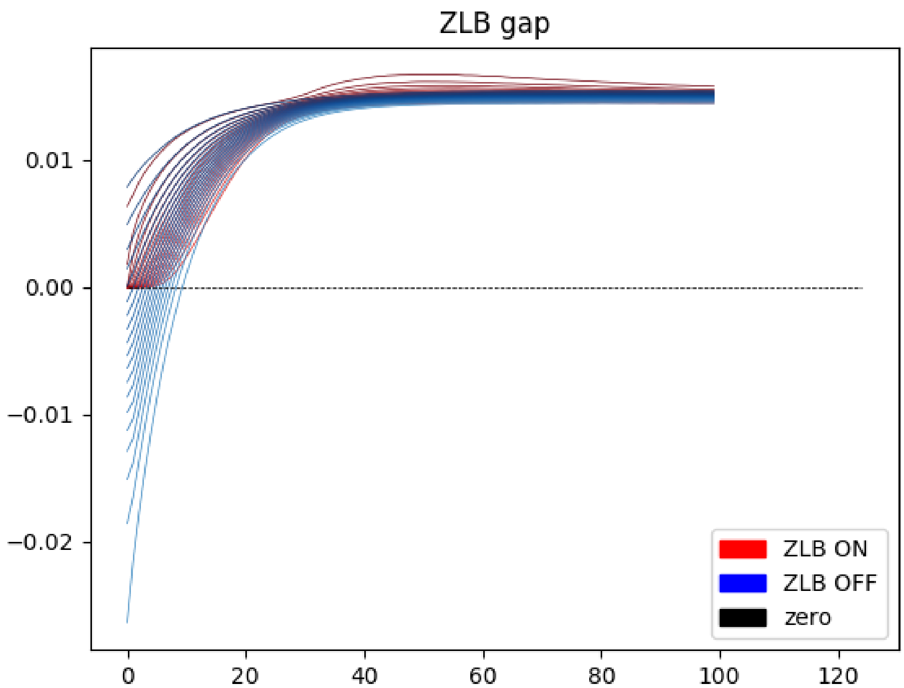

Figure 20: The whole distribution of interest rates normalized so that the ZLB constraint equals zero (i.e. ZLB gap is defined as the interest rate minus the value of ZLB); ZLB vs non-ZLB economy. Using **mean** as the future aggregator.

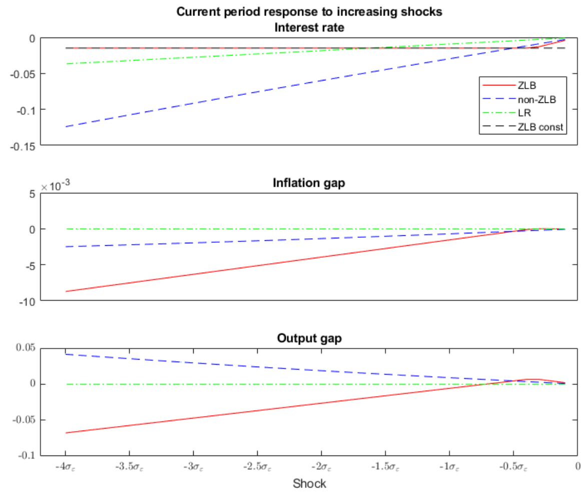

(a)

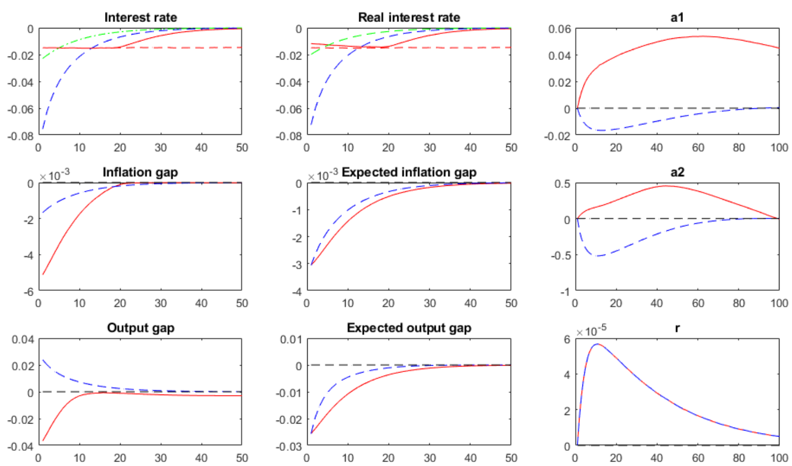

(b)

Figure 21: Impulse response to a shock occuring after a long liquidity trap. Red – constrained policy, blue – unconstrained policy, and green – long-run policy. The initial condition is the state of a constrained economy which has been in a liquidity trap for the last 15 periods (except the counterfactual shock is at its empirical median). The economy had been in the median state before the liquidity trap occured. Panel (a) – current period response to shocks of varying sizes. Panel (b) – response to a large shock ($-2.5\sigma_\varepsilon$).

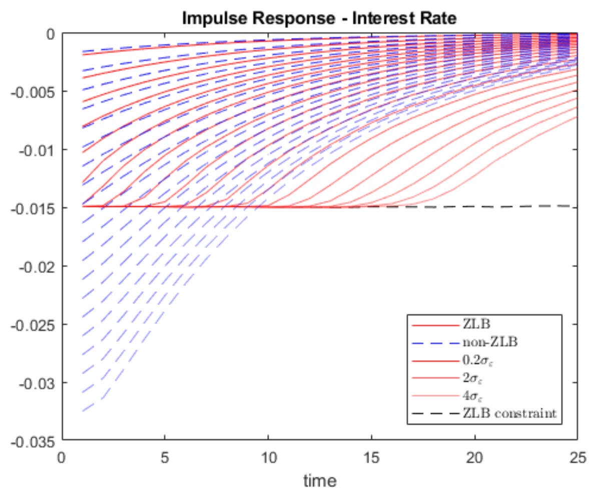

Figure 22: Non-linear impulse response to shocks of varying sizes.
The size of the shock ranges from $-0.2\sigma_\varepsilon$ to $-4\sigma_\varepsilon$ with increment of $0.2\sigma_\varepsilon$. The initial condition is the median state from the ergodic set of the constrained economy.

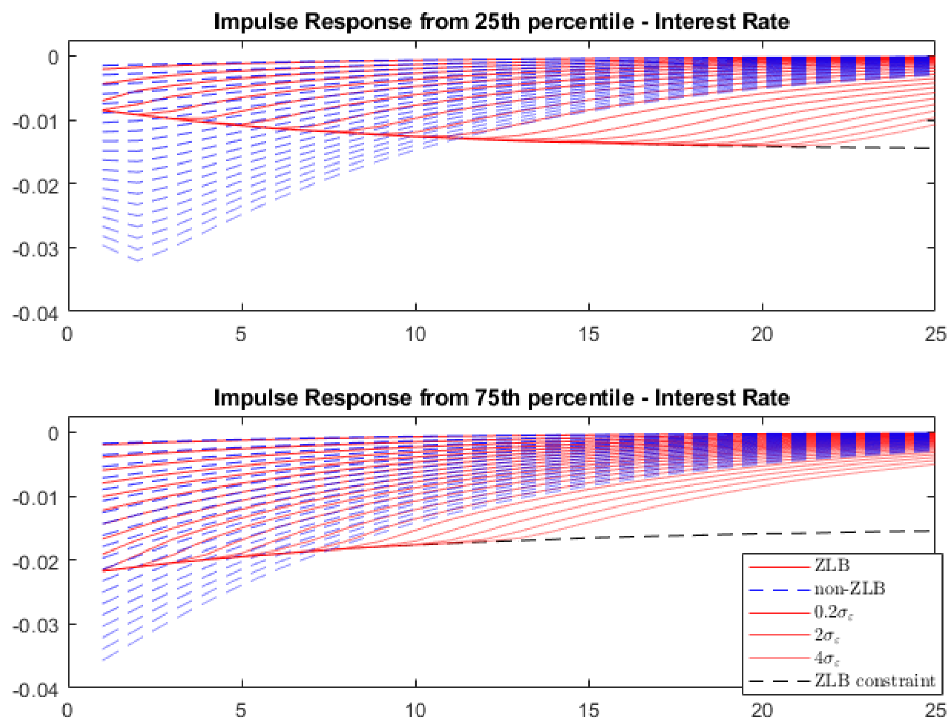

Figure 23: Non-linear impulse response to shocks of varying sizes.
The size of the shock ranges from $-0.2\sigma_\varepsilon$ to $-4\sigma_\varepsilon$ with increment of $0.2\sigma_\varepsilon$. The initial condition are the lower and upper quartile states from the ergodic set of the constrained economy. Note that the ZLB constraint is defined as the maximum allowable negative deviation from the counterfactual policy without the impulse and as such it is not a horizontal line.

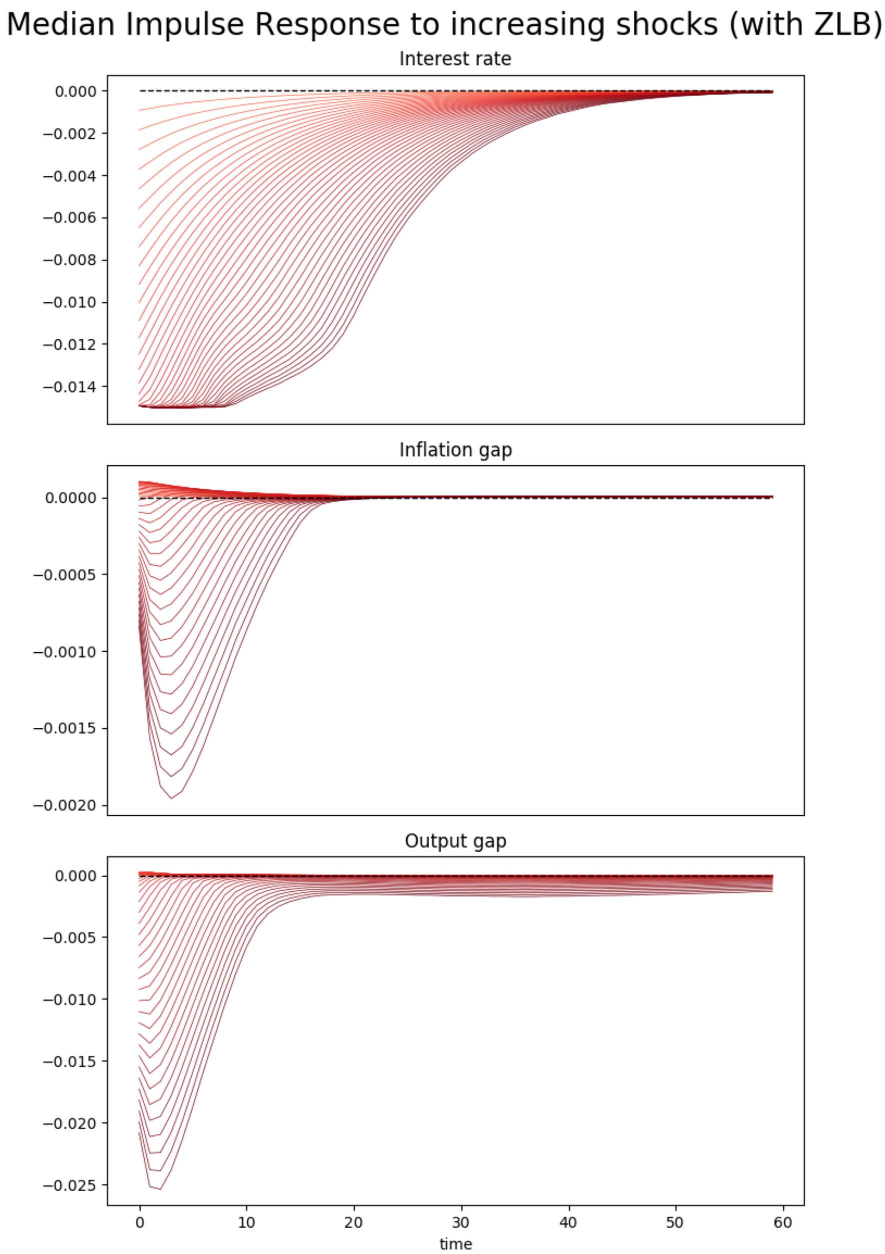

Figure 24: Median non-linear impulse response to shocks of varying sizes. The size of the shock ranges from $-0.1\sigma_\varepsilon$ to $-4\sigma_\varepsilon$ with increment of $0.1\sigma_\varepsilon$ (horizontal axis shows time).

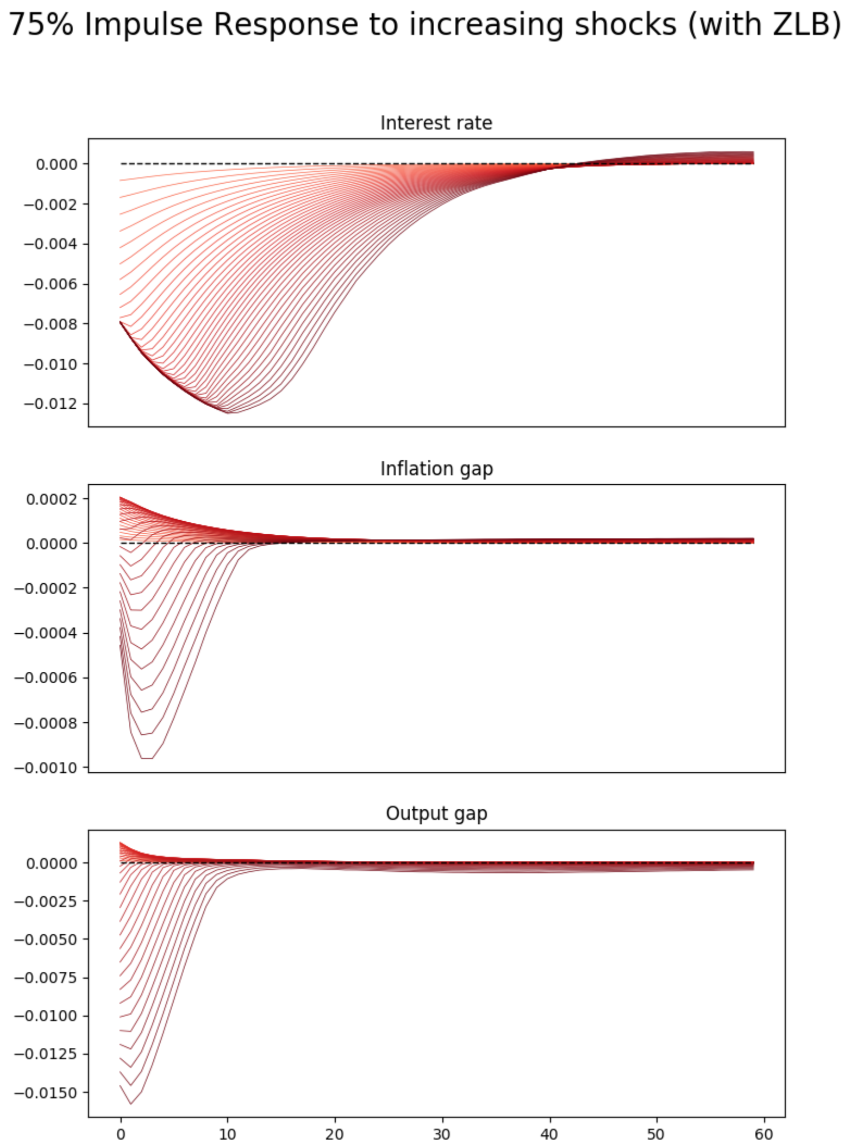

Figure 25: 75% non-linear impulse response to shocks of varying sizes.
Future aggregator is the median. Compared to Figure 24 this chart shows overshooting of the interest rate and little more overshooting of inflation.

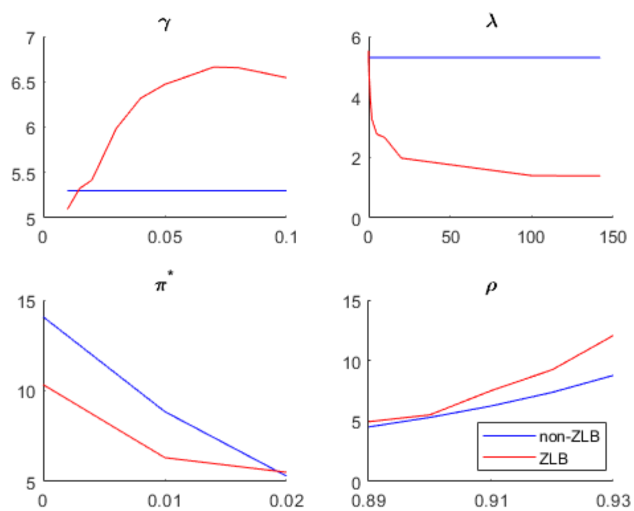

Figure 26: The precentage of periods in which ZLB binds.

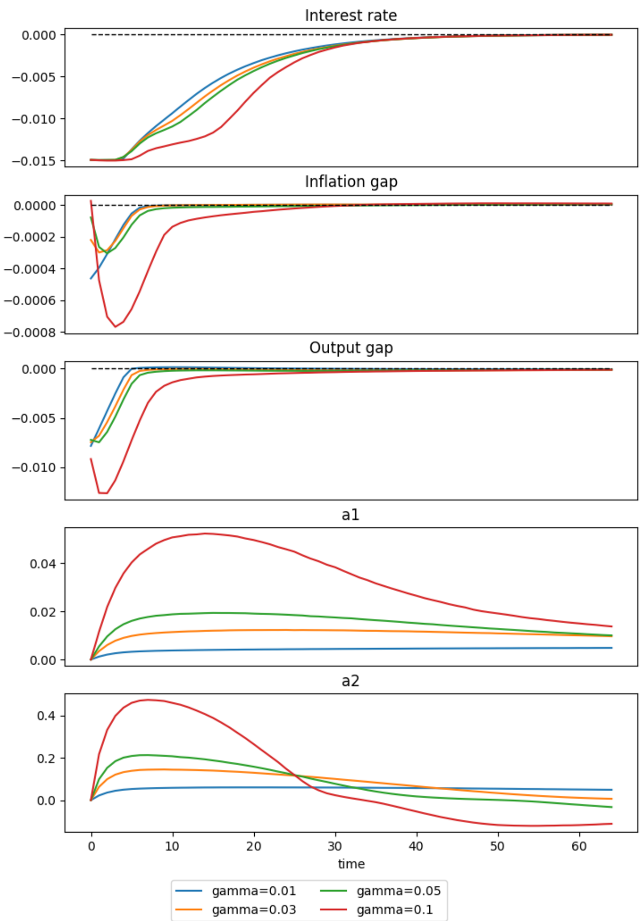

Figure 27: Generalized impulse response for different γ.
Median responses only.

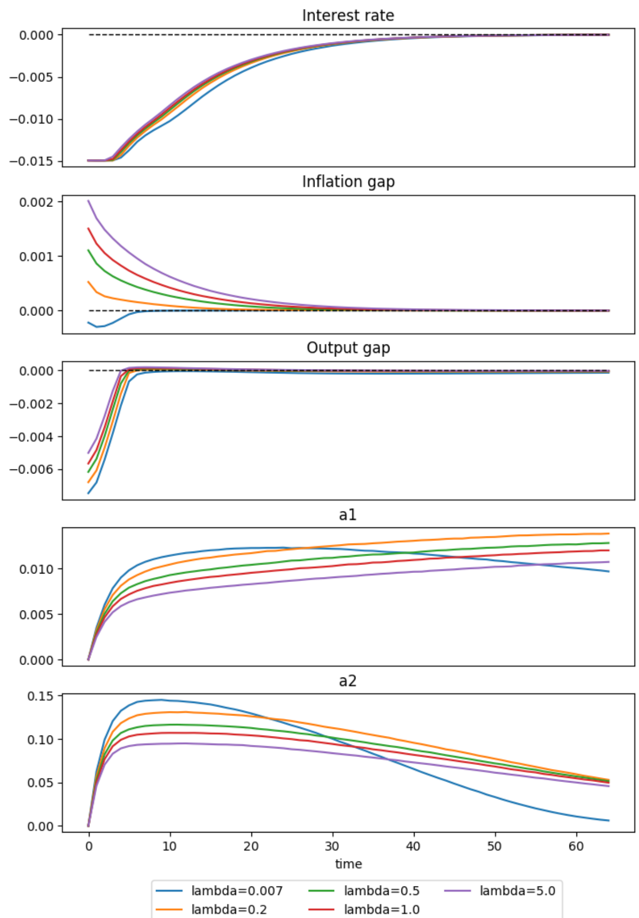

Figure 28: Generalized impulse response for different λ.
Median responses only.

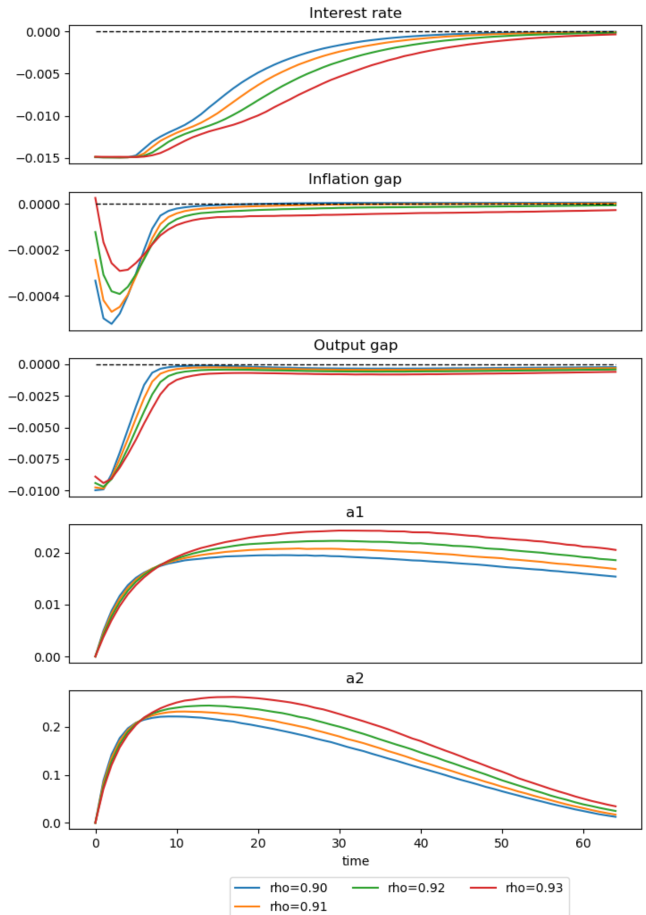

Figure 29: Generalized impulse response for different ρ.
Median responses only.

References

Adam, Klaus, and Roberto 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.

———. 2007. "Discretionary monetary policy and the zero lower bound on nominal interest rates." *Journal of Monetary Economics* 54:728–752.

Ascari, Guido, and Argia M Sbordone. 2014. "The Macroeconomics of Trend Inflation." *Journal of Economic Literature* 52 (3): 679–739.

Benhabib, Jess, Stephanie Schmitt-Grohé, and Martín Uribe. 2001. "The Perils of Taylor Rules." *Journal of Economic Theory* 96 (1): 40–69.

Bray, Margaret. 1982. "Learning, estimation, and the stability of rational expectations." *Journal of Economic Theory* 26 (2): 318–339.

Calvo, Guillermo A. 1983. "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics* 12, no. 3 (September): 383–398.

Clarida, Richard, Jordi Gali, and Mark Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37, no. 4 (December): 1661–1707.

de Boor, Carl. 1978. *A Practical Guide to Splines.,* 27:XVIII, 348. New York: Springer-Verlag New York, January. ISBN: 978-0-387-95366-3.

Druedahl, Jeppe, and Thomas Hogholm Jorgensen. 2017. "A general endogenous grid method for multi-dimensional models with non-convexities and constraints." *Journal of Economic Dynamics and Control* 74:87–107.

Eggertsson, Gauti, and Michael Woodford. 2003. "The Zero Bound on Interest Rates and Optimal." *Brookings Papers on Economic Activity,* no. 1: 139–233.

Evans, George W., Eran Guse, and Seppo Honkapohja. 2008. "Liquidity traps, learning and stagnation." *European Economic Review* 52 (8): 1438–1463.

Evans, George W., and Seppo Honkapohja. 2001. *Learning and Expectations in Macroeconomics.* Princeton University Press. ISBN: 9780691049212.

———. 2005. "Policy interaction, expectations and the liquidity trap." Monetary Policy and Learning, *Review of Economic Dynamics* 8 (2): 303–323.

Farhi, Emmanuel, and Iván Werning. 2017. *Monetary Policy, Bounded Rationality, and Incomplete Markets.* NBER Working Papers 23281. National Bureau of Economic Research, Inc, March.

Fernández-Villaverde, Jesús, Grey Gordon, Pablo Guerrón-Quintana, and Juan F. Rubio-Ramírez. 2015. "Nonlinear adventures at the zero lower bound." *Journal of Economic Dynamics and Control* 57:182–204.

Gallant, A Ronald, Peter E Rossi, and George Tauchen. 1993. "Nonlinear Dynamic Structures." *Econometrica* 61 (4): 871–907.

Gaspar, Vitor, Frank Smets, and David Vestin. 2006. "Adaptive Learning, Persistence, and Optimal Monetary Policy." *Journal of the European Economic Association* 4, nos. 2-3 (May): 376–385.

Gavin, William T., Benjamin D. Keen, Alexander W. Richter, and Nathaniel A. Throckmorton. 2015. "The zero lower bound, the dual mandate, and unconventional dynamics." *Journal of Economic Dynamics and Control* 55:14–38.

Gourieroux, Christian, and Joann Jasiak. 2005. "Nonlinear Innovations and Impulse Responses with Application to VaR Sensitivity." *Annales d'Économie et de Statistique,* no. 78: 1–31.

Iovino, Luigi, and Dmitriy Sergeyev. 2017. *Central Bank Balance Sheet Policies without Rational Expectations.* Unpublished. December.

Judd, Kenneth L., Lilia Maliar, Serguei Maliar, and Inna Tsener. 2017. "How to solve dynamic stochastic models computing expectations just once." *Quantitative Economics* 8, no. 3 (November): 851–893.

Judd, Kenneth L., Lilia Maliar, Serguei Maliar, and Rafael Valero. 2014. "Smolyak Method for Solving Dynamic Economic Models: Lagrange Interpolation, Anisotropic Grid and Adaptive Domain." *Journal of Economic Dynamics and Control* 44:92–123.

Koop, Gary, Hashem M Pesaran, and Simon M Potter. 1996. "Impulse response analysis in nonlinear multivariate models." *Journal of Econometrics* 74 (1): 119–147.

Kopecky, Karen A., and Richard M H Suen. 2010. "Finite state Markov-chain approximations to highly persistent processes." *Review of Economic Dynamics* 13 (3): 701–714.

Krueger, Dirk, and Felix Kubler. 2004. "Computing equilibrium in OLG models with stochastic production." *Journal of Economic Dynamics and Control* 28, no. 7 (April): 1411–1436.

Krugman, Paul R. 1998. "It's Baaack: Japan's Slump and the Return of the Liquidity Trap." *Brookings Papers on Economic Activity* 29 (2): 137–206.

Marcet, Albert, and Thomas J. Sargent. 1989. "Convergence of least squares learning mechanisms in self-referential linear stochastic models." *Journal of Economic Theory* 48 (2): 337–368.

Molnár, Krisztina, and Sergio Santoro. 2014. "Optimal monetary policy when agents are learning." *European Economic Review* 66:39–62.

Orphanides, Athanasios, and John C. Williams. 2008. "Learning, expectations formation, and the pitfalls of optimal control monetary policy." *Journal of Monetary Economics* 55 (October): S80–S96.

Potter, Simon M. 2000. "Nonlinear Impulse Response Functions." *Journal of Economic Dynamics & Control* 24:1425–1446.

Press, William H., Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. 1986. *Numerical Recipes: The Art of Scientific Computing.* Princeton University Press. ISBN: 0521308119.

Rotemberg, Julio J. 1982. "Monopolistic Price Adjustment and Aggregate Output." *The Review of Economic Studies* 49 (4): 517–531.

Rotemberg, Julio J., and Michael Woodford. 1998. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy: Expanded Version."

Sims, Christopher A. 1980. "Macroeconomics and Reality." *Econometrica* 48, no. 1 (January): 1–48.

Tauchen, George. 1986. "Finite state markov-chain approximations to univariate and vector autoregressions." *Economics Letters* 20 (2): 177–181.

Taylor, John B. 1979. "Estimation and Control of a Macroeconomic Model with Rational Expectations." *Econometrica* 47, no. 5 (September): 1267–1286.

Werning, Ivan. 2011. "Managing a Liquidity Trap: Monetary and Fiscal Policy."

Woodford, Michael. 2003. *Interest and Prices : Foundations of a Theory of Monetary Policy.* 804. Princeton University Press. ISBN: 9781400830169.