The Past is Present:
Optimal Monetary Policy at the Effective Lower Bound

Fernando Duarte* and Benjamin Marrow†

September 2020

Abstract

We use a New Keynesian model with an effective lower bound (ELB) and a general stochastic process for the natural rate to study optimal monetary policy. The central bank has perfect commitment and an interest rate smoothing term in its loss function. Despite the ELB binding occasionally and endogenously, we can derive a closed-form solution for the optimal interest rate: it is the maximum of zero and a weighted average of all past realizations of the output gap. This implies that the optimal interest rate (i) takes a simple form, (ii) is path dependent at all times, (iii) should be pre-emptively lowered when close to the ELB –or kept at zero if at the ELB– if and only if the weighted average of past output gaps is negative, and (iv) behaves very differently from the Taylor rule. We illustrate these insights by solving for key variables in the New Keynesian model using a neural network.

1 Introduction

For more than a decade, short-term nominal interest rates in many advanced economies —–including Japan, the US and Europe—– have been against or close to their effective lower bound. There is ample awareness by both academics and policymakers that visits to the effective lower bound (ELB) may be more frequent and last longer than in the past. One of the main forces highlighted in the literature in driving interest rates towards their ELB are shocks to the natural rate of interest, especially in an environment in which the overall level of the natural rate is already persistently low.

In this paper, we study how monetary policy should respond to shocks to the natural rate when constrained by the ELB. We use the simplest textbook version of the linearized New Keynesian model (NK), but augmented by adding an ELB and a preference for smooth interest rates for the central bank.

*Federal Reserve Bank of New York

†Booth School of Business, University of Chicago. email bmarrow@chicagobooth.edu. I am grateful to Rohan Kekre for advice. This research was funded in part by the John and Serena Liew Fellowship Fund at the Fama-Miller Center for Research in Finance and by the Center for Research in Security Prices, both located at the University of Chicago Booth School of Business. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

The only exogenous disturbance in the model are shocks to the natural rate. We use this NK model as our starting point not only because we can obtain transparent results, but also because it continues to be the most common workhorse model for macroeconomic analysis especially at central banks. We assume that the interest rate is the central bank's sole policy tool and abstract from unconventional policies that do not involve the interest rate, such as large-scale asset purchases. The central bank's preference for smooth interest rates is consistent with empirical evidence and can be microfounded in several ways, including from transaction frictions.[1]

We also simplify the analysis by assuming prices are fully rigid, with zero inflation throughout, which allows us to focus on how shocks to the natural rate affect the economy through their effect on the output gap as governed by the IS equation. In models without the ELB, the IS equation is usually not a binding constraint for the central bank, and all monetary policy tradeoffs stem from the Phillips Curve. Instead, with an ELB, the IS equation can be binding. Indeed, for a given path of the output gap, the behavior of inflation is governed by the Phillips Curve and is the same with or without the ELB, and with fully or only partially rigid prices. Therefore, our strategy of considering fixed prices allows us to concentrate on what we consider is one of the primary differences between models with and without the ELB. Of course, studying inflation with partially rigid prices remains an important question that we hope can be informed by our results with fully rigid prices.

Despite the importance of the question and the copious research on how to deal with the ELB, an explicit characterization of how optimal policy responds to shocks to the natural rate has not yet been achieved with a high degree of generality. Analytical results usually abstract from uncertainty altogether, and numerical studies tend to be calibration-specific.

In this paper, We take a step towards greater generality by using a mathematical approach not yet used in the literature, the continuous time stochastic maximum principle. This maximum principle allows us to derive necessary and sufficient first order conditions (FOC) for the central bank's optimal policy problem under commitment. We derive these FOC using a generic Ito stochastic process for the natural rate. By the Martingale Representation Theorem, this means that we consider all measurable square integrable processes for the natural rate, a very general class.

Using the FOC and the maximum principle, we can derive a closed-form expression for the optimal interest rate as a function of the output gap, and a single backward stochastic differential equation (BSDE) for the output gap. A BSDE is similar to a (forward) stochastic differential equation, only that instead of the initial condition being given, it is the terminal value of the process that is prescribed, reflecting the forward-looking nature of the households' rational expectations. The BSDE for the output gap in our model is quite simple, with a drift that is piecewise linear as a function of the nominal rate and linear as a function of the natural rate, allowing for transparent interpretation. Coupled with the the (forward) stochastic differential equation for the natural rate, the FOC characterize optimal monetary policy as the solution to a system of forward-backward differential equations (FBSDEs).

We show how to solve this FBSDE using a neural network, an efficient solution method that essentially only requires the use of existing standard software packages (in this case, Tensorflow for Python). Neural

[1] Eggertsson and Woodford (2003), chapter 6.

networks offer a number of advantages over other approaches. Because neural networks can be scaled without much effort to hundreds of state variables, the method we use can be immediately applied to big-scale dynamic stochastic general equilibrium (DSGE). One implication of this approach is that it is straightforward to add additional states, like inflation, to the model. By contrast, certain closed-form approaches to continuous time optimal control with constraints, like the application of Tanaka's formula, cannot incorporate more than one state variable (Chen and Zhou, 2004).

Additionally, neural networks make it possible to solve for all endogenous variables not only as functions of the state variables, but also treating parameters as additional "state variables" of the neural network. In this way, after the neural network is trained, it provides all endogenous variables of the model as a function not only of the model's state variables, but also of the underlying parameters. The implication is that the solution of the model can be evaluated in a matter of seconds for any parameter values and any value of the state variables. Although not exactly an analytical solution, this strategy nevertheless significantly alleviates the limitations imposed by numerical solutions that rely on specific calibrations, as the entire parameter space can be explored at close to zero marginal cost. For larger models, adding parameters as inputs to the neural network may provide a substantial speed advantage when estimating parameters, since the solution to the model does not need to be re-computed every time a new parameter combination is tried. We illustrate this idea by making the initial level of the natural rate, one of the parameters of our model, an input to the neural network.

Armed with the closed-form expression for the optimal interest rate and the solution to the BSDE that governs the behavior of the output gap, we provide five results that are, to the best of our knowledge, new to the literature.

Our first result is that the optimal interest rate takes a simple form. It can be written as the maximum of zero and a weighted average of all past realizations of the output gap, with weights that discount past realizations exponentially at the same rate as the representative household's. It can be communicated in one sentence, and without any reference to the natural rate of interest.

Our second result is that optimal interest rates are path dependent at all times. Until now, only other elements of policy, such as the time-varying targets in targeting rules, had been shown to be path-dependent,[2] but not interest rates. Many authors consider rules for nominal interest rates that closely track either the natural rate (in a path-independent way) or a past average of inflation (in a path-dependent way). Instead, we show that tracking a weighted average of past output gaps can also be beneficial.

Our third result is that, whether a central bank close to the ELB should "keep its powder dry" (keep rates higher than otherwise so they can be lowered more in the future if needed, for example, to avoid hitting the ELB) or engage in "insurance cuts" (keep rates lower than otherwise to stimulate the economy now so as to minimize the probability of hitting the ELB) depends not on the current level of the natural rate or the output gap, but on the entire path of output gaps. Our model results imply that for the same level of the natural rate and for the same level of the output gap, the central bank may sometimes find it optimal to keep its powder dry, and do insurance cuts at other times.

[2]Eggertsson and Woodford (2003).

Fourth, there is considerable debate regarding the appropriate duration of forward guidance once at the ELB. In particular, results regarding how the length of the promise to keep nominal rates at zero should change when shocks hit the economy, or as a function of the parameters of the model, have not been thoroughly investigated. We provide a simple answer: interest rates should be kept at zero as long as the exponential weighted average of output gaps remains negative, and not any longer. In addition, given the weighted average nature of the interest rate, the expected "liftoff" date changes slowly over time, so that any one shock only has a muted impact on the optimal duration of forward guidance.

Fifth, we show that the the truncated Taylor rule (the maximum of zero and the interest rate prescribed by the Taylor rule) is a poor approximation to optimal policy for any choice of Taylor rule coefficients despite being fully optimal if the ELB were removed.

Finally, our paper serves as a proof-of-concept for the application of machine learning methods to the analysis of continuous-time, stochastic, non-linear New Keynesian models. Given the burgeoning research on ELB-related phenomena, our paper provides a clear working example of an important addition to the toolkit available to economists.

The remainder of the paper is structured as follows. Section 2 discusses related literature. Section 3 formulates the model. Section 4 present our computational approach using neural networks. Section 5 discusses the results and provides intuition. Section 6 concludes.

2 Related literature

Our paper adds to research on monetary policy at the effective lower bound. Previous analytic results typically assume a combination of: very simple processes for the natural rate such as two-state Markov processes, deterministic, steady-state-only or perfect foresight economies, or that the natural rate reverts back to its steady-state level in a deterministic manner and stays there forever. For example, the seminal work of Eggertsson and Woodford (2003) studies a process for the natural rate of interest that is unexpectedly negative in period 0 and stays constant unless it reverts back to its steady-state level forever, which happens with a fixed probability in every period. Werning (2011) considers a deterministic economy in which the natural rate goes from a constant negative value to its steady state level at a known future time. This allows him to employ a deterministic version of the maximum principle to solve for the optimal interest rate in closed-form. Other papers partially or fully relax these simplifying assumptions, but must then resort to numerical methods on calibrated models, or to the study of exogenously specified monetary policy rules (rather than optimal policy). Although a numerical approach can be very insightful, it is ill-suited to answer questions with high degree of generality. For any given set of numerical results, one can always wonder whether they are robust to a different specification of the exogenous processes, or under all reasonable parameter combinations of the model.

Adam and Billi (2006) develop a numerical algorithm to solve the NK model by adding two state variables that allow them to re-write the non-recursive problem of the central planner in recursive form, so that standard dynamic programming techniques can be used. Their main findings are that, first, "nominal interest rates may have to be lowered more aggressively in response to shocks than what is instead suggested by a model without lower bound" and, second, that "the presence of shocks that lead to zero nominal rates alters also the optimal policy response to non-binding shocks". Eggertsson *et al.* (2020) develop a computationally efficient numerical algorithm to solve for models with an ELB under the special assumption that the underlying shock process is a two-state Markov process with an absorbing state, similar to Eggertsson and Woodford (2003). They find that "previously suggested policy rules – such as price level targeting and nominal GDP targeting – do not perform well when there is a small drop in the price level, as observed during the Great Recession, because they do not imply sufficiently strong commitment to low future interest rates". Other papers use simplifying assumptions to obtain analytical answers. Mertens and Williams (2019) consider a version of the NK model with one demand shock and one supply shock, both of which follow uniform, iid, distributions. Rather than studying the optimal policy, they compare outcomes across different exogenously given monetary policy rules. The combination of uniformly distributed iid shocks and exogenous rules provides enough tractability to derive solutions in closed form. They find that average-inflation targeting and price-level targeting can both reduce the adverse effects of the lower bound on the economy.

Bilbiie (2019) similarly considers a two-state Markov process for the natural rate in a discrete time, infinite horizon NK model. In particular, the natural rate takes on one of two values, and when it attains its steady state (positive) value, it remains there in perpetuity. The paper illustrates the path-dependence of monetary policy by providing a closed-form solution for optimal forward guidance, i.e., the length of time that interest rates ought to remain at zero after a period of negative natural rates. He approximates this policy rule as half the duration of the liquidity trap times the "disruption" (the depth of the fall in the natural rate below 0).

Nakata and Schmidt (2019) consider a discrete-time stochastic NK model where the natural rate follows an $AR(1)$ process. They estimate this model numerically to show that incorporating a smoothing term for the interest rate, even if not justified by the microfoundations of the model, reduces the central bank's losses for a central bank without commitment. The induced preference for gradual changes in interest rates makes it easier to have longer forward guidance periods that mimic commitment. They solve for the optimal numerical weights on the smoothing term in the central bank's loss function taking their model calibration as given.

Regarding the debate of whether interest rates should be cut pre-emptively when close to the ELB, the literature provides conflicting advice. Duarte *et al.* (2020) and Williams (2009), for example, argue for the insurance cut view, while Taylor (2017) argues for the alternate view.

Our paper is also related to a growing literature on machine learning and neural networks. While the primary applications of neural networks are to classification problems – the problem of correctly assigning inputs to output categories – recent work has applied them to more complex, non-discrete problems. In the vein of our study, previous papers have employed neural networks to solve ordinary differential equations (e.g., Malek and Shekari Beidokhti (2006)) and high-dimensional stochastic differential equations (Raissi (2018), Chen and Wan (2020)). A more recent literature has applied neural networks specifically to the solution of difference equations that characterize state dynamics in economic models. Ashwin (2020) uses neural networks to classify the determinacy of rational expectations equilibria in a discrete, stochastic New-Keynesian model with a Taylor rule. By having agents "learn" (i.e., form expectations of the state process) using neural networks, he analyzes which equilibria yield explosive expectations. Fernández-Villaverde *et al.* (2019), meanwhile, applies neural network approximation to an infinite-dimensional heterogeneous agent model, to solve for agents' conditional expectation of the perceived law of motion of aggregate debt. Our paper illustrates another example of how machine learning methods can be efficiently used to solve a large class of multi-state, non-linear, stochastic model with occasionally binding constraints that are common to monetary policy and analysis.

3 A Benchmark New Keynesian Model with a ELB

Our starting point is the standard textbook version of the New Keynesian (NK) model of Galí (2015) and Woodford (2011). Households have risk averse utility over differentiated products and supply labor to an intermediate goods producing sector. Intermediate goods have a constant returns to scale technology with exogenous productivity and labor as only input. These intermediate goods producing firms maximize profits subject to a demand curve for differentiated products and Calvo style price stickiness. Their output is sold to the final goods producers in a monopolistically competitive way. Government spending, transfers, and taxes are all equal to zero. The central bank sets the short-term (instantaneous) nominal interest rate by paying interest on base money in the cashless limit. For simplicity, we assume nominal prices are fully fixed, so inflation is constant and identically zero.

Time is continuous and indexed by $t \in [0, T]$. There is a single source of uncertainty modeled as a one-dimensional Brownian motion $\{W_t, 0 \leq t \leq T\}$ defined on a probability space $(\Omega, \mathcal{F}, P)$. We let $\{\mathcal{F}_t, 0 \leq t \leq T\}$ be the natural filtration of $\{W_t\}$, where $\mathcal{F}_0$ contains all the P-null sets of $\mathcal{F}$, and denote the conditional expectation $E[\cdot \mid \mathcal{F}_t]$ by $E_t[\cdot]$. The Brownian motion W_t will represent shocks to the natural rate of interest, defined below.

We study the problem of a central bank that chooses the path of short-term nominal interest rates i_t in order to minimize a quadratic approximation to the social welfare function, subject to the equilibrium conditions of private sector optimality and market clearing. We assume that the central bank has a perfect commitment technology and commits to implementing its optimal policy from the point of view of $t = 0$. The central bank solves:

$$\min_{\{i_t\}_{t=0}^{T}} \frac{1}{2} E_0 \int_0^T e^{-\beta t} \left(x_t^2 + \alpha i_t^2 \right) dt \tag{1}$$

$$dx_t = \frac{1}{\sigma} \left(i_t - r_t \right) dt + z_t dW_t \tag{2}$$

$$x_T = 0 \tag{3}$$

$$dr_t = \mu \left(t, r_t \right) dt + v \left(t, r_t \right) dW_t \tag{4}$$

$$r_0 = r \tag{5}$$

$$i \geq 0 \tag{6}$$

The variable x_t is the output gap, the log-deviation of output from the hypothetical output that would prevail in the flexible price, efficient allocation. Henceforth, for brevity, we refer to the output gap simply as output. The variable r_t is the exogenous natural rate of interest, defined as the real interest rate that would prevail in the flexible price, efficient economy with $x_t = 0$ for all t. The natural rate is stochastic because it reflects the stochastic productivity shocks to firms. The nominal interest rate i_t is constrained to remain above some effective lower bound that we take to be zero without loss of generality, as shown in equation (6).

Equation (1) states that, at time $t = 0$, the central bank minimizes a quadratic approximation of the true social welfare loss function by choosing the path of nominal interest rates $\{i_t, 0 \leq t \leq T\}$. The loss function penalizes expected deviations of output x_t and the nominal interest rate i_t from zero – the interest rate smoothing term– with future losses discounted at a rate $\beta > 0$. The constant $\alpha > 0$ measures the relative weight that the social welfare function places on deviations from zero in the interest rate relative to output. The central bank minimizes its loss function subject to two constraints. The first constraint in equation (2) is the IS curve or, equivalently, the representative agent's Euler equation (in discrete time, the IS equation is only an approximation of the Euler equation, but in continuous time the Euler equation is already linear). The constant $\sigma^{-1} > 0$ is the elasticity of intertemporal substitution of the representative agent, and $z_t dW_t$ is an expectational error term, with the variable z_t determined endogenously. Equation (3) is a terminal condition for output which states output must be zero at T. This terminal condition is the linearized counterpart of the optimality condition of the representative household that states that financial wealth must be zero at time T (having positive wealth at T is suboptimal, as it could have been consumed without changing any prior decisions). The second constraint, in equation (4), gives the stochastic process for the natural rate r_t. We assume that r_t follows an Ito diffusion with drift and volatility given by the functions $\mu(\cdot, \cdot)$ and $v(\cdot, \cdot)$, respectively. Equation (5) gives the initial condition for r_t, where r is a constant.

Mathematically, equation (2) is a backward stochastic differential equation (BSDE) with unknowns (x_t, z_t) and terminal condition given by equation (3). It reflects the forward-looking rational expectations nature of the Euler equation of the representative agent. It is called a backward equation because the terminal condition is given, and the equation must be solved backward from $t = T$, when the value of x_t is known, to $t = 0$. Equation (4) is a forward stochastic differential equation (FSDE) in the unknown r_t with initial condition given by (5). Therefore, the central bank solves a finite horizon stochastic control problem constrained by a system of forward-backward stochastic differential equations (FBSDE).

3.1 The Stochastic Maximum Principle

In this section, we derive the central bank's necessary and sufficient conditions for a solution to the central bank's optimization problem using the stochastic maximum principle.[3] We define the Hamiltonian by

$$H(t, r, x, z, i, \lambda, p, q) = \frac{1}{2} e^{-\beta t} \left(x^2 + \alpha i^2\right) - \frac{1}{\sigma}(i - r)\lambda + \mu(t, r)\, p + v(t, r)\, q \tag{7}$$

[3] There are many versions of the stochastic maximum principle. We use the one in Øksendal and Sulem (2015).

The adjoint processes are λ_t and (p_t, q_t), and satisfy

$$\begin{aligned} d\lambda_t &= \left.\frac{\partial H}{\partial x}\right|_t dt + \left.\frac{\partial H}{\partial z}\right|_t dW_t & (8) \\ \lambda_0 &= 0 & (9) \\ dp_t &= -\left.\frac{\partial H}{\partial r}\right|_t dt + q_t dW_t & (10) \\ p_T &= 0 & (11) \end{aligned}$$

where we use the shorthand notation $\left.\frac{\partial H}{\partial x}\right|_t$ to denote $\frac{\partial H}{\partial x}$ evaluated at $(t, r_t, x_t, z_t, i_t, \lambda_t, p_t, q_t)$, and similarly for the other derivatives. The system (8)-(11) is a system of FBSDE, with forward process λ_t and backward process (p_t, q_t). The adjoints λ_t, p_t and q_t have the economic interpretation of shadow prices analogous to Lagrange multipliers in discrete-time constrained optimization – λ_t is associated with the IS equation (2), and (p_t, q_t) with, respectively, the drift and stochastic part of the process for the natural rate in (4). The initial condition $\lambda_0 = 0$ in (9) reflects the intuition that the initial level of output is unconstrained, unlike the terminal value x_T, which is constrained to be zero. The terminal condition $p_T = 0$ in (11) reflects the intuition that at the terminal time T, the central bank is unconstrained by the dynamics of the natural rate, since at T they do not affect the future of the economy any longer.

Computing the derivatives in (8)-(11), we get

$$\begin{aligned} d\lambda_t &= e^{-\beta t} x_t dt & (12) \\ \lambda_0 &= 0 & (13) \\ dp_t &= -\left(\frac{1}{\sigma}\lambda_t + \frac{\partial \mu}{\partial r}(t, r_t)\, p_t + \frac{\partial \mu}{\partial v}(t, r_t)\, q_t\right) dt + q_t dW_t & (14) \\ p_T &= 0 & (15) \end{aligned}$$

We now minimize the Hamiltonian with respect to the control $i \geq 0$. If the solution is interior, the FOC is

$$0 = \frac{\partial H}{\partial i} = e^{-\beta t}\alpha i - \frac{1}{\sigma}\lambda$$

which gives

$$i = \frac{e^{\beta t}}{\alpha\sigma}\lambda$$

If the solution is a corner solution, $i = 0$. It follows that the optimal interest rate i_t^* satisfies

$$i_t^* = \max\left\{0, \frac{e^{\beta t}}{\alpha\sigma}\lambda_t\right\} \tag{16}$$

Using (16) in (2)-(5) and (12)-(15), we get the following system of FBSDE

$$
\begin{aligned}
\text{Forward SDEs} & : \\
dr_t &= -\rho_r\left(r_t-\mu_r\right)dt+\sigma_r dW_t \\
d\lambda_t &= e^{-\beta t}x_t dt \\
r_0 &= r \\
\lambda_0 &= 0 \\
\text{Backward SDEs} & : \\
dx_t &= \frac{1}{\sigma}\left(\frac{1}{\sigma\alpha}\max\left\{0,\frac{e^{\beta t}}{\alpha\sigma}\lambda_t\right\}-r_t\right)dt+z_t dW_t \\
dp_t &= -\left(\frac{1}{\sigma}\lambda_t+\frac{\partial\mu}{\partial r}\left(t,r_t\right)p_t+\frac{\partial\mu}{\partial v}\left(t,r_t\right)q_t\right)dt+q_t dW_t \\
x_T &= 0 \\
p_T &= 0
\end{aligned}
\tag{17}
$$

where a solution is a tuple $(r_t, \lambda_t, x_t, p_t, i_t, z_t, q_t)$ such that the system in (17) holds, and each variable is adapted to the filtration induced by (W_t, x_T, p_T) (i.e. does not use future information about W_t, x_t, p_t).

This system of FBSDEs are the necessary and sufficient conditions for optimality of the central bank's problem (1)-(6). We assume enough regularity conditions for the functions $\mu(\cdot,\cdot)$ and $v(\cdot,\cdot)$ such that the system (17) has a unique solution.[4]

The process for λ_t can be solved explicitly in terms of x_t

$$
\lambda_t = \int_0^t e^{-\beta s} x_s ds \tag{18}
$$

so that

$$
i_t^* = \frac{1}{\alpha\sigma}\max\left\{0,\int_0^t e^{-\beta(s-t)}x_s ds\right\} \tag{19}
$$

$$
dx_t = \frac{1}{\sigma}\left(\frac{1}{\sigma^2\alpha^2}\max\left\{0,\int_0^t e^{-\beta(s-t)}x_s ds\right\}-r_t\right)dt+z_t dW_t \tag{20}
$$

Equation (20) and its terminal condition $x_T = 0$ provide a full characterization of the solution to the central bank's problem. If x_t is known, then i_t^* can be determined using equation (19), λ_t using equation (18), and (p_t, q_t) using equation (14). We solve equation (20) numerically using a neural network.

4 Neural Networks

In this section we provide an overview of neural networks and frame the solution to the FBSDE from (17) as a supervised learning problem.

[4]For example, they are uniformly Lipschitz .

4.1 Neural Network Overview

Neural networks refer to a class of semi-parametric models used to approximate nonlinear functional relationships, and more generally map input data to a given output. The theoretical basis for such networks stems from the Universal Approximation Theorem, which holds that any measurable function can be approximated arbitrarily well by some feed-forward (e.g. neural network) architecture (Cybenko (1989)).

The basic structure of a neural network involves of a series of "layers", a collection of nodes known as "neurons" that constitute each layer, and an architecture that determines how the neurons connect across layers. An initial layer inherits the input data and transforms it for input into the second layer; each successive layer operates on the outputs (neurons) from previous layer until the final layer returns a predicted output. Since intermediate layers are learned, non-parametric compositions of previous layers' neurons, they have no readily available interpretation and are often known as "hidden layers."

Within each layer are a chosen number of quasi-affine functions called neurons. The network assigns each neuron a scaling coefficient ("weight"), a linear shift ("bias"), and a nonlinear activation function that transforms the weighted sum of neurons from the antecedent layer. The connection between neurons is illustrated in Figure 1. The three neurons in layer n connect to neuron X^{n+1} in layer $n+1$. Given the weights (w_i) and biases (b_i) assigned to neurons X_i^n, neuron X^{n+1} takes a weighted sum of all neurons in the antecedent layer that fed into it, and applies the activation function $\sigma(\cdot)$. All neurons in layer $n+1$ will then feed forward into neurons in layer $n+2$ in a similar fashion, until the final layer is reached.

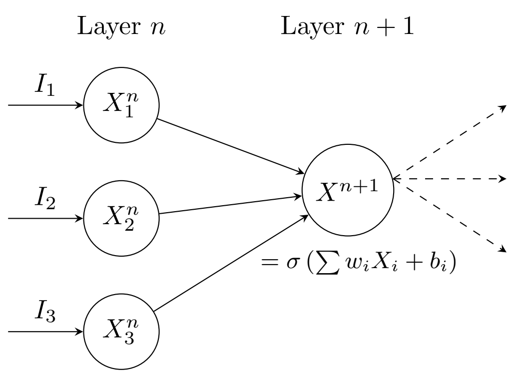

Figure 1: Connection Between Neurons

The architecture and hyperparameters of the model – the depth (in layers) of the network, the number of neurons in each layer, the structure of connections between neurons, and the activation functions – are fixed *ex ante*. Meanwhile, the weights and biases of each neuron are what is determined through model training. Training neural networks involves the definition of a loss function over the predicted output. Given the network-constructed output, the loss function computes the associated error, and the neural network updates the weights and biases of each neuron. One common approach (and the one used in this paper) is to update the parameters in the direction of the negative gradient, by computing the partial derivatives of the loss function with respect the parameters of each neuron, a process known as backpropagation. To reduce the computational burden of computing a full set of partial derivatives on each instance of training, it is also common to compute the loss and implement backpropagation using a randomly selected “batch” of the data in each iteration (stochastic gradient descent).

4.2 FBSDE Neural Networks

We construct our neural network as a variation of Raissi (2018). For a given system of FBSDEs, with $X_t \in \mathbb{R}^D$, $Y_t \in \mathbb{R}$, $Z_t \in \mathbb{R}^D$, $W_t \in \mathbb{R}^D$.

$$\begin{aligned} dX_t &= \mu\left(t, X_t, Y_t, Z_t\right) dt + \sigma\left(t, X_t, Y_t\right) dW_t, \qquad t \in [0, T] \\ X_0 &= \xi \\ dY_t &= \varphi\left(t, X_t, Y_t, Z_t\right) dt + Z_t' \sigma\left(t, X_t, Y_t\right) dW_t, \qquad t \in [0, T] \\ Y_T &= g\left(X_T\right) \end{aligned}$$

Raissi employs a 5-layer neural network to parametrize the quasi-linear partial differential equation, $u_t = f\left(t, x, u, Du, D^2u\right)$, with the property that

$$Y_t = u\left(t, X_t\right) \tag{21}$$

$$Z_t = Du\left(t, X_t\right) \tag{22}$$

In our neural network,

$$\begin{aligned} X_t &= \begin{pmatrix} r_t & \lambda_t \end{pmatrix}' \\ Y_t &= x_t \\ Z_t &= \begin{pmatrix} z_t & q_t \end{pmatrix}' \end{aligned}$$

A $D+2$-dimensional first layer accepts the D-dimensional vector of forward states (the output gap X_t), the time domain (t), and the initial condition for the natural rate (r_0). The subsequent four layers contain 40 neurons each, with sinusoidal activation functions. A final, one-dimensional layer returns the prediction of u_t. The backpropagation method employs the ADAM optimizer, with 100,000 total iterations at various learning rates.[5] While the input draws of the Brownian motion W_t are always random, we use stochastic gradient descent on the r_0 input by drawing random batches of r_0 from a uniform interval in each iteration of the training. Figure 2 provides a representation of a densely connected neural network similar to the one we use, but with only 10 neurons per layer (our model uses 40 neurons in each intermediate layer).

As in Raissi, the loss function is computed as the error to the Euler-Maryama approximation method, given the Z_t, Y_t implied by the current iteration of u_t. Recall that the Euler-Maryama scheme expresses

$$X^{n+1} \approx X^n + \mu\left(t^n, X^n, Y^n, Z^n\right) \Delta t^n + \sigma\left(t^n, X^n, Y^n\right) \Delta W^n \tag{23}$$

$$Y^{n+1} \approx Y^n + \varphi\left(t^n, X^n, Y^n, Z^n\right) \Delta t^n + \left(Z^n\right)' \sigma\left(t^n, X^n, Y^n\right) \Delta W^n \tag{24}$$

[5]Following Raissi, we use, sequentially, 20,000 iterations at learning rate $10e^{-3}$; 30,000 iterations at learning rate $10e^{-4}$; 30,000 iterations at learning rate $10e^{-5}$; and 20,000 iterations at learning rate $10e^{-6}$.

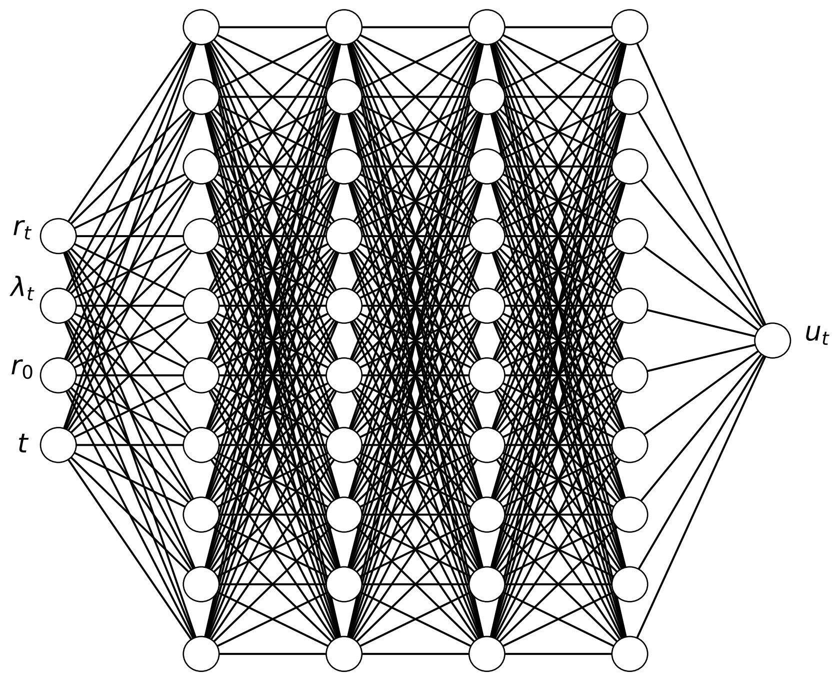

Figure 2: FBSDE Architecture (10 Neurons per Layer)

with

$$X_m^0 = \xi$$

The loss function is computed as

$$Loss = \sum_{k=1}^{K}\sum_{m=1}^{M}\sum_{n=0}^{N-1}\left|Y_{m,k}^{n+1} - Y_{m,k}^{n} - \Phi_{m,k}^{n}\Delta t^{n} - \left(Z_{m,k}^{n}\right)' \Sigma_{m,k}^{n}\Delta W_{m,k}^{n}\right|^2 + \sum_{k=1}^{K}\sum_{m=1}^{M}\left|Y_{n,k}^{N} - g\left(X_{m,k}^{N}\right)\right|^2,$$

where M is the batch size of the Brownian motion, K is the batch size of the initial condition r_0, N is the number of discretized time periods,[6] and

$$\Phi_{m,k}^{n} := \varphi\left(t^n, X_{m,k}^{n}, Y_{m,k}^{n}, Z_{m,k}^{n}\right)$$
$$\Sigma_{m,k}^{n} := \sigma\left(t^n, X_{m,k}^{n}, Y_{m,k}^{n}\right)$$

More specifically, we let

$$\begin{aligned}
\mu\left(t, X, Y, Z\right) &= \begin{bmatrix} -\rho_r\left(X_1 - \mu_r\right) \\ e^{-\beta t}X_2 \end{bmatrix} \\
\varphi\left(t, X, Y, Z\right) &= \frac{1}{\sigma}\left(\frac{1}{\sigma\alpha_i}\max\left\{0, \frac{e^{\beta t}}{\alpha\sigma}X_2\right\} - X_1\right) \\
\sigma\left(t, X, Y\right) &= \begin{bmatrix} \sigma_r & 0 \\ 0 & 0 \end{bmatrix} \\
g\left(X\right) &= 0
\end{aligned}$$

where X_1 and X_2 refer to the first and second entries of the vector $X = (r, \lambda)'$ respectively, and $\sigma(\cdot)$ is such that that the FBSDE loads on only a single Brownian motion.

Given current network estimate, $\hat{u}_t$, of $u_t = f(t, X, u, Du, D^2u, r_0)$, the steps to computing the loss are as follows. In each iteration,

1. Initialize model with $X_0 = (r_0, \lambda_0)$.
2. Compute $\hat{Y}_0$, $\hat{Z}_0$ by applying (21) and (22) to $\hat{u}_0$.
3. Construct $\tilde{X}_1$ and then $\tilde{Y}_1$ using Euler-Maryama method on inputs $\hat{Y}_0$, $\hat{Z}_0$, and W_0.
4. Directly compute $\hat{Y}_1$ by applying (21) to $\hat{u}_1$.
5. Add to loss $(\hat{Y}_1 - \tilde{Y}_1)^2$
6. Repeat steps (2)-(5) for remaining $N - 1$ time periods, letting $\hat{Y}_N = g(\hat{X}_n)$.
7. Repeat steps (1)-(6) for M paths of the Brownian motions and K draws of r_0 in the batch.

Steps 1 to 7 return a single scalar loss value. After the loss is computed, the optimizer computes the gradient of the loss with respect to the neurons, and updates the parametrization of $\hat{u}_t$ accordingly, by varying the weights and biases.

[6] We let $M = 100$ and $K = 12$, with the training interval for r_0 in $[-0.1, 0.1]$.

For the natural interest rate, we pick $\mu(t, r_t) = -\rho_r(r_t - \mu_r)$ and $v(t, r_t) = \sigma_r$, so that the interest rate follows an Ornstein–Uhlenbeck process, the continuous time analog of an AR(1) process. We use the following value for the model's parameters:

Parameter	Value	Description
β	0.02	Discount rate
α_i	1	Weight on interest rate gap
$1/\sigma$	1	Elasticity of intertemporal substitution
ρ_r	1	Mean reversion speed, r_t
μ_r	0.02	Long-run mean, r_t
σ_r	0.05	Stochastic volatility, r_t

Figure 3 illustrates the convergence of our neural network over its 100,000 iterations. To smooth out the noise induced by the randomness of the batches, we plot rolling averages over windows of two different sizes.

Figure 3: Neural Network Training Loss

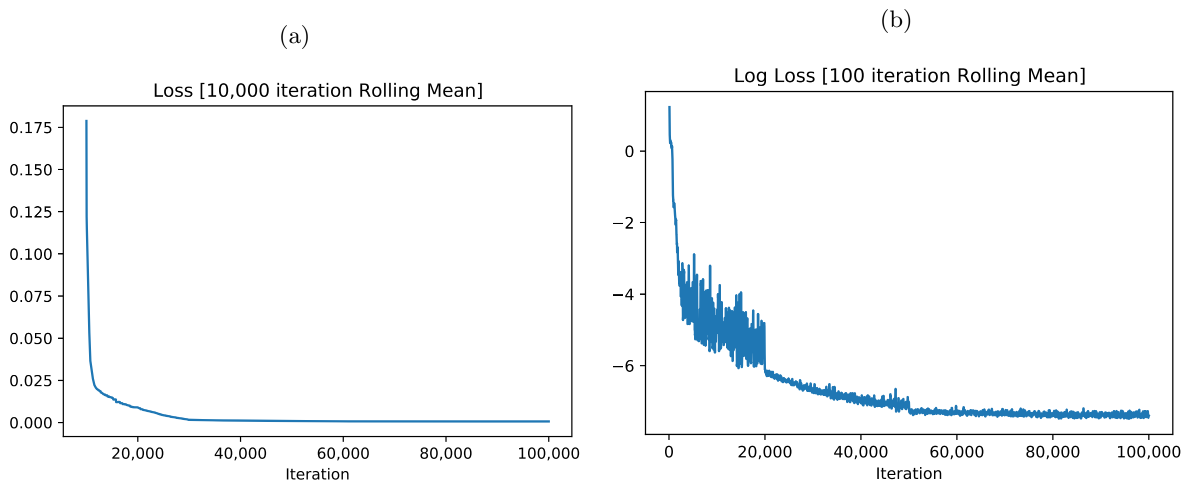

5 Optimal Monetary Policy

5.1 Intuition

Equation (18) shows that λ_t, the Lagrange multiplier associated with the IS equation, is an exponential moving average of past realizations of output x_t. If x_t is not identically zero for all t — a zero-probability event — then the IS equation is a binding constraint for the central bank. This is true at all times, even when the nominal rate is unconstrained, away from the ELB. There are two reasons. First, even without an ELB, the IS equation introduces a tradeoff between x_t^2 and i_t^2, the two terms in the objective function. The IS equation implies that in order to stabilize output and make the x_t^2 term small in the loss function, the central bank must change the interest rates more, which makes the i_t^2 term in the loss function higher. Second, mere possibility of nominal rates being constrained by the ELB in the future makes the IS equation binding today. This is in contrast to models without an interest smoothing term and without the ELB, in which the IS equation is never binding.

The optimal nominal interest rate is completely determined by λ_t, as can be seen from equation (16). When λ_t is positive, i_t^* is proportional to λ_t, with constant of proportionality $(\sigma\alpha)^{-1} e^{-\beta t}$. The intuition is that the central bank would like to equate the marginal benefit of relaxing the IS equation constraint, given by $\frac{1}{\sigma}\lambda_t$, with the marginal cost of changing i_t, given by $e^{-\beta t}\alpha i$. When λ_t is negative, the closest the central bank can get to equalizing marginal benefit and marginal cost is to set i_t^* to zero, as i_t^* it is constrained by the ELB. In a model without a ELB, i_t^* would be proportional to λ_t at all times. If, in addition to not having an ELB, there were no interest rate smoothing in the loss function, it would be optimal to track the natural rate and set $i_t = r_t$.

Combining equations (16) and (18) gives equation (19). The weighted average nature of i_t makes it smooth over time, consistent with the central bank's interest rate objective. Using (19) in the IS equation gives equation (20), which describes the evolution of output in terms of the level of the exogenous natural rate, the initial endogenous level of output x_0, and the endogenous volatility term z_t.

The left panel of Figure 4 plots four sample paths for z_t where the only difference between paths is that they use four different realizations of the path of the Brownian motion W_t. First, we note that all paths for z_t have a terminal value of $z_T = 0$. In order for X_t to satisfy its terminal condition $X_T = 0$ with certainty, it must be that $z_T = 0$, or else x_T would be stochastic. For values of t close to zero, z_t has higher values, reflecting the higher uncertainty that prevails early. Second, we note that the four paths are essentially identical, suggesting that z_t does not depend strongly on the realization of the Brownian motion.

The right panel of Figure 4 also plots four paths for z_t but, this time, the four paths were constructed by using the same path of W_t but four different values for r_0. We see that there is a small dependence of z_t on r_0, especially for small t. Overall, the two panels of Figure 4 imply that z_t is almost deterministic.

Figure 5 plots the relationship between the starting condition for the natural rate, r_0, and the endogenously determined initial condition for the output gap, x_0. Higher initial natural rates imply larger initial output gaps. Figure 6 is useful to understand why. We plot the dynamics of r_t, x_t, i_t and λ_t when $W_t = 0$ for all t. Each of the four lines in the panels only differ in that a different starting point for r_0 was chosen. All paths must have $x_T = 0$. Now consider the blue path, which has the highest r_0. Because r_t is persistent, r_t stays high and positive throughout. But, keeping i_t fixed, a high and positive r_t implies a *negative* growth rate of output: low real rates make the household save less and consume more today, which leads to lower expected consumption in the future, i.e., a negative expected growth rate for consumption and, since output equals consumption, a negative expected growth rate for x_t. Nominal interest rates start at zero and increase slowly, as can be seen in the bottom left panel of the Figure, so they do not increase enough to undo the reduction in real rates. The combination of a negative expected growth rate and a terminal condition of zero implies that x_0 must be high enough to hit zero after it has negative growth. As r_0 decreases, the incentive of households to save increases, increasing the growth rate of x_t.

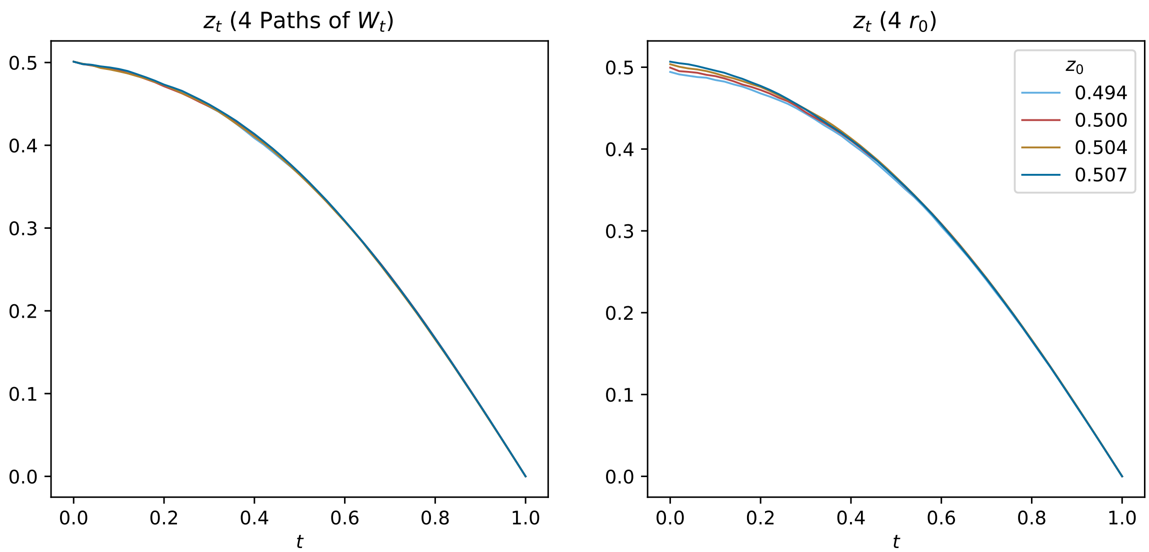

Figure 4: Endogenous volatility term z_t across paths of W_t and across r_0

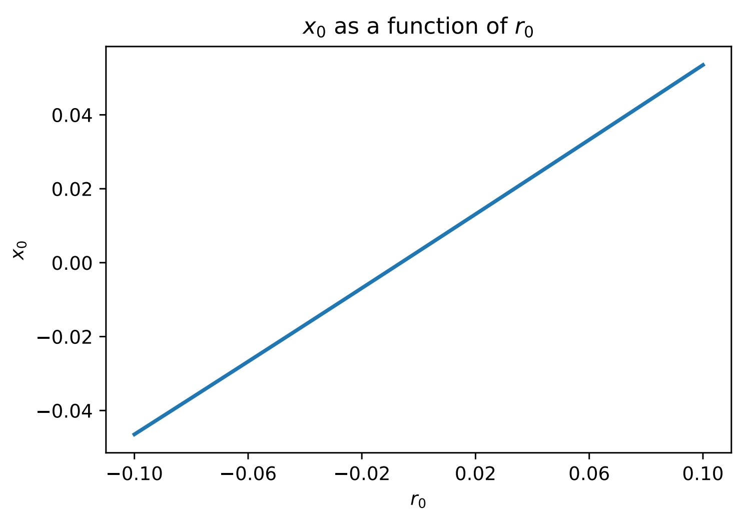

Figure 5: Initial output gap, x_0, as a function of initial natural rate r_0

The light blue line shows a path with the smallest (most negative) initial r_0, while the orange and red lines show intermediate values. The same logic applies to these paths: as we consider lower r_0, the incentive to save increases because the real rate increases. The smooth path for i_t does not reverse this effect. In fact, for the two paths with the lowest r_0, nominal rates are expected to be at the ELB at all times, so that the only changes in $i_t - r_t$ stem from r_t.

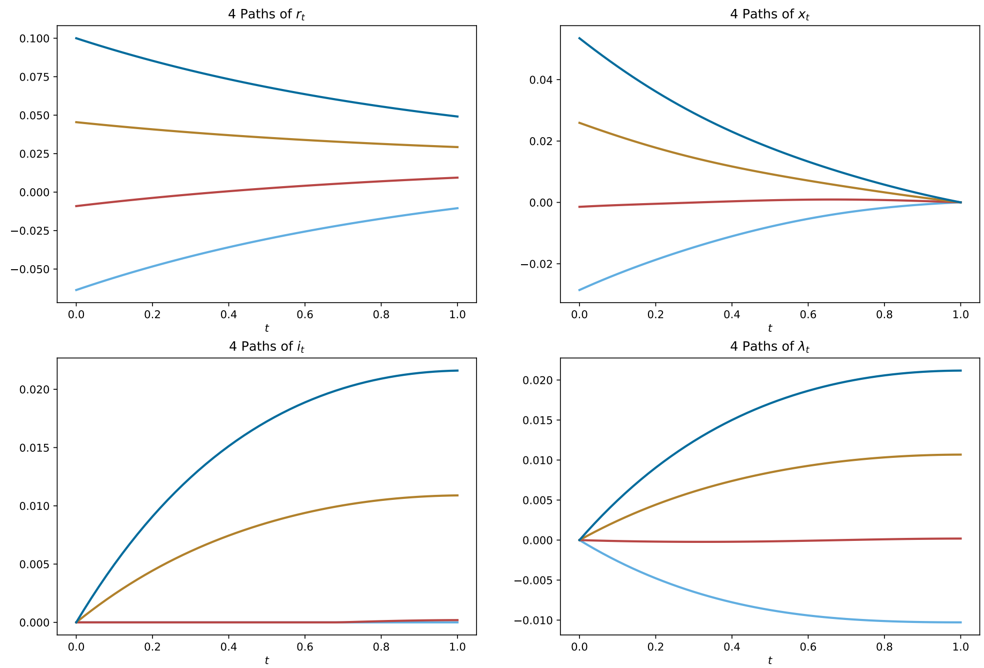

Figure 6: Model Paths: $W_t = 0, \ \forall t$

5.2 Impulse Response Functions

Figure 7 shows the impulse response functions (IRF) of x_t (first column), r_t (second column) and i_t (third column) to a shock of size 1 at $t = 0$. The first, second, and third rows use a value of r_0 of 0.01, -0.06, and -0.10, respectively. Because the model is nonlinear, different initial conditions can lead to different IRF.

To construct IRFs we first simulate 500,000 paths of x_t, r_t, and i_t for a given r_0, and take the median of each. We then compute the model-predicted path when $W_0 = 1$, and $W_t = 0, \ \forall t > 0$. The impulse response function is given as the "shocked" path less the median path.[7]

The IRF of the natural is the same in all cases, since the process for r_t is homoskedastic. The natural rate jumps at $t = 0$ and decays exponentially, consistent with the Ornstein-Uhlenbeck processes for r_t. The process for x_t is also essentially the same for all cases, a consequence of z_0 being essentially constant in r_0, as shown in Figure 4.

[7] An alternative definition of impulse response, where the benchmark is the simulated path when $W_t = 0, \ \forall t$ yields near identical results.

The nominal rate, however, shows very different behavior depending on the value of r_0. For $r_0 = 0.01$, the nominal rate is unconstrained and it smoothly increases, balancing the need to keep interest rates smooth while also stabilizing output by increasing i_t to gradually counter the increase in r_t

For r_0 sufficiently negative, as in the second row, interest rates do not immediately jump in response to the shock but stay 0 until the output gap has been positive for a long enough period (in this case until $t = 0.4$). This reflects the notion of forward guidance highlighted by Eggertsson and Woodford (2003), wherein the central bank maintains interest rates at the ELB longer than the natural rate states negative. When interest rates are very negative, as in the third row of Figure 7, the shock is not large enough to induce interest rates to escape their lower bound at any point in the interval.

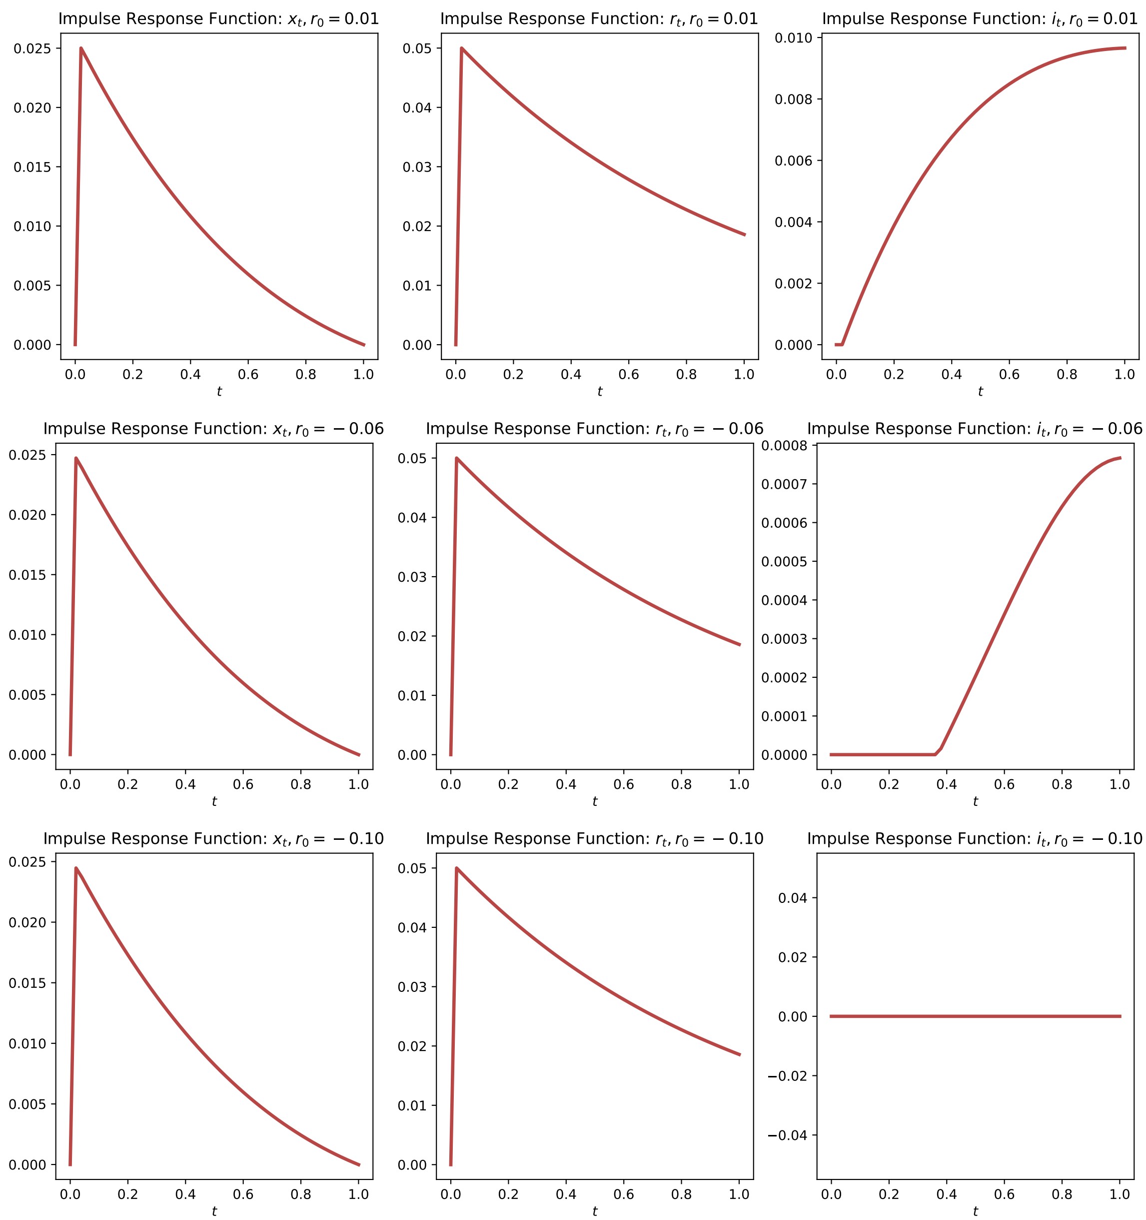

Figure 7: Impulse Response Functions

5.3 Forward Guidance

To understand forward guidance more deeply, and as a way to compare with the existing literature, we construct in Figure 8 the optimal policy when the natural rate follows a path similar to that in Werning (2011). That is, we input the precise Brownian motions such that r_t adheres to $r_\ell < 0$ until some time t_1, at which point it jumps up to $r_h > 0$ for the remainder of the interval. Of course, this jump is *not* anticipated in our model, unlike in 8 where everything is deterministic and known.

Figure 8 shows the optimality of forward guidance in his environment, with interest rates not rising when the natural rate returns to positive, but at some later date $t \approx 0.65$.

Figure 9 constructs an alternative scenario in which the natural rate is initially positive and later becomes negative, beginning at $r_h > 0$ and dropping unexpectedly to $r_\ell < 0$. As noted, interest rates do not drop immediately to 0, but gradually decrease. If the jump had been anticipated, as would be the case in a deterministic economy, the interest rate would have changed before the jump in the natural rate.

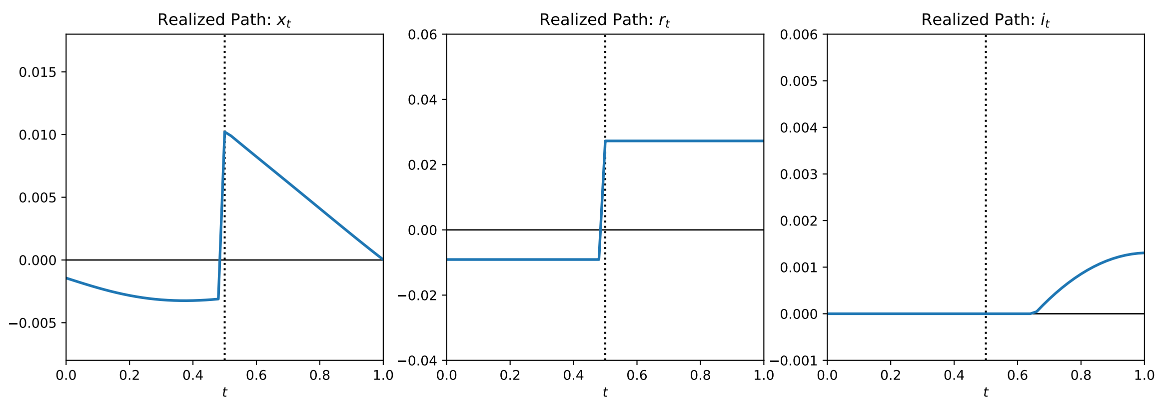

Figure 8: Realized Paths: r_ℓ Discrete Jump to r_h

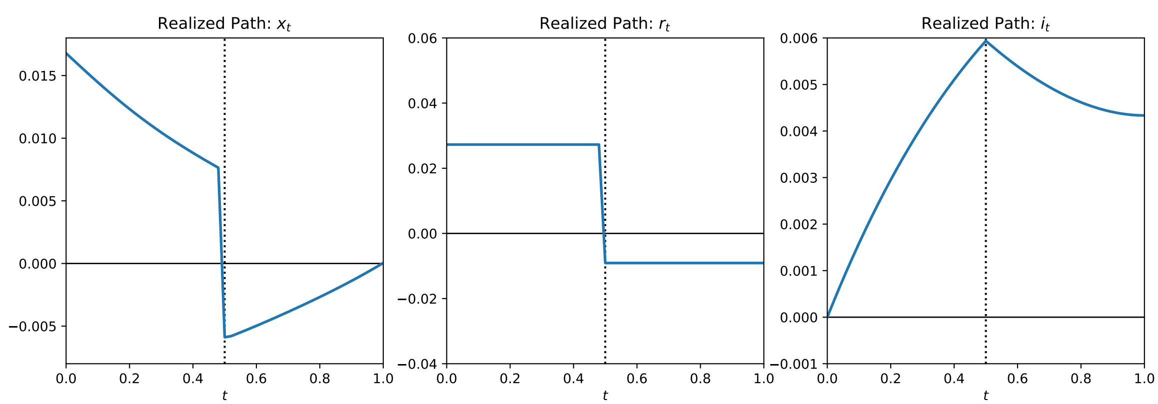

Figure 9: Realized Paths: r_h Discrete Jump to r_ℓ

5.4 Model Dynamics

The solution provided by the neural network allows us to compute the paths of variables for any desired path of the Brownian shocks. In this section, we illustrate how the optimal policy in response to various shocks depends on the path history.

Figures 10-12 chart several paths of the model variables from different initial conditions for r_0. When the natural rate starts out large, as in Figure 10, it trends down toward its long run mean of 0.02. Since natural rates tend to be positive along this path, the central bank raises interest rates, driving down the output gap x_t from its initial condition. When the natural rate starts out negative as in Figure 12, natural rates tend to remain negative along the path toward their long run mean, and so nominal rates adhere to their lower bound for the duration of the policy period. Figure 11 shows four paths of the variables when natural rates begin close to μ_r. The output gap closely mirrors the path of the natural rate, and interest rates exhibit persistent behavior, reflecting exponentially weighted averages of the output gap.

In all cases, the nominal rate i_t starts at 0. This is due to the fact that i_t is a linear transformation of λ_t. Since i_0 is free, optimality implies that $\lambda_0 = 0$, and so $i_0 = 0$ as well. The initial condition of the output gap, x_0, does however vary with model parameters.

5.5 Path Dependence

The dependence of λ_t on the entire past path of x_t introduces path-dependence in the optimal monetary policy and the resulting process for output. That the optimal policy is path-dependent is already documented in Eggertsson and Woodford (2003), Jung *et al.* (2005). There, the central bank chooses i_t so as to achieve an output-adjusted price level that is equal to a time-varying price-level target. If it is not possible to achieve the target, then the central bank chooses $i_t = 0$. In their implementation of optimal policy, the path dependence is reflected in the time-varying price targets, which are path-dependent because they are determined by the shortfalls in achieving the target in the last two periods. One remarkable aspect of this proposal is that it does not require any estimate or knowledge of the process for the natural rate, making it robust to the specification for the natural rate. How the central bank translates the price-level targeting policy into a particular value of the interest rate, and whether the resulting interest rate is path-dependent or depends on the natural rate, is left only implicitly characterized by a system of non-linear difference equations involving Lagrange multipliers and inequality constraints.

To observe the role that weighted averages of x_t play in the determination of optimal policy, we consider two model-generated dynamics. Figure 13 plots two paths of x_t that begin at the same r_0 but face difference realizations of Brownian shocks along the way. For $t < 0.6$, the blue-shaded line has primarily positive output gaps while the orange-shaded line exhibits negative output gaps. As a result, when r_t jumps from a positive to negative value near $t = 0.6$, the optimal policy differs. For the positive output gap path (blue), interest rates do *not* attain their lower bound, but rather stay positive, reflecting its weighted history of positive output gaps. Meanwhile, the negative path (orange) maintains interest rates at 0 even as r_t fluctuates between positive and negative values in the interval $t \in [0.5, 0.8]$. This behavior is consistent with its history of large negative output gaps for $t < 0.5$. A crucial, perhaps

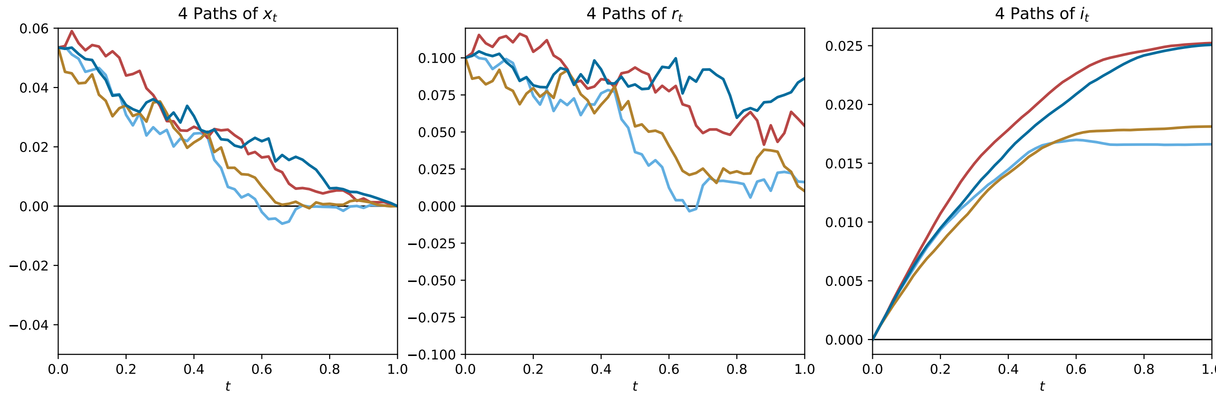

Figure 10: Model Paths: High r_0 ($r_0 = 0.10$)

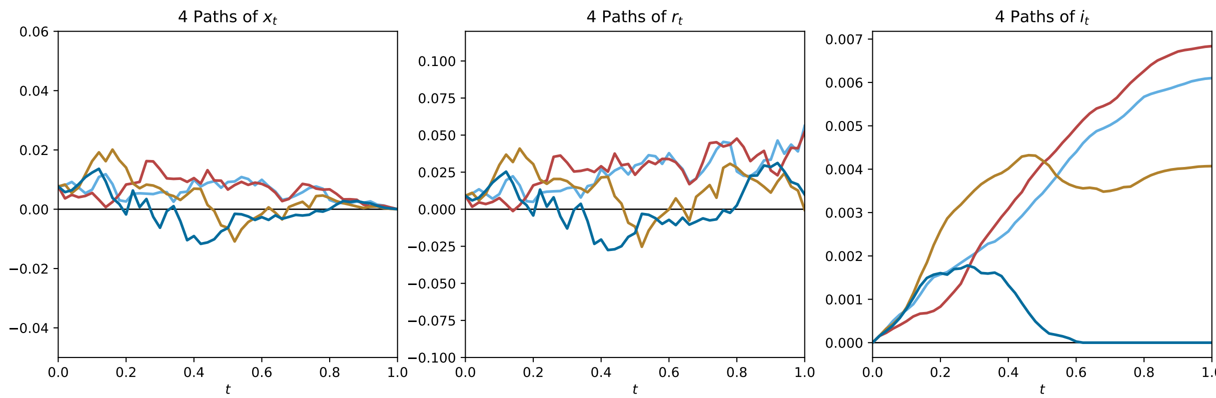

Figure 11: Model Paths: Medium r_0 ($r_0 = 0.01$)

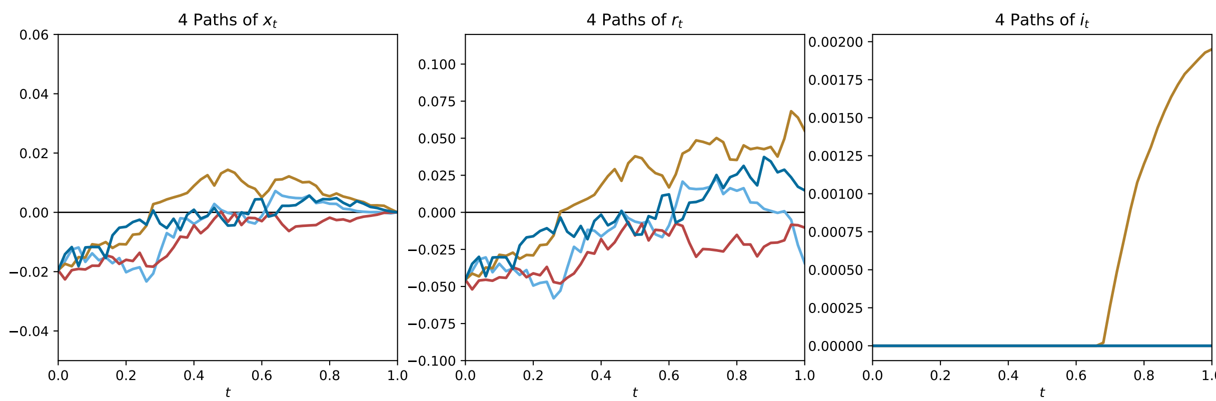

Figure 12: Model Paths: Low r_0 ($r_0 = -0.05$)

surprising, policy implication is that central banks ought not necessarily reduce interest rates to 0 when entering a liquidity trap. Rather, such a policy action is path dependent.

Similarly, Figure 14 shows how the history of x_t matters even when the shock to the natural rate is positive. In the two paths illustrated, both see an increase in r_t from approximately -0.01 around $t = 0.4$, to 0.02 around $t = 0.6$. But over this interval, only the path that featured a sufficiently positive output gap over its history – the blue-shaded line – escapes the lower bound in response to the positive shock to r_t. Moreover, in both figures, the *level* of r_t upon the realization of the shock is the same, confirming that the current level of the natural rate is not important to the determination of optimal policy

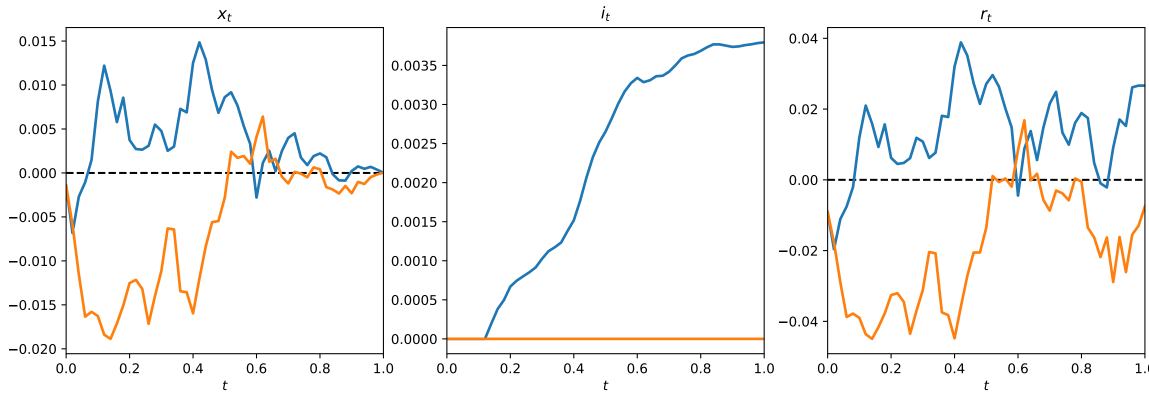

Figure 13: Interest Rates and the History of x_t: r_t falls below 0

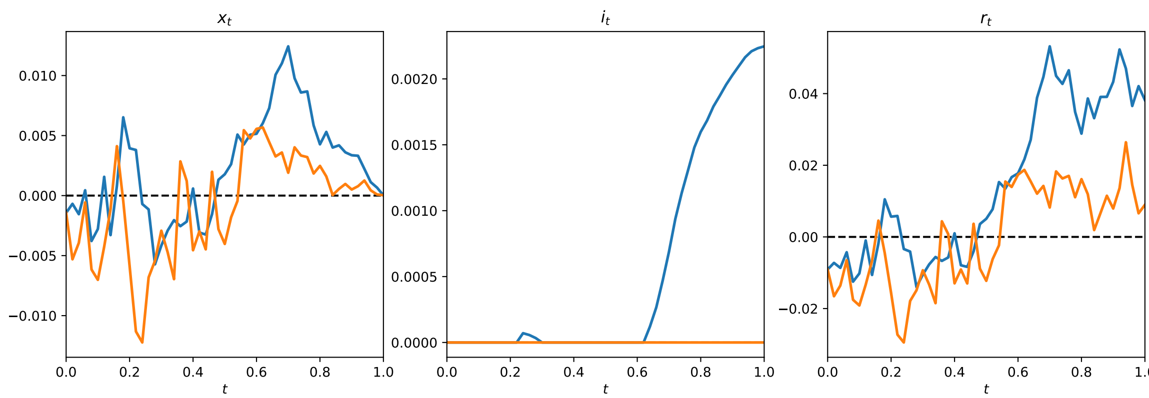

Figure 14: Interest Rates and the History of x_t: r_t rises above 0

Figures 15, 16, and 17 show how path-dependence influences optimal policy in a different way. In all three figures, the bottom panel shows simulated paths that use two different paths of the Brownian motion (blue and orange lines). These two Brownian motions are randomly drawn following the true distribution of the Brownian motion between times $t = 0$ and some t_1. At t_1, we set $W_{t_1} = 1$ and then, for $t > t_1$, we set $W_t = 0$. This exercise allows us to understand the effects of a shock on impact (at t_1) and the expected effects after impact (after t_1, when we shut down the shocks) after there is some meaningful history of the economy. The top panel in all figures plots the IRF of the t_1 shock from the point of view of t_1, that is, it plots the realized paths from the bottom panel minus its expectation conditional on time t_1- information. In particular, for $t < t_1$, the paths of all variables are known so the paths are all equal to zero.

In Figure 15, we pick the two paths of the Brownian motion such that r_0 and r_{t_1} are the same for both paths, but r_t between times 0 and t_1 is mostly positive for one path and mostly negative for the other path. Output and interest rates do not show path-dependence. However, nominal interest rates responds strongly at t_1 when the output gap has been mostly positive (blue line) while it does not respond at all when the output gap has been mostly negative (orange line).

We repeat the same exercise if Figure 16, but now pick different initial levels of the natural rate r_0. The blue lines use $r_0 > 0$ and the orange use $r_0 < 0$. Before t_1, the blue path is never against the ELB, while the orange path is always at the ELB. Even though both paths undergo the same shock at t_1, and have the same level of natural rates r_{t_1}, the orange path does not respond at all for an interval of time after t_1 (forward guidance), while the blue path responds immediately.

Last, Figure 17 considers a negative shock at t_1. After the shock, the orange path hits the ELB and stays there forever after, while the blue path, despite a larger reduction in the interest rate, only hits the ELB later on. This case illustrates how, depending on the path the economy has taken, it is sometimes better to lower interest rates towards the ELB immediately while it is better to leave room for further cuts, even when the shock was the same.

6 Conclusion

This paper characterizes optimal monetary policy with commitment in a textbook New Keynesian model that respects the ELB. While monetary policy at the ELB is an extensive area of research, ours is the first paper, to our knowledge, to explicitly address a continuous-time, stochastic model with general processes for the natural rate.

Using the stochastic maximum principle, we uncover several new insights into optimal monetary policy. The key feature of optimality – that interest rates depend on an exponentially-weighted average of past output gaps – formalizes many of the results anticipated in previous literature, including forward guidance, interest rate gradualism, and price-level targeting. It also provides a new way to communicate policy (in particular, how interest rates respond to shocks) in a simple manner that, crucially, does not need to make any reference to the natural rate. Among the novel implications of this result are that a central bank should *not* necessarily reduce rates to 0 when the natural rate turns negative, and that individual, idiosyncratic shocks to the economy have comparatively small effects on the optimal duration of forward guidance.

Our paper also demonstrates the use of neural networks for solving NK problems of this sort, and models with endogenously binding constraints more generally. Neural networks provide a fast and scalable method for solving and analyzing the model's solution without resorting to simplifications such as linearization. This approach thus allows a much broader set of policy questions to be answered in a quasi-analytic and computationally efficient manner.

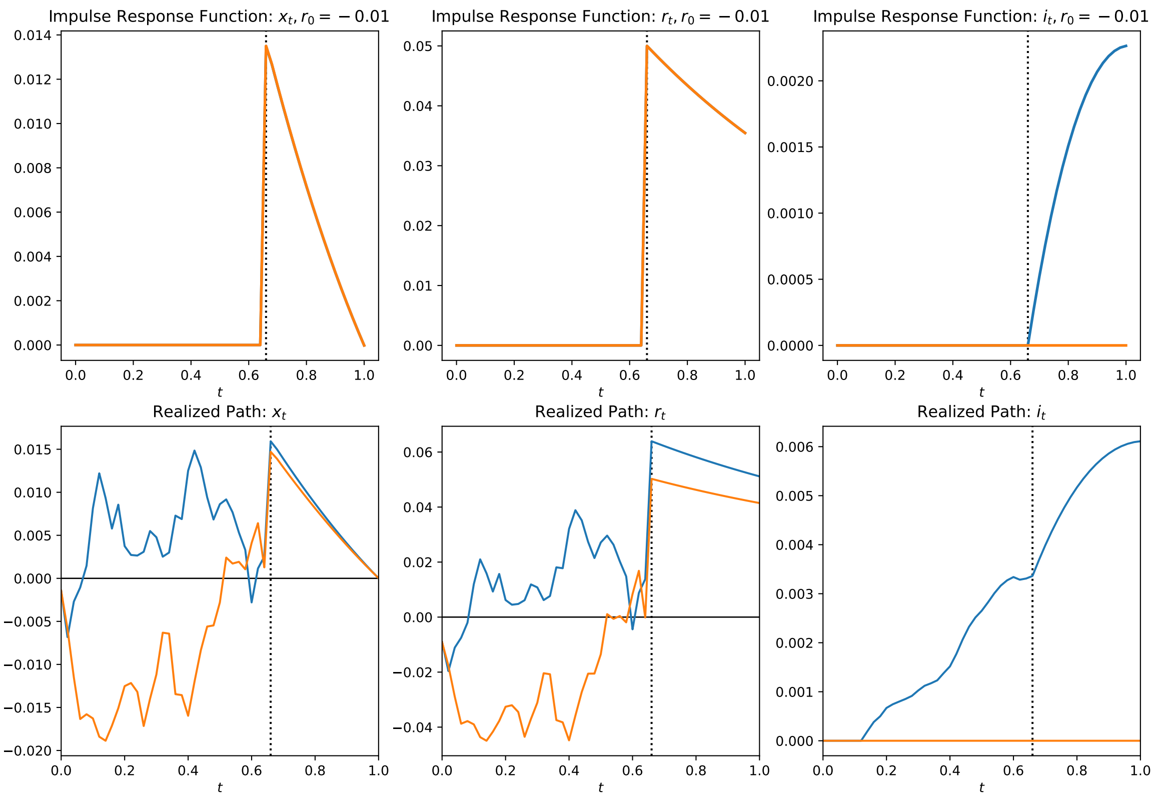

Figure 15: Impulse Response Functions: Positive shock to r_t with similar departure (r_0) and similar arrival (r_{t_1})

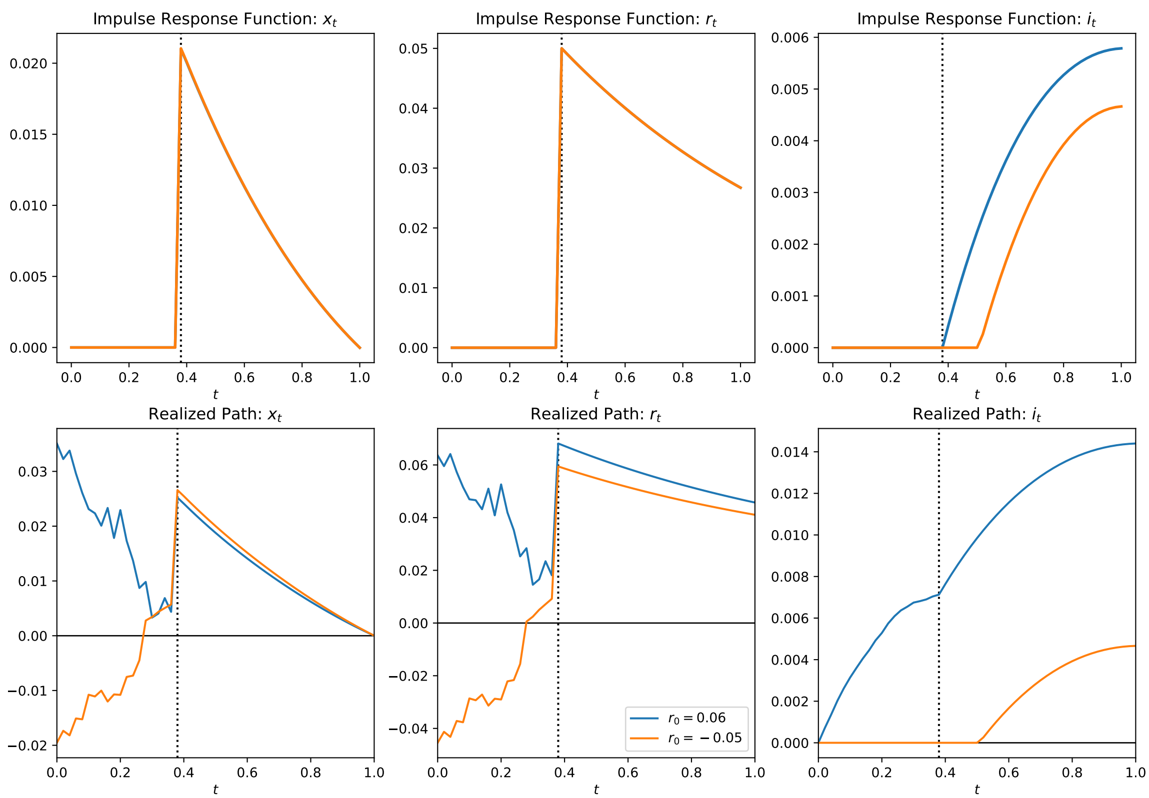

Figure 16: Impulse Response Functions: Positive shock to r_t with different departure (r_0) and similar arrival (r_{t_1})

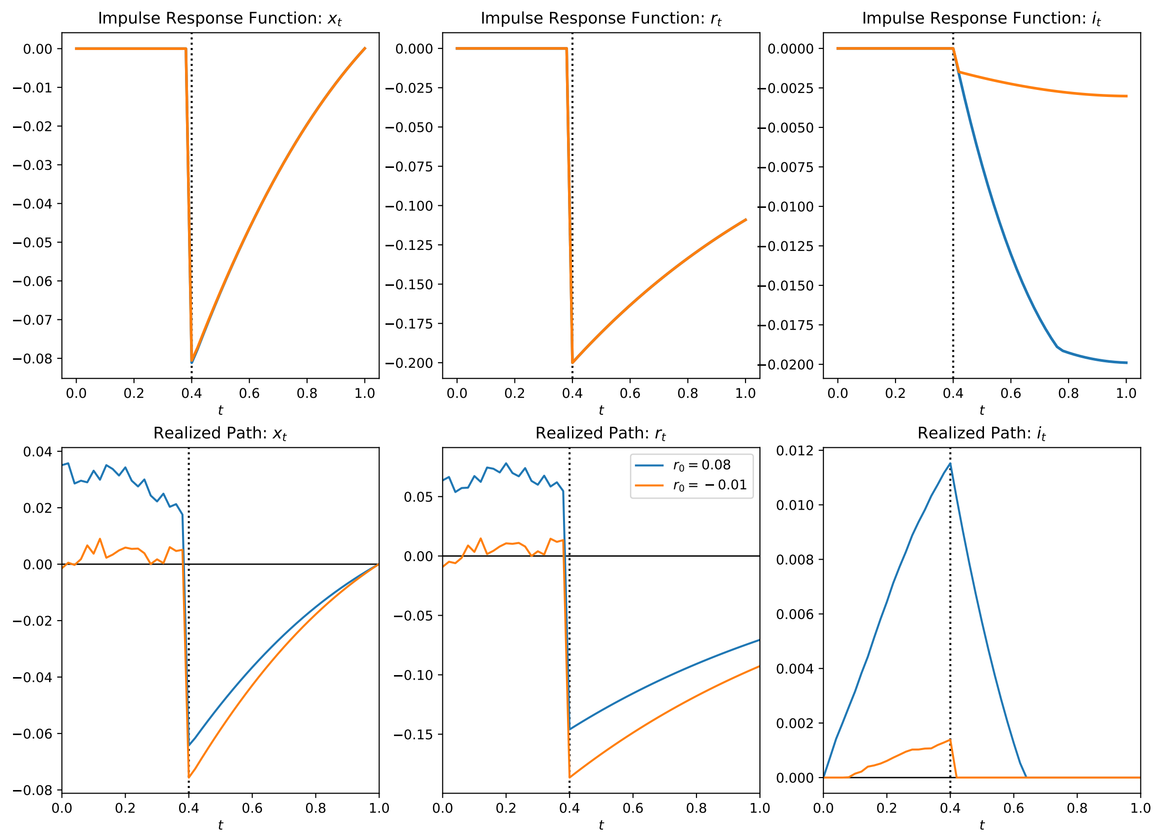

Figure 17: Impulse Response Functions: Negative shock to r_t with different departure (r_0) and arrival (r_{t_1})

While the model considered in this paper abstracted from price changes, future research would benefit from explicitly including more realistic nominal rigidities. Given the scalability of neural networks to additional states and dimensions, incorporating inflation as a state variable and a Phillips Curve as an additional constraint for the central bank is computationally straightforward. Other potential avenues for future research include changing the central bank's loss function to omit the smoothing term, or modifying the commitment technology to allow for more discretionary policy.

References

Klaus Adam and Roberto M. Billi. Optimal monetary policy under commitment with a zero bound on nominal interest rates. *Journal of Money, Credit and Banking*, 38(7):1877–1905, 2006.

Julian Ashwin. Resolving indeterminacy with neural network learning: sinks become sources, 2020.

Florin O. Bilbiie. Optimal forward guidance. *American Economic Journal: Macroeconomics*, 11(4):310–45, October 2019.

Yangang Chen and Justin W. L. Wan. Deep neural network framework based on backward stochastic differential equations for pricing and hedging american options in high dimensions. *Quantitative Finance*, 0(0):1–23, 2020.

Xi Chen and Xun Yu Zhou. Stochastic linear-quadratic control with conic control constraints on an infinite time horizon. *SIAM J. Control Optim.*, 43(3):1120–1150, March 2004.

George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

Fernando M. Duarte, Benjamin K. Johannsen, Leonardo Melosi, and Taisuke Nakata. Strengthening the FOMC's Framework in View of the Effective Lower Bound and Some Considerations Related to Time-Inconsistent Strategies. Finance and Economics Discussion Series 2020-067, Board of Governors of the Federal Reserve System (U.S.), August 2020.

Gauti B. Eggertsson and Michael Woodford. The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 2003(1):139–211, 2003.

Gauti Eggertsson, Sergey K. Egiev, Alessandro Lin, Josef Platzer, and Luca Riva. A Toolkit for Solving Models with a Lower Bound on Interest Rates of Stochastic Duration. Technical report, 2020.

Jesús Fernández-Villaverde, Samuel Hurtado, and Galo Nuño. Financial frictions and the wealth distribution. Working Paper 26302, National Bureau of Economic Research, September 2019.

Jordi Galí. *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press, 2015.

Taehun Jung, Yuki Teranishi, and Tsutomu Watanabe. Optimal monetary policy at the zero-interest-rate bound. *Journal of Money, Credit and Banking*, 37(5):813–835, 2005.

A. Malek and R. Shekari Beidokhti. Numerical solution for high order differential equations using a hybrid neural network—optimization method. *Applied Mathematics and Computation*, 183(1):260 – 271, 2006.

Thomas M. Mertens and John C. Williams. Monetary policy frameworks and the effective lower bound on interest rates. *AEA Papers and Proceedings*, 109:427–32, May 2019.

Taisuke Nakata and Sebastian Schmidt. Gradualism and liquidity traps. *Review of Economic Dynamics*, 31:182 – 199, 2019.

Bernt Øksendal and Agnes Sulem. Risk minimization in financial markets modeled by itô-lévy processes. *Afrika Matematika*, 26(5-6):939–979, 2015.

Maziar Raissi. Forward-backward stochastic neural networks: Deep learning of high-dimensional partial differential equations, 2018.

John B Taylor. Rules versus discretion: assessing the debate over the conduct of monetary policy. Technical report, National Bureau of Economic Research, 2017.

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

John C. Williams. Heeding Daedalus: Optimal Inflation and the Zero Lower Bound. *Brookings Papers on Economic Activity*, 40(2 (Fall)):1–49, 2009.

Michael Woodford. *Interest and prices: Foundations of a theory of monetary policy*. princeton university press, 2011.