DISCUSSION PAPERS IN ECONOMICS

Working Paper No. 21-02

Globally Optimal Monetary Policy

Christian Payne Hennigan
University of Colorado Boulder

October 13, 2021

Revised November 4, 2021

Revised November 9, 2021

Department of Economics

University of Colorado Boulder
Boulder, Colorado 80309

Globally Optimal Monetary Policy

Christian Payne Hennigan

University of Colorado Boulder

November 9, 2021

Latest Version

Abstract

Optimal monetary policy is important on a practical level for central banks. Computational barriers, however, have limited research into important optimal monetary policy questions. With some important exceptions using computational techniques tailored to specific cases, most of the applications addressed have been simplified to make them amenable to solving analytically or to easy computation. This paper brings into the optimal monetary policy literature recent machine learning techniques that apply to a wide range of practical applications for which computational barriers have previously been a problem. I illustrate these techniques as applied to the question of how firm expectations and price distortions should jointly influence optimal monetary policy. In a fully non-linear New Keynesian Model with price and labor distortions, I find that price level stabilization around a long-run value is best when distortions are small. However, the farther we start from the long-run value, the policy response should be more nonlinear and more aggressive. I show that interest rate policy should take into consideration both price dispersion and firm expectations on future costs, the latter directly relating to distortions from monopolistic competition.

JEL Classification: E52

Keywords: Monetary Policy, Interest Rate Policy, Macroeconomic Policy Formation, Computational Economics, Deep Learning, Machine Learning

1 Introduction

There is a long list of open questions in optimal monetary policy that matter to central banks on a practical level. For example, while both the Bank of Japan and Federal Reserve Board have both adopted average inflation targeting, the optimal design of this policy is unclear (Bank of Japan (2016), Federal Reserve Board (2020)). Beyond this, should optimal policy have inertia or adjust rapidly? How should policy account for non-normal uncertainty, for example with inflation expectations changing over time or being skewed?[1][2] Computational barriers have limited research into optimal monetary policy on questions like these.

With some important exceptions using computational techniques tailored to specific applications, most of the applications addressed have been simplified to make them amenable to solving by hand or to easy computation. In line with this, standard monetary policy research has traditionally used a framework of a quadratic objective function and linear dynamics, and it emphasizes a policy of stabilization in the face of small symmetric shocks.[3] Although research on optimal policy in the last two decades, particularly after the Great Recession, has expanded that framework by breaking one or two of the assumptions, these efforts have only resulted in highly case-specific approaches.[4]

This paper brings into the optimal monetary policy literature recent machine learning techniques that apply to a wide range of practical applications for which computational barriers have previously been important. The standard computation strategies in the literature that try

[1]More specifically, how forward guidance should take place given endogenous kurtosis of inflation expectations or uncertainty with regards to the model, or how non-additive uncertainty should influence policy.

[2]Preferences which respond to worst-case scenarios, rational inattention, or endogenous shock variances are also important modeling questions. See Mishkin (2010) for a more detailed list of monetary policy questions that are difficult to address by traditional techniques.

[3]The now-classic reference on this approach is Woodford (2003). See also Rotemberg and Woodford (1997); Clarida, Gali, and Gertler (1999); Giannoni and Woodford (2005); Levin, Onatski, and Williams (2005); and Schmitt-Grohe and Uribe (2005), among many others.

[4]An exhaustive literature review is impossible here, but see for example Kim and Ruge-Murcia (2019) for a model with nonlinearities and assymetric shocks, where higher order pertubation around the steady state is used; Curdia and Woodford (2015) for a modification the NK model to take credit frictions into account without fundamentally altering the LQ framework; Wu and Li (2014) for rational inattention implications in a log-linearized model; Swanson (2006) for monetary policy under parameter uncertainty in a linear model; Faulwasser, T. et al. (2020) for a nonlinear quadratic model used to study unconventional monetary policy; Nobay and Peel (2000) and Huh, Lee, and Lee (2009) for monetary policy in the context of a nonlinear Phillips curve; and Bilbiie and Ragot (2020) and Chale (2020) for an LQ framework modified for heterogenous agents and liquidity constraints, respectively.

to move beyond simplifications themselves often have fatal issues, like being limited in incorporating an effective lower bound and having a hard time dealing with anything beyond a small number of variables in the model.[5] Even when looking at a small-scale New Keynesian Model, the degree of the nonlinearities involved and number of variables makes standard procedures difficult to apply. Deep learning techniques are convenient because they can approximate any unknown function[6], are less sensitive to good initial guesses, and largely avoid the curse of dimensionality (Kang et al 2020, Raissi 2018). This is crucial for them to be practical for other applications like the ones mentioned at the outset, since a central bank may have no idea what the function they are searching for looks like and often deal with large models. Other papers have used neural networks to help solve other economic models,[7] but to my knowledge this is the first paper to make use of these techniques for optimal monetary policy research.

This paper illustrates these techniques as applied to the question of how firm expectations and price distortions should jointly influence optimal monetary policy. Inspired by the literature on optimal policy in the lineage of Ramsey (1927) and Stokey and Lucas (1983), I will have the central bank maximize household welfare, constrained by the optimal choices made by firms and households in a decentralized economy. This allows analysis beyond a small neighborhood of the steady state, in a way synergistic with the computational approach discussed later.[8] To think of how firm expectations and price distortions may influence optimal monetary policy, I will use a small-scale New Keynesian model, since it quite elegantly includes both nominal rigidity through price dispersion and imperfect competition in the form of price staggering by monopolistically competitive firms.

What happens to these two distortions in applications? Usually there is an assumption of a production subsidy provided through a lump sum tax in order to eliminate the mark-up distortion from imperfect competition. When this is done, monetary policy is concerned with the only

[5]The most glaring examples of these problems are higher order pertubation methods being unable to accommodate inequality constraints like an effective lower bound (Swanson et al 2006) and discretation, as well as projection methods, succumbing to sensitivity to initial conditions and the curse of dimensionality.

[6]More specifically, any Borel measurable function, by the universal approximation theorem (Bach 2017)

[7]Recent applications include Scheidegger and Bilionis (2017), Duarte (2018), Maliar et al. (2019), Ferandez-Villaverde et al. (2020), and Azinovic et al. (2020).

[8]I also avoid convoluted construction inherent in traditional methods of deriving the objective function that makes comparisons even between very similar non-standard model environments and substandard policies diffi-

remaining distortion, relative price dispersion. A significant amount of initial price dispersion is usually just assumed away (Gali 2003, Woodford 2003). With these assumptions, even with the countless extensions of the model, the base intuition of optimal monetary policy remains largely the same: the goal of policy should be to stabilize the price level to keep the economy at its natural level of output. In stabilizing the price level, a central bank stamps out the nominal distortions associated with sticky prices, reaching the flexible price equilibrium.[9] However, these are not so innocuous. They require us to already be near the long-run value, and to assume stabilization policy that we find optimal now to have already been used in the past. What happens though when, for whatever reason, we find ourselves stranded away from the long run? You could consider, for example that a central bank decides to change its inflation target. Even if everything else is optimal before and after besides the target, the bank has to decide the best transition path. You could also imagine a shock that is not predicted, such as a financial crisis or a large supply chain blockage, that shifts the environment enough that the bank has to think about transitioning back to "normal".

Some papers have broken these assumptions to study what would happen away from the long-run outcome. Yun (2005), for example, does not assume away initial price dispersion. Though he confirms the result that complete stabilization of the price level is optimal in the absence of initial price dispersion, he finds that optimal inflation targets respond to changes in the level of price distortion otherwise. Because price dispersion in a second order term, Yun cannot rely on a set of linearly approximated equilibrium conditions. He still maintains, however, a subsidy to eliminate the inefficiency associated with monopolistic competition. By introducing an employment subsidy, he ends up with a one-to-one relationship between the growth rates of inflation and price dispersion. This drastically simplifies the first order conditions of the firm and he is able to drop all the forward looking parts of the model, which themselves evolve non-linearly. Yun then drastically simplifies the aggregate supply curve since only current costs matter. Although including price distortions while assuming an efficient steady state gives clean analytical results even with non-linearity, including the forward-looking firm expectations that

cult (Benigno & Woodford, 2008).

[9]See, for example, Goodfried and King (1997), Rotemberg and Woodford (1997); Clarida, Gali, and Gertler

are necessary without that assumption makes this impossible.

To study both channels of distortion in this model, I directly include the nonlinear sources of inefficiency, price dispersion and a distorted stead-state, as well as the forward-looking elements of the firm's pricing decision. I make use of some of the results of Fernandez-Villaverde et al (2012), who analyze the dynamics of a fully non-linear New Keynsian model.[10] With sufficient starting inefficiencies in price dispersion and steady state, I confirm the result that stabilization policy is no longer optimal in the short run, and find that the central bank needs to think in terms of an optimal transition path to a long run value, which I also determine computationally. In the short run, the interest rate should be designed to have firm marginal cost be proportional to a function of both price dispersion and the shadow price on firm expectations on future costs. Although I confirm the broad conclusion that stabilization policy is not optimal in the short run in the presence of large distortions and that the central bank should think in terms of optimal transition, I show that the forward-looking elements of firm decision making are key for determining optimal policy when both model distortions are present. The reason this shadow price on firm expectations, and firm expectations themselves, are usually left out is obvious when looking at my results: they are computationally very difficult to determine. Other papers like Khan et al (2003) have likewise avoided the analytical simplification at first, including full non-linearities,[11] but then they linearize for the sake of computation, making their approach less useful for my question.

I show that the interest rate should guide optimal marginal cost in terms of both price dispersion and the co-state variable associated with firm expectations of future costs. This co-state variable has no closed form solution, and must be approximated. To do this, I first figure out the optimal path of the entire system defined by the first order conditions of the Hamiltonian, which is a system of ordinary differential equations (ODE's). This is an open form solution, in that the choice variables are a function of time and not the other state variables. This technique

(1999), Gali (2003), Woodford (2003).

[10]They are not interested in the question of monetary policy and allow for a Taylor-rule type policy.

[11]Their analysis, furthermore, is for a model with money which is different from what is considered in this application.

and what follows was first done broadly Nakamura-Zimmerman 2020.[12] To get the optimal paths, I approximate every variable by a neural network with input being time (this is equivalent to thinking of the ODE system as a neural network with vector output, input still time). To train this neural network, I take all of the equilibrium conditions of the ODE, express them as $f(x) = 0$, and try to minimize with randomly chosen time points along the path. In other words, I try to get all of the neural network variables to run along the paths as close as possible to the path constraints while satisfying boundary conditions, slowly adjusting the networks to be more optimal with a batch of time points along the path. Because boundary conditions are not handled entirely well here (a common problem in the deep learning literature), I follow the procedure of Lagari et al (2020) to convert the problem into one of "hard boundaries". This is really just nesting the neural network into a temporary function of time to get the process started; it forces the boundary conditions to hold at all times.

I use and modify packages developed out of a team at Brown University (Deep XDE, Lu et al 2021) for solving ODE systems in the way I described. I get multiple paths by randomly selecting initial starting positions (you can think of these as different initial distortions). I collect 150 of these paths and I then train the interest rate as a neural network, this time as a function of the variables in the model. I am converting the open-form solution to a closed-form one. I do this by seeing what the relationships are between the paths. This, again, is similar to the procedure of Nakamura-Zimmerman (2020). Because I have commitment, this is valid, since nothing is a function of time (to use the language of optimal control theory, I have an autonomous solution). In other words, I am seeing what the interest rate would have to be in order to make the model evolve as it does. Again, all of this is necessary because it allows me to handle the non-linearities without a good initial guess. Other global methods fail here, since I have no idea what this value function looks like given that I do not make all the heavy assumptions on the analytical side. It also can handle broader problems because it is not held up by the curse of dimensionality in the same way. Even with my number of variables, things are very computational taxing.

I show, finally, that these two channels of inefficiency, price distortion and monopolistic com-

[12]They, however, focus on the linear-quadratic case so this also serves as an extension of their method.

petition, matter for optimal monetary policy, and, furthermore, that this latter channel can be expressed in terms of firm expectations of future costs. The extent to which one channel matters more than the other will depend on the parameters of the model and initial conditions. I find also that policy should be more aggressive when price dispersion increases. More broadly, in answering the question of how firm expectations and price distortions should influence optimal monetary policy, I am able to avoid the usual assumptions of a quadratic objective function and linear dynamics, without simplification for the sake of computation, and I move away from a view of monetary policy as stabilizing the economy around a long-run result in the face of small Gaussian shocks. In Section 2 of the paper, I will lay out the broad strokes of setting up the central bank's problem and its analytical results. In Section 3, I will demonstrate the machine learning technique used for solving the model. Section 4 concludes.

2 Analytical Approach and Implications for Monetary Policy

2.1 The Model

In the following analysis, I will first illustrate the important elements and equilibrium conditions of the model environment, a nonlinear New Keynsian Model[13] making use of the staggered price-setting of Calvo (1983) with money only playing the role of a unit of account[14]. I will then use those equilibrium conditions to both set up the constraint set and rewrite the reward function in a convenient fashion for the central bank's problem. The policy prescription of the central bank - an optimal interest rate rule - will be inferred from the optimal dynamics of the decentralized economy. I will be freely taking from the results of the full underlying model, only reproducing the results as they are important for the central bank's problem. The full underlying model can be found in Appendix A. I set the modelling environments in continuous time, which helps me characterize much of the equilibrium dynamics analytically and to avoid needing to compute expectations even within a global solution.[15]

After the problem is posed, I will give the analytical results and their implication for monetary policy. I will show how the optimal dynamics should be thought of in terms of an interest rate policy that most optimally structures marginal cost as an equilibrium function of the state variables. Although the full result can only be determined numerically, I will demonstrate through analytical results that with a significant level of price dispersion that stabilization is not optimal in the short run and that a distorted stead state implies that an inflation target is not optimal, though I will confirm the general result of the monetary policy literature that stabilization around some value will be the optimal long run policy. Finally, I will numerically approximate the model solution using deep learning techniques. These techniques, though used for this specific case, can be much more widely applied than what is shown here.

[13]Consult Fernandez-Villaverde et al (2012) for an articulation, which makes use of more types of shocks than the ones analyzed in this paper though does not analyze monetary policy, and uses an inertial type Taylor rule.

[14]This is in the same vein as Woodford (1999) and Gali and Monacelli (2002)

[15]There is nothing essential, however, about continuous time. This method could be done in discrete time, though

2.2 Value Function

The objective function is exactly the consumer's utility function, though this no longer is expressed in terms of consumption and labor, but in terms rather of marginal utility and the product of price dispersion and marginal cost. These are derived from the equilibrium conditions of the underlying model. The central bank is thus trying to maximize:

$$\int_0^\infty e^{-\rho t}\{\ln(c_t) - \psi \frac{n_t^{1+\gamma}}{1+\gamma}\}$$

Of course, in a decentralized economy consumption and labor are both functions of the underlying state variables. We will express them in that form. To do this, consider the first order conditions of the household. For any interior solution, and when $\psi \neq 0$:

$$1/c_t = \lambda_t$$

$$\psi l_t^\gamma = \lambda_t w_t$$

$$\psi l_t^\gamma c_t = w_t$$

We also have an expression for aggregate production:

$$y_t = \frac{An_t}{v_t}$$

where v_t refers to price dispersion, defined mathematically as:

$$v_t = \int_0^1 (\frac{p_{it}}{p_t})^{-\varepsilon} di$$

$$v_t \geqslant 1$$

at a much higher computational cost. See Achdou et al. (2017) for an explanation of the advantages of continuous time in this manner.

Note that $v_t = 1$ would imply efficiency. Price dispersion acts as the point of inefficiency flowing from staggered price setting. Price dispersion can also be thought of as a misalignment between decisions made on the basis of marginal cost and those made on the basis of marginal utility. In practical terms it acts as a wedge between production in terms of inputs and in terms of output after aggregation.

Combining together the first order and market conditions:

$$\begin{aligned}
\text{mc}_t &= \psi n_t^{1+\gamma}/v_t \\
\text{mc}_t &= \psi(A\lambda_t)^{-(1+\gamma)}v_t^{\gamma} \\
c_t &= A(\frac{\text{mc}_t}{\psi v_t^{\gamma}})
\end{aligned}$$

where $\text{mc}_t = \frac{w_t}{A_t}$ refers to real marginal cost.

One important thing to take away here is that the policy functions of the representative agent - in this context, optimal consumption and labor as functions of the underlying state variables - can be analytically expressed in terms of of a function of price dispersion and marginal cost. This is a partial equilibrium result, and of course in the full New Keynesian model this will not be enough because firms also take into consideration expected future marginal costs as well as the current marginal cost when determining their prices. As we will see, the evolution of price dispersion is defined by a path constraint, but there is no such constraint on marginal cost. The equilibrium value for marginal cost will end up being an unknown function of the state variables. Really, marginal cost serves as a still unknown function with solves the equilibrium conditions and the maximized Bellman equation or Hamiltonian. (Fernandez-Villaverde et al (2012)).

The traditional approach is to linearize the equilibrium conditions and solve the system of linear dynamic equations (For an example and discussion of what the traditional approach looks like, see Appendix B). In contrast, my non-linear approach uses the competitive solution and the implied general equilibrium value function, which in turn pins down the unknown marginal costs. Observe that the costate variable depends on the stochastic shocks and price dispersion.

This is why our approach to solve for the general equilibrium values has been to augment the vector of state variables of the household's value function by the law of motions for expectations and price dispersion.

From the above, I will use the conditions:

$$\psi n_t^{1+\gamma} = v_t \mathrm{mc}_t$$
$$c_t = 1/\lambda_t$$

I can then rewrite the value function as:

$$\int_0^\infty e^{-\rho t}[\ln(1/\lambda_t) - \frac{v_t \mathrm{mc}_t}{1+\gamma}]dt$$

2.3 The Problem of the Central Bank

Even though the central bank only directly chooses the interest rate, I expand the choice set to include inflation and marginal cost. Because these are functions of underlying state variables, this is mathematically valid in terms of the optimal control problem (see Appendix C). More generally, you can think of this approach as a Ramsey problem: I will determine the optimal flow of the overall system and then back out the optimal interest rates as the way to decentralize the problem. I am just working in reverse: first taking the equilibrium constraints implied by the decentralized problem and then second having the central bank act as a kind of central planner.

$$\underset{r_t, \pi_t, \mathrm{mc}_t}{max} \int_0^\infty e^{-\rho t}[\ln(1/\lambda_t) - \frac{v_t \mathrm{mc}_t}{1+\gamma}]dt$$

s.t.

$$\dot{\Sigma_R} = (\theta - (\varepsilon - 1)\pi_t)\Sigma_{Rt} - 1 \tag{1}$$

$$\dot{\Sigma_C} = (\theta - \varepsilon\pi_t)\Sigma_{Ct} - \mathrm{mc}_t \tag{2}$$

$$\dot{v} = \theta(1 + \pi_t \frac{1-\varepsilon}{\theta})^{-\frac{\varepsilon}{1-\varepsilon}} + (\varepsilon\pi_t - \theta)v_t \tag{3}$$

$$\dot{\lambda} = (\rho - r_t + \pi_t)\lambda_t \tag{4}$$

$$\mathrm{mc}_t = \psi(A\lambda_t)^{-(1+\gamma)} v_t^{\gamma} \tag{5}$$

$$\pi_t = \frac{\theta}{1-\varepsilon}[(\frac{\varepsilon}{\varepsilon - 1}\frac{\Sigma_{Ct}}{\Sigma_{Rt}})^{1-\varepsilon} - 1] \tag{6}$$

$$\Sigma_{Rt}, \Sigma_{Ct}, \mathrm{mc}_t \geqslant 0, \forall t$$

$$v_t \geqslant 1, \lambda_t > 0, \forall t$$

$$\Sigma_{R0}, \Sigma_{C0}, v_0, \quad \text{given}$$

A quick summary of the variables:

Σ_R : Firm expectations about future aggregate demand conditions.

Σ_C : Firm expectations about future costs

v : Price dispersion

λ : Marginal Utility of Wealth

π : Inflation

mc : Firm marginal costs

A quick summary of parameters:

ρ : Discounting Parameter

θ : Rate of the Calvo Process, an exponential distribution

(1/θ gives the expected wait time until the next price change, $\theta > 0$)

ε : Elasticity of Substitution

ψ : The Disutility of Labor

γ : Inverse of Frisch Labor Supply Elasticity

My approach is different from than of Yun (2005) because I do not subsidize employment. Because of this, I cannot ignore the forward-looking pricing equation of the firm:

$$\max_{p_{it}} \mathrm{E}_t \int_t^{\infty} \frac{\lambda_\tau}{\lambda_t} e^{-\theta(\tau - t)} [\frac{p_{it}}{p_\tau} y_{i\tau} - \mathrm{mc}_\tau y_{i\tau}] d\tau$$

The first order conditions of the firm is as follows. The ratio of the optimal new price, common across all firms able to reset their prices, and the prices of the final good, is given by:

$$\frac{p_{it}}{p_t} = \frac{\varepsilon}{\varepsilon - 1} \frac{\Sigma_{Ct}}{\Sigma_{Rt}}$$

where:

$$\Sigma_{Rt} = \int_t^{\infty} \lambda_\tau e^{-\theta(\tau - t)} (\frac{p_t}{p_\tau})^{1-\varepsilon} y_\tau d\tau$$

represents expected present discounted value of future revenues, which acts as an estimate of future aggregate demand conditions and

$$\Sigma_{Ct} = \int_t^{\infty} \lambda_\tau e^{-\theta(\tau - t)} \mathrm{mc}_\tau (\frac{p_t}{p_\tau})^{-\varepsilon} y_\tau d\tau$$

represents expected present discounted value of future costs. Because I am working in continuous time, how to represent the above pieces of the pricing decision as state variables is more intuitive.

$$\frac{p_{it}}{p_t} = \frac{1}{v_t}$$

$$\pi_t = \dot{v}$$

These auxiliary variables are related to inflation in the following way:

$$\pi_t = \frac{\theta}{1-\varepsilon} [(\frac{\varepsilon}{\varepsilon - 1} \frac{\Sigma_{Ct}}{\Sigma_{Rt}})^{1-\varepsilon} - 1$$

Once we include the definitions of these variables, we can see that the above can be thought of as the Aggregate Supply Curve of the model.

$$\pi_t = \frac{\theta}{1-\varepsilon}[(\frac{\varepsilon}{\varepsilon-1}\frac{\int_t^\infty \lambda_\tau e^{-\theta(\tau-t)}\text{mc}_\tau(\frac{p_t}{p_\tau})^{-\varepsilon}y_\tau d\tau}{\int_t^\infty \lambda_\tau e^{-\theta(\tau-t)}(\frac{p_t}{p_\tau})^{1-\varepsilon}y_\tau d\tau})^{1-\varepsilon} - 1$$

2.3.1 The Hamiltonian

I will solve the central bank's problem using optimal control theory. I construct the Hamiltonian as follows:

$$\begin{aligned}
H =& \ln(1/\lambda_t) - \frac{v_t\text{mc}_t}{1+\gamma} \\
&+ \Lambda_{\Sigma_R}[(\theta-(\varepsilon-1)\pi_t)\Sigma_{Rt} - 1] \\
&+ \Lambda_{\Sigma_C}[(\theta-\varepsilon\pi_t)\Sigma_{Ct} - \text{mc}_t] \\
&+ \Lambda_v[\theta(1+\pi_t\frac{1-\varepsilon}{\theta})^{-\frac{\varepsilon}{1-\varepsilon}} + (\varepsilon\pi_t-\theta)v_t] \\
&+ \Lambda_\lambda[(\rho - r_t + \pi_t)\lambda_t] \\
&+ \mu_{mc}[\text{mc}_t - \psi(A\lambda_t)^{-(1+\gamma)}v_t^\gamma] \\
&+ \mu_\pi[(1+\pi_t\frac{1-\varepsilon}{\theta})^{\frac{1}{1-\varepsilon}} - \frac{\varepsilon}{\varepsilon-1}\frac{\Sigma_{Ct}}{\Sigma_{Rt}}]
\end{aligned}$$

2.3.2 First Order Conditions of the Hamiltonian System

Control variables:

$$\frac{\partial H}{\partial r_t} = 0 = -\Lambda_\lambda\lambda_t \tag{7}$$

$$\begin{aligned}
\frac{\partial H}{\partial \pi_t} = 0 =& -\Lambda_{\Sigma_R}(\varepsilon-1)\Sigma_{Rt} - \Lambda_{\Sigma_C}\varepsilon\Sigma_{Ct} - \Lambda_v[\varepsilon(1+\pi_t(1-\varepsilon)/\theta)^{-\frac{1}{1-\varepsilon}} - \varepsilon v_t] \\
&+ \Lambda_\lambda\lambda_t + \frac{\mu_\pi}{\theta}((1+\pi_t\frac{1-\varepsilon}{\theta})^{\frac{\varepsilon}{1-\varepsilon}}
\end{aligned} \tag{8}$$

$$\frac{\partial H}{\partial \text{mc}_t} = 0 = -\Lambda_{\Sigma_C} + \mu_{mc} - \frac{v_t}{1+\gamma} \tag{9}$$

Equation (7) may seem strange - after all, if the only tool the central bank has is the interest rate and the interest rate has not effect, are we done here? The short answer is no. The interest rate provides structure for the environment. We are finding (through equations (8) and (9)) that optimal structure, and then backing out the interest rate implied by the optimal dynamics of the system. It is important to remember that any other choice of interest rate not given by these implied dynamics will cause the dynamics to be different, and hence worse.

Let us break down equations (8) and (9) further. Equation (8) is the first order condition for inflation:

$$(1+\pi_t\frac{1-\varepsilon}{\theta})^{\frac{1}{1-\varepsilon}}\underbrace{\mu_\pi}_{①} = \underbrace{\Lambda_{\Sigma_R}(\varepsilon-1)\Sigma_{Rt}+\Lambda_{\Sigma_C}\varepsilon\Sigma_{Ct}}_{②} + \underbrace{\Lambda_v\varepsilon(1+\pi_t(1-\varepsilon)/\theta)^{-\frac{1}{1-\varepsilon}}-\Lambda_v\varepsilon v_t}_{③}$$

The weighted benefits from raising inflation ①, represented by the lagrange multiplier μ_π, will have two effects. The first will be through the effect on the optimal firm pricing decision ②. Inflation causes expected real revenues and real costs to both decrease. Whether this will cause the reset price to go up or down will depend on the current values of Σ_R and Σ_C, which represent the weight on the reset price by expected future aggregate demand conditions and marginal costs, respectively. Inflation also has an impact on price dispersion ③. Generally, inflation causes price dispersion to increase, though the overall effect of this will depend on the current levels of inflation and price dispersion. Any model, including Yun (2005), which adds a subsidy to employment will lose the effects represented by ②.

Equation (9) is the first order condition for marginal cost:

$$\underbrace{\mu_{mc}}_{①} = \underbrace{\Lambda_{\Sigma_C}}_{②} + \underbrace{\frac{v_t}{1+\gamma}}_{③}$$

The weighted benefits from raising marginal costs ①, represented by the lagrange multiplier μ_{mc}, will have two effects. The first will be through the effect on the optimal firm pricing decision ②. Increase marginal cost also has an effect on labor ③, since it is equivalent an increase in the equilibrium wage. This effect will be magnified by the amount of price dispersion in the model, which decreases the efficiency of labor with respect to output. Again, any model, including Yun (2005), which adds a subsidy to employment will lose the effects represented by ②.

Now let us look at the other optimality conditions:

State variables:

$$\frac{\partial H}{\partial \Sigma_{Rt}} = -\dot{\Lambda}_{\Sigma_R} + \rho\Lambda_{\Sigma_R} = \Lambda_{\Sigma_R}(\theta - (\varepsilon - 1)\pi_t) + \mu_\pi \frac{\varepsilon}{\varepsilon - 1}\frac{\Sigma_{Ct}}{(\Sigma_{Rt})^2} \tag{10}$$

$$\frac{\partial H}{\partial \Sigma_{Ct}} = -\dot{\Lambda}_{\Sigma_C} + \rho\Lambda_{\Sigma_C} = \Lambda_{\Sigma_C}(\theta - \varepsilon\pi_t) - \frac{\mu_\pi}{\Sigma_{Rt}}\frac{\varepsilon}{\varepsilon - 1} \tag{11}$$

$$\frac{\partial H}{\partial v_t} = -\dot{\Lambda_v} + \rho\Lambda_v = -\frac{\mathrm{mc}_t}{1+\gamma} + \Lambda_v(\varepsilon\pi_t - \theta) - \mu_{mc}\gamma\psi(\lambda_t A)^{-(1+\gamma)}v_t^{\gamma-1} \tag{12}$$

$$\frac{\partial H}{\partial \lambda_t} = -\dot{\Lambda_\lambda} + \rho\Lambda_\lambda = -\frac{1}{\lambda_t} + \Lambda_\lambda(\rho - r_t + \pi_t) + \mu_{mc}(1+\gamma)\psi A^{-(1+\gamma)}\lambda_t^{-(2+\gamma)}v_t^\gamma] \tag{13}$$

As well as the transversality conditions, specified in Appendix D.1.

2.4 Determining Optimal Marginal Cost

From (7) and (13), we have:

$$\frac{1}{\lambda_t} = \mu_{mc}\psi(1+\gamma)v_t^\gamma A(\lambda_t A)^{-(2+\gamma)}$$

$$1 = \mu_{mc}\psi(1+\gamma)v_t^\gamma(\lambda_t A)^{-(1+\gamma)}$$

$$(\lambda_t A)^{1+\gamma} = \mu_{mc}\psi(1+\gamma)v_t^\gamma \tag{14}$$

From (9) and the above, we have:

$$(\lambda_t A)^{1+\gamma} = (\Lambda_{\Sigma_C} + \frac{v_t}{1+\gamma})\psi(1+\gamma)v_t^\gamma$$

$$\lambda_t = \frac{1}{A}[(\Lambda_{\Sigma_C} + \frac{v_t}{1+\gamma})\psi(1+\gamma)v_t^\gamma]^{\frac{1}{1+\gamma}} \tag{15}$$

Taking now the original equilibrium condition for marginal cost (5) and using the results above, we get:

$$\mathrm{mc}_t = \frac{1}{(1+\gamma)\Lambda_{\Sigma_C} + v_t} \tag{16}$$

This is the most fundamental result of the model, and represents the optimal function for marginal cost to take in our system. The way to think about the optimal interest rate is precisely to *struc-*

ture the environment in such a way to make this marginal cost possible. In Ramsey's language, it will be the way to decentralize the solution. It says that marginal firm cost is the inverse of a linear combination of the shadow price on expectations of future cost and price dispersion. The shadow price on expected future costs is weighted by a measure of labor supply elasticity. This effectively shows us that the optimal marginal cost is related to the two inefficiencies of the New Keynesian Model - inefficient production due to monopolistic competition and price dispersion.

Compare the above equation with the results of Yun (2005):

$$\mathrm{mc}_t = \frac{1}{(1+\eta)v_t} \tag{17}$$

where η is the optimal subsidy rate, which would be set to $1/(\varepsilon - 1)$. Note that because of the employment subsidy, the long-run steady state value is now efficient, Yun does not have to take into consideration the evolution of expectations in his model, and the weighted shadow price of future expectations of cost, $(1+\gamma)\Lambda_{\Sigma_C}$, become irrelevant.

2.5 Determining the Interest Rate

Let us get an expression of the interest rate.

Taking the time derivative of λ_t, we arrive at:

$$\dot{\lambda} = \frac{\psi(1+\gamma)}{A}[(\Lambda_{\Sigma_C} + \frac{v_t}{1+\gamma})\psi(1+\gamma)v_t^{\gamma}]^{\frac{-\gamma}{1+\gamma}}[(\Lambda_{\Sigma_C} + \frac{v_t}{1+\gamma})\gamma v_t^{\gamma-1}\dot{v} + v_t^{\gamma}(\dot{\Lambda_{\Sigma_C}} + \frac{\dot{v}}{1+\gamma})]$$

Dividing by λ_t:

$$\begin{aligned}\frac{\dot{\lambda}}{\lambda_t} &= \frac{(\Lambda_{\Sigma_C} + \frac{v_t}{1+\gamma})\gamma v_t^{\gamma-1}\dot{v} + v_t^{\gamma}(\dot{\Lambda_{\Sigma_C}} + \frac{\dot{v}}{1+\gamma})}{(\Lambda_{\Sigma_C} + \frac{v_t}{1+\gamma})v_t^{\gamma}} \\ &= \gamma\frac{\dot{v}}{v_t} + \frac{(1+\gamma)\dot{\Lambda_{\Sigma_C}} + \dot{v}}{((1+\gamma)\Lambda_{\Sigma_C} + v_t)} \qquad (18)\end{aligned}$$

Taking the time derivative of marginal cost:

$$\dot{\text{mc}} = -\frac{(1+\gamma)\dot{\Lambda_{\Sigma_C}} + \dot{v}}{((1+\gamma)\Lambda_{\Sigma_C} + v_t)^2} \qquad (19)$$

From (4), the equilibrium condition for $\frac{\dot{\lambda}}{\lambda_t}$, and (18), and (19), we arrive at the following identity defining the interest rate:

$$r_t = \rho + \pi_t - \gamma\frac{\dot{v}}{v_t} + \frac{\dot{\text{mc}}}{\text{mc}_t} \qquad (20)$$

Note that this is an equilibrium object, not a rule per se. However, we must remember that marginal cost is in our system a function of the underlying state variables in an unknown way. Inflation, also, is a function of state variables, though in a direct, observable way. We can see this through a re-write of the above as:

$$r_t = \rho + \frac{\theta}{1-\varepsilon}[(\frac{\varepsilon}{\varepsilon-1}\frac{\Sigma_{Ct}}{\Sigma_{Rt}})^{1-\varepsilon} - 1] - \gamma\frac{\dot{v}}{v_t} + \frac{(1+\gamma)\dot{\Lambda_{\Sigma_C}} + \dot{v}}{((1+\gamma)\Lambda_{\Sigma_C} + v_t)} \qquad (21)$$

Keep in mind that $\dot{\Lambda_{\Sigma_C}}$ and Λ_{Σ_C} cannot be determined analytically and must be computed. I will later demonstrate machine learning techniques for solving this problem that can be applied to other uses. Before moving on to computation, however, let us analytically compare the dynamics of our result so far with a rule of the Taylor Rule variety.

2.6 Comparison of Dynamics with Taylor Rule

Consider for example the following Taylor rule. To elaborate more on intuition rather than realism, I assume the central bank only cares about inflation and that the target level is 0.

$$\begin{aligned} r_t =& \phi_\pi \pi_t \\ \phi_\pi &> 0 \end{aligned} \tag{22}$$

I can now plug this value in to the equilibrium conditions of the model to arrive at the dynamics of marginal cost, the key variable in our underlying system.

$$\begin{aligned} \frac{\dot{\lambda}}{\lambda_t} &= \rho + \pi_t - \phi_\pi \pi_t \\ &= \gamma \frac{\dot{v}}{v_t} - \frac{\dot{\mathrm{mc}}}{\mathrm{mc}_t} \\ &\Rightarrow \\ \frac{\dot{\mathrm{mc}}}{\mathrm{mc}_t} &= -\gamma \frac{\dot{v}}{v_t} - ((1 - \phi_\pi)\pi_t) \end{aligned}$$

Compare the above equation with the results of Yun (2005):

$$\frac{\dot{\mathrm{mc}}}{\mathrm{mc}_t} = -\frac{\dot{v}}{(1+\eta)v_t} \tag{23}$$

where η is the optimal subsidy rate, which would be set to $1/(\varepsilon - 1)$.

Compare both of these with my fundamental result from before, which was:

$$\frac{\dot{\mathrm{mc}}}{\mathrm{mc}_t} = -\frac{(1+\gamma)\dot{\Lambda_{\Sigma_C}} + \dot{v}}{(1+\gamma)\Lambda_{\Sigma_C} + v_t} \tag{24}$$

2.7 Comparison of Dynamics of Price Stability Rule

We can also determine a rule that seeks perfect price stability, in other words one that sets out to:

$$\pi_t = 0, \forall t$$

From this we get:

$$\dot{v} = \theta(1 - v_t)$$

We also have the following relationship between expectation terms:

$$\frac{\Sigma_{Ct}}{\Sigma_{Rt}} = \frac{\varepsilon - 1}{\varepsilon}$$

This does not imply the two terms are always constant, but it *does* imply that their growth rates must be the same to maintain a constant proportion. From the equilibrium dynamics of each we get that:

$$\dot{\Sigma_R} = \theta\Sigma_{Rt} - 1$$
$$\dot{\Sigma_C} = \theta\Sigma_{Ct} - \mathrm{mc}_t$$

Which means that:

$$\theta\Sigma_{Rt} - 1 = \theta\Sigma_{Ct} - \mathrm{mc}_t$$

So we arrive at an expression of marginal cost:

$$\mathrm{mc}_t = 1 + \theta(\Sigma_{Ct} - \Sigma_{Rt})$$

Taking the time derivative:

$$\dot{\text{mc}} = \theta(\dot{\Sigma_C} - \dot{\Sigma_R}) = 0$$

Intuitively, this must be zero because the goal of marginal cost in this instance is to make the expectation of future costs have a specific value to maintain a certain proportion. Once this is locked into place, there is no more

This allows us to say something about how our result compares. Remember our first order condition:

$$\frac{\dot{\text{mc}}}{\text{mc}_t} = -\frac{(1+\gamma)\dot{\Lambda_{\Sigma_C}} + \dot{v}}{(1+\gamma)\Lambda_{\Sigma_C} + v_t} \tag{25}$$

Note that even mechanically this cannot be equal to 0 except in steady state, since, from the dynamics with inflation set to 0 included:

$$\dot{v} = \theta(1 - v_t)$$

Now, what this does mean is that my optimal result converges to the standard result if $v_0 = 1$, in other words if we start in the long run equilibrium. This is a confirmation of the results of Yun (2005), though without the assumption of a subsidy on employment, meaning that the steady state itself is distorted, then this will not be correct. I will show this explicitly in a later section, 2.9. This, however, requires thinking through possible steady states, which requires me to first lay out the optimal dynamics of the entire system.

2.8 The System

We arrive thus at the following system of differential equations which define the optimal evolution of the economy. The border conditions are given by the steady state values of the relevant variables. Note that our costate variable, the marginal utility of wealth, now vanishes in the system governing our complete results. Because marginal cost is allowed to freely move, the optimal interest rate is what I "back out" of the process controlling the costate term.

$$\dot{\Lambda_{\Sigma_R}} = (\rho - (\theta - (\varepsilon - 1)\pi_t))\Lambda_{\Sigma_R} - \mu_\pi \frac{\varepsilon}{\varepsilon - 1}\frac{\Sigma_{Ct}}{(\Sigma_{Rt})^2} \tag{26}$$

$$\dot{\Lambda_{\Sigma_C}} = (\rho - (\theta - \varepsilon\pi_t))\Lambda_{\Sigma_C} + \mu_\pi \frac{\varepsilon}{\varepsilon - 1}\frac{1}{\Sigma_{Rt}} \tag{27}$$

$$\dot{\Lambda_v} = (\rho - \varepsilon\pi_t + \theta)\Lambda_v + \frac{\mathrm{mc}_t}{1+\gamma} + \frac{\gamma}{1+\gamma}\frac{1}{v_t} \tag{28}$$

$$\dot{\Sigma_R} = (\theta - (\varepsilon - 1)\pi_t)\Sigma_{Rt} - 1 \tag{29}$$

$$\dot{\Sigma_C} = (\theta - \varepsilon\pi_t)\Sigma_{Ct} - \mathrm{mc}_t \tag{30}$$

$$\dot{v} = \theta(1 + \pi_t(1-\varepsilon)/\theta)^{-\varepsilon/(1-\varepsilon)} + (\varepsilon\pi_t - \theta)v_t \tag{31}$$

$$s.t.$$

$$(1 + \pi_t \frac{1-\varepsilon}{\theta})^{\frac{1}{1-\varepsilon}} = \frac{\varepsilon}{\varepsilon - 1}\frac{\Sigma_{Ct}}{\Sigma_{Rt}} \tag{32}$$

$$\mathrm{mc}_t = \frac{1}{(1+\gamma)\Lambda_{\Sigma_C} + v_t} \tag{33}$$

$$\frac{\mu_\pi}{\theta}(1 + \pi_t(1-\varepsilon)/\theta)^{\frac{\varepsilon}{1-\varepsilon}} = \Lambda_{\Sigma_R}(\varepsilon - 1)\Sigma_{Rt} + \Lambda_{\Sigma_C}\varepsilon\Sigma_{Ct} + \Lambda_v[\varepsilon[(1 + (1-\varepsilon)\pi_t/\theta)]^{\frac{-1}{1-\varepsilon}} - \varepsilon v_t] \tag{34}$$

I will use this system in our numerical computation. The steady state values are also computationally determined. To see the steady state implied by a specified parameterization, see 3.4.1.

2.9 Non-Optimality of Zero-Inflation Long-Run Target

I will not solve here for the full steady-state allocation, but instead want to point to one result: the non-optimality of perfect price stability.

From the above, assuming $\pi_{ss} = 0$ we have first from equation (31):

$$0 = \theta(1 - v_{ss})$$

$$\Rightarrow v_{ss} = 1$$

From equations (29) and (30), we have that:

$$\Sigma_{Rss} = \frac{1}{\theta}$$
$$\Sigma_{Css} = \frac{\mathrm{mc}_{ss}}{\theta}$$

From equation (32) and the above results:

$$1 = \frac{\varepsilon}{\varepsilon - 1}\frac{\Sigma_{Css}}{\Sigma_{Rss}}$$
$$\Rightarrow \mathrm{mc}_{ss} = \frac{\varepsilon - 1}{\varepsilon}$$

Using these results, we can take equations (26) and (27) to arrive at:

$$\Lambda_{\Sigma_{Rss}} = \frac{\theta\mu_{\pi}}{(\rho - \theta)}$$
$$\Lambda_{\Sigma_{Css}} = -\frac{\varepsilon\theta\mu_{\pi}}{(\varepsilon - 1)(\rho - \theta)}$$

Substituting into equation (34):

$$\frac{\mu_{\pi}}{\theta} = \frac{(\varepsilon - 1)\theta\mu_{\pi}}{\theta(\rho - \theta)} - \frac{\varepsilon(\varepsilon - 1)\varepsilon\theta\mu_{\pi}}{\theta(\rho - \theta)\varepsilon(\varepsilon - 1)}$$
$$\Rightarrow \frac{1}{\theta} = \frac{(\varepsilon - 1)}{(\rho - \theta)} - \frac{\varepsilon}{(\rho - \theta)}$$
$$\Rightarrow \frac{(\rho - \theta)}{\theta} = -1$$
$$\Rightarrow \rho - \theta = -\theta$$
$$\Rightarrow \rho = 0$$

Thus the parameterization we would need in order to solve the problem would imply that the discount parameter would equal zero, meaning that the representative does not discount the future at all. Most of the time the literature assumes that the subjective discount factor is close enough to zero where approximation does not have to factor this in. Far enough from the steady state, however, this is no longer the case. Also, unlike Yun (2005), because of a distorted steady state we do not converge in the long run to a zero inflation steady state. However, because there exists a steady state, the central bank will always return to stabilizing around some given steady state value. For the computationally determined steady state values of a given parameter set, see 3.4.1.

3 Numerical Solution: A Deep Learning Approach

In recent years deep learning has been applied to problems of PDEs (Raissi et al 2018). More recently it has been applied to explicit economics questions (Duarte 2018, Fernandez-Villaverde et al 2020). My approach follows most closely the method of Nakamura-Zimmerer et. al (2021a, 2021b), who use neural networks to approximate two-part boundary problems by first solving a system of ordinary differential equations, though in that paper the authors limit themselves to an LQ framework. I then use this data to train a neural network for approximating the relevant portions of the value function along the optimal path. The insight is that because I am only interested in the interest rate, I do not need to fully approximate the entire value function. In other words, I first numerically compute the optimal path, and then use the optimal path to train an additional model for the interrelation of state and control variables. This approach of solving the system of equations defining the monetary policy problem as a boundary problem and then using the results to train an additional model of the interest rate significantly aids with computational speed.

As mentioned in the introduction, there are advantages of this deep learning approach compared to more familiar methods. By the universal approximation theorem (Bach 2017), a neural network can approximate any unknown Borel measurable function, and neural networks are less sensitive to good initial guesses than collocation methods. A neural network method (largely) allows one to avoid the curse of dimensionality that define grid based methods, which forms the bulk of economic numerical methods.

3.1 Deep Learning: A Brief Overview

At the lowest level, a neural network is composed of "neurons", functions of the form:

$$n(x;\Theta) \equiv \phi(\theta_0 + \sum_{i}^{N} \theta_i x_i)$$

The function takes input x and is paramaterized by the weight vector Θ. The activation function $\phi(.)$ is a nonlinear function. Common functions include the hyperbolic tangent.

I create a "layer" by stacking N_1 neurons on top of each other:

$$N(x;\Theta) \equiv (n(x;\Theta_1), ..., n(x;\Theta_{N1}))^T$$

3.2 Deep Learning: Application

Let us now look at how to actually implement the deep learning framework. To get a feel for the approach, consider a subset of our equilibrium conditions:

$$\begin{aligned}
0 &= -\dot{\Lambda_{\Sigma_R}} + (\rho - (\theta - (\varepsilon - 1)\pi_t))\Lambda_{\Sigma_R} - \mu_\pi \frac{\varepsilon}{\varepsilon - 1}\frac{\Sigma_{Ct}}{(\Sigma_{Rt})^2} \\
0 &= -\dot{\Lambda_{\Sigma_C}} + (\rho - (\theta - \varepsilon\pi_t))\Lambda_{\Sigma_C} + \mu_\pi \frac{\varepsilon}{\varepsilon - 1}\frac{1}{\Sigma_{Rt}} \\
0 &= -\dot{\Lambda_v} + (\rho - \varepsilon\pi_t + \theta)\Lambda_v + \frac{\mathrm{mc}_t}{1+\gamma} + \frac{\gamma}{1+\gamma}\frac{1}{v_t}
\end{aligned}$$

I will express each variable as a neural network with time as the only input variable. I will define the error associated with each equilibrium condition in the following way:

$$\begin{aligned}
err_{\Lambda_{\Sigma_R}} &= -\frac{\partial \Lambda_{\Sigma_R}(t,\Theta)}{\partial t} + (\rho - (\theta - (\varepsilon - 1)\pi(t,\Theta)))\Lambda_{\Sigma_R}(t,\Theta) \\
&\quad - \mu_\pi(t,\Theta)\frac{\varepsilon}{\varepsilon - 1}\frac{\Sigma_C(t,\Theta)}{(\Sigma_R(t,\Theta))^2} \\
err_{\Lambda_{\Sigma_C}} &= -\frac{\partial \Lambda_{\Sigma_C}(t,\Theta)}{\partial t} + (\rho - (\theta - \varepsilon\pi(t,\Theta)))\Lambda_{\Sigma_C}(t,\Theta) \\
&\quad + \mu_\pi(t,\Theta)\frac{\varepsilon}{\varepsilon - 1}\frac{1}{\Sigma_R(t,\Theta)} \\
err_{\Lambda_v} &= -\frac{\partial \Lambda_v(t,\Theta)}{\partial t} + (\rho - \varepsilon\pi(t,\Theta) + \theta)\Lambda_v(t,\Theta) \\
&\quad + \frac{\mathrm{mc}(t,\Theta)}{1+\gamma} + \frac{\gamma}{1+\gamma}\frac{1}{v(t,\Theta)}
\end{aligned}$$

These error terms relate to deviations in the neural model from the dynamic path constraints and train the model to trace the optimal path as defined by optimality conditions.

I can do the same for path equality constraints defining certain variables and associated with Lagrange multipliers:

$$
\begin{aligned}
err_{mc} =& -\mathrm{mc}(t,\Theta) + \frac{1}{(1+\gamma)\Lambda_{\Sigma_C}(t,\Theta) + v(t,\Theta)} \\
err_{\mu_\pi} =& -\frac{\mu_\pi(t,\Theta)}{\theta}(1+\pi(t,\Theta)(1-\varepsilon)/\theta)^{\frac{\varepsilon}{1-\varepsilon}} \\
& + \Lambda_{\Sigma_R}(t,\Theta)(\varepsilon-1)\Sigma_R(t,\Theta) + \Lambda_{\Sigma_C}(t,\Theta)\varepsilon\Sigma_C(t,\Theta) \\
& - \Lambda_v(t,\Theta)[\varepsilon[\theta(1+(1-\varepsilon)\pi(t,\Theta)/\theta)]^{\frac{2\varepsilon-1}{1-\varepsilon}} + \varepsilon v(t,\Theta)]
\end{aligned}
$$

As well as for the dynamics of each variable and the constraint defining inflation:

$$
\begin{aligned}
err_{\Sigma_R} =& -\frac{\partial \Sigma_R(t,\Theta)}{\partial t} + (\theta - (\varepsilon-1)\pi(t,\Theta))\Sigma_R(t,\Theta) - 1 \\
err_{\Sigma_C} =& -\frac{\partial \Sigma_C(t,\Theta)}{\partial t} + (\theta - \varepsilon\pi(t,\Theta))\Sigma_C(t,\Theta) - \mathrm{mc}(t,\Theta) \\
err_v =& -\frac{\partial v(t,\Theta)}{\partial t} + \theta(1+\pi(t,\Theta)(1-\varepsilon)/\theta)^{\varepsilon/(1-\varepsilon)} \\
& + (\varepsilon\pi(t,\Theta) - \theta)v(t,\Theta) \\
err_\pi =& -(1+\pi(t,\Theta)\frac{1-\varepsilon}{\theta})^{\frac{1}{1-\varepsilon}} + \frac{\varepsilon}{\varepsilon-1}\frac{\Sigma_C(t,\Theta)}{\Sigma_R(t,\Theta)}
\end{aligned}
$$

Finally I define the error at the boundary conditions.

$$
\begin{aligned}
err_{\Sigma_R,0} =& \Sigma_R(0,\Theta) - \Sigma_{R0} \\
err_{\Sigma_C,0} =& \Sigma_C(0,\Theta) - \Sigma_{C0} \\
err_{v,0} =& v(t,\Theta) - v_0 \\
err_{\Lambda_{\Sigma_R},T} =& e^{-pT}\Lambda_{\Sigma_R}(T,\Theta) \\
err_{\Lambda_{\Sigma_R},T} =& e^{-pT}\Lambda_{\Sigma_C}(T,\Theta) \\
err_{\Lambda_{\Sigma_R},T} =& e^{-pT}\Lambda_v(T,\Theta)
\end{aligned}
$$

For estimation another technique was also used to increase efficiency in cases where error at the boundaries was unacceptably large. Instead of directly including error terms for the boundary conditions, I reformulate the neural network as a neural form to incorporate "hard boundaries", as shown in Lagari et al (2020).

A neural form is any construction that is built upon a neural network. For our purposes, consider for example the neural form associated with inflation:

$$(\frac{T-t}{T})\pi_0 + t(t-T)\pi(t,\Theta))) + (\frac{t}{T})\pi_{ss}$$

We can see that for the above at either boundary - the terminal steady state or the initial condition - the neural form is constructed by design to fit the boundary condition with complete accuracy. The central component, the actual neural network within the neural form, is what is trained to fit the path conditions.

The infinite-horizon variation is obtained with the limit $T \to \infty$. I will use the error above defined for a particular value of T, then extend that value as I solve if the terminal errors are above a certain tolerance.

The total loss is defined as:

$$\begin{aligned} loss(t;\Theta) =& err^2_{\Lambda_{\Sigma_R}} + err^2_{\Lambda_{\Sigma_C}} + err^2_{\Lambda_v} + err^2_{mc} + err^2_{\mu_\pi} \\ & + err^2_{\Sigma_R} + err^2_{\Sigma_C} + err^2_{\pi} + err^2_{\Sigma_R,0} + err^2_{\Sigma_C,0} + err^2_{v,0} \\ & + err^2_{\Lambda_{\Sigma_R},T} + err^2_{\Lambda_{\Sigma_R},T} + err^2_{\Lambda_{\Sigma_R},T} \end{aligned}$$

To solve the model, I choose the parameter set Θ to minimize the above global loss function over a set of time points.

$$\frac{1}{|D|}\sum_{i=1}^{|D|} loss(t_i;\theta)$$

The solution will be the open loop solution, in other words I will have the optimal path. Once I have the optimal path, then I can define another neural network to approximate the optimal interest rate as a function of the state variables:

$$r(x) = NN(x;\theta)$$

Note that I had previously defined the optimal interest rate as:

$$r_t = \rho + \pi_t - \gamma\frac{\dot{v}}{v_t} + \frac{\dot{\text{mc}}_t}{\text{mc}_t}$$

I thus already have the optimal interest rate defined on the optimal path. To train the neural network, I first obtain a set of optimal paths with randomly chosen initial points. I then take a set of points along these optimal paths and minimize the error to approximate the interest rate as a function of the state variables. In other words, I convert a set of open loop solutions to a closed loop one.

$$err_r(x,t) = \frac{1}{|D|}\sum_{i=1}^{|D|}[r(x;\theta) - r(t)]^2$$

3.3 Technical Details

All coding was done in python using the Tensorflow library. For the first round, I construct a fully connected neural network with 8 hidden layers of 120 neurons each. The library DeepXDE was used for the first round (Lu et al 2021). The sigmoid activation function and Adam stochastic gradient descent-type algorithm are adopted in the neural network. For the second round, I construct a fully connected neural network with 4 hidden layers of 64 neurons each. The tanh activation function and Adam stochastic gradient descent-type algorithm are adopted in the neural network. This work utilized the Summit supercomputer, which is supported by the National Science Foundation (awards ACI-1532235 and ACI-1532236), the University of Colorado Boulder, and Colorado State University. The Summit supercomputer is a joint effort of the University of Colorado Boulder and Colorado State University.

3.3.1 Parameterization and Steady State Values

Table 1: Parameterization

γ	1	Frisch labor supply elasticity
ρ	0.01	Subjective rate of time preference
ψ	1	Preference for leisure
θ	0.65	Calvo parameter
ε	25	Elasticity of substitution for intermediate goods

The following steady state values are computationally determined.

Table 2: Steady State Values

$\Lambda_{\Sigma_{RSS}}$	-0.0200	Costate, Discounted Future Revenues
$\Lambda_{\Sigma_{CSS}}$	0.0208	Costate, Discounted Future Costs
Λ_{vSS}	-1.4844	Costate, Price Dispersion
Σ_{RSS}	1.5380	Discounted Future Revenues
Σ_{CSS}	1.4765	Discounted Future Costs
v_{SS}	$1.0+1.9472\times10^{-9}$	Price Dispersion
$\mu_{\pi SS}$	0.0197	Lagrange coefficient, inflation and auxiliary variables
mc_{SS}	0.9600	Marginal Cost
π_{SS}	-8.1181×10^{-6}	Inflation
r_{SS}	0.9992%	Interest Rate

3.4 Computational Results

We arrive then at the computational results. First, I want to give a sample set of transition paths with a set of initial conditions. Next, I will show what the interest rate looks like in terms of the other variables of the model.

3.4.1 Sample Transition Paths

The initial conditions are as follows. Note that marginal cost and inflation are functions of the other three true state variables of the model. I include them for clarity.

Table 3: Initial Conditions

Σ_0	1.62	Discounted Future Revenues
Σ_0	1.61	Discounted Future Costs
v_0	1.62	Price Dispersion
mc_0	0.348	Marginal Cost
π_0	2.96×10^{-6}	Inflation

The first thing to notice is that the central bank wants to get price dispersion down, and does so on a gliding path. Price dispersion, though inevitable with greater production, is very damaging to consumer welfare. You can see that the central bank lowers price dispersion nearly to 1, the long run value. This is very close to, if not essentially, efficient.

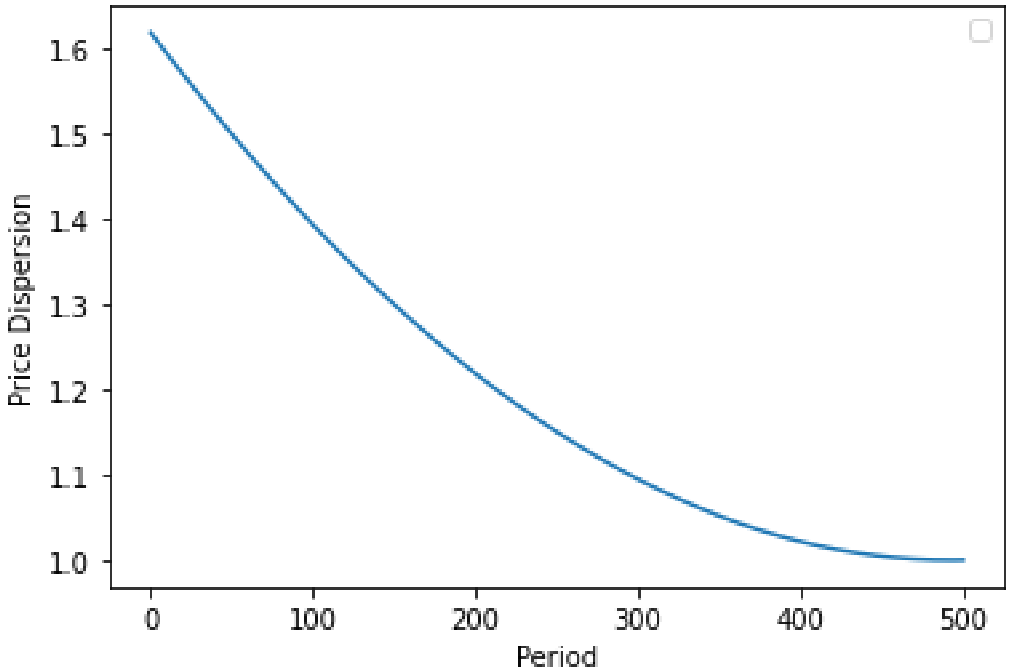

For price dispersion to decrease, inflation must also decrease. The central bank essentially slows down the economy, and does so to such an extent as to cause deflation for a short amount of time. This deflation, however, is not as extreme as the inflation that the economy started with. Note that the central bank's hands are tied with initial inflation. This is the main difference between my results and Yun (2005). When you assume away production distortions, then you do not need to include the firm's pricing decision. When you do not include this pricing decision, then inflation can freely jump. This is not possible in my setup.

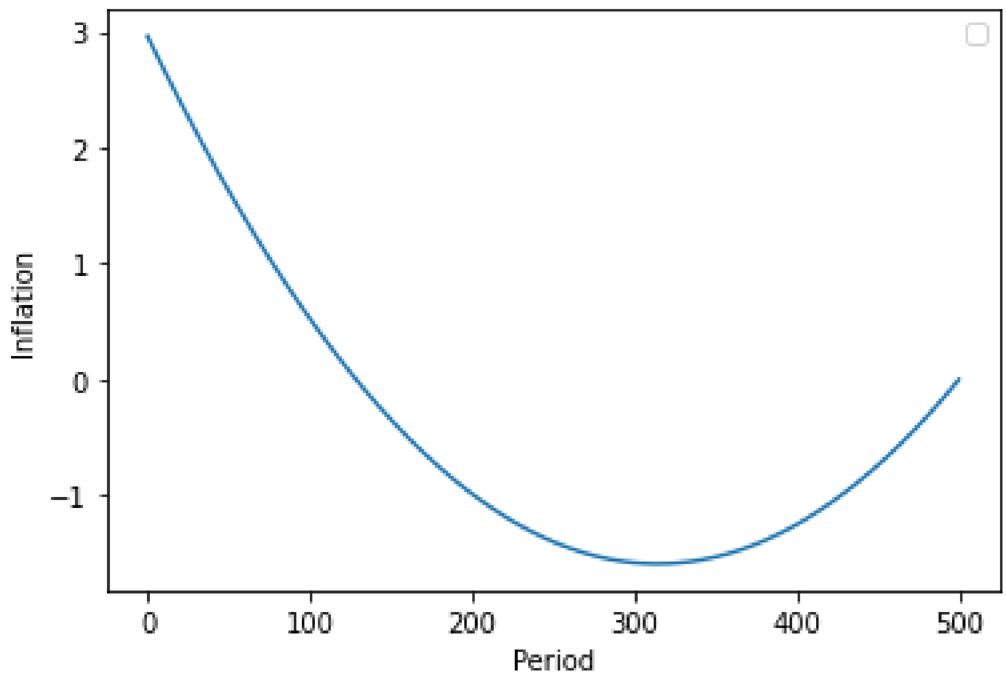

This is done by having the interest be relatively high at the outset, and then lowering it to guide inflation back to its long-run value.

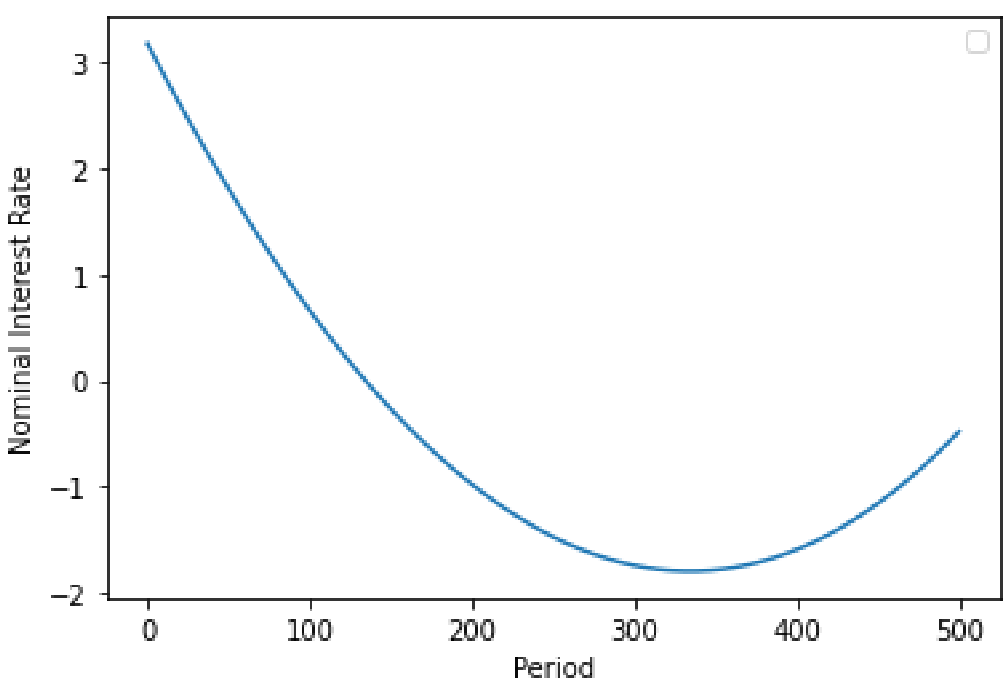

You can see the effect through marginal cost as well, which rises before lowering to its long run value.

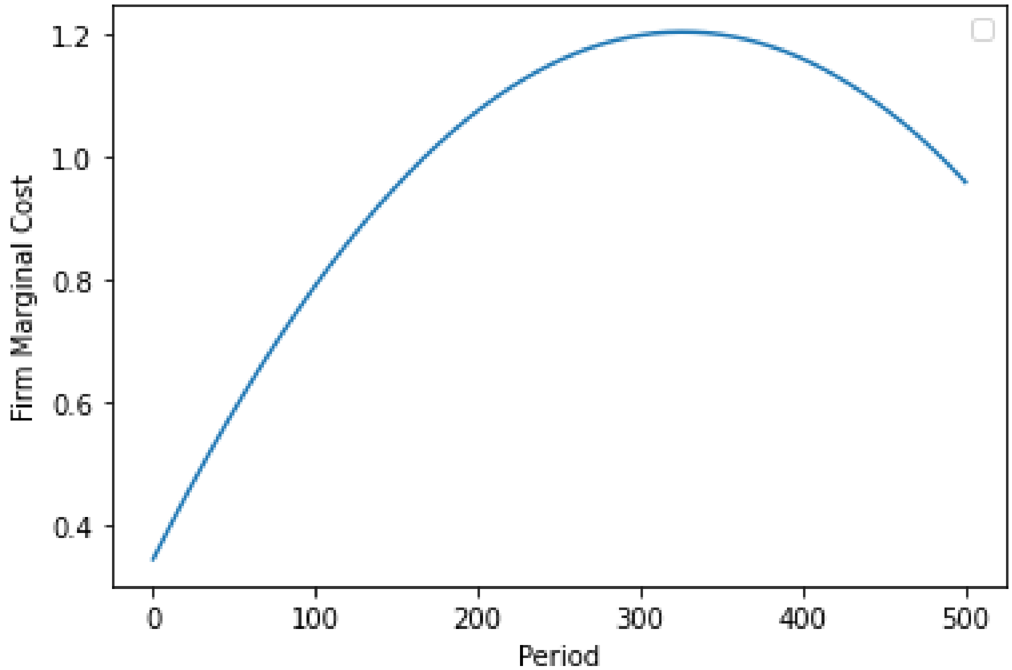

Finally, there are the dynamics of firm expectations. The central bank acts to at first have high expectations of revenues and costs, with a smaller difference between the two. This lowers economic output. Then the central bank causes both expectations to fall, though it does so to a greater extent with expectations of future costs. This causes the gap between the two to increase. These dynamics, again, are lost if linearization is doen to the model.

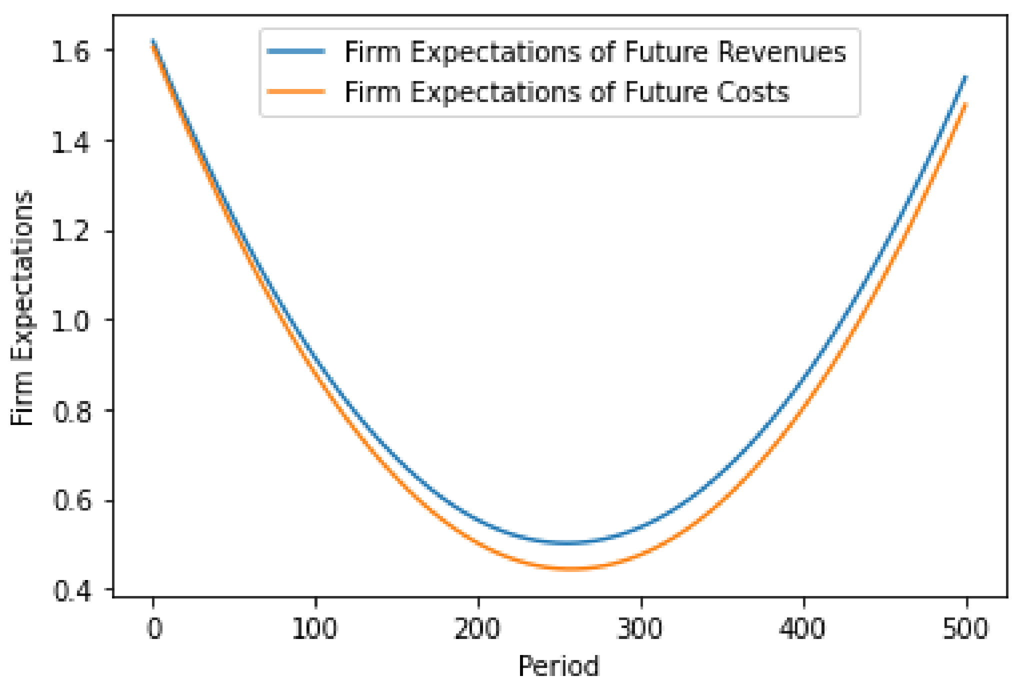

3.4.2 The Interest Rate

I illustrate next the interest rate as a function of relevant state variables. Note that here the interest rate is given as a function of inflation (π) at particular values of price dispersion (v). Remember that v=1 implies total efficiency, with v being bounded below by 1.

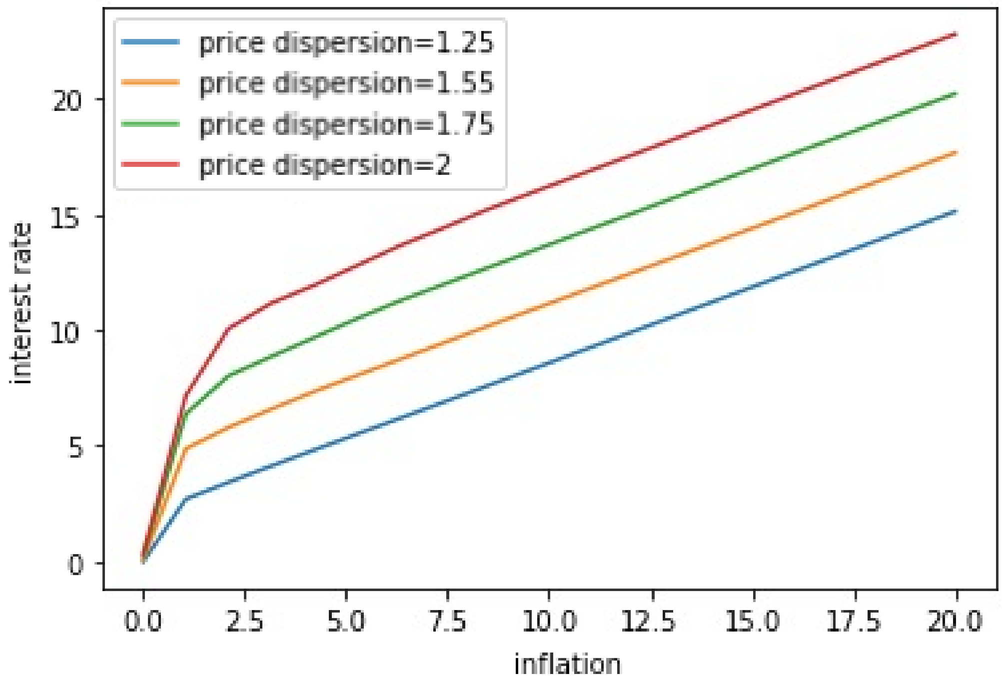

A few things are of note here. First, is the confirmation that nonlinearities matter. It also supports the idea that the price dispersion term does in fact influence optimal monetary policy. Traditionally, this term is considered orthogonal to the policy decision, or at the very least it is discarded as being second order. The relationship is intuitive: as inefficiencies associated with price dispersion become more pronounced, then the central bank should be more aggressive overall.

4 Conclusion

Although optimal monetary policy is important to central banks for the practical problems that they face, computational limits have been a barrier. With some important exceptions that have used computational techniques in a case-by-case basis, most of the applications addressed have been simplified to make them easier to solve by hand or easier to compute. In this paper, I applied recent machine learning techniques that may be used to answer a wide range of other questions where computational limitations have been a major issue. I illustrated these techniques as applied to the question of how firm expectations and price distortions should jointly influence optimal monetary policy. In a fully non-linear New Keynesian Model with price and labor distortions, I found that price level stabilization around a long-run value is best when distortions are small. However, the farther we start from the long-run value, the policy response should be more nonlinear and more aggressive. I showed that interest rate policy should take into consideration both price dispersion and firm expectations on future costs, the latter directly relating to distortions from monopolistic competition. I showed how in the face of a large initial price distortion, the central bank is aggressive enough to temper the price dispersion, even at the cost of temporary deflation. This deflation occurs after the initial shock, however, because the production distortions limit the path inflation can take.

The next step then would be to apply the methods used here to more complex environments to pose questions that are difficult to answer using the normal assumptions of the optimal monetary policy literature. Future research could continue to leverage these machine learning techniques to answer further questions, including, for example, to what extent optimal policy should display inertia or adjust rapidly, how skewed or non-normal uncertainty should matter, and how to incorporate more behavioral-style preferences like rational inattention. While my application is focused on optimal monetary policy, the techniques can be applied to macroeconomics more broadly, industrial organization, finance, labor, and others. It can efficiently solve models with hundreds of state variables. It can be used for both reduced-form and structural dynamic models, and can be used to get information from unstructured sources of data like text. Machine learning methods are particularly promising for multi-agent modeling, efficiently estimating global solutions, business cycle and financial forecasting, and using text data, among many other things.

A Underlying Model

The underlying model is a continuous time New Keynesian model with labor as the production input, Calvo pricing and monopolistic competition, and no uncertainty. I will summarize the important aspects of this model for clarity in the ensuing analysis. For fuller derivations, please reference Fernandez-Villaverde et al (2012), who work out the generaly dynamics of this kind of model in continuous time, though through specifying a Taylor rule type monetary policy.

A.1 Consumption

A representative consumer seeks to maximize lifetime utility, represented by a utility function separable in consumption (c) and hours worked (n).

$$\int_0^\infty e^{-\rho t}\{\ln(c_t) - \psi\frac{n_t^{1+\gamma}}{1+\gamma}\}$$

Where ρ is the subjective rate of time preference, ψ is the disultity of labor, and γ is the inverse of Frisch labor supply elasticity.

The household can trade on Arrow securities and on nominal government bonds b_t at a nominal interest rate r_t. The household earns a disposable income of $r_t b_t + p_t w_t n_t + p_t \Pi_t$, where p_t is the price of the consumption good, w_t is the real wage, and Π_t represents firm profits. Household financial wealth evolves as follows. Note that $\dot{b}$ refers to $\mathrm{d}b/\mathrm{d}t$

$$\dot{b} = r_t b_t - p_t c_t + p_t w_t n_t + p_t \Pi_t$$

Inflation is defined as:

$$\pi_t = \frac{\dot{p}}{p_t}$$

Let us define real financial wealth as $a_t \equiv \frac{b_t}{p_t}$. Real wealth then evolves as follows:

$$\dot{a} = \frac{r_t b_t - p_t c_t + p_t w_t n_t + p_t \Pi_t}{p_t} - \frac{b_t}{p_t^2}\pi_t p_t$$
$$= ((r_t - \pi_t)a_t - c_t + w_t n_t + \Pi_t$$

A.2 Production

Final good production is competitive. A representative producer purchases intermediate goods and produces the final good with the production function:

$$y_t = (\int_0^1 y_{it}^{\frac{\varepsilon-1}{\varepsilon}})^{\frac{\varepsilon}{\varepsilon-1}}$$

where ε is the elasticity of substitution.

The input demand functions associated with the final good producer's problem are given as:

$$y_{it} = (\frac{p_{it}}{p_t})^{-\varepsilon} y_t \qquad \forall i$$
$$p_t = (\int_0^1 p_{it}^{1-\varepsilon} di)^{\frac{1}{1-\varepsilon}}$$

Each intermediate firm i produces differentiated goods out of labor using:

$$y_{it} = A n_{it}$$

where n_{it} is the amount of labor rented by firm i and A is a technology paramter. The intermediate good producer is a monopolistic firm and price setting is carried out via the Calvo formulation. At rate θ, intermediate firm i get the opportunity to reset their price. Any firm which does not receive such signal does not have the opportunity to change their price. The probability of receiving such a signal is independent of the timing of the last signal.

Prices are set to maximize expected discounted profits. Note that an expectation operator is used because although there is no uncertainty in the aggregate, because the timing of individual firm price changes is random there is uncertainty on the individual firm level. Note also that real marginal cost, $\text{mc}_\tau = w_\tau / A$, is common across firms because firms share a common technological parameter.

The intermediate firm's problem is:

$$\max_{p_{it}} \text{E}_t \int_t^\infty \frac{\lambda_\tau}{\lambda_t} e^{-\theta(\tau-t)} [\frac{p_{it}}{p_\tau} y_{i\tau} - \text{mc}_\tau y_{i\tau}] d\tau$$

where λ_τ is the time t value of consumption in period τ to the household.

The first order conditions of the firm is as follows. The ratio of the optimal new price, common across all firms able to reset their prices, and the prices of the final good, is given by:

$$\frac{p_{it}}{p_t} = \frac{\varepsilon}{\varepsilon - 1} \frac{\Sigma_{Ct}}{\Sigma_{Rt}}$$

where:

$$\Sigma_{Rt} = \int_t^\infty \lambda_\tau e^{-\theta(\tau-t)} (\frac{p_t}{p_\tau})^{1-\varepsilon} y_\tau d\tau$$

represents expected present discounted value of total future revenue, and

$$\Sigma_{Ct} = \int_t^\infty \lambda_\tau e^{-\theta(\tau-t)} \text{mc}_\tau (\frac{p_t}{p_\tau})^{-\varepsilon} y_\tau d\tau$$

represents expected present discounted value of total future costs. In both cases the λ term refers to the discount factor in terms of consumer valuation. This will refer to a stochastic discount factor in later sections. I maintain the same framework here for consistency.

This means that the optimal reset price equals the desired markup $\frac{\varepsilon}{\varepsilon-1}$ multiplied by the ratio of the future cost index Σ_{Ct} and future revenue index Σ_{Rt}. Because any firm has virtually no effect on aggregate terms, both of these indexes are exogenous to the firm.

The other variable of interest is that for price dispersion, v, which can be viewed as the inefficiency associated with not all firms having the same price at the same time. In practical terms it acts as a wedge between production in terms of inputs and in terms of output after aggregation.

$$y_t = \frac{An_t}{v_t}$$

where

$$v_t = \int_0^1 (\frac{p_{it}}{p_t})^{-\varepsilon} di$$

Note that $1 \leqslant v_t$, where $v_t = 1$ would imply efficiency. Price dispersion acts as the point of inefficiency flowing from staggered price setting. Price dispersion can also be thought of as a misalignment between decisions made on the basis of marginal cost and those made on the basis of marginal utility. This theme of a wedge between benefits and costs will be revisited when I examine the central bank's problem.

A.3 Equilibrium Results

The equilibrium results from the above model will be used later to formulate the constraints of the central bank's problem. These equilibrium results are expressed by the following differential equations and equality constraints.

$$\begin{aligned}
\dot{\Sigma_R} &= (\theta - (\varepsilon - 1)\pi_t)\Sigma_{Rt} - 1 \\
\dot{\Sigma_C} &= (\theta - \varepsilon\pi_t)\Sigma_{Ct} - \mathrm{mc}_t \\
\dot{v} &= \theta(1 + \pi_t \frac{1-\varepsilon}{\theta})^{-\frac{\varepsilon}{1-\varepsilon}} + (\varepsilon\pi_t - \theta)v_t \\
\dot{\lambda} &= (\rho - r_t + \pi_t)\lambda_t \\
\mathrm{mc}_t &= \psi(A\lambda_t)^{-(1+\gamma)} v_t^{\gamma} \\
(1 + \pi_t \frac{1-\varepsilon}{\theta})^{\frac{1}{1-\varepsilon}} &= \frac{\varepsilon}{\varepsilon - 1} \frac{\Sigma_{Ct}}{\Sigma_{Rt}}
\end{aligned}$$

The above conditions correspond to the development of the firm's revenue and cost expectations, aggregate price dispersion, the household Euler equation, and equations which determine equilibrium marginal cost and inflation. An important point here is that marginal cost acts as a sort of key. Indeed, for any given level of marginal cost and state variables the partial equilibrium is determined for consumption and labor (here essentially the consumption and labor decision is equivalent to a joint determination of marginal utility and marginal cost). This latter point will be exploited for the central bank's problem and will continue to act as a guiding principle throughout all our analysis later on.

There are other points of interest here. The sign of the relationship between the time derivative of the future cost and revenue indexes and the indexes themselves is dependent on the current level of inflation. We see that for very low, near zero or negative inflation levels, that higher current values of these indexes increases the time derivative, and that higher inflation levels flip that relationship. Note also that there are knife-edge cases of inflation where the time derivative loses all relationship with current levels.

B The Traditional Way

Before continuing, it may be useful to compare our approach with what is commonly done in the legacy of Woodford (2003). The central bank would be attempting to minimize a quadratic loss function. The following is a typical case:

$$\min_{r_t} \frac{1}{2}\int_0^\infty e^{-\rho t}[\alpha_\pi \pi_t^2 + \alpha_x x_t^2]dt$$

subject to:

$$\begin{aligned} d\pi_t &= (\rho_\pi(\pi_t - \bar{\pi}) - \kappa_x x_t)dt + \rho_\pi dZ_t \\ dx_t &= \frac{1}{\gamma}[r_t - \tilde{r} - (\pi_t - \bar{\pi})] + \rho_x dZ_t \end{aligned}$$

Where x_t is the output gap and π_t is inflation (here it is assumed that the natural rate of inflation is zero).

A few things are obvious. First, this derivation is dependent on the underlying model in a way that must be derived and is not apparent a priori. It also relies on shocks being relatively small and assumes symmetry of effects as well as Gaussian shocks. The approach of this current paper maintains generality in these aspects. In addition, if one would like to change the underlying model, for example to introduce a distorted steady state, non-standard preferences, or rational inattention, then our approach makes this easier to accommodate in a tractable way.

C Equivalency of Optimal Control Techniques

I will now prove that treating functions of state variables as controls with a Lagrangian is equivalent to working directly with them as state variables. This was used implicitly in the central bank's problem.

Consider the problem:

$$
\begin{aligned}
maxF(X, Y, Z) \\
s.t. \\
\dot{y} = G(X, Y, Z) \\
Z = H(Y)
\end{aligned}
$$

Consider the method of substitution:

$$
\begin{aligned}
maxF(X, Y, H(Y)) \\
s.t. \\
\dot{y} = G(X, Y, H(Y))
\end{aligned}
$$

Take the Hamiltonian:

$$H = F(X, Y, H(Y)) + \Lambda G(X, Y, H(Y))$$

FOC:

$$
\begin{aligned}
x &: F_x = 0 \\
y &: -\dot{\lambda} + \rho\Lambda = F_y + F_{H(y)} H_y + \Lambda[G_y + G_{H(y)} H_y]
\end{aligned}
$$

Let us take the Lagrangian method:

$$H = F(X, Y, Z) + \Lambda G(X, Y, Z) + \mu(H(Y) - Z)$$

FOC:

$$\begin{aligned}
x &: F_x + \Lambda G_x = 0 \\
z &: F_z + \Lambda G_z - \mu = 0 \\
y &: -\dot{\lambda} + \rho\Lambda = F_y + \Lambda G_y + \mu H_y \\
\mu &: G(Y) = Z
\end{aligned}$$

This implies:

$$\begin{aligned}
&F_z + \Lambda G_z = \mu \\
&-\dot{\lambda} + \rho\Lambda = F_y + \Lambda G_y + \mu H_y \\
&-\dot{\lambda} + \rho\Lambda = F_y + \Lambda G_y + [F_z + \Lambda G_z] H_y \\
&-\dot{\lambda} + \rho\Lambda = F_y + F_z H_y + \Lambda(G_y + G_z H_y) \\
&-\dot{\lambda} + \rho\Lambda = F_y + F_z H_y + \Lambda[G_y + G_z H_y] \\
&-\dot{\lambda} + \rho\Lambda = F_y + F_{H(y)} H_y + \Lambda[G_y + G_{H(y)} H_y]
\end{aligned}$$

We see therefore that for the purposes of optimal control analysis that we can consider any function of the state variables as a control variable without loss.

D Transversality Conditions

D.1 Deterministic Case

Transversality Conditions for the Deterministic Case:

$$\lim_{T\to\infty} e^{-pT}\Lambda_{\Sigma_R} \geqslant 0$$

$$\lim_{T\to\infty} e^{-pT}\Lambda_{\Sigma_C} \geqslant 0$$

$$\lim_{T\to\infty} e^{-pT}\Lambda_{vt} \geqslant 0$$

$$\lim_{T\to\infty} e^{-pT}\Lambda_{\lambda t} \geqslant 0$$

$$\lim_{T\to\infty} e^{-pT}\Lambda_{\Sigma_R}\Sigma_{Rt} = 0$$

$$\lim_{T\to\infty} e^{-pT}\Lambda_{\Sigma_C}\Sigma_{Ct} = 0$$

$$\lim_{T\to\infty} e^{-pT}\Lambda_{vt}v_t = 0$$

$$\lim_{T\to\infty} e^{-pT}\Lambda_{\lambda t}\lambda_t = 0$$

References

[1] Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. "Income and wealth distribution in macroeconomics: A continuous-time approach". Working Paper 23732, National Bureau of Economic Research, 2017,

[2] Azinovic, M., Gaegauf, L., and Scheidegger, S. "Deep equilibrium nets". University of Lausanne, 2020

[3] Bach, F. "Breaking the curse of dimensionality with convex neural networks". Journal of Machine Learning Research, 18(1):629â681, 2017.

[4] Benigno, Pierpaolo, & Woodford, Michael, âLinear-Quadratic Approximation of Optimal Policy Problems,â NBER working paper no. 12672, revised August 2008.

[5] Bilbiie, Florin O. & Ragot, Xavier. "Optimal monetary policy and liquidity with heterogeneous households", Review of Economic Dynamics, 2020.

[6] Challe, Edouard, "Uninsured Unemployment Risk and Optimal Monetary Policy in a Zero-Liquidity Economy." American Economic Journal: Macroeconomics, 12 (2): 241-83, 2020.

[7] Christiano, L., Motto, R., and Rostagno, M. "Notes on ramsey-optimal monetary policy". Mimeo, 2007.

[8] Clarida, Richard, Gali, Jordi Gali, and Gertler, Mark, "The Science of Monetary Policy: A New Keynesian Perspective," Journal of Economic Literature, vol. 37 (December), pp. 1661-707, 1999.

[9] Curdia, Vasco & Woodford, Michael, "Credit Frictions and Optimal Monetary Policy," Journal of Monetary. Economics, Elsevier, vol. 84(C), pages 30-65, 2016.

[10] Duarte, V. "Machine learning for continuous-time economics". Working paper, 2018.

[11] El Karoui, N., Peng, S. Peng, and Quenez, M.C. "Backward stochastic differential equations in finance", Mathematical Finance, 7, 1-71, 1997.

[12] El Karoui, N. & Mazliak, L (editors), "Backward stochastic differential equations", Pitman research notes in mathematics series, 1997.

[13] Faulwasser, T. et al. "Unconventional monetary policy in a nonlinear quadratic model." Studies in Nonlinear Dynamics and Econometrics, 2020.

[14] Fernandez-Villaverde, Jesus, Hurtado, Samuel, and Nuno, Galo, "Financial Frictions and the Wealth Distribution," NBER Working Papers 26302, National Bureau of Economic Research, Inc, 2019.

[15] Fernandez-Villaverdez, Jesus, Posch, Olaf, and Rubio-Ramirez, Juan F."Solving the new Keynesian model in continuous time," Meeting Papers 829, *Society for Economic Dynamics*, 2012.

[16] Gali, Jordan. "New Perspectives on Monetary Policy, Inflation, and the Business Cycle," in Mathias Dewatripont; Lars P. Hansen and Stephen J. Turnovsky, eds, *Advances in economics and econometrics: Theory and applications*, Vol. 3. Cambridge: Cambridge University Press, pp. 151-97, 2003.

[17] Gali, Jordi and Monacelli, Tommaso. "Monetary Policy and Exchange Rate Volatility in a Small Open Economy," Review of Economic Studies, Vol. 72, pp. 707-734, 2005.

[18] Giannoni, Marc P., and Woodford, Michael, "Optimal Interest-Rate Rules: I. General Theory," *National Bureau of Economic Research Working Paper Series*, No. 9419, January 2003

[19] Giannoni, Marc P., and Woodford, Michael, "Optimal Inflation-Targeting Rules," in Ben S. Bernanke and Michael Woodford, eds., Inflation Targeting. Chicago: University of Chicago Press, pp. 93-172, 2005.

[20] Goodfriend, Marvin and King, Robert, "The New Neoclassical Synthesis and the Role of Monetary Policy," in Ben S. Bernanke and J. J. Rotemberg, *NBER macroeconomics annual 1997*, Vol. 12. Cambridge, MA: MIT Press, pp. 231-83, 1997.

[21] Huh, H., Lee, H.H., Lee, "N. Nonlinear Phillips curve, NAIRU and monetary policy rules"

Empirical Econ., 37, pp. 131-151, 2009.

[22] Kahn, Aubhik; King, Robert G. and WOlman, Alexander L. "Optimal Monetary Policy." Review of Economic Studies, 70(4), pp. 825-60, 2003.

[23] Kim, Jinill and Ruge-Murcia, Francisco J., "Extreme Events and Optimal Monetary Policy". International Economic Review, Vol. 60, Issue 2, pp. 939-963, May 2019

[24] Lagari, Pola Lydia Tsoukalas, Lefteri Safarkhani, Salar Lagaris, Isaac. (2020). Systematic Construction of Neural Forms for Solving Partial Differential Equations Inside Rectangular Domains, Subject to Initial, Boundary and Interface Conditions. International Journal on Artificial Intelligence Tools. 29. 10.1142/S0218213020500098.

[25] Levin, Andrew T., Onatski, Alexei, Williams, John C., and Williams, Noah, "Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models," in Mark Gertler and Kenneth Rogoff, eds., NBER Macroeconomics Annual 2005. Cambridge, Mass. MIT Press, pp. 229-87, 2005.

[26] Lu, X. Meng, Z. Mao, G. E. Karniadakis. DeepXDE: A deep learning library for solving differential equations. SIAM Review, 63(1), 208â228, 2021.

[27] Lucas, Robert E., Jr. and Stokey, Nancy L., "Optimal Fiscal and Monetary Policy in an Economy without Capital," Journal of Monetary Economics, 12(1), pp. 55-93, 1983.

[28] Maliar, L., Maliar, S., and Winant, P. "Will artificial intelligence replace computational economists any time soon?" CEPR Discussion Papers DP14024, CEPR, 2019.

[29] Mishkin, Frederic, "Monetary policy strategy: Lessons from the crisis", Prepared for the ECB (European Central Bank) Central Banking Conference on Monetary Policy Revisited: Lessons from the Crisis, Frankfurt, 2010.

[30] Nakamura-Zimmerer, Tenavi and Gong, Qi and Kang, Wei, "Adaptive Deep Learning for High-Dimensional Hamilton–Jacobi–Bellman Equations" , SIAM Journal on Scientific Computing, Vol. 43, pp. A1221âA1247, Jan. 2021.

[31] Nakamura-Zimmerer, Tenavi and Gong, Qi and Kang, Wei, “QRnet: Optimal Regulator Design With LQR-Augmented Neural Networks,” in IEEE Control Systems Letters, vol. 5, no. 4, pp. 1303-1308, Oct. 2021

[32] Nobay, A. R. & Peel, D. A., “Optimal monetary policy with a nonlinear Phillips curve,” Economics Letters, Elsevier, vol. 67(2), pages 159-164, 2000.

[33] Pardoux, Etienne & Tang, Shanjian. “Forward-backward stochastic differential equations and quasilinear parabolic pdes.” Probability Theory and Related Fields, 114(2):123–150, 1999.

[34] Pham, Huyen, ”On some recent aspects of stochastic control and their applications”. Probability Surveys 2, 506–549, 2005.

[35] M. Raissi, “Deep hidden physics models: Deep learning of nonlinear partial differential equations”, arXiv preprint arXiv:1801.06637, 2018.

[36] Ramsey, Frank P., “A Contribution to the Theory of Taxation,” Economic Journal, 37(1), pp. 47-61, 1927.

[37] Rotemberg, Julio, & Woodford, Michael, “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy,” in Ben S. Bernanke and Julio J. Rotemberg, eds., NBER Macroeconomics Annual 1997. Cambridge, Mass.: MIT Press, pp. 297-346, 1997.

[38] Scheidegger, S. and Bilionis, I. “Machine learning for high-dimensional dynamic stochastic economies”. UCL, 2017.

[39] Schmitt-Grohe, Stephanie & Uribe, Martin,“Optimal Fiscal and Monetary Policy in a Medium-Scale Macroeconomic Model,” in Mark Gertler and Kenneth Rogoff, eds., NBER Macroeconomics Annual 2005. Cambridge, Mass.: MIT Press, pp. 383-425, 2005.

[40] Shi, Haizhou. “Backward Stochastic Differential Equations in Finance.” Wadham College University of Oxford, 2010.

[41] Socgnia, V. K.& Menoukeu-Pamen, O. ”An infinite horizon stochastic maximum principle for discounted control problem with Lipschitz coefficients”, Journal of Mathematical Analysis and Applications, 422, 684-711, 2015.

[42] Swanson, E. T. “Optimal non-linear policy: signal extraction with a non- normal prior”, Journal of Economic Dynamics and Control 30, 185-203, 2006.

[43] Swanson, Eric & Anderson, Gary & Levin, Andrew. (2006). Higher-Order Perturbation Solutions to Dynamic, Discrete-Time Rational Expectations Models. Federal Reserve Bank of San Francisco, Working Paper Series. 1.000-31.000. 10.24148/wp2006-01.

[44] Woodford, Michael, “Optimal Monetary Policy Inertia,” National Bureau of Economic Research, Inc., NBER Working Papers: No. 7261, 1999.

[45] Woodford, Michael, “Interest and Prices: Foundations of a Theory of Monetary Policy”, Princeton University Press, Princeton, 2003.

[46] Wu, Ho-Mou & Li, Minghao, “Optimal Monetary Policy with Asymmetric Shocks and Rational Inattention” 2016. Working paper

[47] Yong, Jiongmin, & Zhong, Xun Yu. “Stochastic Controls: Hamiltonian Systems and HJB Equations”. Springer, 1999.

[48] Yun, Tack, “Optimal Monetary Policy with Relative Price Distortions,” American Economic Review 95(1), 89-109, 2005.