Optimal Monetary Policy in the New Keynesian Model

Eric Sims
University of Notre Dame

Spring 2024

1 Introduction

These notes describe optimal monetary policy in the basic New Keynesian model. We take the basic model to be characterized by two principal equations: the IS curve and the Phillips Curve, both written in terms of the output gap:

$$\pi_t = \gamma x_t + \beta \mathbb{E}_t \pi_{t+1} \tag{1}$$

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^f\right) \tag{2}$$

In the most basic model from class, $\gamma = \frac{(1-\phi)(1-\phi\beta)}{\phi}(\sigma+\chi)$, where σ is the inverse intertemporal elasticity of substitution and χ is the inverse Frisch labor supply elasticity. r_t^f is the natural rate of interest, and we assume it to be exogenous (this can be written as a function of other exogenous shocks; e.g. productivity or government spending shocks).

The Federal Reserve in the US has a "dual mandate" in that it wants to promote price stability and full employment. In the context of the basic model, we can take price stability to mean minimizing volatility in inflation, and full employment to mean minimizing volatility in the output gap. We can define a loss function for the central bank as follows:

$$\mathcal{L} = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{1}{2}\left[\pi_{t+j}^2 + \omega x_{t+j}^2\right] \tag{3}$$

The flow objective is $\pi_t^2 + \omega x_t^2$. Since these are deviations, we can interpret these squares as variances. So, the loss function is the weighted sum of the variances of inflation and the output gap. $\omega \geq 0$ is simply the relative weight the policymaker puts on output fluctuations. The 1/2 outside is just a scaling parameter that won't affect the optimum.

While a loss function like this seems to accord well with policymaking in practice, it can actually be derived from taking a second-order approximation of the value function of the representative household in the model. In doing so, the relative weight on the output gap, ω, can be pinned down by deeper parameters of the underlying model, rather than just taken as given. In particular, in the most basic version of the NK model, the value of ω ought to be:

$$\omega = \frac{\gamma}{(1+\chi\epsilon)\epsilon}$$

Given that ϵ is likely pretty large, and most standard calibrations would have γ be pretty small, the relative weight on output gap fluctuations ought to be quite small if one is basing it on the approximation to the representative household's value function. But, for what I am going to do, I will for the most part just take ω as given – it could be anything.

1.1 Aside: Monopoly Price-Setting and Steady State Distortions

The loss function I wrote down above can be derived from a second-order approximation to the welfare function of the representative household *in the case in which the non-stochastic steady state of the model is efficient*. If there is any kind of steady state distortion, the loss function we derived above will not be valid. This ought to make some sense: what the planner should care about is bringing output to its efficient level, call it y_t^e in log deviation form. In general, because of the monopoly distortion, $y_t^f \neq y_t^e$.

The steady-state distortion in the basic NK model comes from monopoly power in price-setting. In the steady state of the model, optimal behavior by intermediate firms gives rise to a labor demand condition that satisfies:[1]

$$w = \frac{\epsilon - 1}{\epsilon}$$

Recall that $\epsilon > 1$ and steady state labor productivity (the marginal product of labor) is normalized to one. Labor market-clearing (i.e. combining labor demand with supply) gives:

$$\theta N^{\chi} = C^{-\sigma}\frac{\epsilon - 1}{\epsilon}$$

This is *not* the labor market condition the planner would choose. The planner would equate the marginal rate of substitution between labor and consumption to the marginal product of labor (the latter of which is just unity in the basic model). The planner would therefore have:

$$\theta N^{\chi} = C^{-\sigma}$$

For the quadratic loss function above to be valid, the steady state about which the approximation is taken needs to be efficient. The steady state can be made efficient provided there are appropriate subsidies in place from the government. In particular, suppose that the government levies both a (constant) consumption tax, τ^C, and a (constant) labor tax, τ^N. The labor supply condition of the household would be:

$$\theta N^{\chi} = C^{-\sigma}\frac{1-\tau^N}{1+\tau^C}w$$

[1] For this subsection, and indeed this entire set of notes, I am assuming that steady state inflation is $\Pi = 1$ (gross, so $\pi = 0$, net).

Plugging in the labor demand condition:

$$\theta N^{\chi} = C^{-\sigma}\frac{1-\tau^N}{1+\tau^C}\frac{\epsilon-1}{\epsilon}$$

The fiscal authority can use a consumption or labor tax to implement the efficient allocation provided:

$$\frac{1-\tau^N}{1+\tau^C}\frac{\epsilon-1}{\epsilon} = 1$$

For example, if the fiscal authority only uses the labor tax (so $\tau^C = 0$), we need:

$$1-\tau^N = \frac{\epsilon}{\epsilon-1}$$

This means that $\tau^N < 0$; in other words, the fiscal authority needs to subsidize labor to offset the monopoly distortion. In contrast, suppose that the fiscal authority instead wants to use the consumption tax only (so $\tau^N = 0$). Then we must have:

$$1+\tau^C = \frac{\epsilon-1}{\epsilon}$$

This means that $\tau^C < 0$; in other words, the fiscal authority needs to subsidize consumption to offset the monopoly distortion. In the steady state, a consumption tax and a labor tax are isomorphic. What is important is that the fiscal authority subsidizes consumption and/or labor to bring the economy's steady-state level of production up to the efficient level.

2 Optimal Policy Under Discretion

We will consider two versions of an optimal policy problem: discretion and commitment. Under discretion, the policymaker will choose its instrument to minimize the period loss function, taking all future values in the loss function as given. Basically, the policymaker is unable to commit to anything about the future, and so it just minimizes the loss function in the present, taking the future as given. Under commitment, in contrast, the policymaker picks an entire sequence of its policy instrument to minimize the entire loss function (i.e. the full present discounted value). The problem is much easier to solve under discretion.

The instrument the policymaker has at its disposal is the short-term nominal interest rate, i_t. Rather than closing the model with an instrument rule like a Taylor rule, we are instead going to derive an optimal targeting rule. The central bank wants to pick i_t to minimize the loss function, taking the Phillips Curve and IS equation as constraints. In particular, under discretion the problem is:

$$\min_{i_t} \ \mathcal{L} = \frac{1}{2}\pi_t^2 + \frac{\omega}{2}x_t^2$$

s.t.

$$\pi_t = \gamma x_t + \beta \mathbb{E}_t \pi_{t+1}$$
$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^f\right)$$

The policymaker takes r_t^f, which is exogenous, as given. It also takes $\mathbb{E}_t \pi_{t+1}$ and $\mathbb{E}_t x_{t+1}$ as given as well, since the policymaker will re-optimize in the future. The IS equation therefore already expresses x_t as a function of things the policymaker takes as given as well as its instrument. Plugging that into the Phillips Curve to eliminate x_t gives:

$$\pi_t = \gamma \mathbb{E}_t x_{t+1} - \frac{\gamma}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^f\right) + \beta \mathbb{E}_t \pi_{t+1}$$

We can therefore write the policy problem as an unconstrained problem of choosing i_t:

$$\min_{i_t} \ \mathcal{L} = \frac{1}{2}\left[\gamma \mathbb{E}_t x_{t+1} - \frac{\gamma}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^f\right) + \beta \mathbb{E}_t \pi_{t+1}\right]^2 + \frac{\omega}{2}\left[\mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^f\right)\right]^2$$

The derivative of the loss function with respect to i_t is:

$$\frac{\partial \mathcal{L}}{\partial i_t} = -\frac{\gamma}{\sigma}\pi_t - \frac{\omega}{\sigma} x_t$$

This follows because the terms in brackets are just π_t and x_t, respectively. Focusing on the derivative with respect to i_t in the inflation term, for example, we bring the exponent down (the 2s cancel) and we subtract one from the exponent, which leaves π_t, and we multiply by the derivative of the inside with respect to i_t, which is $-\frac{\gamma}{\sigma}$. Setting the derivative equal to zero and simplifying yields:

$$x_t = -\frac{\gamma}{\omega}\pi_t \tag{4}$$

(4) is the optimality condition. It has a "lean against the wind" flavor – when inflation is high, the policymaker allows the output gap to go negative, where the amount depends on the relative weight on the gap in the objective function, ω, as well as the slope of the Phillips curve, γ. One can include (4) as the "policy rule" and the third equation in the model instead of a Taylor rule. (4) is an implicit inflation target: each period, the policymaker targets a negative relationship between the current inflation rate and the current output gap.

To solve for the required path of the interest rate to implement this lean against the wind condition, plug it into Phillips curve to eliminate x_t:

$$\pi_t = -\frac{\gamma^2}{\omega}\pi_t + \beta \mathbb{E}_t \pi_{t+1}$$

Or:

$$\mathbb{E}_t \pi_{t+1} = \frac{\left(1 + \frac{\gamma^2}{\omega}\right)}{\beta}\pi_t$$

This is an explosive difference equation in inflation. The coefficient $\frac{\left(1+\frac{\gamma^2}{\omega}\right)}{\beta} > 1$. This means that the only non-explosive solution for inflation is $\pi_t = 0$. But from lean against the wind condition, if $\pi_t = 0$, then $x_t = 0$. This means that, under discretion in this model, we can get to the global minimum of the loss function at 0! What does this imply about the behavior of the nominal interest rate? If $\pi_t = x_t = 0$ in period t, then agents will expect the same for $t+1$: $\mathbb{E}_t \pi_{t+1} = \mathbb{E}_t x_{t+1} = 0$. For the IS equation to hold, this requires:

$$i_t = r_t^f \tag{5}$$

In other words, optimal monetary policy under discretion entails moving the interest rate one-for-one with the natural rate of interest. This results in both inflation and the output gap being completely stabilized.

3 Optimal Policy Under Commitment

The problem under commitment gets messier. This is because the central bank doesn't just get to choose the current instrument, i_t, it gets to choose future instruments, like i_{t+1} (and so on). This is relevant because the central bank doesn't need to take future inflation and the output gap as given when choosing the current interest rate.

To make things easier, let's write the problem in the following way:

$$\mathcal{L} = \frac{1}{2}\pi_t^2 + \frac{1}{2}\omega x_t^2 + \beta \mathbb{E}_t \left[\frac{1}{2}\pi_{t+1}^2 + \frac{1}{2}\omega x_{t+1}^2\right] + \beta^2 \mathbb{E}_t \left[\frac{1}{2}\pi_{t+2}^2 + \frac{1}{2}\omega x_{t+2}^2\right] + \dots$$

s.t.

$$\pi_{t+j} = \gamma \mathbb{E}_{t+j} x_{t+j+1} - \frac{\gamma}{\sigma}\left(i_{t+j} - \mathbb{E}_{t+j}\pi_{t+j+1} - r_{t+j}^f\right) + \beta \mathbb{E}_{t+j}\pi_{t+j+1} \tag{6}$$

$$x_{t+j} = \mathbb{E}_{t+j} x_{t+j+1} - \frac{1}{\sigma}\left(i_{t+j} - \mathbb{E}_{t+j}\pi_{t+j+1} - r_{t+j}^f\right) \tag{7}$$

The problem of the central bank is to pick i_t, i_{t+1}, i_{t+2}, ... to minimize $\mathcal{L}$, subject to inflation and the output gap being determined in each period according to the IS equation and Phillips curves.

When choosing i_t, there is no direct effect on future inflation or the future gap (since there are no endogenous state variables). The derivative of the loss function with respect to i_t is:

$$\frac{\partial \mathcal{L}}{\partial i_t} = \frac{\partial \mathcal{L}}{\partial \pi_t}\frac{\partial \pi_t}{\partial i_t} + \frac{\partial \mathcal{L}}{\partial x_t}\frac{\partial x_t}{\partial i_t}$$

This is making use of the chain rule – $\frac{\partial \mathcal{L}}{\partial \pi_t}$ is the derivative of the loss function with respect to π_t, and $\frac{\partial \pi_t}{\partial i_t}$ is the derivative of inflation with respect to i_t (i.e. the derivative of (6) with respect to i_t), and similarly for x_t. These partial derivatives are:

$$\frac{\partial \mathcal{L}}{\partial \pi_t} = \pi_t$$
$$\frac{\partial \mathcal{L}}{\partial x_t} = \omega x_t$$
$$\frac{\partial \pi_t}{\partial i_t} = -\frac{\gamma}{\sigma}$$
$$\frac{\partial x_t}{\partial i_t} = -\frac{1}{\sigma}$$

Plugging these in, we have:

$$\frac{\partial \mathcal{L}}{\partial i_t} = -\frac{\gamma}{\sigma}\pi_t - \frac{\omega}{\sigma}x_t$$

Setting this equal to zero and simplifying yields:

$$x_t = -\frac{\gamma}{\omega}\pi_t \tag{8}$$

(8) is the *same* as the first order condition under discretion – a "lean against the wind" condition. One would be tempted to conclude that there is no difference between commitment and discretion. But one would be wrong. Under commitment, the central bank doesn't just choose i_t in period t, it chooses an expected *path* of future policy rates, too. The choice of i_{t+1}, for example, influences $\mathbb{E}_t \pi_{t+1}$ and $\mathbb{E}_t x_{t+1}$, which are both relevant for the loss function in period t. We need to take this into account when choosing a path of *future* interest rates.

This is going to make the problem more cumbersome. Consider the FOC with respect to i_{t+1}. When choosing i_{t+1}, we need to take into account how it affects the loss function in period t as well as how it impacts the loss function in $t+1$. Formally, the derivative of the loss function (with respect to i_{t+1}, from the perspective of period t) is:

$$\frac{\partial \mathcal{L}}{\partial i_{t+1}} = \frac{\partial \mathcal{L}}{\partial \pi_t}\left[\frac{\partial \pi_t}{\partial \pi_{t+1}}\frac{\partial \pi_{t+1}}{\partial i_{t+1}} + \frac{\partial \pi_t}{\partial x_{t+1}}\frac{\partial x_{t+1}}{\partial i_{t+1}}\right] + \frac{\partial \mathcal{L}}{\partial x_t}\left[\frac{\partial x_t}{\partial \pi_{t+1}}\frac{\partial \pi_{t+1}}{\partial i_{t+1}} + \frac{\partial x_t}{\partial x_{t+1}}\frac{\partial x_{t+1}}{\partial i_{t+1}}\right] + \frac{\partial \mathcal{L}}{\partial \pi_{t+1}}\frac{\partial \pi_{t+1}}{\partial i_{t+1}} + \frac{\partial \mathcal{L}}{\partial x_{t+1}}\frac{\partial x_{t+1}}{\partial i_{t+1}}$$

This looks kind of nasty. And it is. There is the chain rule all over the place. As above, $\frac{\partial \mathcal{L}}{\partial \pi_t}$ is the derivative of the loss function with respect to current inflation. The term $\left[\frac{\partial \pi_t}{\partial \pi_{t+1}}\frac{\partial \pi_{t+1}}{\partial i_{t+1}} + \frac{\partial \pi_t}{\partial x_{t+1}}\frac{\partial x_{t+1}}{\partial i_{t+1}}\right]$ measures how next period's interest rate impacts current inflation. It doesn't do this directly (i_{t+1} doesn't directly show up (6)). It does through its impact on $\mathbb{E}_t \pi_{t+1}$ and $\mathbb{E}_t x_{t+1}$. That's what those terms in the brackets are picking up. Same story for the next overall term in the derivative – this is picking up the effect on the loss function of i_{t+1} coming through x_t (i.e. (7)), which again comes through future inflation and the future output gap. In addition to the impacts on the current loss function, the choice of i_{t+1} impacts the future loss function, as well. This is measured by the terms

$\frac{\partial \mathcal{L}}{\partial \pi_{t+1}} \frac{\partial \pi_{t+1}}{\partial i_{t+1}}$ and $\frac{\partial \mathcal{L}}{\partial x_{t+1}} \frac{\partial x_{t+1}}{\partial i_{t+1}}$.

Let's collect all these derivatives together in one place.

$$\frac{\partial \mathcal{L}}{\partial \pi_{t+j}} = \beta^j \pi_{t+j}$$
$$\frac{\partial \mathcal{L}}{\partial x_{t+j}} = \omega \beta^j x_{t+j}$$
$$\frac{\partial \pi_t}{\partial \pi_{t+1}} = \frac{\gamma}{\sigma} + \beta$$
$$\frac{\partial \pi_t}{\partial x_{t+1}} = \gamma$$
$$\frac{\partial x_t}{\partial \pi_{t+1}} = \frac{1}{\sigma}$$
$$\frac{\partial x_t}{\partial x_{t+1}} = 1$$
$$\frac{\partial \pi_t}{\partial i_t} = -\frac{\gamma}{\sigma}$$
$$\frac{\partial x_t}{\partial i_t} = -\frac{1}{\sigma}$$

Now plug these in to the general expression for the derivative above:

$$\frac{\partial \mathcal{L}}{\partial i_{t+1}} = \pi_t \left[-\left(\frac{\gamma}{\sigma}\right)^2 - \frac{\gamma}{\sigma}\beta - \frac{\gamma}{\sigma} \right] + \omega x_t \left[-\frac{\gamma}{\sigma^2} - \frac{1}{\sigma} \right] - \beta \pi_{t+1} \frac{\gamma}{\sigma} - \omega \beta x_{t+1} \frac{1}{\sigma}$$

This can be simplified a bit:

$$\frac{\partial \mathcal{L}}{\partial i_{t+1}} = -\frac{\gamma}{\sigma} \left[\frac{\gamma}{\sigma} + 1 + \beta \right] \pi_t - \frac{\omega}{\sigma} \left[\frac{\gamma}{\sigma} + 1 \right] x_t - \frac{\beta\gamma}{\sigma} \pi_{t+1} - \frac{\omega\beta}{\sigma} x_{t+1} \tag{9}$$

Note that we know something about the relationship from x_t and π_t. In particular, the FOC for i_t is $x_t = -\frac{\gamma}{\omega}\pi_t$. If we plug this in above, we have

$$\frac{\partial \mathcal{L}}{\partial i_{t+1}} = -\frac{\gamma}{\sigma} \left[\frac{\gamma}{\sigma} + 1 + \beta \right] \pi_t + \frac{\gamma}{\sigma} \left[\frac{\gamma}{\sigma} + 1 \right] \pi_t - \frac{\beta\gamma}{\sigma} \pi_{t+1} - \frac{\omega\beta}{\sigma} x_{t+1}$$

But this is now:

$$\frac{\partial \mathcal{L}}{\partial i_{t+1}} = -\frac{\gamma\beta}{\sigma} \pi_t - \frac{\gamma\beta}{\sigma} \pi_{t+1} - \omega\beta x_{t+1}$$

Setting this equal to zero and simplifying yields:

$$x_{t+1} = -\frac{\gamma}{\omega} \left(\pi_t + \pi_{t+1} \right)$$

In other words, the FOC fo i_{t+1} results in a condition where the central bank will plan to implement policy so that the output gap in $t+1$ is negatively related to the *sum* of inflation in period t and $t+1$. What is the sum of inflation (note that these are net inflation rates, expressed in deviation form)? Well:

$$\pi_t = \ln P_t - \ln P_{t-1}$$
$$\pi_{t+1} = \ln P_{t+1} - \ln P_t$$

Therefore, we have:

$$\pi_t + \pi_{t+1} = \ln P_{t+1} - \ln P_{t-1}$$

This means that the FOC for i_{t+1} can be written:

$$x_{t+1} = -\frac{\gamma}{\omega}\left(\ln P_{t+1} - \ln P_{t-1}\right) \tag{10}$$

In other words, optimal policy for commitment means setting the *future* output gap as a "lean against the wind condition" as a function of the *price level* (relative to the period before policy was set), rather than inflation rate.

Now, to think about this moving forward, things would extremely messy – we would have all sorts of derivatives floating around. This is because, for example, i_{t+2} is going to be relevant for the loss function in t. How? Through its impact on π_{t+2} and x_{t+2}, which are relevant for π_{t+1} and x_{t+1}, which are in turn relevant for π_t and x_t. Proceeding via the "plug and chug" method (i.e. solving for π_t and x_t in terms of i_t and future variables, and doing an unconstrained problem) quickly becomes infeasible.

An easier, though to me less intuitive, way to do this is to set the problem up as a Lagrangian (instead of plugging in the constraints). We treat the IS and Phillips Curves as constraints with Lagrange multipliers. The Lagrangian would be:

$$\mathbb{L} = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \left[\frac{\pi_{t+j}^2}{2} + \omega \frac{x_{t+j}^2}{2} + \psi_{1,t+j}\left(\gamma x_{t+j} + \beta \pi_{t+j+1} - \pi_{t+j}\right) + \dots \right.$$
$$\left. \psi_{2,t+j}\left(x_{t+j+1} - \frac{1}{\sigma}\left(i_{t+j} - \pi_{t+j+1} - r_{t+j}^f \right) - x_{t+j} \right) \right]$$

The policymaker gets to choose i_{t+j}, π_{t+j}, and x_{t+j}. For $j = 0$, the FOC are:

$$\frac{\partial \mathbb{L}}{\partial \pi_t} = \pi_t - \psi_{1,t} = 0$$
$$\frac{\partial \mathbb{L}}{\partial x_t} = \omega x_t + \gamma \psi_{1,t} - \psi_{2,t} = 0$$
$$\frac{\partial \mathbb{L}}{\partial i_t} = -\frac{1}{\sigma}\psi_{2,t} = 0$$

Setting these equal to zero means:

$$\psi_{1,t} = \pi_t$$

$$\omega x_t + \gamma \psi_{1,t} = \psi_{2,t}$$

$$\psi_{2,t} = 0$$

Combining these together:

$$x_t = -\frac{\gamma}{\omega}\pi_t$$

This the same condition we had before for period t. Now consider the derivatives with respect to inflation, the gap, and the interest rate led forward one period. We have:

$$\frac{\partial \mathbb{L}}{\partial \pi_{t+1}} = \beta\psi_{1,t} + \frac{1}{\sigma}\psi_{2,t} + \beta\pi_{t+1} - \beta\psi_{1,t+1} = 0$$

$$\frac{\partial \mathbb{L}}{\partial x_{t+1}} = -\psi_{2,t} + \beta\omega x_{t+1} + \gamma\beta\psi_{1,t+1} - \beta\psi_{2,t+1} = 0$$

$$\frac{\partial \mathbb{L}}{\partial i_{t+1}} = -\beta\frac{1}{\sigma}\psi_{2,t+1} = 0$$

Now, for sake of completeness, let us stop for a moment and think about why these derivatives are what they are. Focus on the FOC with respect to π_{t+1}. π_{t+1} shows up in the period t constraint on the Phillips Curve (the term $\beta\psi_{1,t}$), the period t constraint on the IS equation (the term $\frac{1}{\sigma}\psi_{2,t}$), the period $t+1$ loss function (the term $\beta\pi_{t+1}$, where the β shows up because the $t+1$ loss is discounted relative to period t), and the period $t+1$ constraint on the Phillips Curve (the term $-\beta\psi_{1,t+1}$, where the discounting is because the constraint is led forward one period). Similarly, for the derivative with respect to x_{t+1}: it shows up in the period t IS constraint ($\psi_{2,t}$), the $t+1$ objective function (the term $\beta\omega x_{t+1}$), the period $t+1$ constraint on the Euler equation (the term $\beta\gamma\psi_{1,t+1}$), and the $t+1$ constraint on the IS equation (the term $-\beta\psi_{2,t+1}$). The future interest rate, i_{t+1}, only appears in the $t+1$ IS equation, and so the derivative of the Lagrangian is $-\beta\frac{1}{\sigma}\psi_{2,t+1}$. We therefore must have $\psi_{2,t+1} = 0$. But using this, and setting the other two derivatives to zero, we have:

$$\pi_{t+1} = \psi_{1,t+1} - \psi_{1,t}$$

$$\psi_{1,t+1} = -\frac{\omega}{\gamma}x_{t+1}$$

Combining these, we have:

$$\pi_{1,t+1} = -\frac{\omega}{\gamma}x_{t+1} - \psi_{1,t}$$

But we know from the period t FOC that $\psi_{1,t} = \pi_t$. Hence, we have:

$$\pi_{t+1} = -\frac{\omega}{\gamma} x_{t+1} - \pi_t$$

Or:

$$x_{t+1} = -\frac{\gamma}{\omega} (\pi_t + \pi_{t+1})$$

This is exactly what we had above using the "plug and chug" method!

Now, let's take the derivatives with respect to period $t+2$ stuff:

$$\frac{\partial \mathbb{L}}{\partial \pi_{t+2}} = \beta^2 \psi_{1,t+1} + \frac{1}{\sigma} \beta \psi_{2,t+1} + \beta^2 \pi_{t+2} - \beta^2 \psi_{1,t+2} = 0$$

$$\frac{\partial \mathbb{L}}{\partial x_{t+2}} = -\beta \psi_{2,t+1} + \beta^2 \omega x_{t+2} + \gamma \beta^2 \psi_{1,t+2} - \beta^2 \psi_{2,t+2} = 0$$

$$\frac{\partial \mathbb{L}}{\partial i_{t+2}} = -\beta^2 \frac{1}{\sigma} \psi_{2,t+2} = 0$$

We again have $\psi_{2,t+2} = 0$. This means that:

$$\psi_{1,t+2} = -\frac{\omega}{\gamma} x_{t+2}$$

And:

$$\pi_{t+2} = \psi_{1,t+2} - \psi_{1,t+1}$$

But this is:

$$\pi_{t+2} = -\frac{\omega}{\gamma} x_{t+2} - \psi_{1,t+1}$$

But from above, we know that $\psi_{1,t+1} = \pi_{t+1} + \psi_{1,t}$. So we have:

$$\pi_{t+2} = -\frac{\omega}{\gamma} x_{t+2} - \pi_{t+1} - \psi_{1,t}$$

But from the period t FOC, we know that $\psi_{1,t} = \pi_t$. Hence, we have:

$$x_{t+2} = -\frac{\gamma}{\omega} (\pi_t + \pi_{t+1} + \pi_{t+2})$$

But, just as above, this can be written:

$$x_{t+2} = -\frac{\gamma}{\omega} (\ln P_{t+2} - \ln P_{t-1})$$

If one keeps going, for any $j \geq 1$, the first order conditions boil down to:

$$x_{t+j} = -\frac{\gamma}{\omega} \sum_{s=0}^{j} \pi_{t+s} = -\frac{\gamma}{\omega} (\ln P_{t+j} - \ln P_{t-1}) \tag{11}$$

In other words, under commitment the central bank commits to an implicit *price level target* wherein the output gap is negatively related to the price level gap (relative to some initial condition). This differs from the solution under discretion, where each period the output gap is negatively related to the inflation rate in that period.

3.1 Time Inconsistency

The solution under commitment is plagued by the problem of *time inconsistency.* The basic idea of time inconsistency is that a policymaker wants to plan to do something in the future, but when the future comes around, if the policymaker can re-optimize, she will want to deviate from the plan.

A simple real-world example of time inconsistency involves diet. In period t, one might wish to plan to eat poorly in the present, but go on a diet in the next period $(t+1)$. But when the future rolls around, one might want to eat poorly again for one more period $(t+1)$, deferring the diet until $t+2$. The problem would recur again in $t+2$.

The time inconsistency problem in the basic New Keynesian model is as follows. In period t, the planner wants to set:

$$x_t = -\frac{\gamma}{\omega}\pi_t$$

From the perspective of period t, the policymaker wants to plan to implement policy in $t+1$ as follows:

$$x_{t+1} = -\frac{\gamma}{\omega}\left(\pi_t + \pi_{t+1}\right)$$

In other words, policy in $t+1$ will be *backward-looking* in the sense that it will depend on what inflation was in the previous period (t). But here's the problem. If the planner can re-optimize in period $t+1$, she will not choose the above solution! Rather, she would implement policy such that:

$$x_{t+1} = -\frac{\gamma}{\omega}\pi_{t+1}$$

Unless the policymaker can commit to not re-optimizing in the future, it will not be possible to implement the optimal solution under commitment. "Bygones will be bygones" – the policymaker won't care what inflation was in period t when setting policy in $t+1$ if she can re-optimize. But this inability to commit might make the policymaker worse off from the perspective of period t, as we shall see below.

4 Commitment vs. Discretion and the Divine Coincidence

The optimal targeting rule under discretion is a "lean against the wind" condition in terms of the output gap and the inflation rate:

$$x_t = -\frac{\gamma}{\omega}\pi_t \tag{12}$$

The optimal targeting rule under commitment looks similar and has a similar feel, but it is an implicit price level target: the policymaker is targeting a relationship between the output gap (in any period) and the price level (relative to some benchmark, which we will normalize to zero in the log, or unity in the level).[2]

$$x_t = -\frac{\gamma}{\omega} \ln P_t \tag{13}$$

How do these targeting rules differ, and what are the gains from commitment over discretion?[3]

4.1 Divine Coincidence

In the basic New Keynesian model, it turns out that there are no gains from commitment over discretion. This is not a general result, but is specific to the model as written down. This boils down to a term deemed the "Divine Coincidence" (Blanchard and Gali, 2007), which points out that there actually is no tradeoff between stabilizing inflation and the output gap in the basic New Keynesian model. We saw this above when discussing discretion. Suppose that, each period, the central bank is following discretion, which means $x_{t+j} = -\frac{\gamma}{\omega}\pi_{t+j}$ for any $j \geq 0$. Plugging this condition into the period t Phillips Curve, we have:

$$\pi_t = -\frac{\gamma^2}{\omega}\pi_t + \beta \mathbb{E}_t \pi_{t+1}$$

Therefore:

$$\mathbb{E}_t \pi_{t+1} = \frac{1 + \frac{\gamma^2}{\omega}}{\beta}\pi_t$$

If we solve this difference equation forward, we get:

$$\mathbb{E}_t \pi_{t+j} = \left(\frac{1 + \frac{\gamma^2}{\omega}}{\beta} \right)^j \pi_t$$

Since $\frac{1+\frac{\gamma^2}{\omega}}{\beta} > 1$, in the limit as j gets big, $\mathbb{E}_t \pi_{t+j}$ explodes *unless* $\pi_t = 0$. So we must have $\pi_t = 0$. But, from the FOC, this means $x_t = 0$. Since this will be true each period, we will have $\pi_{t+j} = x_{t+j} = 0$ as the unique, non-explosive to the policy problem. This means that $\mathcal{L} = 0$, which is the lowest value it can take (given that it is quadratic). Put another way: even under discretion, there is no tradeoff between inflation and the output gap. A policymaker can set both equal to zero.

[2] A subtle issue at play here is that the New Keynesian model (without money) does not determine the price level, but rather the inflation rate. For a more in-depth discussion, see Cochrane (2011, *Journal of Political Economy*), "Determinacy and Identification with Taylor Rules." Hence, I need to impose a normalization on the $t-1$ price level, so, for simplicity, I'm going to assume $P_{t-1} = 1$ (so $\ln P_{t-1} = 0$. Since the model determine π_t, given P_{t-1}, I can determine P_t. But without some nominal anchor like money formally modeled, I cannot determine P_{t-1}.

[3] Note that I am referring to a *targeting rule* as a policy rule that shows some relationship between targets, like π_t and x_t, without explicit reference to the instrument (i_t in this case). In contrast, an *instrument rule* is a rule that expresses the instrument as a function of targets – e.g. a Taylor-type rule.

Since discretion achieves $\mathcal{L} = 0$, there can be no gain from commitment. Under commitment, we will get exactly the same solution: $\pi_{t+j} = x_{t+j} = 0$, which will imply (see above) that $i_{t+j} = r^f_{t+j}$ each period.

4.2 Determinacy, Instrument Rules, and Stochastic Intercepts

In the basic NK model, there is no distinction between optimal monetary policy under discretion or commitment. The Divine coincidence holds, and it is possible to completely stabilize both inflation and the output gap at all times, achieving the global minimum of the loss function.

As shown above, the optimal targeting rule (either $x_t = -\frac{\gamma}{\omega}\pi_t$ under discretion, or $x_t = -\frac{\gamma}{\omega}\ln P_t$ under commitment) requires that the nominal interest rate (the policy instrument) track the natural rate of interest one-for-one. Formally:

$$i_t = r^f_t \tag{14}$$

One would be tempted to conclude that an instrument rule like (14) would be equivalent to the optimal targeting rules derived above. *This conclusion is wrong.* (14) is an *equilibrium outcome* for the policy instrument under a targeting rule. If one tries to solve the model with (14) as the policy rule, there will exist no determinate equilibrium.

As we saw in our notes on determinacy with interest rate rules, for determinacy (i.e. uniqueness of the equilibrium), the nominal interest rate needs to react sufficiently strongly to *endogenous* variables (like inflation). Taking the equilibrium outcome in (14) as an instrument rule would be a problem, since r^f_t is exogenous.

To see this concretely, suppose that the central bank adopted (14) as the policy rule. Plug this into the IS equation:

$$x_t = \mathbb{E}_t x_{t+1} + \frac{1}{\sigma}\mathbb{E}_t \pi_{t+1}$$

Suppose further that prices are flexible, so $\phi = 0$ (and therefore $\gamma \to \infty$). This means $x_t = \mathbb{E}_t x_{t+1} = 0$. Then, from above, we see that $\mathbb{E}_t \pi_{t+1} = 0$. *Expected* inflation is pinned down, but what about current inflation? There is no unique solution. From the Phillips Curve, we would have $\pi_t = \gamma x_t$. But ∞ times 0 could be anything. So, an exogenous interest rate rule results in indeterminacy. Note that this indeterminacy would result even if $\gamma < \infty$, which would mean that not only would π_t be indeterminate, x_t would be as well.

It is possible to follow an instrument rule that reacts to the natural rate of interest. In particular, suppose the central bank follows an instrument rule with a "stochastic intercept":

$$i_t = r^f_t + \phi_\pi \pi_t \tag{15}$$

Plug this into the IS equation. We get:

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left(\phi_\pi \pi_t - \mathbb{E}_t \pi_{t+1}\right)$$

Again assume that $\phi = 0$, which means that $x_t = \mathbb{E}_t x_{t+1} = 0$. Then, for the IS equation to hold, we must have:

$$\mathbb{E}_t \pi_{t+1} = \phi_\pi \pi_t$$

Solving this forward, we have:

$$\mathbb{E}_t \pi_{t+T} = \phi_\pi^T \pi_t$$

If $\phi_\pi > 1$, then the only non-explosive solution is $\pi_t = 0$. But if $\phi_\pi < 1$, any value of π_t is consistent with inflation not exploding. This will also be the case even when $\gamma < \infty$ (so that prices are sticky and the Phillips Curve is non-vertical).

Now, let's drop the assumption that prices are flexible (so $\gamma < \infty$). As long as $\phi_\pi > 1$, then a rule like (15) will result in an equilibrium in which $\pi_t = x_t = 0$. To see this, guess that the solution is $x_t = \mathbb{E}_t x_{t+1} = 0$. Then the IS equation again gives us $\mathbb{E}_t \pi_{t+1} = \phi_\pi \pi_t$. The only non-explosive solution possible is $\pi_t = 0$, which is also consistent with the Phillips Curve holding.

There is something a bit odd going on here. A rule like (15) with $\phi > 1$ produces an equilibrium in which $i_t = r_t^f$ at all times. But an instrument rule that pre-specifies $i_t = r_t^f$ will produce a non-unique equilibrium. The key point here is that $i_t = r_t^f$ has to emerge as *an equilibrium outcome*, not a pre-specified instrument rule.

Cochrane (2011) makes a big point about this issue. A rule like (15), in his terminology, is essentially a "threat" to "blow up the world." Forward-looking models need to have the right number of explosive roots for uniqueness. We pin down unique equilibrium via terminal conditions. If $\phi_\pi < 1$, then the terminal condition (that inflation not explode) will hold for any π_t (and, consequently, any x_t). We need $\phi_\pi > 1$ so that the terminal condition holding gives us unique initial conditions. In Cochrane's words, $\phi_\pi > 0$ in (15) is a threat to let inflation go to $+/-\infty$ *unless* $\pi_t = 0$.

5 Cost-Push Shocks and a Monetary Policy Tradeoff

To make monetary policy "difficult," and for there to be some potential gain from commitment over discretion, we need something that can break the Divine Coincidence. The simplest way to do this is to introduce a "cost-push" shock into the Phillips Curve:

$$\pi_t = \gamma x_t + \beta \mathbb{E}_t \pi_{t+1} + u_t$$

If $u_t \neq 0$, then it is not going to be possible for $\pi_t = x_t = 0$ at all times.

The FOC under either commitment or discretion will be the same as above. Let's focus on discretion first. Plug in the FOC into the Phillips Curve to eliminate x_t:

$$\pi_t = -\frac{\gamma^2}{\omega} \pi_t + \beta \mathbb{E}_t \pi_{t+1} + u_t$$

Or:

$$\mathbb{E}_t \pi_{t+1} = \frac{1 + \frac{\gamma^2}{\omega}}{\beta} \pi_t - \frac{1}{\beta} u_t$$

This may be written a little more compactly as:

$$\mathbb{E}_t \pi_{t+1} = \frac{\omega + \gamma^2}{\omega \beta} \pi_t - \frac{1}{\beta} u_t$$

This is a forward-looking difference equation. The coefficient on π_t is > 1, which means it is going to be explosive (just like before). But now there is a term involving u_t. Let's assume that u_t follows an exogenous AR(1) process:

$$u_t = \rho_u u_{t-1} + s_u \varepsilon_{u,t}$$

This means that $\mathbb{E}_t u_{t+j} = \rho_u^j u_t$. As long as $0 < \rho_u < 1$, which we shall assume, in expectation u_t eventually returns to zero. Let's solve the difference equation forward. Going forward to $t + 2$, we have:

$$\mathbb{E}_t \pi_{t+2} = \frac{\omega + \gamma^2}{\omega \beta} \mathbb{E}_t \pi_{t+1} - \frac{1}{\beta} \mathbb{E}_t u_{t+1}$$

Define the auxiliary parameter:

$$a = \frac{\omega + \gamma^2}{\omega \beta}$$

So:

$$\mathbb{E}_t \pi_{t+2} = a \mathbb{E}_t \pi_{t+1} - \frac{1}{\beta} \mathbb{E}_t u_{t+1}$$

Which is:

$$\mathbb{E}_t \pi_{t+2} = a^2 \pi_t - a \frac{1}{\beta} u_t - \frac{1}{\beta} \rho_u u_t$$

Now, similarly, going forward another period, we have:

$$\mathbb{E}_t \pi_{t+3} = a \mathbb{E}_t \pi_{t+2} - \frac{1}{\beta} \mathbb{E}_t u_{t+2}$$

Which can be written:

$$\mathbb{E}_t \pi_{t+3} = a^3 \pi_t - \frac{a^2}{\beta} u_t - \frac{a}{\beta} \rho_u u_t - \frac{1}{\beta} \rho_u^2 u_t$$

Which may be written:

$$\mathbb{E}_t \pi_{t+3} = a^3 \pi_t - \frac{1}{\beta} \left(\rho_u^2 + a \rho_u + a^2 \right) u_t$$

Go forward another period. WE have:

$$\mathbb{E}_t\pi_{t+4} = a\mathbb{E}_t\pi_{t+3} - \frac{1}{\beta}\rho_u^3 u_t$$

But this is:

$$\mathbb{E}_t\pi_{t+4} = a^4\pi_t - \frac{1}{\beta}\left(a\rho_u^2 + a^2\rho_u + a^3\right)u_t - \frac{1}{\beta}\rho_u^3 u_t$$

But this can be written:

$$\mathbb{E}_t\pi_{t+4} = a^4\pi_t - \frac{1}{\beta}\left(\rho_u^3 + a\rho_u^2 + a^2\rho_u + a^3\right)u_t$$

You can see a pattern. Going forward some arbitrary number of periods, T, we have:

$$\mathbb{E}_t\pi_{t+T} = a^T\pi_t - \frac{u_t}{\beta}S$$

The term S is shorthand for:

$$S = \rho_u^{T-1} + a\rho_u^{T-2} + a^2\rho_u^{T-3} + \ldots a^{T-1}$$

Divide both sides by a^{T-1}, so $\widehat{S} = \frac{S}{a^{T-1}}$. We have:

$$\frac{S}{a^{T-1}} = \widehat{S} = \left(\frac{\rho_u}{a}\right)^{T-1} + \left(\frac{\rho_u}{a}\right)^{T-2} + \left(\frac{\rho_u}{a}\right)^{T-3} + \ldots$$

Now, to ease notation a bit more, define $b = \frac{\rho_u}{a}$, which is less than one (given that $\rho_u < 1$ and $a > 1$). We have:

$$\widehat{S} = 1 + b + b^2 + \cdots + b^{T-1}$$

So:

$$b\widehat{S} = b + b^2 + \cdots + b^T$$

Therefore:

$$\widehat{S} - b\widehat{S} = 1 - b^T$$

So:

$$\widehat{S} = \frac{1-b^T}{1-b}$$

Now, you might ask: where is this getting us? Well, we can write:

$$\mathbb{E}_t\pi_{t+T} = a^T\pi_t - u_t\frac{a^{T-1}}{\beta}\widehat{S}$$

Which is:

$$\mathbb{E}_t\pi_{t+T} = a^T\pi_t - u_t\frac{a^{T-1}}{\beta}\frac{1-b^T}{1-b}$$

To have $\mathbb{E}_t\pi_{t+T} \to 0$ as $T \to \infty$, we must therefore have the right hand side equal to zero. But this gives π_t as a function of u_t!:

$$a\pi_t = u_t\frac{1}{\beta}\frac{1-b^T}{1-b}$$

But as $T \to \infty$, this becomes:

$$\pi_t = (\beta a(1-b))^{-1}\, u_t$$

Now, let's write this back in terms of the underlying parameters. Plugging in for b first:

$$(\beta a(1-b))^{-1} = \left(\beta a\left(1 - \frac{\rho_u}{a}\right)\right)^{-1} = (\beta(a-\rho_u))^{-1}$$

Now plugging in for a:

$$(\beta(a-\rho_u))^{-1} = \left(\beta\left(\frac{\omega+\gamma^2}{\omega\beta} - \rho_u\right)\right)^{-1} = \left(\frac{\omega+\gamma^2-\rho_u\omega\beta}{\omega}\right)^{-1} = \frac{\omega}{\omega(1-\rho_u\beta)+\gamma^2}$$

In other words, under discretion, the solution for π_t is:

$$\pi_t = \frac{\omega}{\omega(1-\rho_u\beta)+\gamma^2}u_t \tag{16}$$

And, therefore, the solution for x_t is:

$$x_t = -\frac{\gamma}{\omega(1-\rho_u\beta)+\gamma^2}u_t \tag{17}$$

(16) and (17) are actually pretty intuitive expressions. Suppose that $\omega \to \infty$. This means that the policymaker really wants to stabilize the output gap. Then, the coefficient for x_t will go to 0, so that $x_t = 0$. Alternatively, suppose that $\gamma \to \infty$ (so that prices are flexible). Then both coefficients go to zero, so neither inflation nor the output gap react to the cost-push shock.

Now let's think about what things look like under commitment. We could proceed as we did under discretion – take the Phillips Curve, and plug in the first order condition in period t, solve forward, impose a terminal condition, and then get an analytic expression for endogenous variables. But, because the optimal solution under commitment is backward-looking, things get very messy very quickly.

For this reason, it is much simpler to just solve the model in a program like Dynare with $x_t = -\frac{\gamma}{\omega}\ln P_t$, where I am normalizing the initial price level to unity (zero in the log). Below, I show impulse responses to a cost-push shock under both discretion (solid lines) and commitment for different values of ω, keeping everything else the same.

Figure 1: Cost-Push Shock: Commitment vs. Discretion ($\omega = 0.01$)

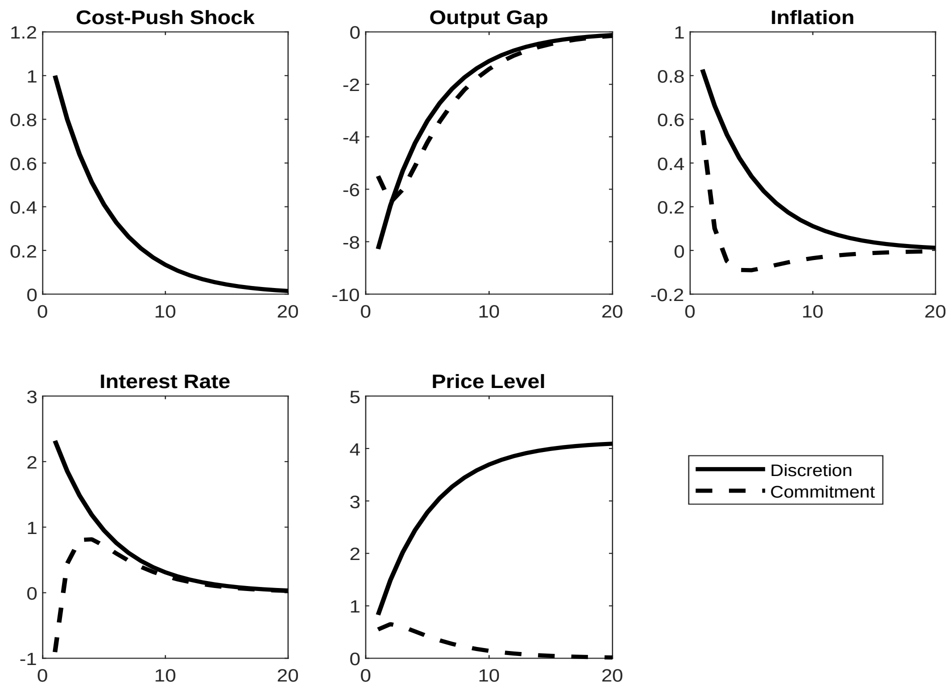

Figure 2: Cost-Push Shock: Commitment vs. Discretion ($\omega = 0.5$)

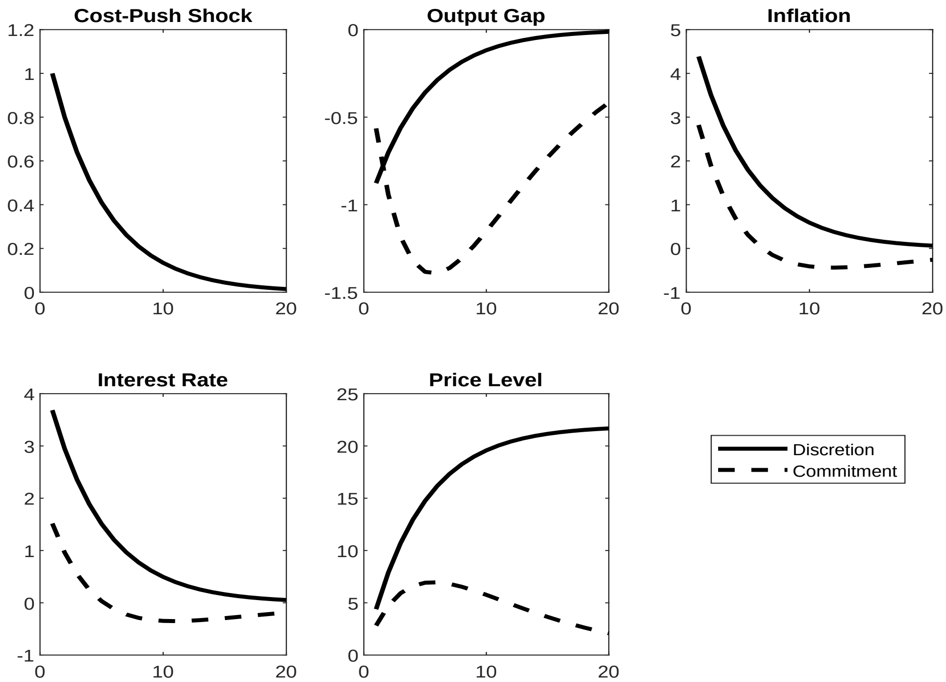

Figure 3: Cost-Push Shock: Commitment vs. Discretion ($\omega = 1$)

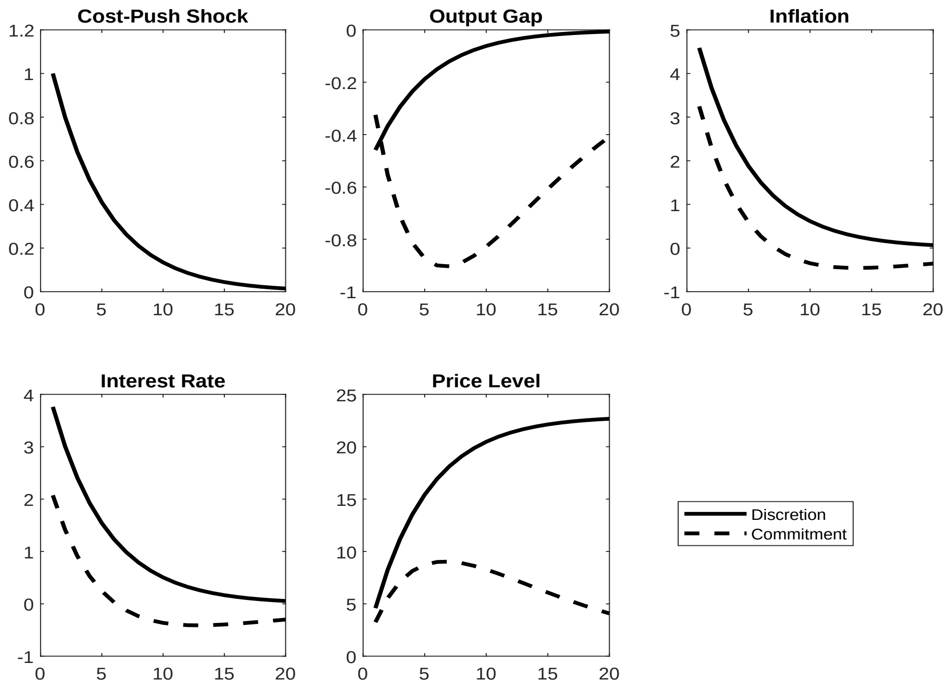

The most noticeable thing in these IRFs should be obvious from looking at the lean against the wind condition. Under commitment, the price level will be stationary, whereas under discretion it will not. We see these in all of the graphs – the price level eventually always returns to its pre-shock value under commitment, but not under discretion. The exact paths of inflation, the output, and the interest rate differ depending on the value of ω, but the price level always returns to trend.

The absolute units are not very meaningful and are subject to scale effects, but I can calculate the welfare loss under discretion and commitment when there are a cost-push shocks. I set the variance of the shock equal to one. When $\omega = 0.01$, the welfare loss under discretion is 190.35, whereas it is 110.18 under commitment. When $\omega = 0.5$, the losses are 2725 under discretion and 1300 under commitment. When $\omega = 1$, the losses are 2952 under discretion and 1659 under commitment. When cost-push shocks are present, the policymaker clearly does better (regardless of the weight, ω, that it chooses) under commitment.