Graduate Macro Theory II:
A High-Level Overview of the New Keynesian Model

Eric Sims
University of Notre Dame

Spring 2024

The simplest New Keynesian model boils down to three equations: a supply relationship, a demand relationship, and a monetary policy rule.

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

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

$$i_t = \phi_\pi \pi_t + \phi_x x_t + u_t \tag{3}$$

The variables are all log-linear deviations about steady state, and the model is linearized about a zero inflation steady state (so $\Pi = 1$, where Π is steady state gross inflation). In (1), the coefficient $\gamma = \frac{(1-\phi)(1-\phi\beta)}{\phi}(\sigma + \chi)$ in the most basic model. $\phi \in [0,1]$ is the Calvo parameter – this is both the fraction of intermediate firms who *cannot* adjust their prices in a given period and the probability that a price chosen today will still be in effect tomorrow. β is the subjective discount factor. σ is the inverse elasticity of substitution. χ is the inverse Frisch elasticity. r_t^f is the natural rate of interest and can be thought of as exogenous. (3) is a policy rule for the nominal rate in terms of targets and an exogenous shock, u_t. I'm thinking of u_t as potentially following an AR(1) process to get some persistence; it is possible to write the interest rate rule with smoothing and an iid shock. We could consider alternative policy rules:

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

$$x_t = -\frac{\gamma}{\omega}\ln P_t \tag{5}$$

$$\pi_t = 0 \tag{6}$$

$$x_t = 0 \tag{7}$$

These are all targeting rules (not instrument rules like a Taylor rule). The first is the optimal targeting rule under discretion; the second is the optimal targeting rule under commitment (where I have normalized the log price level prior to the beginning of time to 0, so unity in the level, and we would need an additional equation defining $\pi_t = \ln P_t - \ln P_{t-1}$). The third and fourth are strict inflation and output targets, respectively. It is also possible to close the model with a rule for the money supply, but we would then need to specify a money demand curve (which we could do with money in the utility function in a way that wouldn't otherwise impact the properties of the model).

1 The NK Model has a RBC Core

The core of a NK model is a real business cycle model (without capital). A RBC model without capital (but endogenous labor supply) can actually be thought of as an endowment economy. Here's how.

The household problem is:

$$\max_{C_t, N_t, B_t} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma}}{1-\sigma} - \theta \frac{N_t^{1+\chi}}{1+\chi} \right)$$

s.t.

$$P_t C_t + B_t \leq W_t N_t + P_t D_t + (1 + i_{t-1}) B_{t-1}$$

Here P_t is the price level, and B_t are nominal bonds. W_t is the nominal wage, i_t is the nominal interest rate, and D_t are real dividends from ownership in firms. The FOC are:

$$1 = \mathbb{E}_t \Lambda_{t,t+1} (1 + i_t) \Pi_{t+1}^{-1} \tag{8}$$

$$\Lambda_{t-1,t} = \beta \left(\frac{C_t}{C_{t-1}} \right)^{-\sigma} \tag{9}$$

$$\theta N_t^{\chi} = C_t^{-\sigma} w_t \tag{10}$$

The Fisher relationship defines the real interest rate as $1 + r_t = (1 + i_t) \mathbb{E}_t \Pi_{t+1}^{-1}$; or, approximately, $r_t \approx i_t - \mathbb{E}_t \pi_{t+1}$. Other than writing the Euler equation in nominal terms, these are the same FOC as in a RBC model (without capital).

In the RBC core, we can think about a representative firm that produces output according to:

$$Y_t^e = A_t N_t \tag{11}$$

I'm using a superscript e to denote that this is the efficient allocation. The profit maximization condition is to hire labor up until the point at which the wage equals the marginal product of labor:

$$w_t = A_t \tag{12}$$

The firm earns no profit, so $D_t = 0$ and $w_t N_t = Y_t$. This implies the resource constraint:

$$Y_t^e = C_t \tag{13}$$

Because there is no capital (i.e. no endogenous state variables), it is easy to solve for the equilibrium level of output. Combine (13), (12), (11) and (10). We have:

$$\theta\left(\frac{Y_t^e}{A_t}\right)^{\chi} = \left(Y_t^e\right)^{-\sigma} A_t$$

This implies:

$$Y_t^e = \theta^{-\frac{1}{\sigma+\chi}} A_t^{\frac{1+\chi}{\sigma+\chi}} \tag{14}$$

(14) is the efficient level of output – the level of output we would get with (i) no monopoly distortions and (ii) no price rigidity. Since A_t is exogenous, we can think about Y_t^e as effectively endogenous, and imposing $C_t = Y_t$ with the Euler equation would determine the equilibrium r_t^e:

$$1 = \mathbb{E}_t \beta \left(\frac{Y_{t+1}^e}{Y_t^e}\right)^{-\sigma} (1 + r_t^e) \tag{15}$$

We could, in addition, determine the inflation rate by adding a policy rule that satisfies the Taylor principle, i.e. $i_t = \phi_\pi \pi_t$, with $\phi_\pi > 1$.

1.1 Monopolistic Competition

To the RBC core, the basic NK model adds (i) monopolistic competition and (ii) price stickiness. (i) is an input into getting (ii): we need firms to have price-setting power to think about price rigidity. On its own, monopolistic competition doesn't do much – it just distorts the competitive equilibrium relative to the efficient equilibrium by a fixed amount.

I'm going to skip most of the math. If all we have is monopolistic competition, the labor demand condition becomes:

$$w_t = \frac{\epsilon - 1}{\epsilon} A_t \tag{16}$$

$\epsilon > 1$. It measures the elasticity of substitution among intermediate goods. Imperfect substitutability gives these firms pricing power. They desire a constant markup of price over marginal cost. Since all firms face the same productivity and same wage, they all have the same marginal cost, and hence all choose exactly the same price:

$$P_t = \frac{\epsilon}{\epsilon - 1} \frac{W_t}{A_t} \tag{17}$$

Nominal marginal cost is W_t/A_t, where W_t is the nominal wage. We can define real marginal cost as $mc_t = \frac{W_t}{A_t}\frac{1}{P_t} = \frac{w_t}{A_t}$. This gives us (16). If we take this labor demand condition under flexible prices (but with monopolistic competition), we can derive a simple expression for the *flexible price level* of output. We have:

$$\theta\left(\frac{Y_t^f}{A_t}\right)^{\chi} = \left(Y_t^f\right)^{-\sigma} \frac{\epsilon - 1}{\epsilon} A_t$$

Which implies:

$$Y_t^f = \theta^{-\frac{1}{\sigma+\chi}} \left(\frac{\epsilon - 1}{\epsilon} \right)^{\frac{1}{\sigma+\chi}} A_t^{\frac{1+\chi}{\sigma+\chi}} \tag{18}$$

Note that flexible price output, Y_t^f, only differs by a *constant*, $\left(\frac{\epsilon-1}{\epsilon}\right)^{\frac{1}{\sigma+\chi}}$, relative to efficient output, (14). If $\epsilon \to \infty$, this constant is one, and the equilibria are identical. $\epsilon < \infty$ makes the term $\left(\frac{\epsilon-1}{\epsilon}\right)^{\frac{1}{\sigma+\chi}} < 1$, so the flexible price equilibrium output level is distorted (by a constant amount) relative to the efficient equilibrium. This distortion could be "undone" with a subsidy to labor. In linearized terms, constants drop out, so flexible price output and efficient output move one-to-one.

The "natural rate of interest" is the real interest rate that would be consistent with the flexible price allocation of output. From the Euler equation, this satisfies:

$$1 = \mathbb{E}_t \beta \left(\frac{Y_{t+1}^f}{Y_t^f} \right)^{-\sigma} (1 + r_t^f) \tag{19}$$

Note that the natural rate of interest depends on expected (flexible price) output growth. Since flexible price output differs from efficient output by a constant, this drops out, and $r_t^f = r_t^e$.

2 Adding Sticky Prices

To the model with monopolistic competition, we add sticky prices. ϕ measures the probability that a price chosen today remains in effect tomorrow. It is also the fraction of firms who cannot adjust their price in a given period. This all gets a bit messy, and I don't want you to get stuck in the algebra. What happens is the following. Firms who can adjust their price (a fraction $1-\phi$ of them) choose the same price, $P_t^{\#}$. This price is chosen so that, roughly, these firms get the right markup on average (recall the desired markup is $\epsilon/(\epsilon-1)$), taking as given the probability that a price chosen today will still be in effect in the future. The firms who do not get to adjust (a fraction $1-\phi$), charge whatever price they last posted.

The Calvo assumption is really just an aggregation trick. With Calvo, the aggregate price level evolves according to:

$$P_t^{1-\epsilon} = (1-\phi) \left(P_t^{\#} \right)^{1-\epsilon} + \phi P_{t-1}^{1-\epsilon} \tag{20}$$

In other words, $\phi > 0$ (i.e. some price stickiness), makes the aggregate price level a state variable. The aggregate price level (raised to the $1-\epsilon$) is a convex combination of the reset price (raised to the $1-\epsilon$) with the lagged price level (raised to the $1-\epsilon$). The bigger is ϕ, the slower-moving is the aggregate price level. Although we write this condition in terms of the gross inflation rate (and the relative reset price, $p_t^{\#} = P_t^{\#}/P_t$), this is what is going on.

3 Verbal Intuition for Effects of Shocks

We can think about *aggregate* labor demand in this model as being given by:

$$w_t = mc_t A_t$$

mc_t is real marginal cost – this is the inverse of the markup of price over marginal cost. When prices are flexible, the markup is constant, so $mc_t = \frac{\epsilon-1}{\epsilon}$ (which is less than one, so monopolistic competition distorts labor demand).

Let's suppose that there is an aggregate productivity shock – A_t goes up. For a given real wage, this makes marginal cost go down (the price markup goes up). Firms would like to *lower* their prices so that they have their desired markups. But only a fraction $1 - \phi$ of firms can lower their prices – the other fraction ϕ are stuck. This makes the aggregate price level fall *less* than it otherwise would in the absence of price rigidity. Hence, the aggregate price markup goes *up*, or equivalently real marginal cost (the inverse price markup) goes *down*. This is like a negative labor demand shock (equivalent to an increase in the labor tax rate, or in wedges terminology, a labor wedge shock). Hence, in aggregate, hours, N_t, and therefore output, Y_t, go up less than they would if prices were flexible.

Now suppose there is a demand shock – something like a *negative* u_t that lowers the interest rate for a given inflation rate. This would make the real interest rate lower (holding everything else fixed), which would stimulate demand. The only way to produce more is for people to work more, which would require a higher real wage. A higher wage would mean that marginal cost is higher – firms would like to raise their prices to keep their markup constant. But only a fraction can. Some firms raise their prices, which pushes aggregate inflation up. But the aggregate price level under-reacts, the aggregate markup goes down, and therefore real marginal cost goes up. But in terms of labor demand, this raises aggregate labor demand. Firms that can't adjust their prices have to produce more than they would otherwise like, because their markups are too low. Hence, N_t and therefore Y_t go up.

The output gap is defined as the difference between actual and flexible price output (where the latter is proportional to efficient output):

$$X_t = \frac{Y_t}{Y_t^f} \tag{21}$$

Or, in log deviations:

$$x_t = y_t - y_t^f \tag{22}$$

Figure 1 plots a hypothetical evolution of efficient, flexible price, and actual output over time. The efficient level of output fluctuates due to changes in A_t. The flexible price level of output differs from efficient by a constant gap, which is a function of the inverse markup, but its movements are the same. Efficient output is shown in black and flexible price output is shown in blue. Actual output (shown in red) fluctuates around the flexible price level of output. When output is above flexible price output, it means that markups are too low (real marginal cost is too high) relative to a world with flexible prices and vice-versa.

Figure 1: Efficient, Flexible, and Actual Output Over Time

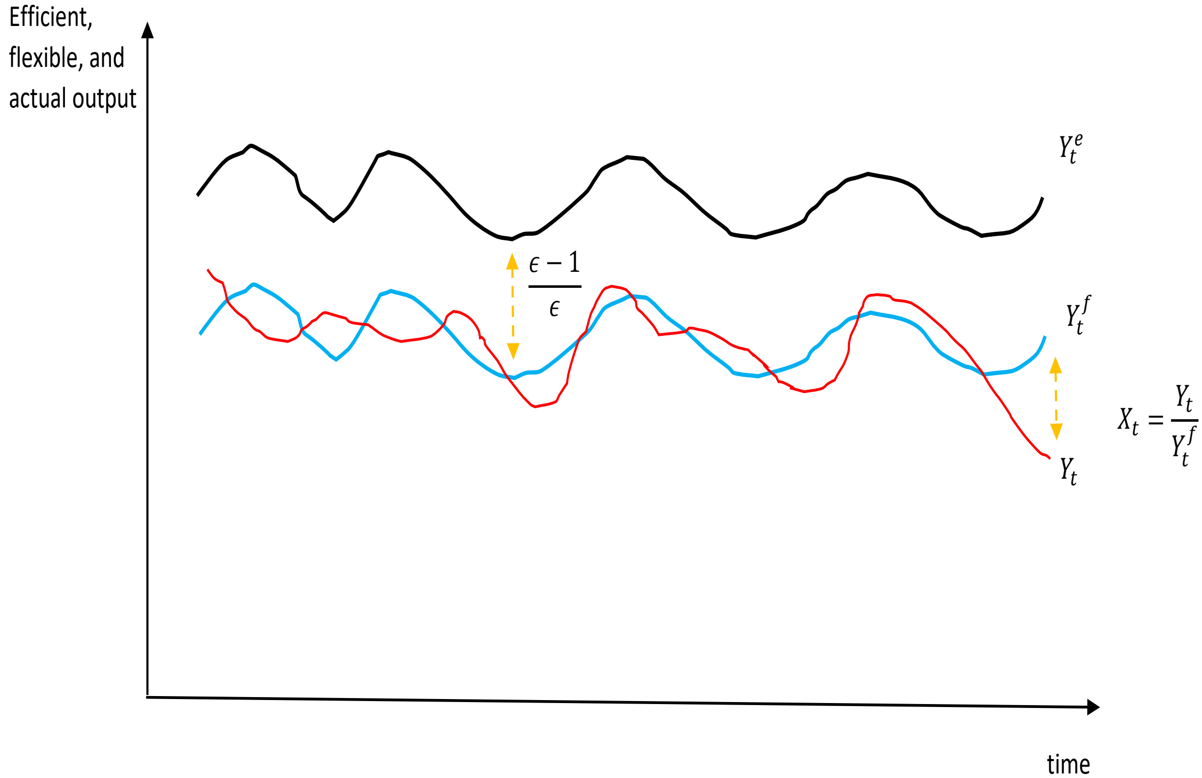

Monetary policy cannot affect the time paths of Y_t^e or Y_t^f. It can influence Y_t (the red line). Fiscal policy, via something like a labor subsidy, could affect Y_t^f (and hence the point about which Y_t fluctuates). Monetary policy can affect the path of Y_t. In periods where $Y_t > Y_t^f$, monetary policy could "tighten" (raise interest rates) and vice-versa.

4 Graphical Intuition in the Linearized Model

After jumping through a bunch of hoops, one can get the linearized Phillips Curves and IS equations above, (1) and (2). Note that there are different ways to write these, such as:

$$\pi_t = \frac{(1-\phi)(1-\phi\beta)}{\phi}\widetilde{mc}_t + \beta\mathbb{E}_t\pi_{t+1} \tag{23}$$

$$\pi_t = \gamma(y_t - y_t^f) + \beta\mathbb{E}_t\pi_{t+1} \tag{24}$$

$$y_t = \mathbb{E}_t y_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1}\right) \tag{25}$$

Depending on the application, different ways of writing the model can be helpful. For now, let's go with the "gap" formulation.

We can combine the IS and policy rules to have an aggregate demand curve. For simplicity, I am going to assume that $\phi_x = 0$. This means that the combined IS/policy rule equation is:

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left(\phi_\pi \pi_t + u_t - \mathbb{E}_t \pi_{t+1} - r_t^f\right) \tag{26}$$

As long as $\phi_\pi > 0$, this is downward-sloping. We can plot both the Phillips Curve and the combined IS/policy rule together in a graph with π_t on the vertical axis and x_t on the horizontal. The Phillips Curve crosses through the point where $x_t = \pi_t = 0$. I show a hypothetical vertical supply curve there, which occurs when $\gamma \to \infty$ (equivalently, when $\phi \to 0$).

Figure 2: AD-AS Representation

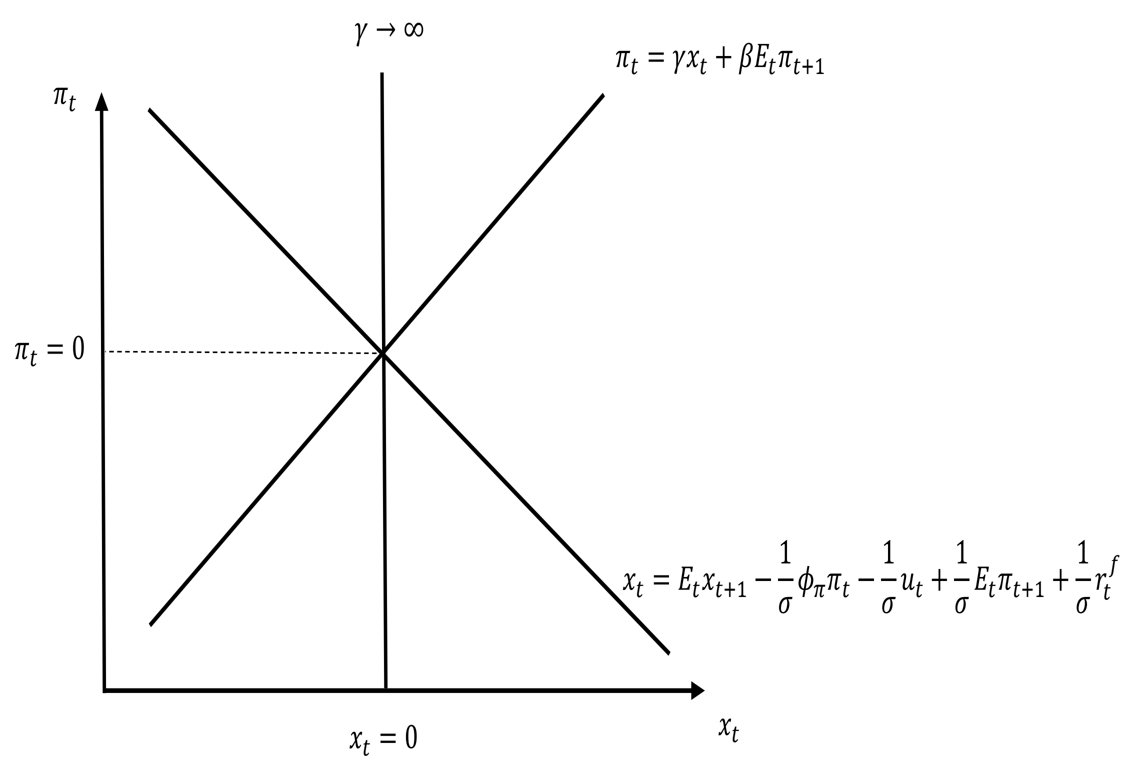

The graph is a little tricky (even dangerous) to use because I'm treating future endogenous variables as given, and the model is (very) forward-looking. I am ignoring issues about determinacy and uniqueness in drawing this in static form, and am explicitly not thinking about dynamics. But we can still use this to think about intuition.

Suppose that there is a monetary shock: u_t goes down (the central bank cuts interest rates exogenously). A reduction in u_t shifts the AD curve to the right, as shown below:

Figure 3: AD-AS Representation, Monetary Shock

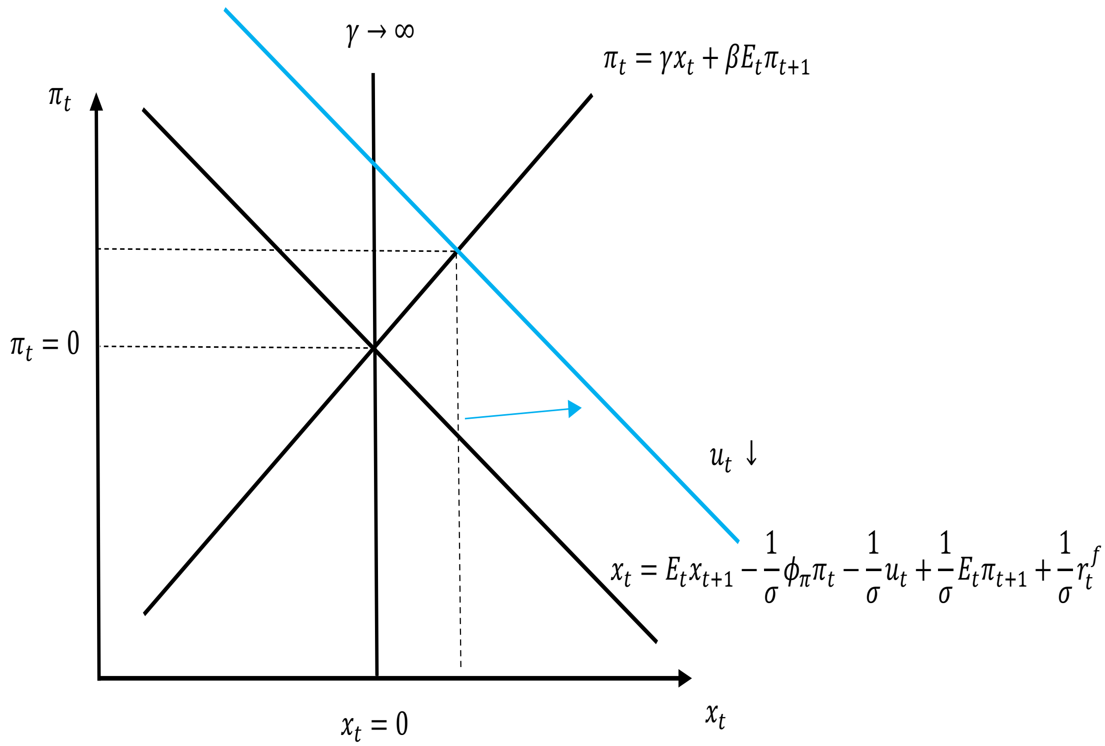

As long as the Phillips Curve is not vertical, this will cause the output gap to increase (and hence output, y_t, to increase, since y_t^f isn't affected) and inflation to rise. If prices were flexible, inflation would rise even more and output would be unaffected. If one wanted to think about things in terms of the money supply, an increase in M_t causes an increase in $m_t = M_t/P_t$ to the extent to which P_t cannot fully adjust (in the short run). When m_t goes up, $C_t = Y_t$ has to go up as well.

Now consider a "supply" shock – this is a bit weird, because it's also going to shift what I'm calling the AD curve. But that is a function of the "gap formulation" (see more below). In particular, suppose that r_t^f declines (this is what would happen after a *positive* shock to productivity).

Figure 4: AD-AS Representation, Productivity Shock

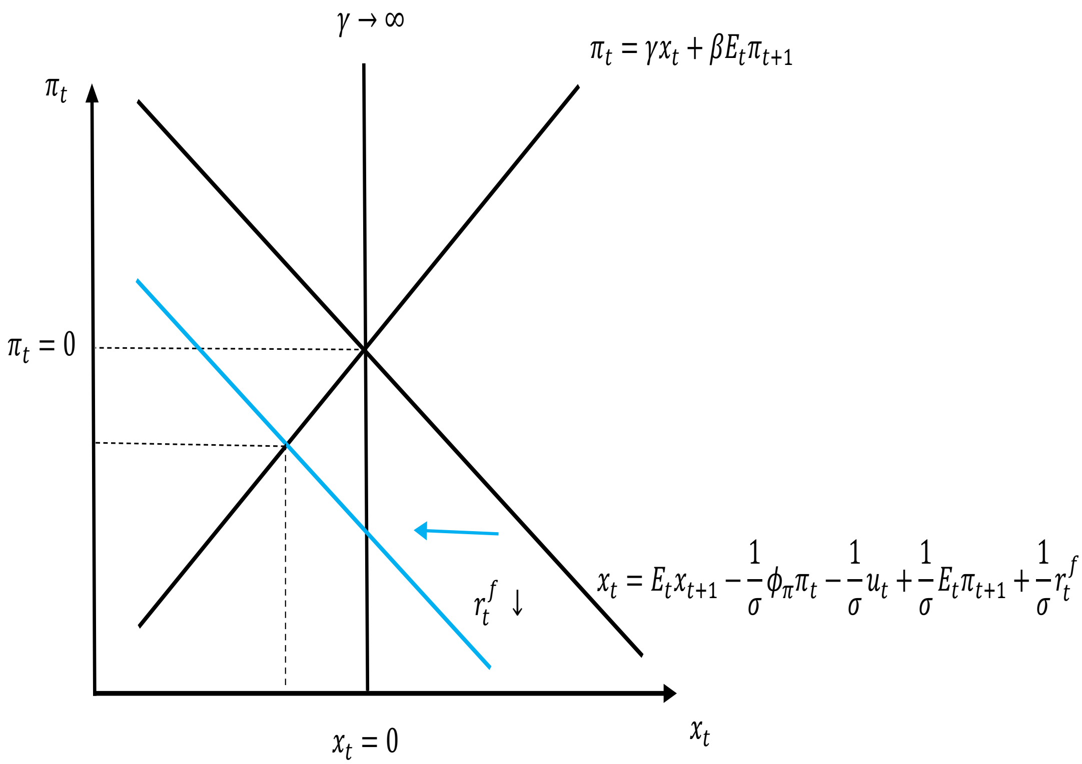

The demand curve shifts in. If the supply curve isn't vertical, inflation doesn't fall enough, and the output gap declines. Since $x_t = y_t - y_t^f$, this doesn't necessarily mean that y_t is falling – it's just that it's not rising as much as y_t^f is. Inflation undershoots (falls less than it would if prices were flexible). Again, if one wanted to think about this in terms of money supply and demand, $m_t = M_t/P_t$ is falling, but not by as much as it would need to (because P_t is sticky) to implement the flexible price allocation.

Note: one could add other shocks (like government spending or preference shocks) to this basic model. It makes things more complicated but the basic intuition for how things are working still obtains.

4.1 Optimal Policy

As we see above, effectively monetary policy controls the slope of the AD curve via ϕ_π. If $\phi_\pi \to \infty$, the AD curve becomes horizontal (shown in the green line below). If the objective is to limit fluctuations in both π_t and x_t, this is good! With a flat AD curve, neither r_t^f nor u_t cause it to shift. And, without a cost-push shock, the Phillips Curve doesn't shift. So committing to a strict inflation target stabilizes *both* π_t and x_t – both will be constant at 0, regardless of what is happening to r_t^f (or u_t). We can think about the policy rule being a targeting rule like $\pi_t = 0$ or a Taylor rule with $\phi_\pi \to \infty$. Pretty cool!

Figure 5: Optimal Policy: Horizontal AD Curve

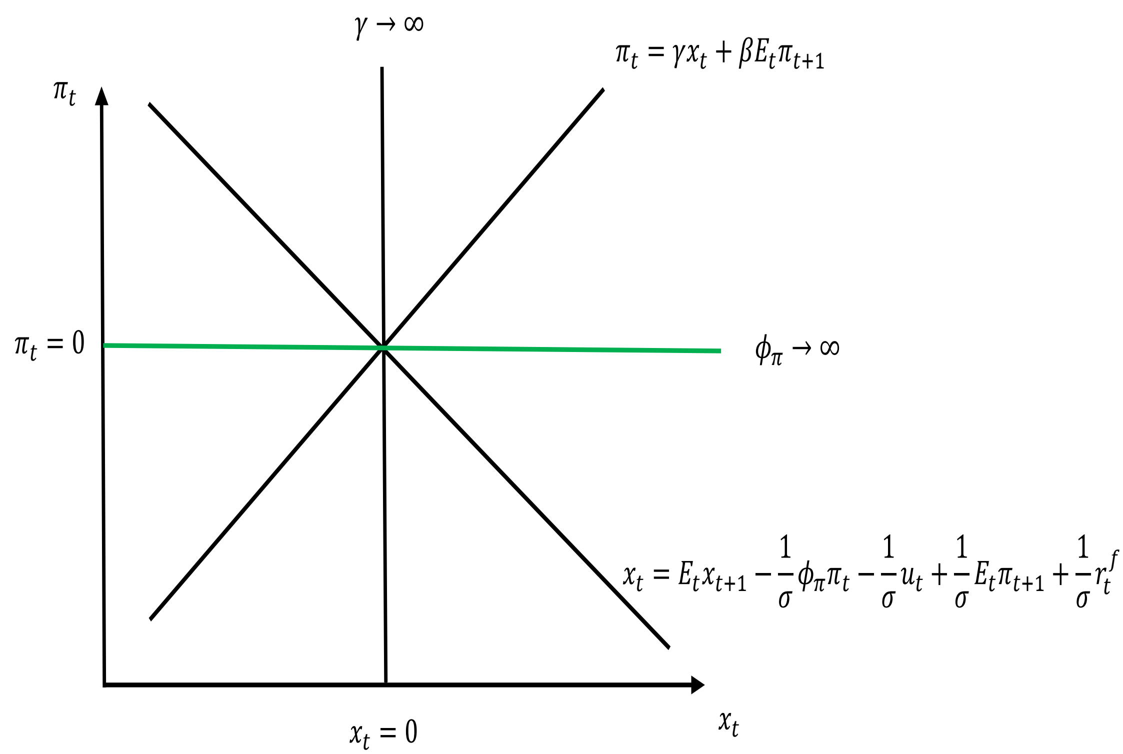

This picture provides the basic gist of the concept of the "Divine Coincidence." Given that shocks only load onto the demand curve (as written, via r_t^f), implementing policy so as to make the demand curve horizontal (a strict inflation target) completely neutralizes the effects of those shocks on *both* the output gap and inflation. If we added a cost-push shock, the Phillips Curve would shift independently of r_t^f. A horizontal demand curve would not necessarily be optimal – this would entail relatively large fluctuations in x_t in response to the cost-push shock with no change in π_t. With a downward-sloping AD curve with a finite ϕ_π, cost-push shocks would move both inflation and output. How steep (or not) the central bank would want to make the AD curve would depend on ω (its relative weight attached to fluctuations in the output gap).

4.2 Alternative Graphical Representation

The above is very helpful, I think. But it is a bit weird – "supply" shocks impact the demand curve. We could add a cost-push shock, which would only shift the NKPC, but it's a bit weird to think about r_t^f as shifting demand, not supply.

It's easy to deal with this. Don't write the equations in terms of the gap; use what we had above:

$$\pi_t = \gamma(y_t - y_t^f) + \beta \mathbb{E}_t \pi_{t+1} \tag{27}$$

$$y_t = \mathbb{E}_t y_{t+1} - \frac{1}{\sigma}\left(\phi_\pi \pi_t + u_t - \mathbb{E}_t \pi_{t+1}\right) \tag{28}$$

The graphs look the same, except we have y_t on the horizontal axis (instead of x_t). The Phillips Curve crosses through the point $y_t = y_t^f$ when $\pi_t = 0$, and r_t^f does not appear in the demand curve.

Figure 6: Alternative AD-AS Representation

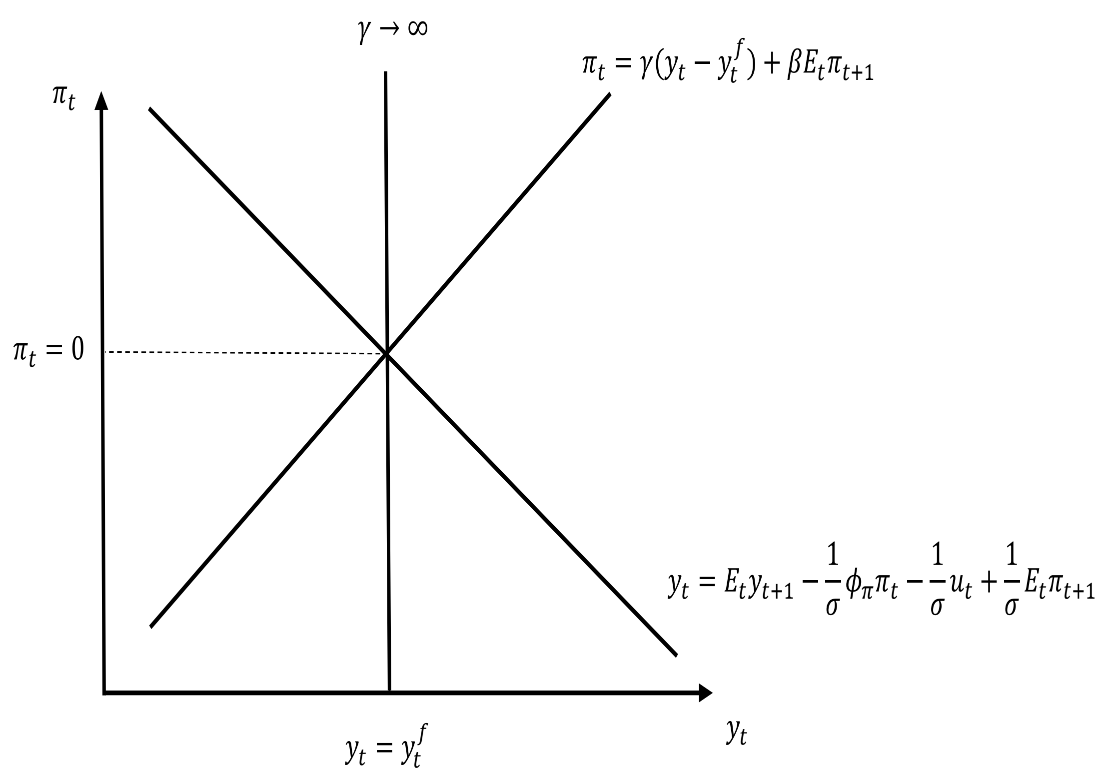

Consider now a monetary shock – a reduction in u_t (a loosening of interest rate policy). This shifts the AD curve out to the right, just as it did above in the gap formulation. This causes output, y_t, and π_t to rise (provided the Phillips Curve is not vertical). How the shock transmits into output and inflation depends on the slope of the Phillips Curve – the flatter it is, the more output rises (and the less inflation rises) and vice-versa). Since y_t^f is unaffected by u_t, we see that the increase in y_t causes the output gap to become positive.

Figure 7: Alternative AD-AS Representation, Monetary Shock

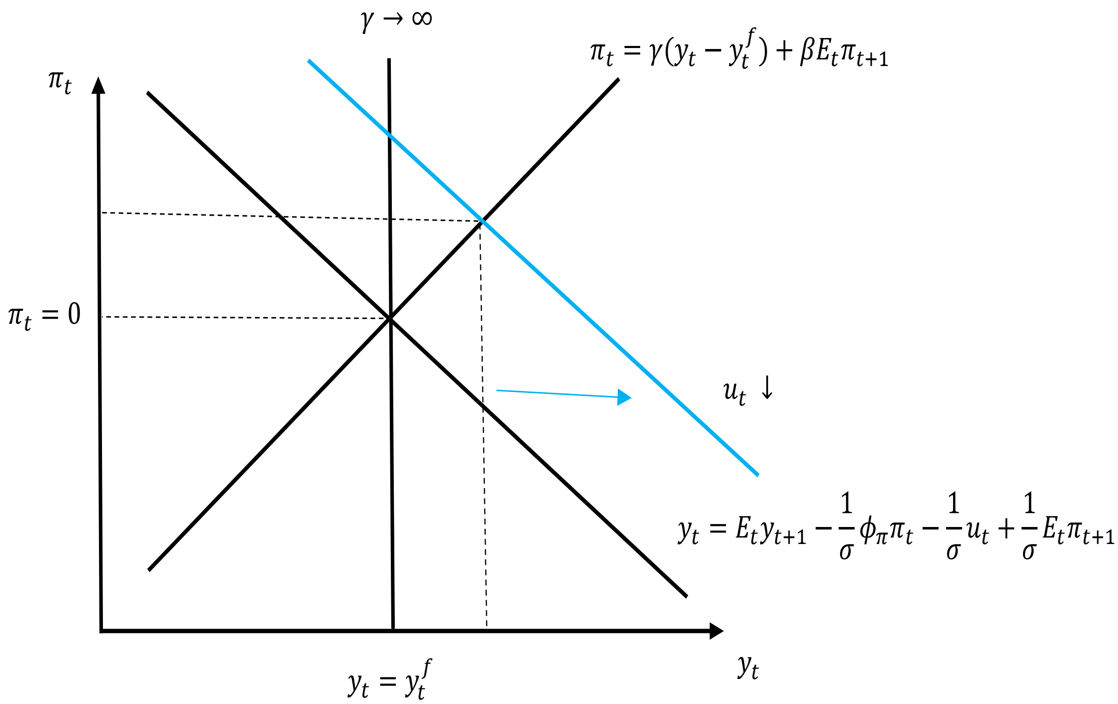

Now let's think about a productivity shock. This is a bit more natural when the equations are written this way and the graphs are drawn with y_t, instead of x_t, on the horizontal axis. An increase in a_t causes y_t^f to increase (and r_t^f to decline, but r_t^f no longer appears in the equations). The increase in y_t^f causes the Phillips Curve to shift to the right, horizontally by the amount of the increase in y_t^f. This is shown with the vertical Phillips Curve (corresponding to $\gamma \rightarrow \infty$) shift horizontally by the same amount. Output rises and inflation falls. But – importantly – output rises *by* less than the horizontal shift of the Phillips Curve (unless the Phillips Curve is vertical). In other words, output *under-reacts* to the productivity shock; this under-reaction will be bigger the flatter is the Phillips Curve (i.e. the smaller is γ).

Figure 8: Alternative AD-AS Representation, Productivity Shock

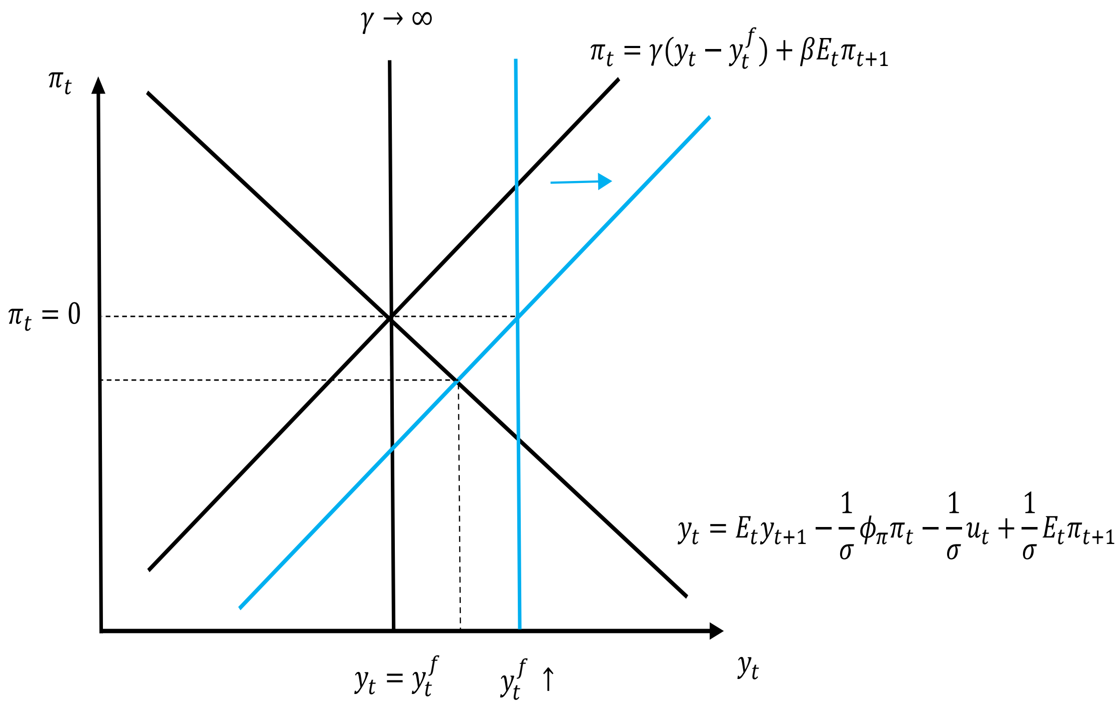

These last couple of pictures nicely make a point I've made in class. By making the Phillips Curve non-vertical, price stickiness *amplifies* the effects of demand shocks on output (which otherwise would do nothing) and *dampens* the effects of supply shocks on output. This is all coming through the slop of the Phillips Curve. Optimal monetary policy entails trying to limit the effects of demand shocks and trying to stabilize the output gap, and policy can impact things by affecting the slope of the AD curve.

The notion of the Divine Coincidence and the desirability of inflation targeting still holds here. We can think about a strict inflation target as having $\phi_\pi \to \infty$, which makes the aggregate demand curve horizontal. This is good – it doesn't shift when u_t changes, and output changes by the full amount of whatever shift in y_t^f there is, so $x_t = y_t - y_t^f = 0$ and inflation remains constant.