Chapter 17
Optimal Monetary Policy

We now begin thinking about optimal monetary policy. Our notion of "optimality" will be from the perspective of the representative consumer's utility function. That is, we will suppose that the monetary authority's objective in setting policy is to maximize the utility of the representative consumer. This seems like a natural notion of optimality – it builds in the idea that the "government," here in the guise of a central bank, exists to try to make people as well off as possible. There may be other notions of optimality one might want to consider, as well. For example, perhaps central bankers are primarily interested in being re-appointed by the government, and perhaps being re-appointed involves different incentives than simply maximizing the utility of the consumers in the economy. Such political considerations are interesting ones to think about when studying the determination of policy, but given our focus thus far, we will only consider the first notion of optimality – that of maximizing the utility of the representative consumer.

This look at optimal policy is an introductory look. Two of the limitations we impose on our analysis here is that we only consider steady-state optimal monetary policy, and we do not consider fiscal policy and its possible impacts on the conduct of monetary policy. Limiting our scope this way will allows us hone in on core principles of monetary policy; it will also allow us to develop the basic mode of analysis for all optimal policy problems without too many extraneous issues. After we have developed the basic results and intuition, we will later study to what extent the lessons learned here carry over to richer environments in which we study dynamically-optimal (as opposed to just steady-state) policy as well as interactions between fiscal and monetary policy.

The structure of any optimal policy problem is the following. We must first specify how households make optimal choices (including, as a preliminary step, what sorts of assets are available to consumers). We must also specify how production occurs and how firms make optimal choices. We must then consider *simultaneously* the optimal choices of both households and firms *along with* the resource constraint of the economy; together, all of these elements comprise the **equilibrium** of the economy. The way we will think of policy-makers (here, monetary policy-makers, but later fiscal policy-makers as well) is that they sit "above" the economy, watching how equilibrium unfolds. Policy-makers understand that for any **given** policy they choose, the private sector (consumers and firms) will make optimal choices that will result in *some* equilibrium. The various equilibria that result for any given, arbitrary, policy can be welfare-ranked according to the respresentative agent's utility function. That is, we can evaluate the welfare of any given policy by simply inserting the resulting equilibrium levels of consumption (and leisure, if we allow for leisure in our model) into consumers' utility functions. We can think of *optimal policy problems* as problems of choosing the *best* equilibrium, where "best equilibrium" means the one that maximizes the utility of the representative consumer.

We will develop the details of the analysis using the basic cash-in-advance model. The basic results and intuition carry over to other typical monetary models as well, including the MIU model and the shopping-time model.

Firms

The way in which we model firms here is the simplest possible way: the representative firm simply hires labor each period in perfectly-competitive labor markets and sells its output. The production technology we assume here is also as simple as possible, linear in labor: $y_t = f(n_t) = n_t$. Firms' profits in period t (in nominal terms) are thus simply $P_t y_t - W_t n_t$, where the notation is standard: P_t is the nominal price of goods, W_t is the nominal wage, and n_t is the quantity of labor. When the firm is maximizing profits, we assume it takes as given both the nominal price P and the nominal wage W.[132] Substituting the linear production technology into the profit function and optimizing with respect to n_t (the only thing the firm decides here is how many units of labor to hire on a period-by-period basis) yields the firm first-order condition $P_t - W_t = 0$. If we define, as usual, the real wage as $w_t = W_t / P_t$, the result of firm profit-maximization is

$$w_t = 1 \tag{1.25}$$

Condition (1.25) is one of the equilibrium conditions of the simple model we are developing, and is the only one that arises from the firm (supply) side of the model.

Consumers

We will model consumers using our cash-in-advance (CIA) specification. The representative consumer begins period t with nominal money holdings M_{t-1}, nominal bond holdings B_{t-1}, and stock (a real asset) holdings a_{t-1}. The period-t budget constraint of the consumer is

$$P_t c_t + P_t^b B_t + M_t + S_t a_t = W_t n_t + M_{t-1} + B_{t-1} + (S_t + D_t) a_{t-1} + \tau_t, \tag{1.26}$$

[132] Nothing more than our usual assumption of price-taking behavior; here, price-taking describes the firm's behavior in both output markets and input markets.

where the notation again is as in the MIU model of Chapter 14: S_t is the nominal price of a unit of stock, D_t is the nominal dividend paid by each unit of stock, and P_t^b is the nominal price of a one-period nominal bond with face-value $1. Because we continue to assume that all bonds are one-period bonds and the face value of each bond is $FV = 1$, we have that $P_t^b = \frac{1}{1+i_t}$ (which you should recall), where i_t is the net nominal interest rate on a nominal bond held from period t to period t+1. Note the term $W_t n_t$ in the budget constraint: it represents total labor income in period t.[133] The consumer takes the wage W_t as given.[134]

The term τ_t in the budget constraint is a lump-sum amount that consumers receive from (or must pay to) the central bank. This τ_t is the means by which the monetary authority achieves changes in the money supply: a positive value of τ_t means the government is expanding the money supply in period t, and a negative value of τ_t means the government is contracting the money supply in period t. We return to how τ_t is set when we discuss below what the monetary authority does.

The consumer also faces in each period the CIA constraint

$$P_t c_t = M_t . \qquad (1.27)$$

The instantaneous utility function of the consumer is $u(c_t, 1-n_t)$.[135] Note that the second argument is indeed *leisure* (as in our basic consumption-leisure model of Chapter 2), which equals the total time endowment minus the amount of time spent working.

---

[133] Thus, what we have here on the consumer side (demand side) of the model is the "intertemporal consumption-leisure model" studied earlier, in the form of an infinite-period model. That is, we have a dynamic setting in which consumers repeatedly (sequentially) make consumption-leisure decisions along with consumption-savings decisions. By now, this sort of idea should be straightforward.

[134] We, "the modeler," know from the firm optimality condition (1.25) that it will (in equilibrium) be the case that $W_t = P_t$; however, the consumer need not "understand" this; all the consumer does is take *whatever* W_t is as given.

[135] The (1-n) term is essentially just the (168 – n) term we considered in our basic consumption-leisure model. There, we simply supposed (arbitrarily) that a week was the relevant period of time to consider, in which there are 168 total hours. Our measure of a period is completely arbitrary, hence our "total number of hours in a period" is also completely arbitrary. Here, we have adopted the more usual convention in macro models that total time available in a period is "one unit." With total time of one unit, consumers work a *fraction* n of their total time and hence spend the fraction $1-n$ of their total time in leisure; that is, having normalized total time to one unit, n ends up measuring the fraction of the consumer's hours that are spent working. If we wanted to normalize things back to hours in a week, we would simply multiply n by 168.

Therefore, with subjective discount factor β, the consumer's lifetime utility function beginning from period t is given, as usual, by $\sum_{s=0}^{\infty}\beta^s u(c_{t+s},1-n_{t+s})$.

We will use the sequential Lagrangian approach to solve the consumer's utility maximization problem. Let λ_t be the Lagrange multiplier on the period-t budget constraint, and let μ_t be the Lagrange multiplier on the period-t CIA constraint. Writing out the first couple of terms of the Lagrangian,[136] we have

$$\begin{aligned}
&u(c_t,1-n_t)+\beta u(c_{t+1},1-n_{t+1})+\beta^2 u(c_{t+2},1-n_{t+2})+...\\
&+\lambda_t\left[W_t n_t + M_{t-1}+B_{t-1}+(S_t+D_t)a_{t-1}+\tau_t-P_t c_t-P_t^b B_t-M_t-S_t a_t\right]\\
&+\mu_t\left[M_t-P_t c_t\right]\\
&+\beta\lambda_{t+1}\left[W_{t+1}n_{t+1}+M_t+B_t+(S_{t+1}+D_{t+1})a_t+\tau_{t+1}-P_{t+1}c_{t+1}-P_{t+1}^b B_{t+1}-M_{t+1}-S_{t+1}a_{t+1}\right]\\
&+\beta\mu_{t+1}\left[M_{t+1}-P_{t+1}c_{t+1}\right]\\
&+....
\end{aligned} \tag{1.28}$$

In period t, the consumer chooses (c_t,n_t,M_t,B_t,a_t). Proceeding mechanically, the first-order-conditions with respect to each of these five choice variables, respectively, are:

$$u_1(c_t,1-n_t)-\lambda_t P_t-\mu_t P_t=0 \tag{1.29}$$

$$-u_2(c_t,1-n_t)+\lambda_t W_t=0 \tag{1.30}$$

$$-\lambda_t+\mu_t+\beta\lambda_{t+1}=0 \tag{1.31}$$

$$-\lambda_t P_t^b+\beta\lambda_{t+1}=0 \tag{1.32}$$

$$-\lambda_t S_t+\beta\lambda_{t+1}(S_{t+1}+D_{t+1})=0 \tag{1.33}$$

You should be able to recognize that these first-order conditions are essentially identical to those from our study of the infinite-period MIU model, with of course the exception that there we ignored the labor-leisure dimension.[137]

[136] By now, formulating the Lagrangian has to be essentially automatic. If it is not, now is certainly the time to go back and review this type of formulation.

[137] Indeed, comparing these FOCs with the FOCs in our discussion of monetary policy shows that the FOCs on consumption, stock, and bonds are very similar; the FOC on money is identical once one recognizes that the multiplier μ_t in our CIA model here is effectively just the term $\dfrac{u_2\left(c_t,\dfrac{M_t}{P_t}\right)}{P_t}$ (which involves the marginal utility of real money balances) from the MIU model.

Conditions (1.29) through (1.33) describe how consumers make optimal choices; as such, they represent equilibrium conditions. As usual, though, it is instructive to not work with these raw first-order conditions directly, but instead combine them into interpretable expressions of the form "MRS equals a price ratio" which are the cornerstone of consumer theory. Begin by re-writing the FOC on consumption as

$$\frac{u_1(c_t, 1-n_t)}{P_t} = \lambda_t + \mu_t. \tag{1.34}$$

Next, note that (1.31) and (1.32) can be combined to give $\mu_t = \lambda_t(1-P_t^b)$, a relationship between Lagrange multipliers. For the most part, we have tried to avoid thinking directly in terms of multipliers, but for optimal policy issues, it turns out to often be useful to think directly in terms of multipliers.

Very briefly (and somewhat informally), a Lagrange multiplier measures the marginal utility value of relaxing a particular constraint by a small (marginal) amount. Take the budget constraint, for example. If somehow we (the modeler or, more pertinently, the policy-maker) could add a little extra income or wealth (which could be in the form of money but need not be) to the consumer's budget, the multiplier ends up measuring the marginal utility of those extra resources to the consumer. Similarly, if we (the modeler or the policy-maker) could hand a little *money* (**specifically** money in this case) to the consumer, *both* of the multipliers λ_t *and* μ_t would somehow be involved in determining the marginal utility of this to the consumer.[138]

Regardless, using $\mu_t = \lambda_t(1-P_t^b)$ in condition (1.34) gives us

$$\frac{u_1(c_t, 1-n_t)}{P_t} = \lambda_t + \lambda_t(1-P_t^b). \tag{1.35}$$

Next, because $P_t^b = \frac{1}{1+i_t}$, we can write this as

$$\frac{u_1(c_t, 1-n_t)}{P_t} = \lambda_t + \lambda_t\left(\frac{i_t}{1+i_t}\right). \tag{1.36}$$

Next, note that condition (1.30) yields $\lambda_t = u_2(c_t, 1-n_t)/W_t$; inserting this in (1.36) gives

$$\frac{u_1(c_t, 1-n_t)}{P_t} = \frac{u_2(c_t, 1-n_t)}{W_t} + \frac{u_2(c_t, 1-n_t)}{W_t}\left(\frac{i_t}{1+i_t}\right). \tag{1.37}$$

[138] You may have encountered this notion of multipliers measuring marginal utility in an intermediate microeconomics course; any more advanced study of this idea, however, we leave for a more advanced course in economic theory.

Divide this through by $u_1(c_t, 1-n_t)$ and then multiply through by W_t to arrive at

$$\frac{u_2(c_t,1-n_t)}{u_1(c_t,1-n_t)} + \frac{u_2(c_t,1-n_t)}{u_1(c_t,1-n_t)}\left(\frac{i_t}{1+i_t}\right) = \frac{W_t}{P_t}. \tag{1.38}$$

The right-hand-side of this is simply the real wage, w_t. The left-hand-side resembles the MRS between consumption and leisure; indeed, if the nominal interest rate were zero ($i_t = 0$) it would be exactly that MRS, and condition (1.38) would be exactly the consumption-leisure optimality condition we have already studied, which would state, as usual, that at the consumer's optimal choice, the MRS between consumption and leisure (which, again, is what $1-n$ is) equals the real wage.[139]

Combining the terms on the left-hand-side and re-arranging slightly, we can write this optimality condition as

$$\frac{u_2(c_t,1-n_t)}{u_1(c_t,1-n_t)} = w_t\left[1+\frac{i_t}{1+i_t}\right]^{-1}, \tag{1.39}$$

which shows exactly how, at the optimal choice, the MRS between consumption and leisure depends on the real wage and the nominal interest rate.

The nominal interest rate is linked to the multipliers μ_t and λ_t via

$$\mu_t = \lambda_t\left(\frac{i_t}{1+i_t}\right), \tag{1.40}$$

which, recall, comes from combining (1.31) and (1.32) and using the relationship between the nominal price of bonds and the nominal interest rate. We could thus alternatively express the optimality condition (1.38) as

$$\frac{u_2(c_t,1-n_t)}{u_1(c_t,1-n_t)} + \frac{u_2(c_t,1-n_t)}{u_1(c_t,1-n_t)}\left(\frac{\mu_t}{\lambda_t}\right) = w_t. \tag{1.41}$$

What condition (1.39) or, equivalently, (1.41), reveals is that the presence of money in this economy somehow throws a "wedge" into the consumption-leisure optimality condition. If the nominal interest rate were zero, condition (1.39) tells us the consumption-leisure margin would be set according to just the real wage. Condition

[139] Here, with no taxes, which we allowed for in the basic consumption-leisure model. We will re-introduce taxation in our second, more in-depth study of optimal policy when we consider the joint effects of fiscal and monetary policy.

(1.40) tells us that if the nominal interest rate were zero, then the multiplier μ_t would be zero; condition (1.41) then tells us (again) that the consumption-leisure margin would be set according to just the real wage.

The wedge being thrown into the consumption-leisure margin stems from the CIA constraint itself. Indeed, one way of thinking about an economy *without a CIA constraint* is to assume that $\mu_t = 0$ always, in which case clearly (1.41) shows us that the consumption-leisure margin would depend on just the real wage. Apparently, though, the presence of this CIA wedge means that at the consumer's optimal choice, the MRS between consumption and leisure depends on *both the real wage and the nominal interest rate.*

Thinking of money, and more generally, government policy variables, as throwing "wedges" into consumer optimality conditions is an important way of understanding the effects of (both monetary and fiscal) policy, as well as understanding how to design *optimal* policies. Indeed, we already know from our study of the basic consumption-leisure model that labor taxes throw a wedge (in the form of one minus the labor tax rate) into the consumption-leisure optimality condition. We now have encountered another policy variable that potentially throws a wedge into this margin: money.

Finally, as usual, the asset a_t in our model allows us to construct an intertemporal optimality condition, linking the real interest rate to consumers' intertemporal marginal rate of substitution. Express condition (1.33) as $\frac{\lambda_t}{\lambda_{t+1}} = \beta \frac{S_{t+1} + D_{t+1}}{S_t}$. Multiply and divide the numerator by P_{t+1}, and multiply and divide the denominator by P_t to get

$$\frac{\lambda_t}{\lambda_{t+1}} = \beta \frac{(S_{t+1} + D_{t+1}) / P_{t+1}}{S_t / P_t} \frac{P_{t+1}}{P_t}. \tag{1.42}$$

Recall from Chapter 14 that we can define the real interest rate as

$$1 + r_t = \frac{(S_{t+1} + D_{t+1}) / P_{t+1}}{S_t / P_t}; \tag{1.43}$$

using this, the previous expression is

$$\frac{\lambda_t}{\lambda_{t+1}} = \beta (1 + r_t) \frac{P_{t+1}}{P_t}. \tag{1.44}$$

Next, recall that the first-order condition on labor can be written $\lambda_t = u_2(c_t, 1 - n_t) / W_t$; using this along with the time t+1 version in (1.44) allows us to express things as

$$\frac{u_2(c_t,1-n_t)}{u_2(c_{t+1},1-n_{t+1})} = \beta(1+r_t)\frac{W_t}{W_{t+1}}\frac{P_{t+1}}{P_t}. \tag{1.45}$$

This is one way of expressing the consumption-savings optimality condition in this model; note that on the left-hand-side, the intertemporal MRS (note the different t and t+1 subscripts!) is expressed in terms of the marginal utility of leisure rather than in terms of the marginal utility of consumption, but this is no problem. The interpretation is the same as any "standard" consumption-savings optimality condition: it describes the consumer's optimal tradeoff *over time.*

Conditions (1.45) and (1.39) (or, equivalently, condition (1.41)) summarize the optimization problem of the representative consumer; understand well that (1.39) condenses conditions (1.29) through (1.32). Condition (1.39) is the model's consumption-leisure (intratemporal) optimality condition, and condition (1.45) is the model's consumption-savings (intertemporal) optimality condition.[140] Also, the CIA constraint itself, condition (1.27), is a description of the household's choices. Thus, condition (1.27), condition (1.45), and condition (1.39) are equilibrium conditions of the model we are developing; they describe the consumer (demand) side of the model.

Government

When studying optimal (government) policy, we need to specify what the government *does.* In the model here, the only thing the government, which is just a monetary authority thus far, *does* is print money and hand it to consumers (or, if shrinking the money supply, "ask for money back from consumers").[141] This handing over of (or "asking to return") money occurs via the τ_t introduced above. The government's (monetary authority's) budget constraint is thus simply

$$M_t = M_{t-1} + \tau_t. \tag{1.46}$$

[140] Recall from our brief consideration of the interemporal consumption-leisure model that *both* a consumption-leisure and a consumption-savings tradeoff arises in a dynamic model once we "glue together" the simple consumption-savings and the simple consumption-leisure models; although couched inside a more complicated model, this idea underpins things here as well.

[141] Obviously, this is not literally true. Remember from basic macroeconomics that in practice central banks use open market operations to expand or contract the money supply. Open-market operations are conducted with the banking (financial) sector; in our model here, we do not include a banking sector, hence our metaphor that the central bank deals directly with consumers. In reality, the banking sector effectively just acts as an intermediary between the central bank and consumers (and firms) in the conduct of monetary policy. Given that our model is already become large (and will soon become even larger as we continue to enrich the scope of issues we want to consider), it seems worthwhile to not try to model a banking sector.

If $\tau_t > 0$, the central bank is expanding the nominal money supply in period t, while if $\tau_t < 0$, the central bank is contracting the nominal money supply in period t.

We can express this transfer τ_t in terms of the growth rate of money. Letting g_t be the net growth rate of nominal money during period t, we can write $\tau_t = g_t M_{t-1}$. Thus, for example, if the central bank expands the money supply by 10 percent in period 16, we would write $g_{16} = 0.10$. With this, the government budget constraint can be written as

$$M_t = (1 + g_t) M_{t-1}. \qquad (1.47)$$

This government budget constraint seems trivial here, but conceptually it will be crucial when we consider the interactions of fiscal and monetary policy.

Resource Constraint

The resource constraint of the economy describes all of the different uses of total output (GDP) of the economy. In our economy, output is produced by the linear-in-labor production technology and there is no other use for output other than consumption. Hence the resource constraint is simply

$$c_t = n_t \qquad (1.48)$$

In richer models that we will develop, government spending and investment will also be components of the resource constraint. Informally, you should think of the resource constraint as the "GDP accounting equation" from basic macroeconomics.[142]

Equilibrium and Steady-State Equilibrium

Any (well-specified) macroeconomic model must have a notion of *equilibrium.* Like in basic microeconomics, equilibrium is a collection of prices and quantities that in concert make all markets clear, given that both demand (consumer choices) and supply (firm choices) decisions in the economy are made optimally. When constructing a macroeconomic model, making sure to identify properly the equilibrium conditions is crucial, and this step can only be done after setting up and solving the household and firm optimization problems In our model here, equilibrium is described by the firm optimality condition (1.25); the household optimality conditions, which, recall, we were able to condense into (1.27), (1.39), and (1.45); and the resource constraint (1.48).

[142] As on the first day of basic macroeconomics, the GDP accounting identity states that GDP = C + I + G + NX. In our model here, I = G = NX = 0, hence all output is simply consumed by consumers.

Because conditions (1.25) and (1.48) are so simple, let's simply substitute the w_t and n_t terms in the other equilibrium conditions, in which case we can represent the equilibrium conditions of the model as

$$c_t = \frac{M_t}{P_t}, \tag{1.49}$$

$$\frac{u_2(c_t, 1-c_t)}{u_1(c_t, 1-c_t)} = \left[1 + \frac{i_t}{1+i_t}\right]^{-1}, \tag{1.50}$$

and

$$\frac{u_2(c_t, 1-c_t)}{u_2(c_{t+1}, 1-c_{t+1})} = \beta(1+r_t). \tag{1.51}$$

Note that in writing (1.51) we are using the fact that in equilibrium, $\frac{W_t}{W_{t+1}}\frac{P_{t+1}}{P_t} = \frac{w_t}{w_{t+1}} = \frac{1}{1}$.

Next, if we take the time t-1 version of (1.49) and divide (1.49) by it, we would have

$$\frac{M_t / M_{t-1}}{P_t / P_{t-1}} = \frac{c_t}{c_{t-1}}. \tag{1.52}$$

Because, by assumption, $M_t = (1+g_t)M_{t-1}$ and, by earlier definition, inflation between period t-1 and t is given by $1+\pi_t = P_t / P_{t-1}$, we can express this as

$$\frac{1+g_t}{1+\pi_t} = \frac{c_t}{c_{t-1}}, \tag{1.53}$$

which shows that the inflation rate between t-1 and t is linked to *both* the rate of money growth between t-1 and t *and* by how much consumption grows between t-1 and t. Condition (1.53) is not a new equilibrium condition of the model; rather it arose from combining the t and t-1 versions of (1.49) with the money supply rule being followed by the government.

Recall the notion of a steady-state equilibrium: a steady state is a condition in which all **real** quantity variables are constant. In terms of the preceding two expressions, steady state would involve constant c and constant i over time.[143] Denote by $\bar{c}$ and $\bar{i}$ the steady-state values of consumption and the nominal interest rates. Imposing steady-state on (1.53) immediately shows us that **in steady state (i.e., in the long run),** the rate of

[143] Note that even though i is the *nominal* interest rate, it is not a *quantity* variable; it is more akin to a price (i.e., the opportunity cost of holding money).(1.56)

inflation **is equal to** the rate of money growth. That is, in this economy, the quantity-theoretic prediction that inflation is governed by money growth *in the long run* shines through. But, note well that *outside* a steady-state, condition (1.53) shows that inflation and the money growth rate *need not be identical.*

Next, the steady-state version of the consumption-savings optimality condition reveals that, in a steady-state equilibrium, $1+r=1/\beta$, a condition we have encountered before: in steady state (again, read this as "in the long run" or "on average"), the real interest rate is determined by the representative consumer's discount factor. Furthermore, because we know the Fisher relation also exists in the background (though we did not formally derive it here, we could have), the steady-state version of (1.51) can also be written

$$1+\pi=\beta(1+i). \tag{1.54}$$

Continuing to link the conditions we are deriving: we just concluded that in steady state, $\pi = g$, that is, the inflation rate equals the money growth rate. Thus, the last expression can be written

$$1+g=\beta(1+i). \tag{1.55}$$

This of course means that $1+i=\frac{1+g}{\beta}$ and $i=\frac{1+g-\beta}{\beta}$. Finally, using these expressions in the steady-state version of the consumption-leisure optimality condition (expressed in (1.50)), we have that

$$\frac{u_2(\bar{c},1-\bar{c})}{u_1(\bar{c},1-\bar{c})}=\frac{1+g}{1+g+1+g-\beta}. \tag{1.56}$$

Let's take stock of where we've arrived. After setting up and solving the firm profit-maximization problem and the consumer utility-maximization problem as well as specifying how monetary policy is conducted (a simple money growth rule), we defined equilibrium. We then condensed the equilibrium conditions into a more compact set of conditions. We then imposed steady-state on the equilibrium conditions, which has allowed us to express the steady-state equilibrium *of the entire private sector of the model* in the single, compact form of equation (1.56). What condition (1.56) describes is *how the steady-state equilibrium level of consumption depends on the steady-state rate of growth of the nominal money supply.* We've compacted the entire model (i.e., its setup and solution) into a single expression. There is no (reliable) shortcut for all of the analysis we have done; one *must* go through the entire solution of the demand and supply sides of the model, description of the equilibrium, and then (*and only then*) can one impose steady-state. For the purpose of what we now (finally) turn to, the optimal (steady-state) policy problem, condition (1.56) is crucial.

Formulation and Solution of Optimal Policy Problem

The reason condition (1.56) is crucial is that *it describes how the private sector of the economy responds (in steady-state) to monetary policy.* That is, it encapsulates the decision-making of consumers and firms, who all *take monetary policy as given.* Recall our heuristic description at the outset of how to think of the optimal policy-maker (or, at least, what one useful way of thinking about policy-makers might be): policy-makers understand how the economy responds to various policy settings, and (if they are benevolent) set policy to bring about the highest possible welfare for the economy. In our model, welfare is naturally measured by the representative consumer's utility, and the idea of "how the economy responds to various policy settings" is captured by the equilibrium of the model, which in turn we have been able (after quite a bit of effort, to be sure) to condense down to the single restriction (1.56). In the parlance of the optimal policy literature, we will say that policy-makers respect the equilibrium of the economy when choosing policy.

Note also the game-theoretic undertones here: it is as if the private sector (both firms and consumers) "moves second," after policy has been set. The "first move" belongs to the policy-maker, who takes into account the "optimal response" of the private sector when deciding its "move" (its policy setting). If you have studied the idea of Bertrand or Cournot competition in intermediate microeconomics, there is some similarity in idea here: here, the policy-maker takes into account the response function of the entire economy when choosing the optimal policy. This idea is an important one, one that will carry over into the richer and richer optimal policy questions we consider.

In terms of setting up the optimal policy problem, condition (1.56) defines the steady-state equilibrium $\bar{c}$ *as a function of* g. To emphasize this **functional** dependence, let's from now on write $\bar{c}(g)$; thus, expression defines the function $\bar{c}(g)$. If you were given a particular utility function, you could solve for the function $\bar{c}(g)$.

Finally, then the optimal (steady-state) policy problem is to choose a (constant) growth rate of money g that maximizes consumer's (steady-state) utility subject to the equilibrium of the economy described by (1.56). In a steady-state equilibrium, consumers' lifetime utility (formally, still starting at date t, but "dates" have less meaning once the economy has arrived in steady state – nonetheless, we will keep this formalism) is given by

$$\sum_{s=0}^{\infty} \beta^s u\left(\bar{c}(g), 1-\bar{c}(g)\right) = \frac{u\left(\bar{c}(g), 1-\bar{c}(g)\right)}{1-\beta}. \tag{1.57}$$

So the optimal policy problem boils down to choosing a constant growth rate of money that maximizes (1.57); mathematically, no constraints are required on this optimization problem *because we have already built all constraints imposed by equilibrium into* (1.57). Being careful in using the chain rule of calculus, the first-order condition of (1.57) with respect to g is

$$u_1\left(\bar{c}(g),1-\bar{c}(g)\right)\cdot\bar{c}'(g)-u_2\left(\bar{c}(g),1-\bar{c}(g)\right)\cdot\bar{c}'(g)=0, \qquad (1.58)$$

where $\bar{c}'(g)$ is how steady-state equilibrium consumption responds to a marginal change in the money growth rate (i.e., it is the derivative of the function $\bar{c}(g)$ with respect to g).[144] As it stands, this condition may not seem particularly illuminating, but it is the solution to the policy problem: the value of g that solves (1.58) *is* the optimal steady-state money growth rate, the one that maximizes consumers' utility, taking into account how the private sector responds to monetary policy.

Notice that we can cancel the $\bar{c}'(g)$ terms from (1.58); doing so and rearranging gives us that under the optimal policy,

$$\frac{u_2\left(\bar{c},1-\bar{c}\right)}{u_1\left(\bar{c},1-\bar{c}\right)}=1. \qquad (1.59)$$

This condition states that *when policy is chosen optimally,* the representative consumer's MRS between consumption and leisure equals one.

Our final task (and most important task conceptually) is to compare (1.59) with (1.56): they look very similar, but they have importantly different interpretations. Condition (1.56) describes, *for an arbitrary money growth rate,* how steady-state equilibrium consumption responds; condition (1.59) instead describes how the marginal utilities of consumption and leisure should relate to each other *when policy is set optimally,* in the sense we have defined. Clearly, the two conditions coincide only if $g=\beta-1$. Thus, the optimal steady-state growth rate of money is $\beta-1$; this is the final solution of the optimal policy problem. **Note that this conclusion has absolutely nothing to do with the precise functional form of the utility function – we have made absolutely no assumptions about what the utility function is.**

With this steady-state optimal growth rate of the money supply in hand, we can also determine what the optimal steady-state nominal interest rate is. To do so, return to the steady-state equilibrium condition (1.55), which, recall, is the Fisher equation with the steady-state version of the CIA constraint imposed on it. From this condition and the result $g=\beta-1$ if policy is set optimally, it immediately follows that

---

[144] Note that we've dropped the $1-\beta$ term from this first-order condition because it does not affect the solution of the policy problem – i.e., we're just dropping a constant.

$$i = 0 \tag{1.60}$$

when monetary policy is set optimally. That is, the optimal nominal interest rate is zero. This result, originally due to Friedman (1969), is a hallmark result in the theory of optimal policy and has come to be known as the **Friedman Rule.** Because of the steady-state equilibrium relationship, $1+g=\beta(1+i)$, we can think alternatively (and equivalently) of the Friedman Rule as $i=0$ and/or $g=\beta-1$.

The Friedman Rule

The Friedman Rule is a celebrated result in monetary theory. Let's first explicitly note the implication of what it states: with $g=\beta-1$ and with $\beta\in(0,1)$, clearly the Friedman Rule means $g<0$. *The optimal money growth rate is negative* – the central bank should steadily contract the nominal money supply. In turn, because in steady state $g = \pi$, if the central bank is running the Friedman Rule, *the optimal inflation rate is negative as well.*

There are a few alternative (and ultimately interrelated) ways to think about this result. One way to think about it is to think about what purpose money serves in this model. The CIA constraint is meant to represent some transactions motive for holding money; it is intended as a way of capturing money's medium of exchange function. Let's step back from the CIA model for a moment, though. The consumers (the representative consumer) in the dynamic macro model we had been building *before* we introduced the topic of CIA were perfectly well able to acquire consumption goods. In the basic dynamic macro model we had built *without* the CIA constraint, consumers were not clamoring for a body called a "central bank" to provide them with an object called "money." In introducing the CIA structure, we as the modelers are *forcing* the consumers in our model to hold and use money. But they didn't need it in the first place. So the optimal thing for the central bank to do it to remove it.

Here is another way to think about the Friedman Rule. There are two nominal assets in the model, money and bonds. Stocks, as we have mentioned before, are fundamentally a real asset. Bonds potentially pay interest at a nominal rate i, but they do not (by assumption) serve a medium of exchange function. Money pays no interest, but consumers must carry it in order to purchase goods. Every dollar of his resources that the consumer holds as money is a dollar he cannot hold in the form of an interest-bearing bond. Thus, there is an opportunity cost of holding money, which is the foregone interest earnings on nominal bonds. The benevolent central bank, realizing that consumers *must* hold money (due to the CIA constraint, which was *not* put into place by the central bank) seeks to make it as costless as possible for them to do so. Making money costless to hold means making the opportunity cost of holding money zero, hence conducting monetary policy in such a way as to make $i = 0$.

Yet another way of considering the optimality of the Friedman Rule is using the basic tools of consumer theory. Recall from our study of the simple consumption-leisure model that at the consumer's optimal choice, the MRS between consumption and leisure is equated to the real wage.[145] Yet, here in the CIA model, we see that consumer' optimal choice depends not only on the real wage but also on the nominal interest rate – condition (1.39) shows us this. Condition (1.39) is the consumption-leisure optimality condition of this model: a positive nominal interest rate is interfering with the (unfettered) optimal choice along the consumption-leisure margin. That is, a positive nominal interest rate introduces a wedge in consumers' work/leisure choice, a wedge that, from the point of view of optimality, shouldn't be there: work/leisure choices "should" be made (again, from the point of view of optimality) according to *only* the real wage.[146] The benevolent central bank does not want there to be such a wedge in consumers' decisions, hence it conducts monetary policy in such a way as to make $i = 0$. Once the government chooses to conduct policy in such a way as to make $i = 0$, the steady-state version of the Fisher equation immediately tells us that that way to do so it to make the inflation rate equal to $\beta - 1$; but because $\pi = g$ in steady-state, this means a gradual withdrawal of money from the economy.

The optimal rate of money contraction depends on β. Suppose consumers were completely patient, meaning $\beta = 1$ (i.e., no discounting of future utility at all). Then, the Friedman Rule (which, note, is still optimal) means $g = 0$ -- the central bank should not change the money supply at all. The reason why things depend so critically on β is that every unit of money held means one less unit of an interest-bearing bond that is held. If the consumer *had* held that bond, he would have waited one period before receiving the principal plus interest on the bond back. Waiting one period entails an "impatience cost" measured by β. So the higher is β the less "bad" from a welfare perspective is *not* removing money and its attendant wedge from the economy.

As we noted, in steady-state the Friedman Rule can either be stated as $i = 0$ or $g = \beta - 1$. This is **not** true when we consider things outside of steady state. When we later turn to dynamically-optimal monetary policy, the correct notion of the Friedman Rule is as $i_t = 0$, which does **not** automatically translate into $g_t = \beta - 1$.

145 More precisely, it is equated to the after-tax real wage. Here we ignore labor taxes, so the prescription from the basic consumption-leisure model is that consumers choose consumption and leisure to set their MRS equal to the real wage.

146 We will study in more depth the content of this statement that work/leisure choices "should" be made only according to the real wage; the issue has to do with social efficiency/Pareto optimality.