10 Imperfect Competition and Real and Nominal Price Rigidity

Chapter 10 Contents

A. Topics and Tools

Romer's Sections 6.5 through 6.8 examine the "new Keynesian" approach to aggregate supply, focusing on the microeconomic basis for price stickiness. The new Keynesian paradigm first emerged in the late 1970s in response to the "neoclassical" models of Robert Lucas and others. This "neoclassical counterrevolution" (to the original "Keynesian revolution" after the 1930s) focused on market-clearing models with well-established microeconomic foundations and rationally formed expectations. These models usually concluded that countercyclical fiscal and monetary policy were at best useless and at worst destabilizing, calling into question the long-established Keynesian prescriptions for macroeconomic policy.

The first models to be labeled as new Keynesian assumed that (for whatever reason) firms and workers set labor contracts that specify the nominal wage in advance for more than one period. Stanley Fischer (1977) and John Taylor (1979) were able to show that one could derive a Keynesian policy result (*i.e.*, that countercyclical monetary policy can improve welfare) while maintaining fairly solid microeconomic foundations and the assumption of rational expectations.

The first wave of new Keynesian analysis was devoted to examining the implications of different assumptions about how wages and prices are set. Since everyone agrees that wages and prices are unlikely to be *perfectly* flexible (for example, no one denies the existence of labor contracts in the union sector), these models were quite popular as alternatives to the neoclassical school. However, as Romer notes in Chapter 6, Part A (pages 245 and 246), wage-contract models predict a strongly countercyclical real wage, which is not consistent with actual observations.

Part B of Chapter 6 lays out the basic microfoundations of the new Keynesian approach to price stickiness; Chapter 7 then applies these models in a macroeconomic context under several different assumptions about the frequency and nature of price changes.

We focus in this coursebook chapter on Sections 6.5 through 6.8, which is one of the most exciting sections of the course, but also one of the most challenging. On center stage are the concepts of ***nominal and real price rigidity***. Both of these concepts have to do with resistance to price changes. As you might expect, nominal rigidity occurs when a firm exhibits reluctance to change its price in nominal (dollar) terms. Real rigidity is a situation when a firm does not want to change its price relative to some other prices, for example when a firm wants to keep its price in line with those of other firms in the market.

We begin in Section 6.5 by introducing a model in which firms are imperfectly competitive. Imperfect competition is necessary for models with price stickiness because perfectly competitive firms must always charge the market equilibrium price, eliminating any possibility for rigidity. Sections 6.6 and 6.7 explore the microeconomic implications of nominal and real rigidities. Romer demonstrates that nominal rigidities of sufficient magnitude to explain price stickiness in the real world are implausible.

Real rigidities alone do not introduce non-neutral effects of monetary shocks; everyone will adjust together to the new market-clearing equilibrium. Only by combining nominal and real rigidities so that the latter amplify the non-neutral effects of the former are new Keynesians able to achieve a realistic depiction of non-neutral monetary effects.

Section 6.8 discusses the possibility that real rigidities can lead ***to coordination failures*** in the macroeconomy. This literature employs game theory to examine interactions between firms. Under certain circumstances, rigidities can lead to multiple equilibria in the macroeconomy. If there are two or more points of equilibrium, it is often possible to demonstrate that one is Pareto-superior to the other(s). In such a situation, the economy can become trapped at an inferior equilibrium where a change in policy could potentially push the economy to the best equilibrium.

B. What's New and Keynesian about "New Keynesian" Economics

The label "new Keynesian economics" has been given to the broad class of models that we are studying in Romer's Chapter 6. Two natural questions to ask before we study their details are "In what ways are these models Keynesian?" and "In what ways do these models differ from 'old Keynesian' models?"

The first new Keynesian models evolved as a reaction to the "neoclassical revolution" a set of models represented by Lucas (1972), Sargent and Wallace (1975), and Barro (1976). These neoclassical models were based on aggregation of a standard microeconomic general-equilibrium model in which all market clear. The only exception to perfect competition in these models was that agents did not have complete information. (Lucas's model is discussed in Romer's Section 6.9 and Coursebook Chapter 11.) Prior to 1970, the intellectual skirmishes between the post-war Keynesian and monetarists had been fought on the turf of aggregate models such as *IS/LM.* Lucas and his followers transformed the rules of play to require rigorous specification of the maximization decisions made by individual agents.

Natural extension of the basic microeconomic theories of utility and profit maximization under perfect competition leads to a model with a strongly classical flavor, so this change of venue posed a particular challenge to Keynesians. In particular, traditional Keynesian economics was based on simple assumptions that prices or wages are sticky and that markets do not clear. The neoclassical revolution put the onus on Keynesians to prove that such stickiness and non-market-clearing behavior can be justified in a world where individual agents maximize utility and profit.

The new Keynesian school arose in response to that challenge. Thus, new Keynesians are Keynesian in that they believe that wage and price stickiness are important features of the economy and that this implies a positive role for countercyclical policy. They differ from "old Keynesians" in that rather than simply asserting that prices or wages are sticky, they seek a microeconomic framework in which the maximizing decisions of rational agents lead to stickiness.

To see why this challenge is important, consider the simple microeconomic market whose demand and supply curves are shown in Figure 1. If the price of the good adjusts freely to clear the market, it will be at p^* with quantity demanded equal to quantity supplied at q^*. Suppose instead that the price were to be at the lower level p_1. Since there is an excess demand at this lower price, quantity exchanged is now

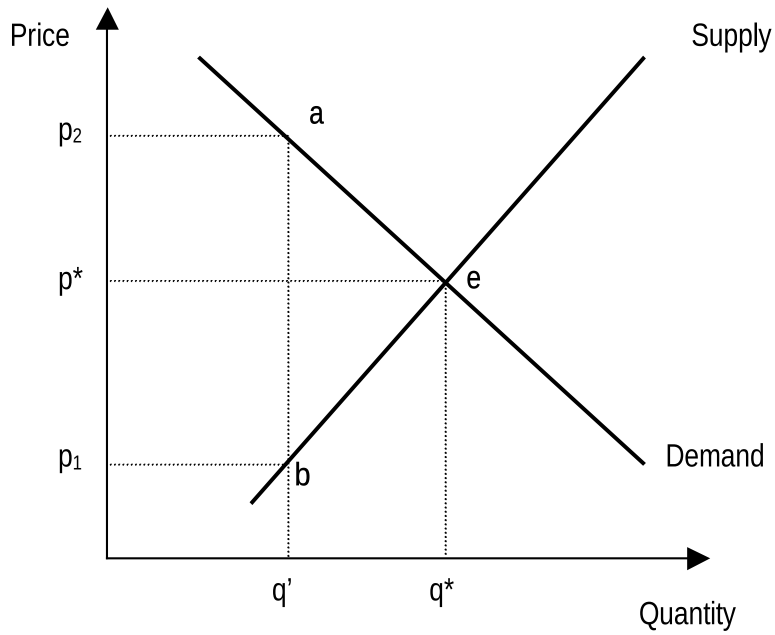

Figure 1. Simple supply/demand model

only q_1. Comparing the welfare of market participants in these two cases using the sum of consumer and producer surplus, it is clear that there are net losses to having the price level at p_1. These deadweight losses are measured by the area of the triangle *abe*. (The same losses would be incurred if the price were to be at p_2 higher than p^*.)

Dissatisfied potential buyers have an incentive to bid the price up from p_1. Similarly, if the price were above equilibrium, sellers would have incentive to undersell their competitors who are selling at price p_2. The challenge to new Keynesian economists is thus to explain why markets would adopt price-setting institutions that fail to capture the extra welfare that would be gained it the price was at the market-clearing level. Robert Barro has compared the absence of price adjustment by agents to "leaving $50 bills lying on the sidewalk."[1]

[1] An interesting irony is that Barro's first major work was a series of papers and a book he co-authored in the early 1970s. (See Barro and Grossman (1971).) The model developed in this early work was a comprehensive microeconomic general-disequilibrium model of how maximizing firms and households would behave *in a world with perfectly rigid wages and prices*. By the late 1970s, when he had jumped on the neoclassical bandwagon, Barro regarded this work as not useful.

Institutions of price setting

But how, exactly, do those $50 bills get picked up? Who is in a position to pick them up? Is it possible that it might cost them more than $50 to pick them up? A closer look at these questions provides a response to Barro's puzzle.

We are all familiar with the idea that prices will adjust to the "equilibrium," market-clearing level. We teach this in every economics textbook, yet we spend very little time talking about why or how this happens. In fact, we rarely consider the question of *who* sets prices in actual market.

In the neoclassical perfect-competition model, price setting is assumed to be done by a hypothetical "Walrasian auctioneer." The auctioneer solicits demand and supply information *a priori* from all market participants. She then calculates the equilibrium price and calls it out to all traders. All exchanges then occur at the market-clearing equilibrium price. The Walrasian auctioneer does not charge for her services. She incurs no cost in finding out how much every prospective buyer or seller is willing to trade at each possible price. Moreover, her services are performed immediately and continuously. Any time that a change in demand or supply occurs, she knows immediately and she immediately informs all traders. Truly a superhuman creature this auctioneer must be!

But of course there is no Walrasian auctioneer in actual markets for goods and services. Price setting is performed by mere mortals who could only at great cost (if at all) acquire the knowledge we attribute to the auctioneer. The identity of the price setter varies from market to market. In some markets, brokers or dealers facilitate trade and set prices. Although they do many of the same things as a Walrasian auctioneers, they incur costs in doing them and therefore charge fees for their services. These fees are often implicitly collected as the spread between the prices at which the dealer sells and buys: the "bid-ask spread" or the dealer's markup.

Real-world auction markets may also approximate Walrasian outcomes. In the double-oral auction experiment of Econ 201, you saw how prices tended to converge to the market-clearing level without any explicit coordination from an auctioneer. Both buyers and sellers can call out bids and offers in a double-oral auction. Convergence occurs naturally as a result of having all buyers and sellers interacting in a single location where they can easily compare the prices of prospective partners and where information about transaction prices is readily available.

Some important markets function by auction or with extensive dealer or broker coordination. Financial markets, real-estate markets, and markets for important raw materials such as metals and agricultural products often have a high degree of coordination. They may approximate Walrasian outcomes if the coordination is effective and inexpensive.

However, two of the most important macroeconomic markets seem very different. Wages in labor markets are set by individual or collective negotiations, which cost the participants time and raise the possibility of acrimonious work stoppages if agreement cannot be reached quickly. Thus, most wages are set for a period of a year or more in an explicit or implicit contract. The contract wages may be fixed in nominal terms or they may incorporate an "indexing rule" by which changes in prices are automatically reflected (at least partially) in the nominal wage.

The market for retail goods and services also differs greatly from the Walrasian model. Here, individual sellers almost always set prices. Most sellers face some competition in the market for their products, but not textbook "perfect" competition. In all cases, there are real costs associated with figuring out the sales implications of alternative pricing strategies. Firms may also incur adjustment costs when they change prices, where price adjustment can be thought of either in nominal terms or in relation to other producers.

As we shall see, the main way in which new Keynesians rationalize imperfect price adjustment is to consider the process of price setting and the costs of price adjustment. While there are certainly opportunity costs to market participants of trading at non-market-clearing prices (the $50 bills *are* on the sidewalk), there may be other benefits associated with keeping prices at a non-market-clearing level (there are also costs of picking the bills up).

One of the key tenets of new Keynesian macroeconomics is that if a firm's price is "close to" the market-clearing price, then the gains to the firm in moving to the exact market-clearing price may be quite small. Because of this, even relatively small adjustment costs such as the printing of new menus or catalogs may be sufficient to delay price adjustment by a profit-maximizing firm.

Market structure and price adjustment

Remember that sluggish price adjustment cannot occur in a world of perfect competition. If the products of all firms are perfect substitutes for one another, as is assumed in the perfectly competitive model, then any firm that sets its price even a penny above its competitors will see its sales drop to zero. Clearly, in that situation no firm can afford to have any (downward) price rigidity at all. Because of this, new Keynesian models have incorporated ***imperfect competition*** as the fundamental micro-level market structure.

The market model that is usually chosen in new Keynesian models is a simple version of ***monopolistic competition***. You may recall from your introductory economics course that monopolistic competition exists when there are many sellers of a good and entry is free but when the good itself is differentiated. Free entry assures that economic profits are zero in the long run. (The macroeconomic version does not incorporate entry.) Product differentiation implies that one firm's product is not a perfect substitute for the goods produced by others, so firms face downward-sloping demand curves rather than being price-takers.

In many ways, imperfect competition is much more difficult to handle than perfect competition. Because firms are not price-takers, we cannot rely on the usual apparatus of supply and demand as a framework for analysis. Instead we must look specifically at the price-setting behavior of each firm, subject to the constraints of demand. This is the agenda that new Keynesian macroeconomists have undertaken.

C. Romer's Model of Imperfect Competition

Part B describes some basic microeconomic models in which prices fail to adjust completely. Instead, they have local monopolies over the production of their individual goods. We then consider the conditions under which these monopolies will choose to adjust their prices fully in response to a shock vs. keeping their prices unchanged. We shall see that each firm's decision about whether to adjust prices may depend on *other* firms' decisions, which creates a strategic interaction among firms and increases the likelihood of inefficient non-adjustment of prices. The possibility that producers may individually fail to adjust prices when they would be collectively better off if they did adjust is called ***coordination failure*** and is a central concept of new Keynesian macroeconomics.

As in monopolistic competition, we think of the economy as being composed of a large number of small firms, each of which has a monopoly on its own variant of a product. These variants are close, but not perfect, substitutes for one another. This means that each producer has a downward-sloping, elastic demand curve and maximizes profit like a monopoly: by setting marginal revenue equal to marginal cost.

Household utility maximization

Household utility depends on consumption and (negatively) on labor effort:

$$U = C - \frac{1}{\gamma} L^{\gamma} .$$

The marginal utility of labor is

$$\frac{\partial U}{\partial L} = -L^{\gamma - 1} < 0 .$$

The assumption that $\gamma > 1$ assures that the marginal disutility of labor increases as the household works more.

The concept of C in this equation is something new to us. Unlike the models that we have previously studied, this model has many consumption goods being produced by many firms. They are not perfect substitutes so we cannot just add them up with a simple sum. Instead, we assume that utility depends on an *index C* of the household's consumption of the various goods in the model.

The index is constructed using a single parameter that determines how closely households are willing to substitute the products of various firms. (For simplicity, the products of all firms are assumed to be equally close substitutes for the products of all other firms.) The most convenient way to express this substitution parameter is as η (eta), the absolute value of the elasticity of demand for the product of each individual firm. Under perfect competition, each firm's demand is infinitely elastic ($\eta \to \infty$). If $\eta \leq 1$, then the firm faces inelastic or unit-elastic demand, and can increase profits infinitely by raising its price without bound. Thus, for the model to describe monopolistic competition, we restrict the range to $1 < \eta < \infty$.

If there are N firms in the economy, then we can construct a "Dixit-Stiglitz" index of average consumption across the N goods as

$$C = \left(\sum_{i=1}^{N} C_i^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}.$$

Notice that the exponent outside the summation is the reciprocal of the exponent of each consumption quantity inside the summation. This is a critical property of the Dixit-Stiglitz index because it leads to a property analogous to constant returns to scale. Notice that if each C_i inside the summation is doubled, it increases each $C_i^{\frac{\eta-1}{\eta}}$ term in the sum by a factor of $2^{\frac{\eta-1}{\eta}}$. This increases the sum by that same factor, which means that the index C increases by a factor of $2^{\frac{\eta-1}{\eta}\frac{\eta}{\eta-1}} = 2$.

We want the number of firms to be large so that, as under perfect competition, each firm is an infinitesimally small part of the overall market. The limiting case as the number of firms/products $N \to \infty$ is Romer's equation (6.39):

$$C = \left(\int_{i=0}^{1} C_i^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}}$$

We refer to this situation as having a "continuum of firms and products." Firms and products are indexed by a real number varying over the unit interval [0, 1], rather than by an integer running from 1 to N. While we never really have infinitely many firms and products in an economy, it is a convenient approximation to the situation where the size of each firm/product relative to the whole economy is vanishingly small.

Romer proceeds to analyze households' utility-maximizing behavior on pages 270 and 271. The equations here seem complex, but the underlying principles are straightforward. He defines $S \equiv \int_{i=0}^{1} P_i C_i di$ to be total nominal spending by household i. In the equilibrium of the model, this will be determined by the level of aggregate demand. At the individual level it will depend on the amount that the household works. Initially we take S as given and use the utility function do determine the household's consumption of the individual products C_i. Because the utility function is additive in functions of C and L, we can analyze the decisions to consume and to work separately. (This is possible because the second cross-partial derivatives of the utility function are zero: the marginal utility of consumption does not depend on labor and the marginal (dis)utility of labor does not depend on consumption.)

The use of the Lagrangian in (6.42) will be familiar to those who have taken Econ 313. For others, the method of Lagrange multipliers is a way of maximizing a function subject to a constraint. In this case, the household wants to maximize its utility from consumption subject to the constraint that given total spending must equal the integral (summation) of nominal spending on each good. The Lagrangian expression in (6.42) is formed by adding to the utility function a "Lagrange multiplier" λ multiplied by the constraint, which is in the form of an expression that equals zero when the constraint is satisfied.

To find the first-order conditions for utility maximization subject to the constraint, we take the (partial) derivative of $\mathcal{L}$ with respect to each of the choice variables (in this case, the C_i values) and setting it equal to zero. This may seem a bit daunting here because there are infinitely many C_i variables, but each one enters the equation in exactly the same way as all of the others, so we can examine one "representative" good i and then generalize the analysis to apply to all the others as well. Differentiating (6.42) with respect to one specific good C_i and setting the derivative equal to zero yields (6.43). Don't be confused by the sudden appearance of j as an index in the integral on the left-hand side of (6.43). We are using i to refer to the specific good with respect to which we are differentiating, so we need another index in the integral/summation over all goods.

Proceeding from (6.43), how does Romer know that the C_i function must have the form of (6.44)? If we solve (6.43) for C_i we get

$$C_i^{-\frac{1}{\eta}} = \frac{\lambda}{\left(\int_{j=0}^{1} C_j^{\frac{\eta-1}{\eta}} dj\right)^{\frac{1}{\eta-1}}} P_i$$

$$C_i = \frac{\lambda^{-\eta}}{\left(\int_{j=0}^{1} C_j^{\frac{\eta-1}{\eta}} dj\right)^{-\frac{\eta}{\eta-1}}} P_i^{-\eta} = C\lambda^{-\eta} P_i^{-\eta}.$$

Thus, the constant A in (6.44) is $C\lambda^{-\eta}$. We don't know the value of λ so we still have work to do, but we have verified that the solution is of the form of (6.44) as Romer asserts. We stated above that $-\eta$ was the price elasticity of the household's demand for good i. Equation (6.44) shows that this is the case. Finding the elasticity,

$$\frac{\partial C_i}{\partial P_i}\frac{P_i}{C_i} = -\eta A P_i^{-\eta-1}\frac{P_i}{C_i} = -\eta\frac{AP_i^{-\eta}}{AP_i^{-\eta}} = -\eta.$$

To derive (6.45), we perform the substitution that Romer suggests:

$$\int_{i=0}^{1} P_i A P_i^{-\eta} di = S$$

$$A\int_{i=0}^{1} P_i^{1-\eta} di = S$$

$$A = \frac{S}{\int_{i=0}^{1} P_i^{1-\eta} di}.$$

This leads to (6.46) with Romer showing all the intermediate steps. The definition of the price index P in (6.47) follows naturally from its position in the denominator of (6.46). Let's think a bit about the intuition here. We started by defining nominal expenditures to be $S \equiv \int_{i=0}^{1} P_i C_i di$. We have defined the index of total consumption C (averaged across goods) to be the expression in (6.39). It makes sense to define the price index P to be S/C, so that total spending $S = P \times C$. Equation (6.46) shows that the price index in (6.47) is precisely the price index for which this is true.

How does the price index P compare with common price indexes such as the CPI? It basically differs in only one way: it averages a power $1-\eta$ of the prices rather than the prices themselves, then returns it to the original units by taking the average to the $1/(1-\eta)$ power. If $\eta = 0$, then both exponents equal one and the price index in (6.47) becomes an (equally weighted) average of the prices across all the goods.

This is what the CPI would do if all goods in the market basket were purchased in equal quantities.[2]

Equation (6.48) completes the consumption choice (for now) by substituting. Consumption of each good i is equal to the average-consumption index C times the relative price of good i raised to the $-\eta$ power. Equation (6.48) can be rewritten as

$$\frac{C_i}{C} = \left(\frac{P_i}{P}\right)^{-\eta}.$$

In this form, it shows that the ratio of consumption of good i to the average-consumption index depends negatively on the relative price of i.

The short section following (6.48) on page 271 describes the household's labor-supply choice. The only aspect of this that is likely to be confusing is the nature of R, the household's "profit income." We need this here because firms in this model are monopolies earning positive economic profits. These profits must accrue to the households that own the firms, hence the R term. Equation (6.51) shows that the elasticity of labor supply is $1/(\gamma - 1) > 0$. This elasticity turns out to be very important in determining the properties of the model.

Firms' behavior

Firm behavior in our growth models involved setting the marginal product of labor equal to the real wage, so that principle should be familiar. Because firms here are monopolies, we must distinguish the marginal *revenue* product of labor from the marginal product. Recall from Econ 201 that MRP = MP × MR. For competitive firms, MR = P, so setting MRP = W is equivalent to setting MP = W/P. In the growth models, the marginal product of labor decreased as more labor is hired (given the fixed amount of capital), so the labor-demand curve sloped downward.

In this model, we choose the simplest possible production framework: there is no capital and no diminishing marginal returns to labor—output is proportional to labor input. Moreover, since the units in which we measure output and labor input are arbitrary, we can choose them so that one unit of labor input creates one unit of output. Thus, the production function for good i is $Y_i = L_i$. We shall see that this simple model delivers a well-defined production decision and a unique labor-market equilibrium.

---

[2] All goods in this model are "symmetric," though they are not "identical" because they are not perfect substitutes). All have the same elasticity parameters and so all will end up being consumed in equal quantities.

Equation (6.52) defines the real profits of the firm that has a monopoly on the production of product i. The only source of demand for good i is consumption, so in equilibrium $Y_i = C_i$ and, in the aggregate, $Y = C$, with the aggregate consumption and output both being the Dixit-Stiglitz indexes given by equation (6.39) for consumption and

$$Y = \left(\int_{i=0}^{1} Y_i^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}. \tag{1}$$

From the consumption-demand equation for good i we know that the monopoly firm's sales of i must be

$$Y_i = \left(\frac{P_i}{P} \right)^{-\eta} Y,$$

where Y is the aggregate output index given by equation (1).

Substituting this into (6.52) and noting that $Y = L$ in the aggregate gives us equation (6.53) for real profits. Having substituted for Y_i and L_i, the firm's choice variable is its relative price (P_i/P). Setting the derivative of real profit with respect to the relative price equal to zero gives equation (6.54), which solves to the optimal pricing equation (6.55).

Equation (6.55) shows that the firm's optimal relative price is a markup of the ratio $\frac{\eta}{\eta - 1}$ over the firm's marginal cost, which equals the real wage rate because each unit of output requires one unit of labor to produce. This is a standard result from monopoly theory. Note that the markup ratio is greater than one and decreases in product-demand elasticity η. In a perfectly competitive market, $\eta \to \infty$ and the markup ratio approaches one: perfectly competitive firms earn no profit and sell at a price equal to marginal cost. As $\eta \to 1$ from above, the markup ratio grows without bound. Firms with price elasticity barely greater than one have enormous market power and can set prices many times higher than marginal cost. The parameter η is another key parameter that will determine the characteristics of our equilibrium.

Equilibrium

The previous sections have described the microeconomics of the model, first from the household's perspective and then from the firm's. What remains is to consider the aggregate equilibrium conditions: the macroeconomics of the model. In particular, we need to specify the determination of aggregate demand in the model. This is the missing puzzle piece that we left open when we took S to be given in the derivation of the household's equilibrium consumption behavior. Because consumption is equal to total output, S is nominal GDP: $S = PC = PY$.

When we first introduced the idea of aggregate demand at the beginning of the course (and in Coursebook Chapter 2), we discussed the quantity theory of money, in which nominal expenditures are represented as $MV = PY$, where M is the money supply and V is the income velocity at which money circulates. This theory leads to a downward-sloping AD curve similar to that of the *IS/LM* model, but is much easier to work with. We will use this theory to represent aggregate demand in the models in the second half of Chapter 6 and Chapter 7.

We simplify the analysis even more by setting letting $S = M$, where M can be interpreted narrowly as the money supply (with velocity fixed at one) or more broadly as a measure of aggregate demand. I prefer the latter interpretation, with changes in M reflecting not just shifts in monetary policy but also changes in fiscal policy or other shocks to desired expenditures. Given this simple AD specification, real aggregate demand is given by

$$Y = \frac{M}{P}.$$

From the production function, $Y = L$, so we can use the labor-supply equation for the representative household (6.38) to get Romer's (6.56) and then the markup equation to get (6.57).

Using (6.57), we apply the final, crucially important equilibrium condition to complete the solution. All firms and goods are symmetric. They all face the same demand elasticity and the same real wage, so they all set the same markup ratio and the same relative price. Given that every P_i is the same, the aggregate average price level from equation (6.47) must be equal to the level of each P_i, as Romer notes in the text below (6.47). That means that the final equilibrium condition is $P_i = P$, or

$$\left(\frac{P_i}{P}\right) = 1.$$

Applying this condition in (6.57) and solving for Y gives the equilibrium level of output in the model as equation (6.59):

$$Y = \left(\frac{\eta - 1}{\eta}\right)^{\frac{1}{\gamma - 1}}.$$

The price level is then

$$P=\frac{M}{Y}=\frac{M}{\left(\frac{\eta-1}{\eta}\right)^{\frac{1}{\gamma-1}}}.$$

Properties of the model

We have now solved a model in which producers are imperfectly competitive, but markets clear with perfectly flexible wages and prices. What are the properties of this model?

First, as Romer notes on page 271, the equilibrium, "natural" level of output in this model is less than the optimal amount. In our model, households interact in a market economy in which firms have monopoly power. As an alternative setup, we can think of the each household operating in total isolation. In this case, each household would choose an amount $\bar{L}$ to work, producing $\bar{Y}=\bar{L}$ units of output which it then consumes. With consumption and labor effort equal to $\bar{L}$, the household's utility is

$$U=C-\frac{1}{\gamma}L^{\gamma}=\bar{L}-\frac{1}{\gamma}\bar{L}^{\gamma}.$$

Choosing $\bar{L}$ to maximize the household's utility,

$$\frac{\partial U}{\partial \bar{L}}=1-\bar{L}^{\gamma-1}=0,$$

which implies $\bar{L}=(1)^{1/(\gamma-1)}=1$. Therefore the optimal level of output, employment, and consumption in this model is $Y = C = L = 1$.

From equation (6.59), the level of output under imperfect competition is less than one, therefore less than the optimal amount. This occurs due to the monopoly power of firms. Firms with monopoly power have marginal revenue less than price, so they set their level of output where price exceeds marginal cost—this is the markup equation (6.55)—and they pay labor its marginal revenue product, not price times marginal product as under perfect competition. Because these monopolies overcharge for their output, they sell less than the ideal amount, causing equilibrium GDP in this economy to fall short of the optimal level.

Notice from equations (6.59) and (6.55) that when $\eta \to \infty$, $\frac{\eta - 1}{\eta} \to 1$, which means that the markup ratio is one (firms price at marginal cost), and the equilibrium level of real GDP equals one, which means that the optimal level of output is achieved. Infinite demand elasticity means that firms are perfectly competitive, as markets become more competitive the monopoly distortion goes away.

The fact that natural output is less than optimal means that booms in the economy (where output is above its natural level) may bring the economy closer to its optimal output and recessions may move it farther away. This corresponds to the usual popular opinion that booms are good and recessions bad.

A second conclusion is that "aggregate-demand externalities" are present in this model. Lowering her price not only implies that the individual agent sells more, but that everyone else sells more (at a given relative price) through the effect that her increase in income has on aggregate demand. That means that by lowering price she would make everyone else's utility higher. We shall see later on that these externalities (or "spillovers") represent failures of the price system to coordinate agents' decisions in an optimal way.

A third implication of the model is that imperfect competition alone does not lead to non-neutrality of money and a role for aggregate demand in the determination of output. The aggregate supply curve is still vertical: the aggregate-demand variable M does not appear in equation (6.59), which determines Y. From equation (6.60), it is clear that increases in aggregate demand (the money supply if we so identify M) result in proportional changes in the price level. Thus money is neutral in this model; in order for money to have real effects, we will need to introduce additional imperfections into the model, usually in the form of wage or price stickiness.

D. Nominal and Real Rigidities

As Romer points out, most of the evidence suggests that the magnitude of nominal rigidities is quite small.[3] They are often called ***menu costs***, after the idea that a restaurant incurs a small but significant cost when it changes prices because it must print new menus. This example is illustrative in a number of ways. First, although the cost of new menus is not trivial, it is small relative to the overall costs of running a restaurant. It is easy to imagine a restaurant deciding not to change its prices by a penny or two because doing so would require printing new menus. But if the optimal

[3] However, see Levy et al. (1997) for some contradictory evidence.

price deviates from the price printed in the current menus by several dollars, it seems implausible that the restaurant would not go ahead and make the change.

A second consideration is that literal menu costs are endogenous to the firm. A restaurant that expects to change its menu prices frequently can minimize its menu costs by using paper inserts that are easily changed or even forsaking a printed menu altogether in favor of listing the day's (or even hour's) prices on a blackboard. This kind of response leads to an effect analogous to the effect of high and low inflation variance on the slope of the aggregate-supply curve in the Lucas model. If it is desirable or necessary to adjust prices frequently, firms will find ways to reduce menu costs and prices will be more flexible.

A third characteristic of menu costs is that they are a fixed cost. If you are printing up new menus, it does not cost any more to change prices by a lot than to make a small price adjustment. Although some early models incorporated increasing rather than constant adjustment costs (such as making the cost a quadratic function of the amount by which prices are changed), the emphasis in the new Keynesian models has been almost entirely on fixed menu costs.[4]

Because menu costs and other nominal rigidities are most likely small, new Keynesian models have explored "amplification mechanisms" that can cause small nominal rigidities at the micro level to lead to large aggregate price stickiness and therefore significant non-neutralities.

Mankiw's menu-cost model

The origin of the menu-cost model was a paper by N. Gregory Mankiw (1985). In this paper Mankiw argued that small menu costs might lead to considerable stickiness because in the neighborhood of the profit-maximizing price the loss in profit from being away from the optimal price is likely to be small. Mankiw's partial-equilibrium analysis is shown either by Romer's Figure 6.10 or Figure 6.11. In Figure 6.11, the flatness of the profit function near its peak means that a firm's price can be some (horizontal) distance from the optimal price with profit falling (vertically) only a little. Since the gains from changing price may be small, menu costs may not need to be large in order to prevent complete price adjustment.

The following intuitive scenario may illustrate the model: Suppose there is a decline in the money supply that reduces aggregate demand. Each firm will see this as a decline in its demand curve. The firm can lower its price and maintain output or it can reduce output keeping its price fixed.[5] It will choose the strategy that leads to the higher profit.

---

[4] See Rotemberg (1982) for an early model in which price-adjustment costs are quadratic.

[5] A firm that faces fixed menu costs will never choose partial adjustment in response to a one-time change in aggregate demand. It will either adjust fully to the optimal price or keep its

However, the magnitude of the decline in any firm's demand curve depends on what its rival firms do. If its rivals lower their prices, then the original firm will see a much larger decline in demand than if the rivals keep prices fixed. In moving from partial to general equilibrium, we must decide what assumption to make about the adjustment behavior of other firms. Romer looks at the incentive of one firm to adjust its price given that its rivals do not adjust. If the single firm decides not to adjust given that its rivals do not adjust, then the situation of full non-adjustment (fixed prices) is a ***Nash equilibrium***. In game theory, Nash equilibrium is a situation in which each agent is behaving optimally given the behavior of the other agent(s). This case qualifies as a Nash equilibrium because (in a two-firm example) firm B chooses not to adjust price given that firm A does not adjust; and firm A chooses not to adjust if firm B does not.[6]

The behavior of wages is a second important consideration in determining whether the firm chooses to adjust its price. In the model that Romer considers in Section 6.6, wages are perfectly flexible and the labor market clears. As he points out, a rather large change in wages is likely to result from a monetary shock if labor supply is inelastic. Referring to Figure 6.10, a change in wages will lower the marginal-cost curve and increase the costs of non-adjustment. Thus Romer's assumption of no wage rigidity and inelastic labor supply makes the likelihood of price rigidity much lower.

Interaction of nominal and real rigidities

As Romer shows in the numerical example on pages 276 through 278, menu costs alone are not sufficient to explain price rigidity in response to sizable aggregate-demand shocks if there are no other rigidities in the system. In terms of Romer's discussion of Figure 6.11, the high elasticity of the desired relative price with respect to aggregate demand means that the horizontal change in the desired price (C and D in Figure 6.11) will be very large, even for relatively small changes in aggregate demand.

It may be helpful to think intuitively about Romer's "quantitative example." Aggregate demand declines, pushing the demand curve of each firm downward. Given that no other firm has yet reduced its price, each firm is faced with a choice: (1) cut price in nominal and relative terms so that you can sell at your optimal production level or (2) keep its nominal and relative price fixed and lower production to the smaller sales level supported by the reduced demand at the original price.

---

price fixed. As we will see in the Caplin-Spulber model, the expectation of *steady* increases in aggregate demand can lead to *over-adjustment*.

[6] There can be multiple Nash equilibria. In this context, full adjustment by both firms would also be an equilibrium if A would decide to adjust its price if B did, and vice versa.

Non-neutrality results only if firms choose option 2. If one firm chooses option 1, then *all* firms will (since they are identical) and the aggregate price will adjust in proportion to the monetary contraction. In this case, aggregate-demand shocks do not affect output and money is neutral.

So how costly is it to choose option 2? As Romer points out, significant real rigidity makes option 2 less costly. If firms' products are close substitutes, so that their demand is quite elastic, then each firm will want its price to stay close to those of its competitors. That makes option 1 unattractive *if* other firms are *not* expected to change prices. Romer notes that strong real-price rigidity reduces the optimal amount of price adjustment (*CD* in Figure 6.11) if the firm *does* adjust, and thus reduces the benefits to adjustment.

However, the competitive labor market in Romer's quantitative example makes it almost impossible for firms not to adjust prices (and thus absorb the fluctuation by changing output) because labor supply tends to be inelastic. A labor-supply elasticity of 0.1, which is roughly consistent with empirical evidence, implies that a 3% decline in output. This corresponds to the 3% reduction in aggregate demand if no firm changed its price, requiring a 30% decrease in the real wage for people to still be on their labor-supply curves. Such a drastic wage cut pushes the MC curve in Figure 6.11 down strongly and greatly increases the size of the triangle that measures the lost profit from non-adjustment.

Is this reasonable? Probably not. As we will discuss in the section on unemployment, there are good reasons for abandoning the Walrasian paradigm as a model of the labor market. If wages are sticky enough that they do not fall 30%, then the forces encouraging the firm not to reduce production are muted and non-adjustment is more likely. Romer illustrates this in his "second quantitative example" starting on page 284. This is an example of a common conclusion: small rigidities in different parts of the model can build on each other and lead to much more substantial rigidity in the aggregate model.

Coordination failures

In Walrasian models, the auctioneer acts to coordinate the decisions of individual producers and consumers by assuring that prices send correct scarcity signals. However, once we banish the Walrasian auctioneer and rely on less-than-omniscient firms to set prices, we must worry about whether the price-setters will accomplish coordination efficiently.

Perhaps the idea of coordination failures can be most easily introduced by a now-familiar example. If real price rigidities are sufficiently important that no firm is willing to change its price unless the other firms change theirs, who is going to be the first to change price? It may be optimal for all firms to adjust if there is an effective coordination mechanism: the increase in profits from *all* firms being at the new optimal price may exceed collective adjustment costs. However, if each firm's increase in profit is smaller than the menu cost *given that the others do not adjust*, then no firm will want to "break the ice" and be the first to change its price.

In such a situation, the economy may have multiple alternative equilibria. In this case, there is both an "adjustment equilibrium" and a "non-adjustment equilibrium." Either is an equilibrium: if everyone adjusts then there is no incentive for any individual to change her decision, but if everyone does not adjust there is also no incentive for the individual to change. The adjustment equilibrium may have higher profit (and general welfare) than the non-adjustment equilibrium, but the price system without a Walrasian auctioneer may not provide appropriate signals and incentives to move the economy from the inefficient equilibrium to the efficient one. This is what we mean by a coordination failure.

Cooper and John (1988) set up a useful framework for the general analysis of models with coordination failures. They introduce several very important concepts in defining the conditions under which coordination failures can occur. ***Spillovers*** are nothing more or less than externalities. Positive spillovers occur when one individual's action makes someone else's utility higher; negative spillovers occur when an action lowers someone else's utility. Spillovers alone can lead to inefficiency of equilibrium (as you know from your study of externalities in Econ 201), but are not sufficient to cause multiple equilibria.

Strategic complementarities lie at the heart of the possibility of multiple equilibria. They occur when one person taking an action makes that action more attractive for others. The application to price rigidity is obvious: if real rigidities are important, then firm A adjusting its price makes it more beneficial for firm B to adjust its price. A far less important case for our purposes is ***strategic substitutability***, which occurs when one person acting makes it less desirable for others to act in the same way. For example, if someone else closes the window, you do not have to do so. Strategic substitutability is common in traditional game-theoretic models of firm interaction, but does not lead to coordination failures of the kind discussed here. An example of strategic substitutability occurs in the Cournot model of oligopoly, where if one firm increases production, then the demand for the product of other firms falls so they will probably have incentive to reduce output.

Strategic interaction among symmetric agents is usually analyzed with game theory. Each firm has a ***reaction function*** that describes the optimal value of its decision variable as a function of the value that others choose. Because we usually assume that all agents are identical, it makes sense to look for a "symmetric" equilibrium in which all make the same decision. A ***symmetric Nash equilibrium*** occurs where the representative firm's reaction function crosses the 45-degree line, since that is where it makes the same decision as everyone else.

When strategic complementarity is present, each firm's reaction function is upward-sloping. If the reaction function has portions that have a slope greater than one and others where the slope is less than one, then it is possible that it can intersect the 45-degree line more than once, leading to multiple symmetric Nash equilibria. This is what must happen to support the example of adjustment and non-adjustment equilibria considered above.

E. Works Cited in Text

Barro, Robert J. 1976. Rational Expectations and the Role of Monetary Policy. *Journal of Monetary Economics* 2 (1):1-32.

Barro, Robert J., and Herschel I. Grossman. 1971. A General Disequilibrium Model of Income and Employment. *American Economic Review* 61 (1):82-93.

Cooper, Russell, and Andrew John. 1988. Coordinating Coordination Failures in Keynesian Models. *Quarterly Journal of Economics* 103 (3):441-63.

Fischer, Stanley. 1977. Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule. *Journal of Political Economy* 85 (1):191-205.

Levy, Daniel, Mark Bergen, Shantanu Dutta, and R. Venable. 1997. The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains. *Quarterly Journal of Economics* 112 (3):791-825.

Lucas, Robert E., Jr. 1972. Expectations and the Neutrality of Money. *Journal of Economic Theory* 4 (2):103-124.

Mankiw, N. Gregory. 1985. Small Menu Costs and Large Business Cycles: A Macroeconomic Model. *Quarterly Journal of Economics* 100 (1):529-38.

Rotemberg, Julio J. 1982. Monopolistic Price Adjustment and Aggregate Output. *Review of Economic Studies* 49 (4):517-531.

Sargent, Thomas J., and Neil Wallace. 1975. 'Rational' Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule. *Journal of Political Economy* 83 (2):241-254.

Taylor, John B. 1979. Staggered Wage Setting in a Macro Model. *American Economic Review* 69 (2):108-13.