

11 IMPERFECT COMPETITION AND REAL AND NOMINAL PRICE RIGIDITY

Chapter 11 Contents

A. Topics and Tools	1
B. What's New and Keynesian about "New Keynesian" Economics.....	3
<i>Institutions of price setting.....</i>	<i>4</i>
<i>Market structure and price adjustment.....</i>	<i>6</i>
C. Understanding Romer's Chapter 6, Part B	6
<i>Equilibrium with imperfect competition.....</i>	<i>7</i>
<i>Nominal rigidities</i>	<i>10</i>
<i>Mankiw's menu-cost model.....</i>	<i>11</i>
<i>Nominal and real rigidities.....</i>	<i>12</i>
<i>Coordination failures.....</i>	<i>13</i>
D. Works Cited in Text	15

A. Topics and Tools

The second and third sections of Romer's Chapter 6 examine the "new Keynesian" response to the Lucas model. The microeconomic basis of Lucas's model was reasonable and quite appealing, but Keynesians found the assumption of perfect market-clearing and the conclusions that resulted from it unreasonable. Thus macroeconomists who maintained Keynesian beliefs had strong incentives to develop a macroeconomic model that preserved desirable analytical features of Lucas's model but incorporated wage or price stickiness and led to opposite conclusions about macroeconomic policy.

The first models to be labeled as new Keynesian began from the assumption 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 key

assumption for Lucas's policy-ineffectiveness result, it turned out, was the assumption of continuous market clearing via perfectly flexible wages and prices.

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 Lucas imperfect-information framework. However, as Romer noted in Chapter 5 (pages 242 and 243), these wage-contract models predict a strongly counter-cyclical real wage, which is not consistent with actual observations.

Part B of Chapter 6 lays out the basic microfoundations of the new Keynesian approach; Part C 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 Part B, 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 6.4 by introducing a model in which firms are imperfectly competitive. Imperfect competition is necessary for models with price stickiness because perfectly competitive firms always charge the market equilibrium price, eliminating any possibility for rigidity. Sections 6.5 and 6.6 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 6.7 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 the last two sections of 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?"

As noted in the overview above, the first new Keynesian models evolved as a reaction to the "neoclassical revolution" represented by Lucas (1972), Sargent and Wallace (1975), and Barro (1976). Whereas the earlier 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 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

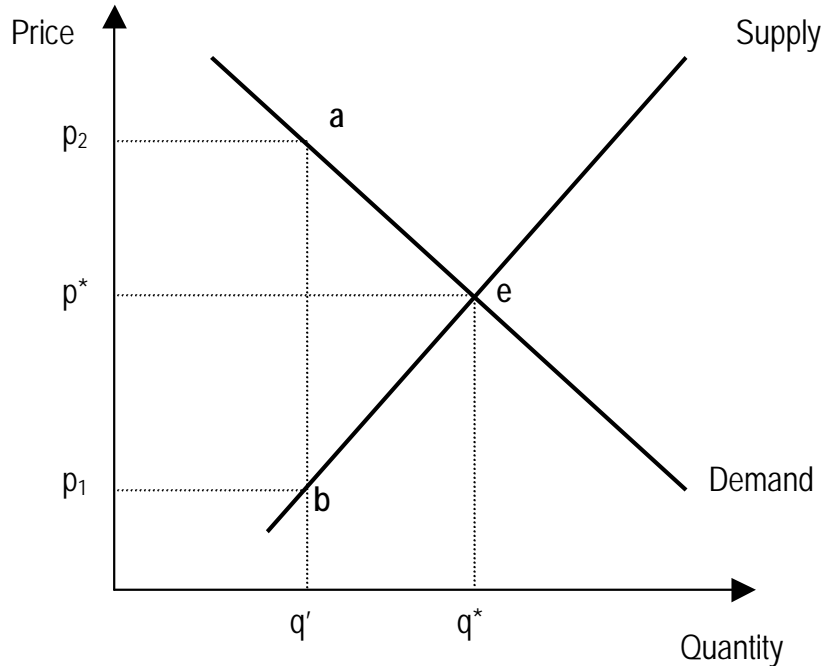


Figure 1. Simple supply/demand model

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 if 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

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.

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 had repudiated this work as not useful.

In the neoclassical perfect-competition model, price setting is 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 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 costs associated with keeping prices at that 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. 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. Understanding Romer's Chapter 6, Part B

Part B describes some basic microeconomic models in which prices fail to adjust completely. It begins by examining a macroeconomic model that is identical to the

perfect-information “island” model developed in Part A except that firms are no longer perfectly competitive. 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 will see 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.

Equilibrium with imperfect competition

Before we consider the analysis of the model in section 6.4, we will review exactly how this model differs from the full-information version of the Lucas model that Romer presents in section 6.1. First of all, there is now a competitive labor market in which households can exchange their labor for money, rather than each household producing output using only its own labor. This assumption appears to have no effect whatsoever on the equilibrium of the model—the same conclusion would result if households produced goods directly. Once Romer solves for the equilibrium equations (6.45) through (6.47), he never returns to examine how real wages and the level of employment behave. Since the model is symmetric across agents, each agent ends up supplying exactly the same amount of labor as she uses in production and it really doesn’t matter if she supplies it to herself (outside of the labor market) or to another firm.

A second difference between the full-information Lucas model and the imperfect-competition model is the absence of a product-specific (or island-specific, industry-specific, etc.) demand shock in the latter. Recall that the Lucas model required the product-specific shock in order to generate confusion in the minds of agents between aggregate and relative price changes. Since the present model does not rely on imperfect information or confusing price signals, there is no need for a product-specific demand shock. Moreover, since there is no need for imperfect information, this model assumes that the money supply is immediately observed.

The third and most significant difference between the two models is that in the model of Section 6.4 the local market for each individual good is not perfectly competitive. Referring again to the island paradigm, instead of each island having a competitive market in which many producers supply a homogeneous good, each firm now has a monopoly on production of the good on its island. While there is competition among producers in the sense that the demand for each one’s product is sensitive to changes in others’ prices, the products of different producers are not perfect substitutes. This means that each producer has a downward-sloping demand curve and

maximizes profit like a monopoly: by setting marginal revenue equal to marginal cost.

The demand function for the i th agent's product is quite important in this model. It is given in log form as

$$q_i = y - \eta(p_i - p)$$

in text on page 286. The parameter η is the elasticity of demand for each individual good with respect to its relative price. This parameter is assumed to be the same across goods, so each good is an equally satisfactory substitute for others. The intercept term y is the log of aggregate demand. As in the Lucas model, aggregate demand depends on a monetary or aggregate-demand variable m .

As is typical in macro models, agents are totally symmetric, so all will end up supplying the same amount of labor. That means that no agent will be a *net* supplier or demander of labor in the labor market, so the amount of labor used in her production will be equal to the amount she supplies, or $Q_i = L_i$ for all i . Notice that if we substitute L_i for Q_i in Romer's equation (6.37), the WL_i terms cancel and we are left with (6.3) from the Lucas model. This shows that the introduction of the labor market does not really change anything relative to the Lucas model.

The crucial difference between perfect and imperfect competition is whether firms behave as price-takers. When competition is perfect, each agent takes the price of her good as given—determined by market forces beyond her control and “called out” by the Walrasian auctioneer. In the Lucas model, there are many agents producing *each* good, so that each agent is a small part not only of the macroeconomy but also of the market for her own product (on her own island). Thus, the utility function in the Lucas model (6.3) is maximized only with respect to labor effort to get the single first-order condition (6.4).

With imperfect competition, each agent is the *sole* producer of her particular commodity, which is an imperfect substitute for the products of others. In this case, each agent has some monopoly power and, as in the standard micro theory of monopoly, sets *both* price and quantity produced subject to the constraint of the demand function. Romer substitutes P_i/P from the demand function into utility function (6.37) to get (6.38). This utility expression is a function of two separate choice variables of the agent: price and labor effort. By differentiating the utility function with respect to both choice variables, we end up with two first-order conditions.

Differentiating (6.38) with respect to P_i and setting the resulting expression equal to zero gives (6.39), which can be solved as Romer suggests to yield (6.40). This equation reflects a standard result from the micro theory of monopoly—that the ratio of price to marginal cost (the “markup”) is a decreasing function $\eta/(\eta - 1)$ of the elasticity of the monopolist's demand. If the monopolist's good were to become a closer

substitute for other goods, demand elasticity η would increase and the markup ratio would get smaller. If the monopolist's good approaches being a perfect substitute for the goods of others, then $\eta \rightarrow \infty$ and the markup ratio approaches one. This is the (limiting) case of perfect competition.²

Differentiating (6.38) with respect to L_i and setting the resulting expression equal to zero gives the labor-supply equation (6.41) or (6.42). Thus (6.40) and (6.42) are the two equations that describe the behavior of the individual agent: what price she sets and how much she works. Quantity produced can be read from the demand equation given the price. This completes our microeconomic analysis at the level of the individual agent. Next we must examine how these agents will interact together to determine equilibrium in the aggregate markets.

Romer invokes the symmetry condition (that each individual supplies the same amount of labor) to determine that $L = L_i = Y_i = Y$. Recall that the aggregate variables Y and L are the averages of the corresponding individual variables, not the sums. You can think of them as per-capita values rather than aggregate ones. Straightforward substitution gets us to the results (6.46) and (6.47).

Romer characterizes (6.47) as the “equilibrium” price. However, we must think about this use of the term very carefully. The most common use of equilibrium to characterize a market price is to use it synonymously with “market-clearing.” However, this “equilibrium price” is somewhat different. The concept of market clearing is not quite the same in an imperfectly competitive market, since there is no supply curve to balance with demand. Instead, the imperfectly competitive firm chooses the profit-maximizing combination of price and quantity from the “menu” of alternatives offered by the demand curve. The equilibrium price in this case is the monopoly's profit-maximizing price given its demand curve. This price is an equilibrium in the sense that there is no incentive for it to change once that price has been established. It is also consistent with the basic notion of market-clearing because buyers are able to obtain as much of the good as they want at the equilibrium price, while suppliers are content to produce and sell that much and no more. It is clear that, to answer Barro's criticism, there are no \$50 bills lying on the sidewalk in the model of section 6.4.

Romer draws three important conclusions from (6.46) and (6.47) about the implications of imperfect competition for macroeconomics. First, he notes that the equilibrium quantity of output produced depends on the degree of imperfect competition,

² At the other extreme, if $\eta \rightarrow 1$ then the demand for the monopoly's good is no longer elastic. A monopoly with an inelastic demand curve can always increase profit by increasing price because revenues go up while costs go down (along with output). If the demand curve of a monopoly were everywhere inelastic, there would be no finite price at which it would maximize profit. If price increases enough, the demand for most goods would become elastic, so it is reasonable to assume that firms will set prices where $\eta < 1$.

as measured by the elasticity of demand for each individual product η . The lower the demand elasticity, the less competitive the market and the smaller is the equilibrium level of production. This mirrors the standard monopoly result that monopoly power leads to “contrived scarcity.” The welfare-maximizing quantity of output occurs under perfect competition. When $\eta \rightarrow \infty$, the equilibrium level of output approaches the optimal level, which, as in the Lucas model, is one.

This result implies that output and employment in the imperfectly competitive economy will settle at levels lower than those that maximize collective utility. Equilibrium here is not Pareto optimal. Everyone would be better off working harder and producing more, but the incentives of the price system under imperfect competition lead to a lower level of work and output. This conclusion supports the common layman’s perception that increases in output above its normal level are good and declines in output are bad. By contrast, in the Lucas model *all* deviations of output from the optimal value are equally undesirable.

The 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.

The third important conclusion follows from equation (6.47), which establishes that money and aggregate-demand shocks are neutral in the imperfect-competition model. The aggregate-supply curve is still vertical. Thus, imperfect competition with flexible prices is not sufficient to reverse the monetary neutrality result of the standard classical model. We shall see below that price-stickiness is the ingredient that leads directly to monetary non-neutrality.

Nominal rigidities

As Romer points out in his introduction to Part B, 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 price deviates from the price printed in the

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

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 6.2 or Figure 6.3. In Figure 6.3, 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 6.5, 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 6.2, 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.

Nominal and real rigidities

As Romer shows in the numerical example on pages 292 through 294, 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 6.3, 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 6.3) 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.

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 pro-

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.

portion 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 6.3) 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 6.3 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 299. 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.

D. 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.