

Chapter 12

New Keynesian Economics

The two current leading views of business cycles are Real Business Cycle (RBC) Theory and New Keynesian Economics. Each of these schools of thought has a rich history marked by frequent vigorous debate between them. While a terse summary of their general views does not do them justice, we first briefly highlight the main difference between the two points of view. RBC Theory views periodic expansions and recessions as natural, indeed efficient, responses of the economy to the ups and downs of the state of “technology” of the economy.⁶³ As such, recessions are not dire events for the economy, but rather natural slowdowns which are preceded by an expansion and which will again be followed by future expansions. In terms familiar from microeconomics, pure RBC Theory maintains that the aggregate economy operates perfectly competitively on both the demand side and the supply side.⁶⁴ An important implication of RBC Theory, therefore, is that the government has no role to play in the macroeconomy – that is, neither fiscal policy nor monetary policy can be used to improve the macroeconomic condition. In contrast, New Keynesian Economics adopts the point of view that there are fundamental market failures in the aggregate economy which render business cycle fluctuations, specifically periods of lower-than-potential GDP, inefficient.⁶⁵ The important implication of this point of view is that the government may indeed have a role to play in improving macroeconomic conditions.

Here we will consider just one strand of New Keynesian theory. As we will see, microeconomic analysis is at the heart of New Keynesian theory. Unlike most of our discussion of representative-agent macroeconomics, however, the focus of New Keynesian theory is not on the microeconomics of consumer behavior but rather on the microeconomics of firm behavior.

Differentiated Goods and the Consumption Aggregator

In our study of the representative consumer, we supposed that there was only one object (i.e., “all stuff”) which the consumer purchased in order to obtain utility. This was true in both the consumption-leisure model (in which there was literally only one object called

⁶³ “Technology” here is broadly defined – specifically, the Solow Residual (on which we will have more to say when we study RBC theory) is the most often-used measure of technology.

⁶⁴ Even the most ardent RBC macroeconomist does not literally believe the economy is perfectly competitive in the pure textbook sense, but rather that in aggregate market failures tend not to be so catastrophic as to make perfect competition a terrible approximation.

⁶⁵ The notion of “efficiency” you should have in mind throughout our discussion here is exactly that from microeconomics: a market (in our case, the entire macroeconomy) is operating efficiently if there is no deadweight loss, which means that no trades between suppliers and demanders which could increase overall utility go unconsummated. Another familiar characterization of economic efficiency is that price equals marginal cost.

“all stuff” which the consumer could purchase) as well as the consumption-savings model (in which there was one object called “all stuff” in each of the two periods of the economy which the consumer could purchase).

The use of a single consumption good is obviously a theoretical simplification. In reality, consumers purchase a vast number of goods and services from which they obtain utility. And in reality, these goods are somewhat substitutable for each other. For example, when making your decision about where to spend your Saturday evening, you may decide to go to the movies or go to a Mets game. Both options are forms of entertainment, but clearly they are not perfect substitutes for each other. Even if you decide to go to the movies, you will have to choose between the latest Harry Potter movie and the newest action thriller – clearly these two movies are also imperfect substitutes for each other.

If we believe that there are a great many options for consumption, each of which is at least a little different from every other option, available, then one way to reconcile our previous use of a single consumption good with this fact is to suppose that “all stuff” is composed of these great many differentiated goods. Specifically, we will now suppose that our usual notion of “all stuff” consumption is a function,

$$c = c(c^1, c^2, c^3, \dots, c^N) \quad (14)$$

where c^1 denotes the type-1 consumption good (perhaps a movie), c^2 denotes the type-2 consumption good (perhaps a Mets game), c^3 is the type-3 consumption good (perhaps dinner at a fancy restaurant), etc. If there are N different goods, then we have that “all stuff” consumption is a function of the N different types of consumption goods.

Note that we have written this function – formally called the **consumption aggregator function** – in abstract form. Especially note that we do **not** necessarily mean the simple sum $c = c^1 + c^2 + c^3 + \dots + c^N$. In fact, in theoretical New Keynesian models, the function $c(\bullet)$ is usually assumed to satisfy the following two properties: the first partial derivative with respect to consumption of good type i satisfies

$$\frac{\partial c(\bullet)}{\partial c^i} > 0 \quad (15)$$

and the second partial derivative with respect to consumption of good type i satisfies

$$\frac{\partial^2 c(\bullet)}{\partial (c^i)^2} < 0 \quad (16).$$

These two conditions state, respectively, that total consumption c is an increasing function of consumption of type i and that total consumption c increases at an ever-decreasing rate as consumption of type i increases. Think of this simply as (if you’re a

movie-buff) the more movies you see, the more total consumption (not just of movies but of “all stuff”) you enjoy – but the more and more movies you see, the less and less extra total consumption you gain. These properties should remind you of the general properties we imposed on the representative consumer’s utility function – note well, however, that the consumption aggregator function is not a utility function. Indeed, the utility function still takes total consumption c as an argument and continues to have the usual properties we have discussed at length. It is still utility, and not consumption, that is the maximization goal of the representative consumer. A very common functional form assumed for the consumption aggregator in New Keynesian models is

$$c(c^1, c^2, c^3, \dots, c^N) = \left[(c^1)^{1/\varepsilon} + (c^2)^{1/\varepsilon} + (c^3)^{1/\varepsilon} + \dots + (c^N)^{1/\varepsilon} \right]^\varepsilon, \quad (17)$$

where $\varepsilon \geq 1$. The value ε (the Greek letter “epsilon”) has very important economic meaning in these kinds of New Keynesian models. It governs how substitutable, from the point of view of the consumer, the different goods are for each other. At one extreme is the value $\varepsilon = 1$, which, when substituted into the above expression, yields the simple sum $c^1 + c^2 + c^3 + \dots + c^N$. With $\varepsilon = 1$, each consumption good is just as good as any other from the point of view of the representative consumer – that is, the goods are perfect substitutes for each other. With $\varepsilon > 1$, however, goods are only imperfect substitutes for each other, meaning that they are differentiated to a degree depending on the exact value of ε . A basis for all of New Keynesian economics is the assumption that $\varepsilon > 1$.

Monopolistically Competitive Firms

The heart of New Keynesian Economics lies not in the representative consumer, but rather with firms. Each of the N differentiated goods is assumed to be produced by a distinct monopolistically competitive firm. Recall from basic microeconomics that a fundamental feature of monopolistic competition is that goods are similar to each other but not completely identical. Continuing the example from above, movies and baseball games are similar goods (forms of entertainment) but obviously not identical.

Also recall from basic microeconomics that a firm that produces a differentiated good possesses market power. In terms of analysis that should be familiar, this market power manifests itself in the fact that the firm faces a downward-sloping demand curve and hence a marginal revenue curve that lies strictly below its demand curve. This in turn implies that the firm’s profit-maximizing choice features price greater than marginal cost – algebraically, $P^i > MC^i$, where P^i denotes the nominal profit-maximizing price of the firm that produces good i and MC^i is the nominal marginal cost at the profit-maximizing quantity. These features are summarized in Figure 53.




Figure 53. A monopolistically competitive firm faces a downward-sloping demand curve for its product. The marginal revenue curve thus lies strictly below the demand curve, and the firm’s profit-maximizing choice of output occurs where $MR = MC$. At this optimal quantity, price exceeds marginal cost.

Recall from expression (17) above that $\varepsilon = 1$ in the consumption aggregator implies that the goods are perfect substitutes for each other. The implication for firms of $\varepsilon = 1$ is that they do **not** have any market power (and thus face perfectly elastic demand curves). Thus, setting $\varepsilon = 1$ (in New Keynesian models that use this channel to introduce market failures) is one way of “shutting off” the New Keynesian elements of a New Keynesian model.

For the remainder of our discussion, we will assume that $N = 2$ for simplicity, so that the representative consumer’s total consumption is a function of two differentiated goods, each of which is produced by a distinct firm. Also, we assume $\varepsilon > 1$ (strictly), except where noted to highlight some issues, so that the model we are considering is indeed a New Keynesian model.

The Aggregate Price Level

In our simple models which featured only one homogenous consumption good, the nominal price level of the economy was a simple object – it was just the nominal price of the single consumption good. In our New Keynesian model here, however, even specializing to the case of just $N = 2$ differentiated goods renders the nominal price level of the economy a somewhat more complicated notion to consider. Clearly the aggregate price level, which we will denote by P , should depend somehow on the nominal prices of the two distinct goods, which we denote by P^1 and P^2 . That is, P is some function

of P^1 and P^2 . There are many possible ways of aggregating the individual prices into a single measure of the price level of the economy. We will refrain from putting a particular functional form on this **price-level aggregator** (even though New Keynesian models offer a great many from which to choose) and simply use the abstract function

$$P = P(P^1, P^2). \quad (18)$$

To re-emphasize, the unsuperscripted P denotes the nominal price level of the economy (the price of a “market basket” of goods) while P^1 and P^2 denote, respectively, the nominal prices of good type 1 and good type 2. For our purposes the important feature of this function is that P is strictly increasing in each of its arguments. In calculus notation, this means

$$\frac{\partial P(\bullet)}{\partial P^1} > 0 \quad (19)$$

and

$$\frac{\partial P(\bullet)}{\partial P^2} > 0. \quad (20)$$

The assumption that the price level P is strictly increasing in each of the individual prices should strike you as reasonable. To draw an analogy with the Consumer Price Index (CPI), if the price of any of the goods in the CPI basket rises, the aggregate price level rises even though the relationship between the CPI and the individual price is usually not a straightforward one. Similarly here in our New Keynesian model, the functional relationship between the price of any individual good and the aggregate price level is in general a complicated one, but the aggregate price level depends positively on each individual price in the basket.

Aggregate Consumption Demand

Each of the two firms (in our case where we have specialized to $N = 2$) faces a downward-sloping demand curve for its product. Because there is more than one good, we now need to define aggregate consumption demand. A simple definition of aggregate consumption demand may seem to be that aggregate consumption demand is the sum of the demands for each of the differentiated products. Graphically, this latter would mean that aggregate consumption demand is the “horizontal summation” of each individual good’s demand curve.⁶⁶ However, such a procedure would be incorrect in our New Keynesian model.

Horizontal summation of demand curves is incorrect here because goods that are different cannot be summed. For example if we have 20 apples and 10 bananas, we have – well,

⁶⁶ This notion of horizontally summing demand curves should be familiar to you from microeconomics.

20 apples and 10 bananas. That is, they inherently cannot be summed because they are different objects.⁶⁷ More fundamentally, such a procedure turns out to approach the problem of constructing the aggregate consumption demand from the wrong point of view. Recall in our study of the consumption-leisure model that we derived the (aggregate) consumption demand function. This consumption demand function remains the correct notion of consumption demand, because it is total consumption, which in turn is the consumption aggregator, which still is the direct argument of the representative consumer's utility function.

The consumption aggregator then shows the relationship between aggregate consumption demand and demand for each of the differentiated goods. Recall from above that total consumption (the “all stuff” object) is an increasing function of each type of differentiated consumption good. This fact is all we need to conclude that there is a positive relationship between total consumption and consumption of each differentiated good. We thus diagram the relationship between the demand for the two goods and aggregate consumption demand in Figure 54.




Figure 54. There is a positive relationship between aggregate consumption demand and demand for each differentiated good. Aggregate consumption demand is derived by using the consumption-leisure model.

The way we will interpret the relationship depicted in Figure 54 is the following: when some event (the usual events are preference shocks and government policy) shifts the consumption demand function, the demand functions facing the two individual firms also shift in the same direction. Because the individual firms' demand functions shift, the associated marginal revenue functions shift as well, implying a new profit-maximizing choice of price and quantity for each individual firm.⁶⁸

To preview the comparison of New Keynesian theory and RBC theory we will make after studying both points of view, notice that the exogenous event – that is, the “unexplained” event, or shock – that begins the thought experiment is a shift in a component of demand (here, consumption). New Keynesian theory holds that “demand shocks,” coupled with the rest of the theoretical apparatus we are describing in this chapter, are the predominant

⁶⁷ We could then define the term “fruit” to be either a banana or an apple, and assume the conversion one apple = one fruit and one banana = one fruit. With these assumptions (which notice probably come very naturally to you as you think about this example), we would then say we have 30 pieces of fruit.


⁶⁸ Convince yourself of this last point by diagramming it.

factor that causes macroeconomic fluctuations. In contrast, RBC theory holds that it is “supply shocks” that are the predominant source of macroeconomic fluctuations.

Staggered Price-Setting

Consider the case in which the consumption demand function shifts outwards for some reason. As just described, this event in turn causes the demand functions, and hence the associated marginal revenue functions, facing each individual firm to shift outwards as well. Assuming the marginal cost functions do not shift, these events imply new profit-maximizing choices of price and quantity for each of the two firms in Figure 54. In particular, both quantity c^i and price P^i rise.⁶⁹ Because the price of each good rises, clearly the aggregate price level P rises because of the properties of the price-level aggregator we described above.

However, suppose instead that only one of the two firms can change its price at the time of the event that shifts the aggregate consumption demand curve. Let’s suppose it is firm 2 (the firm that produces consumption good c^2) that has the ability to change price. Firm 1 (the firm that produces consumption good c^1), on the other hand, cannot change its price at all. The situation facing firm 1 is illustrated in Figure 55.



⁶⁹ Again, diagram it to convince yourself.

Figure 55. After a shift outwards of the consumption demand function, the demand function facing firm 1 and hence firm 1's marginal revenue function both shift outwards. If firm 1 could change its price, it would raise both its price and quantity produced because the new MR function intersects the MC function at a higher quantity. Associated with such a price increase is a rise in the quantity produced. However, if firm 1 continues with its initial optimal price, then the quantity of good 1 that it produces rises by even more, as can be read off the new demand function.

Figure 55 shows the events facing firm 1 following a rise in consumption demand. Firm 1's demand function and marginal revenue function both shift outwards. If Firm 1 were to change its price, it would choose the price labeled "optimal P^1 if no price stickiness" because that price yields the quantity at which marginal cost equals marginal revenue. We can see from the diagram that quantity produced would rise. However, if firm 1 does not change its price and is forced to continue using the initial optimal P^1 , then its quantity produced will rise by even more, as can be read off the new demand function.

Firm 2 is assumed to be able to change its price immediately. It thus faces a similar situation as firm 1, except it will raise its price P^2 . Consequently, its quantity produced also rises, but by less than it would if were also unable to change its price.

Implications for Government Policy

Now we consider how the macroeconomic effects of policy differ depending on whether or not firm 1 is able to change its price. First consider the case where firm 1 can change its price. We can conclude that the aggregate price level P rises because the aggregate price level is an increasing function of both P^1 and P^2 . Because both P^1 and P^2 rise, P clearly rises. On the other hand, suppose firm 2 can raise its price but firm 1 does not raise its price. In this case, P^2 still rises by the same amount as in the case that firm 1 could change its price because firm 2's decisions do not depend on firm 1's actions.⁷⁰ However, because P^1 does not change, the aggregate price level P does not rise by as much as in the case that firm 1 did change its price. Simultaneously, overall consumption c rises by more with firm 1's price held constant because c^1 rises by more in this case.

In the preceding conclusion lies the important implications of New Keynesian theory for government policy. In the presence of staggered price-setting, government policy that raises aggregate consumption demand generates less inflation and a larger increase in production than the same government policy in the absence of staggered price-setting. Thus, staggered price-setting gives economic policy-makers leverage over real quantities in the economy, more leverage than they would have if all firms could adjust prices simultaneously.

Critique of New Keynesian Theory

⁷⁰ That is, there are no game-theoretic interactions here.

The obvious criticism of the strand of New Keynesian Theory that we have developed is the assumption of staggered price-setting. In particular, we have offered no explanation why some subset of firms cannot or does not change its price in the face of an increase in its demand while some other subset does change its price.

One common justification given for staggered price-setting is the presence of asynchronous menu costs. Recall that menu costs are the costs incurred by a firm simply by the act of changing prices.⁷¹ If different sectors of the economy experience menu costs of different magnitudes at different times during the business cycle, then staggered price-setting may arise. For example, perhaps it is easier (i.e., less costly) for military-equipment producers to raise their prices during war time than it is for producers of entertainment simply because of the political and/or cultural environment of the time. This example suggests that “menu costs” need not be interpreted as only direct costs of changing prices but also can include more intangible costs such as lost goodwill, etc. While there is little empirical evidence suggesting the magnitude of these intangible menu costs, it is at least a plausible story.

At a more realistic level, it does seem true that firms do not all change prices simultaneously. If we accept this, then staggered price-setting seems less strange of an assumption and perhaps we can accept not having a stronger microeconomic foundation for it.

Finally, strong empirical support for the New Keynesian view comes from data that suggest that government policy, both fiscal and monetary, does have important short-run effects on output. Data generally support the view that government policy is more effective than real business cycle (RBC) theory predicts. This last observation alone may be enough justification for studying New Keynesian theory – and indeed, the debate between the RBC view and the New Keynesian view has been one of the most vigorous areas of debate in macroeconomics over the past few years.

⁷¹ The terminology itself implies its meaning: a restaurant that wants to change its prices must print new menus, which itself has a cost associated with it.

Appendix: Theories of Price Stickiness

During the normal ups and downs of the economy (termed “business cycles,” a topic which is usually in the domain of macroeconomics), demand tends to fluctuate – even demand for broadly-defined categories of goods. In terms of a supply and demand diagram, this means the demand curve shifts in and out during the “normal” course of economic events. We will consider “good economic times” to be periods of relatively high economy-wide income and “bad economic times” to be periods of relatively low economy-wide income. If we further adopt the simplifying assumption that all goods are normal goods, this means that in good economic times demand curves tend to shift out, while in bad economic times demand curves tend to shift in. With an unshifting supply curve, these shifts of the demand curve imply fluctuations of equilibrium price.

At the other extreme, price fluctuations could arise solely due to changes in costs for firms, holding the demand function constant. Indeed, many industries are subject to such cost “shocks” from time to time. For example, the production costs of manufacturing firms, whose output makes up a sizable fraction of the U.S. economy, generally rise when the price of steel rises.⁷² When the general wage level in the economy rises, as occurred during much of the 1990’s in the U.S., all firms’ costs rise. Profit maximization by firms implies that firms would try to pass along most if not all of any increase in production costs to consumers in the form of higher prices.

Thinking then of our usual downward-sloping demand curves and upward-sloping supply curves, these fluctuations in demand and costs should imply that (equilibrium) prices fluctuate a fair amount during the course of business cycles. But much empirical evidence has shown that prices do not in fact fluctuate very much over short periods of time. The phenomenon that prices do not fluctuate as much as standard economic theory predicts is known in academic circles as “price stickiness.”

Here we will consider five simple theories of price stickiness and also visit some of the real-world evidence surrounding these theories. The discussion here follows the basic structure presented by Blinder and his coauthors in *Asking About Prices* (1998), a very interesting exploration of the origins of price stickiness that proceeds, rather than by simply presenting economic theories in a vacuum, by asking managers of businesses in various industries about the relative importance of various postulated theories of price stickiness.

Theory I: Constant Marginal Cost of Production

Suppose markets are perfectly competitive and that all firms have constant marginal cost (MC) curves over the relevant range of outputs. That is, for a single firm, if the “normal” range of output is between 100 units and 500 units, suppose the MC curve is perfectly

⁷² As occurred, for example, when the Bush Administration passed a set of steel tariffs in March 2002.

horizontal over this range. A constant MC schedule is not so difficult to rationalize. For example, if a firm has a machine that makes between 100 and 500 units of output and the electricity, water, etc. necessary to operate the machine has constant per-unit cost (perhaps because the firm has a contract with the utility company), then over the range 100-500 units, the extra total cost incurred in producing one more unit could be constant – that is, MC is constant. Because the market supply curve is simply the horizontal sum of individual firms' marginal cost curves, this implies that the market supply curve in such an industry is perfectly elastic.

Figure 56 illustrates how prices can be constant despite fluctuations in the demand schedule during the course of the business cycle. Thus, constant marginal cost is one reason why prices may be sticky.




Figure 56. If all firms in a perfectly competitive industry have constant marginal cost, then the market supply curve is perfectly elastic, which in turn implies that fluctuations in the demand curve leave prices unchanged.

Evidence about marginal cost curves in U.S. industries provides surprisingly more support for constant marginal cost than many economists may think. According to interview data compiled by Blinder et al (1998), 40% of U.S. goods and services are produced by firms that report that they have constant marginal cost over “usual” ranges of production. Only 11% report rising marginal cost over “usual” ranges of production, while 33% report **declining** marginal costs over “usual” ranges of production (but here there seems to be reason to suspect that respondents were confusing marginal cost with average total cost).⁷³ This evidence seems at odds with the usual textbook assumption that marginal costs rise as output expands.

⁷³ Recall that it is possible for marginal cost to be rising even as average total cost is falling – this is true when the MC curve is below the ATC curve.

Theory II: Variable Elasticity of Demand

Instead of considering perfectly competitive industries, let us think of monopolistically competitive industries. Recall that in monopolistically competitive industries, each firm, because it produces a slightly different good than every other firm in the industry, faces a downward-sloping demand curve.⁷⁴ Here we will suppose that each individual firm's MC curve has the usual upward slope. Profit maximization by the individual firm implies choosing the quantity such that marginal revenue equals marginal cost ($MR = MC$), as shown in Figure 57, in which q^* is the profit-maximizing quantity and P^* the associated profit-maximizing price.




Figure 57. A monopolistically competitive firm faces a downward-sloping demand curve, and profit maximization implies producing that quantity that equates marginal revenue to marginal cost.

In good economic times, when incomes are relatively high, the demand curve, and hence the marginal revenue curve, in Figure 57 would shift outwards. If this shift out is a parallel shift, the new profit-maximizing quantity and price are as shown in Figure 58. The profit-maximizing price rises due to the increase in demand.

The shift out in the demand curve may not be a parallel shift, however. Figure 59 shows a shift out of the demand curve in which the entire demand curve becomes more elastic (that is, flatter) as economy-wide income rises. Demand curve Q_2^D is shifted out relative to demand curve Q_1^D but is flatter than Q_1^D . At high prices, the demand curve has shifted out by less than at lower prices. Such a nonparallel shift may seem reasonable on the grounds that as incomes rise, the “newest” units sold in the market are sold to those consumers who have just entered the market. These newest consumers are likely to be the most price-sensitive consumers – that is, the newest consumers are likely to be the ones who have the most elastic demands in the first place. Thus, the overall market demand curve becomes more elastic as it shifts out.

⁷⁴ As opposed to the perfectly horizontal demand curve that an individual firm in a perfectly competitive industry faces.

The new profit-maximizing quantity in Figure 59 is q_2^* , the point at which the new MR curve intersects the unchanged MC curve. But the profit-maximizing price, read off of the new demand curve, is unchanged. Of course, the case illustrated is a very special case – the more general point is that because of possibly varying elasticity of demand as demand curves shift, price volatility may be much smaller than would otherwise be the case.




Figure 58. When the demand curve facing a monopolistic competitor shifts out (with an associated shift out of the marginal revenue curve) in parallel manner, the profit-maximizing quantity and price both increase, as long as the marginal cost curve remains stable.




Figure 59. If the demand curve facing a monopolistic competitor shifts out in a non-parallel manner, it is possible that the profit-maximizing price may be unchanged.

To summarize, the basic idea of this theory of price stickiness is that demand curves becomes less elastic as they shift in and more elastic as they shift out.

To operationalize this theory in their interview study, Blinder et al (1998) posed the idea as one in which when business turns down a company loses its *least* loyal customers first and retains its *most* loyal customers. Their presumption in stating the theory in this way to business managers was that the most loyal customers are the least price-sensitive ones. When asked whether this phenomenon was an important one in their own pricing decisions, 56% responded that it was totally unimportant, 13% responded that it was of minor importance, 22% responded that it was of moderate importance, and 9% responded that it was very important. Thus, their survey indicates that while this theory of price stickiness cannot be claimed to be *the* theory of price stickiness, it does not seem completely irrelevant.⁷⁵

Theory III: Kinked Demand Curve

Consider again monopolistically competitive firms, which face downward-sloping demand curves. One classical theory of how an individual monopolistically-competitive firm behaves is that it believes any price cut it initiates will be matched by its competitors while any price hike it initiates will not be matched by its competitors. The idea is simple: if a firm lowers its price, other firms (which offer similar but not identical products) will find it in their best interest to lower their price as well so as not to lose

⁷⁵ Notice also that for this theory of price stickiness to be an important one, much of the economy must be characterized by monopolistic competition rather than perfect competition. The issue of whether perfect competition is descriptive of most of the economy is a matter of ongoing debate.

their customers to the first firm. Thus, lowering price will lead to an increase in quantity demanded, but not a very large increase. On the other hand, if a firm raises its price, other firms will find it in their best interest to keep their prices unchanged, under the belief that customers of the first firm will switch to buying their products instead. Thus, a price increase will lead to a large decrease in quantity demanded. The preceding discussion can be formalized in the demand curve shown in Figure 60, in which the demand curve above the price P^* is relatively elastic (flat) while below the price P^* is relatively inelastic (steep).

This “kinked demand curve” leads to a break in the MR curve. This break is illustrated as a jump down in the MR curve from c_H to c_L in Figure 60. The reason for the jump in the MR curve is simply the fact that the slope of the **total revenue (TR)** function changes abruptly at the kink price – and **marginal revenue is simply the slope of the total revenue function**. Let’s understand this a bit more carefully: recall that total revenue is defined as $TR = P \cdot Q$. But from the perspective of a price-setting firm such as a monopolistic competitor, price is related to quantity by the demand function, so we can write the *function* $P(Q)$ to stand for price. With this, we have that total revenue is given by $TR(Q) = P(Q) \cdot Q$, where again the functional dependences are emphasized by the notation. When graphing TR versus Q, then, the slope is $TR'(Q) = P'(Q) \cdot Q + P(Q)$, which follows from the chain rule of calculus. But, as mentioned above, marginal revenue is the slope of the total revenue function, so $MR(Q) = TR'(Q) = P'(Q) \cdot Q + P(Q)$. At the kink point in Figure 60, $P'(Q)$ changes abruptly (i.e., discontinuously), causing a discontinuity in the MR function.




Figure 60. In the kinked demand curve model, the MC curve could cross the MR curve anywhere in the region between c_H and c_L and the profit-maximizing price would be P^* regardless. Thus, the price P^* is a “sticky” price for this firm.

Superimposing an MC curve in the kinked demand curve model, as Figure 60 does, shows that if the MC curve crosses the MR curve anywhere in the region between c_H and c_L , the profit-maximizing quantity and price remain unchanged. Thus, P^* is a “sticky” price in the kinked demand curve model because marginal costs can shift over a broad range with no attendant change in price.

In their interview study, Blinder et al (1998) posed the kinked demand curve model to business managers as one of “price leadership.” As discussed above, the underlying idea of the kinked demand curve model is that competitors will match price changes in the downward direction but not in the upward direction. Thus, if they could, firms would find it profitable to all agree simultaneously to raise prices.⁷⁶ Because such collaboration is usually not feasible, some firms (perhaps the biggest ones, say, or the most popular ones) may naturally emerge as the “price leaders” of the industry: if those leader firms raise prices, that sends a signal to the rest of the firms in the industry that price hikes will be matched. The specific question the Blinder study posed to business managers concerned “the importance of price coordination failure” as the reason why a firm held back on price changes. A very large 62% of firms responded that this factor was a very important one in their pricing decisions. Of the thirteen different theories of price stickiness the Blinder study explored, this one received the most support from business managers, lending support to the importance of the kinked demand curve model as a theory of price stickiness.

Theory IV: Psychological Pricing Points

Casual introspection of your own experiences as a consumer will probably convince you that certain prices occur more often than others. For example, a \$999 computer is probably more common than a \$1,000 computer. Or a \$9.99 pizza pie is probably more common than a \$10 pizza pie. Many sellers, especially in the retail business, apparently believe that certain threshold prices have special psychological significance to consumers: a \$999 computer may seem “a lot cheaper” than a \$1,000 computer, for instance. In terms of our model of supply and demand, the demand curve would be extremely elastic at prices just above these threshold prices, as illustrated in Figure 61. That is, prices just above the threshold pricing points would cause many consumers to exit the market.

⁷⁶ Such collaboration on price-setting, called “collusion,” is illegal in the U.S. A number of laws, including the Sherman Antitrust Act and the Clayton Act, outlaw collusive behavior among firms.

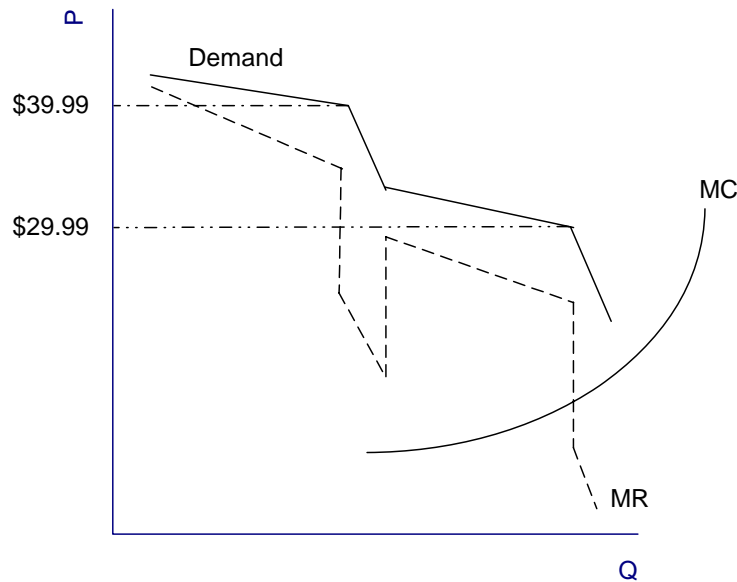


Figure 61. With psychological pricing points (here, \$39.99 and \$29.99), the demand curve is very elastic at prices just above those threshold prices and less elastic at prices just below those threshold prices. Thus, the monopolistic competitor’s marginal revenue curve has multiple jumps in it, one corresponding to each kink point on the demand curve. Thus, the MC curve can intersect the MR curve at multiple vertical sections of the MR curve, giving rise to the “pricing points” as the profit-maximizing prices.

The demand curve in Figure 61 has multiple kinks and therefore multiple jumps in the MR curve. Thus, the analysis here proceeds similar to the case of the kinked demand curve discussed above. As shown in the figure, the MC curve can move up and down over a broad range and still cross the MR curve at the same quantity, at which the profit-maximizing price is \$29.99. The MC curve could also shift so far up as to cross the MR curve in its leftmost vertical region, at which the associated profit-maximizing price is \$39.99. Once the MC curve moves into this latter region, it can again move up and down over a broad range in which the profit-maximizing price becomes “stuck” at \$39.99.

The interview study by Blinder et al (1998) asked business managers several different questions relating to psychological pricing points, the most relevant of which to our discussion here is “How important are psychological price points in deterring price increases in your company?” As with the theories above, the responses yield mixed evidence: 59% responded price points were totally unimportant, 16% responded they were of minor importance, 15% responded they were of moderate importance, and 10% responded they were very important. So this is yet another theory of price stickiness that cannot be claimed to be *the* theory of price stickiness but does seem to have some relevance.

Theory V: Menu Costs

Perhaps the predominant modern theory of price stickiness⁷⁷ is that the very act of changing prices itself entails costs. Indeed, this is also the simplest of theories of price stickiness. The basic idea is most easily illustrated with an example. Suppose a restaurant is considering increasing the prices of some of the items on its menu. Presumably, price increases are being considered because they would be in the best interest of the restaurant – that is, the price increases would presumably increase total profit. To make the example concrete, suppose that at current demand conditions, if the restaurant could costlessly change its prices, \$1000 in extra total profit would be generated. However, in order to implement its price changes, the restaurant would have to print new menus. If the restaurant had to pay its printer \$2000 to print new menus, it clearly is not in the interest of the firm to change its prices – indeed, changing prices would cause total profit to decrease by \$1000, so the firm instead chooses to hold its prices steady. This example suggests the terminology: a **menu cost** is a cost incurred by a firm due to the price-adjustment process itself – in our example, it is literally the price of printing new menus.

Let us now consider more generally how the presence of menu costs may deter price changes and hence induce price stickiness. For this discussion, we again must suppose that firms are monopolistically competitive.⁷⁸ We also make two further simplifying assumptions:

1. Fixed costs are zero⁷⁹
2. Marginal cost is constant (which along with the assumption that fixed costs are zero implies that average total cost is identical to marginal cost)⁸⁰

With these assumptions, Figure 62 illustrates the usual case of a monopolistic firm's profit-maximizing price and quantity both rising if the demand for its good shifts to the right.⁸¹

Now, with the aid of Figure 63, we will consider how much extra profit accrues to the firm if it changes price to the new optimal price. In Figure 63, the area labeled π^{initial} is

⁷⁷ At least in modern macroeconomic theory – part of the foundation of a class of macroeconomic theories which go under the name “New Keynesian” theories are menu costs. You may encounter this topic in a course on macroeconomics.

⁷⁸ Challenge question: why does this theory become meaningless if markets are assumed to operate perfectly competitively?

⁷⁹ The basic results we will now develop do not depend on this assumption, although the graphical analysis is somewhat more complicated with fixed costs.

⁸⁰ Recall from our first theory of price stickiness that 40% of firms in the Blinder et al study reported having constant marginal cost over their usual levels of production, so this assumption seems not too outlandish – again, though, the main results to follow do not depend on this assumption, but making this assumption makes the graphical analysis clearer.

⁸¹ Thus, here we are assuming that any change in elasticity of demand as demand shifts out is not enough to render the optimal price sticky.

the total profit the firm earned under the initial demand curve at the initial optimal price.⁸² Once the demand curve shifts out, the firm would like to raise its price. The area labeled $\pi^{\text{additional, adjust}}$ is the additional profit generated if the firm does adjust its price. However, suppose that in order to adjust its price, the firm incurs a one-time fixed menu cost in the amount C . That is, if it wants to change its price, the cost of printing new price lists, etc. is C . Thus, the total profit the firm earns if it chooses to adjust its price following the shift in demand is given by

$$\pi^{\text{adjust}} = \pi^{\text{initial}} + \pi^{\text{additional, adjust}} - C \quad (21)$$

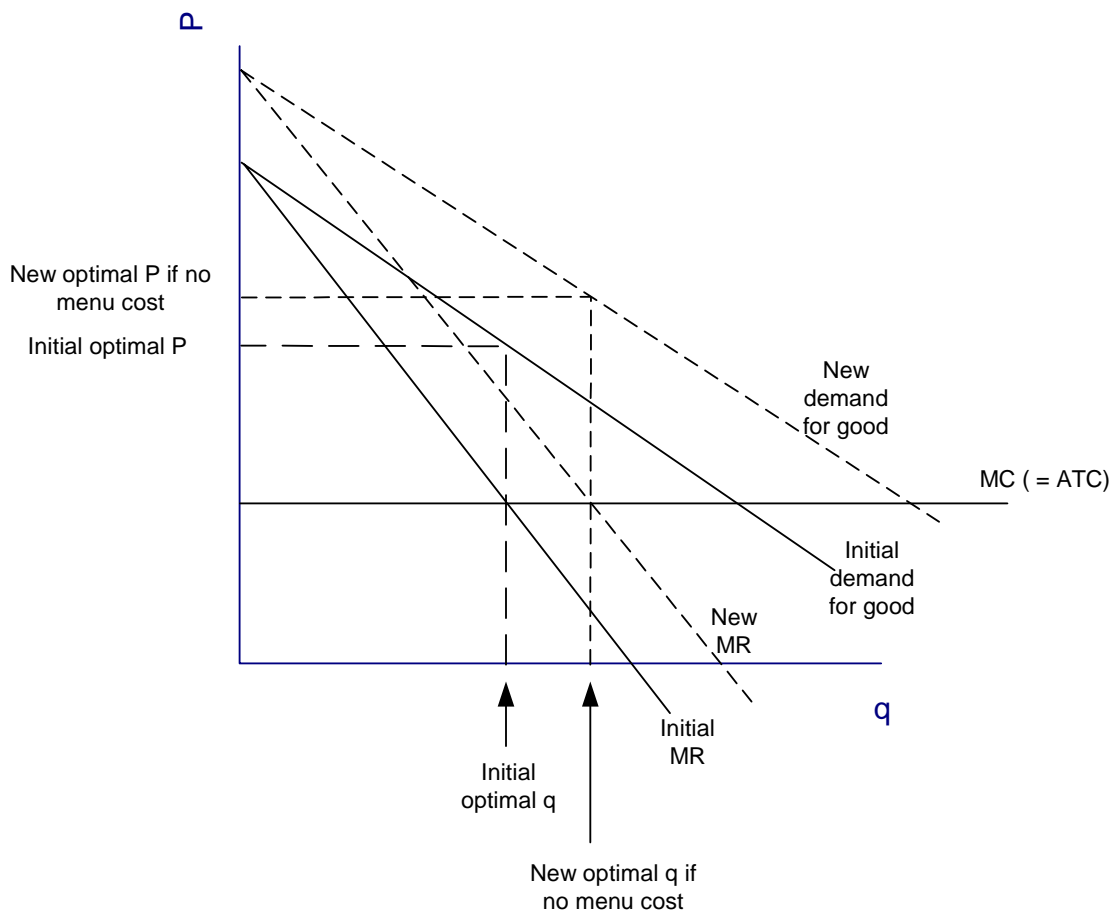


Figure 62. In the absence of menu costs, an outward shift of demand for a good generally causes a monopolistically competitive firm to increase both its price and quantity.

⁸² This is simply an application of the definition of the profit function, $\pi(q) = (P(q) - ATC(q))q$, at the profit-maximizing level of output.

If this menu cost C is large, the firm may be deterred from implementing the price change. With the aid of Figure 64, we can now compute how large a value of C would persuade the firm to hold its price constant despite the shift in demand.




Figure 63. Following an outward shift of the demand function (and hence an outward shift of the marginal revenue function) and in the absence of any menu costs, profits are the sum of initial profits and additional profits generated by the rise in price and quantity. If there is a menu cost, however, this menu cost must be subtracted from profits.




Figure 64. If the firm does not adjust its price following a shift in demand, then it must produce according to the new demand function at the price it charged before the shift in demand. Because this price in general is no longer the profit-maximizing price, for the last several units of output it produces the firm will be incurring a loss, represented graphically by the fact that for the last several units produced marginal cost exceeds marginal revenue. Total profit is given by the sum of initial profit and additional profit generated by the increase in quantity.

If the firm holds its price constant at the initial optimal P , then it will produce at the quantity labeled “new q if menu cost” in Figure 64 – that is, it must produce according to the new demand function.⁸³ For every unit produced between the points “new optimal q if no menu cost” and “new optimal q if menu cost,” the firm is incurring a loss because for each of these units marginal cost exceeds marginal revenue.⁸⁴ Thus, the total profit the firm earns if it holds its price constant following the shift in demand is given by

⁸³ Because, after all, the demand function is a relationship between price and quantity.

⁸⁴ If this point is not clear, now is the time to review the idea that profit-maximization implies producing until the point at which marginal revenue equals marginal cost.

$$\pi^{\text{don't adjust}} = \pi^{\text{initial}} + \pi^{\text{additional, don't adjust}}, \quad (22)$$

in which the menu cost does not appear because the menu cost is incurred only if the firm adjusts its price. Comparing the profit expressions in equation (21) and equation (22), it should be clear that $\pi^{\text{adjust}} > \pi^{\text{don't adjust}}$ if and only if

$$\pi^{\text{additional, adjust}} - \pi^{\text{additional, don't adjust}} > C.$$

That is, only if the menu cost is smaller than the **extra additional** profit the firm would enjoy by changing price should it incur the menu cost and adjust its price. Indeed, this result is sensible – only if the menu cost is smaller than the benefit of changing cost (which is the extra profits generated by the price change) should the firm choose to adjust its price.

The interview study by Blinder et al (1998) posed several questions to managers regarding the presence and importance of price adjustment costs in pricing decisions. Out of over 200 managers that responded to the question “Does your company have explicit costs of price adjustment?”, 36% responded that they had no costs of price adjustment, 21% responded that they did but that their size was trivial, and 43% responded that they had adjustment costs which were non-trivial. The very general way in which this question was posed, however, left open the possibility that the “price adjustment costs” which managers had in mind were not the “menu costs” we have described here. Thus, to further investigate this, managers who responded that they had at least “trivial” costs of price adjustment were asked what the nature of the most important of those price adjustment costs were. These responses are presented in Table 2.

	Nature of Cost	Percent of businesses with price adjustment costs citing item as most important type of cost
1	Printing new catalogs, new price lists, new packaging, etc.	28.4%
2	Loss of future sales by antagonizing customers	27.4
3	Informing salespeople and customers	24.0
4	Decision making time of executives	15.6
5	Getting the sales force to cooperate	4.5

Table 2. Of the 64% of all managers who responded that their businesses had at least trivial costs of price adjustment, the five types of price adjustment costs cited as the most important. (Source: Blinder et al (1998), p. 234)

Of the five types of price adjustment costs cited in Table 2, item 1 (printing new catalogs, etc) and item 3 (informing people of price change) accord most with the notion of menu costs we developed above.

References

Blinder, Alan S., Elie R. D. Canetti, David E. Lebow, and Jeremy B. Rudd. 1998. *Asking About Prices: A New Approach to Understanding Price Stickiness*. Russell Sage Foundation, New York.

