

Chapter 13

Real Business Cycle Theory

Real Business Cycle (RBC) Theory is the other dominant strand of thought in modern macroeconomics. For the most part, RBC theory has held much less sway amongst policy-makers than has New Keynesian theory. Among theoretical macroeconomists, however, RBC theory is very well-known and well-understood and even provides the foundations for some of New Keynesian theory.⁹⁴ Although there are a number of ways in which RBC theory differs from New Keynesian theory, we will focus on two differences. The most important difference by far is that RBC theory eschews the idea of sticky prices, while New Keynesian theory embraces it. RBC theory views prices as fully flexible – that is, all prices can be and are re-set very frequently. Even more precisely, RBC theory supposes that perfect competition in all markets is a good starting point for analyzing the macroeconomy. Second, RBC theory does not view exogenous shifts of consumption demand as a good description of data, but rather “shifts in supply” as the predominant reason for macroeconomic fluctuations.

The basic mechanics that we will use to sketch out the main elements of RBC theory are the theory of the representative firm, the simple consumption-savings model, and the static consumption-leisure model. We could instead employ the intertemporal consumption-leisure model, rather than the simple consumption-savings model and the static consumption-leisure model in tandem. As we saw earlier, though, the algebra becomes quite messy and graphical tools become difficult to use. RBC theorists do in fact use the intertemporal consumption-leisure model in their workhorse models, but we will be able to develop the basic results using the two models together.

The RBC Technology Shock

Recall from our discussion of the aggregate production function $f(k, n)$ that we could augment it with a technology parameter A , so that total output is given by $A \cdot f(k, n)$. This technology parameter is usually identified with the Solow Residual, which is a measure constructed from data on output, capital, and labor. We describe how to compute Solow Residuals soon. This way of measuring technology has the virtue that it does not require taking a stand on what constitutes “technology” – i.e., it does not require

⁹⁴ Although not the staggered price-setting that we emphasized. RBC theory has made important contributions to the understanding of macroeconomics despite never having taken center stage in policy debates. The pioneering work of Ed Prescott and Finn Kydland, widely viewed as the “fathers” of RBC theory, was finally widely-recognized in 2004 when they were jointly awarded the Nobel Prize in Economic Sciences.

identifying the state of “technology” of an economy with, say, the number of computers it uses or with the number of Ph.D.’s it employs or with how many people use wireless internet connections or any number of other measurements you might be able to think of that somehow capture how “technologically advanced” an economy is.

The most commonly-used production function in RBC theory is the **Cobb-Douglas production function**

$$f(k, n) = k^\alpha n^{1-\alpha}, \quad (11)$$

in which the parameter α measures the percentage of total GDP that goes towards paying for the costs of capital used in production and the value $1-\alpha$ measures the percentage of total GDP that goes towards paying for the costs of labor used in production. More common terminology is that “ α is capital’s share of output” and “ $1-\alpha$ is labor’s share of output.” Empirical evidence shows that for the U.S. α is about 0.33. Thus, about one-third of the total value of goods and services produced (i.e., GDP) pays for the capital used in production, while the remaining two-thirds pays for labor costs.

To illustrate how to compute Solow Residuals (and hence the level of technology), consider the following example. Suppose that it is known (i.e., can be measured) that in the year 2003 the capital stock of the economy was $k=1000$, and the quantity of labor used was $n=8$. Suppose also that the Cobb-Douglas function $k^\alpha n^{1-\alpha}$ describes the economy in question, and it is known (or at least estimated) that $\alpha=1/3$. Finally, total output (GDP) in the year 2003 was $y=50$. Using this information, it is possible to compute the level of technology as that amount left “unaccounted for” in the transformation of inputs into outputs. Because we know that $y = A \cdot f(k, n) = A \cdot k^\alpha n^{1-\alpha}$, we can back out the value of A during this period. Using the given data, it follows that $A=1.2$. Now suppose that in the year 2004, the capital stock, the quantity of labor used, and the production function (including capital’s share α) all remained unchanged, but total output was $y=60$. Again using the production function, we can conclude that in 2004, $A=1.5$, meaning that technology improved between 2003 and 2004. This notion of “technology” is a very broad and in some sense vague one – it simply identifies “technology” as some unexplained factor that changes the nature of the production process. In our simple example, the capital stock and quantity of labor did not change between 2003 and 2004, yet output increased. The reason for this may be many-fold: the quality of computers and machines used in production may have improved; decreased government regulation may have removed hindrances on companies’ practices; the state of knowledge of workers in the economy may have advanced (i.e., people may have become more educated), etc. As this list suggests, this macroeconomic notion of

“technology” need not correspond literally to the usual notion of technology, that of computers and the Internet, etc.⁹⁵

Technology Shocks and Aggregate Fluctuations: An Overview

Unexplained variations in the technology parameter A should be viewed as supply shocks because they affect the production function, which ultimately determines the supply function of an economy. This stands in contrast to New Keynesian theory, which holds that shocks to aggregate demand – in the form of shocks to government policy or consumer preferences – provide the important impetus for business cycles. In the influential study that effectively launched the RBC school of thought, Kydland and Prescott (1982) found that fluctuations in the Solow Residual accounted for well over half of fluctuations in GDP, leading them to conclude that a theory of business cycles could be built with technology as its centerpiece.

Before we study in more depth how RBC theory works, we sketch the basic outline of the theory. With perfect competition, the real wage rate and the real rental rate of capital in the economy respond to technology shocks immediately.⁹⁶ For simplicity, assume that it is consumers rather than firms that own the capital that is used in production and that consumers rent their capital on a period-by-period basis to firms, in addition to supplying labor to firms.⁹⁷ Consider a temporary positive rise in productivity. In terms of total output for a given quantity of k and n , the rise in A causes the production function to rotate upwards around the origin when viewed in both $y-n$ space and in $y-k$ space, as we saw in our earlier study of the theory of the representative firm and investment demand. The marginal product of each factor of production is thus larger, holding all else constant. With perfectly competitive factor markets, this means that the price of each factor (in this case, the wage and rental rate) rises. Intuitively, the usual notion of the “law of supply” tells us that when price increases, supply will increase. The rise in capital and labor, coupled with the initial rise in A unambiguously causes total output to rise both in the present *and* the future.

In what follows, we use w to denote the real wage and r to denote the real rental rate of capital. Also, for simplicity, we assume that the labor tax rate is always zero, so that $t = 0$.

⁹⁵ Indeed, when Nobel-Prize-winning economist Robert Solow first proposed this way of measuring technology, he likened it to “a measure of ignorance,” because ultimately it is simply an unexplained (a “residual”) aspect of the production process, one that we do not understand. At the time, it was effectively just an accounting exercise. But RBC theory re-cast the Solow Residual as the centerpiece of a new view of macroeconomics.

⁹⁶ That is, no prices – and the real wage and the real rental rate are, after all, prices – are “sticky” at all.

⁹⁷ You may not think this is a simplification at all, given that in reality firms are usually thought of as “owning” their capital. Ultimately, however, it is the stakeholders of the firms that own the firm and hence the capital – in our theoretical model, this reduces to the representative consumer.

Technology, Factor Prices, and Output

With the production function $A \cdot f(k, n)$, the marginal product with respect to capital is given by $A \cdot f_k(k, n)$ and the marginal product of labor is given by $A \cdot f_n(k, n)$. Here, $f_k(k, n)$ and $f_n(k, n)$ denote the derivatives of the function f with respect to k and n , respectively. Notice that these derivatives are in general themselves functions of both k and n , as suggested by the notation. From our study of firms, we know that these marginal product functions determine the demand functions for labor and capital. The important feature now to consider is that changes in A shift these demand functions.

Suppose A rises suddenly, for example. Then, for any given quantity of labor and capital, the output function becomes steeper (that is, the slope increases) – which graphically is what it means for the marginal product to rise. This situation is depicted in **Figure 70**.




Figure 70. A rise in A_t causes output to rise for any given quantity of capital and labor.

Figure 71 illustrates the effect on the marginal product of labor from a different perspective, plotting the marginal product directly on the vertical axis, rather than leaving it implied by a plot of the output function as in **Figure 70**. The rise in A shifts the marginal product of labor outwards. A profit-maximizing firm will hire labor only to the point at which the marginal product equals the wage. For a given wage, a rise in A raises the profit-maximizing quantity of labor any given individual firm desires. But the wage itself will rise, as **Figure 72** illustrates. **Figure 72** shows the aggregate labor market. The labor demand function is a horizontal summation of each individual firm's demand function (which in turn is simply the marginal product function), and the labor supply function is that derived from the consumption-leisure model. Note that here we are assuming that the representative consumer is in the upward-sloping portion of the labor supply curve. As **Figure 72** shows, the equilibrium real wage rises, and the aggregate

quantity of labor hired increases. Thus, returning to **Figure 71**, the wage from the perspective of the price-taking representative firm also rises, but not enough to prevent it from, on net, hiring more labor than before the increase in A .

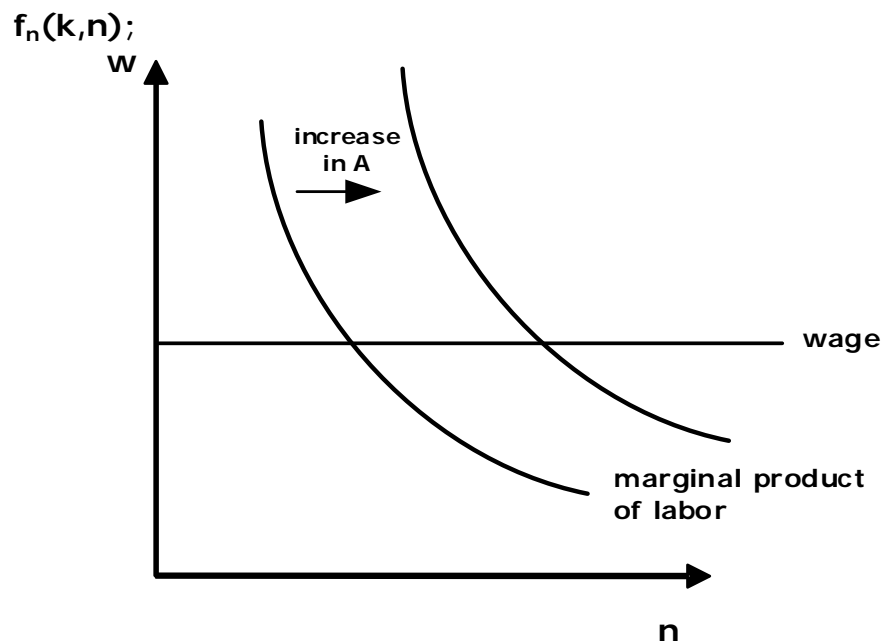


Figure 71. The marginal product of labor rises when A rises. For a given wage, the optimal quantity of labor demanded by a firm rises.

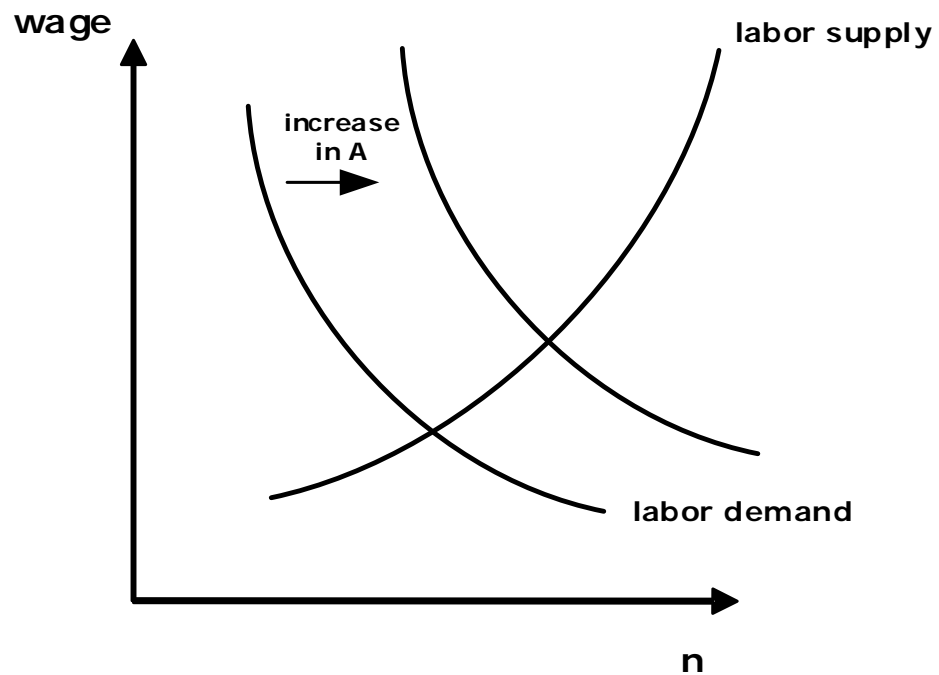


Figure 72. In the aggregate labor market, the equilibrium real wage rises when technology improves.

The effects on the market for capital are qualitatively similar. **Figure 73** shows that the marginal product of capital shifts out due to the rise in A . A profit-maximizing firm will hire (future) capital only to the point at which the marginal product equals the rental rate. For a given rental rate, a rise in A raises the profit-maximizing quantity of (future) capital any given individual firm desires. But the rental rate itself will rise, as **Figure 74** illustrates. **Figure 74** shows the aggregate market for savings and investment.⁹⁸ As **Figure 74** shows, the equilibrium rental rate rises, and the aggregate quantity of investment undertaken increases. Thus, returning to **Figure 73**, the rental rate from the perspective of the price-taking representative firm also rises, but not enough to prevent it from, on net, undertaking more investment than before the increase in A . **Figure 74** also shows that in funds-market equilibrium, the representative consumer saves more due to the rise in A .

⁹⁸ Because there is an increasing relationship (in fact, linear in the way that we studied firm behavior) between demand for future capital and current investment, we can make the jump from the market for (future) capital to the funds market of **Figure 74**.




Figure 73. The marginal product of capital rises when A rises. For a given rental rate, the optimal quantity of (future) capital demanded by a firm rises.

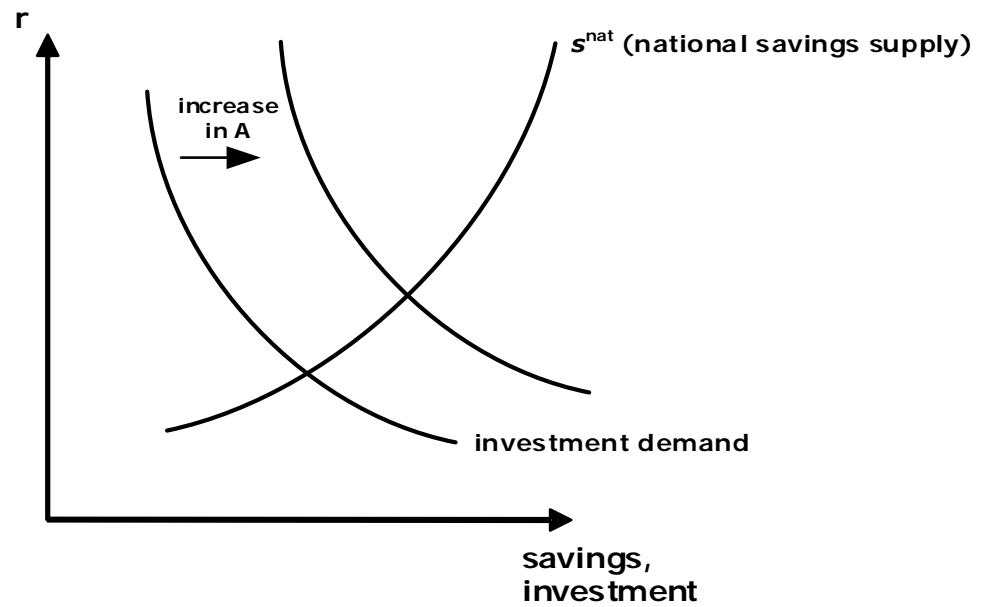


Figure 74. In the aggregate funds market, the equilibrium rental rate rises when technology improves.

Given all the effects of an increase in A that we have traced out, the effect on total output y is clear. The increase in A led to an increase in both *future* k and current n through its effects on factor prices. The function $k^\alpha n^{1-\alpha}$ is strictly increasing in both arguments, so total output $y = A \cdot k^\alpha n^{1-\alpha}$ unambiguously increases in both the current period as well as the future.

Now that we have traced out the aggregate effects, we examine the representative consumer's response along both the static consumption-leisure margin and the consumption-savings margin. This analysis actually becomes simple because we already know what the aggregate effects are – in this sense, the rest of the analysis is simply “looking under the hood.”

Effects on Consumption-Leisure Margin

The effects on the consumer's optimal choice along the consumption-leisure margin are particularly simple to describe – indeed, they are identical to what we have already studied. Recall from above that we are assuming that the economy is in the upward-sloping portion of the labor supply curve. This means that as the real wage rises and the budget constraint steepens as a result, the new optimal choice of (c, l) in the current period features higher consumption and *less* leisure.⁹⁹

Effects on Consumption-Savings Margin

The effects on the consumer's optimal choice along the consumption-savings margin are more subtle than we presented earlier, because, unlike in our earlier analysis, the consumer's income y (which equals GDP if consumers are the owners of both labor *and* capital) does not remain constant when the real rental rate increases.¹⁰⁰ We will frame our discussion here in terms of the real LBC of the consumer.

⁹⁹ The slope of the budget constraint in the static consumption-leisure model is, recall, $-(1-t)W/P$. In our discussion here we have assumed that the labor tax rate is zero; and recall that the real wage is simply the nominal wage divided by the aggregate price level W/P . Thus, a rise in the real wage causes the budget line to steepen.

¹⁰⁰ Recall from our earlier presentation of the simple consumption-savings model that the consumer had no control over his income – that is no longer the case here.




Figure 75. Initial optimal choice, before the rise in A . The slope of the LBC is $-(1+r)$ (that is, assume that there are no distortionary taxes).

Figure 75 shows the consumer's initial consumption-savings decision, before the improvement in technology which raises the rental rate. We assume no distortionary taxes, so the slope of the LBC is $-(1+r)$. As we have already traced out, the rise in A causes the rental rate r to rise. We know that this causes the LBC to steepen. However, rather than pivoting around the point marked (y_1, y_2) in **Figure 75**, the ordinates y_1 and y_2 both themselves increase because total output increases due to the rise in A .

Let's make this idea more precise. We've already seen above how and why a rise in A leads to a rise in total output. Suppose that the rise in A is purely temporary – it occurs unexpectedly in period 1, and then in period 2 A reverts to its normal value. Even though the rise in A occurs only for period 1, its effects on total output are felt in both periods 1 and 2 due to the effect on investment.

With consumers owning both labor *and* capital, total output in any given is paid to consumers. Total output is higher in period 1 due to improved technology *and* the rise in labor supplied, so the representative consumer's income in period 1, which we denote y_1 , rises. In period 2, total output, and hence the consumer's income, is higher, as well. The

reason that total output is higher in period 2, even though by then A has reverted to its previous value, is that the increased investment in period 1 means that the capital stock in period 2 is higher than it otherwise would have been. The increase in the capital stock means that period-2 output will rise as well. Graphically, the point through which the consumer's LBC must pass moves as shown in **Figure 76**.




Figure 76. The temporary rise in A in period 1 leads to higher income for the consumer in both period 1 and period 2.

Figure 77 then adds the LBCs to the diagram in **Figure 76**. Note that the new LBC is steeper because r is larger and passes through the point marked $(\text{new } y_1, \text{new } y_2)$ rather than the point marked $(\text{initial } y_1, \text{initial } y_2)$. That is, the LBC both shifts and rotates, rather than just rotating as in the simple consumption-savings model.




Figure 77. Following the rise in A and the resulting rise in r , the consumer's income rises in both period 1 and period 2. The new LBC is steeper than the initial LBC and passes through the new income point.

Finally, **Figure 78** illustrates how the optimal choice of consumption across time changes. We see that consumption rises in both period 1 and period 2. However, consumption in period 1 does not rise by as much as income rises in period 1 – we can conclude this because the horizontal distance between the new c_1^* and the initial c_1^* is smaller than the distance between the new y_1 and the initial y_1 .¹⁰¹ The consumer thus optimally spends only part of the gain in period-1 income and saves the rest for period 2, when he can again consume more than originally planned. The consumer thus smooths the gain in period-1 income over time – this illustrates the important principle of **consumption-smoothing**. The intuition for this result is reasonable – when faced with a rise in current

¹⁰¹ Actually, this need not always be the case – it actually depends critically on the shape of the indifference curves (i.e., it depends critically on the exact functional form of the utility function). We've illustrated here the most usual assumption made in RBC models, however.

income, individuals typically increase their current spending less than one-for-one with the rise in current income.¹⁰² With consumption-smoothing, private savings in period 1 rises – which we already knew from our analysis of the aggregate funds market above. Now we have also analyzed the same effect from the representative consumer’s perspective.

Putting it Together – Business Cycle Fluctuations

With the above descriptions of the effects of a change in technology, we are ready to understand how fluctuations in technology lead to the periodic ups and downs, termed business cycles, of the economy. A temporary rise in A makes the two factors of production, labor and capital, more productive on the margin. The increased productivity leads to increases in the real wage and the real rental rate, which induces both increased labor supply and increased private savings. In equilibrium in the funds market, increased savings means increased investment, which in turn means a higher *future* (i.e., period 2) capital stock. Thus, output (equivalently, income) rises in both period 1 and 2 – in the parlance of the RBC literature, the temporary (i.e., for only one period) technology shock leads to a persistent (i.e., for more than one period) change in output and consumption. To test your understanding of the basics of RBC theory, it is useful for you to trace out for yourself the effects of a temporary decline in A .

¹⁰² Another way to state this result, which should be familiar from introductory macroeconomics, is that the marginal propensity to consume (out of current income) is typically taken to be less than one. A one-dollar rise in current income thus leads to a less-than-one-dollar rise in current consumption.

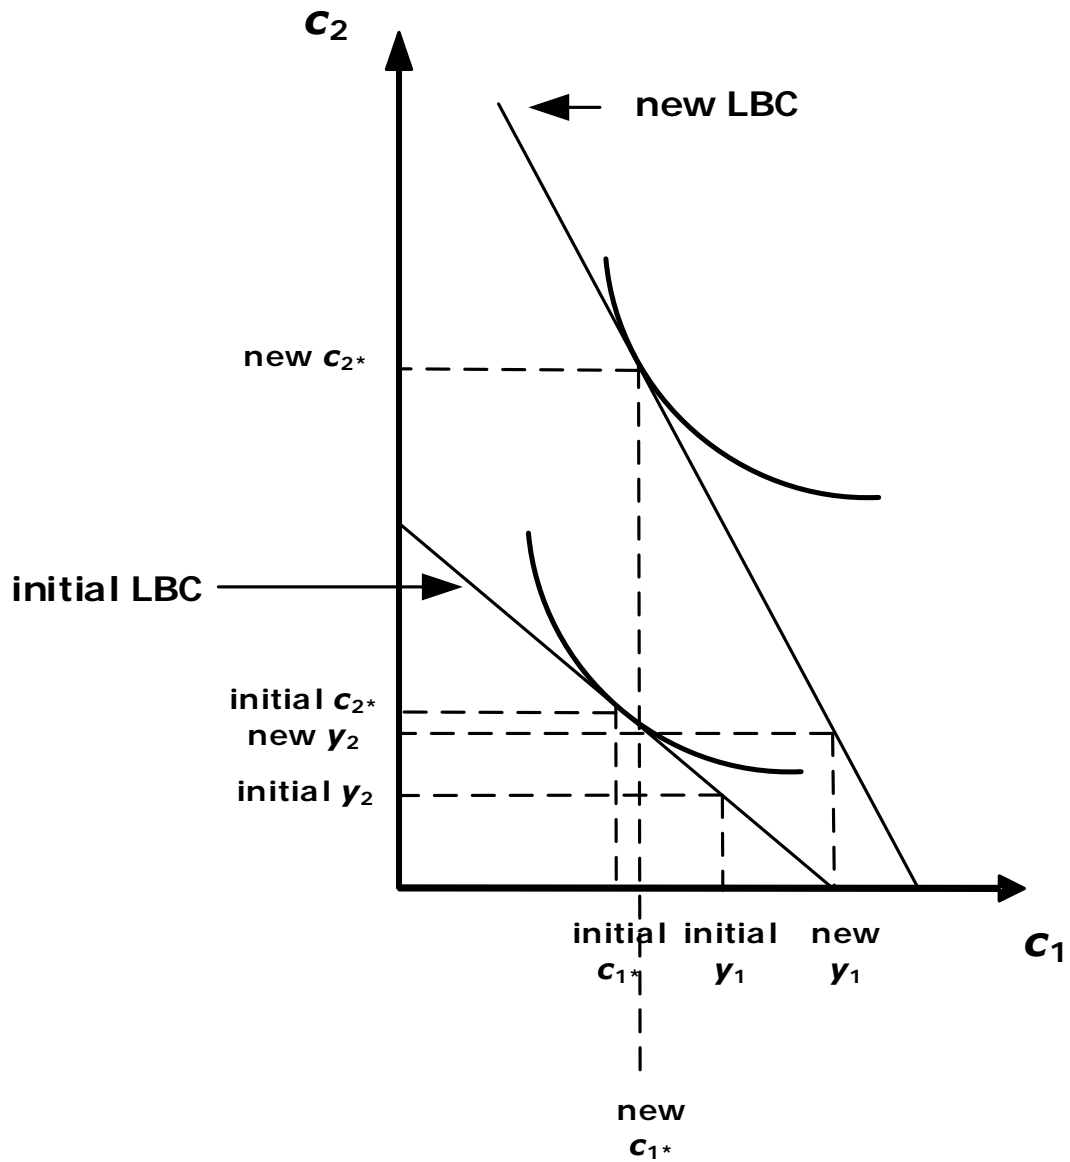


Figure 78. Optimal consumption in both period 1 and period 2 rise following the temporary rise in A.

