# Business Cycles

Jón Steinsson*
UC Berkeley

May 29, 2026

Economies have always been subject to fluctuations. The Genesis talks about seven years of famine following seven years of plenty (Genesis 41:53-54). Throughout most of history, economies were predominantly agricultural. Specialization and trade was relatively limited. In this era, economic fluctuations were mostly due to weather, disease, and warfare.

Over the course of the 18th and 19th centuries, the Industrial Revolution gradually transformed the economies of Western Europe, North America, and a growing list of other regions, raising living standards for a larger and larger fraction of their populations. Industrialization involved a gradual shift away from agriculture, towards industry (as the name suggests), but also towards trade, finance, and services.

Industrialization involved greater and greater division of labor. This meant that each person's work became more specialized and they relied more heavily on trade with others. As a consequence, the role of money and banking in the economy expanded enormously. Money and banking facilitated trade and were arguably an important contributor to the stunning economic growth that industrialization brought. However, money and banking were also new sources of economic fluctuations.

Heavy reliance on trade also meant that people's beliefs about the actions of others became central to their own decision making. Each person came to rely heavily other people's demand for the goods that they produced, and on other people to supply the goods that they wanted to buy. But what if these other people didn't do these things. More to the point, what if these other people didn't do these things because they were worried that yet others wouldn't do these things. Perhaps economic fluctuations could arise simply due to fluctuations in confidence.

*I would like to thank Joshua Zheng for excellent research assistance and Shigeru Fujita and Emi Nakamura for valuable discussions. First posted in May 2026.

Industrialization led to the rise of organized labor, which fought for better working conditions and higher pay. An important weapon of organized labor was the strike. In some cases, strikes were large enough or hit important enough nodes of the economy (such as coal production) that they resulted in aggregate fluctuations. Industrialization also led to a heavy reliance on fossil fuels (coal, oil, and natural gas). Disruptions in the supply of these—sometimes due to strikes but other times due to war or other forms of international disputes—have been a major source of economic fluctuations.

The dominant current of economic thought during the 18th and 19th centuries—a school of thought we now refer to as Classical Economics—emphasized the self-equilibrating nature of free-market capitalism. The economy was thought to "bob to the surface, like a cork held under water" of its own accord after a crisis, with output equilibrating at an appropriate level (Tarshis, 1987). Adam Smith argued that

> Little else is requisite to carry a state to the highest degree of opulence from the lowest barbarism, but peace, easy taxes, and a tolerable administration of justice; all the rest being brought about by the natural course of things. (quoted in Stewart, 1793)

This laissez faire view of economic policy was prevalent among economists. Another central pillar of Classical Economics was that the value of money should be tied to that of a commodity. Adherence to a commodity standard was seldom questioned and often taken to be unalterable.

The Great Depression dealt a massive blow to this set of beliefs. Output collapsed in many countries and unemployment soared to levels not seen before. Faith in the ability of market economies to self-equilibrate was badly shaken. A growing chorus argued that capitalism was doomed and some form of planned socialism was the way of the future. Arguments to this end were aided by strong growth in the Soviet Union under Stalin's five-year plans.

It was against this backdrop that John Maynard Keynes published his *General Theory of Employment, Interest, and Money* in 1936. In this book, Keynes presented a theory of aggregate economic fluctuations in which market economies did not self-equilibrate absent ongoing government intervention in the form of monetary and fiscal policy. Many of the elements of Keynes' theory had earlier antecedents in economic thought (Laidler, 1999, 2008). But Keynes' book catalyzed a major shift in economic thought towards greater emphasis on short-run fluctuations and government policy to help dampen these fluctuations. This shift is often referred to as the Keynesian Revolution and marks the birth of macroeconomics as a distinct subdiscipline of economics.

Keynes had been a critic of his colleagues' focus on self-equilibration in the "long run" even before the Great Depression. In 1923, he famously wrote:

> But this *long run* is a misleading guide to current affairs. *In the long run* we are all dead. Economists set themselves too easy, too useless a task if in tempestuous seasons they can only tell us that when the storm is long past the ocean is flat again. (Keynes, 1923, p. 80, italics in original)

Short-run fluctuations in economic activity are often referred to as *business cycles*. In this chapter, we introduce the notion of business cycles and discuss some of their basic features. The next few chapters then seek to understand business cycles through the lens of Keynesian theory.

# 1 Defining the Business Cycle

Business cycles are typically defined as fluctuations in economic activity around some notion of potential or trend. Figure 1 presents a stylized illustration. Potential output is trending upward in the figure and actual output is fluctuating around potential output. A business cycle is one such fluctuation, for example, from peak to peak.

The business cycle has two phases: an expansion phase and a contraction phase. The expansion phase is from the trough of the cycle (denoted by T in Figure 1) to the next peak (denoted by P in Figure 1). The contraction phase is from the peak to the next trough. The economy is said to be in *recession* when it is in the contraction phase of a business cycle, i.e, when it is contracting. This definition means that recessions last from peak to trough.

The difference between actual output and potential output is referred to as the *output gap*. In Figure 1, this is the vertical distance between the solid wavy line (output) and the broken trending line (potential output). The output gap is positive when output is above potential output and it is negative when output is below potential output.

Notice that periods of negative output gap and periods of recession do not coincide. If output peaks above potential output (as it does in Figure 1), the output gap

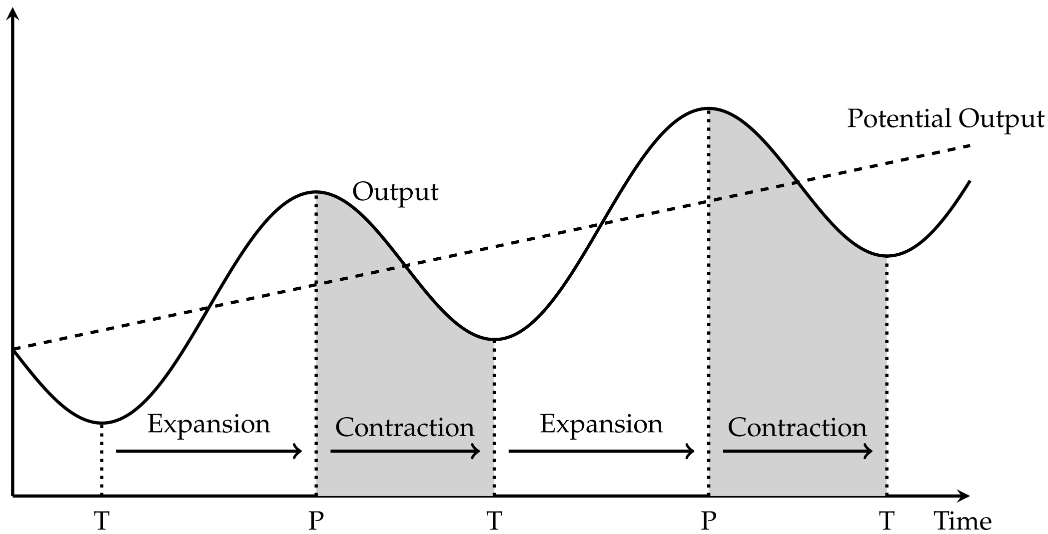

Figure 1: Stylized Business Cycles

will be positive in the early part of the contraction phase, i.e., at the beginning of a recession. Likewise, if output troughs below potential output (as it does in Figure 1), the output gap will be negative in the early part of the expansion phase, i.e., when the economy is no longer in recession.

## 1.1 The Duration of Business Cycles

Figure 1 is stylized in several different ways. First, business cycles differ substantially in duration. Table 1 lists the peaks and troughs of business cycles in the United States from 1927 to 2020. We can use these dates to calculate the duration of each expansion and contraction. This is done in columns 3 and 4 of the table.

We see that expansions vary a great deal in duration. Some are as short as one year, while others last longer than 10 years. The average duration of an expansion is about 5 years and the standard deviation of the duration of expansions is about 3 years. Contractions are much shorter than expansions on average. Their average duration is only about 1 year. Furthermore, the duration of contractions is much less variable. If we exclude the Great Depression (1929-1933)—a massive outlier in terms of duration—the standard deviation of the duration of recessions is only about 4 months.

The semi-official arbiter of peaks and troughs of the business cycle in the United States is the Business Cycle Dating Committee of the National Bureau of Economic Research (NBER). The dates listed in Table 1 are those determined by the NBER.

Table 1: U.S. Business Cycles

| NBER Dates | | Duration | | Size | |
|---|---|---|---|---|---|
| Peak | Trough | Exp. | Cont. | Exp. | Cont. |
| – | Nov 1927 | | | | |
| Aug 1929 | Mar 1933 | 21 | 43 | 3.0 | 23.3 |
| May 1937 | Jun 1938 | 50 | 13 | 17.1 | 4.9 |
| Feb 1945 | Oct 1945 | 80 | 8 | 12.4 | 3.9 |
| Nov 1948 | Oct 1949 | 37 | 11 | 1.4 | 4.5 |
| Jul 1953 | May 1954 | 45 | 10 | 5.4 | 3.6 |
| Aug 1957 | Apr 1958 | 39 | 8 | 2.4 | 3.8 |
| Apr 1960 | Feb 1961 | 24 | 10 | 2.7 | 2.3 |
| Dec 1969 | Nov 1970 | 106 | 11 | 3.7 | 2.7 |
| Nov 1973 | Mar 1975 | 36 | 16 | 1.5 | 4.4 |
| Jan 1980 | Jul 1980 | 58 | 6 | 3.4 | 2.2 |
| Jul 1981 | Nov 1982 | 12 | 16 | 0.6 | 3.6 |
| Jul 1990 | Mar 1991 | 92 | 8 | 5.8 | 2.8 |
| Mar 2001 | Nov 2001 | 120 | 8 | 4.0 | 2.5 |
| Dec 2007 | Jun 2009 | 73 | 18 | 1.9 | 5.6 |
| Feb 2020 | Apr 2020 | 128 | 2 | 6.5 | 11.3 |

*Note:* "Exp." is short for "Expansion", while "Cont." is short for "Contraction." The third column reports the duration of the expansion ending with the peak listed in that row in months. The fourth column reports the duration of the contraction ending with the trough listed in that row in months. The fifth column measures the peak to trough decrease in the unemployment rate associated with the expansion, while the sixth column measures the trough to peak increase in the unemployment rate associated with the contraction. The dates of the peaks and troughs of for the unemployment rate do not coincide exactly with the dates of the NBER peaks and troughs. I choose the peaks and troughs of the unemployment rate that best line up with the NBER business cycle dates. In most cases, I use the algorithm of Dupraz, Nakamura, and Steinsson (2005) to choose these dates. In some cases their algorithm and the NBER do not identify the same cycles. In these cases, I choose peaks and troughs of the unemployment rate manually. Business cycle dates are from the website of the NBER. Data on unemployment from 1948 onward is from the U.S. BLS, while data for unemployment before 1948 is from Petrosky-Nadeau and Zhang (2021).

When determining the timing of peaks and troughs, the NBER looks for "significant decline in economic activity that is spread across the economy and lasts more than a few months." But the three key criteria—depth, diffusion, and duration—can be met to varying degrees and some judgment is employed in the determination. For example, the Covid recession from February 2020 to April 2020 did not last "more than a few months" but it was sufficiently sharp for that part of the definition to be overridden.

The NBER considers a number of economy-wide measures of economic activity when determining whether a peak or trough has occurred. These include real personal income, employment, unemployment, consumption expenditures, and industrial production. There is no specific formula used. Others sometimes use simpler definitions to determine whether the economy is in recession, such as whether the economy has experienced two consecutive quarters of negative growth in GDP.

Importantly, the NBER defines contractions in terms of absolute declines in the level economic activity. There are various other ways contractions could be defined. For example, they could be defined in terms of declines in economic activity relative to a trend, or as periods when growth is lower than average (but not necessarily negative), or as periods when economic activity declines in per capita terms. There is no right or wrong way to define economic contractions. Different definitions will be more or less useful in different circumstances.

The NBER's approach has arguably been a useful approach for the U.S. economy for the past 100 years. But other definitions may be more useful for other countries and time periods. Consider a country with a very high average growth rate (such as Japan between 1950 and 1980 or China between 1980 and 2010). For such a country, even a very large reduction in growth will not count as a recession according to the NBER's definition as long as growth remains positive. Or consider a country experiencing a substantial decline in its population due to low fertility (something that is likely to occur in many wealthy countries in the coming decades). Such countries may be in recession according to the NBER's definition much of the time even if per capita growth is positive.

The NBER has dated business cycle peaks and troughs in the U.S. all the way back to 1854. In Figure 1, I only use the NBER dates from 1927 onward. The reason for this is that the methodology used by the NBER to date business cycles prior to 1927 was fundamentally different from the methodology used after that date. Romer (1994) has argued convincingly that NBER business cycle peaks and troughs prior to 1927 measure peaks and troughs in economic activity relative to trend as opposed to in levels. The pre-1927 NBER dates therefore measure something quite different from the post-1927 dates. This means that they are not comparable.

## 1.2 Do Economic Expansions Die of Old Age?

Why do expansions end? An enduring idea is that expansions end partly because of a buildup of imbalances and distortions that make the economy more and more frail as the expansion becomes longer. Some even view recessions as a necessary cleansing mechanism for the economy.

The notion that expansions become more frail as they age can be analyzed using a branch of statistics called survival analysis. The central concept in survival analysis is the hazard function of failure: the probability that failure occurs in the next period given survival up to that point. One example is the probability that a person dies over the next year conditional on having survived to a certain age. This hazard of dying varies with age. It is quite low for people that are young but rises as people get older. We say that the hazard function of death is upward sloping in age. In this sense, people become more frail as they age.

For economic expansions, failure is the onset of a recession (i.e., the "death" of an expansion). We are interested in whether the hazard of an economic expansion ending is increasing in the age of the expansion. In other words, we are interested in whether the hazard function for an expansion ending is upward sloping in the age of the expansion.

Rudebusch (2016) analyzes this question using the data on business cycle peaks and troughs in Table 1 for the period 1945 to 2016. He concludes that there is little evidence to support the notion that expansions become more likely to end the longer they last. In fact, he cannot reject the notion that the hazard function of expansions ending is completely flat. This conclusion lines up with earlier analysis by Diebold and Rudebusch (1990) and Sichel (1991).

One possible reason for Rudebusch's conclusion is that he has relatively little data to analyze. There were only 12 expansions between 1945 and 2016. Rudebusch, therefore, only had 12 data points to estimate the hazard function for the end of expansions. It is not surprising that this yields imprecise estimates. Analyzing a longer sample period would be valuable because it would be likely to yield more precise estimates. Romer's (1994) work, which we discussed above, suggests that 1927 may be a good starting point for this type of analysis since this is when the NBER settled on its current conception of how business cycle peaks and troughs should be dated. Starting in 1927 rather than 1945 adds three additional expansions to the sample. Extending the sample to 2026 adds a fourth additional expansion, bringing the total to 16.

Table 2 presents results using the sample period 1927 to 2026. Using the same relatively standard statistical model for hazard function estimation as Rudebusch (see the table note for details), I estimate that the slope of the hazard function of expansions ending is positive. In the model I estimate, the slope of the hazard func-

Table 2: Do Expansions Die of Old Age?

| | | |
|---|---|---|
| Shape Parameter ($k$) | 1.85 | [1.45, 2.88] |
| Probability of Recession within Next Year: | | |
| Expansion two years old | 13.6% | [4.7%, 23.5%] |
| Expansion four years old | 21.4% | [12.9%, 35.4%] |
| Expansion six years old | 28.0% | [20.8%, 49.7%] |
| Expansion eight years old | 33.7% | [25.1%, 63.8%] |
| Expansion ten years old | 38.9% | [28.2%, 74.5%] |

*Note:* These results are from a model in which expansions are assumed to have a Weibull distribution implying that survival to time $t$ is given by $S(t) = P(T > t) = \exp[-(t/\lambda)^k]$, where $T$ is the time of failure (business cycle peak) and $\lambda$ and $k$ are the scale and shape parameters, respectively. This model implies a hazard function $h(t) = f(t)/S(t) = (k/\lambda)(t/\lambda)^{k-1}$ where $f(t)$ is the density of failure. The hazard function is upward sloping if $k > 1$, flat if $k = 1$, and downward sloping if $k < 1$. The model is estimated on NBER business cycle expansions from 1927Q4 until 2026Q1. The last expansion is treated as left censored in March 2026. I report 95% confidence intervals in brackets. These are calculated using a weighted bootstrap where 500 reweighted replicate samples are drawn with weights drawn using the multinomial distribution.

tion is governed by a parameter that is called the shape parameter and is often denoted by $k$. The hazard function is upward sloping if this parameter is larger than one. I estimate a value of 1.85 for this parameter with a 95% confidence interval of $[1.45, 2.88]$. This implies that I can reject that the hazard function is flat ($k = 1$). (My point estimate of 1.85 is not very different from Rudebusch's. But my estimate of the confidence interval is smaller given my longer sample period.)

An upward-sloping hazard function implies that the probability that an expansion will end ('die') rises as the expansion becomes older. Table 2 reports the estimated probability that an expansion will end within one year for expansions of different age. My estimates imply that an expansion that is two years old will 'die' with 13.6% probability over the next year, while an expansion that is ten years old will 'die' with 38.9% probability over the next year. This result suggests that there is, in fact, some evidence that recessions become more frail as they age.

The statistical uncertainty associated with all of the estimates in Table 2 is quite large (the estimated confidence intervals are large). Despite this fact, the upward slope of the hazard function is statistically significantly different from zero, as noted earlier. I wonder, however, whether this is due to the fact that three of the expansions in the sample ended after roughly ten years (the 1961-1969, 1991-2001, and 2009-2020 expansions). Perhaps it is a bit of a fluke that our sample has three ex-

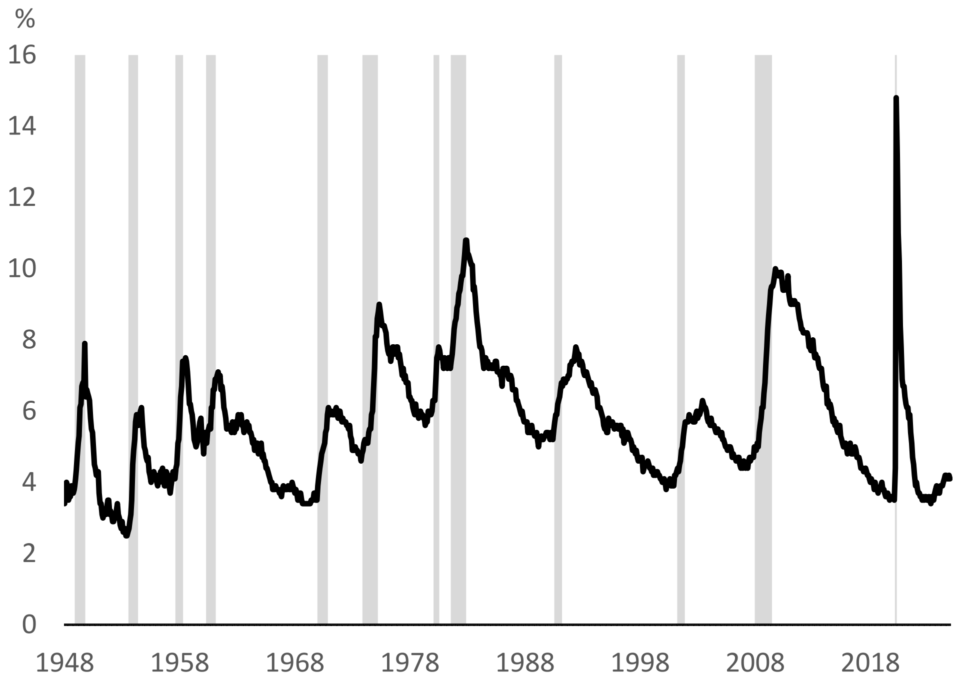


Figure 2: Unemployment in the United States, 1948-2024

*Note:* The figure plots the unemployment rate in the United States from 1948 to 2024. The source of these data are the Bureau of Labor Statistics. The shaded periods are NBER recessions.

pansions so closely clustered in age at a relatively high age but none that reached a higher age. Perhaps this is what is pushing the statistical model to estimate an upward sloping hazard, even though it is really just a fluke. This is hard to know.

## 1.3 The Size of Business Cycles

Figure 2 plots the evolution of the unemployment rate in the United States from 1948 to 2024. We see that the unemployment rate has experienced large swings over this period. These are the business cycles we have been discussing. The unemployment rate rises rapidly during recessions and falls (more slowly) during expansions. This cyclicality of the unemployment rate makes it a particularly useful statistic for visualizing business cycles.

A simple (if crude) measure of the size of recessions is the percentage point increase in the unemployment rate associated with the recession. The right-most column of Table 1 reports this measure for U.S. recessions since 1927. By this metric, business cycles vary a great deal in size. The Great Depression is far and away the largest recession. The unemployment rate rose by 23.3 percentage points over the

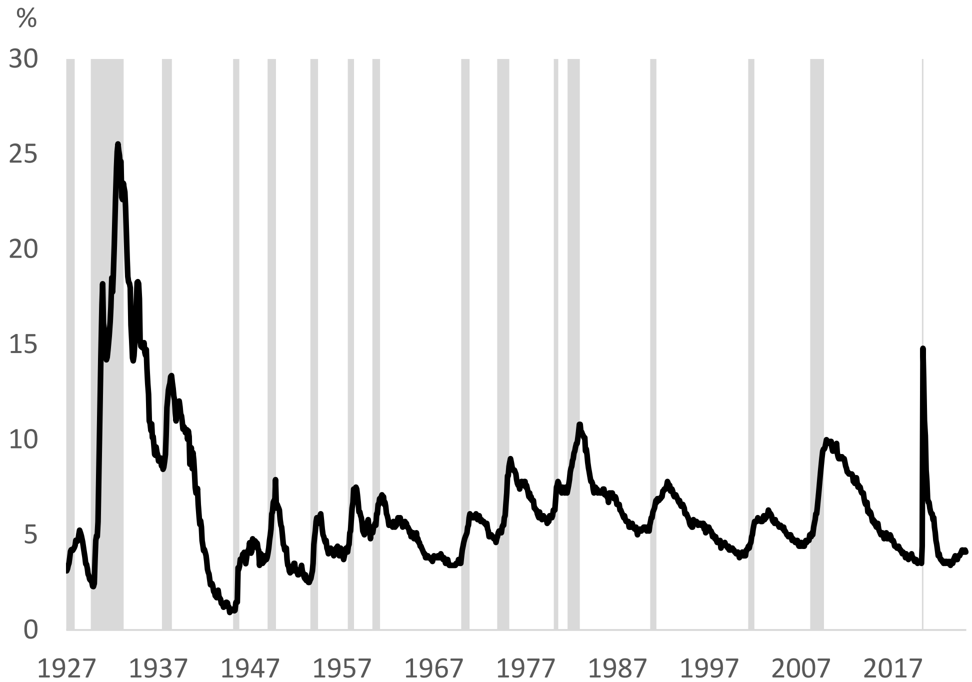


Figure 3: Unemployment in the United States, 1927-2024

*Note:* The figure plots the unemployment rate in the United States from 1927 to 2024. The source of the data from 1948 onward are the Bureau of Labor Statistics. The source of the data prior to 1948 is Petrosky-Nadeau and Zhang (2021). The shaded periods are NBER recessions.

course of the Great Depression from a low of 2.3% to a high of 25.5%. The Great Depression is so large that I have excluded it from Figure 2 because it would completely dominate the figure. Figure 3 extends the the sample period back to 1927 to include the Great Depression.

In the post-WWII period, the Covid recession of 2020 is by far the largest recession if measured peak to trough (trough to peak for the unemployment rate). However, the recovery after the Covid recession was unusually rapid. This made the period of high unemployment associated with Covid much more transient than for many other recessions.

Contrast the Covid recession with the (so called) Great Recession of 2007-2009. The increase in the unemployment rate over the course of the Great Recession was 5.6 percentage points. This makes it the third most severe recession by the peak-to-trough metric, but only half as large as the Covid recession. However, the unemployment rate stayed high for much longer and fell more slowly after the Great Recession.

Perhaps we should consider an alternative size metric for recessions that takes the persistence of high unemployment into account. A crude measure that does this is the integral of the area under the unemployment rate from its trough to the point where the unemployment rate reaches that level again. On this metric, the Great Recession is much larger than the Covid recession. The Great Recession led to roughly 26 percent-years of elevated unemployment (by this metric), while the Covid recession led to (only) about 6.5 percent-years of elevated unemployment. (This metric is harder to calculate for some earlier recessions since unemployment sometimes does not reach its earlier trough before a new recession hits the economy. This is true, for example, for the 1980 recession.)

## 1.4 The Seasonal Cycle

Another way in which Figure 1 is stylized is that it abstracts from seasonality in economic activity. The economy experiences substantial seasonal ups and downs every year. In contrast to business cycle fluctuations, these seasonal fluctuations are quite regular and can therefore be anticipated quite well. Employment usually falls sharply in January, then rises until May or June, falls over the summer, and then rises through the fall until November or December. Unemployment does the reverse. Interestingly, interest rates and asset prices do not experience seasonal fluctuations, and prices of goods and services experience only very mild seasonality. It is mostly the quantities of goods and services produced and consumed that experience seasonal cycles.

A number of factors contribute to seasonality in economic activity. Construction and agriculture tend to contract over the winter. Schools are are not in session over the summer. Holidays also play a role. For example, a disproportionately large fraction of consumer purchases occur in the weeks before Christmas, which leads firms to produce a disproportionate number of goods in the fall in anticipation of the Christmas shopping season.

The data we have considered up until this point in this chapter (and in this book more generally) has been *seasonally adjusted*. This means that an algorithm has been used to smooth out the seasonal cycle. There are different ways this can be done. A simple method is to regress the series in question on a set of month dummies and subtract the estimated coefficients on these month dummies from the observations for that month of the year. However, seasonality changes over time and the estimated coefficients on such month dummies can be influence by when in the year large recessions happen to occur. This makes seasonal adjustment somewhat

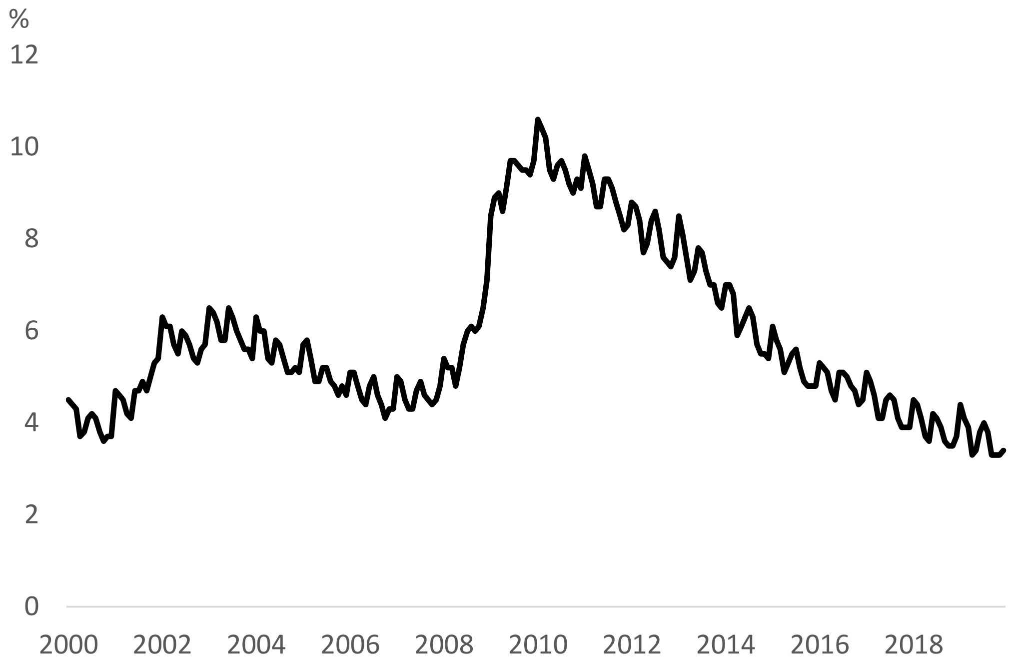


Figure 4: Unemployment Not Seasonally Adjusted in the United States, 2000-2019

*Note:* The figure plots the unemployment rate not seasonally adjusted in the United States from January 2000 to December 2019. The source of the data is the Bureau of Labor Statistics.

complicated and has led statistical agencies in different countries to develop more sophisticated algorithms for seasonal adjustment. The algorithm that is most commonly used in the United States (as of this writing) is called the X13 algorithm.

Figure 4 plots the raw (i.e, not seasonally adjusted) unemployment rate in the United States between January 2000 and December 2019. The seasonal cycle is evident with unemployment rising each winter and again (less so) each summer. Figure 5 provides a clearer view of the seasonal cycle for the unemployment rate. It plots the average difference between the seasonally adjusted and not seasonally adjusted unemployment rate in the United States by month for the sample period 1950 to 2019. The seasonal cycle in the unemployment rate is quite substantial. The difference between the peak and trough of the seasonal cycle is about 1 percentage point of unemployment.

## 2 The Output Gap

In section 1, we defined the output gap as the difference between output and potential output. It is useful to be a bit more precise and define the output gap as the

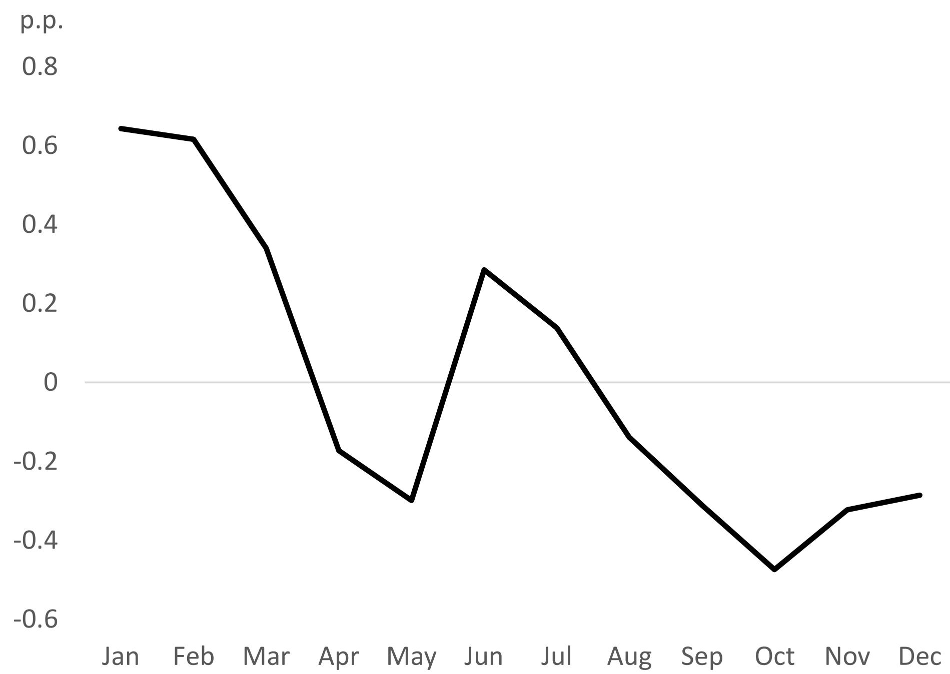

Figure 5: Seasonals in Unemployment in the United States

*Note:* The figure plots the average difference between the seasonally adjusted and not seasonally adjusted unemployment rate in the United States by month. The sample period is 1950 to 2019. The source of the data is the Bureau of Labor Statistics.

percentage difference between output and potential output. We denote the output gap as $\tilde{Y}_t$ (read 'Y-tilde'). Our definition implies that

$$\tilde{Y}_t = \frac{Y_t - Y_t^*}{Y_t^*}, \tag{1}$$

where $Y_t$ is output and $Y_t^*$ (read 'Y-star') is potential output.

The output gap is very tricky to measure because potential output is fundamentally unobservable. Roughly speaking, potential output is the level of output that would prevail if the economy were operating at, but not beyond, its natural capacity. Sometimes scholars use the term the *natural rate of output* to refer to roughly the same concept. Other terms also used in a similar fashion are the *efficient level of output* and the *full employment level of output*. Sometimes scholars also include in their definition of potential output the requirement that it be consistent with a stable inflation rate. We will come back to this idea in chapter XX [Phillips curve chapter].

Importantly, potential output is the level of output that would prevail in a counterfactual situation. Consider the following thought experiment: Suppose this-or-that problem had not occurred (e.g., the Great Depression or Covid) and the economy had instead hummed along at capacity over that period (i.e., at its "natural" or “efficient” or “full employment” rate) rather than doing what it actually did. What would the level of output have been in that case? That alternative (counterfactual) level of output is potential output.

Since we don’t observe the counterfactual situation (we only observe what actually occurred), potential output is fundamentally unobservable. This means it must be estimated, which is not an easy task. The economy is constantly affected by various developments. This means that it is difficult to identify times when output is clearly equal to potential output. Different scholars (and institutions) take different approaches to the problem of estimating potential output. Sometimes these different approaches yield quite different answers, and which approach is most appropriate can be quite controversial.

Perhaps the simplest approach to estimating potential output is to equate it with the trend in actual output. In this case, estimating potential output boils down to estimating a trend for output. While very simple in principle, this approach raises a number of practical and conceptual issues. First, there is the issue of what sample period to use to estimate the trend. This can matter quite a bit. For relatively short sample periods, it matters whether the start and end points are at business cycle peaks or troughs (or somewhere in between). Over longer sample periods, the trend can change.

Output per capita in the United States grew faster on average in the first few decades after World War II than it has since. Figure 6 plots log real GDP per capita in the United States for the sample period 1947 to 2025. I have also plotted two linear trend lines in this figure. One is fit through the first part of the sample (1947 to 1973) and extended on to 1995. The other is fit through the second part of the sample (1974 to 2025) and extended back to 1960. The two trend lines are clearly very different. The slope of the first is 2.4% per year, while the slope of the second is only 1.7% per year. This reflects the substantial reduction in trend growth of the U.S. economy that occurred in the 1970s. Doing the same exercise for countries in Western Europe or for Japan yields substantially larger trend breaks.

One approach to circumventing this problem is to fit a more flexible trend to output, i.e., a trend that is not completely linear. There are quite a few different methods for doing this. Two popular approaches are the Hodrick and Prescott (1980, 1997) filter (HP-filter for short) and the Baxter and King (1999) band-pass filter. Figure 7 plots HP-filtered log real GDP for the United States from 1950 to 2022.

An alternative is to use a production function approach to estimate potential output. This approach typically models output as depending on labor, capital, and

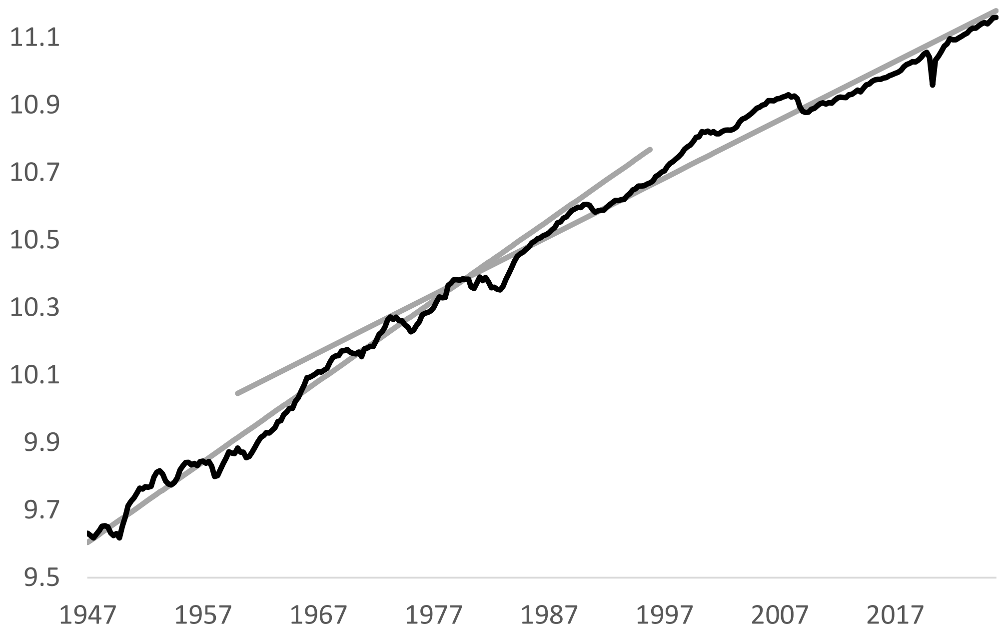

Figure 6: Log Real GDP per Capita in the United States

*Note:* The figure plots the logarithm of real GDP per capita in the United States from 1947 to 2025 along with two trend lines. The first trend line is estimated from 1947 to 1973 but plotted until 1995. The second trend line is estimated from 1974 to 2025 but plotted back to 1960. The source of the data is the Bureau of Economic Analysis.

total factor productivity (TFP). Estimates of potential output are constructed from estimates of potential levels of these inputs. For labor, this typically involves estimating the "natural" rate of unemployment and labor force participation. For capital it might involve adjusting for cyclical fluctuations in capacity utilization. Typically potential TFP is modeled as a smoothed version of realized TFP in this type of analysis.

The Congressional Budget Office (CBO) is an example of an institution that largely uses the production function approach to estimate potential output. Figure 7 presents the output gap implied by the CBO's measure of potential output. While not identical, the two output gap measures in Figure 7 are quite similar. This implies that the CBO's measure of potential output is similar to the HP-filtered trend in log real GDP most of the time. Both measure are highly negatively correlated with the unemployment rate.

The output gap measures presented in Figure 7 are retrospective in nature. This means that these entire series are constructed at the end of the sample period (in 2025 to be precise). For example, the CBO estimate of the output gap for 1975 that is plotted in Figure 7 is what the CBO believed in 2025 that the output gap in 1975 was.

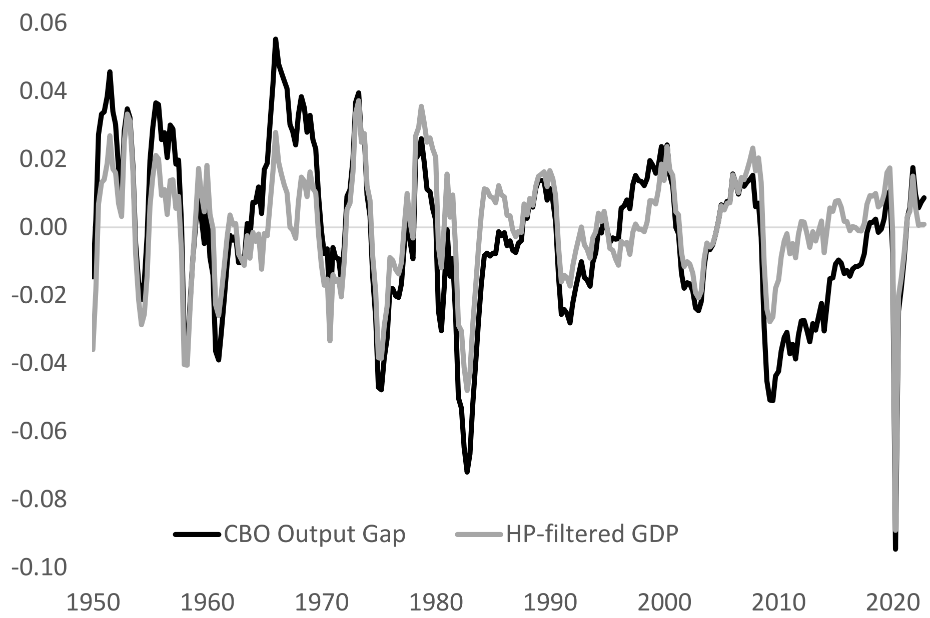


Figure 7: CBO Output Gap and HP-Filtered GDP

*Note:* The figure plots the CBO output gap (the difference between log real GDP and the logarithm of the CBO's estimate of potential output) and HP-filtered log real GDP (the difference between log real GDP and its HP-filter trend). I HP-filter with a smoothing parameter $\lambda = 1600$. The sample period is 1950 to 2022. The HP-filter was applied to real GDP data over the sample period 1947Q1 to 2025:Q2.

This may differ from the CBO's estimate in 1975 about the then current output gap. Estimates made in 2025 about 1975 benefit from hindsight: it is useful to know what subsequently occurred when estimating the output gap in 1975. Retrospective estimates also reflect improvements in estimation methods and changes in intellectual fads.

When studying past economic policy, it is important to have an accurate sense for what policy-makers knew (or thought they knew) *at the time* they were making policy. For example, when assessing the monetary policy of the Federal Reserve in 1980, it is important to consider what GDP, unemployment, inflation, and the output gap were estimated to be at that time. This may differ non-trivially from what we currently estimate GDP, unemployment, inflation, and the output gap to have been in 1980.

Economists have constructed *real-time* datasets to conduct this type of analysis. The primary real-time dataset used in macroeconomics for the United States is compiled by economists at the Federal Reserve Bank of Philadelphia. This dataset contains all vintages of GDP data and inflation data (and much more) several decades back in time. For example, all vintages of GDP data back to the 1965:Q4 vintage are available in this dataset. (Each such vintage then contains data back in time from that date.)

The Philadelphia Fed's real-time dataset also contains real-time estimates of the output gap constructed by the staff of the Board of Governors of the Federal Reserve. Estimates of the output gap are included in the briefing material the Fed Board staff prepared for each meeting of the Federal Open Market Committee, the committee that decides on monetary policy at the Fed. These estimate are often referred to as Greenbook estimates since the briefing book the staff prepare with information about the state of the economy and the staff's forecast was traditionally called the Greenbook (before it was merged with another staff report called the Bluebook to create the Tealbook in 2010).

Figure 8 plots retrospective and real-time Greenbook estimates of the output gap. We see that these estimate differ quite substantially before 1990 with the real-time estimate systematically indicating that the economy was weaker than the retrospective estimate. As we will discuss in more detail in chapter XX [Phillips curve chapter], the Federal Reserve's monetary policy in the 1970s is widely considered to have been too loose. One contributing factor was a belief at the time that output was far below potential. The retrospective estimates used today indicate that this belief was off by a very substantial amount.

## 3 The Natural Rate of Unemployment

We saw in the last section that estimating the output gap is quite tricky. This motivates using the unemployment rate as a gauge of the business cycle, as we did earlier in this chapter. The unemployment rate does not have a trend over long periods of time. This means that no detrending is needed when using the unemployment rate to think about the business cycle.

It, nevertheless, is useful to define a concept analogous to potential output for the unemployment rate. The most common term for this concept is the "natural rate" of unemployment, and the most common mathematical notation for it is $u^*$ (u-star). As with output, we can think of the natural rate of unemployment (roughly speaking) as the level of unemployment that would prevail if the economy were operating at, but not beyond, it natural capacity. Other terms that are sometimes used for this concept are the full-employment level of unemployment, the efficient level of

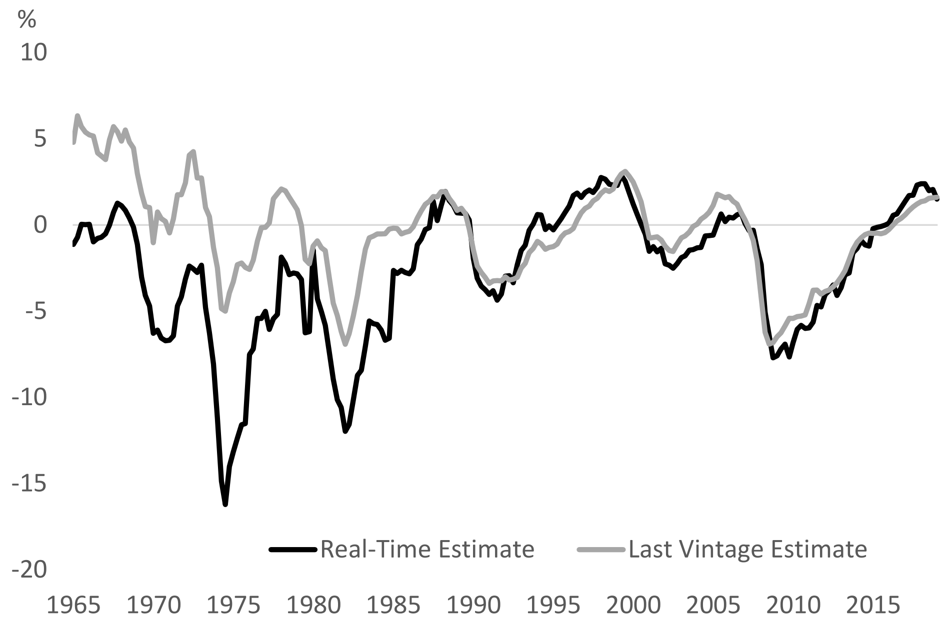


Figure 8: Real-Time and Retrospective Greenbook Output Gap Estimates

*Note:* The figure plots estimates of the output gap made by the Federal Reserve in its Greenbook (and Tealbook more recently). The figure plots both the real-time Greenbook estimate and a retrospective Greenbook estimate from the last vintage for the sample period. The sample period is 1965Q4 to 2019Q4. The source of these data are Edge and Rudd (2016) for data up to 2008Q4 and the Federal Reserve Bank of Philadelphia's Real-Time Data Research Center for more recent years.

unemployment, and the NAIRU (non-accelerating-inflation rate of unemployment).

Many students' first reaction when the idea of a natural rate of unemployment is introduced is: Shouldn't this be zero? Don't we want everyone to be employed who wants to work? Why would there be anything "natural" about unemployment? But this perspective is simplistic. It fails to appreciate the fact that there is naturally a great deal of churn in the labor market and it takes time for workers to find jobs.

Every month, some new workers enter the labor force, some fraction of employed workers are fired, and some fraction quits their jobs, for various reasons. These workers must search for new jobs. But they don't want just any job. They want to find a job that fits their interests and their level of skill. They want to find a job that isn't located too far away from where they live, and that pays well. Hence, finding a job takes time. Likewise, firms don't just want any worker. They want a worker that they think is well-qualified and a good match for their position.

For these reasons, the unemployment rate will not be zero even when the economy is doing very well. In other words, the natural rate of unemployment is not zero.

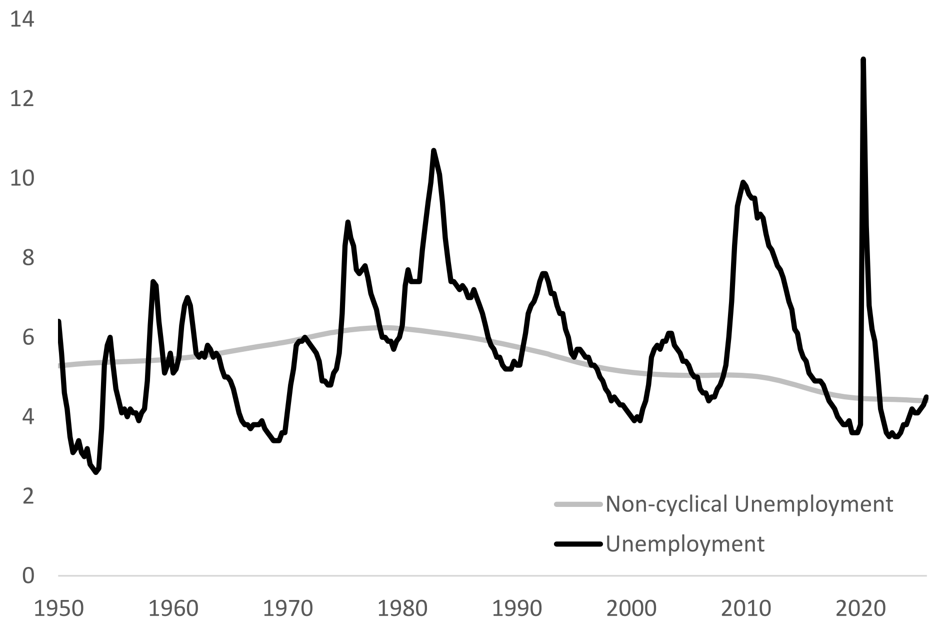


Figure 9: Unemployment and CBO Natural Rate Estimate

*Note:* The figure plots the U.S. unemployment rate and the CBO's estimate of the natural rate of unemployment. The CBO renamed this estimate the non-cyclical unemployment rate in 2021. Prior to that, it was called the "natural rate of unemployment (long-term)". The sample period is 1950 to 2025.

But how high is it? This is not entirely clear. As with potential output, the natural rate of unemployment is the unemployment rate in a counterfactual scenario. It is therefore fundamentally unobservable and can only be estimated.

Figure 9 plots an estimate of the natural rate of unemployment in the United States produced by the U.S. Congressional Budget Office (CBO) along with the actual unemployment rate. According to the CBO's estimate, the natural rate of unemployment has hovered between 5% and 6% for much of the last 75 years. At the end of the sample period (2025 Q4), the CBO estimates the natural rate to have been 4.4%. This was very close to the actual unemployment rate at that time (4.5%). According to the CBO, therefore, the U.S. economy was close to full employment at the end of 2025.

## 3.1 Why Does the Natural Rate of Unemployment Change?

While it is much smoother than the actual unemployment rate, the CBO's estimate of the natural rate of unemployment has varied by non-trivial amounts over time. It was 5.3% in 1950, gradually rose to 6.2% in the late 1970s, and has been falling since then. Interestingly, these movements correspond roughly to movements in the average actual unemployment rate. The average unemployment rate by decade was: 1950s: 4.5%, 1960s: 4.8%, 1970s: 6.2%, 1980s: 7.3%, 1990s: 5.8%, 2000: 5.5%, 2010: 6.2%. (The high average over the 2010s is, clearly, heavily influenced by the Great Recession.)

Why might the natural rate of unemployment have been higher in the late 1970s than in the 2020s or the 1950s? One reason is demographics. When workers are young, they are still finding their place in the workforce and learning about their skills and preferences for different jobs. They therefore change jobs more often than older workers. Young workers may also on average be less serious about keeping their jobs for various reasons. These factors imply that young workers tend to have a higher unemployment rate than older workers. As a consequence, periods when an unusually large fraction of the labor force is young will tend to have a higher natural rate of unemployment.

The late 1970s were a period with an unusually large fraction of young workers in the labor force. The reason for this was the *baby boom*: a sharp increase in fertility in the United States between 1940 and 1960. Figure 10 plots the total fertility rate in the U.S. from 1933 to 2024. U.S. fertility was about 2.1 in 1940 (down from about 7.0 in 1800). It then rose to 3.75 in 1960 before falling back to 2.0 in the early 1970s. This sharp increase and subsequent decrease in fertility between 1940 and 1970 is called the baby boom.

The unusually large cohorts from the baby boom started to enter the labor force in the 1960s and 1970s. This meant that by the late 1970s an unusually large fraction of the labor force was made up of young workers with higher than average propensities to be unemployed. Shimer (1998) argues that the aging of the baby boom generation can explain about a 0.75% fall in unemployment between the late 1970s and the late 1990s. This lines up well with the fall in the CBOs estimate of the natural rate of unemployment over this period.

### 3.2 The Plucking View of Business Cycles

In Figure 9, the unemployment rate fluctuates around the natural rate. These fluctuations are relatively symmetric. The unemployment rate is above the natural rate about half of the time and below the natural rate about half of the time. Also, the average natural rate (5.4%) is very close the average unemployment rate (5.7%), implying that the amplitude of the fluctuations above and below the natural rate are

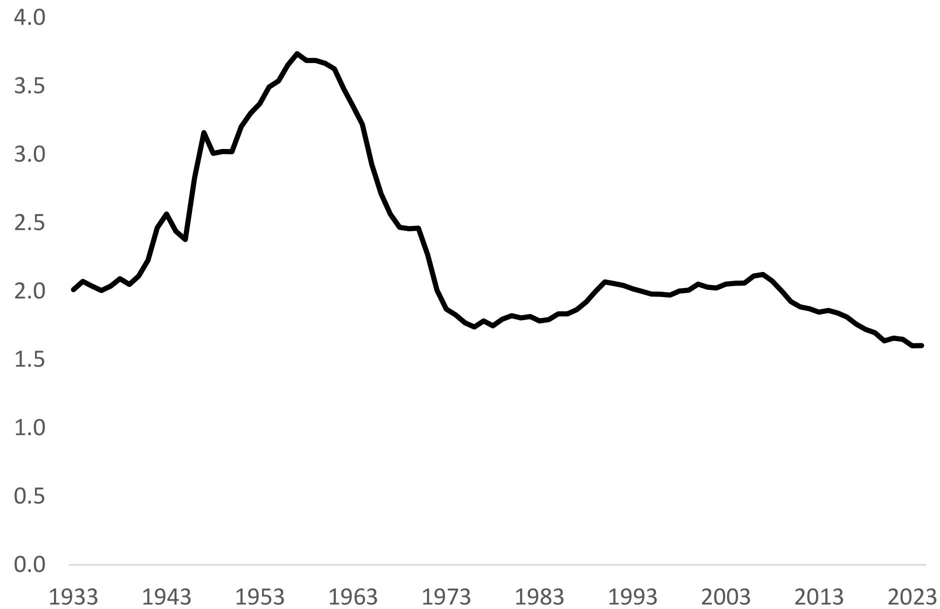


Figure 10: The Total Fertility Rate in the United States

*Note:* The figure plots the U.S. total fertility rate from 1933 to 2024. The total fertility rate is calculated by adding the fertility rates of women of different ages in a particular year. The source of these data is the Human Fertility Database.

close to being equally large.

These features make Figure 9 a good illustration of the conventional "natural rate view" of business cycles. This view holds that the unemployment rate and output fluctuate symmetrically around their natural rates.

An alternative to this natural rate view is the "plucking view" of business cycles. According to the plucking view, potential output is not an average level around which output fluctuates symmetrically but much closer to being a maximum attainable level, with fluctuations largely occurring below this level.

The plucking view draws its name from an analogy made by the economist Milton Friedman:

> Consider an elastic string stretched taut between two points on the underside of a rigid horizontal board and glued lightly to the board. Let the string be plucked at a number of points chosen more or less at random with a force that varies at random, and then held down at the lowest point reached. The result will be to produce a succession of apparent cycles in the string whose amplitudes depend on the force used in plucking the string. (Friedman, 1964)

In this analogy, the string is meant to represent output and the plucks are recessions.

Friedman argued that the plucking view could be distinguished empirically from the natural rate view by considering two characteristics of business cycles: 1) the correlation of the amplitude of a contraction and the amplitude of the subsequent expansion, and 2) the correlation of the amplitude of an expansion and the amplitude of the subsequent contraction. If the plucking view is correct, he argued:

> The cycles are symmetrical about their troughs; each contraction is of the same amplitude as the succeeding expansion. But there is no necessary connection between the amplitude of an expansion and the amplitude of the succeeding contraction. (Friedman, 1964)

In other words, the amplitude of a contraction should be highly correlated with the amplitude of the succeeding expansion, while no such correlation should exist between the amplitude of an expansion and the succeeding contraction. In a plucking model, this asymmetry arises because expansions are recoveries from earlier contractions (plucks), but the size of each pluck is independent from what came before.

If the conventional natural rate view is true, the asymmetry discussed above should not appear. This asymmetry also sheds like on the potential validity of a third view: that large booms sow the seeds of subsequent recessions. This third view implies the opposite asymmetry: that the amplitude of an expansion is correlated with the amplitude of the succeeding contraction, while the amplitude of a contraction is not necessarily correlated with the amplitude of the succeeding expansion.

Figure 11 presents two scatterplots. The left panel plots the amplitude of contractions on the horizontal axis and the amplitude of the subsequent expansion on the vertical axis. We see a strong positive correlation. The dashed line is a best fit line. This line has a slope of 1.06, which is highly statistically significant (the standard error is 0.18). The right panel plots the amplitude of expansions on the horizontal axis and the amplitude of the subsequent contraction on the vertical axis. Here, we see no relationship. Again, the dashed line is a best fit line. This line has a slope of 0.21, which is not statistically significantly different from zero (standard error of 0.39).

Figure 11 strongly supports the plucking view of business cycles. But why does this matter? One reason is that it suggests that the natural rate of unemployment is much lower than the CBO estimates it to be: it suggests that the natural rate is towards the very lowest or even below the very lowest unemployment rate we

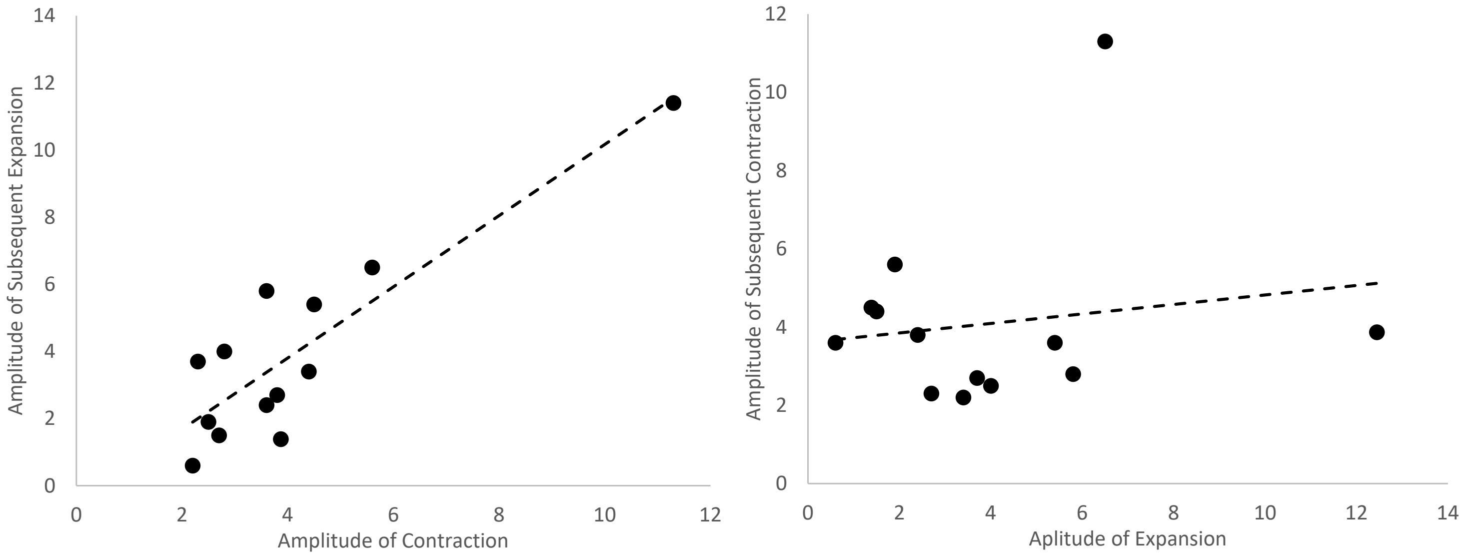


Figure 11: The Plucking Property of U.S. Business Cycles

*Note:* The data in this panel is drawn from Table 1 with the addition that the amplitude of the expansion after the Covid recession is added (using April 2023 as the peak). The sample period for the left panel starts with the contraction that began in February 1945, while the sample period for the right panel starts with the expansion that ended in February 1945.

have observed in the data. Unemployment reached a low of 3.8% in 1999, before the onset of the 2001 recession, a low of 3.5% in 2019, before the onset of the Covid recession, and a low of 3.4% in 2023. Perhaps the true natural rate of unemployment was not 4.4% in 2025 but 3.5% or even lower.

A characteristic of expansions that supports this view is that unemployment keeps falling during expansions until a recession occurs. Even towards the end of very long expansions (late 1960s, late 1990s, late 2010s) the unemployment rate kept falling. Were it to reach the natural rate, one might expect the decrease in the unemployment rate to peter out and the unemployment rate to start fluctuating around the natural rate. But this seems not to happen (with the possible exception of the last few years of the sample) which suggests that perhaps the unemployment rate has never actually reached the natural rate.

The plucking view suggests an alternative explanation for the high average unemployment rate in the 1970s and 1980s than the one given above having to do with demographics. If plucks (recessions) happen frequently, the economy will not have enough time between the plucks to fully recover. This will then lead the unemployment rate to ratchet up. The 1970s and early 1980s fit this explanation: several recessions occurred in relatively rapid succession (1969, 1973, 1980, 1981). Since unemployment usually falls quite a bit more slowly during expansions than it rises during recessions, the economy did not have enough time between these recessions to recover, and drifted further and further away from the natural rate. It was only with the advent of two long expansions (1980s and 1990s), with only a mild recession in between, that the unemployment rate was able to recover back to the levels seen in the 1950s and 1960s.

Another reason why it matters to know whether the plucking view or the natural rate view is correct is that the two views have quite different implications about the value of stabilization policy. According to the natural rate view, stabilization policy reduces the magnitude of booms and busts without affecting the average level of output. In other words, it reduces the variance of output, but doesn't change its mean. According to the plucking view, however, fewer and shallower recessions raise the average level of output (since recessions are shortfalls below a maximum). This means that stabilization policy can potentially result in much larger welfare gains if the plucking view is correct than if the natural rate view is correct.

# 4 Okun's Law

We have discussed two business cycle indicators in this chapter: the output gap and the unemployment rate. It will prove convenient over the next few chapters to be able to easily move between these two indicators. Sometimes it is most natural to discuss the business cycle in term of the output gap, while in other circumstances the unemployment rate is more natural.

In 1962, the economist Arthur Okun described a simple empirical regularity that links these two concepts in the United States. Using data from 1947 to 1960, Okun argued that when the unemployment rate fell by one percentage point the output gap rose by 3% (Okun, 1962). Later analysis has revised this estimate down. A common rule of thumb is that the output gap increases by 2% for each percentage point that unemployment falls. This 2-to-1 relationship is often referred to as Okun's Law. Mathematically, it takes the form

$$\tilde{Y}_t = -2(u_t - u^*), \tag{2}$$

where $u_t$ denotes the unemployment rate and $u^*$ denotes the natural rate of unemployment.

The simplest way to evaluate Okun's Law empirically is to consider the relationship between the change in GDP over some period and the change in the unemploy-

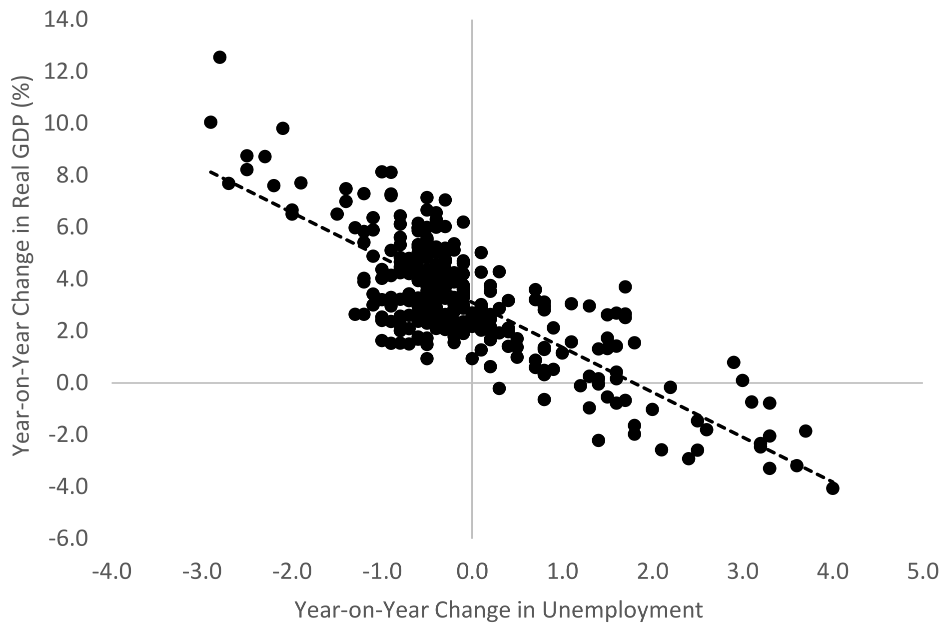


Figure 12: Okun's Law in the United States

*Note:* This plot has the year-on-year change in the unemployment rate on the horizontal axis and the year-on-year change in the natural logarithm of real GDP on the vertical axis. These data are from the U.S. Bureau of Labor Statistics and the U.S. Bureau of Economic Analysis, respectively.

ment rate over the same period. This allows one to avoid having to use estimates of potential output and the natural rate of unemployment. These concepts are "differenced out" when one considers Okun's Law in terms of changes over time.

Figure 12 plots Okun's Law (in terms of changes) for the period 1950 to 2019. There is clearly a fairly tight negative relationship between the change in unemployment and the change in output. Over this sample period, the best-fit line though the points has a slope of -1.75. This is a slightly smaller slope than the traditional rule of thumb of 2-to-1. When we evaluate the relationship over the period 1950 to 1999, we get a best-fit slope of -1.93. For simplicity, we will use a coefficient of -2.0 for Okun's Law in the rest of the book.

## 4.1 Why is the Okun's Law Coefficient So Large?

From a simple-minded production function perspective, the size of the coefficient in Okun's Law may seem surprisingly large. Consider the Cobb-Douglas production function $Y_t = A_t K_t^a L_t^{1-a}$. If we take the natural logarithm of this production function and consider the difference between the resulting equation at time $t$ and time $-1$ we get

$$\ln Y_t - \ln Y_{t-1} = (\ln A_t - \ln A_{t-1}) + a\,(\ln K_t - \ln K_{t-1}) + (1-a)\,(\ln L_t - \ln L_{t-1})$$

or equivalently

$$\Delta \ln Y_t = \Delta \ln A_t + a\Delta \ln K_t + (1-a)\Delta \ln L_t,$$

where $\Delta$ refers to a change over time, e.g., $\Delta \ln Y_t = \ln Y_t - \ln Y_{t-1}$. If we furthermore equate $L_t$ with employment, we get that $L_t = 1 - U_t$, where we assume for simplicity that the size of the labor force is 1. This means that $\Delta \ln L_t = \Delta \ln(1-U_t)$. Since $U_t$ is relatively small and $Y_t$ is relatively close to $Y_{t-1}$, we can approximate $\ln(1-U_t) \approx -U_t$ and $\Delta \ln Y_t \approx (Y_t - Y_{t-1})/Y_{t-1}$. Using these expressions, we can rewrite the above equation as

$$\frac{Y_t - Y_{t-1}}{Y_{t-1}} \approx \Delta \ln A_t + a\Delta \ln K_t - (1-a)\Delta U_t. \tag{3}$$

One way to read this equation is that it suggests that an increase in demand in the economy that increases output and decreases the unemployment rate should give rise to an Okun's Law type relationship but with a coefficient equal to $1-a$. Recall from Chapter XX [Production Chapter] that $1-a$ is equal to the labor share of income, which is on the order of 2/3 (or a bit less). In other words, a simple reading of equation (3) suggests that the coefficient in Okun's Law should be 0.67 (or less) rather than 2.0 (or 1.75), a pretty large discrepancy.

Why is the simple framework of equation (3) so far off when it comes to Okun's Law? There are quite a few possible reasons for this. Each of these has to do with a particular way in which firms can change their level of production by other means than a change in employment. Hiring and firing workers in response to changes in demand is costly both for firms and for workers. Workers value jobs with stable employment and firms would like to avoid the costs of hiring and firing.

This logic suggests that firms will limit changes in employment to some extent over the business cycle. When demand is low, workers will work fewer hours and the hours they work will involve less effort (there will be less for them to do). On the other hand, when demand is high, workers will work more hours and put in more effort to meet the high demand. This type of behavior by firms when it comes to hiring and firing over the business cycle is sometimes referred to as labor hoarding.

In addition, a strong economy can draw more people into the labor force. The simple math above assumed that the labor force was constant over the business cycle. This may not be the case. Likewise, the population may vary with output.

A larger population results in a larger economy and more output. But a strong economy also induces more immigration. So, causation likely goes both ways when it comes to the population. And then there are the other two terms in equation (3)—productivity and capital. They may covary with output in ways that result in larger movements in output for a given movement in employment.

Fujita, Ramey, and Roded (2026) propose a useful way to account for the role of these different factors. Consider first that the amount of labor (hours worked) in an economy $L_t$ is equal to employment $E_t$ times hours worked per worker $H_t$. That is, $L_t = H_t E_t$. Employment can be decomposed as follows:

$$E_t = \left(\frac{E_t}{LF_t}\right)\left(\frac{LF_t}{N_t}\right) N_t,$$

where $LF_t$ denotes the labor force and $N_t$ denotes the population. This equation then says that employment is equal to the employment rate $E_t/LF_t$ times the labor force participation rate $LF_t/N_t$ times the population. The employment rate is equal to 1 minus the unemployment rate: $E_t/LF_t = 1 - U_t/LF_t = 1 - u_t$, where $u_t$ is the unemployment rate. Finally, output is equal to output per hour worked times hours worked: $Y_t = \frac{Y_t}{L_t} L_t$. Putting all of this together, we have that

$$Y_t = \left(\frac{Y_t}{L_t}\right) H_t \left(1 - u_t\right) \left(\frac{LF_t}{N_t}\right) N_t. \tag{4}$$

This equation is an identity. It can therefore be used to decompose changes in output into different factors in an accounting sense. We take a natural logarithm of equation (4):

$$\ln Y_t = \ln\left(\frac{Y_t}{L_t}\right) + \ln H_t + \ln\left(1 - u_t\right) + \ln\left(\frac{LF_t}{N_t}\right) + \ln N_t.$$

We then subtract the time $t-1$ version of this equation from the time $t$ version, write the results in terms of changes

$$\Delta \ln Y_t = \Delta \ln\left(\frac{Y_t}{L_t}\right) + \Delta \ln H_t + \Delta \ln\left(1 - u_t\right) + \Delta \ln\left(\frac{LF_t}{N_t}\right) + \Delta \ln N_t.$$

Next we make use of the fact that small changes in the natural logarithm of a variable are approximately equal to the growth rate of the variable and the fact that $\Delta \ln(1 - u_t)$ is approximately equal to $-\Delta u_t$ (these are first-order Taylor approximations). This then yields

$$g_{Y,t} = g_{Y/L,t} + g_{H,t} - \Delta u_t + g_{LF/N,t} + g_{N,t}, \tag{5}$$

where we use $g_{X,t}$ to denote the growth rate of variable $X$ and time $t$.

We next take the covariance of both sides of equation (5) with $\Delta u_t$ and use the fact that $\text{Cov}(\Delta u_t, \Delta u_t) = \text{Var}(\Delta u_t)$ to get:

$$\begin{aligned}\text{Cov}\left(g_{Y,t}, \Delta u_t\right) = \; & \text{Cov}\left(g_{Y/L,t}, \Delta u_t\right) + \text{Cov}\left(g_{H,t}, \Delta u_t\right) - \text{Var}\left(\Delta u_t\right) \\ & + \text{Cov}\left(g_{LF/N,t}, \Delta u_t\right) + \text{Cov}\left(g_{N,t}, \Delta u_t\right).\end{aligned}$$

Finally, we divide through by $\text{Var}\left(\Delta u_t\right)$ and rearrange to arrive at:

$$\begin{aligned}\frac{\text{Cov}\left(g_{Y,t}, \Delta u_t\right)}{\text{Var}\left(\Delta u_t\right)} = & -1 + \frac{\text{Cov}\left(g_{Y/L,t}, \Delta u_t\right)}{\text{Var}\left(\Delta u_t\right)} + \frac{\text{Cov}\left(g_{H,t}, \Delta u_t\right)}{\text{Var}\left(\Delta u_t\right)} \\ & + \frac{\text{Cov}\left(g_{LF/N,t}, \Delta u_t\right)}{\text{Var}\left(\Delta u_t\right)} + \frac{\text{Cov}\left(g_{N,t}, \Delta u_t\right)}{\text{Var}\left(\Delta u_t\right)}.\end{aligned} \tag{6}$$

The left-hand side of equation (6) is the formula for the regression coefficient when output growth is regressed on the change in unemployment. In other words, it is the formula for the slope of the best-fit line through the cloud of points in Figure 12. The right-hand side of equation (6) then says that the Okun's Law slope is equal to -1 plus four other regression coefficients from regressions of the growth rates of output per hour, hours worked per worker, the labor force participation rate, and population on the change in unemployment. (In practice this will not work perfectly because the data available for these different variables are not fully comparable.)

Figure 3 presents the Fujita-Ramey-Roded decomposition from equation (6) for the period 1950-2019 and for the first and second half of this period. Let's start from the bottom of the table. The coefficients on both labor force participation and population are quite close to zero. This is true for all three sample periods. It seems that neither the labor force or population is very cyclically sensitive. These variables therefore do not contribute much to the size of the Okun's Law coefficient in this sample period.

The coefficient on hours worked per worker is quite a bit larger in absolute size. This coefficient is negative, implying that hours worked per worker tend to rise when unemployment falls (i.e., in economic upswings). We say that hours worked per worker are pro-cyclical (increase in an upswing). The coefficient is rather stable across these three subsample at roughly -0.40. So, it contributes substantially to the size of the Okun's Law coefficient.

That brings us to the coefficient on output per hour. For the whole sample, this coefficient is very close to zero. In other words, output per hour (also referred to as labor productivity) has been virtually acyclical on average over our 70 year sample

Table 3: Fujita-Ramey-Roded Decomposition

| | 1950-2019 | 1950-1984 | 1985-2019 |
|---|---|---|---|
| Okun's Law | -1.75 | -2.00 | -1.27 |
| | (0.08) | (0.08) | (0.10) |
| Output per Hour | 0.02 | -0.28 | 0.62 |
| | (0.08) | (0.10) | (0.09) |
| Hours per worker | -0.39 | -0.37 | -0.42 |
| | (0.02) | (0.04) | (0.04) |
| Labor force participation rate | -0.10 | -0.08 | -0.16 |
| | (0.03) | (0.04) | (0.05) |
| Population | 0.07 | 0.10 | 0.04 |
| | (0.03) | (0.03) | (0.03) |

*Note:* The table presents the Okun's Law regression coefficient as well as the four regression coefficients on the right-hand side of equation (6). In all cases, the year-on-year change in the unemployment rate is the regressor. The dependent variable for the Okun's Law regression is the year-on-year percentage change in GDP. The dependent variables for the other four regressions are the year-on-year percentage change in the variable listed in the left-most column. Adding the bottom four coefficients and subtracting one does not yield the top coefficient as equation (6) would suggest because the variables in the different regressions are drawn from different datasets that are not fully comparable and because of sampling error in the data. The sample period for each case is listed at the top. Standard errors are in parentheses.

period. This acyclicality, however, masks quite a bit of heterogeneity across subsamples. For the 1950-1984 sample, the coefficient is -0.28. This means that labor productivity was quite pro-cyclical in the early post-WWII period. In the latter half of our sample, however, the coefficient is 0.62, implying that labor productivity was strongly counter-cyclical from the mid-1980s until 2019. We also see that the Okun's Law coefficient is substantially smaller for this more recent sample period.

When we think about the contribution of labor productivity to the Okun's Law coefficient, it is also important to consider what the baseline should be. If equation (3) is our baseline, labor productivity should have a coefficient of about 1/3. That would result in an Okun's Law coefficient of 2/3 with the other three coefficients in equation (6) being zero. So, even a coefficient of zero, as occurs over the entire sample, implies that some combination of $A_t$ and $K_t$ are pro-cyclical.

Fujita, Ramey, and Roded consider a longer sample period than I have considered including the 1940s and the period around Covid. They find very large deviations from an Okun's Law coefficient of 2.0 both in the 1940s and around Covid. For the 1940s, they estimate an Okun's Law coefficient of -4.3, implying that unemployment moved less than half as much as normal when compared to movements in GDP. Why did unemployment move so little in this case? They point to demobilization after World War II as having been very unusual with both hours and labor force participation dropping substantially as the economy demobilized. Around Covid, however, they estimate an Okun's Law coefficient of only -1.1. In this case, labor productivity was highly counter-cyclical (fell substantially during the pandemic and recovered as the economy recovered) while labor force participation was highly pro-cyclical (many people left the labor force temporarily during the pandemic). These cases remind us that Okun's Law is not a law of nature, even if it is a useful regularity during most times.

# 5 Our First Business Cycle Model

In chapter XX [Quantity Theory], we developed our first (and simplest) monetary business cycle model. I refer to this model as the medieval economy. The (pedagogical) idea is that this model might represent the economy back in the middle ages when the only medium of exchange was gold coins. The model has two equations. The first of these is the quantity equation:

$$M_t V = P_t Y_t, \tag{7}$$

where $M_t$ denotes the money supply (gold coins), $V$ denotes the "velocity" of money (how often money changes hands), $P_t$ denotes the price level, and $Y_t$ denotes output. The other equation of the medieval economy model is the price setting equation:

$$\frac{P_{t+1}}{P_t} = \left(\frac{Y_t}{Y^*}\right)^{\theta}, \tag{8}$$

where $Y^*$ denotes the desired level of output and $\theta$ is a parameter representing the sensitivity of price changes to the gap between output and the desired level of output. In chapter XX [Quantity Theory], we used this model to explore how changes in the money supply affect output and prices in the short run and the long run.

## 5.1 Allowing Potential Output to Change Over Time

We now augment this model slightly by allowing desired output—which we now also refer to as potential output and the natural rate of output—to change over time. Recall from chapter XX [Quantity Theory] that we defined desired output as the output that results when people work the amount they desire: $Y^* = AL^*$. In that chapter, we assumed, for simplicity, that productivity $A$ was constant. Now we allow productivity to changes over time. In this case, potential output will also change over time: $Y_t^* = A_t L^*$.

This means that the price setting equation in our medieval economy becomes

$$\frac{P_{t+1}}{P_t} = \left(\frac{Y_t}{Y_t^*}\right)^\theta. \tag{9}$$

The only difference versus equation (8) is that $Y_t^*$ replaces $Y^*$.

The simplest way in which potential output $Y_t^*$ can change over time is for it to have a trend. This is what we assumed in Figure 1 and this is what is implicitly assumed when the output gap is computed by removing a trend, as in Figure 6. We will often make this simplifying assumption.

However, if productivity does not increase smoothly along a trend, potential output will not increase smoothly along a trend. In this case, some portion of the business cycle will be due to fluctuations in productivity. Since changes in productivity result in changes in desired output, it will be desirable for actual output to fluctuate with potential output. In this sense, some business cycle fluctuations may be optimal.

The desired level of labor supply may also fluctuate over time. One example of this is the Covid recession: it was optimal for economic activity to contract during the Covid pandemic since economic activity involved people interacting and people wanted to avoid interacting so as not to contract the virus. In other words, people's desired labor supply contracted during Covid. Another possible reason why desired labor may fluctuate is fluctuations in productivity. If productivity is temporarily low, people may want to temporarily reduce their labor supply *because* productivity is low (since low productivity makes it not a good time to work).

There is a school of thought within macroeconomics that argues that most economic fluctuations are due to productivity shocks and the resulting variation in desired output. This branch of the field is called real business cycle theory. If real business cycle theory is correct, most business cycle fluctuations are optimal. Furthermore, the monetary and fiscal policies that we will study in the next few chapters are counterproductive since they aim to stabilize output, while it is optimal for output to fluctuate. Relatively few macroeconomists adhere to this viewpoint. In its strongest form, it holds that the large fall in output during the Great Depression was optimal, i.e., that the 25% unemployment rate during the Great Depression was a Great Vacation.

## 5.2 Supply and Demand in the Medieval Economy

Can we use the concepts of supply and demand to think about the medieval economy model? Let's start with the quantity equation—equation (7). If we take the natural logarithm of the quantity equation and consider the difference between the resulting equation at time $t$ and time $t-1$ we get

$$\Delta \ln M_t = \Delta \ln P_t + \Delta \ln Y_t. \tag{10}$$

We can then manipulate $\Delta \ln P_t$ as follows:

$$\Delta \log P_t = \log\left(\frac{P_t}{P_{t-1}}\right) = \log(1+\pi_t) \approx \pi_t, \tag{11}$$

where the last step is a first-order Taylor series approximation (valid when $\pi_t$ is relatively small). And we can manipulate $\Delta \ln Y_t$ as follows:

$$\begin{aligned}\Delta \ln Y_t &= \ln Y_t - \ln Y_{t-1} + \ln Y_t^* - \ln Y_t^* + \ln Y_{t-1}^* - \ln Y_{t-1}^* \\ &= \ln\left(\frac{Y_t}{Y_t^*}\right) - \ln\left(\frac{Y_{t-1}}{Y_{t-1}^*}\right) + \ln\left(\frac{Y_t^*}{Y_{t-1}^*}\right) \\ &\approx \tilde{Y}_t - \tilde{Y}_{t-1}.\end{aligned} \tag{12}$$

where the last step uses a Taylor series approximation and drops the $\ln\left(Y_t^*/Y_{t-1}^*\right)$ term. Think of $\ln\left(Y_t^*/Y_{t-1}^*\right)$ as a small constant. We could keep it. But it would not contribute anything interesting to the analysis, and it would clutter the math. So, we drop it.

If we plug equations (11) and (12) into equation (10) we get

$$\Delta \ln M_t = \pi_t + \tilde{Y}_t - \tilde{Y}_{t-1}. \tag{13}$$

We can rewrite this equation as

$$\pi_t = -\tilde{Y}_t + \tilde{Y}_{t-1} + \Delta \ln M_t. \tag{14}$$

Here the quantity equation is written so that it gives inflation $\pi_t$ as a function of the current output gap $\tilde{Y}_t$, last period's output gap $\tilde{Y}_{t-1}$, and the change in the money supply $\Delta \ln M_t$.

We can plot the points that satisfy this equation in $(\tilde{Y}_t, \pi_t)$ space. This is the black curve in Figure 13. This curve is downward-sloping: the higher is the output gap, the lower is inflation. The reason for this is that higher output implies more demand

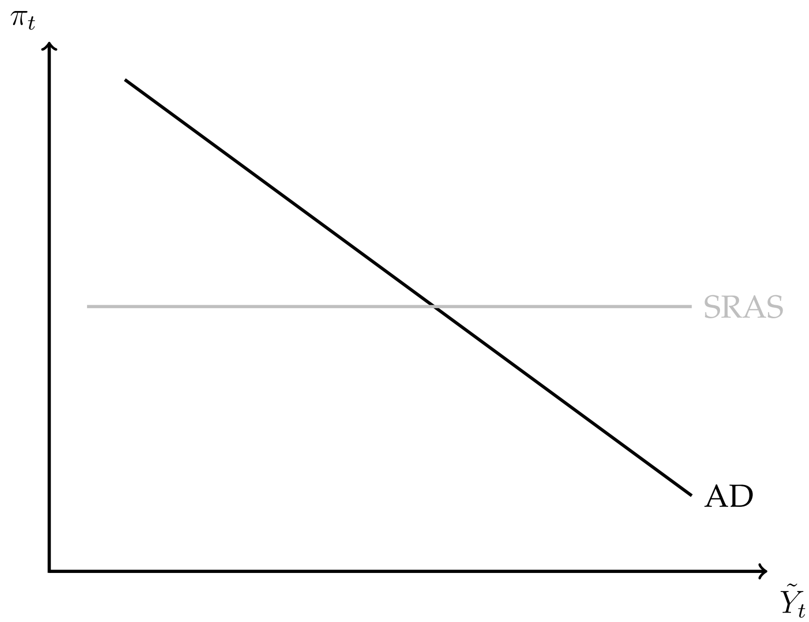

Figure 13: Aggregate Demand

for money. For a given money supply, this means lower prices (and lower inflation). We refer to this relationship as the aggregate demand curve of our medieval economy.

Next consider the price setting equation. We can lag this equation by one period and take natural logarithms to get

$$\ln\left(\frac{P_t}{P_{t-1}}\right) = \theta \ln\left(\frac{Y_{t-1}}{Y^*_{t-1}}\right). \tag{15}$$

We can then use the same types of manipulations as in equations (11) and (12) to get

$$\pi_t = \theta \tilde{Y}_{t-1}. \tag{16}$$

Just as with the quantity equation, we can plot the set of points that satisfy this equation $(\tilde{Y}_t, \pi_t)$ space. This is the gray line in Figure 13. We refer to it as the short-run aggregate supply curve of the medieval economy.

Notice, that the short-run aggregate supply curve is horizontal, as opposed to upward sloping. The reason for this is that the lagged output gap $\tilde{Y}_{t-1}$ appears in equation (16) as opposed to the contemporaneous output gap $\tilde{Y}_t$. It is the contemporaneous output gap that is on the horizontal axis of the figure (not the lagged output gap). In the medieval economy, prices are time $t$ are determined not by the output gap at time $t$, but rather by the output gap at time $t-1$ (the lagged output gap). So, inflation at time $t$ is not a function of the output gap at time $t$.

Together, the aggregate demand curve (quantity equation) and short-run aggregate supply curve (price setting equation) yield equlibrium values for inflation and the output gap in the medieval economy. This equilibrium is represented by the intersection of the two curves in Figure 13. One can then use Okun's Law to determine the equilibrium level of unemployment.

# References

BAXTER, M. AND R. G. KING (1999): "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *Review of Economics and Statistics*, 81, 575–593.

DIEBOLD, F. X. AND G. D. RUDEBUSCH (1990): "A Nonparametric INvestiation of Duation Dependence in the American Business Cycle," *Journal of Political Economy*, 98, 596–616.

DUPRAZ, S., E. NAKAMURA, AND J. STEINSSON (2005): "A Plucking Model of Business Cycles," *Journal of Monetary Economics*, 152, 103766.

EDGE, R. M. AND J. B. RUDD (2016): "Real-Time Properties of the Federal Reserve's Output Gap," *Review of Economics and Statistics*, 98, 785–791.

FRIEDMAN, M. (1964): "Monetary Studies of the National Bureau," in *The National Bureau Enters Its 45th Year*, New York, NY: National Bureau of Economic Research, 7–25, 44th Annual Report, available at http://www.nber.org/nberhistory/annualreports.html.

FUJITA, S., V. RAMEY, AND T. RODED (2026): "Why Didn't the U.S. Unemployment Rate Rise at the End of WWII," Working Paper, Stanford University.

HODRICK, R. J. AND E. C. PRESCOTT (1980): "Post-War U.S. Business Cycles: An Empirical Investigation," Working Paper, Carnegie-Mellon University.

——— (1997): "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit and Banking*, 29, 1–16.

KEYNES, J. M. (1923): *A Tract on Monetary Reform*, London, UK: Macmillan.

——— (1936): *The General Theory of Employment, Interest and Money*, London, UK: Macmillan.

LAIDLER, D. (1999): *Fabricating the Keynesian Revolution*, Cambridge, UK: Cambridge University Press.

——— (2008): "Keynesian Revolution," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, London, UK: Palgrave Macmillan.

OKUN, A. (1962): "Potential GNP: Its Measurement and Significance," Cowles Foundation Paper No. 190.

PETROSKY-NADEAU, N. AND L. ZHANG (2021): "Unemployment Crises," *Journal of Monetary Economics*, 117, 335–353.

ROMER, C. D. (1994): "Remeasuring Business Cycles," *Journal of Economic History*, 54, 573–609.

RUDEBUSCH, G. D. (2016): "Will the Economic Recovery Die of Old Age?" FRBSF Economic Letter, 2016-03.

SHIMER, R. (1998): "Why Is the U.S. Unemployment Rate So Much Lower?" in *NBER Macroeconomics Annual 1998*, ed. by B. S. Bernanke and J. J. Rotemberg, Cambridge, MA: MIT Press, vol. 13, 11–74.

SICHEL, D. E. (1991): "Business Cycle Duration Dependence: A Parametric Approach," *Review of Economics and Statistics*, 73, 254–260.

STEWART, D. (1793): "Account of the Life and Writings of Adam Smith, LL.D." Transactions of the Royal Society of Edinburgh. Read by Stewart January 21 and March 18, 1793.

TARSHIS, L. (1987): "Keynesian Revolution," in *The New Palgrave: A Dictionary of Economics*, ed. by J. Eatwell, M. Milgate, and P. Newman, London, UK: Macmillan, vol. 2, 47–50.