Output Gaps: Uses and Limitations*

BY ROC ARMENTER

The concept of resource slack is central to understanding the dynamics between employment, output, and inflation. But what amount of slack is consistent with price stability? To answer this question, economists define baseline values for unemployment and output known as the natural rate of unemployment and potential output. The concepts of output and employment gaps can be useful to economists in several ways. First, they often guide the inflation forecasts of Federal Reserve staff and other researchers and market participants. Second, some economists argue that employment gaps are a useful guide for policy aimed at achieving maximum sustainable employment and price stability. In this article, Roc Armenter briefly discusses two important examples of sophisticated measures of resource slack that are grounded in economic theory: the nonaccelerating-inflation rate of unemployment and the output gap measure published quarterly by the Congressional Budget Office.

Economists have greatly improved their understanding of the dynamics between employment, output, and inflation since the 1970s. The concept of resource slack is central to these dynamics. For example, economists track the share of unemployed workers, allowing them to estimate how quickly firms may be able to expand employment without having to raise wages to attract workers. Other measures of slack are the percentage of industrial capacity available or the ratio of inventories to sales.

It is unreasonable to expect all workers to be employed or industrial capacity to be at 100 percent. For example, unemployed workers will take the time to find the best job for them and perhaps may need time to relocate.

What amount of slack is consistent with price stability? To answer this question, economists define baseline values for unemployment and output known, respectively, as the natural rate of unemployment and potential output. These are the levels of employment and output consistent with the economy operating with stable prices. The output gap and employment gap are defined as the differences between the actual level for each variable and the baseline value. The actual level may be lower than the baseline level, and thus the output or employment gap can be negative.

The output and employment gap concepts can be useful to economists in several ways. First, output and employment gaps often guide the inflation forecasts of Federal Reserve staff, as well as those of researchers and market participants. For example, if the output gap is positive — that is, output is above its baseline level — firms will be operating close to full production capacity. Hence, firms will not be able to increase production further without significant investments. These investments are costly so it will take a while for firms to increase production. In the meantime,

Roc Armenter is a senior economist in the Research Department of the Philadelphia Fed. This article is available free of charge at www.philadelphiafed.org/research-and-data/publications/.

*The views expressed here are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

firms will respond to increased demand by raising prices. Thus, positive output gaps can signal future inflationary pressures. Since monetary policy operates with significant lags, it is important that policymakers have an accurate inflation forecast.

Second, some economists argue that employment gaps are a useful guide for policy aimed at achieving maximum sustainable employment and price stability, which are the mandated objectives of monetary policy in the United States. Over the medium term, employment is driven by fundamentals such as productivity and labor supply growth, and these medium-term measures are used to infer simultaneously the natural rate of unemployment and the unemployment gap. Most economists do not think monetary policy is part of these medium-term fundamentals. Instead, attempts to drive employment or output above their fundamental levels would result in unwanted inflation and no employment gains.

Economists have developed sophisticated measures of resource slack that are grounded in economic theory, yet remain workable in practical terms. We will briefly discuss two important examples: the nonaccelerating inflation rate of unemployment (NAIRU) model and the output gap measure published quarterly by the Congressional Budget Office (CBO). However, even the latest models recognize that there is a large amount of uncertainty about output and employment gaps. Moreover, there remain competing definitions of the output and employment gaps. Alternative measures sometimes offer contrasting implications for monetary policy. Therefore, it is important to understand the limitations of current output gap estimates for both inflation forecasting and output stabilization.

MILTON FRIEDMAN AND THE NATURAL RATE OF UNEMPLOYMENT

Economists have long believed in a relationship between money, prices, and employment — some say since the 18th century![1] However, it was not until 1958 that A.W. Phillips provided the first statistical analysis comparing wage inflation and unemployment, using data for the United Kingdom since 1861.[2] Phillips found that when unemployment was high, inflation was low. This negative relationship now bears his name: the Phillips curve. A few years later Paul Samuelson and Robert Solow imported the Phillips curve to the United States.[3] Samuelson and Solow used price inflation instead of wage inflation, a choice now preferred by most researchers. Figure 1 displays a typical Phillips curve plot for the period 1948-1965. Each dot corresponds to the inflation and unemployment rate during a quarter. The solid line displays the statistical relationship.

Nowadays economists recognize that the Phillips curve is more than a statistical relationship between two variables. The modern view of the Phillips curve is rooted in the ideas that Milton Friedman developed at the University of Chicago during the late 1960s.[4] Friedman believed that attempts to increase employment by increasing inflation were misguided. Only unanticipated inflation, Friedman argued, has the ability to stimulate employment. For example, if households and firms expect an inflation rate of 2 percent, a 3 percent inflation rate would effectively increase output and employment by boosting real demand. However, as workers and firms came to expect a 3 percent inflation rate, they would embody such expectations in wage demands and price setting. As a result, an inflation of 3 percent would increase nominal output compared with 2 percent inflation, but since all prices and wages adjusted by 3 percent as well, there would be no change in real output and employment.

Nowadays economists recognize that the Phillips curve is more than a statistical relationship between two variables.

Friedman's view was validated when inflation rose persistently in the 1970s despite a marked slowdown in employment growth. Indeed, analysts coined the term stagflation to describe the combination of stagnant growth and inflation.

The old view of the Phillips curve could not explain the stagflation phenomenon. Figure 2 plots unemployment and inflation rates for the period 1970-1979. We include the Phillips curve as given in Figure 1 for the period 1948-1965. The actual observations are all northeast of the Phillips curve: Both inflation and unemployment rates were higher than the theory predicted. This was not immediately recognized, and many policymakers mistakenly believed that inflation would reverse course and moderate.[5]

[1] Robert Lucas, in his Nobel lecture, traced the observation back to the writings of David Hume.

[2] See the article by A.W. Phillips.

[3] See the study by Samuelson and Solow.

[4] For a detailed discussion of the so-called New Keynesian Phillips curve, see the article by Keith Sill on page 17.

[5] The book by Thomas Sargent discusses the experience of the 1970s in rich detail.

FIGURE 1

The Classical Phillips Curve 1948-1965

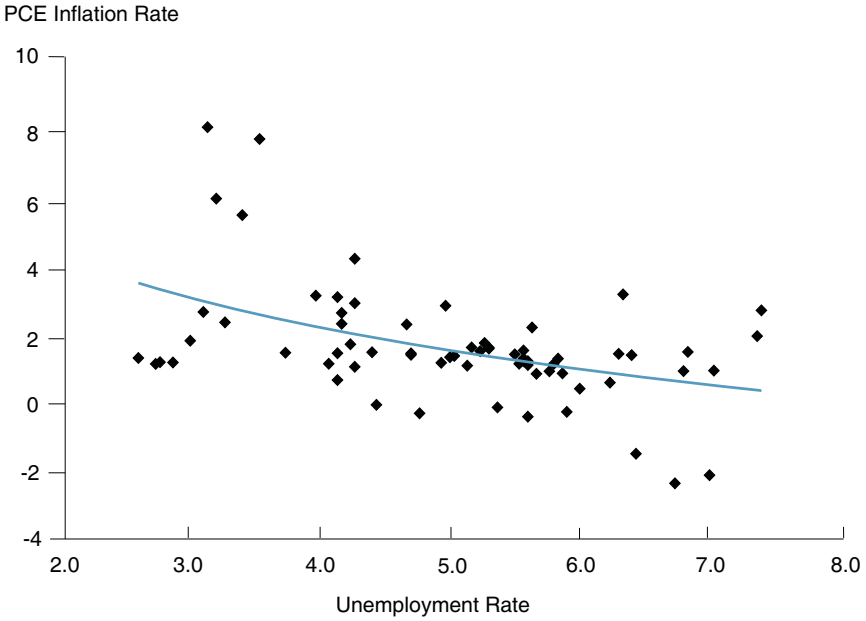

Quarterly data, 1948-1965, seasonally adjusted
Sources: BLS/Haver for unemployment rate; BEA/Haver for PCE inflation rate

FIGURE 2

Stagflation 1970-1979

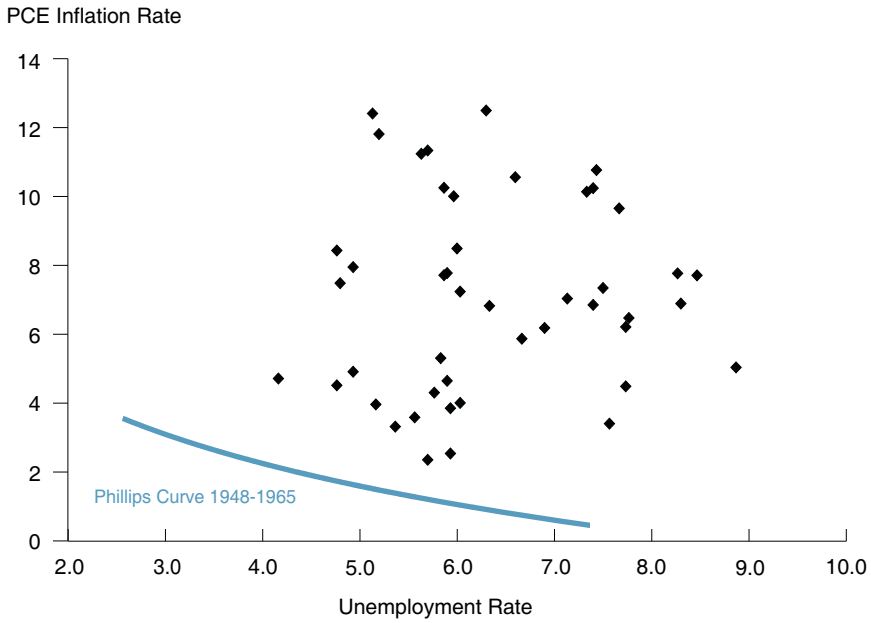

Quarterly data, 1970-1979, seasonally adjusted
Sources: BLS/Haver for unemployment rate; BEA/Haver for PCE inflation rate

(See *Did Oil Prices Cause the Inflation in the 1970s?*)

Friedman defined a baseline value for employment in his theory and thus postulated employment gaps as we know them today. The so-called natural rate of unemployment is the rate we would observe if inflation were exactly as expected. This definition is mainly theoretical, but, as we shall see later, some current measures of employment gaps are inspired by this definition.

Friedman's view came with some key policy "lessons" that would be learned the hard way. First, attempts to exploit the Phillips curve would bring about lower unemployment temporarily at best; further increases in the money supply would be met by rising inflation. Second, in order to be able to stabilize output in the short term without causing rising inflation, policymakers need to know what the natural rate of unemployment is. Friedman himself was deeply skeptical that this could be done effectively. To do so, policymakers would need to accurately forecast the state of the economy. Finally, Friedman's analysis highlighted the importance of inflation expectations in achieving price stabilization. Nowadays central banks around the world realize that any attempt to exploit the trade-off between inflation and employment will be short-lived, at best. Thus, central bank policy emphasizes price stabilization in order to anchor inflation expectations.[6]

SOME OUTPUT AND EMPLOYMENT GAP MEASURES

Currently, there is an array of statistical procedures to approximate the natural rate of unemployment or its

[6] The study by Jeffrey Lacker and John Weinberg contains further discussion of the Phillips curve and the research on inflation and unemployment.

Did Oil Prices Cause the Inflation in the 1970s?

Twice in the 1970s the price of oil soared. The first oil crisis started in October 1973 when the Organization of the Petroleum Exporting Countries (OPEC) declared it would stop oil shipments to the United States. Oil prices tripled in a year. The second oil crisis occurred in 1979 as a result of the Iranian revolution. Although some OPEC countries increased their production, oil prices more than doubled. The United States would have to wait until 1986 to see oil prices drop below $20 again.

Some of the most dramatic inflation rates recorded were associated with these oil crises. However, oil prices alone cannot explain the extraordinary behavior of prices and employment in the 1970s. First, the increase in the inflation rate was noticeable well before 1973, while oil prices remained low. The figure plots the inflation rate (left axis) and the oil-price level (right axis). By 1970, inflation was above 4 percent and never dropped below 3 percent. Meanwhile, oil prices were completely flat.

Second, the crises of 1973 and 1979 were one-time price increases. As such, they can explain only temporary increases in inflation, for example, the spike in the inflation rate in 1974. However, one-time price increases cannot explain the persistent rise in inflation that is evident in the figure.* Inflation was close to 8 percent in 1978, five years after the oil crisis.

Finally, inflation was widespread. Oil prices clearly affected the price of gas, transportation, and some other goods that use oil in their production. But oil prices should have only a small impact on other goods, such as food or apparel. Yet these goods also showed persistent inflationary behavior.

For further reading on oil prices and economic activity, see the *Business Review* article by Sylvain Leduc and the one by Keith Sill.

FIGURE

Oil Prices and Inflation in the 1970s

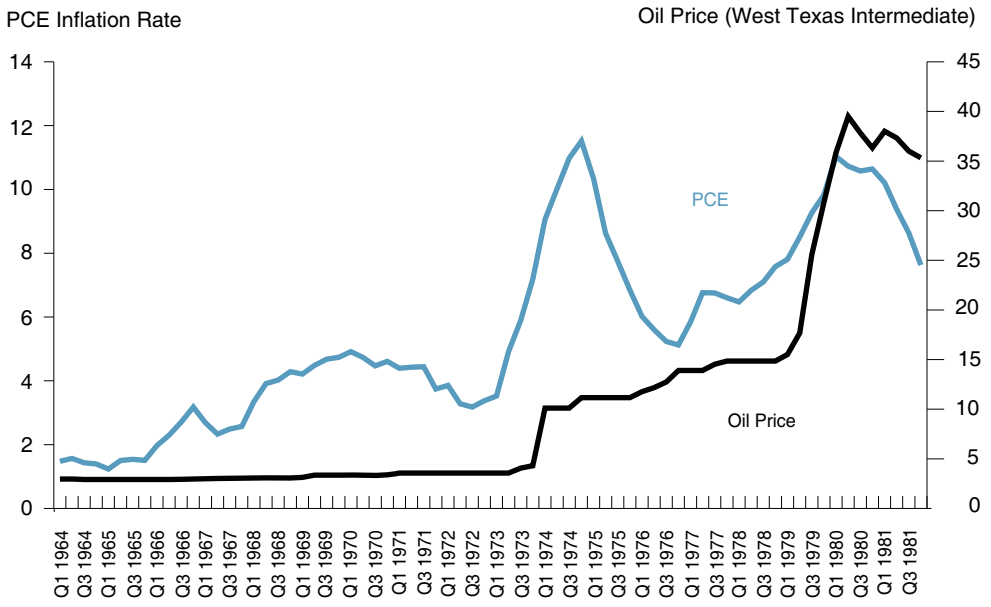

Quarterly data, 1964-1981
Sources: BEA/Haver for PCE inflation rate; Wall Street Journal/Haver for oil prices

* Cost increases, such as the oil shock, take some time to pass through to consumer prices. However, the rise in inflation was too large and, more important, too prolonged for any reasonable estimate of pass-through.

equivalent concept in terms of GDP: potential output. There are many differences across these measures. For example, some models are evaluated monthly, others quarterly. Models also often differ in the variables taken into consideration. However, all of the gap measures discussed here share two defining properties. First, baseline levels are defined by their neutral stance on inflation; that is, they all try to measure the level of output or employment consistent with price stability. Second, researchers assume that the natural rate and potential output move slowly; that is, their determinants (such as labor supply) operate exclusively in the medium to long term.

We will now briefly discuss two leading models for employment and output gaps, respectively.

The NAIRU. The very name of the NAIRU — the nonaccelerating-inflation rate of unemployment — spells out the model. The NAIRU is the level of unemployment consistent with inflation behaving as expected. In the model, researchers also specify how inflation and unemployment are related in the short term.

The NAIRU closely tracks Friedman's ideas. In this model, expected inflation is usually given by an average of past inflation rates. The idea is that firms and households create their inflation expectations on the basis of their recent experience with prices. If inflation is unexpectedly high in one period, it will drive unemployment down, but it will also drive inflation expectations up, as the average of past inflation rates increases accordingly. If inflation stays at the same high rate, the effect on unemployment will be less. Hence the trade-off between inflation and unemployment is necessarily short-lived.

NAIRU models regularly incorporate information about real factors, such as commodity prices, that may affect unemployment or inflation in the short run. Newer NAIRU models allow the natural rate of unemployment to vary over time.[7] The fluctuations in the natural rate of unemployment, though, are assumed to be slow moving and very persistent. The assumption mostly conforms to our intuition. For example, the natural rate of unemployment reflects long-term changes in labor productivity growth, brought about by the introduction of new technologies.

Figure 3 plots the unemployment rate and an estimate of the NAIRU for the period 1949-2005. The NAIRU increased in the 1970s and fell for most of the 1990s. However, the NAIRU is much smoother than the actual unemployment rate, capturing only the persistent swings in the data. The flip side, as shown in Figure 3, is that the employment gap — the difference between the NAIRU and the unemployment rate — is relatively short-lived; that is, it does not stay positive or negative for long periods of time.

CBO Output Gap. The Congressional Budget Office (CBO) publishes a quarterly estimate of potential output that is a key input to the fiscal policy outlook and is widely tracked by professional forecasters. The CBO defines potential output as the level of output that is neither adding to nor subtracting from inflationary pressures. Thus, it is an equivalent concept to the NAIRU in terms of output rather than employment.

Because of its focus on fiscal policy, the CBO is interested in forecasting output in the medium term in addition to inflation. The longer horizon forecast is needed in order to evaluate

[7] For an overview of NAIRU models, see the 1997 study by Douglas Staiger, James Stock, and Mark Watson.

FIGURE 3

NAIRU and the Unemployment Rate

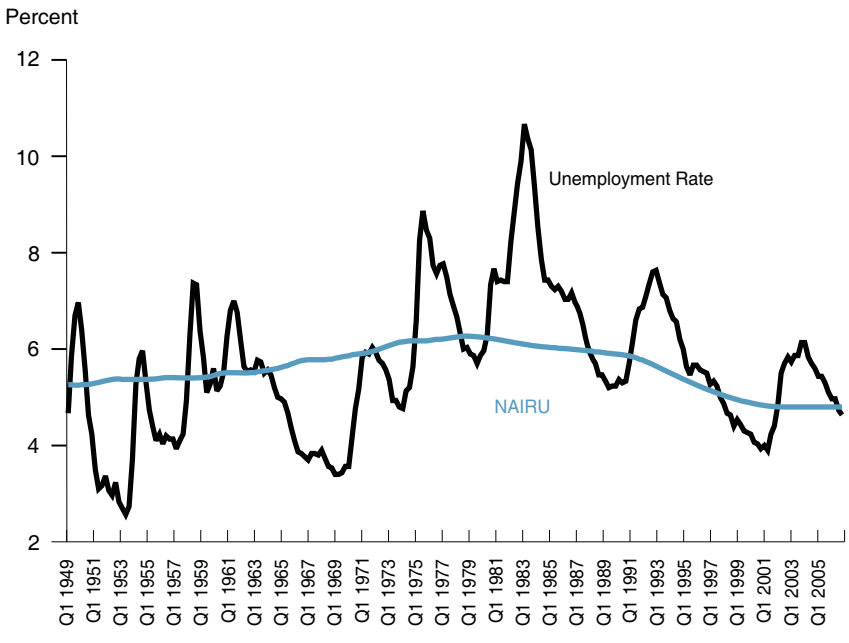

Quarterly data, 1949-2005
Sources: BLS/Haver for unemployment rate; Haver for NAIRU

the fiscal cost of different programs and the debt and tax changes needed to finance them. To this end, the CBO makes an effort to incorporate additional information into its computation of potential output. For example, the CBO pays special attention to demographic and educational trends. These trends are important for forecasting the labor supply over the next five to 10 years, but they are unlikely to affect inflation.

The CBO uses a production function approach to track gross domestic product.[8] The production function is very good for combining diverse data sources. For example, the CBO uses data on labor supply, capital utilization rates, industrial capacity, and electricity consumption. Researchers then break down these data series into transitory and persistent components; then the latter is used to construct the estimate of potential output. A key input in the estimate is productivity or, more broadly speaking, technology. For this the CBO must rely on estimates of worker productivity and judgment calls. And, once again, the assumption is that these fundamentals are slow moving.[9]

Figure 4 plots the output gap series, as computed by the CBO, for the period 1956-2009. Whenever the output gap is positive, the economy is running above potential — for example, in the second half of the 1990s. In contrast, the economy was below potential for most of the 1970s and early 1980s, as well as in 2009. With some exceptions, output gaps close quite fast, meaning the economy reverts to the path of potential output.

[8] A production function states a relationship between inputs (like labor and capital) and output (goods and services combined).

[9] For further reading on the computation of the CBO output gap, see the background paper published by the CBO.

RECENT DEVELOPMENTS

Economists continue to work on models of potential output and the natural rate of unemployment. Newer models seek a more flexible specification of potential output, allowing for short-term fluctuations, or incorporate additional variables in the specification, such as interest rates or aggregate consumption.

The second half of the 1990s also presented a challenge. As in the 1970s, the observed inflation and unemployment rates did not square with existing models. This time, though, the situation was the opposite and thus more benign: Output and employment growth were high, yet inflation remained low and stable. Figure 5 plots the output gap, as computed by the CBO, and inflation, computed from personal consumption expenditures (PCE), for the second half of the 1990s. The output gap turned positive in early 1996 and became very large by the end of the decade, as one can see from Figure 5 as well. Not only was the output gap large, but, by 2000, it had not shown any signs of moderation. In this context, many economic models would predict that inflation would rise sharply. Yet inflation actually declined from 1996 to 1998, as can also be seen in Figure 5. Only in 1998 and 1999 did inflation show a very modest pick-up.[10]

SHORTCOMINGS OF STATISTICAL GAPS

Despite the sophistication of these models, it turns out that they are limited in their ability to forecast inflation in the short term. It is important to understand the limitations of the use-

[10] The 2002 study by Staiger, Stock, and Watson contains an extensive discussion of the experience in the 1990s for NAIRU models.

FIGURE 4

Output Gap

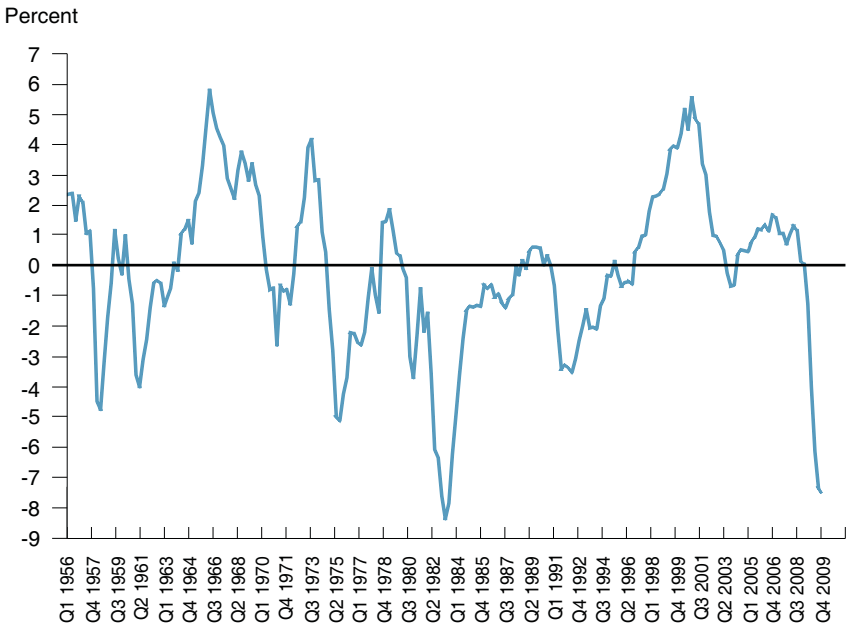

Quarterly data, 1956-2009
Sources: CBO/Haver

fulness of output and employment gaps as a guide for setting monetary policy.

The first set of shortcomings of statistical models involves the data. Most macroeconomic data are released with a significant lag and are subject to revision. Moreover, forecasts must rely on the most recent releases, which are indeed the most likely to be revised. This is often known as the "end-of-sample" problem, and there is little researchers can do about it. For example, models that use GDP data must always rely on the Bureau of Economic Analysis' (BEA) "advance" estimate for the last quarter, which is usually heavily revised.[11]

Another shortcoming arises from the assumption that estimates of potential output and the natural rate of unemployment move smoothly. While we usually view the technology diffusion process as slow and smooth, some factors can have a sudden impact on supply. For example, extreme weather conditions can lead to the disruption of services and, in some cases, have a very persistent effect. More important, this assumption implies that fluctuations in output or unemployment are always initially counted as changes in the gap, not as changes in potential output or the natural rate of unemployment. In more technical terms, errors in the short-term forecast of output or employment gaps accumulate, and it takes a while before these errors are corrected. For example, say a new technology brings a significant improvement in productivity and thus simultaneously increases potential and actual output. Because researchers do not observe productivity directly, the increase in output will be initially viewed as a temporary deviation and ascribed to an increase in the output gap. Only when researchers observe that the higher output level persists over the medium term will the estimate of potential output be updated.

We return to the experience of the second half of the 1990s for a real-life example regarding estimates of potential output. Figure 6 plots (in black) the level of real GDP from 1996 to 2006, with the level normalized to 100 in 1996. As mentioned earlier, this was a period of rapid and persistent economic growth: Real GDP grew 20 percent in four years. Let us see how estimates of potential output caught up with the experience. The gray line in Figure 6 plots the estimates of potential output provided by the CBO in 1996; they are thus a forecast for the period. The dark blue line reports the estimates of potential output at the end of 1998. After three years of actual data pointing to strong GDP growth, economists barely nudged up their estimates of potential output.[12] In other words, the model suggested a huge output gap. Finally, the light blue line plots the latest estimates of potential output. Economists now recognize that most of the growth in the second half of the 1990s was due to fundamentals: Estimates of potential output track real GDP much more closely than the initial estimates.

There are also reasons to think that the statistical relationship between inflation and output varies over time. Policymakers are aware of this and often have additional information about the likely path of infla-

FIGURE 5

Output and Inflation in the 1990s

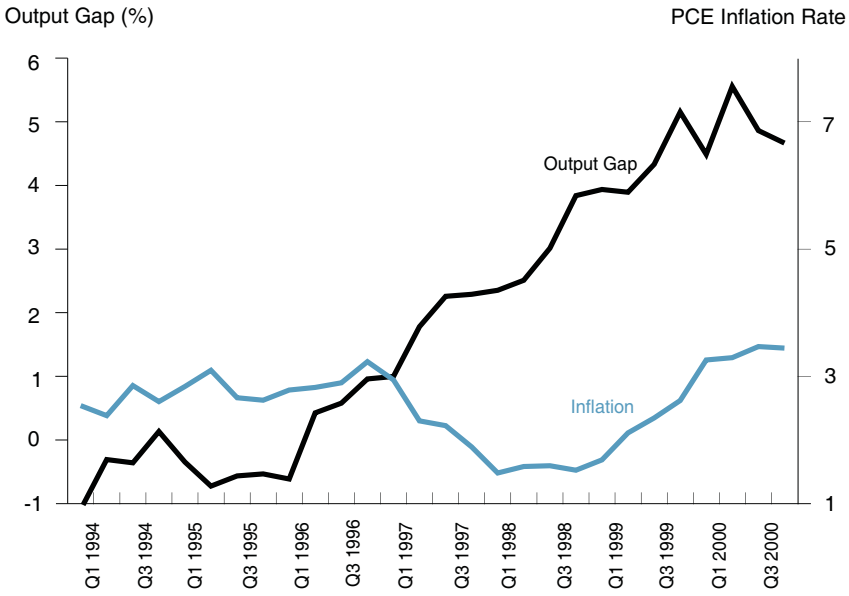

Quarterly data, 1994-2000
Sources: CBO/Haver for output gap; BEA/Haver for PCE inflation

[11] Advanced models are careful to use real-time data. The Real-Time Data Research Center at the Federal Reserve Bank of Philadelphia has developed a real-time data set for macroeconomists, available at http://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data/.

[12] This is a problem common to all statistical models: It is difficult to distinguish temporary from permanent fluctuations at the end of the sample.

FIGURE 6

GDP and Potential Output Estimates

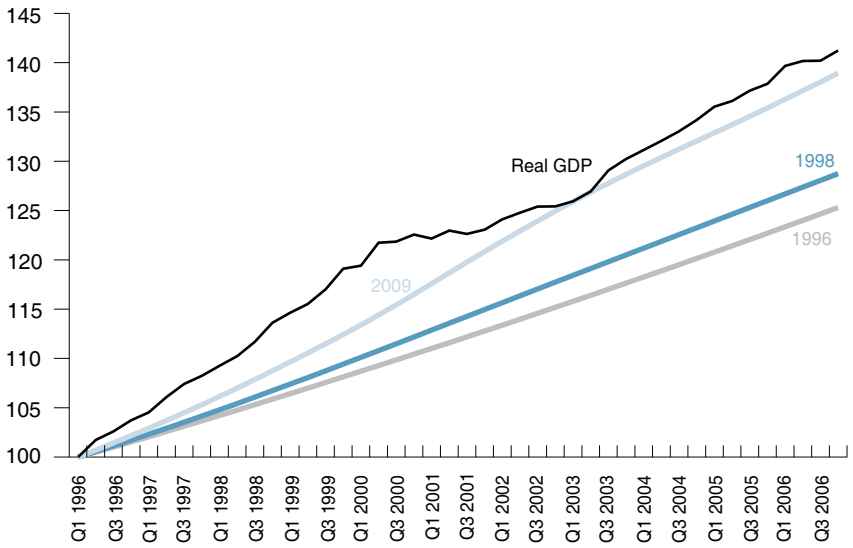

Quarterly data, all series set equal to 100 in 1996
Sources: CBO for potential output estimates; BEA/Haver for real GDP

tion or output that cannot be readily quantified. Statistical models provide little guidance on how to modify the resulting estimates in order to incorporate the additional information. For example, the current downturn has affected the manufacturing and construction sectors disproportionately. It is clear that some workers will have to find jobs in new sectors, and this will take some additional time. The natural rate of unemployment will have to be adjusted upward at least for a time to account for the reallocation. However, the statistical models do not contain a breakdown by sector, and thus it is not clear how to modify the model's estimates.

CONCLUSION

We have seen how economists have developed the concepts of output and employment gaps as key tools for understanding the long-term movements in unemployment and output. These long-term measures are often used to provide information on economic activity relative to trend and thus help us gauge how relatively strong or weak current economic activity is. Understandably, such knowledge is often important for implementing economic policies.[13] However, as this article has shown, measures of the NAIRU and potential output are subject to severe measurement problems that detract from their application. Therefore, it is important that we understand their limitations.

[13] The article by Anthony Santomero discusses how the availability and reliability of economic information affects policymaking.

REFERENCES

Congressional Budget Office. "A Summary of Alternative Methods for Estimating Potential GDP," Background Paper (March 2004).

Lacker, Jeffrey, and John A. Weinberg. "Inflation and Unemployment," Federal Reserve Bank of Richmond, Annual Report (2006).

Leduc, Sylvain. "Oil Prices Strike Back," Federal Reserve Bank of Philadelphia, *Business Review* (First Quarter 2001), pp. 21-30.

Lucas, Robert. "Monetary Neutrality," Nobel Prize lecture, December 7, 1995, available at: http://nobelprize.org/nobel_prizes/economics/laureates/1995/lucas-lecture.pdf.

Phillips, A. W. "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," *Economica*, 25:100 (1958), pp. 283-99.

Samuelson, Paul A., and Robert M. Solow. "Analytical Aspects of Anti-Inflation Policy," *American Economic Review,* 50:2 (1960), pp. 177-94.

Santomero, Anthony M. "Making Monetary Policy: What Do We Know and When Do We Know It?" Federal Reserve Bank of Philadelphia *Business Review* (Fourth Quarter 2005), pp. 1-8.

Sargent, Thomas J. *The Conquest of American Inflation.* Princeton: Princeton University Press, 2001.

Sill, Keith. "The Macroeconomics of Oil Shocks," Federal Reserve Bank of Philadelphia, *Business Review* (First Quarter 2007), pp. 21-31.

Sill, Keith. "Inflation Dynamics and the New Keynesian Phillips Curve," Federal Reserve Bank of Philadelphia *Business Review* (First Quarter 2011).

Staiger, Douglas, James H. Stock, and Mark Watson. "The NAIRU, Unemployment, and Monetary Policy," *Journal of Economic Perspectives*, 11:1 (1997), pp. 33-49.

Staiger, Douglas, James H. Stock, and Mark Watson. "Prices, Wages, and the U.S. NAIRU in the 1990s," in Alan Krueger and Robert Solow, eds., *The Roaring Nineties*. Russell Sage Foundation, New York, 2002.

Urban Productivity Advantages from Job Search and Matching*

BY JEFFREY LIN

Densely populated areas tend to be more productive. Of course, the cost of living and producing in these locations is higher because congestion raises the cost of scarce fixed resources such as land. But despite the higher prices, many people and businesses continue to live and work in these areas. Why? One explanation is that these locations have natural advantages, such as proximity to a river. Another says that this concentration of households and businesses by itself generates productivity advantages in the form of agglomeration economies. In studying these agglomeration economies, economists have pursued two other questions. Do agglomeration economies exist and how big are they? And what are the precise sources of these agglomeration economies? In this article, Jeffrey Lin describes the evidence for agglomeration economies from job search and matching and then asks whether it may be large enough to offer meaningful explanations for differences in productivity and density.

Why do people in densely populated areas tend to be more productive? In countries like the U.S., places dense in workers, machines, firms, and households also tend to be places where people are able to produce more things. Of course, these places are also usually more expensive to produce in and to live in because congestion raises the price of scarce fixed resources such as land. Despite these high prices, many businesses and people continue to choose these locations.

Jeffrey Lin is a senior economist in the Research Department of the Philadelphia Fed. This article is available free of charge at www.philadelphiafed.org/research-and-data/publications/.

*The views expressed here are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

A typical first explanation is that these densely populated areas enjoy intrinsic natural advantages, such as Philadelphia's proximity to a navigable waterway and a relatively deep harbor. Advantages like these can reduce the costs of shipping and the price of traded goods, attracting both businesses and households. This story can often be compelling, even though, today, many people in the Philadelphia region do not experience direct benefits from the Delaware River. An intriguing alternative explanation is that bringing together workers, businesses, and households can, by itself, generate these productivity advantages. These kinds of advantages are often called agglomeration economies, and they describe situations in which geographic concentrations of economic activity allow businesses and households to save on the costs of transporting people, materials, and ideas.

Urban economists have pursued two related research questions. First, do these agglomeration economies exist, and, if so, how big are they? Second, what are the precise sources of these agglomeration economies?

Many researchers have already discovered evidence that these agglomeration economies do exist and that they are big enough to offer meaningful explanations of present-day differences in productivity and density. For example, in an attempt to answer the first question, economists Antonio Ciccone and Robert Hall, using data for U.S. states, found that a doubling of employment density increased average labor productivity by about 6 percent. Although other studies have provided different estimates of the exact mag-

nitude of this effect, many have noted that agglomeration economies make an important contribution to differences in productivity across locations.[1] In addition, research by Satyajit Chatterjee (discussed in his 2003 *Business Review* article) also suggests that agglomeration economies play some explanatory role in these differences, even after accounting for natural advantages.

For both academic and policy reasons, an important next step is to investigate the specific *sources* of agglomeration economies. In this article, I will discuss some of my recent research on one potential source: opportunities to better match workers' skills to job requirements. Dense urban areas have thick labor markets —– that is, markets with many different kinds of workers and jobs —– and might therefore benefit from improved job search and matching. This idea — that markets with more participants can offer better matches — is typically attributed to Alfred Marshall, and the idea was formalized in economist Peter Diamond's "coconut" model. (If consumers have tastes for a particular variety of "coconut," they are more likely to find the one they prefer in a large market where more types of coconuts are sold.) Intuitively, we know that workers have varying skills and jobs have varying skill requirements. From the perspective of a worker, searching for a suitable job may be easier in a large city with many potential employers. Put another way, workers in large cities may find a job that is better matched to their talents, for the same search costs. This is a potential source of agglomeration economies; geographic concentration increases productivity because workers need not let their acquired skills lapse by taking less-suitable jobs.

It is important to note that, in theory, there are a number of different sources of agglomeration economies. In a 2005 *Business Review* article, Jerry Carlino discusses a few of the many possible economic mechanisms responsible for agglomeration economies. His 2001 *Business Review* article talks about one possible mechanism — knowledge spillovers — related to the increased production and flow of (new) ideas and information in dense cities. In a later *Business Review* article (2009), he describes his paper in which he evaluates another potential mechanism: Urban population density may increase the amount and variety of goods and services available for households to consume. As another example, I show evidence for yet another mechanism in a recent working paper: Geographic concentrations of skilled workers and potential users of new products or processes can increase the rate of adaptation to new technologies. In general, as explained by Gilles Duranton and Diego Puga, agglomeration economies might arise from mechanisms related to sharing, learning, or matching. Sharing refers to advantages that arise from distributing the costs of large indivisible investments across many producers or consumers, as might be the case with a large factory or consumption amenities, as in Carlino's article. Learning refers to advantages in either the creation of new technologies, as described by Jane Jacobs; the formation of human capital, as described by Edward Glaeser and David Maré; or adaptation to new technologies, as in my working paper.

In order to evaluate alternative proposals, policymakers concerned with city growth, the productivity of local workers, or the welfare of local residents need to understand the specific economic forces that generate productivity advantages and attract businesses and households to certain places. Should local leaders sponsor arts and cultural programs or invest in transportation infrastructure? What kinds of businesses should cities be interested in attracting? The answer to these questions depends on the relative strength of different kinds of agglomeration economies. In other words, for both intellectual and practical reasons, it is useful to know what is happening inside the "black box" of agglomeration economies.

However, finding evidence that distinguishes one kind of agglomeration economy from another can be challenging. Different mechanisms often have similar predictions for aggregate city-level data. For example, most (if not all) kinds of agglomeration economies predict higher wages and higher land prices in denser cities.

In order to evaluate alternative proposals, policymakers concerned with city growth, the productivity of local workers, or the welfare of local residents need to understand the specific economic forces that generate productivity advantages and attract businesses and households to certain places.

[1] See the paper by Gerald Carlino and Richard Voith; the recent working paper by Morris Davis, Jonas Fisher, and Toni Whited; and the 2004 article by Stuart Rosenthal and William Strange.

(These facts are in line with conventional wisdom and easily confirmed using aggregate census data.) Therefore, looking inside the "black box" of agglomeration economies often requires creative research strategies. Recent work in this area, including my own, has been made possible by the increasing availability of large data sets that contain detailed information at the plant, household, or worker level. Using micro-data, it is sometimes possible to test predictions that are unique to one kind of agglomeration economy and not associated with another kind. In this way, it becomes possible to highlight variables that should be of interest to policymakers.

I will describe the evidence for agglomeration economies from job search and matching using just such a strategy. An important caveat is that the research strategy described here does not rule out other sources of agglomeration economies. Instead, I evaluate whether there is evidence for this source of agglomeration economies and then ask whether it may be large enough to offer meaningful explanations for differences in productivity and density.

JOB SEARCH AND MATCHING IN CITIES

In my recent working paper with Hoyt Bleakley, we test for agglomeration economies from job search and matching. The intuition for our test is as follows. Consider a worker in a small city who loses her job. She has some specialized skills (either innate or gained through experience) suited to the activities she performed or the output she produced in her previous job. If the separation from her previous job is permanent, the worker now faces a choice: She could wait a long time before finding employment performing similar tasks but at a different firm. Or, because waiting is costly, it may make more sense to accept a job elsewhere in the local economy that is less suited to her unique skill set. (Alternatively, she might choose to move to a location where there is greater demand for her skills, but of course, moving is also costly.) Since her skills are less suited to this job, some of her skills go unused, and she may be less productive.

This worker, in a small city, faces a "small numbers" problem: She happens to be without a job, but does there happen to be another firm that needs a worker with her skill set? On the other hand, workers in dense cities benefit from market thickness: They are less likely to be in a narrow labor market at a moment in which their skills are in excess supply. This potential source of agglomeration economies yields an interesting, and potentially unique, prediction: Workers should choose to eschew their specialized skills *less frequently* in large, dense cities, where they are more likely to find job openings suited to their talents.

We evaluate this prediction by examining the likelihood that workers change *occupations* or *industries*. These job classifications, characterizing either the tasks or activities performed or the kinds of output produced, have been used in a number of labor-market studies on specific human capital.[2] We expect that in the presence of agglomeration economies from job search and matching, workers should choose to change occupations and industries less frequently in denser labor markets.

Further, this agglomeration economy should also affect workers' early decisions about skill specialization. In separate studies, economists Kevin Murphy and Sunwoong Kim have proposed how density might change the market for specialized skills. In Kim's model, sparsely populated areas have fewer firms in each sector, and therefore, a worker might have invested less in narrow skills because she anticipated that there would be fewer potential employers in the event of a separation.[3] Therefore, in large cities, workers choose to invest more in specialized skills, making it even less likely that they would want to change occupations or industries in dense cities and compounding density's effect on productivity.[4]

Using data from the decennial U.S. census and the monthly Current Population Survey (CPS), Bleakley and I confirm this prediction. We find that workers are less likely to change occupation or industry in metropolitan areas with high population density (Figure 1). The data are at the worker level, and the key outcome of interest is a change in each worker's reported occupation or industry.[5] Respondents to the 1970 census reported these changes for 1965 and 1970. The CPS samples in the 1990s and 2000s reported these changes for individual workers, both for the year of the survey and up to three years earlier. The key explanatory variable is local population density, measured for each worker's metropolitan area of residence. Figure 1 summarizes our main result. Here, each point represents a metropolitan

[2] For example, see the study by Derek Neal and the one by Daniel Parent on industry-specific skills; see Gueorgui Kambourov and Iourii Manovskii's recent paper on occupation-specific skills.

[3] Alternatively, workers in small cities with specialized skills might choose to move to denser cities.

[4] For example, James Baumgardner found that doctors are more specialized in big cities; similarly, Luis Garicano and Thomas Hubbard found more specialization among lawyers in larger markets.

[5] We obtain similar results whether our outcome of interest measures a change in each worker's reported occupation, a change in reported industry, or a change in either reported occupation or reported industry.

FIGURE 1

Occupation and Industry Switching and Local Population Density

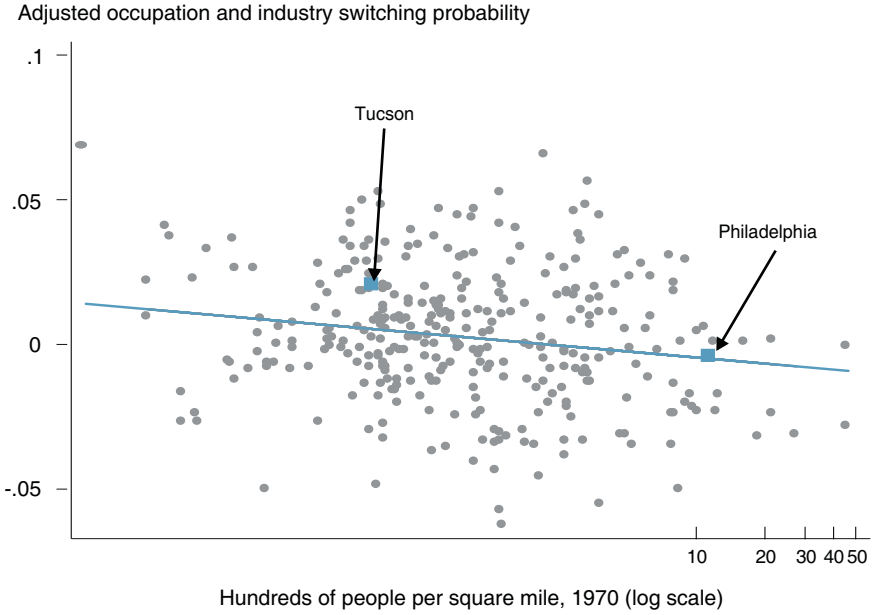

Source: Author's calculations and the 1970 U.S. census

area or a group of co-terminous counties in 1970, and population density is measured on the horizontal axis. The vertical axis measures the probability that a worker in each location changed either occupation or industry between 1965 and 1970. The fitted line shows that workers in locations with higher population densities are less likely to switch occupations or industries. Further, the magnitude of this thick-market effect is large enough to be relevant in understanding differences across locations. For example, a change in density from, say, Tucson, Arizona, to Philadelphia, is associated, on average, with a decrease of 1 percent in occupation or industry switching over a five-year period.

This negative correlation between switching and local population density supports the existence of agglomeration economies in job search and matching. But we also rule out other important alternative explanations. For example, we compare similar workers by controlling for characteristics such as gender, age, race, ethnicity, and educational attainment, and whether or not they have moved recently. We also control for fixed characteristics of a worker's previous occupation and industry, so that our comparison is among workers sharing the same initial occupation and industry. Jobs in different occupations and industries may require different levels of specialized skills. If we control for previous occupation and industry, the results do not simply reflect differences in the composition of occupation or industry across cities. The graph in Figure 1 already controls for all of these effects.

Metropolitan areas are also different along a lot of other dimensions. We control for other characteristics of cities, such as industry composition (e.g., the relative size of the manufacturing sector), average educational attainment, and climate, with little impact on our main result. There is an additional issue of potential measurement error associated with using metropolitan-area-level population density. Since metropolitan areas are based on county boundaries, we are more likely to mis-measure local density in western states that feature relatively large counties. For example, the Los Angeles metropolitan area includes counties that stretch to the Arizona and Nevada borders, including desert lands that are sparsely populated. Our results are similar when we adjust our density measure using census tract data.

Another story to consider is that changing jobs or employers by workers (as opposed to changing occupation or industry) may also depend on the size of the local labor market. Other studies have found mixed evidence of density's effect on job switching.[6] One way we can check to see how this might affect our results is to use information available in the U.S. CPS supplements. This is the survey conducted every month to estimate important statistics such as the unemployment rate. In addition, the CPS also periodically includes supplemental questions of interest to researchers or policymakers. In January and February, these supplements usually include questions related to job changing. In these supplements, the CPS reports workers' reasons for changing jobs; many lost their jobs because their plant or firm closed. Thus, increased opportunities due to population density probably did not cause them to change jobs, since they lost their jobs involuntarily. These workers also change occupation or industry less frequently in larger cities, so job

[6] See the papers by Bruce Fallick, Charles Fleischman, and James Rebitzer; Jeffrey Groen; Guido de Blasio and Sabrina Di Addario; and Jeremy Fox for conflicting evidence on this question.

changing is probably not an important explanation of our main result.

Some workers may have innate specialized skills and may also "sort" themselves into large metropolitan areas. The fact that they have innate specialized skills implies that they may choose to switch occupations or industries less frequently. However, in this story, these workers choose to live in large labor markets for reasons other than improved opportunities for job search and matching. For example, they may be interested in the consumption amenities available in such cities. If this is an important explanation for our main result, workers whose location choice is not influenced by such considerations should not experience a similar pattern relating density to occupation or industry switching. In fact, using information on workers' places of birth, we find that our results are similar for those workers whose choice of location was influenced by the state in which they were born. Taking all of these pieces of evidence together, we argue that agglomeration economies from job search and matching are the likeliest explanation for our results.

YOUNGER WORKERS

An additional piece of evidence weighs in favor of agglomeration economies from job search and matching. If job searching is less costly in large cities, we can make another interesting prediction: People may find it easier to shop around for a good occupation or industry match in a dense city. Of course, it makes sense to do this for younger workers who are just starting their careers: They have fewer specialized skills accumulated, and they have the rest of their careers to gain from great matches. In contrast, older workers have spent many more years accumulating specialized skills: Instead of sampling different occupations, these workers choose jobs more closely matched to their existing skills.

Following this logic, the correlation between changing occupation and industry and population density may depend on workers' potential experience. (Potential experience measures how long workers have potentially been in the labor market: their age, minus the number of years they spent in school, minus six, the number of years between birth and school.) We find that this is indeed the case. Figure 2 shows the effect of density on occupation and industry switching for different levels of potential labor market experience.

For young workers with less than 10 years of potential experience, being in a large city actually increases the likelihood that they will change occupations or industries. (In Figure 2, this can be seen in the positive estimated effect of density on occupation and industry switching.) In contrast, for older workers, density lowers the likelihood of such changes. (On average, the effect due to older workers dominates the overall effect seen in Figure 1, since older workers constitute much of the total workforce.) This positive effect of density on switching early in workers' careers provides further support for the thick-market matching hypothesis, but it is harder to reconcile with other stories of how density might affect occupation and industry switching. If there are benefits from matching in dense cities, workers could take advantage of low search costs to search more intensively for the right occupation or industry match. This occupation and industry shopping could potentially be greater than the negative effect of density on switching shown in the previous section (and thus be, on net, positive). However, since search intensity is like an investment whose gains are realized throughout the working lifetime, this new, positive effect should be strongest at younger ages. Compare this with a story in

FIGURE 2

Effect of Density on Occupation and Industry Switching Depends on Potential Experience

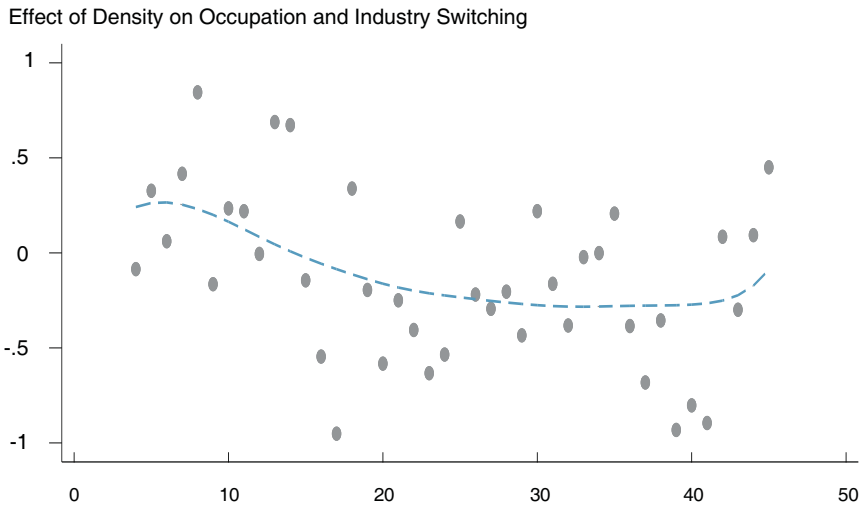

Source: Author's calculations and the 1970 U.S. census

which workers in dense cities are more specialized for some other reason (not better job search and matching), such as faster learning or greater returns to specialization because of improved opportunities for the division of labor. If there are no differences in search costs across cities, it is unlikely that we would observe more occupation and industry switching in dense cities among the youngest workers.

POTENTIAL IMPLICATIONS FOR PRODUCTIVITY AND WAGES

Finally, our estimated differences in occupation and industry switching could be large enough to offer meaningful explanations of differences in productivity. We can get a feel for what our estimates might mean for the relationship between density and wages by doing some quick calculations. First, in small cities, specialized skills fall into disuse faster, as workers churn through more occupations and industries. There are earlier estimates by Derek Neal (1995) and Daniel Parent (2000) on how much of a worker's wage is due to industry-specific skills. Neal estimates that 10 percent of income is derived from industry-specific skills for men with 10 years of experience; Parent estimates that 10 to 20 percent of workers' income is derived from industry-specific skills. To span the range of likely possibilities, say that the fraction is somewhere between 5 and 25 percent. We multiply this by our own estimates of density-driven differences in industry switching — approximately 0.6 percent measured over a five-year horizon or about 4.8 percent over a 40-year career. These calculations suggest that, over 40 years, a doubling of labor market density implies somewhere between 0.2 percent and 1.2 percent higher wage growth through this mechanism. In comparison, the extra growth in wages in dense areas, in the same units, is about 2 percent over 40 years.

Second, in small cities, workers might be less inclined to invest in specialized skills. Note that the previous calculation does not account for differences in behavior that might result from expectations about the usefulness of specialized skills in big cities. Calculating the potential effect on wages is difficult, since it depends on how costly it is to acquire specialized skills and how quickly those skills fall into disuse, even without changing occupation or industry. In our related working paper, we find that, for reasonable values of these variables, this mechanism can explain nearly all of the observed differences in productivity levels across locations. To sum up, our back-of-the-envelope calculations suggest that the relationship between density and occupation and industry switching can account for most of the differences across cities in workers' income growth and nearly all of the differences in income levels.

> Overall, workers in metropolitan areas with lower population density tend to be more likely to change occupations.

PHILADELPHIA AND THE THIRD FEDERAL RESERVE DISTRICT

These differences in occupation changing can be seen even among the handful of metropolitan areas within the Third District. The Table displays population density, taken from recent U.S. Census Bureau estimates, and occupation switching in Third District and selected nearby metropolitan areas, calculated using recent samples from the CPS. Overall, workers in metropolitan areas with lower population density tend to be more likely to change occupations. (Of course, these are raw numbers, without some of the controls for other factors that vary across cities used in creating Figure 1.) For example, in our District, the Altoona, Vineland–Millville–Bridgeton, and Johnstown metropolitan areas have the highest average occupation-changing rates and also relatively low population densities. In contrast, the Trenton–Ewing metropolitan area has both the lowest rate of occupation changing and the highest population density of any metropolitan area in the Third District. Even within our region, some of the differences in density and productivity seem to be related to differences in the accumulation and preservation of specialized skills.

CONCLUSION

In this article, I have discussed new evidence for one potential source of agglomeration economies: better job search and matching. The broader agenda for this kind of work is to provide support for appropriate local policy choices. If urban productivity advantages are due mostly to job matching advantages, that may suggest that local development strategies that don't take advantage of these thick-market effects may not be effective. An important caution is that policy effects are likely to be small relative to the magnitudes needed for noticeable changes in local productivity. This can be seen in the persistence of city characteristics: Places that are densely populated or that have highly educated workforces also had similar characteristics in decades or even centuries past.

Finally, an important further step is to understand the relative importance of different sources of agglom-

TABLE

Occupation Switching in Third District Metropolitan Areas, 2005-2009

	Persons per square mile, 2007	Percent of workers switching occupations last year, 2005-2009 average
Third District Metropolitan Areas		
Trenton-Ewing, NJ	1,617.5	6.4
Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	1,258.8	10.9
Allentown-Bethlehem-Easton, PA-NJ	550.8	10.4
Atlantic City, NJ	482.4	8.8
Reading, PA	468.0	9.6
Harrisburg-Carlisle, PA	324.7	13.9
Vineland-Millville-Bridgeton, NJ	317.9	14.6
Scranton-Wilkes-Barre, PA	314.6	11.4
Lancaster, PA	267.4	8.7
Dover, DE	258.2	10.0
Altoona, PA	238.7	15.4
Johnstown, PA	210.7	14.3
Metropolitan Areas Outside the Third District		
New York-Northern New Jersey Long Island, NY-NJ-PA	2,797.6	10.1
Boston-Cambridge-Quincy, MA-NH	1,278.3	10.3
Cleveland-Elyria-Mentor, OH	1,045.9	9.1
Baltimore-Towson, MD	1,022.6	9.7
Washington-Arlington-Alexandria, DC-VA-MD-WV	943.0	10.5
Cincinnati-Middletown, OH-KY-IN	485.1	11.1
Pittsburgh, PA	446.2	12.7

Source: Author's calculations, U.S. Census Bureau, and the 2005-09 Current Population

eration economies. Stuart Rosenthal and William Strange, in their 2001 study, and Glenn Ellison, Edward Glaeser, and William Kerr have some intriguing early results in this area. Using industry locations as observations, Rosenthal and Strange compare a measure of spatial concentration with industry-location characteristics that proxy for the presence of knowledge spillovers, input sharing, natural advantages, and other types of agglomeration economies. Their results indicate that industry concentrations are correlated with a number of these measures, in particular, measures related to labor market concentration. Ellison, Glaeser, and Kerr adopt a similar methodology but use industry pairs as the unit of observation. Their results suggest that linkages between industries are an important reason for co-location patterns. Despite these early efforts, much remains unknown about this important question. One of the priorities for future work should be to assess the relative importance of different mechanisms.

REFERENCES

Baumgardner, James R. "Physicians' Services and the Division of Labor across Local Markets," *Journal of Political Economy*, 96:5 (1988), pp. 948-82.

Bleakley, Hoyt, and Jeffrey Lin. "Thick-Market Effects and Churning in the Labor Market: Evidence from U.S. Cities," Federal Reserve Bank of Philadelphia Working Paper 07-25 (2007).

Carlino, Gerald. "Knowledge Spillovers: Cities' Role in the New Economy," Federal Reserve Bank of Philadelphia *Business Review* (Fourth Quarter 2001), pp. 17-26.

Carlino, Gerald. "The Economic Role of Cities in the 21st Century," Federal Reserve Bank of Philadelphia *Business Review* (Third Quarter 2005), pp. 9-14.

Carlino, Gerald. "Beautiful City," Federal Reserve Bank of Philadelphia *Business Review* (Third Quarter 2009), pp. 10-17.

Carlino, Gerald, and Richard Voith. "Accounting for Differences in Aggregate State Productivity," *Regional Science and Urban Economics*, 22:4 (1992), pp. 597-617.

Chatterjee, Satyajit. "Agglomeration Economies: The Spark That Ignites a City?" Federal Reserve Bank of Philadelphia *Business Review* (Fourth Quarter 2003), pp. 6-13.

Ciccone, Antonio, and Robert E. Hall. "Productivity and the Density of Economic Activity," *American Economic Review*, 86:1 (1996), pp. 54-70.

Davis, Morris, Jonas D.M. Fisher, and Toni M. Whited. "Agglomeration and Productivity: New Estimates and Macroeconomic Implications," Working Paper, University of Wisconsin (2009).

De Blasio, Guido, and Sabrina Di Addario. "Do Workers Benefit from Industrial Agglomeration?" *Journal of Regional Science*, 45:4 (2005), pp. 797-827.

Diamond, Peter. "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90:5 (1982), pp. 881-94.

Duranton, Gilles, and Diego Puga. "Micro-foundations of Urban Agglomeration Economies," in J. Vernon Henderson and Jacques-Francois Thisse, eds., *Handbook of Regional and Urban Economics*, Volume 4. Amsterdam: North-Holland, 2004.

Ellison, Glenn, Edward L. Glaeser, and William Kerr. "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns," *American Economic Review*, 100:3 (June 2010), pp. 1195-1213.

Fallick, Bruce, Charles A. Fleischman, and James B. Rebitzer. "Job Hopping in Silicon Valley: Some Evidence Concerning the Micro Foundations of a High Technology Cluster," *Review of Economics and Statistics*, 88:3 (2006), pp. 472-81.

Fox, Jeremy T. "Labor Market Competition Using Compensation and Dynamic Incentive Schemes," manuscript, Stanford University (September 2002).

Garicano, Luis, and Thomas Hubbard. "Managerial Leverage Is Limited by the Extent of the Market: Hierarchies, Specialization, and the Utilization of Lawyers' Human Capital," *Journal of Law and Economics*, 50:1 (February 2007), pp. 1-43.

Glaeser, Edward L., and David C. Maré. "Cities and Skills," *Journal of Labor Economics*, 19:2 (2001), pp. 316-42.

Groen, Jeffrey. "Occupation-Specific Human Capital and Local Labor Markets," *Oxford Economic Papers*, 58:4 (2006), pp. 722-41.

Jacobs, Jane. *The Economy of Cities*. New York: Random House, 1969.

Kambourov, Gueorgui, and Iourii Manovskii. "Occupational Specificity of Human Capital," *International Economic Review*, 50:1 (2009), pp. 63-115.

Kim, Sunwoong. "Labor Specialization and the Extent of the Market," *Journal of Political Economy*, 97:3 (1989), pp. 692-705.

Lin, Jeffrey. "Technological Adaptation, Cities, and New Work," Federal Reserve Bank of Philadelphia Working Paper 09-17 (October 2009).

Murphy, Kevin M. "Specialization and Human Capital," Ph.D. thesis, University of Chicago (1986).

Neal, Derek. "Industry-Specific Capital: Evidence from Displaced Workers," *Journal of Labor Economics*, 13:4 (1995), pp. 653-77.

Parent, Daniel. "Industry-Specific Capital and the Wage Profile: Evidence from the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics," *Journal of Labor Economics*, 18:2 (2000), pp. 306-23.

Rosenthal, Stuart S., and William C. Strange. "The Determinants of Agglomeration," *Journal of Urban Economics*, 50 (2001), pp. 377-93.

Rosenthal, Stuart S., and William C. Strange. "Evidence on the Nature and Sources of Agglomeration Economies," in J. Vernon Henderson and Jacques-Francois Thisse, eds., *Handbook of Regional and Urban Economics*, Volume 4. Amsterdam: Elsevier, 2004, pp. 2119-71.

Inflation Dynamics and the New Keynesian Phillips Curve*

BY KEITH SILL

A 1977 amendment to the Federal Reserve Act states that the Fed's mandate is "to promote effectively the goals of maximum employment, stable prices, and moderate long-term interest rates." Moderate long-term interest rates require low and stable inflation. Monetary policymakers use instruments such as a short-term interest rate to guide the economy with the aim of achieving an inflation objective. To help guide their decisions, monetary policymakers benefit from having a reliable theory of how inflation is determined, one that relates the setting of their instrument to the unexpected events that hit the economy and consequently to the rate of inflation and other economic variables. In this article, Keith Sill examines a prominent theory of how inflation is determined, as articulated in what is called the New Keynesian Phillips curve. He also investigates some of the implications of the theory for the conduct of monetary policy.

Policymakers, economists, and the public generally agree that low and stable inflation is beneficial to the economy. Low and stable inflation makes it easier for households to plan their savings and investments and for firms to make production and investment decisions. It also helps to promote equity across members of society, since low-income households often do not have access to the financial instruments that help guard savings from being eroded by inflation.[1] Also, households and firms often write contracts that are stated in dollar amounts (nominal terms). A worker may, for example, sign a contract to work over the upcoming year for a fixed dollar amount. If inflation turns out to be higher than what was expected at the time the contract was made, the worker may find he is unable to purchase as many goods and services as planned because his inflation-adjusted income is lower than expected. Stable inflation would help mitigate such problems.

A 1977 amendment to the Federal Reserve Act codified the importance of low and stable inflation as a goal for monetary policymakers. The amendment states that the Fed's mandate is "to promote effectively the goals of maximum employment, stable prices, and moderate long-term interest rates." Moderate long-term interest rates require low and stable inflation, on average. But how does the Fed control inflation? It cannot simply dictate that the rate of price increase will be, say, 2 percent. Rather, monetary policymakers use instruments such as a short-term interest rate to guide the economy with the aim of achieving an inflation objective. To help guide their decisions, monetary policymakers benefit from having a reliable theory of how inflation is determined: a theory that relates the setting of their instrument to the unexpected events

Keith Sill is an assistant vice president and the director of the Real-Time Data Research Center in the Research Department of the Philadelphia Fed. This article is available free of charge at www.philadelphiafed.org/research-and-data/publications/.

*The views expressed here are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

[1] See the April 2007 speech by then-Governor Frederic S. Mishkin.

that hit the economy and consequently to the rate of inflation and to other economic variables of interest. With such a model in hand, policymakers can make informed decisions about the likely course of inflation and how to set an instrument such as the federal funds rate to achieve their inflation objectives.

In this article, we will examine a prominent theory of how inflation is determined, as articulated in what is called the New Keynesian Phillips curve. The theory ties current inflation to expected future inflation, a measure of firms' cost of production, and shocks that hit the economy. When embedded in a larger model of the economy that determines how inflation expectations are formed, the theory gives guidance to policymakers on how to meet their inflation goals. Consequently, we will also investigate some of the implications of the theory for the conduct of monetary policy.

A LITTLE HISTORY: INFLATION AND EMPLOYMENT

There is a long and storied history in macroeconomics about the relationship between inflation and real economic activity. In 1958, William Phillips wrote a paper on the empirical relationship between wage inflation and unemployment in the U.K. over the period 1861-1957. Phillips observed that when wage inflation was high, the unemployment rate tended to be low, and vice versa. This inverse empirical relationship seemed to suggest that there might be a stable, permanent tradeoff between wage inflation, or price-level inflation more generally, and the unemployment rate. If so, policymakers could stimulate the economy and lower the unemployment rate at the expense of somewhat higher inflation. Indeed, for the U.S. economy, there appeared to be a stable tradeoff between inflation and the unemployment rate in the 1960s (Figure 1).[2]

FIGURE 1

Phillips Curve in the 1960s

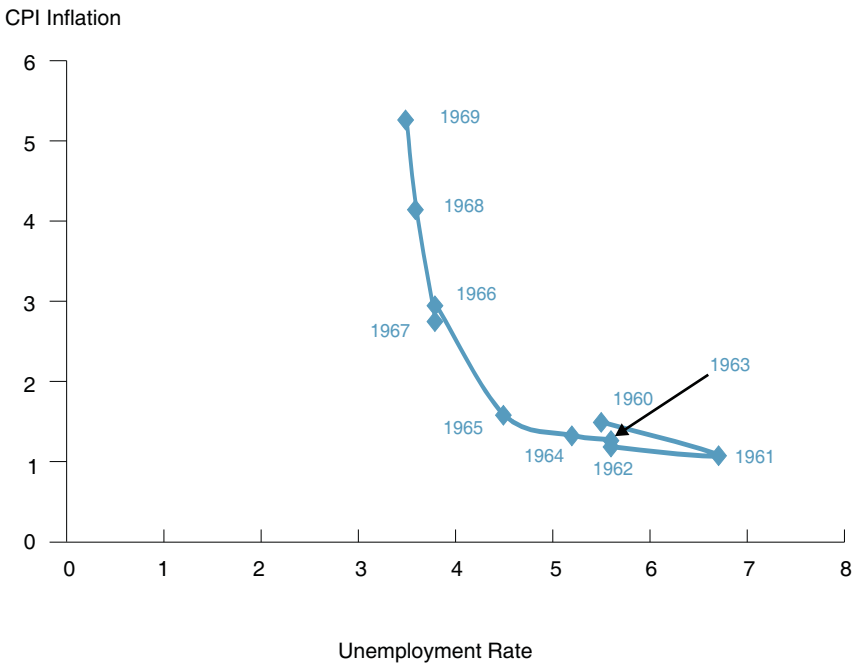

Unfortunately, the Phillips curve turned out to be not as stable as was first believed. The 1970s were a decade during which the economy experienced both high inflation and high unemployment rates, a development that came to be known as stagflation. Indeed, examining the entire span of data from the 1960s to the present, it is difficult to discern a tradeoff between inflation and unemployment. Rather than a negative one, the relationship between inflation and unemployment does not appear to be stable, and if anything, there seems to be a positive relationship between inflation and the unemployment rate (Figure 2).[3] Clearly, the relationship between the unemployment rate and inflation is not as simple as was first believed.

A key insight into the problem with the original Phillips curve's implication of a tradeoff between inflation and unemployment was made by Milton Friedman in his presidential address at the American Social Sciences Association meeting in 1968. Friedman observed that although the original Phillips curve traced out a relationship between money wages and the unemployment rate, what workers

[2] The relationship between inflation and the unemployment rate was an especially tight one in the 1960s. For another perspective on the Phillips curve that uses a longer history of data, see Figure 1, in Roc Armenter's article. See also the article by Jeffrey Lacker and John Weinberg for an accessible discussion of inflation, unemployment, and the Phillips curve.

[3] The episode of high inflation together with high unemployment during the 1970s (the black dots in Figure 2) came to be known as "stagflation." This period led to the recognition that the Phillips curve might not be stable. See Armenter's article for additional discussion.

FIGURE 2

Inflation and Unemployment 1960-2008

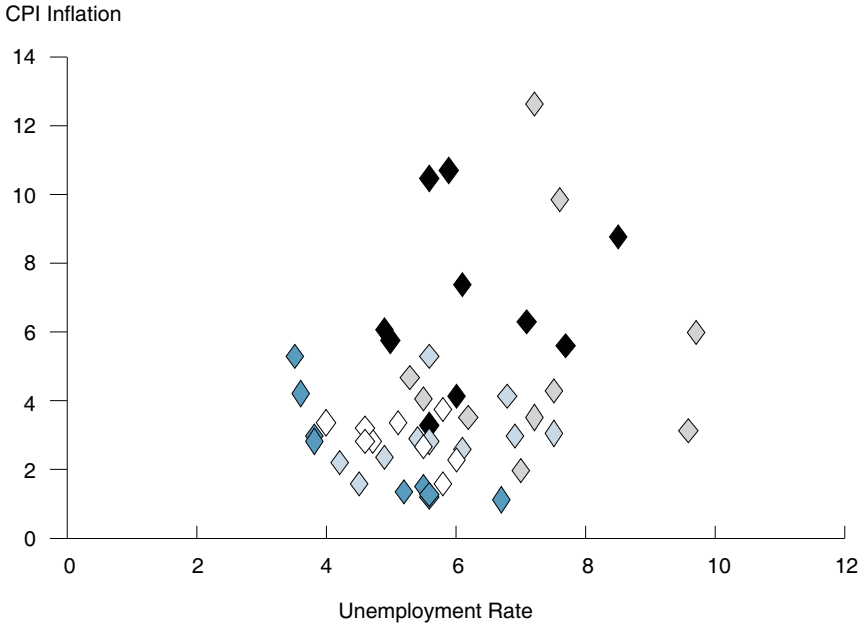

Color coded by decade: 1960s blue, 1970s black, 1980s gray, 1990s white, 2000s light blue

really cared about was their real wage — the wages they were paid relative to the prices they paid for goods and services. This implies that workers care about the expected rate of price increase, or inflation. If everyone expected prices to rise by 10 percent over the coming year, workers would try to negotiate a wage contract that called for at least a 10 percent increase in wages so that, in real terms, they would not be any worse off. Firms would be happy to pay the 10 percent increase because the real cost of labor is unchanged. Consequently, firms would not have an incentive to change employment. One would then expect to see money wages rising by 10 percent with no accompanying decrease in the unemployment rate. The implication is that in the long run, when expectations about price increases are factored in, there should be no exploitable tradeoff between inflation and unemployment.

Note, though, that if inflation turns out to be different than expected, the situation changes. If inflation rises by 15 percent when workers thought it would rise by 10 percent, workers would experience a decline in their inflation-adjusted wages and so would wish they had worked less. On the other hand, firms would have liked to hire additional workers at the lower real wage. If we assume that firms prevail and hire more workers at the existing wage, employment would increase, unemployment would decrease, and we get the Phillips curve relationship. But if expectations are correct, that inverse relationship between unemployment and inflation breaks down.

The view that a stable, inverse relationship between the unemployment rate and inflation disappears once a role for inflation expectations is introduced has gained support from empirical work that tries to predict future inflation using measures of economic activity such as the unemployment rate. The traditional Phillips curve suggests that inflation is related to the unemployment rate (actually its deviation from the economy's normal rate of unemployment). The implication of that theory is that unemployment rates will help to predict future inflation. Statistical analysis indicates that prior to the 1980s, such a relationship appeared to hold in the data: Measures of economic activity such as unemployment rates or fluctuations in output did help predict future inflation.[4] However, since the end of the 1970s, this no longer seems to be the case. Indeed, using data from 1980 onward, it appears that simply predicting that inflation in the next quarter will be what it is this quarter gives a forecast that is very hard to improve upon.[5] This finding is at least consistent with the view that inflation expectations are an important factor to consider when assessing the link between economic activity and inflation.

A NEW PHILLIPS CURVE

The traditional Phillips curve suggested that inflation and measures of economic activity were correlated, although the evidence for that theory now appears weak. Because the posited positive relationship between inflation and the unemployment rate

[4] Often, the output measure used is an output gap, which measures the deviations of real output from some proxy for potential real output, such as a long-term trend output. For more details about the output gap, see the article by Roc Armenter.

[5] See, for example, the papers by Andrew Atkeson and Lee Ohanian and James Stock and Mark Watson.

was based on historical correlations in the data, it faces several potential problems. For one, correlations in the data are likely to change if the structure of the economy changes. For example, if the Federal Reserve were to change the way it conducts monetary policy, it may well turn out that the correlation between inflation and economic activity in the data would change as well.[6] Indeed, such a change might be reflected in a shift in expected inflation. Furthermore, to predict how a change in monetary policy affects correlations in the data, we need a model of the economy that explicitly accounts for how the correlations among economic variables depend on the way monetary policy is set. With such a model in hand, the effects of a change in monetary policy (modeled as a change in the way policymakers respond to information) can be analyzed because the linkages between actions and outcomes are made explicit. One could then examine the model's predictions both before and after the monetary policy change to gauge the likely effects of the policy change on the economy. The key point is that simple correlations in the data are likely to change (and so become unstable) in response to a change in the fundamental structure underlying the economy.

We have also seen that the empirical evidence suggests that while the Phillips curve may have helped predict inflation prior to the 1980s, that relationship appears to have broken down since then. Obviously, models that predicted well in the past need not do so in the future, especially if there is a change in a fundamental factor such as monetary policy. To understand how structural changes to the economy affect empirical correlations, we need a theory of how the economic environment translates into correlations in the data.

The now dominant and workhorse model of monetary policy and business cycles is called the New Keynesian model. It is a structural model that delivers a theory of inflation that bears some resemblance to the traditional Phillips curve, but nonetheless, it has some significant differences. In principle, the model can help policymakers see how shocks to the economy and changes in the economic environment can translate into correlation in the data. In practice, however, this theory, like all economic theories, is a simplification of the actual economy and thus misses many potentially important linkages that are features of the real world. For example, the standard New Keynesian model does not have a well-developed financial sector and therefore has difficulty accounting for economic fluctuations prompted by financial crises.

The now dominant and workhorse model of monetary policy and business cycles is called the New Keynesian model.

THE NEW KEYNESIAN PHILLIPS CURVE MODEL

The New Keynesian Phillips curve is derived from a structural model of the economy that features two key elements. First, firms have some pricing power. That is, they can choose to sell more of their product by setting a lower price, or they can choose to sell a little less but at a higher price. (This is known as imperfect competition.) Second, firms choose to, or are only able to, adjust prices infrequently (sticky prices). They do not adjust their prices fully and immediately to every unexpected event that affects the economy. These two features of the model allow monetary policy to affect more than just prices and inflation in the short run.

Imperfect Competition and Sticky Prices. Imperfect competition means that firms have some power over their price-setting. This contrasts with perfect competition, a situation in which firms have no power to set prices. For example, a farmer bringing wheat to the market will have to take the price offered by buyers; he has virtually no power to demand a price higher than the prevailing market price and hope to attract customers. This is, in part, because he represents a small part of the overall supply of wheat and, in part, because other suppliers of wheat are selling a similar, if not identical product. If the farmer raised his price above the market price, his product would go unsold.

Contrast this with a large firm, such as Honda, that represents a significant share of its market. Honda is a relatively large part of the automobile industry and offers products distinct from those offered by other automakers. Consequently, Honda can set a price for its cars and see what the quantity of cars demanded is at that price. If Honda wants to sell more cars, it can lower the price per car. If it wants more profit per car, it can raise the price. The key point is that Honda has some pricing power, and it can use that power to gauge market demand for a car at a particular price point.

Imperfect competition is an important feature of models that

[6] This is an example of the "Lucas critique." See the article by Robert E. Lucas.

embed a New Keynesian Phillips curve. It allows firms to set a price for their products. The second key component of the models is that although firms can choose the price they set, they can only do so infrequently. This means that at least some prices are unable to immediately adjust in response to the shocks hitting the economy.

Suppose that, contrary to the assumptions of the New Keynesian model, all firms were able to, and did, adjust their prices instantaneously in response to shocks. Then monetary policy would have little influence on the nonmonetary, real side of the economy — consumption, output, and investment. Instead, monetary policy would only be able to affect the general price level, even in the short run. To see this, suppose monetary policy is implemented using an interest rate policy, such as is done in the U.S. If the Fed raised the short-term nominal interest rate and prices adjusted instantly, the rise in the nominal interest rate would be matched by a rise in expected inflation that would keep the real interest rate unchanged.[7] With an unchanged real interest rate, households and firms have no incentive to change their planned consumption and investment, and so the real side of the economy would be unaffected. The Fed controls inflation by changing the amount of liquidity in the economy, but it cannot influence real economic activity.

Suppose, though, that not all prices adjusted instantly in response to an unexpected event that hits the economy or a change in the monetary policy interest rate. This could happen, for example, if contracts are written in nominal terms for a fixed duration or if firms face costs of adjusting the prices they charge. In addition to adjusting the general amount of liquidity in the economy, the Fed now has an additional channel through which to influence inflation. If prices are sticky, expected inflation will not rise one-for-one with an increase in the nominal interest rate and as a consequence the real interest rate would rise too. The rise in the real interest rate leads households to boost their savings, since the return to savings is higher (and so households cut back a little bit on their consumption). Similarly, with higher real interest rates, firms want to borrow less to fund investment, since the cost of funds is now higher (consequently investment falls). With less demand for consumption and investment, real output for the economy is lower. Thus, with sticky prices, a Fed-induced rise in the nominal interest rate is contractionary for the real economy, at least in the short run when some prices do not fully adjust.

With sticky prices, a Fed-induced rise in the nominal interest rate is contractionary for the real economy, at least in the short run when some prices do not fully adjust.

HOW IS INFLATION DETERMINED?

We can now put together the two pieces — imperfect competition and infrequent price adjustment — to show how inflation is determined in the structural model according to the New Keynesian Phillips curve. With imperfect competition, firms seek to maximize profit by setting a price that is marked up over the marginal cost of producing an additional unit of product. So, with imperfect competition we might find that firms maximize profits when they set their prices 20 percent higher than the marginal cost. If firms set prices below this optimal price, quantity demanded rises and revenue increases, but that increase is outweighed by the rise in production costs and profits fall. If they set prices above the optimal price, the quantity demanded falls and the decline in revenue outweighs the decline in production costs, so again profits fall.

Now consider the implications of infrequent price adjustment. Firms recognize that the price set today will prevail for some time; they will not be able to reset prices in response to every development in the economy. Consequently, a firm that is trying to figure out the optimal price to set today will want to take into account not only what today's marginal cost of production is but also what the marginal cost of production is likely to be for the entire time frame over which it expects the price to prevail. For example, if the firm anticipated that it would not reset its prices until one year from now, it would want to estimate what marginal costs would be over the next year when setting prices today.

Thus, the interaction of imperfect competition and infrequent price adjustment leads firms to set product prices taking into account the expected future behavior of marginal costs. This requires firms to project future expected demand, marginal cost, and

[7] The real interest rate is equal to the nominal interest rate less expected inflation. Consequently, it is the expected return to savings after accounting for expected inflation.

future price levels (or inflation). In the stylized model, the solution to this problem is that a firm sets a price that is its desired markup over a weighted average of current and expected nominal marginal costs.

How is aggregate inflation — the change in the overall level of prices between two periods — determined? The price level in this period will be a combination of prices set by firms that are adjusting prices today and of prices set by firms that are not adjusting their prices in this period. This means that the level of inflation is determined by the fact that firms that reset their prices today choose a different price from the one they charged yesterday. Since firms that reset prices set them as a markup over marginal cost, we find that when we add up across firms to get the economy-wide price level with which to calculate inflation, it must reflect the anticipated path of future real marginal cost for the economy as a whole. The mathematics of the New Keynesian Phillips curve allows us to express the deviation of inflation from its long-run expected value, as a weighted sum of three components: (1) the expected deviation of next period inflation from its long-run expected value; (2) the deviation of real marginal cost from its long-run expected value; and (3) an error term representing unexpected events that lead firms to change their markups over marginal cost.[8]

[8] Derivations of the New Keynesian Phillips curve can be found in many advanced macroeconomic textbooks and survey articles. For one such derivation, see the book by Jordi Gali listed in the references. The form of the New Keynesian Phillips curve is given by: $\pi_t = \beta E_t \pi_{t+1} + kmc_t + \varepsilon_t$ where π_t is the deviation of inflation from its expected long-run value, $E_t \pi_{t+1}$ is the expected value today of the deviation of inflation tomorrow from its long-run expected value, mc_t is the deviation of marginal cost from its long-run expected value, and ε_t represents unanticipated events that cause firms to change their markup.

IMPLICATIONS OF THE NEW KEYNESIAN PHILLIPS CURVE FOR INFLATION AND ECONOMIC ACTIVITY

Both the New Keynesian Phillips curve and the traditional Phillips curve provide theories of how inflation is determined. However, the two theories differ in the role they assign to expected inflation as a determinant of current inflation and in the nonmonetary economic variables that are the important drivers of inflation.

Consider first how the theories differ in the economic activity variable that drives short-run movements in inflation. The New Keynesian Phillips curve suggests that the short-run dynamics of inflation are driven by the expected path of marginal cost. But remember that in the traditional Phillips curve, it is the unemployment rate that is driving inflation. While it's at least conceivable that the unemployment rate is correlated with marginal cost and thus serves as a good empirical proxy, it turns out that, based on the empirical evidence, the unemployment rate does not appear to be highly correlated with measures of marginal cost.

Figure 3 presents some evidence on this. As shown, the unemployment rate and real unit labor cost, a measure of marginal cost, do not exhibit a great deal of co-movement. Indeed, the simple correlation between the two series is about zero. Under the New Keynesian Phillips curve model, looking at unemployment rates as indications of inflation pressure is not the obvious thing to do.[9]

[9] There could be a higher correlation between unit labor costs and an unemployment rate gap measure if we defined the gap in a way such that the difference between actual unemployment rates and the economy's normal rate of unemployment moves in the right way. But usually we think of the normal rate of unemployment as being a slow-moving object (which is itself subject to great measurement uncertainty). Consequently, it is unlikely that the unemployment rate gap is highly correlated with unit labor costs.

FIGURE 3

Unit Labor Cost and Unemployment Rate

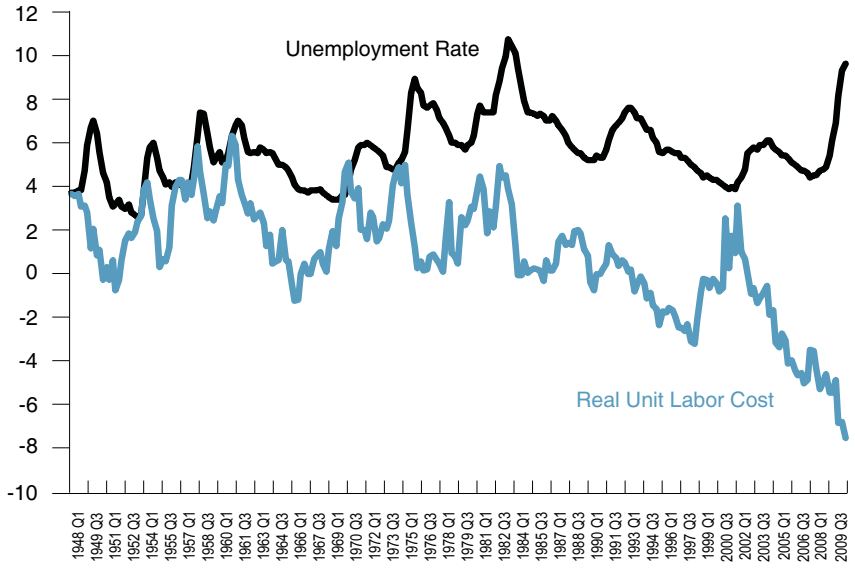

Unit labor is defined as total labor compensation divided by real output. We then deflate unit costs by the GDP implicit price deflator to translate it into real terms and take logs (and scale up by a factor of 100).

The evidence on the correlation of output gaps, which are measures of the level of real output less a measure of the level of potential real output, with unit labor costs is a bit more nuanced. Figure 4 shows a plot of the Congressional Budget Office measure of the output gap and real unit labor costs. In this figure, though, we have removed long-run fluctuations from the data and we focus instead on fluctuations over the span of the typical business cycle's duration (which is eight years or less).[10] The figure shows that at this "business cycle frequency," the correlation was negative up until the 1990s. However, over the past 15 years or so the correlation looks positive.

This may be somewhat encouraging for the use of output gaps in accounting for inflation. But there are several important measurement issues with these series. First, economists disagree on the best way to measure the output gap and different methods give rise to very different estimates of the size of the gap at a point in time. Furthermore, we can extract information about fluctuations in a series over business cycle frequencies only long after the fact — real-time measures of the business cycle component of a series are highly uncertain.

Another important difference between the two versions of the Phillips curve is the role they assign to expected inflation as a determinant of movements in inflation today. A key feature determining inflation under the New Keynesian Phillips curve theory is the implication that inflation anticipates, or leads, measures of economic activity. Inflation responds to higher levels of expected marginal cost and so rises today in anticipation of that higher cost. In contrast, empirically estimated traditional Phillips curves are often specified to include lagged values of economic activity. Such a specification could be justified in the New Keynesian Phillips curve framework if the lagged values were useful for predicting marginal cost in the future.

It is important to note that the basic New Keynesian Phillips curve as described above does not imply a high degree of correlation over time in inflation rates: The inflation process is not very persistent. Indeed, there is no persistence over and above that which would be associated with marginal cost. As an empirical matter, though, there does appear to be more inflation persistence in the U.S. data than what would be implied by the baseline New Keynesian Phillips curve model.[11] One way in which persistence can be introduced into the model is to assume that prices are indexed to inflation. Thus, firms that don't re-optimize their prices in a given period nonetheless move their prices up with the general level of inflation that prevails in the economy. This is a bit of a shortcut, since we might reasonably ask why firms would not just take the time to set prices optimally, since they are going to reset them in line with inflation anyway.

We can also introduce additional inflation persistence into the model by assuming that the expected long-run average rate of inflation changes slowly over time, as opposed to being constant. If the rate of inflation that policymakers are comfortable with changes over time, it would introduce a slow-moving component into the

[10] More technically, the plot shows the cycle component of the two series after the Hodrick-Prescott-filtered trend is removed from the data. The Hodrick-Prescott-filtered cycle represents fluctuations in the series at frequencies from zero to eight years.

[11] However, inflation persistence does not appear to be a pervasive feature of economies. See the paper by Luca Benati, who shows that the degree of inflation persistence varies across countries and within countries according to the monetary policy regime that is in place.

FIGURE 4

Output Gap and Unit Labor Costs

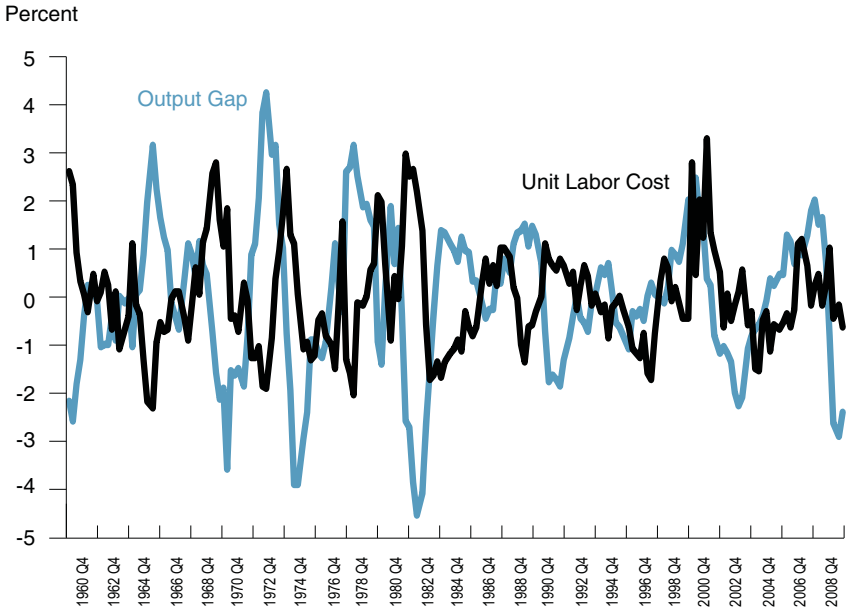

Output gap as measured by the Congressional Budget Office. Real unit labor costs are as described in the footnote to Figure 3. Both series are Hodrick-Prescott filtered, and the business cycle component is plotted.

inflation process and make actual inflation more persistent.[12]

MONETARY POLICY IMPLICATIONS OF THE NEW KEYNESIAN PHILLIPS CURVE: LESSONS AND CAVEATS

What can policymakers learn from the New Keynesian Phillips curve theory? One immediate implication is that unemployment rate gaps and output gaps should be used with caution when trying to assess inflation pressures in the economy. The theory implies that inflation is determined by expected future real marginal cost, and marginal cost does not appear to be highly correlated with unemployment rates or output gaps, as conventionally defined. Indeed, the paper by Jordi Gali and Mark Gertler argues that the New Keynesian Phillips curve with marginal cost as the measure of economic activity fits the data better than a traditional Phillips curve specification that uses output gaps.[13]

In addition to real activity measures, the New Keynesian Phillips curve suggests that expectations about the future are important for determining inflation today. For example, the theory indicates that monetary policy that is expected to be stimulative in the future can lead to higher inflation today. How does monetary policy end up being inflationary in this baseline model? Recall that inflation is given by the weighted sum of future real marginal costs. When monetary policymakers stimulate the economy by lowering interest rates, this action also stimulates demand. For firms to meet the higher demand, they must hire additional workers. Attracting additional workers requires a higher real wage — which raises the marginal cost of production for firms. Hence, firms that are re-optimizing their prices raise their prices today, and inflation ensues. The key point of contrast with the traditional Phillips curve model is that expectations of the future are an important component for inflation today.

Economic models that embed a New Keynesian Phillips curve tend to suggest that monetary policy can achieve about the best outcome possible when the policy interest rate responds aggressively to current or expected inflation: rising more than one-for-one when inflation rises, and falling more than one-for-one when inflation falls. The models also tend to suggest that the economy will be more stable if policymakers respond more aggressively to inflation developments than to developments in real activity such as unemployment rates and output gaps. The models do not suggest that developments in the real economy should necessarily be ignored, but policy should not respond too aggressively to them in a direct manner, since an aggressive policy response tends to promote further economic instability.[14]

Clearly, in the real world, monetary policymakers pay careful attention to developments in inflation and in output and employment. New Keynesian Phillips curve economic models make many simplifying assumptions, so their implications should be viewed with care. For example, firms' price-setting behavior, which, as we have seen, is a key component of the inflation process, is not very well understood and so is not modeled at a very deep level. The New Keynesian Phillips curve models tend to be at their most accurate when the economy is in "normal times" and behavior is not too far from average behavior. The models will not predict well, for example, in times of financial crisis, since the baseline New Keynesian model has no meaningful financial sector. This is not to suggest that New Keynesian Phillips curve models are not a useful part of the toolkit for monetary policymakers. They can help to clarify ideas about the transmission of shocks through the economy and point to likely determinants of economic variables such as inflation. However, empirically reasonable medium- and large-scale equilibrium models that embed the New Keynesian Phillips curve are at an early stage of development. Consequently, policymakers continue to be informed by a variety of models — both empirical and theoretical — as they consider how policy should best react to changes in the economy.

[12] This shows up in the New Keynesian Phillips curve as a persistent change in the long-run expected value of inflation. (Recall that the New Keynesian Phillips curve is an expression about the deviation of inflation from its expected long-run average value.)

[13] Note, though, that Galí and Gertler's study looked at the performance of the Phillips curve only up until the mid 1990s.

[14] See, for example, the article by Stephanie Schmitt-Grohe and Martin Uribe. In the standard New Keynesian model, targeting inflation helps to stabilize the impact of unexpected events on the economy and so leads indirectly to more stable output.

REFERENCES

Armenter, Roc. "Output Gaps: Uses and Limitations," Federal Reserve Bank of Philadelphia *Business Review* (First Quarter 2011).

Atkeson, Andrew, and Lee Ohanian. "Are Phillips Curves Useful for Forecasting Inflation?" Federal Reserve Bank of Minneapolis *Quarterly Review*, 25:1 (Winter 2001), pp. 2-11.

Benati, Luca. "Investigating Inflation Persistence Across Monetary Regimes," *Quarterly Journal of Economics* 123 (2008), pp. 1005-60.

Friedman, Milton. "The Role of Monetary Policy," *American Economic Review*, 58 (1968), pp. 1-17.

Galí, Jordi. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton: Princeton University Press, 2008.

Galí, Jordi, and Mark Gertler. "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics*, 44 (1999), pp. 195-222.

Lacker, Jeffrey, and John Weinberg. "Inflation and Unemployment: A Layperson's Guide to the Phillips Curve," Federal Reserve Bank of Richmond *Economic Quarterly*, 93 (2007), pp. 201-27.

Lucas, Robert E. "Econometric Policy Evaluation: A Critique," in K. Brunner and A.H. Meltzer, eds., *The Phillips Curve and Labor Markets*. Amsterdam: North-Holland, 1976, pp. 19-46.

Mishkin, Frederic S. "Monetary Policy and the Dual Mandate," speech (April 2007).

Phillips, A.W. "The Relationship Between Unemployment and the Rate of Change of Money Wages in the United Kingdom, 1861-1957," *Economica*, 25 (1958), pp. 283-99.

Schmitt-Grohe, Stephanie, and Martin Uribe. "Optimal Simple and Implementable Monetary and Fiscal Rules," *Journal of Monetary Economics*, 54 (2007), pp. 1702-25.

Stock, James, and Mark Watson. "Phillips Curve Inflation Forecasts," NBER Working Paper 14322 (September 2008).

RESEARCH RAP

Abstracts of research papers produced by the economists at the Philadelphia Fed

You can find more Research Rap abstracts on our website at: www.philadelphiafed.org/research-and-data/publications/research-rap/. Or view our working papers at: www.philadelphiafed.org/research-and-data/publications/.

EXTENDED BENEFITS AND UNEMPLOYMENT TRANSITION RATES

Using the monthly CPS, the author estimates unemployment-to-employment (UE) transition rates and unemployment-to-inactivity (UN) transition rates by unemployment duration for male workers. When estimated for the period of 2004-2007, during which no extended benefits are available, both of the transition-rate profiles show clear patterns consistent with the expiration of regular benefits at 26 weeks. These patterns largely disappear in the profiles for the period of 2009-2010, during which large-scale extensions have become available. The author conducts counterfactual experiments in which the estimated profiles for 2009-2010 are replaced by the hypothetical profiles inferred from the ones for 2004-2007. The results indicate that the benefit extensions in recent years have raised male workers' unemployment rate by 0.9-1.7 percentage points. Roughly 50-60 percent of the total increase is attributed to the effects on UE transition rates and the remaining part is accounted for by the effects on UN transition rates.

Working Paper 10-35, "Effects of the UI Benefit Extensions: Evidence from the Monthly CPS," Shigeru Fujita, Federal Reserve Bank of Philadelphia

HOW INVENTORIES AFFECT TRADE, INFORMATION, AND PRICES

The authors study trade between a buyer and a seller when both may have existing inventories of assets similar to those being traded. They analyze how these inventories affect trade, information dissemination, and price formation. The authors show that when the buyer's and seller's initial leverage is moderate, inventories increase price and trade volume, but when leverage is high, trade may become impossible (a "market freeze"). Their analysis predicts a pattern of trade in which prices and trade volume first increase, and then markets break down. The authors use their model to discuss implications for regulatory intervention in illiquid markets.

Working Paper 10-36, "Market Run-Ups, Market Freezes, and Leverage," Philip Bond, University of Minnesota, and Yaron Leitner, Federal Reserve Bank of Philadelphia

EXPLAINING LIFE-CYCLE PATTERNS OF HOUSEHOLDS' TIME USE AND CONSUMPTION

The authors incorporate home production in a dynamic general equilibrium model of consumption and saving with illiquid housing and a collateralized borrowing constraint. They show that the model is capable of explaining life-cycle patterns of households' time use and consumption of different categories. Specifically, households' market hours and home hours are fairly stable

early in the life cycle. Market hours start to decline sharply at age 50, while home hours begin to increase at age 55. Households' consumption of the market good, home input, and housing services all exhibit hump shapes over the life cycle, with the market good having the most pronounced hump, followed by the home input, and then housing services. A plausibly parameterized version of the authors' model predicts that the interaction of the labor efficiency profile and the availability of home production technology explain households' time use over the life cycle. The resulting income profiles, the endogenous borrowing constraint, and the presence of home production account for the initial hump in all three consumption goods. The consumption profiles in the second half of the life cycle are mostly driven by the complementarity of home hours, home input, and housing in home production.

Working Paper 0-37, "Consumption and Time Use over the Life Cycle," Michael Dotsey, Federal Reserve Bank of Philadelphia; Wenli Li, Federal Reserve Bank of Philadelphia; and Fang Yang, State University of New York at Albany

DATA REVISIONS AND THE STATISTICAL (UN)RELIABILITY OF MEASURES OF PRODUCTIVITY GROWTH

Productivity growth is carefully scrutinized by macroeconomists because it plays key roles in understanding private savings behavior, the sources of macroeconomic shocks, the evolution of international competitiveness, and the solvency of public pension systems, among other things. However, estimates of recent and expected productivity growth rates suffer from two potential problems: (i) recent estimates of growth trends are imprecise, and (ii) recently published data often undergo important revisions.

This paper documents the statistical (un)reliability of several measures of aggregate productivity growth in the U.S. by examining the extent to which they are revised over time. The authors also examine the extent to which such revisions contribute to errors in forecasts of U.S. productivity growth.

The authors find that data revisions typically cause appreciable changes in published estimates of productivity growth rates across a range of different productivity measures. Substantial revisions often occur years after the initial data release, which they argue contributes significantly to the overall uncertainty policymakers face. This emphasizes the need for means of reducing the uncertainty facing policymakers and policies robust to uncertainty about current economic conditions.

Working Paper 11-1, "Lessons from the Latest Data on U.S. Productivity," Jan P.A.M. Jacobs, University of Groningen, and Simon van Norden, HEC Montreal and Visiting Scholar, Federal Reserve Bank of Philadelphia

TERMS OF CREDIT IN A COMPETITIVE MARKET

The author studies the terms of credit in a competitive market in which sellers (lenders) are willing to repeatedly finance the purchases of buyers (borrowers) by engaging in a credit relationship. The key frictions are: (i) the lender is unable to observe the borrower's ability to repay a loan; (ii) the borrower cannot commit to any long-term contract; (iii) it is costly for the lender to contact a borrower and to walk away from a contract; and (iv) transactions within each credit relationship are not publicly observable. The lender's optimal contract has two key properties: delayed settlement and debt forgiveness. Asymmetric information gives rise to the property of delayed settlement, which is a contingency in which the lender allows the borrower to defer the repayment of his loan in exchange for more favorable terms of credit within the relationship. This property, together with the borrowers' lack of commitment, gives rise to debt forgiveness. When the borrower's participation constraint binds, the lender needs to "forgive" part of the borrower's debt to keep him in the relationship. Finally, the author studies the impact of the changes in the initial cost of lending on the terms of credit.

Working Paper 11-2, "A Dynamic Model of Unsecured Credit," Daniel R. Sanches, Federal Reserve Bank of Philadelphia

STRATEGIC FACTORS AND THE DECISION TO DEFAULT ON FIRST VS. SECOND LIEN MORTGAGES

Strategic default behavior suggests that the default process is not only a matter of an inability to pay. Economic costs and benefits affect the incidence and timing of defaults. As with prior research, the authors find that people default strategically as their home value falls below the mortgage value (exercise the put option to default on their first mortgage). While some

of these homeowners default on both first mortgages and second lien home equity lines, a large portion of the delinquent borrowers have kept their second lien current during the recent financial crisis. These second liens, which are current but stand behind a seriously delinquent first mortgage, are subject to a high risk of default. On the other hand, relatively few borrowers default on their second liens while remaining current on their first. This paper explores the strategic factors that may affect borrower decisions to default on first vs. second lien mortgages. The authors find that borrowers are more likely to remain current on their second lien if it is a home equity line of credit (HELOC) as compared to a closed-end home equity loan. Moreover, the size of the unused line of credit is an important factor. Interestingly, they find evidence that the various mortgage loss mitigation programs also play a role in providing incentives for homeowners to default on their first mortgages.

Working Paper 11-3, "Strategic Default on First and Second Lien Mortgages During the Financial Crisis," Julapa Jagtiani, Federal Reserve Bank of Philadelphia, and William W. Lang, Federal Reserve Bank of Philadelphia

OPTIMAL MONETARY POLICY WHEN FIAT MONEY AND PRIVATE DEBT COEXIST

The authors study optimal monetary policy in a model in which fiat money and private debt coexist as a means of payment. The credit system is endogenous and allows buyers to relax their cash constraints. However, it is costly for agents to publicly report their trades, which is necessary for the enforcement of private liabilities. If it is too costly for the government to obtain information regarding private transactions, then it relies on the public information generated by the private credit system. If not all private transactions are publicly reported, the government has imperfect public information to implement monetary policy. In this case, the authors show that there is no incentive-feasible policy that can implement the socially efficient allocation. Finally, they characterize the optimal policy for an economy with a low record-keeping cost and a large number of public transactions, which results in a positive long-run inflation rate.

Working Paper 11-4, "Optimal Monetary Policy in a Model of Money and Credit," Pedro Gomis-Porqueras, Australian National University, and Daniel R. Sanches, Federal Reserve Bank of Philadelphia