Labor Market Anxiety and the Downward Trend in the Job Separation Rate*

BY SHIGERU FUJITA

Anecdotal evidence suggests that labor market conditions surrounding American workers had been worsening in recent decades, even before the severe recession in 2007-2009. However, studies by academic researchers have not found clear evidence that worker turnover has increased over time. In this article, Shigeru Fujita shows that there is a long-run downward trend in the separation rate into unemployment and examines several factors that help account for this long-run decline. He argues that the aging of the labor force has played an important role in the trend. He also explains, using an economic model, how the declining separation rate can result from workers' response to the increased sense of job insecurity.

Anecdotal evidence suggests that labor market conditions surrounding American workers had been worsening in recent decades, even before the severe recession in 2007-2009. The following quote from an article in the *New York Times* characterizes the sentiment of American workers: "As workers' job security has evaporated, so has their bargaining power — their ability to ask for more money, more vacation time, more health benefits. Across the nation, and across industries, employees perceive that they are more vulnerable to dismissal now than in the past" (July 3, 1995).

A notable thing about this quote is that this article was published in 1995, nearly four and half years after the shallow and short recession in 1990-91. The average unemployment rate was 5.6 percent in 1995, and thus, the labor market in 1995 was by no means weak from the viewpoint of the level of the unemployment rate.

Academic researchers have also studied this issue of job security more formally.[1] One intuitive approach they've taken is to examine whether there is any upward trend in worker turnover rates. The idea is that increased job insecurity should be reflected in higher worker separations in the data. Interestingly, however, these studies have not found clear evidence that worker turnover has increased over time, despite the view exemplified in the above quote.

There are a number of ways to measure worker turnover, but one relevant measure for the issue of job security is the separation rate into unemployment. This measure is constructed by calculating the number of people who lost their jobs in a given month as a fraction of the total number of employed workers.

Figure 1 presents the separation rate over the last three decades. There are several interesting patterns. First, the separation rate into unemployment increases during recessions. This is not surprising given that firms shed more workers during recessions.[2] Second, while this "counter-cyclicality" is clear in the data, the separation rate has been gradually declining over time. Third, even though the separation rate increased sharply during the Great Recession, its peak was lower than the level we saw during the

Shigeru Fujita 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.

[1] See, for example, the special issue of the *Journal of Labor Economics* in 1999. The entire issue is devoted to job security.

[2] See my 2007 *Business Review* article for a summary of fluctuations of the job separation rate and job finding rates over the business cycle. The focus of this current article is on the long-term trend of the job separation rate.

FIGURE 1

Aggregate Separation Rate

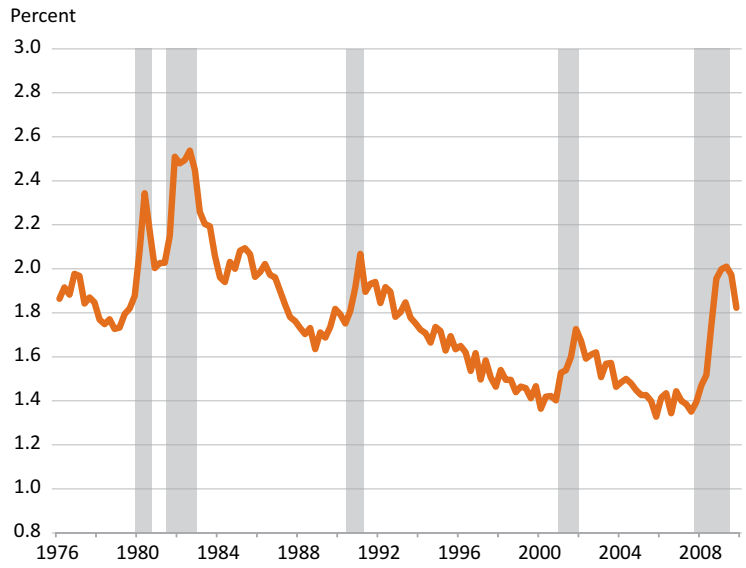

Notes: Author's calculations using CPS basic monthly data. The numbers plotted represent the rate at which employed workers become unemployed per month, expressed as quarterly averages. Grey bars represent NBER recession dates. Last date plotted: 2009/Q4.

recessions in the early 1980s. This is quite surprising given the severity of the Great Recession.

The focus of this article is on the secular decline in the separation rate. Does it imply that labor market conditions concerning job security have improved over time, as opposed to the view often found in the popular press, such as the one quoted at the beginning of this article?

In what follows, I will examine several factors that help account for the long-run decline in the separation rate. The first is the aging of the workforce. I will show empirically that aging has contributed significantly to the declining separation rate. The second explanation is based on a declining trend in business volatility. I present a popular labor market model, called the labor-matching model, to describe how the decline in business volatility lowers the separation rate. These two explanations, however, do not directly speak to the increased sense of job insecurity. The last explanation, which directly addresses this issue, argues that the lower separation rate is actually a result of an increased sense of job insecurity. This somewhat counterintuitive result is explained in an extended version of the labor-matching model.

AGING OF THE LABOR FORCE

Let's start with the aging of the labor force. The share of older workers in the workforce has increased in the last three decades. Aging affects the separation rate because older workers tend to have a stronger attachment to their employers. In other words, workers "shop around for jobs" when they are young, until they eventually settle into a job they like. This career pattern implies that a larger share of older workers reduces the separation rate in the aggregate.

The table on page 3 presents the average separation rate by demographic groups together with the employment share of each group. It presents the numbers for each of the three decades starting from the 1980s. First, let's compare separation rates across different demographic groups. Throughout the 30-year period, young workers (that is, workers younger than 25 years old) always have the highest separation rate. This is true for both genders. Second, one can see that the employment share of older workers has increased since the 1980s. Although employment shares of prime-age workers (workers who are between 25 and 54 years old) declined in the 2000s after increasing in the 1990s, the employment share of young workers declined and that of old workers (that is, workers who are older than 54 years) increased consecutively over the three decades. These changes in the employment shares by themselves reduce the aggregate separation rate. However, one important point to recognize here is that even if one focuses on the trend within each demographic group, the separation rate has been on a declining trend over this 30-year period, save for the separation rate of female workers older than 55 between the 1990s and 2000s. The fact that separation rates are declining even within demographic groups implies that the aging of the labor force cannot be the sole reason for the declining separation rate over the last three decades, as displayed in Figure 1.

But how much of the decline in the aggregate separation rate can be explained by the aging of the labor force? To get a sense, we can calculate the so-called "fixed-weight" separation rate. Note that the observed aggregate separation rate can be thought of as a weighted average of the separation rates of the six demographic groups, where employment shares at each moment are used as weights. In the fixed-weight separation rate, the em-

ployment shares are fixed at the levels at one particular time throughout the sample period. Because employment shares are fixed, this measure is not influenced by the changing demographic composition.[3]

Figure 2 plots the fixed-weight separation rate by fixing the employment shares at the level in 1976-78, together with the observed aggregate separation rate that was also plotted in Figure 1. Figure 2 indicates that the separation rate would have stayed higher than the actual level if the employment share did not change over the past three decades. Therefore, the difference between the two series can be thought of as the effect of the changes in demographics. According to this comparison, roughly one-half of the decline in the aggregate separation rate can be attributed to the aging of the labor force. This is arguably a large contribution, but it also implies that there are other causes as well.

CHANGES IN INDUSTRY STRUCTURE

Another important thing that has changed significantly in the U.S. labor market is that the employment share of the manufacturing sector has shrunk significantly, while service-sector employment has increased its share. This can also explain the declining separation rate if the separation rate in the manufacturing sector tends to be higher than that in the nonmanufacturing sector. Figure 3 presents the separation rates for the two sectors. One can see from the difference between the two series that the separation rate of the manufacturing sector responds more sharply to busi-

[3] Note that this measure is not insensitive to which period is used to fix the employment shares. But we can also calculate a more sophisticated measure, the so-called chain-weighted index, which does not have this problem. Using the chain-weighted separation rate gives the same result.

TABLE

Separation Rate and Employment Share by Age and Gender

	Male			Female		
	16 - 24	25 - 54	55 -	16 - 24	25 - 54	55 -
1980 - 1989	4.79	1.91	0.99	3.23	1.37	0.82
	(10.11)	(37.89)	(8.04)	(9.17)	(29.24)	(5.55)
1990 - 1999	4.18	1.59	0.99	2.99	1.19	0.82
	(8.07)	(39.15)	(6.87)	(7.31)	(33.19)	(5.40)
2000 - 2009	3.75	1.54	1.01	2.66	1.13	0.88
	(7.20)	(37.52)	(8.68)	(6.74)	(32.33)	(7.53)

Notes: Both separation rates and employment shares are expressed as percent. The employment share of each demographic group in parenthesis is based on the monthly CPS Table A-1.

FIGURE 2

Separation Rates: Effect of Demographic Changes

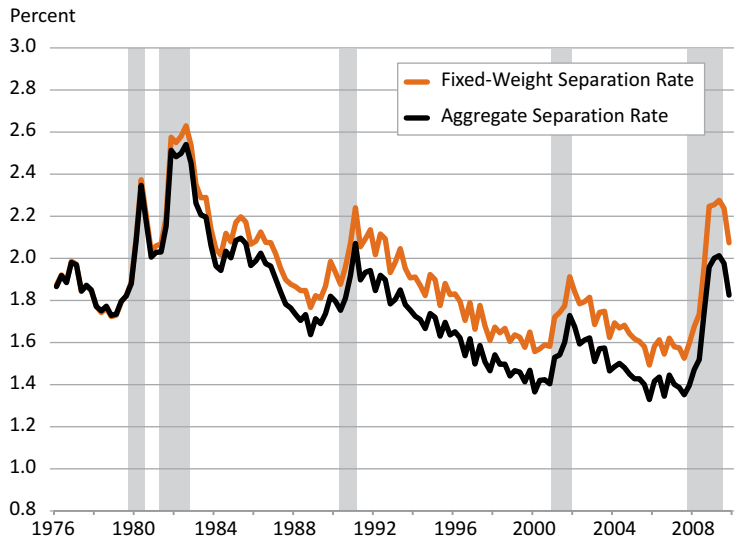

Notes: Author's calculations using CPS basic monthly data. See notes to Figure 1 for the definition of the separation rate. In constructing the fixed-weight separation rate, employment shares of six demographic groups are fixed at the average levels in 1976-1978. Grey bars represent NBER recession dates. Last date plotted: 2009/Q4.

FIGURE 3

Separation Rates by Industry

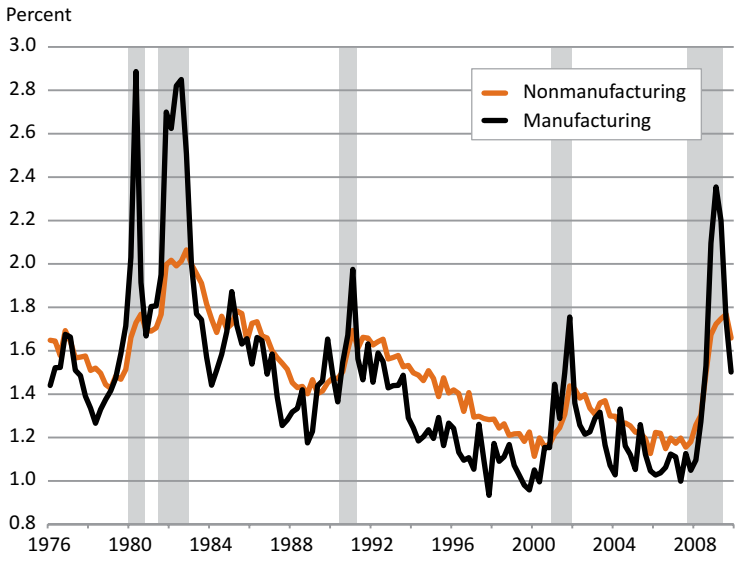

Notes: Author's calculations using CPS basic monthly data. See notes to Figure 1 for the definition of the separation rate. This figure breaks down the aggregate separation rate into separation rates for manufacturing and nonmanufacturing sectors. Grey bars represent NBER recession dates. Last date plotted: 2009/Q4.

ness cycles. However, the overall levels of the separation rates of the two sectors are quite similar and the separation rates of both sectors have been on a similar downward trend. This implies that the shrinking employment share of the manufacturing sector by itself does not constitute a major reason for the downward trend of the aggregate separation rate.[4]

[4] Yet another possibility is an increase in educational attainment. In particular, college enrollment rates have increased significantly. Unfortunately, it is not possible to conduct the same accounting exercise for this dimension of the data, as pointed out by Robert Shimer (see his paper published in 1998). The reason is that the characteristics of the workforce within the same education group are unlikely to be the same between the early 1980s and 2000s. For example, college degrees may have been valued more in the early 1980s, a time when fewer workers were college graduates. However, this may no longer be true in the 2000s, when a much larger fraction of the population graduates from college, which implies that the average characteristics of the workforce that comes under the category of college graduates have changed over time. This last fact makes it difficult to interpret the long-run trend of the separation rate even within the same educational group. Assessing the effects of increasing educational attainment on the separation rate requires an in-depth analysis based on an economic model.

DECLINES IN BUSINESS VOLATILITY

A recent paper by Steven Davis, Jason Faberman, John Haltiwanger, Ron Jarmin, and Javier Miranda provides an alternative story. These authors relate the declining separation rate to the decline in business volatility. Here "business volatility" can be thought of as uncertainty facing firms, and their explanation is based on the idea that the uncertainty has declined over time and thus the separation rate has also declined.[5]

[5] The idea that uncertainty has declined over time may sound odd to some readers given the current economic conditions in the U.S. Their argument, however, is based on a long-run decline in uncertainty, and their paper was written before the Great Recession.

They construct several different measures of business volatility, and one of them is constructed as a dispersion (standard deviation) of employment growth rates across firms. More specifically, they first calculate employment growth between two consecutive years at each establishment and then calculate how dispersed growth rates are across establishments by calculating the standard deviation. This dispersion measure can be computed for each year to obtain a time series of business volatility.[6] They find that the dispersion measure did indeed decline over the period 1977-2005.

Interestingly, the downward trend that both the dispersion measure and the separation rate have been on also holds at the industry level. That is, the authors calculate the dispersion measure and the separation rate for eight different industries and find that the relationship holds in most of these industries.[7]

The economic mechanism relating business volatility and the separation rate can be understood intuitively. When "shocks" facing businesses become smaller, job destruction is less likely to occur, thus reducing the separation rate. The mechanism can be described more formally in an economic model called a labor-matching model, developed by Dale Mortensen and Christopher Pissarides, two of the three Nobel Prize winners in economics in 2010. In the following section, I will use a version of this model again, so let me spend some time explaining the basic structure of the model.

[6] Note that establishment-level employment growth rates are weighted by using the number of employees in the establishment. Their measure also incorporates entry and exit of establishments as well.

[7] These eight industries are mining, construction, nondurable goods manufacturing, durable goods manufacturing, transportation and utilities, retail and wholesale trade, FIRE (finance, insurance, and real estate), and services.

LABOR-MATCHING MODEL

This model analyzes a situation in which there are many employment relationships between an employer and a worker, called "matches." Each match's profitability changes over time, say, due to changing demand. In this model, termination of an employment relationship occurs when the profitability of the match goes below a certain threshold level. An important thing to notice is that this decision to terminate a job takes into account future possibilities. For example, the firm may not let the worker go even if profits temporarily turn negative because finding a new worker is time-consuming and costly. One can show in this model that when uncertainty regarding future demand decreases, the likelihood that job separation will occur declines, a result that translates into a decline in the observed separation rate. The main reason is that the decreased uncertainty makes it less likely that profits will fall below the threshold level.

The explanation by Davis and co-authors is certainly plausible in light of the so-called Great Moderation, a term that refers to the period of low volatility from the early 1980s through the mid-2000s. However, Davis and co-authors' main focus is on the uncertainty facing firms and does not directly examine the uncertainty facing workers, as indicated by the increased sense of job insecurity alluded to in the introduction.

WAGES AND JOB SEPARATION IN A JOB-MATCHING MODEL

My recent working paper proposes the explanation that the increased sense of job insecurity is actually a source of the declining separation rate.[8] Before getting into the details, let me first discuss how wages are determined and how that interacts with the job separation decision in the basic labor-matching model.

As briefly mentioned above, the key idea of the labor-matching model is that it takes time for the worker to find a new job and for the firm to find a new worker. This is called the search friction. An important implication of the search friction is that wages can deviate from workers' productivity. That is, in a hypothetical economy without the search friction, workers can find a better job opportunity immediately if the current wage is lower than their productivity. Similarly, the firm will never pay wages higher than the worker's productivity because it can immediately find a similarly productive worker who is willing to work at a wage lower than this, that is, a wage equal to her productivity.

A slack labor market or a worse alternative opportunity makes workers feel insecure about separating into unemployment, and consequently, they stay with their current employer longer.

In the presence of the search friction, wages can be lower than the worker's productivity when the worker has a strong desire to stay with his or her current employer. For example, when it takes a long time for a worker to find the next job, the worker wants to stay with the current employer rather than become unemployed and search for the next job. This implies a lower separation rate. Moreover, it also means that the worker is willing to accept a lower wage. The same thing could happen when the alternative opportunity to his or her current job (for example, the wage that he or she can expect from a future employer) is not good for the worker.

A slack labor market or a worse alternative opportunity makes workers feel insecure about separating into unemployment, and consequently, they stay with their current employer longer. We will see that when workers face a higher possibility of losing their skills by separating from their current employer, this fear of losing their skills and suffering a wage drop translates into a lower separation rate, which seems consistent with the data we discussed at the beginning of the article.

SKILL LOSS IN THE LABOR-MATCHING MODEL

The extension of the labor-matching framework to include the possibility of skill loss takes the form of workers losing their skills during the period of job search (that is, while they're unemployed). Prominent examples in this vein include papers by Lars Ljungqvist and Thomas Sargent and by Wouter Den Haan, Christian Haefke, and Garey Ramey. Introducing this feature is important in that it allows researchers to replicate a well-known empirical fact: that wages tend to be lower at a worker's new job after he or she has gone through a period of unemployment (see, for example, the paper by Louis Jacobson, Robert LaLonde, and Daniel Sullivan). Furthermore, the literature has shown that declines in wages are often associated with a loss of skills. In my earlier *Business Review* article with Vilas Rao, we studied the experience of workers who lost their jobs around the 2001 recession, and we found that those workers who switched occupations or industry suffered a

[8] See my working paper.

particularly large drop in wages. Our result is consistent with the findings in the existing literature that worker's skills are tied closely to the experience in a certain occupation or industry.[9] Using the model that includes the possibility of skill loss, my working paper analyzes how the fear of losing skills can interact with the job separation decision and wage determination. Specifically, workers accumulate the skill that is specific to their job, but they may lose the skill once they are out of work. The key experiment in my paper is to see the effects of a higher risk of skill loss. What does the higher risk of skill loss represent in the real world? The labor-matching model I used in my experiment does not specify the underlying sources, but these sources can readily be associated with familiar phenomena, such as a rising tide of globalization or rapid technological progress, resulting in more jobs being outsourced to low-wage countries. The question is: How does workers' behavior change when facing a new environment in which workers can lose their skills faster when they are out of work?[10]

TRADE-OFF BETWEEN JOB SECURITY AND WAGE INCREASE

The result of higher skill loss is that both the job separation rate and wages decline. Recall that an important determinant of job separation and wages in the labor-matching framework is the value of opportunities available to the worker outside the current employment relationship. A lower value of outside opportunities lowers wages and the chance of job separation in the current employment relationship. The higher risk of losing skills means that the value of outside opportunities for currently employed workers is smaller. Because workers face an increased chance of ending up in a job that pays less, they become more willing to accept lower wages (or to give up a pay raise) in exchange for keeping their current job.

Recall that the model with skill loss replicates the aforementioned empirical fact that workers often end up with a job that pays less than their previous job. Workers who accumulated experience in a certain occupation or industry lose skills after a job loss and are hired only as inexperienced workers in a different industry or occupation. However, in a new environment in which the risk of skill loss has increased, experienced workers will have to accept lower wages in their current match, and consequently, there will be a smaller drop in wages should a separation occur.

> Fear of losing their skills makes workers reluctant to separate from their current employer and more willing to forgo wage growth.

A recent paper by Henry Farber computes the average earnings losses of job losers using a data set called the Displaced Workers Survey.[11] He presents the average earnings losses since the early 1980s. He calculates the average decline in real weekly earnings in each of the 14 surveys since 1984, including the 2010 survey. The result is that the series does not show an easily discernible downward or upward trend. Thus, the evidence on wage loss is not completely consistent with the model's prediction. However, what is somewhat surprising is not the lack of a downward trend in the size of earnings losses but the lack of an upward trend, which could be due to the mechanism highlighted in the model that there is less room for wages to drop further.

In summary, the explanation I have discussed emphasizes the trade-off between workers' willingness to accept wage cuts (or slow wage growth) and keeping their job: By accepting lower wages, workers can hold on to their jobs. Importantly, it is consistent with the fact that real wages have been stagnant during the period of declining separation rates. One may recall a puzzle in the late 1990s that, even though the labor market appeared to be tight, real wage growth was quite subdued. The following quote from a speech by former Fed Chairman Alan Greenspan offers a clear intuition that corresponds to the implications of the model: "A sense of increasing skill obsolescence has also led to an apparent willingness on the part of employees to forgo wage and benefit increases for increased job security. Thus, despite the incredible tightness of labor markets, increases in compensation per hour have continued to be relatively modest" (October 1, 1998).

CONCLUSION

This article discussed possible sources of the long-run downward trend in the job separation rate. First, the aging of the workforce is one of the main reasons for the trend: An older labor force implies that the labor force, on average, has a stronger attachment

[9] In other words, the skills can be useful as long as a worker stays in the same occupation, even if the worker changes jobs. See, for example, the papers by Derek Neal, and by Gueorgui Kambourov and Iourii Manovskii.

[10] An alternative interpretation is that workers face a higher risk that they will not find a job that uses the skills familiar to them.

[11] The Displaced Workers Survey is conducted every two years. The purpose of the survey is to study the experience of displaced workers, including earnings before and after the displacement.

to employers and thus lowers the separation rate. The second source studied by Davis and co-authors is declining business volatility: Decreased uncertainty makes it less likely that job separation occurs. These two explanations do not directly address the increased sense of job insecurity among American workers.

The third explanation is based on the trade-off between wages and job security: Fear of losing skills makes workers more willing to accept lower wages in exchange for keeping their current jobs. This explanation reconciles the coexistence of stagnant wage growth and the lower job separation rate. An important general point of this last explanation is that gauging job insecurity based solely on the level of labor turnover can be a misleading exercise.

REFERENCES

Davis, Steven, Jason Faberman, John Haltiwanger, Ron Jarmin, and Javier Miranda. "Business Volatility, Job Destruction and Unemployment," *American Economic Journal: Macroeconomics*, 2:2 (2010), pp. 259–87.

Den Haan, Wouter, Christian Haefke, and Garey Ramey. "Turbulence and Unemployment in a Job Matching Model," *Journal of the European Economic Association*, 3:6 (2005), pp. 1360–85.

Farber, Henry. "Job Loss in the Great Recession: Historical Perspective from the Displaced Workers Survey, 1984-2010," Princeton University Industrial Relations Section Working Paper 564 (2011).

Fujita, Shigeru. "What Do Worker Flows Tell Us About Cyclical Fluctuations in Employment?," Federal Reserve Bank of Philadelphia *Business Review* (Second Quarter 2007).

Fujita, Shigeru. "Declining Labor Turnover and Turbulence," Federal Reserve Bank of Philadelphia Working Paper 11-44 (2011).

Fujita, Shigeru, and Vilas Rao. "Earnings Losses of Job Losers During the 2001 Economic Downturn," Federal Reserve Bank of Philadelphia *Business Review* (Fourth Quarter 2009).

Jacobson, Louise, Robert LaLonde, and Daniel Sullivan. "Earnings Losses of Displaced Workers," *American Economic Review*, 83:4 (1993), pp. 685–709.

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

Ljungqvist, Lars, and Thomas Sargent. "The European Unemployment Dilemma," *Journal of Political Economy*, 106:3, (1998). pp. 514-50.

Mortensen, Dale, and Christopher Pissarides. "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61 (1994), pp. 397–415.

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

Shimer, Robert. "Why Is the U.S. Unemployment Rate So Much Lower?," in Ben Bernanke and Julio Rotemberg, eds., *NBER Macroeconomics Annual* 1998, 13 (1998) pp. 11–61.

The Optimum Quantity of Money*

BY DANIEL SANCHES

A central premise of monetary policy in the U.S. throughout the first decade of the 21st century has been a firm commitment to avoid deflation. Indeed, it is the consensus view of policymakers and most economists. Nonetheless, Nobel laureate Milton Friedman proposed that optimal monetary policy should lead to a steady rate of deflation. For some economists, the Friedman rule is mainly a benchmark for thinking clearly about the assumptions underlying our models and a systematic guide for deciding how to modify our models, that is, a way of making scientific progress. However, it is not an exaggeration to say that most of the work in the field of monetary theory has focused on identifying situations in which Friedman's insight does not apply. In this article, Daniel Sanches discusses the Friedman rule and the main arguments that have been made against it.

A central premise of monetary policy in the U.S. throughout the first decade of the 21st century has been a firm commitment to avoid deflation, that is, a persistent fall in the price level. Indeed, it is the consensus view of policymakers and most economists.[1]

Daniel Sanches is an economist in the Research Department of the Philadelphia Fed. This article is available free of charge at www.philadelphiafed.org/research-and-data/publications/.

Nonetheless, in an influential 1969 article, Nobel laureate Milton Friedman proposed that optimal monetary policy should lead to a steady rate of deflation. Since the article was published, his notion of the optimum quantity of money has become one of the most widely celebrated and debated propositions in monetary economics. In large measure, this is because a broad class of monetary models has confirmed that deflation should be part of the best monetary policy.

For some economists, the Friedman rule is mainly a benchmark for thinking clearly about the assumptions underlying our models and a systematic guide for deciding how to modify our models, that is, a way of making scientific progress.[2] In fact, it is not an exaggeration to say that since Friedman proposed his rule for monetary policy decisions, most of the work in the field of monetary theory has focused on identifying situations in which Friedman's insight does not apply. This article discusses the Friedman rule and the main arguments that have been made against it.

WHAT IS MONEY AND WHY DO WE NEED IT?

To understand Friedman's ideas about the best monetary policy, we first need to understand why people need money. This lies at the heart of any theory of monetary policy. The most obvious answer is that people need money to conduct their daily transactions. In principle, money can be any object that serves as a means of payment as long as people believe that it will be widely accepted as a means of payment in future trades. For a long time, commodities such as gold and silver were used as a means of settling transactions; now dollar bills, checks, and debit cards serve this function. In

[1] See, for example, the 2002 and 2010 speeches by Ben Bernanke.

[2] Viewed this way, the optimum quantity of money in monetary theory stands with the Modigliani-Miller theorem in corporate finance and the Coase theorem in contract and bargaining theory.

* 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.

other words, you and I need tangible objects that help us pay for things at the grocery store, at a restaurant, online, etc. While currency and checking accounts are assets specially created for the purpose of serving as a means of payment, other assets, such as stocks and corporate bonds, are not typically used in this way.

Because we need a medium of exchange to pay for things, we can say that households and firms demand some convenient form of money to help them with their daily transactions. But who supplies money? Private entities such as commercial banks offer us checking accounts that allow us to write checks or use a debit card to conveniently pay for things. The Federal Reserve System (the U.S. central bank) also creates money. It supplies U.S. currency and reserve balances that help households, firms, and financial institutions make payments and settle debts. Thus, an important aspect of monetary policy is to control the amount of money in the economy, taking into account people's need for a means of payment. For instance, a central banker would certainly be concerned if there was too little money in the economy relative to the number of transactions. This would certainly cause problems for shoppers, workers, traders, and others.

Money is only one item among a large menu of assets held by households and businesses, so it is helpful to think of the demand for money as part of a broader portfolio problem. For instance, every month you have to decide how much to spend and save out of your income. After making this decision, you have to think about the kinds of assets you want to hold to achieve your monthly goals. You have to decide what fraction of your income you want to keep in your bank account and use your debit card for your daily purchases. You may also put some of your savings into higher-yielding assets. What is important for our discussion is the decision about the kinds of assets you think are useful for helping you pay for transactions.

The Transactions Role for Money. Economists usually define money as something that serves essentially three purposes: a unit of account, a medium of exchange, and a store of value. As a unit of account, money gives us a convenient way to measure the relative values of apples, oranges, and laptops. As a medium of exchange, money is an asset that facilitates transactions. Money allows two complete strangers to engage in trade even though neither party knows anything about the other. When the buyer hands his money to the seller, the transaction is immediately settled, and no further interaction is required. An obvious example is U.S. currency, i.e., the dollar bills you carry in your pocket, which are widely accepted as a means of payment in the U.S. and in some other countries. Other examples of money are checking accounts and some savings accounts that permit you to write a check or use your debit card to pay for your purchases.

Money typically pays a low rate of return, as anyone with a checking account or currency in his or her pocket knows. Why? Since assets that can be used as money provide a transaction service, money issuers (such as commercial banks and the Federal Reserve System) need to pay only a low rate of return in order to induce people to hold their money. And money holders (such as households and firms) are willing to give up some interest income in exchange for the convenience of having an asset that helps them pay for things. In other words, the transaction service that money provides comes at a cost: the low interest income the money holder receives.

The Precautionary Motive for Holding Money. In addition to holding money to conduct transactions, people also hold money as a store of value. In other words, they hold money as a way of transferring purchasing power from today to some future date. Why would people want to hold part of their savings in the form of money if other assets, such as government bonds and certificates of deposit (CDs), typically offer a higher rate of return? People may want to hold some of their savings in the form of money because some unanticipated events, such as an unexpected bill, may make them spend more in a given month than they had initially planned. For instance, if your after-tax monthly income is $3,000 and you decide at the beginning of the month to save $500 and spend $2,500, you may choose to keep, say, $2,800 in your checking account because it could be that you end up spending more on restaurant meals or taxi rides than you had initially planned. In principle, you could handle these unexpected expenses by cashing in bonds or CDs, but brokerage costs, explicit penalties, and uncertainty about the ability to sell securities for full value at short notice make these other assets less than perfect substitutes for unexpected needs. Thus, you keep more money in your checking account than what you actually plan to spend over the month.

Because we need a medium of exchange to pay for things, we can say that households and firms demand some convenient form of money to help them with their daily transactions.

Economists refer to this reason for holding money as the *precautionary motive.*

THE FRIEDMAN RULE

When the Nominal Interest Rate Is Positive, Households and Firms Hold Too Little Money. Now that we understand what money is and why people and firms need it, we can turn to our initial question: What is the best policy concerning money creation? In a 1969 article, Milton Friedman proposed a very simple rule for guiding monetary policy decisions. His goal was to overcome a basic inefficiency in monetary exchange: Households and firms tend to hold excessively small money balances when the nominal interest rate on short-term government bonds or CDs is positive. The nominal interest rate refers to the yield that an investor obtains in terms of dollars. For instance, if a bank lends $1,000 to an individual in exchange for a repayment of $1,050 one year later, then the *nominal* yield on the loan is 5 percent. In contrast, the *real* yield is 5 percent minus the expected rate of inflation. So if the expected rate of inflation is 3 percent, the real yield on the loan is 2 percent.

For the discussion that follows, it is important to distinguish between the nominal and real interest rate. The key point to keep in mind is that individuals and firms are primarily concerned about their purchasing power over goods and services. When an investor holds a bond, the real rate of interest tells you the increase in purchasing power over goods and services that accrue to the bondholder. The real interest rate is determined mainly by households' preferences and firms' production technologies, the main underlying factors of what economists refer to as the *real economy.* (See *The Nominal Interest Rate, the Real Interest Rate, and the Fisher Effect.*)

Why do households hold excessively small money balances? Even though money facilitates transactions, households and firms want to keep their money balances as small as possible. After all, money pays little or no interest. An economist would say that there is an *opportunity cost* of holding money: the interest that the household or firm could have earned by holding a nonmonetary but interest-bearing asset such as a 90-day CD or a Treasury bill. As a result, at any point in time, households and firms will choose to hold only a small fraction of their wealth in the form of monetary assets. In particular, they choose to hold some money to cover their planned expenditures or perhaps somewhat more because of the precautionary motive.[3]

Note that the cost to society of printing more paper money or allowing banks to create new checking accounts – the *social marginal cost* of producing money – is essentially zero. The social marginal cost refers to the additional resources required for the central bank to produce paper money or for a bank to create a deposit account when it makes a loan. Once the central bank has set up the printing press to create notes and once a commercial bank has hired its loan officers, set up its accounting system, bought computers, etc., the actual resource costs of creating additional units of paper money or deposits are negligible.

[3] To be more precise, households hold the right amount of money balances given the prevailing prices and interest rates. As will become clear, Friedman argues that households hold too little money because nominal interest rates are wrong — they are too high — from society's point of view.

Let's say that these costs are zero. So, households' and businesses' marginal cost of holding money (the forgone interest) is greater than the marginal social cost of supplying more money (which equals zero). This means that society would be better off if each household was holding a larger fraction of its wealth in the form of monetary assets, which would permit it to carry out a larger volume of useful transactions, that is, purchases of goods and services.

Even though money facilitates transactions, households and firms want to keep their money balances as small as possible. After all, money pays little or no interest.

Friedman came up with a straightforward way to overcome this inefficiency: Eliminate the opportunity cost of holding money by lowering the *nominal* interest rate until it was equal to the social marginal cost of producing money, that is, zero. In this case, since there is no opportunity cost of holding money, households and firms will not inefficiently economize on their money holdings.

It is important to note that Friedman was not proposing that monetary policy should drive a household's *real* return on its CDs and other nonmoney assets to zero, which would certainly not be a good thing. He was proposing to drive the nominal return, which equals the household's real return – the return that savers care about – plus the expected inflation rate, to zero. (If this distinction between nominal and real returns isn't obvious to you, take another look at *The Nominal Interest Rate, the Real Interest Rate, and the Fisher Effect.*)

Predictable Deflation Will Remedy the Problem. How can the central bank achieve Friedman's prescription?

The Nominal Interest Rate, the Real Interest Rate, and the Fisher Effect

To understand how to implement the Friedman rule, it is important to distinguish between the nominal interest rate and the real interest rate. The nominal interest rate tells you how fast the number of dollars in your account will increase over time if you acquire a certificate of deposit (CD) from a commercial bank or a three-month Treasury bill. For example, suppose you want to purchase a CD from your local bank. The bank will promise to repay the principal amount plus the interest agreed on at the time you acquire the CD. For instance, if the bank offers you a 5 percent annual nominal interest rate for a CD with a face value of $1,000, then at the end of one year, the bank will pay back the principal amount of $1,000 plus $50, which is the interest earned. Thus, the yield on your investment in terms of dollars is exactly 5 percent.

The real interest rate corrects the nominal interest rate for the effects of inflation, so that it tells you how fast the purchasing power of your savings will increase over time. Going back to our previous example, suppose that at the end of one year the inflation rate is 2 percent. This means the real yield on your investment is only 3 percent. In other words, the acquisition of the bank's CD increases the purchasing power of your savings by 3 percent at the end of one year. To compute the real interest rate, we can use the following formula:

Real Interest Rate = Nominal Interest Rate – Inflation Rate

Notice that we can rewrite this equation as follows:

Nominal Interest Rate = Real Interest Rate + Inflation Rate

Thus, we can split the nominal interest rate into two components: the real interest rate and the inflation rate. This allows us to examine the different economic forces that determine the nominal interest rate.

The real interest rate is determined by factors such as individuals' preferences and firms' production technologies. Think about individuals and firms deciding the rate at which they are willing to lend and borrow. Individuals' willingness to postpone current consumption and their projected future consumption needs will determine the interest rate at which they are willing to loan out funds to firms. And firms' expected profits, determined mainly by the marketability of their products and the productivity of their plants, will determine the rate they are willing to pay. If all prices double, that is, if individuals' incomes, the prices of goods and services, and the firms' profits all double, individuals' preferences and firms' productivity haven't fundamentally changed. Monetary policy certainly affects prices, but, at least to a first approximation, some economists often argue that monetary policy doesn't permanently affect the real rate of interest.

The second component that determines the nominal interest rate is the inflation rate. On many occasions, we do not know what the future inflation rate will be when we need to make our investment decisions today. Thus, if we want to understand the determinants of the nominal interest rate today, we should look at a measure of people's expectations about the rate of inflation over the investment period. Thus, we should rewrite the equation above as:

Nominal Interest Rate = Real Interest Rate + *Expected* Rate of Inflation

This means that the nominal interest rate depends on the real interest rate and the *expected* rate of inflation. This expression is usually referred to as the *Fisher relation* or *Fisher equation*, after economist Irving Fisher (1867-1947), who first studied it.

Using the Fisher relation, we can also define what is known as the *Fisher effect*. The Fisher effect says that there is a one-for-one adjustment of the nominal interest rate to the expected rate of inflation. It is important to note that the Fisher effect does not say that the nominal interest rate moves one-for-one with actual inflation. At the time you and I agree on a loan, we both have an expectation of what the inflation rate will be over the contract period so that we can compute the real interest rate, which gives the real cost of the loan for the borrower and the real gain for the lender. But if inflation catches the borrower and lender by surprise, the real cost and gain they initially thought they were going to get are not realized. Thus, the Fisher effect states that the nominal interest rate adjusts one-for-one to *expected* inflation (the inflation rate that both parties thought was going to be realized at the time they signed the contract).

According to Friedman, and many other economists, monetary policy affects the nominal rate of return, but not the real rate of return, at least in the long term. So if the central bank can ensure that the expected rate of inflation equals the negative of the real rate of return, the nominal interest rate will equal zero, according to the Fisher equation. If, for example, the real rate of return is 2 percent, then a 2 percent deflation would mean that the nominal rate of return is zero. Households and firms will be happy to hold assets paying a zero nominal rate of return. Intuitively, when prices of goods and services are falling at 2 percent per year, households' and firms' money balances (and their command over goods and services) are increasing in value at 2 percent per year.

This is exactly what Friedman proposed. He said that the central bank should generate a sustained deflation in the economy to drive the nominal interest rate on short-term securities such as Treasury bills and CDs to zero.

How can the central bank achieve this goal? If we think about an economy in which the average rate of growth of output is zero, then the way to achieve a sustained deflation is to reduce the money supply at a constant rate. Specifically, the central bank should contract the money supply at a rate equal to the economy's real rate of return. The prices of goods and services will fall as the money supply declines.[4]

The rule is slightly more complicated in a growing economy. If the central bank were to keep the supply of money constant in a growing economy, nominal prices would automatically fall, although not necessarily at the rate that would lead households to hold the right amount of money.[5] The specific rule for the rate of growth of the money supply the central bank should target to implement the Friedman rule also depends on the economy's average rate of growth, and the target growth in the money supply might even be positive. While this should be kept in mind, it is probably easiest to think about the Friedman rule for an economy that is not growing.

For Friedman, it was essential that the central bank make a commitment to act in a predictable way. The predictability of a central bank's policy rule is essential because the rule works through people's expectations about how prices will change. People must firmly believe that prices will fall in a predictable way, and this requires that they expect the money supply to shrink at a steady rate. As long as individuals expect prices to fall steadily at a constant rate and the central bank contracts the money supply at the promised rate – so that individuals' expectations are met – expected inflation will equal actual inflation and nominal interest rates will be driven to zero.

The Friedman Rule Without Deflation. While some critics have rejected the Friedman rule out of hand, other policies that do not involve a sustained deflation would also work. For example, David Andolfatto proposes an alternative way of implementing Friedman's prescription. His idea is to make money itself an interest-bearing asset. In this case, money holders would receive interest payments on their currency holdings and their checking accounts. If the nominal interest rate on money holdings equals the nominal yield on other (riskless) nonmonetary securities, then households bear no opportunity cost of holding money. Thus, we accomplish the same outcome without having to engineer a sustained deflation.[6]

CRITICISMS OF THE FRIEDMAN RULE

Even though the logic behind the Friedman rule is very simple and applies to a broad class of economic models, many monetary economists have argued that it is not the appropriate principle to guide monetary policy decisions. These criticisms come in five main varieties.

The Welfare Loss from Holding Too Little Money Is Small. The first criticism does not question Friedman's logic, but it does question whether Friedman has identified an important problem. Some critics of the Friedman rule have argued that in standard monetary models, deviations from the Friedman rule do not matter much for households' well-being even though the Friedman rule allows society to achieve the highest level of welfare within these simple models. To these critics, Friedman may have been correct logically, but the actual costs of holding too little money are small. Hence, some economists have argued that monetary policy should not place an excessive weight on the goal of eliminating the opportunity cost of holding monetary assets.

For instance, Thomas Cooley and Gary Hansen and, later, Robert Lucas have quantified the welfare consequences of having an inflation rate above that prescribed by the Friedman rule. These authors use models in which money is required to settle transactions. They conclude that

[4] Friedman argued that there was a fairly tight relationship between the rate of growth of the money supply (which the central bank could control) and the rate of inflation, although they might diverge for a time.

[5] By nominal prices I mean the price of goods and services in terms of dollars.

[6] While it is important to realize that Friedman's logic does not live or die with his proposal for deflation, the reader should note that policies like Andolfatto's involve a range of implementation issues. Any serious monetary policy prescription needs to take a wide range of practical complications into account.

people would be willing to give up only about 1 percent of their consumption to get rid of a 10 percent inflation, which is viewed as a small cost to society as a whole. According to their model simulations, the opportunity cost of holding money is not really large enough for policymakers to worry about.

Even though these studies have shown that the welfare cost of inflation is quantitatively small in the models they examine, it is hard to avoid concluding that their models must be missing something important. Think of what would happen in the U.S. if the average annual inflation rate were 10 percent. It is hard to believe that most people would not mention inflation as one of their main concerns. Having this in mind, subsequent researchers have shown that realistic additions to standard monetary models can lead to bigger effects. For example, the Cooley and Hansen model and the Lucas model do not include expenditures on machines and equipment. It is natural to ask whether their estimates of the welfare costs of inflation are low because their models have left out something important.

Benjamin Craig and Guillaume Rocheteau argue that firms make smaller capital expenditure decisions when the inflation rate is higher. A higher anticipated inflation rate reduces capital expenditures because it reduces firms' expected real revenue. Firms must decide today how much capital they should purchase to use to produce goods and services in the future. If a firm anticipates a high inflation rate by the time the machines and equipment are ready to produce, then it will probably decide to purchase fewer machines and less equipment today. This reduces the production of goods and services in the capital-intensive sectors and drives up their prices. If this effect is taken into account, the model predicts that households are willing to give up more than 5 percent of their consumption to get rid of a 10 percent inflation rate, more in line with the perception that inflation is one of people's main concerns.

The Friedman Rule Conflicts with Other Objectives. Some economists argue that monetary policy has more important things to do than reduce the opportunity cost of holding money. They argue that the main role of monetary policy is to respond to shocks that hit the economy, for example, a sudden rise in the price of oil or a decline in the demand for housing. Why? Many economists – notably, economists known as *New Keynesians* – believe that some prices in the economy are *sticky*; that is, they do not respond immediately to sudden changes in the economic environment. An online retailer or a restaurant may hesitate to change prices because of the costs of changing advertisements or menus or because they are worried about a negative reaction from consumers. As a result, only some producers change their prices immediately in response to unexpected changes in economic conditions. Other firms will wait until their actual price has moved too far out of line from the price that maximizes profits.

But in many economic models, the economy works best when prices respond flexibly to shocks. To see this, consider a simple example. Suppose that the price of oil suddenly rises 10 percent on a given day and remains at its higher level for some time. A rise in the price of oil certainly increases a manufacturer's costs because oil is usually an important input for the production process. But if the firm cannot increase the prices of its products in line with its higher costs, it will suffer a decline in profitability and may have to decrease production for some time. So production will be lost until the firm is able to change its price. Thus, sticky prices can result in inefficient outcomes. In this sense, we can think of an economy in which prices respond flexibly and immediately to changing conditions as a benchmark to guide policy.

The problem is that eliminating the opportunity cost of holding money, as prescribed by the Friedman rule, may be inconsistent with the goal of mitigating the inefficiency arising from price stickiness. For instance, it could be desirable to have a *positive* nominal interest rate to mitigate the effects of price stickiness. New Keynesian economists believe that this type of inefficiency is more important, so monetary policy should target it.

Note that even if monetary policy has other objectives, it is an open question whether policymakers should ignore Friedman's concern altogether. For example, Aubhik Khan, Robert King, and Alexander Wolman show that in a model with both a transactions role for money and costly price adjustments, the best monetary policy is, in fact, not far from the Friedman rule: In their model, they find that the average level of the nominal interest rate should be close to zero. This means that the transactions role for money is as important as the inefficiencies due to price rigidity empha-

This means that the transactions role for money is as important as the inefficiencies due to price rigidity emphasized in the New Keynesian literature.

sized in the New Keynesian literature. Boragan Aruoba and Frank Schorfheide have estimated a similar model using postwar U.S. data. They find that the inefficiency due to reduced money holdings and the inefficiency due to sticky prices are of similar magnitude. These two studies suggest that even in the presence of sticky prices, the transactions role for money is quantitatively important, so they argue that Friedman's concerns should be taken seriously.

The Recent Japanese Experience. Central bankers usually mention Japan's experience of the last 20 years as a reason to be concerned about deflationary policies. The Japanese economy appears to be stuck in what economists call a *liquidity trap*, a situation in which we observe a very low level of the nominal interest rate. In the last 10 years, the average level of the nominal interest rate has remained below 0.5 percent in Japan, and the inflation rate, as measured by the consumer price index, was positive in only three years. Despite many attempts to stimulate the economy, the average growth rate of output was 1.15 percent from 1997 to 2007, a very slow pace of economic growth.

This combination of deflation and low nominal interest rates creates problems for monetary policy when the economy is in a recession. In this case, any attempt to stimulate the economy by injecting more money through open market operations may have little or no effect on output. Thus, monetary policy should avoid a liquidity trap. In response to these concerns, many economists have devoted a lot of effort to analyzing the best policy responses that would release the economy from the liquidity trap.[7] For this reason, central bankers have been reluctant to consider deflationary policies such as the Friedman rule, especially when they look at the Japanese experience as an example of an economy that appears to be stuck in a liquidity trap.

When Money Is Held for Precautionary Purposes, Some Inflation May Be Good. Another criticism of the Friedman rule is that taking into account the precautionary motive for holding money may lead to prescriptions different from those of the Friedman rule. Some economists argue that precautionary motives are very important for households that have limited ability to insure themselves against sudden declines in income or unexpected expenses. In addition to holding money for transactions purposes, these households also hold money because the boiler or the car may break down unexpectedly. More seriously, many households in the U.S. do not have health insurance or other forms of insurance to protect themselves against unexpected health-care expenses. Thus, holding money balances is a form of self-insurance.

Central bankers usually mention Japan's experience of the last 20 years as a reason to be concerned about deflationary policies.

But insuring yourself by holding money balances costs you something: the interest income you could have obtained by holding a less liquid but interest-bearing asset. People hold more money than they need for transaction purposes because of the precautionary motive. Edward Green and Ruilin Zhou have shown that a mild inflation guarantees that people do not hold too much money for insurance purposes. In other words, a mild inflationary policy balances the costs and benefits of holding money for insurance purposes. This result goes against the Friedman rule because it usually implies a positive level for the nominal interest rate, while the Friedman rule, remember, proposes a zero nominal interest rate. But the extent to which a mild inflation is socially beneficial crucially depends on the extent to which private and public insurance markets do not provide enough protection against unexpected events.

Technological Change Has Made "Money" Obsolete. Even though money is a convenient way to pay for things, there are substitutes for money. Credit cards are a good example. When a buyer enters a store and uses his credit card to pay for his purchases, he does not need any money. The buyer's credit card company keeps track of his balance and authorizes any transaction that does not exceed his credit limit.

The merchant also has an agreement with the credit card company to accept the cards the company issues. Even though credit arrangements of this kind work well, notice that some form of money is still necessary to settle debts among the parties involved in the credit network. For instance, the credit card company pays the merchant on the settlement date usually by transferring money from its checking account to the merchant's account. Also, the buyer needs to pay the credit card company on the due date, usually by making an electronic transfer from his checking account to the credit card company's account.

In this respect, Friedman's argument remains valid (even in an economy in which credit prevails as a means of payment for retail trades) if we broadly interpret transactions to include all kinds of transactions.

[7] For a discussion of the role of monetary and fiscal policy in avoiding a liquidity trap, see Michael Dotsey's *Business Review* article.

CONCLUSION

Many economists have criticized Friedman's notion of the optimum quantity of money, despite its being a fairly robust conclusion across a wide range of models. Although Friedman proposed a monetary policy that leads to steady deflation, subsequent researchers have shown alternative ways to get the same result. In addition, models that take explicit account of how households and firms use money for both transactions and insurance and models in which firms are slow to adjust prices show that Friedman's insights need to be supplemented. While few economists or policymakers would prescribe the Friedman rule as a literal guide to policy, this does not mean that Friedman's insight is irrelevant. The rule has been useful in spurring serious thoughts about the role of money in the economy and has helped economists make scientific progress in the search for more accurate models of the economy.

REFERENCES

Andolfatto, David. "Essential Interest-Bearing Money," *Journal of Economic Theory*, 145 (2010), pp. 1495-1507.

Aruoba, Boragan, and Frank Schorfheide. "Sticky Prices Versus Monetary Frictions: An Estimation of Policy Trade-offs," *American Economic Journal: Macroeconomics*, 3 (2011), pp. 60-90.

Bernanke, Ben. "Deflation: Making Sure 'It' Doesn't Happen Here," remarks made before the National Economists Club, Washington, D.C., November 21, 2002, available at: http://www.econ.fudan.edu.cn/userfiles/file/20110503035359687.pdf.

Bernanke, Ben. "The Economic Outlook and Monetary Policy," speech given at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, WY, August 27, 2010, available at: http://federalreserve.gov/newsevents/speech/bernanke20100827a.htm.

Cooley, Thomas, and Gary Hansen. "The Inflation Tax in a Real Business Cycle Model," *American Economic Review*, 79:4 (1989), pp. 733-48.

Craig, Benjamin, and Guillaume Rocheteau. "Inflation and Welfare: A Search Approach," *Journal of Money, Credit and Banking*, 40:1 (2008), pp. 89-119.

Dotsey, Michael. "Monetary Policy in a Liquidity Trap," Federal Reserve Bank of Philadelphia *Business Review* (Second Quarter 2010), pp. 1-15.

Friedman, Milton. "The Optimum Quantity of Money," in *The Optimum Quantity of Money and Other Essays*. Chicago: Aldine Publishing Company, 1969.

Green, Edward, and Ruilin Zhou. "Money as a Mechanism in a Bewley Economy," *International Economic Review*, 46:2 (2005), pp. 351-71.

Khan, Aubhik, Robert King, and Alexander Wolman. "Optimal Monetary Policy," *Review of Economic Studies*, 70 (2003), pp. 825-60.

Lucas, Robert. "Inflation and Welfare," *Econometrica*, 68:2 (2000), pp. 247-74.

Measuring Economic Uncertainty Using the Survey of Professional Forecasters*

BY KEITH SILL

Uncertainty about how the economy will evolve is a key concern for households and firms. People's views on how likely it is that the economy will be growing, stagnating, or in recession help shape the actions they take today. Consequently, how households and firms respond to uncertainty has implications for economic activity. In addition, uncertainty matters to policymakers: Monetary policymakers recognize that if uncertainty about future inflation is high, decision-making by households and firms becomes more complicated. In this article, Keith Sill describes how uncertainty can be measured using data from the Survey of Professional Forecasters and shows how these measures have changed over time for output growth and inflation. He also examines some links between the macroeconomy and measures of output and inflation uncertainty.

Uncertainty about how the economy will evolve is a key concern for households and firms. People's views on how likely it is that the economy will be growing, stagnating, or in recession help shape the actions they take today. For consumers, how much to spend, what to purchase, and how much to save depend in part on how uncertain they are about their future incomes. For firms, how many workers to hire or how much new capacity to invest in depends on expected future demand and how certain they are that forecasted demand will be realized. Consequently, how households and firms respond to uncertainty has implications for economic activity. In addition, uncertainty matters to policymakers: Monetary policymakers recognize that if uncertainty about future inflation is high, decision-making by households and firms becomes more complicated.

Keith Sill is a vice president in the Philadelphia Fed's Research Department and the director of the Real-Time Data Research Center. 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.

The importance of gauging economic uncertainty points to the need for data on economic uncertainty. Forecast surveys are one such source of data, since they can often be used to construct measures of uncertainty about the future paths of key economic variables such as output growth, unemployment, and the inflation rate. The Philadelphia Fed's Survey of Professional Forecasters (SPF) is an important source of data on economic uncertainty, since it has a long history of directly asking its respondents to assess the uncertainty that surrounds their forecasts of key macroeconomic variables. The survey data enable us to evaluate how uncertainty about the future economy has changed over time and whether uncertainty is rising or falling as we look ahead.

In this article we will describe how uncertainty can be measured using the SPF data and show how these measures have changed over time for output growth and inflation. We will also examine some links between the macroeconomy and measures of output and inflation uncertainty.

UNCERTAINTY MATTERS

Uncertainty about the future can have consequences for the decisions we make today. It is not only what we expect will happen in the future that can matter but also how sure we are about the alternatives we face. A simple example can illustrate how uncer-

tainty about an outcome can influence choices. Take the hypothetical case of an employee who gets an annual salary bonus. In the first scenario, the employee is told he will receive a $1,000 bonus for certain at the end of the year. In the second scenario, the employee is told that there is a 40 percent chance that the bonus will be zero, a 40 percent chance that it will be $2,000, and a 20 percent chance that it will be $1,000. The average payoff in both scenarios is $1,000, but most people are probably not indifferent to the two alternatives: Most people prefer getting the $1,000 for certain rather than taking the gamble of the second scenario. For the most part, people try to avoid risk (all else equal) and would prefer low uncertainty surrounding their expected outcome compared with high uncertainty around the same expected outcome. The interaction of disliking risk and the amount of uncertainty about outcomes influences the choices people make.[1]

While the example above is a bit contrived, there is good reason to believe that households' decisions about how much to save and how much to spend are affected by their views about economic uncertainty. The consumption/saving decision depends on a host of factors, including current interest rates, time to retirement, and anticipated future income and expenses. The decision about how much to save would be easier if there were no uncertainty. If the household were sure of its future income, of its future expenses, of how long it would live, and of future asset prices and returns, it would face a fairly straightforward calculation to figure out how much to save and spend so that its wealth is spent down in the best possible way. However, if the future is uncertain, the nature of the calculation becomes more subtle. For example, if someone becomes very worried about his future employment prospects, even though he anticipates the most likely outcome is that he will keep his job, he may consume less today and try to build up a savings buffer to help maintain consumption during potential bad times.[2] If there were less uncertainty about the future, households would save less and average consumption would be higher.

It's not just households that are influenced by uncertainty; firms' views on uncertainty may affect their current decisions as well.

Indeed, this is a real concern for workers during the recovery. A recent *New York Times* report on a USA Today/Gallup poll showed that in 2011 the fraction of workers who reported being worried about being laid off was about 30 percent. This was substantially higher than the 20 percent or so who reported being worried over the period from 1998 to 2005. Given this uncertainty about their jobs, we might expect that households are being conservative about spending and are trying to build a savings buffer.[3]

It's not just households that are influenced by uncertainty; firms' views on uncertainty may affect their current decisions as well. A firm that expects demand for its products to increase in the future will need to consider expanding production capacity today. Suppose the investment in a new plant is irreversible in the sense that once the capacity is built, it cannot be used for anything other than its intended use. However, a decision to delay the investment until the future is reversible: The firm could go ahead and start the investment project next month if it decides not to start it today. When there is uncertainty about the expected future benefits and costs of the investment project, often the best choice for a firm is to undertake the investment only when the expected benefits exceed the expected costs by a large enough amount. If there were no uncertainty about expected future benefits and expected future costs of the investment, the firm should instead undertake the investment whenever the expected benefits just exceed the expected costs. This phenomenon is sometimes referred to as the option value of waiting. By waiting, the firm might find that its future path is clearer and the investment should then be undertaken.[4] This theory suggests that greater uncertainty about future conditions will lead to fewer investment projects being undertaken today.

Monetary policymakers consider economic uncertainty when designing policy as well. In a 2008 speech, then-Federal Reserve Governor Frederic Mishkin discussed inflation and inflation dynamics.[5] Mishkin noted that policymakers are concerned not just with forecasts of inflation but also with inflation uncertainty. In particular, "Policymakers need to be concerned about any widening of inflation uncer-

[1] See Pablo Guerron's *Business Review* article for a discussion of how uncertainty can affect the macroeconomy.

[2] See the papers by Christopher Carroll and Angus Deaton on the buffer stock model of consumption.

[3] See Shigeru Fujita's article on pages 1-7 for a discussion of how uncertainty can affect the labor market.

[4] See the paper by Robert McDonald and Daniel Siegel.

[5] See the speech by Mishkin.

tainty. Indeed, an increase in inflation uncertainty would likely complicate decision making by consumers and businesses concerning plans for spending, savings, and investment." Thus, monetary policymakers often strive to set policy in a way that leads to low and stable inflation (and maximum sustainable employment in the case of the U.S.). A history of stable inflation means that uncertainty about future inflation is likely to be lower, since people will perceive the central bank as being credible when it promises to deliver a good inflation outcome.

Since uncertainty seems to be an important component of decision making, are there data we can use to get a handle on uncertainty? Forecast surveys provide such data. In particular, the Philadelphia Fed's SPF was designed in part to give insight into the evolution of uncertainty.

THE SURVEY OF PROFESSIONAL FORECASTERS

The SPF asks professional forecasters to give their forecast for 32 key macroeconomic variables, including gross domestic product (GDP), short-term and long-term inflation, and unemployment. The survey was initiated as a joint product of the National Bureau of Economic Research (NBER) and the American Statistical Association (ASA) in 1968 and was originally known as the NBER-ASA Economic Outlook Survey. The Philadelphia Fed took over the survey in 1990. The SPF is conducted quarterly, and typically, the survey gets responses from 50 or so professional forecasters.[6] In the surveys conducted since the Philadelphia Fed took over, the forecasters provide quarterly forecasts for five quarters and annual forecasts for the current year and the following year. (See *Data on Forecast Uncertainty at the Federal Reserve Bank of Philadelphia* for links to various data from the Real-Time Data Research Center.)

To illustrate how the SPF can be used to gauge uncertainty, we will work with a survey that was published in 2011. The table shows the median forecast for real GDP growth, the unemployment rate, and payroll employment from the third quarter 2011 SPF released on August 12, 2011. The columns labeled "New" represent the latest forecast, and the columns labeled "Previous" represent the forecast provided in the second quarter of 2011. Looking across the columns, we see that forecasters were a bit more pessi-

[6] See the article by Dean Croushore for a description of the SPF. More information about the SPF, including the history of the survey, can be found on the Philadelphia Fed's website at: http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/.

TABLE

Survey of Professional Forecasters - Q3 2011

	Real GDP (percent)		Unemployment Rate (percent)		Payrolls (000s/month)	
	Previous	New	Previous	New	Previous	New
Quarterly data:						
2011:Q3	3.4	2.2	8.7	9.1	194.5	105.3
2011:Q4	3.5	2.6	8.5	9.0	173.9	148.7
2012:Q1	2.9	2.2	8.4	8.8	219.4	180.3
2012:Q2	2.5	2.9	8.2	8.7	182.0	138.0
2012:Q3	N.A.	3.2	N.A.	8.6	N.A.	187.0
Annual data (projections are based on annual average levels):						
2011	2.7	1.7	8.7	9.0	130.4	111.5
2012	3.0	2.6	8.1	8.6	194.8	150.1
2013	2.8	2.9	7.5	8.1	N.A.	N.A.
2014	3.3	3.1	7.0	7.6	N.A.	N.A.

Data on Forecast Uncertainty at the Federal Reserve Bank of Philadelphia

The Philadelphia Fed Research Department's Real-Time Data Research Center (RTDRC) makes available on its website data on the Survey of Professional Forecasters (SPF) and Livingston Survey, as well as measures of forecast dispersion for SPF variables.

The home page for the Real-Time Data Research Center is: http://www.philadelphiafed.org/research-and-data/real-time-center/.

The historical data from the SPF are available at: http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/.

Data sets on SPF variable forecast dispersion are available at: http://www.philadelphiafed.org/research-and-data/real-time-center/spf-forecast-dispersion.cfm.

The RTDRC also maintains the Livingston Survey http://www.philadelphiafed.org/research-and-data/real-time-center/livingston-survey/ and provides historical data on the forecasts of Federal Reserve Board of Governors' staff:

http://www.philadelphiafed.org/research-and-data/real-time-center/greenbook-data/.

mistic about their outlook for the U.S. economy compared with the second quarter 2011 survey. The median forecast called for real GDP growth of 1.7 percent in 2011, rising to 3.1 percent in 2014. The unemployment rate was expected to decline slowly to an average of 7.6 percent in 2014. The SPF asks respondents for a payroll employment forecast only for the current year and the next year. Those forecasts indicated a mean forecast of 111,500 jobs per month in 2011 and 150,100 jobs per month in 2012.

The numbers in the table are called point forecasts, since they show a single number for the forecasted variable rather than a range of likely outcomes. That is, each survey respondent gives a specific number representing his or her forecast (expected outcome) for output growth, unemployment, and inflation. The numbers in the table, then, represent the median response of the individual forecasts, but they give us no sense of how uncertain the forecasters are about their individual forecasts. Are they very certain about their forecasts, perhaps more so than usual? Or are they very uncertain about their forecasts? We cannot tell from the information in the table.

Fortunately, the SPF asks each forecaster directly about his or her forecast uncertainty. That is, the SPF respondents are asked to attach a probability to each of a number of pre-assigned intervals over which their forecast may fall. The Philadelphia Fed then takes the mean probabilities over the individual respondents and reports them in the SPF release in the form of a histogram. A histogram is a graphical representation of an estimate of a probability distribution for a variable. That is, a histogram shows the probability that a variable will lie in a certain range. For example, Figure 1 shows the mean probabilities for real GDP growth and core PCE inflation in 2012 as reported in the third quarter 2011 SPF.[7] The figure shows that respondents became somewhat more sure that real GDP growth in 2012 would fall in a range of 2 to 2.9 percent in the third quarter 2011 survey (black bars) compared with what they thought at the time of the previous survey in the second quarter of 2011

[7] Core PCE inflation removes the effects of changes in food and energy prices from the headline PCE measure.

FIGURE 1

Mean Probabilities in 2012

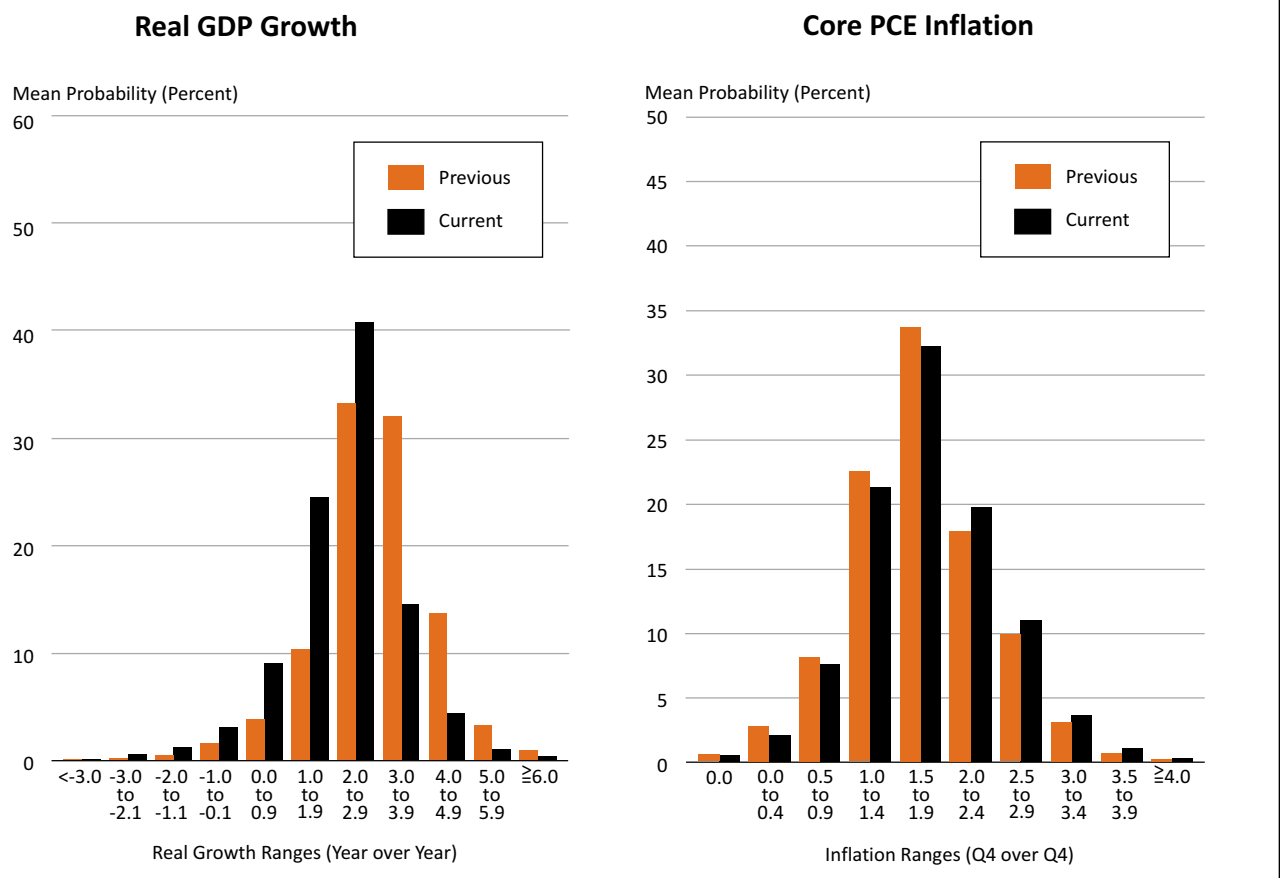

Source: Survey of Professional Forecasters, Third Quarter 2011

(orange bars). The forecasters attach some probability to real GDP growth being less than -1.1 percent, or greater than 5.9 percent, but the probabilities are small. It is clear from the figure that the forecasters see a bit above a 60 percent chance that real GDP growth for 2012 will fall in a range of 1 to 2.9 percent. In addition, the figure shows that forecasters see a greater chance of lower GDP growth compared with the previous forecast. We can see this from the fact that the height of the black bars toward the right side of the chart has shifted down and the height of the black bars toward the left side of the chart has shifted up. This means the forecasters are placing more probability on lower growth outcomes.

For core inflation, the figure suggests that forecasters shifted their views slightly toward the chance of higher inflation in the latest forecast. In particular, the height of the black bars to the right of the 1.5 to 1.9 bin has shifted up relative to the orange bars, and the height of the black bars toward the left end of the chart has shifted down.

What does Figure 1 tell us about forecast uncertainty? Note, first, that if all the SPF respondents were certain that real GDP growth would be in a range of 2 to 2.9 percent, there would be a single black bar at the 2.0 to 2.9 entry on the x axis, and the height of the bar would extend up to 100 percent. Alternatively, if the respondents thought that it was equally likely that real GDP growth would fall in any of the intervals labeled on the x axis, there would be a black bar of the same height (about 9 percent) at each entry on the x axis. In the former case, the respondents have very low (nil) uncertainty about real GDP growth in 2012. In the latter case, the respondents are very uncertain about real GDP growth in 2012. This indicates that a distribution of bars that is very tightly centered indicates low uncertainty compared with a distribution of bars that is very spread out.

One way to quantify the amount of uncertainty represented in Figure 1 is by using a measure of dispersion such as variance. To compute a variance, one calculates the average sum of squared differences of the observations from the mean. The units of measurement attached to variance are a bit awkward to work with, so researchers usually compute the standard deviation, which is the square root of variance. The standard deviation then has the same units of measurement as the data in question. All else equal, when dispersion around the mean is high, the standard deviation is high, and when dispersion around the mean is low, the standard deviation is low. For example, if all the observations of the variable in question were exactly equal to the mean, the standard deviation would be zero.

We can easily compute the standard deviation implied by the survey respondents' views on uncertainty that are embodied in Figure 1 using standard formulas. This gives us a single number for each histogram in the SPF that we can then use to make comparisons across time for uncertainty surrounding the forecasts. The time series of standard deviations from the uncertainty histograms for real GDP growth is shown in Figure 2. We plot the standard deviation for the year-ahead projections of real output growth as of the first quarter SPF for each year since 1981. Prior to 1981 the SPF asked respondents about nominal GNP uncertainty rather than real GDP, so we drop those observations. From 1981 to 1991 the survey asked forecasters to fill in six probability bins (or intervals on the x axis in Figure 1) for real GDP growth. Since 1992, the survey asks forecasters to fill in 10 probability bins. Because of this change in the survey question, we plot the pre-1992 data in black and the post-1992 data in orange. We construct a similar graph for inflation forecasts, where inflation is measured using the GDP deflator. We use this series because of its long history in the SPF (PCE inflation questions were only added to the SPF beginning in 2007). As in the case of GDP, the nature of the questions the forecasters are asked has changed over time. From the third quarter of 1981 to the first quarter of 1985, forecasters were asked to fill in probabilities for six bins (<4, 4 to 5.9, 6 to 7.9, 8 to 9.9, 10 to 11.9, and 12+). We plot the standard deviation from these histograms in black. From the second quarter of 1985 to the fourth quarter of 1991, the size of the bins changed (<2, 2 to 3.9, 4 to 5.9, 6 to 7.9, 8 to 9.9, 10+), and we plot standard deviations for these data in the dotted line. Since the first quarter of 1992, the forecasters have been asked for probabilities over the 10 bins shown in Figure 1, and we plot standard deviations for these data in orange in Figure 2.

The figure shows that there are large shifts in the uncertainty measures when the survey changed the

Especially in the case of inflation, there appears to be a link between the level of inflation and uncertainty as measured by the standard deviation. In particular, when the average forecast for inflation is high, forecast uncertainty tends to be high as well.

FIGURE 2

Output Growth and Inflation Standard Deviations Calculated from SPF Histograms

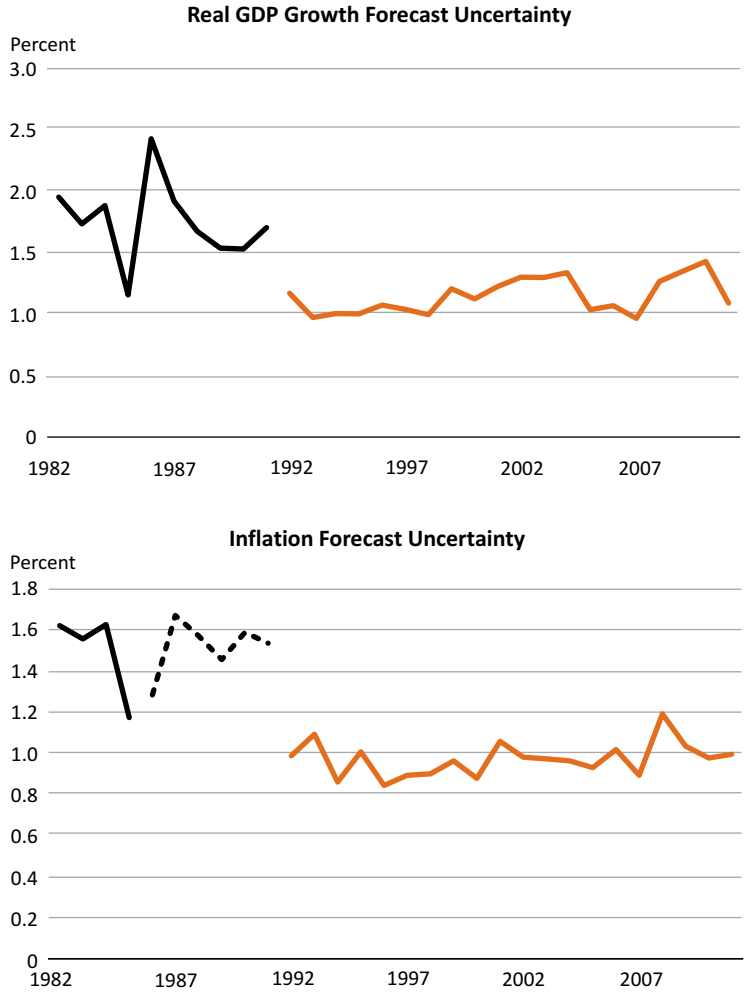

Top panel: black line shows pre-1992 data; orange line shows post-1992 data

Bottom panel: black line shows standard deviations Q3 1981 to Q1 1985; dotted line shows standard deviations Q2 1985 to Q4 1991; orange line shows Q1 1992 to 2011.

Source: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters and author's calculations

number and/or size of the bins that it asked the forecasters to consider. This makes it difficult to compare SPF uncertainty over long spans of time. It is likely, for example, that inflation uncertainty was high in the 1980s, but how high compared to the 1990s and 2000s is difficult to say. Fortunately, researchers such as Robert Rich and Joseph Tracy and Paolo Giordani and Paul Soderlind have used statistical methods to refine the SPF measures of uncertainty and make them more comparable over time.[8] For the most part, their measures do indicate that inflation uncertainty was generally higher in the 1980s than it was in the 1990s. However, it remains a difficult task to assess the magnitude of changes in uncertainty when the survey changes over time.

If we focus on the uncertainty measures in the 1990s and 2000s that are consistently measured, we see that there are fairly sharp movements over the last two decades. Output growth uncertainty rose from the mid-1990s until about 2004 and then moved down sharply. Since the most recent recession, output uncertainty appears to have generally risen. For inflation, it appears that uncertainty has generally been rising since about 1996.

Especially in the case of inflation, there appears to be a link between the level of inflation and uncertainty as measured by the standard deviation. In particular, when the average forecast for inflation is high, forecast uncertainty tends to be high as well. We can see this by looking at a scatter plot of the mean one-year-ahead forecast for inflation and the standard deviation of the one-year-ahead inflation forecasts, both computed from the SPF histograms (Figure 3).[9] From the figure we see that there is a strong tendency for the standard deviation of forecasts for inflation to be high when the mean forecast for inflation is high (that is, the points tend to line up from southwest to northeast). Why might this be? It could be that when expected inflation is high, forecasters are especially unsure about the future course of monetary policy and so are more uncertain about what inflation will be in the future. Since forecasters use different models and beliefs to make their projections, their uncertainty about

[8] Giordani and Soderlind fit normal distribution approximations to the histogram data in the SPF. Rich and Tracy redefine the SPF bins to impose a common 2-percentage-point width throughout the sample period.

[9] A scatterplot is a diagram that displays values for two variables in a data set. The data are shown as a collection of points, each having the value of one of the variables shown on the horizontal axis and the value of the other variable shown on the vertical axis.

FIGURE 3

GDP Deflator Inflation Year-Ahead Mean Forecast vs. Forecast Uncertainty

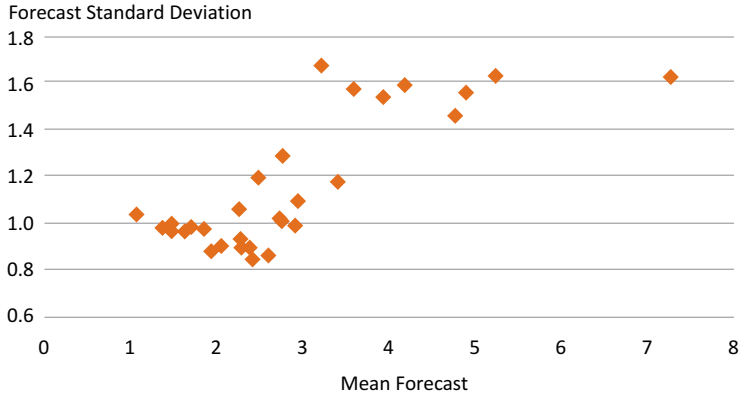

Each point represents the degree of forecast uncertainty for a given mean forecast.

Source: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters and author's calculations

future monetary policy is reflected in a wide range of inflation forecasts. This story is consistent with the episode in the early 1980s when inflation had been running at a high level and inflation expectations were unanchored. Paul Volcker, then-Chairman of the Federal Open Market Committee, engineered the disinflation that began to re-establish the credibility of monetary policymakers as guardians of price-level stability. During this time, forecasters may well have been very uncertain about how credible monetary policy would be and may have reflected this uncertainty in their inflation forecasts.

FORECAST DISAGREEMENT

An alternative measure that has often been used as a proxy for direct measures of uncertainty is called forecast disagreement.[10] Forecast disagreement measures how close the individual forecasters' projections in surveys like the SPF are to each other. The idea is that if all the forecasters are forecasting the same number, there is a sense in which forecast uncertainty may be lower. Similarly, if the forecasters are very far apart in their projections, there is a sense in which forecast uncertainty may be higher. The Philadelphia Fed Research Department's Real-Time Data Research Center (RTDRC) makes available on its website this proxy for uncertainty for selected variables in its SPF database.[11] The RTDRC provides forecast disagreement in the form of the 75th percentile of the point forecasts minus the 25th percentile. That is, we sort the point forecasts from high to low, chop off the top fourth and the bottom fourth, and take the difference of the remaining highest and lowest values. Since this measure removes the top and bottom of the distribution from the computation, it is less sensitive to extreme outliers.

The benefits of using a measure such as forecast disagreement are that such a measure is very easy to compute, it can be computed in a consistent way for the entire history of the survey, and it can be computed for every variable for which respondents provide forecasts. Figure 4 is a plot of inflation forecast disagreement constructed from the data provided on the RTDRC website.

It shows how disagreement about GDP deflator inflation forecasts has evolved over the past 20 years or so. We could examine an even longer history for this series, but we chose to limit it to 1983 for comparability with the measures of uncertainty we presented earlier. We see that there was more disagreement about inflation forecasts in the 1980s and that disagreement gradually declined until the late 1990s. Then, beginning in about 2007, there has been an upward movement in inflation forecast disagreement. Broadly speaking, this is in line with the uncertainty measure we calculated for GDP deflator inflation in Figure 2.

Disagreement measures account for how different the point forecasts of the individual forecasters are. But this is not the same thing as uncertainty about forecasts, and this measure of dispersion as a proxy for uncertainty is not without its problems. In particular, suppose only one forecaster responded to the SPF. In that case, there is no other forecaster with whom to compare her, and so we would conclude, using our forecast disagreement measure, that there was no disagreement; and if disagreement was our proxy for uncertainty, we would have to say that there was no uncertainty. But that lone forecaster who responded to the survey may have been very unsure of her forecast. In fact, she may have had high uncertainty about the future and about the forecast for variables such as output and inflation. We would clearly not be able to uncover information

[10] See, for example, the paper by William Bomberger, which investigates disagreement as a measure of uncertainty. See also the references in Giordani and Soderlind.

[11] See http://www.philadelphiafed.org/research-and-data/real-time-center/spf-forecast-dispersion.cfm.

FIGURE 4

GDP Deflator Inflation Forecast Disagreement

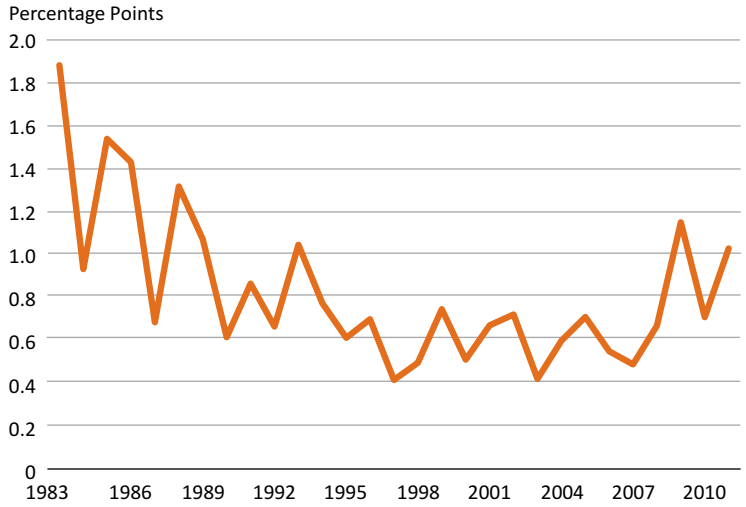

Source: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters and author's calculations

about forecast uncertainty by looking at the disagreement measure. Similarly, it could be that forecast disagreement is not necessarily a good proxy for uncertainty even when we have many forecasters responding to the survey. However, we can compare forecast disagreement with the direct measures of uncertainty in the SPF to get an idea of whether disagreement might be an acceptable proxy for uncertainty.

EVALUATING MEASURES OF UNCERTAINTY

Is uncertainty measured from the SPF histograms the benchmark for measuring economic uncertainty? The SPF allows us to calculate a third measure of uncertainty that has the firmest grounding in terms of economics: We can calculate the standard deviation from each forecaster's histogram and then take the average across forecasters. We call this measure the average dispersion across forecasters.

Note that this measure differs from uncertainty calculated using Figure 1. In that case, we averaged the individual forecasters' views on uncertainty and then calculated a standard deviation, which we plotted in Figure 2. But what if, instead, we calculate the standard deviation for each individual forecaster and then take the average across forecasters? Why might these two measures differ? Because when we first take the average over the individual forecasters reported in the histograms and then compute dispersion, we are, in effect, incorporating information about how their point forecasts differ. That is, we don't account for individuals' mean forecasts when we compute the standard deviation; instead, we account for the mean across all forecasters when we compute the standard deviation. On the other hand, if we first compute the standard deviation for each forecaster, we are, in effect, taking out the mean, or point forecast, for each individual. The average of the individual standard deviations then does not contain information about differences in point forecasts across survey respondents.

This average dispersion measure across forecasters is probably what most people have in mind when they think about economic uncertainty. In effect, it calculates the average level of uncertainty across people. As a practical matter, though, this measure is somewhat difficult to work with. First, the same problem that we had with the survey questions changing over time is present with this measure, as it is with the aggregate measures shown in Figure 1; so a long time series is not readily available. Second, one now has to calculate a dispersion measure from many more histograms that might not have statistical properties as nice as those in the aggregate histograms reported in the SPF.

In part for these reasons, researchers have made use of the link between the uncertainty computed from the average histograms reported in the SPF (and shown in Figure 1) and forecast disagreement to back out average dispersion across forecasters, rather than compute it directly. It can be shown that the variance of the SPF average distribution equals the average variance over the individual forecasters plus forecast disagreement. So, if we want to calculate an uncertainty measure that does not incorporate forecast disagreement, we can simply subtract forecast disagreement from the variance of the aggregate distribution and take the square root to get the units right. This average dispersion across

Is uncertainty measured from the SPF histograms the benchmark for measuring economic uncertainty?

forecasters is probably what we mostly have in mind when we ask whether people are more or less uncertain about economic conditions. Note that if all of the forecasters agreed on their point forecasts, the standard deviation from the aggregate histograms in the SPF would coincide with the average uncertainty across respondents.

Several recent economic studies have examined whether forecast disagreement is a good proxy for average uncertainty, and the studies come to somewhat different conclusions. Giordani and Soderlind find that forecast disagreement is a pretty good proxy for average uncertainty in the case of inflation. Rich and Tracy use different statistical techniques and are more skeptical about how well disagreement proxies for average uncertainty for inflation; Gianna Boero, Jeremy Smith, and Kenneth Wallis are skeptical as well. While average uncertainty is a theoretically more appealing construct, forecast disagreement is easy to compute for any survey of forecasters and so provides a longer history covering more variables than average uncertainty. The European Central Bank is now collecting data on forecast uncertainty in its forecasting survey. In addition, the Bank of England's Survey of External Forecasters has been asking respondents to provide measures of uncertainty similar to those in the SPF. Over time, as the Bank of England's survey and the SPF build up larger data sets on forecaster uncertainty, researchers will have the opportunity to further investigate the extent to which forecast disagreement provides a good proxy for uncertainty.

UNCERTAINTY, DISAGREEMENT, AND AGGREGATE BEHAVIOR

For practical purposes, we have two readily available measures that can potentially serve as proxies for uncertainty: uncertainty measured from the average histograms reported in the SPF (as shown, for example, in Figure 2) and forecast disagreement (as shown, for example, in Figure 4). Our earlier discussion on how uncertainty affects decision-making by households and firms suggested that when uncertainty is high, consumption growth and investment growth might be low. While we do not have a very long time series from the SPF, we can nonetheless examine whether there is a tendency in the data for consumption and investment to be low when uncertainty is high. We can look for this relationship in the data using simple correlations.[12]

However, any such relationships we uncover should not be taken as

[12] The paper by Bachmann, Elstner, and Sims uses survey data to explore the link between uncertainty and economic activity. They find that higher business uncertainty (measured using disagreement in business expectations from the Philadelphia Fed's Business Outlook Survey) leads to declines in economic activity.

FIGURE 5

Forecast Disagreement Versus Consumption and Investment Growth

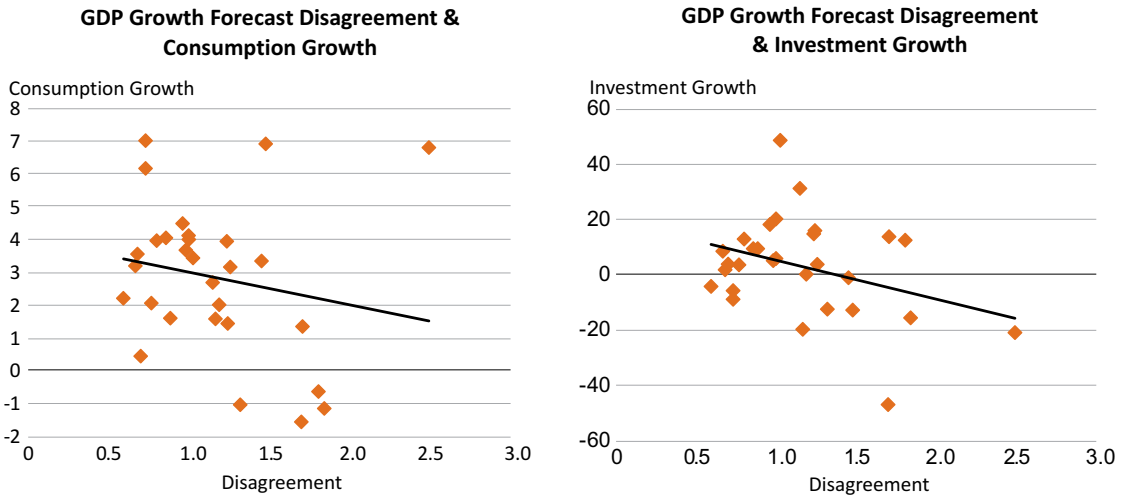

Left panel: Each point measures disagreement computed from the first quarter survey of each year; vertical axis measures consumption growth in quarter in which that survey was taken.

Right panel: Each point measures disagreement for real GDP growth plotted against actual investment growth.

Source: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters and author's calculations

proving or disproving an economic theory that posits a negative relationship between uncertainty and/or disagreement and consumption/income growth: We are instead exploring features of the data that would need to be accounted for by economic theory. Indeed, the causality between growth and uncertainty could go either way: Low consumption growth may indicate to forecasters that the economy is likely to enter a recession and so uncertainty about the future is high; or it may be that uncertainty is high, so consumers save more and consume less in anticipation of tough times ahead. We cannot distinguish between these alternative stories by looking at plots of uncertainty vs. consumption growth.

Figure 5 shows how forecaster disagreement is related to consumption growth and investment growth. The disagreement measure is taken from the RTDRC website and is the difference between the 75th percentile and 25th percentile for forecasts of one-quarter-ahead real GDP growth. We then compare that measure of disagreement to consumption growth and investment growth in the quarter in which the forecasts were made. We do this for the first quarter of each year since 1983 and present the data in the form of a scatter plot. For each point in the figure, the horizontal axis measures disagreement computed from the first quarter survey of each year, and the vertical axis measures consumption growth in the quarter in which that survey was taken. Similarly, the figure shows the scatter plot for disagreement for real GDP growth plotted against actual investment growth.

What we see in both panels is that the points have a tendency to line up down and to the right. This suggests that when disagreement is high, consumption growth and investment growth tend to be low. The regression trend line that is plotted in each figure (the solid black line) confirms this visual impression. This line is the best-fitting line through the points in the figure. The fact that the line in each figure trends down and to the right confirms that when disagreement is high, consumption and investment growth tend to be low.

We construct similar plots in Figure 6, which shows the relationship between uncertainty about real GDP growth and consumption and investment growth. We measure uncertainty using the standard deviation from the histograms reported in the SPF surveys for real GDP growth. Because of the data limitations discussed above, we use data only from 1991 onward for these figures. The uncertainty measure pertains to forecasted annual real GDP growth for the year in which the survey was taken (we again use the SPF from the first quarter of each year), and consumption and investment growth are measured in the quar-

FIGURE 6

Forecast Uncertainty Versus Consumption and Investment Growth

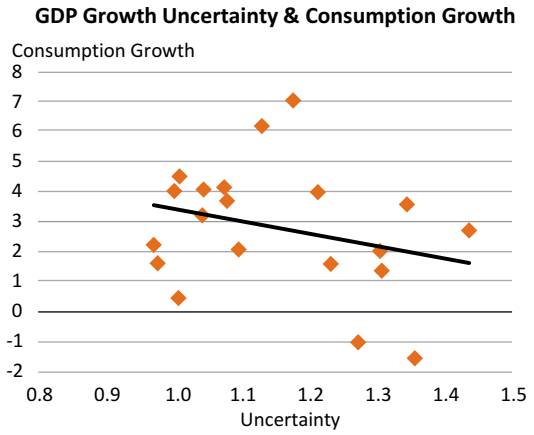

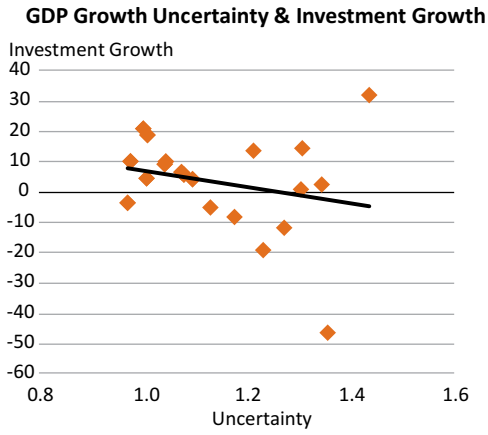

Left panel: Each point measures the relationship between consumption growth and uncertainty about real GDP growth.

Right panel: Each point measures the relationship between uncertainty and investment growth. Uncertainty is measured using the standard deviation from the histograms reported in the SPF for real GDP growth. The uncertainty measure pertains to forecasted annual real GDP growth for the year in which the survey was taken (using the SPF from the first quarter of each year), and consumption and investment growth are measured in the quarter in which the survey was taken.

Source: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters and author's calculations

ter in which the survey was taken.

These figures look quite similar to those that investigated forecast disagreement and growth. In particular, there is a tendency for consumption and investment growth to be low when measured uncertainty is high. As is the case for Figure 5, the best-fitting trend line again slopes down and to the right, confirming a negative relationship between uncertainty and consumption and investment growth.

What about inflation uncertainty? Monetary policymakers care about inflation uncertainty, since it relates to their credibility as guardians of price stability. The Fed's dual mandate includes maintaining low and stable inflation. To the extent that policymakers can achieve this goal, price level changes will be fairly predictable over the medium and long terms for households and firms. This, in turn, should help to make their decision-making somewhat easier. Thus, policymakers care about what level of expected inflation households and firms have and how that expectation changes over time. Is there a relationship between expected inflation and uncertainty? The paper by Rich and Tracy investigates this question using SPF data. What they find is that average uncertainty across forecasters about inflation and expected inflation from the SPF does not appear to be strongly related. However, forecaster disagreement and expected inflation do appear to be related: Higher disagreement about inflation is associated with higher expected inflation.

We can see this relationship in Figure 7, which is a scatter plot of forecaster disagreement about GDP deflator inflation against their forecast of future inflation. The inflation forecast is for quarterly GDP deflator inflation four quarters ahead. The data are annual, measured in the first quarter SPF for each year from 1983 to 2011. The band of high-inflation points, marked in orange, is observations from the 1980s. We again plot the best-fitting trend line to the data, and it shows up as the solid, upward-sloping line in the figure.

The figure shows the tendency found by Rich and Tracy: Higher levels of disagreement about inflation are associated with higher expected inflation. As Rich and Tracy point out, the economic theory behind this apparent relationship is currently a bit thin, especially since their analysis indicates that other uncertainty measures for inflation are not very significantly correlated with expected inflation. It would seem to indicate that forecasters are using quite different models to forecast inflation and that, as inflation rises, those models are leading to quite different predictions about future inflation.

FIGURE 7

Mean Inflation Forecast and Forecast Dispersion

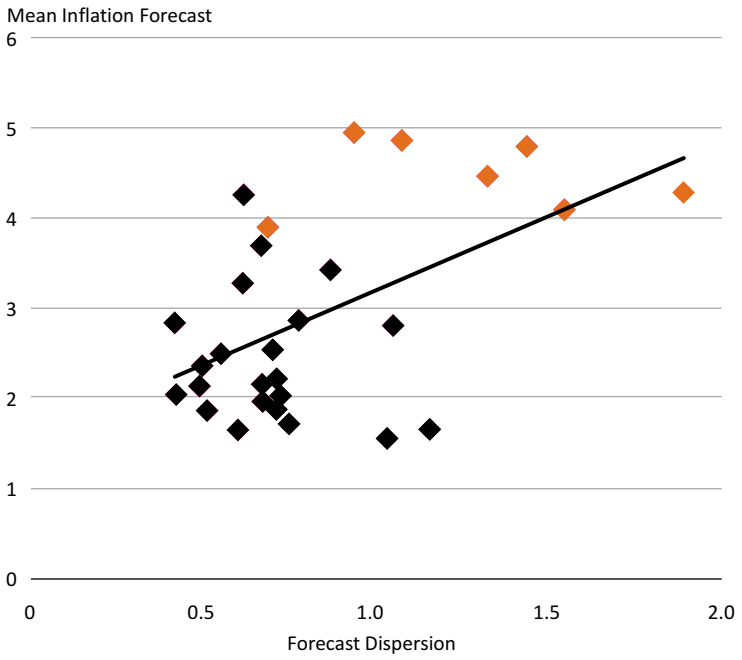

Plot of forecaster disagreement about GDP deflator inflation against forecast of future inflation. Inflation forecast is quarterly GDP deflator inflation four quarters ahead. Annual data, measured in the first quarter SPF for 1983 to 2011. High-inflation points, in orange, are observations from the 1980s.

Source: Federal Reserve Bank of Philadelphia Survey of Professional Forecasters and author's calculations

CONCLUSION

Economic uncertainty is an important facet of decision-making for households, firms, and policymakers. The data on economic uncertainty are not readily available and usually must be gleaned from forecast surveys. The SPF is somewhat unique in that, in addition to standard measures of forecast disagreement, it provides direct measures of uncertainty from its respondents. This has made the SPF a valuable tool for researchers investigating the link between economic uncertainty and economic outcomes.

REFERENCES

Bachmann, Ruediger, Steffen Elstner, and Eric Sims. "Uncertainty and Economic Activity: Evidence from Business Survey Data," NBER Working Paper 16143 (2010).

Boero, Gianna, Jeremy Smith, and Kenneth F. Wallis. "Uncertainty and Disagreement in Economic Prediction: The Bank of England Survey of External Forecasters," *Economic Journal*, 118 (2009), pp. 1107-27.

Bomberger, William A. "Disagreement as a Measure of Uncertainty," *Journal of Money, Credit and Banking*, 28 (1996), pp. 381-92.

Carroll, Christopher D. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112:1 (1997), pp. 1-55.

Croushore, Dean. "Introducing the Survey of Professional Forecasters," Federal Reserve Bank of Philadelphia *Business Review* (November/December 1993), pp. 3-13.

Deaton, Angus. "Saving and Liquidity Constraints," *Econometrica*, 59:5 (1991), pp. 1221-48.

Giordani, Paolo, and Paul Soderlind. "Inflation Forecast Uncertainty," *European Economic Review*, 47 (2003), pp. 1037-59.

Guerron, Pablo. "Risk and Uncertainty," Federal Reserve Bank of Philadelphia *Business Review* (First Quarter 2012).

McDonald, Robert, and Daniel Siegel. "The Value of Waiting to Invest," *Quarterly Journal of Economics*, 101:4 (1986), pp. 707-28.

Mishkin, Frederic S. "Outlook and Risks for the U.S. Economy," presented at the National Association for Business Economics Washington Policy Conference, Washington, D.C.

Rich, Robert, and Joseph Tracy. "The Relationship Among Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts," *Review of Economics and Statistics*, 92:1 (2009), pp. 200-07.

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/.

EXTENDING THE SCOPE OF THE EU'S INNOVATION UNION

The Innovation Union initiative of the European Union focuses on product and process innovation for tangible goods. The authors argue that it is essential to extend the scope of the initiative to include innovation for financial sector products, processes, and regulatory approaches. They make this argument using examples of financial sector innovations in the United States following the Great Depression and on the basis of an examination of the 2008 financial crisis.

Working Paper 12-17, "Building the Innovation Union: Lessons from the 2008 Financial Crisis," Alice O. Nakamura, University of Alberta School of Business; Leonard I. Nakamura, Federal Reserve Bank of Philadelphia; and Masao Nakamura, University of British Columbia, Sauder School of Business

ANALYZING CREDIT RISK UNDER ECONOMIC STRESS CONDITIONS

The authors develop an empirical framework for the credit risk analysis of a generic portfolio of revolving credit accounts and apply it to analyzing a representative panel data set of credit card accounts from a credit bureau. These data cover the period of the most recent deep recession and provide the opportunity to analyze the performance of such a portfolio under significant economic stress conditions. The authors consider a traditional framework for the analysis of credit risk where the probability of default (PD), loss given default (LGD), and exposure at default (EAD) are explicitly considered. The unsecure and revolving nature of credit card lending is naturally modeled in this framework. Their results indicate that unemployment, and, in particular, the level and change in unemployment, plays a significant role in the probability of transition across delinquency states in general and the probability of default in particular. The effect is heterogeneous and proportionally has a more significant impact for high credit score and for high-utilization accounts. The authors' results also indicate that unemployment and a downturn in economic conditions play a quantitatively small, or even irrelevant, role in the changes in account balance associated with changes in an account's delinquency status and in the exposure at default specifically. The impact of a downturn in economic conditions, in particular, changes in unemployment, on the recovery rate and loss given default is found to be large. These findings are of particular relevance for the analysis of credit risk regulatory capital under the IRB approach within the Basel II capital accord.

Working Paper 12-18, "Credit Risk Analysis of Credit Card Portfolios Under Economic Stress Conditions," Piu Banerjee, Federal Reserve Bank of New York, and José J. Canals-Cerdá, Federal Reserve Bank of Philadelphia

PRIVATE VS. PUBLIC MONETARY SYSTEM

The author shows the existence of an inherent instability associated with a purely private monetary system due to the role of endogenous debt limits in the creation of private money. Because the banker's ability to issue liabilities that circulate as a medium of exchange depends on beliefs about future credit conditions, there can be multiple equilibria. Some of these equilibria have undesirable properties: Self-fulfilling collapses of the banking system and persistent fluctuations in the aggregate supply of bank liabilities are possible. In response to this inherent instability of private money, the author formulates a government intervention that guarantees that the economy remains arbitrarily close to the constrained efficient allocation. In particular, the author defines an operational procedure for a central bank capable of ensuring the stability of the monetary system.

Working Paper 12-19, "On the Inherent Instability of Private Money," Daniel R. Sanches, Federal Reserve Bank of Philadelphia

ANALYZING THE GROWTH OF U.S. MANUFACTURING EXPORTS

The authors study empirically and theoretically the growth of U.S. manufacturing exports from 1987 to 2007. They identify the change in iceberg costs with plant-level data on the intensity of exporting by exporters. Given this change in iceberg costs, the authors find that a GE model with heterogeneous establishments and a sunk cost of starting to export is consistent with both aggregate U.S. export growth and the changes in the number and size of U.S. exporters. The model also captures the nonlinear dynamics of U.S. export growth. A model without a sunk export cost generates substantially less trade growth and misses out on the timing of export growth. Contrary to the theory, employment was largely reallocated from very large establishments, those with more than 2,500 employees, toward very small manufacturing establishments, those with fewer than 100 employees. Allowing for faster productivity growth in manufacturing, changes in capital intensity, and some changes in the underlying shock process makes the theory consistent with the changes in the employment size distribution. The authors also find that the contribution of trade to the contraction in U.S. manufacturing employment is small.

Working Paper 12-20, "Do Falling Iceberg Costs Explain Recent U.S. Export Growth?," George Alessandria, Federal Reserve Bank of Philadelphia, and Horag Choi, Monash University

DISTRIBUTIONAL EFFECTS OF MONETARY POLICY ACROSS SOCIO-ECONOMIC GROUPS

The authors build a New Keynesian model in which heterogeneous workers differ with regard to their employment status due to search and matching frictions in the labor market, their potential labor income, and their amount of savings. They use this laboratory to quantitatively assess who stands to win or lose from unanticipated monetary accommodation and who benefits most from systematic monetary stabilization policy. They find substantial redistribution effects of monetary policy shocks; a contractionary monetary policy shock increases income and welfare of the wealthiest 5 percent, while the remaining 95 percent experience lower income and welfare. Consequently, the negative effect of a contractionary monetary policy shock to social welfare is larger if heterogeneity is taken into account.

Working Paper 12-21, "Monetary Policy with Heterogeneous Agents," Nils Gornemann University of Pennsylvania; Keith Kuester, formerly Federal Reserve Bank of Philadelphia; Makoto Nakajima, Federal Reserve Bank of Philadelphia

USING DISTANCE-BASED ECONOMETRIC TECHNIQUES TO ANALYZE THE SPATIAL CONCENTRATION OF R&D LABS

The authors study the location of more than 1,000 research and development (R&D) labs located in the Northeast corridor of the U.S. Using a variety of spatial econometric techniques, they find that these labs are substantially more concentrated in space than the underlying distribution of manufacturing activity. Ripley's K-function tests over a variety of spatial scales reveal that the strongest evidence of concentration occurs at two discrete distances: one at about one-quarter of a mile and another at about 40 miles. They also find that R&D labs in some industries (e.g., chemicals, including drugs) are substantially more spatially concentrated than are R&D labs as a whole.

Tests using local K-functions reveal several concentrations of R&D labs that appear to represent research clusters. The authors verify this conjecture using significance maximizing techniques (e.g., SATSCAN) that also address econometric issues related to "multiple testing" and spatial autocorrelation.

The authors develop a new procedure for identify-

ing clusters – the multiscale core-cluster approach, to identify labs that appear to be clustered at a variety of spatial scales. Locations in these clusters are often related to basic infrastructure such as access to major roads. There is significant variation in the industrial composition of labs across these clusters. The clusters the authors identify appear to be related to knowledge spillovers: Citations to patents previously obtained by inventors residing in clustered areas are significantly more localized than one would predict from a (control) sample of otherwise similar patents.

Working Paper 12-22, "The Agglomeration of R&D Labs," Gerald A. Carlino, Federal Reserve Bank of Philadelphia; Robert M. Hunt, Federal Reserve Bank of Philadelphia; Jake K. Carr, Ohio State University; and Tony E. Smith, University of Pennsylvania

DEVELOPING NARRATIVE MEASURES OF FEDERAL GRANTS-IN-AID PROGRAMS

Because of lags in legislating and implementing fiscal policy, private agents can often anticipate future changes in tax policy and government spending before these changes actually occur, a phenomenon referred to as fiscal foresight. Econometric analysis that fails to model fiscal foresight may obtain tax and spending multipliers that are biased. One way researchers have attempted to deal with the problem of fiscal foresight is by examining the narrative history of government revenue and spending news. The Great Recession and efforts by the federal government through the American Recovery and Reinvestment Act of 2009 (ARRA) to stimulate the economy returned fiscal policy, and in particular the role of state and local governments in such policies, to the center of macro-economic policy-making. In a companion paper, the authors use federal grants-in-aid to state and local governments to provide an evaluation of the effectiveness of the ARRA. The purpose of this paper is to develop narrative measures of the federal grants-in-aid programs beginning with the Federal Highway Act of 1956 through the ARRA of 2009.The narrative measures the authors develop will be used as instruments for federal grants-in-aid in their subsequent analysis of the ARRA.

Working Paper 12-23, "A Narrative Analysis of Post-World War II Changes in Federal Aid," Gerald Carlino, Federal Reserve Bank of Philadelphia, and Robert Inman, Wharton School, University of Pennsylvania, and Visiting Scholar, Federal Reserve Bank of Philadelphia

DO THE SOURCES OF ECONOMIC AND FINANCIAL CRISES DIFFER FROM NONCRISIS BUSINESS CYCLE FLUCTUATIONS?

This paper explores the hypothesis that the sources of economic and financial crises differ from noncrisis business cycle fluctuations. The authors employ Markov-switching Bayesian vector autoregressions (MS-BVARs) to gather evidence about the hypothesis on a long annual U.S. sample running from 1890 to 2010. The sample covers several episodes useful for understanding U.S. economic and financial history, which generate variation in the data that aids in identifying credit supply and demand shocks. The authors identify these shocks within MS-BVARs by tying credit supply and demand movements to inside money and its intertemporal price. The model space is limited to stochastic volatility (SV) in the errors of the MS-BVARs. Of the 15 MS-BVARs estimated, the data favor an MS-BVAR in which economic and financial crises and noncrisis business cycle regimes recur throughout the long annual sample. The best-fitting MS-BVAR also isolates SV regimes in which shocks to inside money dominate aggregate fluctuations.

Working Paper 12-24, "Business Cycles and Financial Crises: The Roles of Credit Supply and Demand Shocks," James M. Nason, Federal Reserve Bank of Philadelphia, and Ellis W. Tallman, Oberlin College and Federal Reserve Bank of Cleveland

DEVELOPMENT CONSTRAINTS AND LAND RENTS

A tractable production-externality-based circular city model in which both firms and workers choose location as well as intensity of land use is presented. The equilibrium structure of the city has either (i) no commuting ("mixed-use" form) or (ii) a central business district (CBD) of positive radius and a surrounding residential ring. Regardless of which form prevails, the intra-city variation in all endogenous variables displays the negative exponential form: $x(r) = x(0)e^{-\phi_x r}$ (where r is the distance from the city center and ϕ_x depends only on preference and technology parameters). An application is presented wherein it is shown that population growth may lead to a *smaller* increase in land rents in cities that cannot expand physically because these cities are less able to exploit the external effect of greater employment density.

Working Paper 12-25, "A Tractable Circular City

Model with an Application to the Effects of Development Constraints on Land Rents," Satyajit Chatterjee, Federal Reserve Bank of Philadelphia, and Burcu Eyigungor, Federal Reserve Bank of Philadelphia

EXPLAINING FIRMS' ENTRY, EXIT, AND RELOCATION DECISIONS IN AN URBAN ECONOMY WITH MULTIPLE LOCATIONS

The authors develop a new dynamic general equilibrium model to explain firm entry, exit, and relocation decisions in an urban economy with multiple locations and agglomeration externalities. The authors characterize the stationary distribution of firms that arises in equilibrium. They estimate the parameters of the model using a method of moments estimator. Using unique panel data collected by Dun and Bradstreet, the authors find that their model fits the moments used in estimation as well as a set of moments that they use for model validation. Agglomeration externalities increase the productivity of firms by about 8 percent. Economic policies that subsidize firm relocations to the central business district increase agglomeration externalities in that area. They also increase economic welfare in the urban economy.

Working Paper 12-26, "Estimating a Dynamic Equilibrium Model of Firm Location Choices in an Urban Economy," Jeffrey Brinkman, Federal Reserve Bank of Philadelphia; Daniele Coen-Pirani, University of Pittsburgh; and Holger Sieg, University of Pennsylvania and NBER