Liquidity:
Meaning, Measurement, Management

Robert E. Lucas, Jr.

"We are Canada's central bank. We work to preserve the value of money by keeping inflation low and stable."
http://bankofcanada.ca/

I like this brief mission statement from the Bank of Canada, both for what it includes and what it leaves out. I think Homer Jones would have liked it too. The maintenance of consistently low inflation rates is an important goal that central banks can actually achieve and, in fact, have been achieving quite well over the past 25 years. One objective of this article is to describe the underlying economic theory that tells us how Federal Reserve actions can attain this goal and the evidence that supports this theory.

Financial stability is another classic responsibility of the Fed. Financial panics are the results of sudden declines in liquidity, and we know that Fed actions can partially offset liquidity declines in an economy on short notice. Is there an underlying economic theory that tells us how this can be done, or better yet, how to arrange matters so that liquidity declines are less likely to occur in the first place? There has been a flood of economic theorizing on this question since 2008, much of it valuable and stimulating, but nothing like the solid evidence that guides the monetary control of inflation has yet been found. My second objective is to explain why short-run monetary stability is so much harder to control than long-run inflation rates.

THE FED AND THE PAYMENTS SYSTEM

The main—really, the *only*—tool the Fed has for influencing the inflation rate is its ability to add to or subtract from the amount of government-issued money—I'll call it cash—in the hands of the public. This is done, of course, by selling government bonds for cash, decreasing

Robert E. Lucas, Jr. is the John Dewey Distinguished Service Professor in Economics at the University of Chicago and received the Nobel Prize in economics in 1995.

Figure 1

Various Monthly Yields: Low-Frequency Only
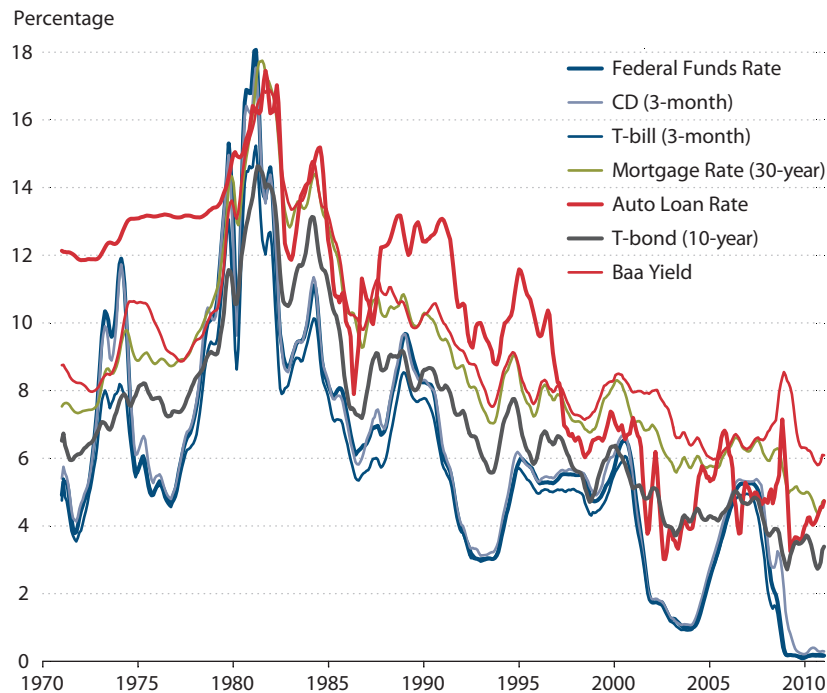
SOURCE: Updated from Kim, Hagen. "Common and Idiosyncratic Fluctuations of Interest Rates from Various Issuers: A Dynamic Factor Approach." Ph.D. dissertation, University of Chicago, 2003. Used with permission.

the cash held by the public, or buying bonds with cash, and so increasing the cash outstanding. Here I take a Monetarist view that (i) open market operations affect the quantity of liquidity in the system—the money supply, if you like—and (ii) changes in this magnitude are what determine inflation rates.

These operations affect interest rates, too—how could they not?—but the assets used in monetary policy are money market instruments with yields that are not closely related to interest rates in general. This is clear from Figure 1, which plots seven different interest rates over the period 1970-2012. The three blue curves on the figure are the federal funds rate and rates for T-bills and 3-month certificates of deposit (CDs). They are all alike. High-grade commercial paper would fit right in too. The other four series are rates for auto loans, home mortgages, long-term bonds, and the yield on Baa-rated securities. One can see from the figure that these four do not closely track the three money market yields. No one uses these money market instruments to finance an investment project, a new house, or a college education. These are specialized assets that are useful in cash management: interest-bearing securities that can be easily converted into non-interest-bearing cash as payments come due. In a cashless economy (if there were such a thing), they would not exist at all.

The only reason anyone ever holds cash is to be ready to make payments to someone. That's it. There is no other reason to hold this dominated asset—an asset that has no risk advantage over other assets that offer a higher return. This fact gives rise to a host of other assets that people believe can be traded for cash on short notice, on predictable terms, and without undue labor costs. These qualities define the terms *liquidity* and *liquid asset*. These terms have no meaning outside a payments system. What we mean by liquid assets, then, is a combination of government-issued cash together with various kinds of promissory notes—IOUs, really—that play a role in the payments system by economizing on cash. Like cash, all these money market instruments have lower yields than other assets with the same risk characteristics: They command a liquidity premium. (Again, look at Figure 1.)

It would be useful to have a theoretical understanding of the relationships of different money market instruments to one another. One might think of these assets as being ordered according to the size of their liquidity premiums: government-issued cash first, then bank demand deposits, then Treasury bills and other money market instruments, and so on. As we go down such a list, the effort involved in converting each asset to cash increases. The convenience of the most-liquid assets offsets their lower yields. Think of a "technology" where labor (ours and that of people we employ) and various liquid assets are combined to generate the highest return consistent with payment objectives without increasing risk.

INFLATION AND THE QUANTITY THEORY OF MONEY

We do not now have an established model of such a payments system (though there is a great deal of current research that seems headed in this direction). So I will do what theorists do and replace the complicated reality sketched above with a much simpler, fictional world that will serve as a useful surrogate for some purposes. The surrogate I describe yields some sharp predictions on observable variables. As I will illustrate shortly, these predictions conform extremely well to historical evidence.

Step back from the complexity of actual money markets, then, and imagine an economy in which every payment is made in cash: There is no choice among payment methods and no one will accept any sort of IOU. In this setting, money holdings are well defined and it is natural to assume that average money holdings of individuals and firms will be *proportional* to spending flows: For some constant K,

$$\text{\$ Spending/year} = K \times (\text{Average \$ money holdings}).$$

If my spending doubles, then I want to hold twice as much cash. If the money supply doubles, so will everyone's spending.

In terms of growth rates, the above formula implies that

$$\text{\% Spending growth/year} = \text{\% Money growth/year}.$$

Finally, we add the hypothesis that real growth is a constant 3 percent per year, unaffected by monetary policy. Then,

Figure 2

Money Growth and Inflation: A High, Positive Correlation
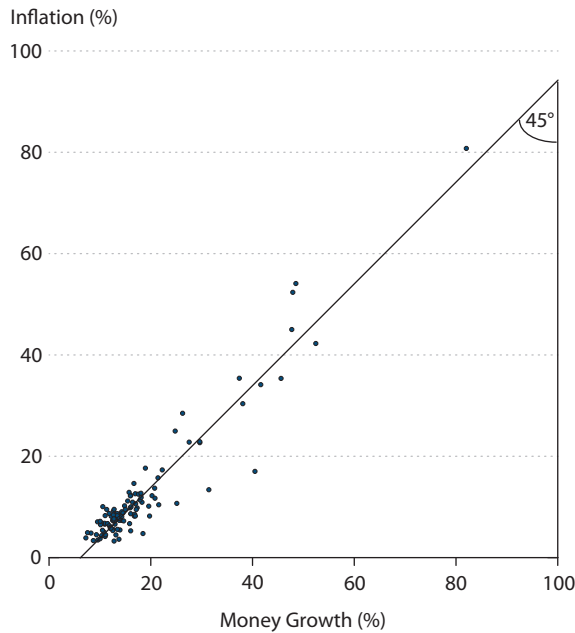
NOTE: The figure shows average annual rates of growth in M2 and in consumer prices during 1960-90 in 110 countries.

SOURCE: Data source IMF. Reproduced with permission from McCandless and Weber (1995, Chart 1, p. 5).

Table 1

Correlation Coefficients for Money Growth and Inflation

Sample	Coefficient for each definition of money		
	M0	M1	M2
All 110 countries	0.925	0.958	0.950
Subsamples			
21 OECD countries	0.894	0.940	0.958
14 Latin American countries	0.973	0.992	0.993

NOTE: Based on data from 1960-90. Inflation is defined as changes in a measure of consumer prices.
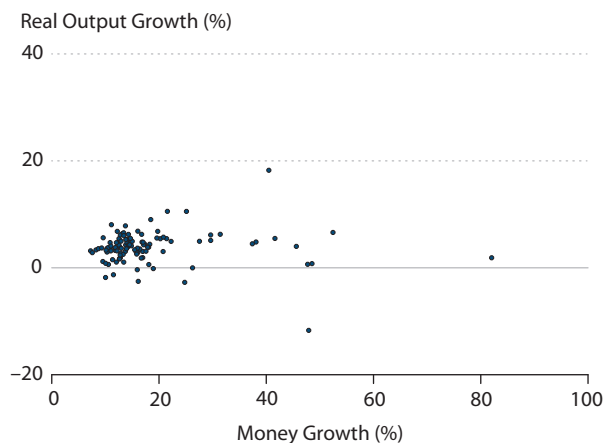
SOURCE: Data source IMF. Reproduced with permission from McCandless and Weber (1995, Table 1, p. 4).

$$\text{Inflation} = \% \text{ Money growth/year} - 3\%.$$

This formula is one expression of what we call the quantity theory of money.

Can these strong assumptions possibly express some useful facts? Let's look at some evidence. To do so, we first need to match objects in the model as best we can with objects we can measure. The National Economic Accounts provide measures of inflation, or we can use the very similar consumer price index (CPI). For the money supply, we have an embarrassment of riches because there is so little agreement on the meaning of the term "money supply" or even whether this idea is worth measuring. (Of course, the same problem arises with the term "liquidity.") Despite this, many measures of money have been proposed and usefully applied over the years.

The measures of money include M0 (actual government-supplied cash), M1 (publicly held currency plus demand deposits), M2 (M1 plus time deposits), and the even broader M3. All of these contain some information about the liquidity in an economy. If we use any but the first of these, we are treating a complicated mix of different liquid assets as though it were a homogeneous "money supply." Clearly there is no right way to do this, but there are useful ways.

Figure 3

Money and Real Output Growth: No Correlation in the Full Sample…

Real Output Growth (%)

Money Growth (%)

NOTE: The figure shows average annual rates of growth in M2 and in nominal GDP, deflated by consumer prices during 1960-90 in 110 countries.

SOURCE: Data source IMF. Reproduced with permission from McCandless and Weber (1995, Chart 2, p. 7).

Figure 4

…But a Positive Correlation in the OECD Subsample

Real Output Growth (%)

Money Growth (%)
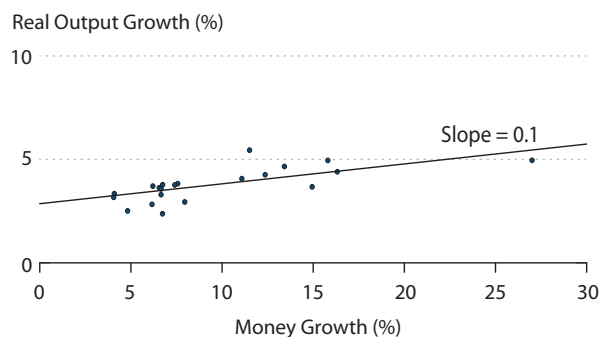
NOTE: The figure shows average annual rates of growth in M0 and in nominal GDP, deflated by consumer prices during 1960-90 in 21 countries.

SOURCE: Data source IMF. Reproduced with permission from McCandless and Weber (1995, Chart 3, p. 8).

We look first at cross-country evidence from International Monetary Fund (IMF) data from 110 countries over the period 1960-90, as presented by George McCandless and Warren Weber (1995). Several different measures of money are used in the figures and tables that follow.

The results are a striking confirmation of the one-for-one effects of money growth on inflation. The money measure used in Figure 2 is M2, but Table 1 shows that M0 and M1 yield very similar results. Figure 3 plots average real gross domestic product (GDP) growth over the same period and the same countries. Using the entire sample, they appear unrelated. For the subsample of members of the Organisation for Economic Co-operation and Development (OECD), there is a slight positive correlation of money growth and real growth (Figure 4). Other studies of money growth and real variables have also been inconclusive. There is no systematic relation between real and nominal growth, but some kind of relation cannot be ruled out.

We have looked at 30-year average inflation rates across countries. Can we learn more from data over time within specific countries? This is a harder question, but worth pursuing. Here I rely on a series of papers by Luca Benati (2005, 2009). Figures 5 and 6 illustrate the methodology. Figures 7 through 9 provide more evidence from several countries and for differ-ent choices of monetary aggregates.

Figure 5 plots the monetary base (M0) and inflation rates for the United Kingdom over the period 1870-2005. The two series tend to move together, but the details are hard to read:
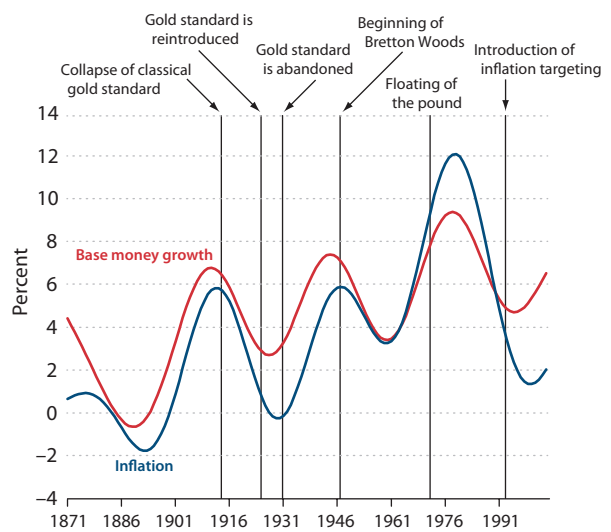
Figure 5

Inflation and Money Growth in the United Kingdom: Raw Data
Composite Index Price* and M0 (annual rates of changes)
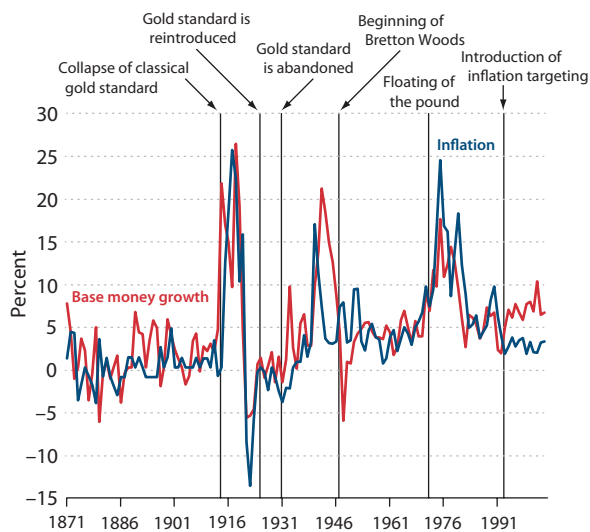
NOTE: *See O'Donoghue, Goulding, and Allen (2004).
SOURCE: Reproduced from Benati (2005, Chart 1a) with permission from the Bank of England.

Figure 6

Inflation and Money Growth in the United Kingdom: Components Beyond 30 Years
Composite Index Price and M0 (annual changes)

SOURCE: Reproduced from Benati (2005, Chart 2a) with permission from the Bank of England.

How similar are these series? Figure 6 offers one way to answer this question. It is based on the same data used in Figure 5 but with the high-frequency movements removed using a statistical filter. The filtering process loses a lot of information, but in return provides a crystal clear picture of the long-run relation between base money and inflation rates. These two filtered time series exhibit exactly the same one-for-one relation seen in the cross sections!

Figures 7 through 9 all apply the same filter used in Figure 6. They differ in the countries studied, the time periods, and the measures of the money supply used. There are interesting differences among these, to be sure, but it is easy to see the predominance of this one-for-one relation in all of them.

These data compare filtered versions of two completely different time series—inflation rates (determined by millions of individual prices) and money growth rates (based on a variety of liquidity measures), over which the Fed has decisive influence. People who do not do empirical work in economics (which includes many economists!) may not appreciate just how good these results are, but I would not recommend a scientific career to anyone who is not impressed by this evidence.

Figure 7

Low-Frequency Components of Inflation and Money Growth Since the Gold Standard Era: Norway, United Kingdom, Sweden, and Switzerland

Norway (1830-2004)

United Kingdom (1871-2005)

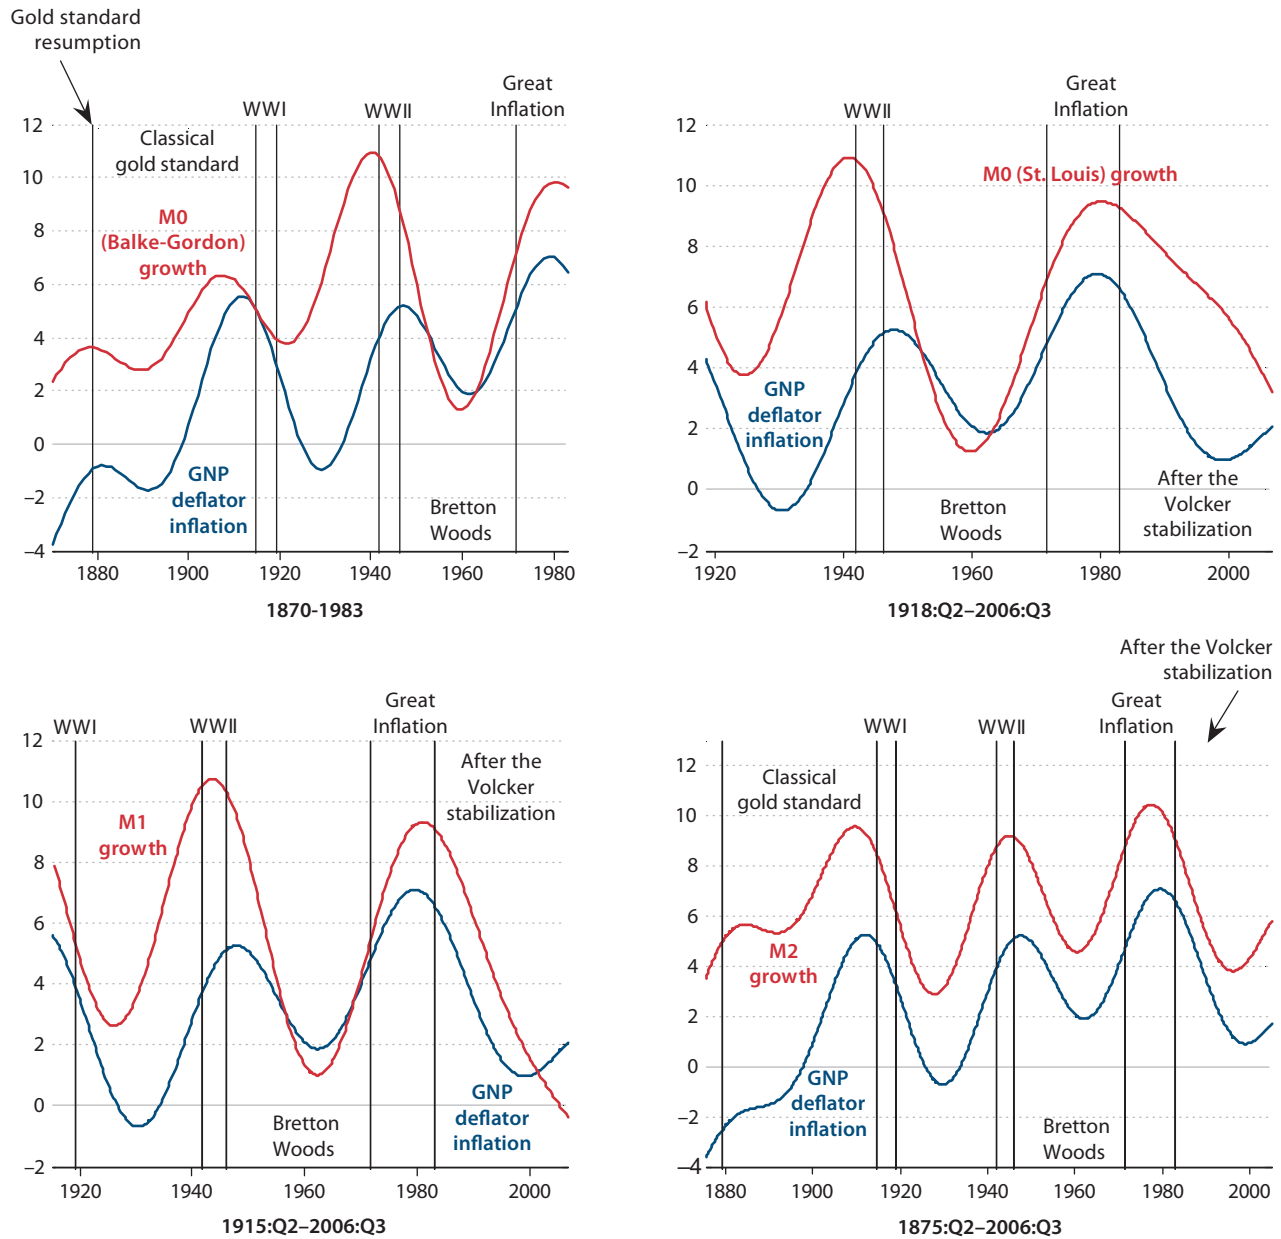**Sweden (1879-1988)**

Switzerland (1908-2005)
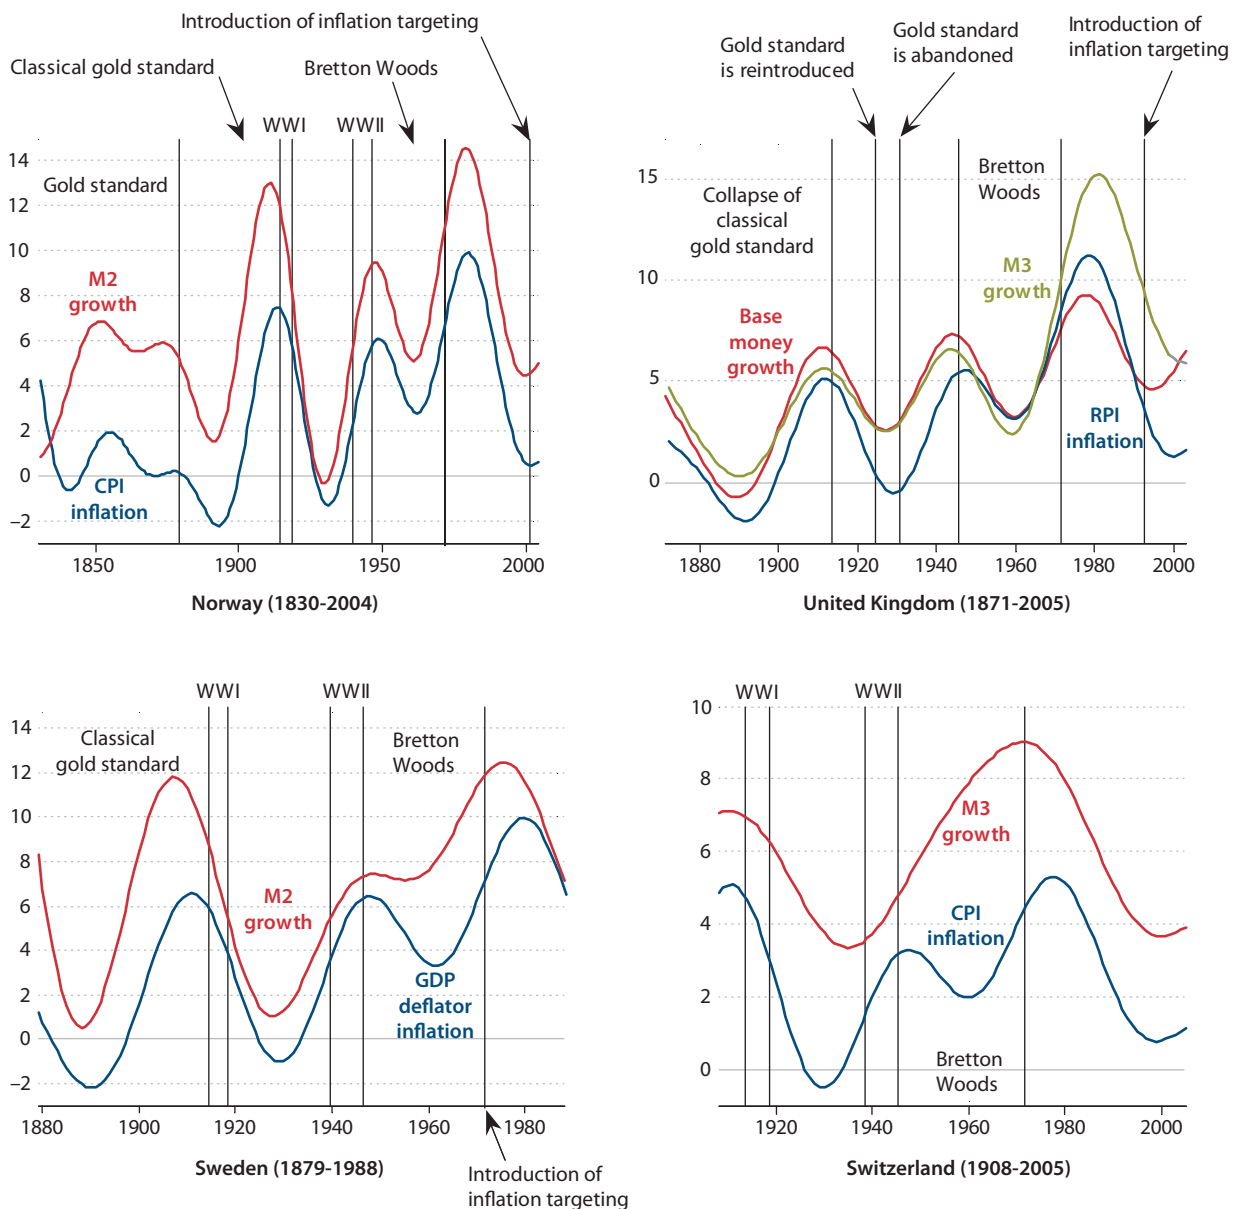
NOTE: RPI, retail price index.

SOURCE: Benati (2009). Used with permission.

Figure 8

Low-Frequency Components of Inflation and Money Growth Since the Gold Standard Era: United States

1870-1983

1918:Q2–2006:Q3

1915:Q2–2006:Q3

1875:Q2–2006:Q3

NOTE: GNP, gross national product.

SOURCE: Benati (2009). Used with permission.

Figure 9

Low-Frequency Components of Inflation and Money Growth Over the Post-WWII Era: Euro Area, United Kingdom, Canada, Sweden, Australia, and Japan

Euro Area (1970:Q2–2006:Q3)

United Kingdom (1963:Q2–2006:Q3)

Canada (January 1946–November 2006)

Sweden (January 1961–November 2006)

Australia (1965:Q2–2006:Q3)

Japan (January 1957–November 2006)
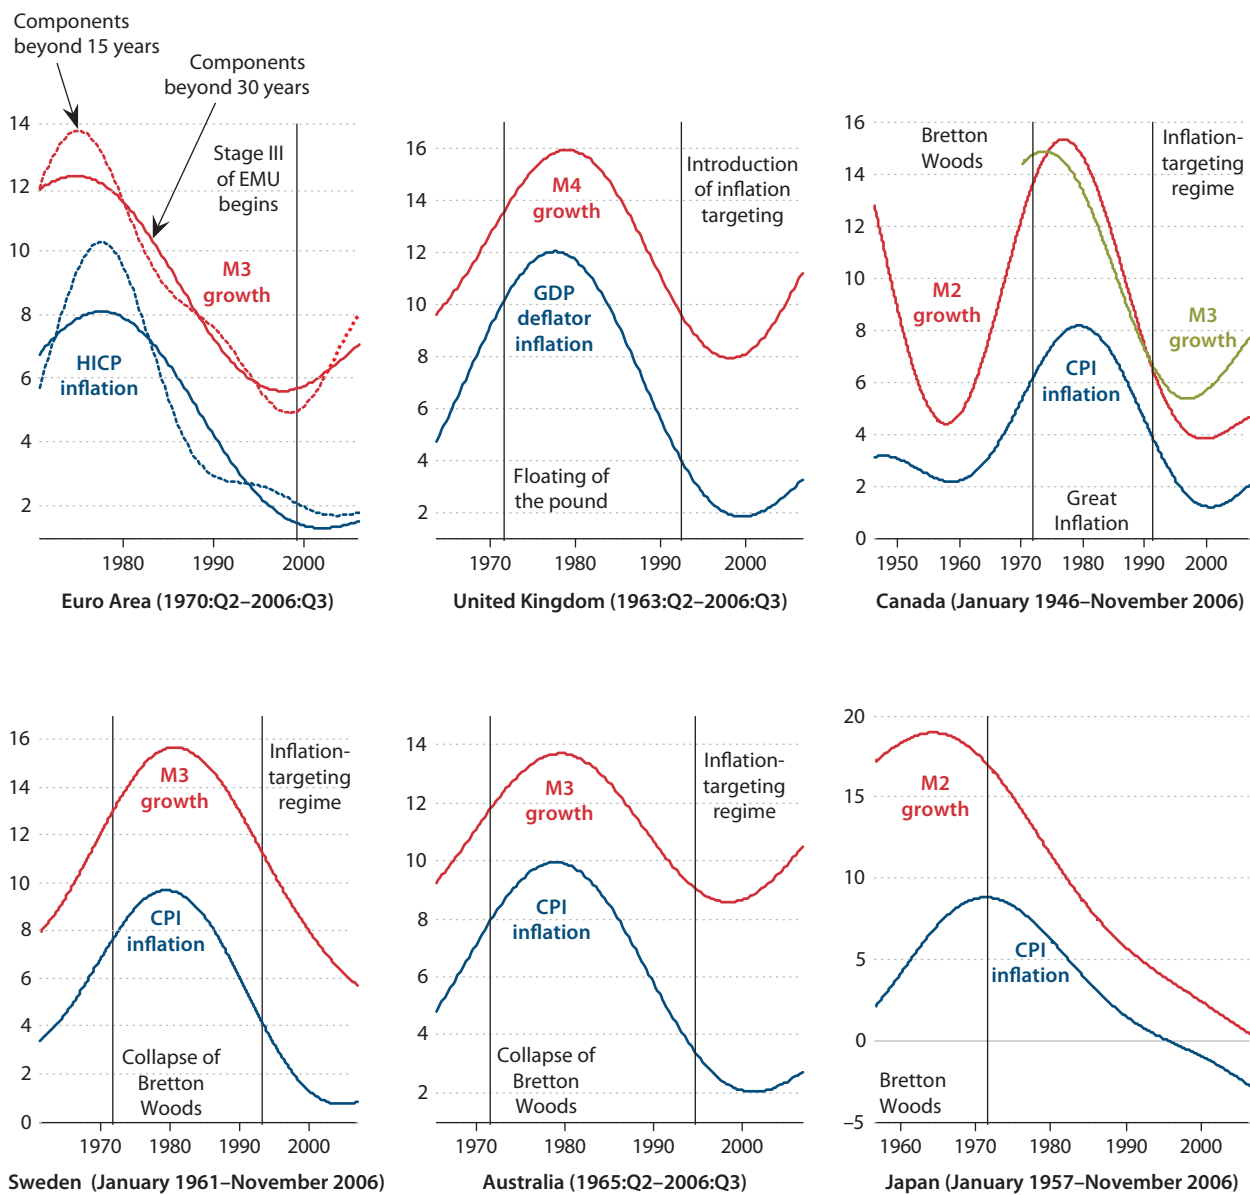
NOTE: EMU, European Economic and Monetary Union; HICP, Harmonised Index of Consumer Prices.

SOURCE: Benati (2009). Used with permission.

But are these results simply too good to be true? We know that payments systems in reality involve a mix of many different liquid assets. How can a theory that ignores these complexities, as the quantity theory does, do so well? It must be the case that, at least in the long run, all these different forms of liquidity grow at something like a common rate, maintaining fairly stable shares in money market portfolios over time. If this is so, it would not be surprising if measures such as M0, M1, and so on simply reflect different collections of liquid assets. This idea is disturbingly vague, I know. It would be a real step forward if it could be replaced by a more precisely stated hypothesis.

Accepting some form of the quantity theory of money, then, we have a theory of inflation. How do we apply this theory to control the inflation rate? One possibility would be to commit to a constant rate of money growth, chosen to imply the desired inflation rate. But what if the theory is not quite accurate or becomes obsolete because of changes in banking practices? A better, more flexible idea, developed in the 1970s, is inflation targeting. The central bank sets a goal, a target of (say) 2 percent—or perhaps a range such as 1 to 3 percent—and if actual inflation goes above the targeted range, the bank reduces the rate of money growth. If inflation continues above 3 percent, the bank reduces the money growth rate some more and keeps taking money out of the system until inflation is brought back within the targeted range. How do we know this will work? All the evidence I have just presented implies that it will. (It also implies that this may take some time.) Central banks all over the world that have used inflation targeting (stated or not) have had remarkable success in keeping inflation low and stable since the early 1980s, even during and after the crisis of 2008.

I have described inflation targeting as though it is implemented by adding to or reducing the money supply. As noted previously, today it is more customary to describe monetary policy as carried out by changes in short-term money market rates, not as changes in the money supply. In normal times, changes in the money supply are effected through trades in money markets; expanding the money supply and reducing the federal funds rate are just two ways of saying the same thing. There is not much reason to argue over the terminology. But since 2008 the funds rate has been so close to zero that reducing the rate is either impossible or meaningless. We talk about "unconventional policy" and "quantitative easing." From the Monetarist view that I am taking, of course, nothing could be more conventional than "quantitative easing." It is just another way to say "increasing the money supply." The Fed does this by buying something other than short-term government securities—long-term governments bonds, for example—but almost any asset would do.

THE NATURE OF FINANCIAL CRISES

I have spent most of my time on the easy parts of monetary policy: the actions that central banks know how to use to stabilize inflation. I have outlined some of the evidence on which these actions are largely based. I have noted the remarkable success of inflation targeting, persisting through the crisis of 2008-09 and the years since. Nevertheless, it must also be said that the financial crisis of 2008 was very costly in terms of lost production, and there is little

Table 2

U.S. Payments Flows (2013)

Cash transfers	Annual flow ($ billions)	Average transaction size ($ thousands)
Automated clearinghouse	24,400	2
Checks processed	8,000	1.3
Fedwire Funds Service	713,000	5,300
Fedwire Securities Service	295,000	15,500
National Settlement Service	17,000	

agreement even among leading economists on how it happened and how it might have been prevented. Some additional discussion of this event seems warranted.

The quantity theory that I have described is consistent with the idea that a sudden reduction in the money supply leads to deflation and reduction in spending. This happens because a sudden loss of liquidity leads people to reduce spending to rebuild a desired ratio of cash to spending flows. But central banks do not just suddenly reduce the money supply. They did not do so in the 1929 crash or in 2008-09. The source of a financial crisis or panic generally lies in a loss of confidence in privately issued promises of cash, not in reductions in government-supplied money. If so, the quantity theory of money as I have applied it here—treating all liquid assets as though they were cash—will not be helpful in the analysis of such crises.

In the 1930s, it was runs on uninsured bank deposits that suffered a loss of confidence. In 2008, it was the "run on repo" (using the terminology Gary Gorton made familiar in his book *Slapped by the Invisible Hand: The Panic of 2007*). In both cases, these runs were in no sense induced by the Fed, but they both demanded a quick response. In the 1930s, the Fed failed to inject liquidity and, as described in Friedman and Schwartz's (1963) monograph, the results were very bad. In 2008, the Fed responded boldly and quickly, but even so this did not completely forestall large reductions in production in the two successive quarters that followed.

To make progress in understanding these events, we need to get into the details of the payments technology, of the mix of liquid assets—part cash but mostly promises of cash—that are used to pay the bills. We begin with some data on U.S. cash flows tabulated by the Fed.

The total annual payments—cash transfers over the year—recorded in Table 2 add up to $1,057,400 billion: roughly $1 quadrillion. Compare this figure with the dollar value of U.S. GDP in 2013, about $17,000 billion: $17 trillion. This means that $62 was transferred from someone to someone else for every dollar of final goods and services produced. Perhaps there are as much as $8 to $10 used in payments flows per dollar of GDP: Payments to factors of production need to be added and so do payments for intermediate goods. But this does not get us close to $62. It is clear that the payments system must cover much more than purchases of goods and services. The rest—most of it—must be settlements of asset exchanges. We can see this in the table.

Look at the differences in transaction sizes. Automated clearinghouse payments and checks processed total about $32 trillion. The average transaction sizes are around $2,000. These

involve payments by households and nonfinancial firms. In contrast, the average transaction size for Fedwire services, which total about $1,000 trillion per year, is about $15 million: Transactions of this size will be made mostly by financial firms, the only firms that routinely trade assets at anything like this volume.

Imagine trying to carry out trading at this volume using government-provided cash, paying no interest or nearly so, to settle accounts. The opportunity cost would be staggering. Financial firms have much larger incentives to economize on cash than do households and nonfinancial firms exactly because their payment rates are orders of magnitude larger. To deal with this problem, they adopt methods to clear among themselves using suitably designed derivative securities. "Suitably designed"? A useful instrument will have both a low trading cost (i.e., labor costs: financial and legal experts) and a solid backup plan if it falls in value— good collateral. There is nothing new about the financial sector's need to economize on cash or on the methods it uses to do so. Promises to pay have substituted for payments in cash since records were kept on cuneiform tablets.

There is nothing necessarily dishonest or fraudulent in this manner of trading, this economizing on cash. We all do it to some degree, and people managing cash flows that differ from ours by many orders of magnitude do it much, much better. They do this today by constructing short-term liquid assets from pieces of illiquid securities. Some of these securities were derived from mortgages. Many other securities could have been and, indeed, have been used.[1] As with any IOU, creditors accept it only when they think they can pass it on. What we mean by a "financial crisis" is an occasion when this confidence vanishes—for whatever reason. The effect of this situation is just like an old-fashioned bank run: Financial firms that were working with minimal amounts of cash are now cash hungry. But there is only so much cash out there, so this can only mean that cash available to households and nonfinancial firms has to shrink.

In quantity-theoretic terms, the money supply available for purchases of goods and services shrinks drastically, implying deflation or losses of production or both. In the 1930s, both of these occurred. In 2008-09, the dramatic, timely injection of new reserves led to a much better outcome. But we can hardly view this as a solution to the problem. In the two quarters following the failure of Lehman Brothers, production flows in the United States fell by something like 6 percent below trend. None of this has been recouped.

CONCLUDING REMARKS

The crisis of 2008-09 revealed weak points in some of our favorite models. Was this a failure of economics? Relative to what we would like to have known, of course. But we can also see enormous progress that would have pleased Homer Jones. The inflation of the 1970s— the largest peacetime inflation in U.S. history—was treated at the time with wage and price controls as the Fed watched inflation climb to 14 percent. The new idea of inflation targeting, focused exclusively on inflation control, has avoided attempts to affect prices by dictating to corporations and labor unions. It has provided us with a feasible monetary policy that has worked well right up to now in the United States and elsewhere.

On the other hand, crisis prevention—the elimination of runs—is still largely an unsolved problem. Any solution must involve a governmental guarantee of deposits, insurance of some kind, and this in turn necessitates regulations on the assets deposit-takers can hold. The 1933 Glass-Steagall Act, based on a mixture of deposit insurance and the regulation of designated commercial banks, was followed by 75 crisis-free years. The act was repealed in 1999, but by then the financial sector and large businesses had moved deposits out of commercial banks.

The Dodd-Frank Act acknowledged this reality and incorporated investment banks under regulation by the Fed. This response completely misses the point. The main reason commercial banks lost deposits in the 1970s and 1980s was the combination of Regulation Q, which prevented interest on deposits, and the onset of inflation. Look again at Figure 1. Why did anyone hold zero-interest deposits in 1980? What if Glass-Steagall had remained in force but with Regulation Q removed? When banks are regulated, there will always be incentives to create liquid assets outside the banking system, however broadly banks are defined. The long, successful history of Glass-Steagall shows that these incentives need not be fatal to regulation, as long as the costs remain reasonable. ■

NOTE

[1] See the long list of examples of securitized asset classes in Gorton (2010).

REFERENCES

Benati, Luca. "Long-Run Evidence on Money Growth and Inflation." *Bank of England Quarterly Bulletin*, Autumn 2005, pp. 349-355; http://www.bankofengland.co.uk/publications/Documents/quarterlybulletin/qb050302.pdf.

Benati, Luca. "Long Run Evidence on Money Growth and Inflation. Working Paper Series No. 1027, European Central Bank, March 2009; http://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1027.pdf.

Friedman, Milton and Schwartz, Anna J. *A Monetary History of the United States, 1867-1960*. Princeton, NJ: Princeton University Press, 1963.

Gorton, Gary. *Slapped by the Invisible Hand: The Panic of 2007*. New York: Oxford University Press, 2010.

McCandless, George and Weber, Warren E. "Some Monetary Facts." Federal Reserve Bank of Minneapolis *Quarterly Review*, Summer 1995, *19*(3), pp. 2-11; http://www.minneapolisfed.org/research/qr/QR1931.pdf.

O'Donoghue, Jim; Goulding, Louise and Allen, Grahame. "Consumer Price Inflation Since 1750." *Economic Trends*, March 2004, *604*, pp. 38-46;
http://inflationdata.com/inflation/economy/UK_Economic_Trends/uk_economic_trends_march_2004.pdf.