A Rule to Stabilize the Price Level

Kevin Dowd

The available evidence suggests that the costs of inflation are very large and the benefits are relatively small.[1] A strong case therefore exists that central banks should make the elimination of inflation and the subsequent maintenance of price-level stability their overriding objective.[2] However, the maintenance of price stability is not an easy technical matter, and if a central bank were to commit itself to price stability, it must have a clear idea of how it would intend to achieve that goal. The question therefore arises: What is the technically most efficient means to achieve price stability? Should the central bank implement a monetary target along conventional monetarist lines? Should it go back to some version of the gold standard? Or is there some other rule it should adopt?

This paper considers the main options facing a central bank and assesses their pros and cons. The most popular solution among modern economists is for the central bank to target the money supply and be left considerable discretion over how it should achieve that target. For a variety of reasons, however, it is an extremely difficult matter in practice to hit money supply targets with any reasonable degree of accuracy, and even successful money supply targeting can still fail to deliver the ultimate objective of a stable price level if velocity is unstable or the economy is hit by shocks or undergoes too much transformation. The alternative favored by older generations of economists was a gold standard. The gold standard is much easier to operate, but the price level under the gold standard is vulnerable to whatever

---

Cato Journal, Vol. 15, No. 1 (Spring/Summer 1995).

The author is Yorkshire Bank Professor of Financial Economics in the School of Financial Studies and Law at Sheffield Hallam University, Sheffield, England. He thanks Mike Belongia, Dave Wheelock, Bill Woolsey, and an anonymous referee for helpful comments.

[1]See, e.g., Cozier and Selody (1991), Dowd (1994b), Howitt (1990), and Selody (1990).

[2]For recent statements of this position, see Angell (1990), Gavin and Stockman (1988), Hoskins (1992), and Jordan (1992).

factors influence the supply and demand for gold, and it is therefore unlikely that the gold standard would deliver the degree of price stability we would want.

The problem with the gold standard has long been recognized, and a solution suggested long ago by Alfred Marshall, symmetallism, was to replace gold with a basket of precious metals. The presumption underlying the symmetallic standard was that the price level would be more stable than it would be under the gold standard because shocks to the individual relative prices of the included commodities would to some extent offset each other. Marshall had in mind a basket consisting of gold and silver only, but the basic principle can be extended to a basket covering many different goods. If individual relative price risks to some extent cancel out, then the larger the basket, the more stable we might expect its relative price, and hence the price level, to be. Taking this principle to its limit, our ideal basket would be a very broad one that covered virtually all the goods and services traded in the economy.

But the larger the basket, the greater the handling, storage costs and other costs involved with using it, and for many goods and services the costs are particularly high. The broad basket approach therefore must be ruled out as imposing unacceptably large storage and handling costs. There is, however, a way round the handling and storage cost problem: One could use a basket of future commodities instead of a basket of current (i.e., spot) ones. Because a commodity future is merely a promise to deliver a commodity (or its cash equivalent), commodity futures do not involve the storage and handling costs of spot commodities. Yet because the prices of spot and future commodities are closely linked, we can go a long way toward stabilizing a price index (i.e., the price of a spot basket) by adopting a symmetallic type of scheme based on a basket of commodity futures. My basic proposal is that the central bank should aim to stabilize the price level by pegging the price of an appropriately chosen basket of future commodities, and the best way to do that would be to peg the price of a price-index futures contract. If the central bank wishes to stabilize prices, this scheme offers it a simple and effective way to do so.

What Do We Mean by Price Stability?

Before discussing the central bank's choice of rules, we must first clarify what we mean by price stability and distinguish it from other related objectives. The term price stability has been used in the literature to refer to at least three different objectives. The first objective, and arguably the one to which the term price stability is least appropriate, is a *steady inflation rate* (i.e., an inflation rate that is roughly the same each year). The second is a *zero inflation rate*, and the third is a *steady level of prices*. The first is the least convincing. The benefits of a steady but positive inflation rate are highly questionable,[3] and there are serious drawbacks. There are good theoretical reasons to believe that even a perfectly anticipated inflation rate is damaging—individuals have to change prices more frequently, and changing prices is costly; inflation leads individuals to overeconomize on their holdings of real balances; and so on—and empirical estimates suggest that the losses involved are large (see, e.g., Driffil et al. 1990; Dowd 1994b). Even leaving aside those losses, there is also considerable evidence suggesting that higher rates of inflation are associated with a higher inflation variability or greater inflation uncertainty. That evidence suggests that one may not be able in practice to pick an inflation rate independently of the variability or uncertainty surrounding that rate, and that in turn suggests that the ideal of a steady but positive inflation rate may be a myth that cannot be realized in practice. There is also the problem that the choice of inflation rate is arbitrary and awkward to defend. If we pick an inflation rate of 5 percent, then why 5 percent and not 6 percent? If 6 percent, why not 7 percent? And so on.

If we reject the first objective, we still have to choose between the other two. Zero inflation and price stability are similar, but they differ in one important respect. If the price level for some reason turned out to rise this period, a price stability objective would require that the price level was subsequently brought down again to its initial level, but a zero inflation objective would "write off" the rise in prices and merely aim to stop prices rising any further in the future. While some economists dislike the idea of correcting price-level shocks, especially if that means that prices have subsequently to be forced down, I

[3]The most commonly cited benefit of a positive inflation rate is the "optimal inflation tax" argument due to Phelps (1973) that maintains that in a world where the government must resort to distortionary taxes, inflation should be positive because the inflation tax should be a component of that package. That argument has been heavily criticized—see, e.g., Dowd (1994b) and the literature cited there—but even if one were prepared to accept it, no one in practice actually knows what the implied optimal inflation rate actually is. Another argument sometimes made is that conventional price indexes are subject to a "quality change bias"—they tend to understate the impact of quality changes—which means that the reported annual inflation rate is typically 1 to 2 percent greater than the true quality-adjusted one. True price stability then requires not zero inflation but a positive inflation rate of 1 to 2 percent. However, the obvious reply to this argument is that government statistical agencies should alter the way price indexes are constructed and build appropriate quality-change allowances into them. But even if one accepts the argument at face value and opts for a positive inflation rate, the argument justifies only a very low inflation rate anyway, and it clearly cannot be used to defend inflation rates much beyond 2 percent.

would nonetheless argue that price stability is the better objective. An important reason we care about the price level (and therefore the inflation rate) is that we believe private individuals need a stable and predictable monetary environment in which to go about their business, and a critical feature of such an environment must be a predictable price level. If prices are subject to shocks of various kinds, then of course perfect predictability will be impossible, but prices will be more predictable under a stable prices objective than under a zero inflation one. Under the former objective, an individual can expect that the central bank will counteract any price-level shocks that occur, and to the extent that such shocks are counteracted, an individual wanting to predict the long-term price level need not worry about interim shocks. Under a zero inflation rule, however, any inflation shock will be built into the price level with no tendency to be counteracted, and the same individual would have to worry equally about all the interim shocks that occur between now and some future period. The price level therefore will be less predictable under the zero inflation rule, and it becomes relatively less predictable the further into the future individuals try to make their predictions (see Gavin and Stockman 1991). If we are concerned about price-level predictability, as I believe we must be, we should choose a stable price-level objective over a zero inflation one.[4]

Monetary Targeting

If we wish to stabilize prices, most modern economists would suggest that the central bank target the rate of growth of the money supply to grow at the same rate as real money demand. If the demand for real balances grows at 2 percent a year, for instance, the money supply also should grow at 2 percent a year, if prices are to remain approximately stable. The central bank would announce a target rate of growth (or growth range) for a particular monetary aggregate, and then manipulate interest rates or the monetary base to try to ensure that the chosen monetary aggregate grows at the desired rate. Monetary growth rules were adopted by many countries from the mid-

[4]The argument in the text is of course by no means exhaustive, but my focus of interest in this paper is with *how* to achieve a particular nominal target, rather than with the *precise* choice of target as such. The text does, however, ignore nominal income targets, and one particular such target, the productivity norm, has recently received considerable attention (see, e.g., Selgin 1990, 1992; Yeager 1992). Under the productivity norm, prices would move inversely with changes in productivity, and Selgin has argued forcefully that the productivity norm would be superior to price stability. I strongly disagree, but refer the interested reader to the references mentioned and to the exchange between Selgin and myself in the *Journal of Macroeconomics* (Dowd 1995b, Selgin 1995).

1970s onward, but their record is patchy, to say the least, and a number of countries moved away from them in the 1980s. In the United States, the adoption of monetary targeting was followed by a fall in inflation in the mid-1970s, but inflation subsequently rose again during the Carter Administration. It fell again after 1979, only to rise yet again in the Reagan years and then fall once more in the early 1990s to its current level of around 3 percent. Canada adopted monetary targeting in 1975, but there was no significant impact on inflation until around 1980, and the Bank of Canada gradually moved away from monetary targets in the years afterward. Britain adopted monetary targeting in 1976, but the choice of the wrong monetary aggregate played a large role in producing a very severe recession in the early 1980s, and monetary targets were subsequently deemphasized and eventually abandoned.

There are good reasons why monetary targeting was not more successful than it was, and they can be illustrated by means of Irving Fisher's equation of exchange in logarithmic form:

$$M + V = P + y, \qquad (1)$$

where M, V, P, and y are respectively the logarithms of the money supply, the velocity of circulation of money, the price level, and real output. If the central bank were to use a monetary target to stabilize prices, it would aim to manipulate M to hit a target for P. A number of problems are then apparent. To begin with, the appropriate definition of the money supply (i.e., the choice of empirical counterpart for M) is far from obvious: Economic theory provides only weak guidance about it, and there is no consensus among economists, some preferring $M1$, others $M2$ or something broader, or a divisia aggregate instead, and others a narrow aggregate such as the monetary base. There are also other problems. Unless the central bank picks the monetary base as the aggregate to be targeted, it does not control its aggregate directly, and it must instead try to influence it by manipulating the instruments at its disposal—the short-term interest rates it controls, or the monetary base. If the money supply function is volatile, or unstable, the central bank may be unable to hit its target within any reasonable margin of error, and the volatility of M will spill over and make P volatile as well. The successful implementation of monetary targeting also requires the demand for money—or, in terms of (1), the velocity of circulation V—to be reasonably stable and predictable. If it is not—and there is much evidence to doubt it is—the price level could still turn out to be a long way from its desired value even if the central bank managed to hit its monetary target reasonably accurately. If the demand for the monetary aggregate is unstable, it becomes difficult for the central bank to tell what values the instruments should take in response to changes in other variables (e.g., y), and it may miss its monetary targets by a considerable margin. Such problems are likely to be particularly acute in periods when the money supply process is subject to rapid technological and regulatory changes that make it very difficult for the central bank to gauge the impact of its policy actions. To make matters even more difficult, central bankers have to make their monetary policy decisions with little knowledge of how long their decisions take to work—the lags involved are variable and often difficult to predict—and with only a very imperfect idea of the state of the economy at the time their decisions are made. Not only do their decisions take time to become effective, but their information is imperfect and typically out of date by the time they get it. In short, monetary targeting is a very imperfect way to try to hit a price-level target.

A Gold Standard

While most modern economists have focused on monetary targets to achieve a desired inflation rate, economists in the past tended to think in terms of commodity price rules for central banks or other issuers of currency to follow. For them the most natural way to generate the price-level behavior one desired was to adopt a suitable commodity standard—a system of rules which imposed on the issuer(s) of currency the obligation to maintain a particular exchange rate between their currency and one or more commodities. If we think of the currency issuer as issuing $1 bills, a commodity standard requires him to ensure that the $1 bills always exchange for a given amount of a particular commodity or basket of commodities. Because the issuer would maintain a fixed exchange rate between the $1 bills and a particular quantity (or quantities) of one or more real goods, the nominal price of that good (or goods) would be fixed. Currency issued under those conditions is often said to be convertible.[5]

Commodity standards of one sort or another were the historical norm and prevailed from very early times until relatively recently. From time to time, a particular government might make the currency

[5]Convertible currency is to be distinguished from fiat currency, which is issued with no such guarantees and which trades in the market for goods and services at whatever exchange rate the market determines. While many modern economists and central banks like the freedom that a fiat currency gives its issuer, economists of older generations tended (in my opinion rightly) to view fiat currencies as harmful because they gave the issuers far too much license. It is no accident that the adoption of fiat currencies has been associated with so much inflation.

No Postage
Necessary
if Mailed
in the
United States

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 13690 WASHINGTON, D.C. 20024

Postage will be paid by addressee

the CATO JOURNAL

1000 Massachusetts Avenue NW
Washington, DC 20077-0172

inconvertible—the United States made the dollar inconvertible when the Civil War broke out in 1861, for example, and the pound sterling became inconvertible during the French Revolutionary War in 1797—but those lapses from convertibility were always perceived as temporary and convertibility was subsequently restored. Thus the U.S. dollar became convertible again in 1879 and the pound in 1821. However, in the 20th century, commodity standards came under attack from economists who believed that it must be possible to improve on them, and from governments and central banks that became increasingly impatient with the constraints they imposed on their ability to issue currency and manipulate interest rates as they pleased. Commodity standards gradually became discredited, the discipline they imposed on the issue of currency was undermined, and the last vestiges of the commodity standard were finally abandoned with the collapse of the Bretton Woods fixed exchange rate system in the early 1970s.

The most well-remembered commodity standard is the international gold standard that prevailed from the late 19th century. The gold standard was a system of rules that required the issuer of currency to maintain the exchange rate of currency against gold (i.e., it fixed the price of gold). If the demand for currency rose, individuals would go to the banks of issue and ask for gold (or, perhaps, some other asset) to be converted into the currency they desired to hold, and the banks would oblige them. If the demand for currency fell, they would go to banks of issue and ask for currency to be converted into gold or some other real asset, and the banks were obliged to satisfy their demands. The supply of currency thus accommodated itself to the demand for it, at prevailing prices, and an excess issue of currency could have only a very limited effect on prices because it would be returned to the banks of issue and retired from circulation.

The price level under the gold standard was then determined by the factors that determined the relative price of gold against goods and services in general. If rp_g is the log of the relative price of gold, the relative price can be decomposed into the nominal price p_g and the general price level P:

$$rp_g = p_g - P. \tag{2}$$

Under the rules of the gold standard, the nominal price of gold p_g is fixed, so the price level P moves inversely with the relative price of gold:

$$P = p_g - rp_g = \text{constant}—rp_g \tag{3}$$

Any factors that alter the relative price of gold therefore alter the price level in the other direction. If there was a shock to the supply of gold such as a major gold discovery, for example, there would now be too much gold to clear the gold market at prevailing prices, and the relative price of gold would have to fall. Since the nominal price of gold would be fixed by the rules of the gold standard, the relative price of gold could only fall and clear the gold market if other prices adjusted upward. The gold discovery would therefore lead to a general rise in prices. If there were some development that led to an increased relative price of gold—increased economic growth leading to an increased demand for gold, perhaps—the gold market could clear only if the prices of goods and services elsewhere fell, and we would get a general fall in prices. Under the gold standard, the relative price of gold can fall if and only if other prices rise, and vice versa.

The gold standard has a number of attractions relative to the fiat (i.e., inconvertible) monetary systems that replaced it. One important feature that it shares along with all convertible monetary systems but not with fiat ones is that it provides a mechanism to ensure an automatic retiring of excess currency from circulation. Under a fiat system, there is no procedure to return excess currency to a central bank that issues more currency than people are willing to hold at prevailing prices. Prices then have to rise to restore equilibrium—and that of course is part of the reason why monetary growth under fiat systems has produced so much inflation. The gold standard also had the attraction relative to fiat systems that it provided the issuer of currency with simple (and manageable!) rules to follow. In effect it merely required that an issuer stand ready to buy back or sell its currency on demand from the public.[6] There was no need for currency-issuing banks to try to estimate the public's demand functions for currency or other monetary aggregates so that they could determine the volume of currency or deposits they should issue. All they had to do was honor their legal commitments and the volumes of currency and deposits, and indeed, interest rates, were determined automatically by market forces.[7] The gold standard was thus very simple to operate.

Nevertheless the gold standard also had its drawbacks. One obvious drawback was that the price level under the gold standard was only

[6]I ignore here any complications resulting from the possibility of indirectly convertible systems in which the banks of issue peg the price of one commodity or asset, but use another to redeem their currency when requested to by the public. The indirect convertibility issue is an important one, but we can ignore it for present purposes and the reader can assume for convenience that banks of issue actually handle the commodity or commodity basket whose price(s) they are fixing. For more on indirect convertibility, see Dowd (1995a), Schnadt and Whittaker (1993) and Woolsey and Yeager (1991).

[7]Central banks such as the Bank of England did have some influence on market interest rates even during the period of the classical gold standard before World War I, but their influence was still limited by the constraint of having to maintain convertibility. As a result, their freedom of maneuver was far more limited than that of a modern central bank operating a fiat monetary system.

as stable as the relative price of gold. If the relative price of gold fell by 10 percent, the price level rose by 10 percent, and so on. If the relative price of gold was unstable, then so too was the price level, and the historical record of the gold standard suggests that this concern about price-level instability is not unfounded. The price level under the historical gold standard was rarely stable from one year to the next, but price changes were seldom as dramatic as those that have characterized the major economies since the collapse of the Bretton Woods system, and over the longer run there was a tendency for price-level movements to reverse themselves (see, e.g., Bordo 1992). The U.S. price level in 1914 was not much different from what it had been a century earlier. This price-level record is better than that of the fiat monetary system that replaced the gold standard, but it is still unsatisfactory if one's objective is a price level that is as stable as possible. In any case, even if one were prepared to accept the short-run price-level instability that characterized the period of the historical gold standard, there is no guarantee if the gold standard were restored that general price movements would continue to cancel out over the longer term as they did before. The gold standard may or may not deliver a more predictable price level than a continuation of the present fiat system, but it is still the case that one of the main arguments against a restoration of the gold standard is the price-level uncertainty that would arise because no one can predict the future relative price of gold with any high degree of confidence.

There is also another drawback. A gold standard involves certain resource costs—the costs of handling, and perhaps mining, the gold that is used in the operation of the gold standard. In the historical gold standard, people used gold coins, and gold coins were a more costly medium of exchange to produce and maintain than pieces of paper or bank deposits. It was also usually the case that banks would keep gold so that they could hand it out over the counter if faced with demands for redemption (i.e., conversion) by the public. Those resource costs gave rise to considerable concern, and their existence led many economists to argue that commodity standards were less efficient than fiat systems because the latter involved no comparable costs and could (it was claimed) deliver otherwise similar outcomes.[8]

[8]That argument is questionable because it ignores the benefits convertibility can bring or implicitly assumes that those benefits can be obtained at no cost by throwing away the convertibility guarantee and pretending that the fiat system can be expected to produce otherwise comparable outcomes. The problem is that a fiat system has a fundamentally different character—the issuer of currency faces very different constraints, and can therefore be expected to behave in a very different manner.

There are reasons to believe the claims are exaggerated,[9] but there is no denying that the resource cost argument was a major reason many economists opposed the gold standard in particular and commodity standards more generally.

Commodity-Basket Monetary Systems

The drawbacks of a gold standard led economists to suggest a new type of commodity standard based on baskets of commodities instead of single commodities only. Concern with the price-level instability it created led Alfred Marshall in the 1880s to suggest a symmetallic system, one in which the value of the currency was tied, not to a particular weight of gold, but to a basket consisting of particular weights of gold and silver (see Marshall [1887] 1926). If p_b is the (log of the) nominal price of the basket, and rp_b (the log of) its price relative to the prices of goods and services generally, the currency issuer would peg p_b in just the same way as it would peg p_g under the gold standard. The (logarithmic) price level would then be given by

$$(4) \qquad P = p_b - rp_b = \text{constant} - rp_b.$$

The distinctive feature of symmetallism is that p_b would be a weighted average of the nominal prices of gold and silver. Thus

$$(5) \qquad p_b = a_g p_g + a_s p_s,$$

where a_g and a_s are given positive weights summing to 1.[10] The relative price of the basket is then

$$(6) \qquad rp_b = p_b - P = a_g p_g + a_s p_s - P = a_g rp_g + a_s rp_s.$$

[9]Where banks were left relatively free under the gold standard, as for example in early 19th century Scotland, the evidence suggests that they could operate on remarkably low reserve ratios (see, e.g., Cameron 1967, 87–88), and those low reserve ratios implied that resource costs were also very low. In any case, economists who have pushed the resource cost argument against convertible currencies have generally overlooked the point that fiat systems have very significant costs of their own. Inflation and uncertainty lead to serious misallocations of resources—markets malfunction, agents feel obliged to invest resources to protect themselves from the consequences of inflation uncertainty, as illustrated by the growth of the financial services industry in the 1970s and 1980s, and so on. As Milton Friedman (1986) points out, those resource costs are highly significant, and my own assessment of the costs (Dowd 1994b) leads me to believe that the resource costs of a convertible monetary system are negligible in comparison.

[10]The text uses a geometric average for expositional purposes, but in the original symmetallist scheme, as proposed by Marshall, the symmetallist basket was to be an arithmetic weighted sum of commodities, so its price was an arithmetic rather than a geometric mean of the prices of the included commodities. The choice of arithmetic or geometric mean is fundamentally one of convenience only. Marshall also envisaged that the weights a_g and a_s would be fixed, but those could in principle be allowed to vary in accordance with rules of their own provided they still sum to one. To anticipate the later discussion, we might want to allow them to vary if the basket was meant to represent a typical consumer's shopping basket and that basket were to change over time.

The symmetallic price level is then

$$(7) \qquad P = \text{constant} - [\, a_g\, rp_g + a_s\, rp_s \,].$$

In comparing (3) to (7), it becomes apparent that the relative price of gold enters the gold-standard price-level equation (i.e., (3)) with a weight of 1, but enters the symmetallic price-level equation (i.e., (7)) with a weight less than 1 (i.e., a_g). Changes in the relative price of gold therefore have a smaller impact on the price level under symmetallism than under the gold standard. The drawback, however, is that symmetallic price level is now also vulnerable to changes in the relative price of silver, whereas the gold-standard price level was not. On the other hand, if changes to the relative prices of gold and silver are not perfectly correlated, then shocks to those relative prices will to some extent offset each other. We might then expect, other things being equal, that a symmetallic system should produce a more stable price level than a gold (or silver) standard. The idea behind symmetallism is thus to make the price level more stable by exploiting any tendency of changes in the relative prices of individual goods to offset each other.

That basic principle can be extended further. If a system based on a basket with two metals produces a more stable price level than a system based on one metal only, we might expect a system based on a basket of three or more metals to produce a more stable price level still, as shocks to the relative prices of the individual commodities included in the basket tend to cancel out even more. Nor is there any reason to restrict the commodities included in the basket to metals. Nonmetallic commodities might also be included, allowing a much larger number of commodities to be included than would otherwise be possible, and in the process allowing one to choose commodities whose relative prices tended to move against each other and produce a more stable relative price for the basket.[11] If we pick a basket with n commodities, the price level is then given by

$$(8) \qquad P = \text{constant} - [\, a_1\, rp_1 + a_2\, rp_2 + \ldots + a_n\, rp_n \,],$$

where the subscripts refer to each of the included commodities and the weights a_i again sum to one. Ideally, we can then produce a stable price level by selecting a broad enough basket (i.e., choosing a high value for n) and including the appropriate commodities in it.

[11]One of the problems with a purely metallic basket is that the shocks to the relative prices of the individual metals will tend to be highly correlated, so the relative price of a metallic basket will be less stable than the relative price of many a basket that included nonmetallic commodities, even if each basket had the same number of included commodities.

There is consequently a good argument that basket-based commodity monetary systems will produce a more stable price level than monetary systems based on single commodities. There is however the problem that it appears to be very difficult to identify a suitable basket in practice. One possibility is Robert Hall's (1982) ANCAP basket—a basket of specified amounts of ammonium nitrate, copper, aluminum, and plywood—and Hall presented evidence to suggest that the purchasing power (i.e., relative price) of this basket in the United States had been stable over the period from 1950 to 1980. Yet it is doubtful that the ANCAP basket continued to have a stable relative price after 1980, and there is in any case no particular reason to have expected it to. Alternatives to the ANCAP have also proved remarkably difficult to find. Karl Brunner (1984), Alan Garner (1985), and James Boughton and William Branson (1989) all searched for but failed to find an equilibrium relationship between commodity price baskets and the U.S. price level, and my own attempts in 1990 to find a relationship between the prices of 11 internationally traded commodities and the (Australian) price level turned out to be a dismal failure. If a suitable basket really exists, no one can yet claim to have found it, and the commodity-basket scheme obviously cannot be operational without a tangible basket to work with.

There is also another serious problem: Commodity-basket systems involve higher resource costs, possibly much higher ones. It is one thing to keep gold coins on hand and pass them over the counter when faced with demands to redeem currency, but it is quite another to keep a basket of goods, particularly a broad one, and hand that over instead. The greater the number of commodities included in the basket, the greater we would normally expect the handling, storage, and other costs involved to be, and the more expensive the monetary system would be to operate. The problem is that each commodity has as it were its own storage and handling needs, and those needs impose certain costs that must be borne if the commodity is to be included in the basket. If a precious metal were included, we would normally need a strongroom of some sort; if a frozen commodity were included, we would need refrigeration facilities; and so on. In some cases, one could perhaps add a new commodity that could be handled with minimal additional costs—if one had the facilities to store gold, one might perhaps be able to add silver at no significant extra cost—but if they existed at all, such cases would be the exception rather than the rule. As we broadened our basket and included more commodities in it, we should therefore expect that the costs of handling a basket of a given nominal value will generally increase. In any case, as the basket gets broader, one will at some point have to deal with commodities that depreciate over time (e.g., perishable goods), or are otherwise relatively costly to store (e.g., because they are bulky), or do not come in easily recognizable or manageable lumps. To some extent, perhaps, we could ameliorate the problems caused by having indivisible goods in the basket by dealing only in high-value baskets—having rules that require the issuers of currency to redeem their currency in units of $1,000, say, to prevent them from being forced to hand over a little basket of goods when someone comes in with a $1 bill—but dealing only in high-value baskets does not fully solve the indivisibility issue and does nothing about some of the other problems (e.g., perishability).

There is, moreover, a deeper aspect to this general problem. The physical characteristics of some commodities mean that the resource costs of including them in the basket are very high (e.g., for highly perishable commodities). Pursuing this line of thought further, many commodities involve resource costs so high that those costs are effectively unlimited. There are commodities whose physical characteristics mean that it is impossible (or almost impossible) to include them in a basket (e.g., nuclear power stations), but the same also applies to many services. It is possible, maybe, to have a basket that included hair-styling of a particular kind, but it would not be easy for a central bank that operated a commodity monetary system based on such a basket to physically deliver the particular services involved. It might keep a hair salon on the premises, perhaps, and employ stylists to stand ready to deliver the service included in the basket, but the same logic that says it should keep a hair salon also suggests that it should keep a bakery, a bar, a public recreation center, and who knows what else. It does not take much imagination to see how high those resource costs might be, and one cannot reasonably to expect the central bank to be able to deliver commodities or services like these. We therefore have little choice in practice but to exclude a large number of goods and services from our basket if we wish to keep our resource costs down to tolerable levels.

The Conflict Between Price-Level Stability and Low Resource Costs

It might help to summarize where we stand. If we want price-level stability from a commodity monetary system, we have good reason to prefer a system based on a basket of some sort, if a suitable one can be found. We also have reason to believe that, other things being equal, a system based on a broader basket would deliver a more stable price level than a system based on a narrower one. If we ignore any resource-cost issues for the moment, the best basket—the one that delivers the most stable price level—is clearly the same basket as the basket whose price is represented by the price index we wish to stabilize. If we wish to stabilize a particular price index, and this index is itself a weighted average consisting of a_1 times the price of good 1, plus a_2 times the price of good 2, and so on, we can think of the price index as representing the price of a basket consisting of good 1, good 2, and so on, all in particular quantities implied by the way the index is constructed. If we wish to stabilize the consumer price index (CPI), for example, we would select the CPI basket, the basket whose price the CPI represents. The relative price of the basket against goods and services in general, as proxied by the price index itself, is consequently 1 (i.e., constant). The price level is

$$P = p_b - \text{constant} \tag{9}$$

and is as stable as the nominal price of the basket itself. The price level can then be stabilized by pegging the nominal price of the anchor.

The problem, of course, is that we have left out the resource costs. Resource costs tend to rise as the basket gets broader, and the resource costs of a CPI basket are effectively unlimited. A broader basket means a more stable relative price, and hence a more stable price level, but it also means higher resource costs, all on an other-things-being-equal basis. We thus face a dilemma—we can choose a broad basket with good price-level stability properties and very high resource costs, or a narrow basket with tolerable resource costs that delivers a poor degree of price level stability. From the point of view of CPI stability, our presumed objective, our ideal basket, the CPI basket, involves resource costs so high that it must be dismissed as unfeasible. To add to our difficulties, if we rule out the CPI basket on resource-cost grounds, it is not clear what other basket we can turn to. No one has yet identified a suitable flesh-and-blood basket that might involve tolerable resource costs. We appear to have gone up a blind alley.

A Commodity Futures Basket?

A possible way out is to think in terms of future commodities (i.e., commodity futures) rather than present (i.e., spot) ones. A commodity in the future is in many ways very similar to a commodity now—the difference is that one has not yet taken delivery of it, so the difference between a futures contract and a spot one is that the former calls for future delivery (or its equivalent) and the latter calls for delivery now—but for our purposes what matters is that their prices will still move very closely with the price of the corresponding spot commodity.

If the price of spot commodities rises, we would expect that prices of the same commodity to be delivered at some time in the future would normally rise as well. The nearer the delivery date, the closer we would expect the correspondence between the two prices to be. If one rises, arbitrage will normally ensure that the other rises as well.

Because these contracts are only promises to deliver, the owner of such a contract does not actually have to handle the commodity concerned, at least until the maturity date arrives, and sometimes not even then.[12] The trading of the contracts therefore involves no resource costs to speak of, and the costs are negligible regardless of the costs of handling the corresponding spot commodities. Agents can hold whatever forward or futures contracts they want and still bear negligible resource costs provided that they do not find themselves forced to take delivery of the commodities concerned when the contracts mature. If the prices of forward or futures contracts move closely with the prices of the corresponding spot contracts, and if handling those contracts involves negligible resource costs, there is a simple solution to our earlier dilemma: we simply replace the spot goods and services in our basket with their futures equivalents.[13]

A futures contract is an agreement that *A* should deliver a particular amount of some commodity to *B*, or make an equivalent payment to him, at a specific future time *T*. In a simple futures contract, the price is agreed now, and some proportion of the price is paid now, and the rest paid on maturity (i.e., at *T*). This proportion is the margin requirement. This margin requirement turns out to be important to ensure that the money supply is appropriately corrected, and it is easiest if we simply assume that the margin requirement is 100 percent. *B* therefore pays all the price now, when the contract is made.[14] If the contract calls for physical delivery of the commodity, *A* delivers the commodity at *T*, but if the contract called for cash settlement, *A* instead pays him whatever is needed to buy the specified amount of the commodity concerned on the prevailing spot market. Since *A* commits to the future physical delivery or equivalent cash payment, *A* is said to sell the futures contract to *B*, and the payment *B* makes to *A* in return is the price of the contract.

[12]Because resource costs as I have defined them only arise when one actually has the physical good, one can still avoid them when the contract matures if the contract calls for settlement in cash of equivalent value.

[13]The idea of using futures prices to help deal with inflation is not new, but it is only very recently that schemes have been suggested to use them to stabilize the price level. For more on this type of scheme, see Sumner (1989) and Dowd (1994a).

[14]I thank Bill Woolsey for very helpful correspondence on this point.

To see how the price of a typical contract is determined, suppose that both agents have the same expectation of the spot price of the commodity at T. Suppose also that both agents have the same attitude toward risk—if one likes, one can think of both agents as being risk-neutral—and there is no significant default risk on the contract. If the price of the contract is too high, in a sense to be explained shortly, a typical agent will prefer to sell the commodity now in the futures market than later on in the spot market. He will judge that he will be better off selling now and investing the proceeds at the going interest rate than he would be waiting until T arrives and selling then. There will then be an excess supply of futures contracts relative to (expected) spot ones, and the price of futures contracts will fall relative to the (expected) future spot price. If the futures price was less than the expected spot price, the same agents would expect to be better off waiting until T arrives and then selling than they would be selling now on the futures market and investing the proceeds until T. There will now be an excess demand for futures contracts relative to (expected) future spot ones, and the price of futures contracts will rise relative to the expected future spot price. The price of futures contracts would therefore keep changing until it was equal to the discounted expected spot price of the commodity concerned. If F is the (log) price of the futures contract, and S the (log) spot price, arbitrage would ensure that

$$F = f(.) + ES, \tag{10}$$

where $f(.)$ is a stable function that makes allowance for the interest cost due to some of the payment taking place when the contract is made,[15] and E is an expectations operator indicating that (10) relates F to the expectation of S. If the term to maturity of the contract is short, then $f(.)$ should be close to zero and

$$F \sim ES. \tag{11}$$

The story has to be modified slightly if agents on either side of the market differ in their attitudes toward risk, but (11) holds even in this case, and the key point is that the price of the futures contract is still closely tied to the expected spot price of the commodity on the maturity date. One price cannot move far without dragging the other with it.

[15]Ignoring risk and related factors, if r is the rate of interest over the term of the contract and θ is the margin requirement, then $f(\ ) \sim \ln[\theta/(1+r)]$. If the term to maturity is short, then r should be small, and $f(\ ) \sim \ln \theta$. It follows, then, that while $f(\ )$ does depend on r, we can normally ignore that dependence if r is small (as it would be, in particular, if the contract has a short term to maturity).

The basic idea behind the proposal is that we should aim to stabilize the price level by exploiting the fact that futures prices move very closely with expected spot prices, and therefore with *ex post* spot prices as well, but handling futures contracts does not involve the resource costs incurred in the handling of many spot commodities and services. One way to implement the idea would be to deal in futures contracts for all the commodities whose prices appear in the price index we wish to stabilize, but doing so would still involve considerable transactions costs—we would have to deal with a large number of different contracts—and there is the additional problem that futures contracts for most of those goods and services do not currently exist anyway. We would therefore have to set up the appropriate markets so that we could deal with the contracts traded in them, and then incur all the transactions costs involved in handling them.

There is, however, a much easier alternative. Instead of trading in a large number of futures contracts for individual commodities, we could trade in a single *price-index futures contract.* A price-index futures contract is a contract to make a payment contingent on the realized value of the appropriate price index at the time stated in the contract. If the earlier futures contracts were contracts to deliver particular individual commodities or make equivalent cash payments, we can think of a price-index futures contract as a contract to deliver the monetary equivalent of the bundle of goods and services whose price the price index represents. If the price index is the CPI, the central bank would therefore trade in CPI futures contracts. The proposal is then very straightforward: *the central bank should be committed to pegging the price of CPI futures contracts of a given relatively short maturity. Since spot prices and futures prices are closely linked, pegging the price of CPI futures would tend to stabilize the spot price of the CPI.* In notation used in the previous two equations, the spot price S is now the CPI (i.e., P), and the futures price F is now fixed. The expected price index is therefore

$$S = EP \approx \text{constant}. \tag{12}$$

The actual price index P is then given by the expectation of P and a random error ϵ that picks up the unanticipated movement in P that was not foreseen when the futures contract was made:

$$P \approx \text{constant} + \epsilon. \tag{13}$$

The actual price level is approximately constant, give or take a random error ϵ. If agents are rational, the price-forecast error should be unpredictable, so ϵ should have a mean of zero. The variance of the price level (i.e., var P) would then be

$$\text{var } P \approx \text{var } \epsilon, \tag{14}$$

that is, the variance of the price level would be approximately equal to the variance of the price-level forecast error.

There are two points to note about the behavior of the price level under the proposed rule. First, the price level should have a constant mean, as given by (13). There should be no trend rate of change in the price level over time, and the mean rate of inflation should be roughly zero. Second, the variance of the price level would be approximately equal to the variance of price-level forecast errors. We cannot predict this variance exactly, but one can form some idea of what it might be from price-level forecast errors from the recent past. To give an illustration, Dennis Jansen (1989, Table 4) presents results for the United States for the period 1959:I to 1988:II that indicate that price-level forecast errors had a variance of about 1.40 times 10^{-5} (i.e., 0.000014), or a standard deviation of about 0.37 percent. If we assume that the price level would be just as predictable under the new rule as it has been in the recent past, then (13) suggests that the price level itself would have a standard deviation in the region 0.35–0.4 percent. However, the absence of monetary policy shocks under the new rule should actually make the price level more predictable under the new regime, so the standard deviation should be somewhat smaller than that range suggests. In short, adopting the proposed rule should lead to a price level that has an approximately constant mean and a standard deviation (to be conservative) of under 0.4 percent.[16]

How the Scheme Would Operate in Practice

To see how the scheme would actually work in more detail, suppose that agents for some reason expected the future value of the price index to rise. If the expected future value of the CPI was too high relative to the (given) value of the CPI futures contract, arbitragers would wish to buy futures contracts now and make provision to sell the commodities spot when the maturity date arrived. The central bank would have to create the additional futures contracts the private agents demanded and sell the contracts to them at the price the central bank was maintaining. Because the purchasers would have to

[16]Apart from trying to tighten up the approximation error in (12) and get a more precise estimate of the variance of the price-level forecast error, it is not clear how much more one can do to predict the behavior of the price level under the proposed rule. The price level forecast shocks ϵ are by their nature unpredictable, and it is not clear (at least to me) how one could confidently simulate them. One possibility, perhaps, would be to use CPI futures data from the U.S. CPI futures market that operated between 1985 and (in effect) 1988, but the market was so thin and data set so short that one must wonder what such an exercise would really tell us. For more on the U.S. CPI futures contract, see the text below and Horrigan (1987).

hand over money at the time they agreed on the contracts—remember the margin requirement[17]—the central bank would be buying back base money in the process and retiring it from circulation. The money supply would fall, and the falling money supply would help reduce the expected value of the future CPI toward its equilibrium value. This equilibrating process on the part of arbitragers would also be reinforced by the behavior of primary suppliers and purchasers. If the future expected spot price was too high relative to the currently prevailing futures price, a typical supplier of goods would expect to make higher profits from selling the good in the future, and so would make plans to increase the supply to be brought to the market. The supplier also might hold back supplies he was planning to sell in the current period to sell them later instead. Similarly, those who purchase goods will anticipate higher future prices and demand more now with a view to demanding less in the future. Both demanders and suppliers will therefore respond in ways that increase excess supply in the future and thus help to bring down expected future prices, and they will continue to alter their plans until expected future prices have been brought down to levels compatible with the pegged price of the CPI futures contract.[18]

The commitment of the central bank to peg the price of CPI futures contracts also would ensure that the new contract found the market it needed to make the central bank's rule deliver the desired outcome. As briefly mentioned in note 16, a CPI futures contract was traded between 1985 and 1988, but trading was so thin that the contract was eventually withdrawn by the exchange that sponsored it, the Coffee, Cocoa, and Sugar Exchange in New York. The reasons for the failure of CPI futures to catch on are not entirely clear. Brian Horrigan (1987) suggests that prices might have been relatively predictable over that period, and so there would have been little incentive to take out the protection that the contract offered. He also suggested that

[17]The margin requirement is crucial. If there is no margin requirement, no base money is retired at the time the contract is made, and the equilibrating mechanism described in the text does not come into play: Significant margin requirements are necessary to ensure that base money is retired or increased at the time that contracts are made. Again, I thank Bill Woolsey for his correspondence on this point.

[18]One objection made to this scheme is that if prices are sticky, it might take time for the price index to adjust, and delays in the response of the CPI might produce destabilizing swings in the money supply and, perhaps, real output. However, the equilibrating process captured by (10) and (11) *requires that it be expected prices that adjust, rather than spot* ones, and there is no reason to expect expected prices to be as sticky as spot prices might be. Provided the term to maturity is not too short, many of the components of the expected CPI will adjust even if their counterparts are sticky, and destabilizing movements in money supply and associated effects should by and large be avoided.

the failure of the CPI futures contract might have had something to do with the absence of an underlying asset, an absence that would have discouraged some traders and therefore kept the market thinner than it might otherwise have been.

I believe the contract proposed here would find the market it needs despite the earlier failure, because of the central bank's obligation to peg the price of the new contracts. (The central bank of course had no such obligation with the previous CPI futures contracts.) The argument is as follows: Suppose the central bank adopted the new policy and there were no takers for the contracts it was offering to buy and sell. The equilibrium condition (10) then tells us that the money supply, interest rates, and so on, must be roughly compatible with the desired outcome (i.e., stable prices). We would therefore have nothing to worry about. If the money supply and so on were to deviate from the levels compatible with the equilibrium conditions, on the other hand, then those conditions would no longer hold, and it would be worthwhile for some private agents to buy or sell contracts at the fixed prices the central bank was offering. The market would then come alive, and it would remain alive for as long as it needed to in order to bring about the changes to base money and so on required to restore equilibrium. In short, the equilibrium condition (10) and the fixed price offered by the central bank would ensure that private agents had the incentive to make the trades necessary to produce the desired outcome.

The commitment of the central bank to peg the price of CPI futures contracts also prevents it from increasing the supply of base money beyond the level compatible with the pegged price of those futures contracts, in much the same way that the commitment to maintain the price of gold prevented banks of issue under the gold standard from issuing excess currency. Like banks of issue under the gold standard, the central bank would effectively commit itself to supply base money on demand. If the public want more base money, they would come to the central bank to sell it futures contracts, and the central bank would buy the contracts from them and give them the additional base money they want. If the public want less base money, they would come to the central bank wanting to exchange base money for futures contracts, and the central bank would oblige them. The central bank's operating rule would be to maintain the price of futures contracts, and it would create or retire base money on demand to ensure that the price of those contracts remained at the level it was supposed to. The commitment to peg the price of futures contracts prevents the central bank from issuing too much base money. Like the gold standard rule, the CPI futures rule would be a very simple one to operate, and the central bank would no longer have to worry about the most appropriate definition of the money supply or whether money demand had shifted.

The rule proposed here could also deal easily and automatically with shocks or other exogenous changes in the economy. Suppose, for example, that there was another oil price hike, as in 1974 or 1979. If non-oil prices remained the same or rose, and that state of affairs was expected to continue, then the expected value of the future CPI would rise and the equilibrating forces discussed earlier would come into play. Agents would demand additional contracts from the central bank, expected future excess supply would rise, and the expected future value of the CPI would fall to a level consistent with the price of the CPI futures contract as pegged by the central bank. If oil prices have risen, however, the value of the CPI can only fall back again if *other prices fall, so non-oil prices would on average have fallen.* What happens is that relative prices adjust to the oil price shock, but they adjust in such a way that the average price level as measured by the CPI is restored pretty much to its initial level.[19]

The monetary system responds in the same way to secular or cyclical changes in the demand for real balances. Should the demand for real balances rise because of economic growth, for instance, we would expect there to be a higher demand for base money as well. Agents who wanted additional base money from the central bank could then create CPI futures contracts to sell for the central bank in exchange for base money, and the money supply would rise to a level consistent with the new demand for real balances at prevailing prices. If income fell and the demand for real balances fell with it, on the other hand, there would be an excess supply of base money at prevailing prices, agents would retire the excess base money in return for futures contracts, and money supply would fall until it was consistent with the

[19]Some economists would question whether it is desirable that the CPI should return to its previous value—for example, one often hears the argument that shocks such as oil shocks should be accommodated. This type of argument is usually based on assertions that certain prices or wages are sticky downwards, and that this stickiness is sufficiently serious that monetary policy should accommodate it. I do not find these assertions convincing. Accommodating monetary policy merely allows the price level to rise, and it is doubtful that countries that pursued accommodating policies in the face of the first oil price shock in 1974, such as the United Kingdom and the United States, fared better than countries that made more of an effort to absorb the shock, such as Japan. I also find it odd that the depression of the 1930s, in which prices in the United States fell by around a third, should have led so many economists to believe that prices are sticky downward. In any case, I would argue that either we want price-level stability or we do not. If we do want it, and certain prices rise, then other prices must necessarily fall, and we can do nothing about it. If we want other prices to rise as well, or we want to resist their falling, we no longer want price-level stability.

lower demand to hold money balances. Whether the demand for money rises or falls, or whether those changes are secular ones or cyclical ones, the supply of money would respond in much the same automatic way to meet the new level of demand at much the same level of prices as that which prevailed before.[20]

Another important point to note is that the rule proposed here is not open to successful speculative attack provided that the central bank commits itself effectively to it. Price-fixing rules can only be open to successful speculative attack if private agents expect the pegged price—in this case, the price of the futures contract on the days when the central bank intervenes in the futures market—to change. If the central bank's commitment were credible, rational private agents would not expect this price to change, and there would be no point speculating against the central bank's price-pegging rule. It would only be worthwhile speculating against it if there were a significant probability the central bank would abandon it—which would not be the case if the commitment were credible, as it should be—or if the rule called for periodic changes in the pegged price, as in a crawling-peg exchange rate system—which the rule does not. There is thus every reason to believe that the rule proposed here would be safe against speculative attack.

There remain only the operational details of how the scheme would actually be implemented.[21] The central bank would not be required to intervene each day to peg the price of futures contracts, as it intervened each day under the gold standard to peg the price of gold. If the CPI is announced once a month, it would suffice for the central bank to intervene in the futures market one day a month to peg the prices of contracts of the chosen maturity.[22] Arbitrage forces could then be relied on to ensure that the prices of futures contracts on other days were roughly consistent with the prices of those contracts on the central bank's intervention days. Creating the market for these CPI futures would be very simple: All the central bank would need to do is decide on the price at which it intended to buy and sell those

[20]One may ask why prices remain much the same. The futures price rule ensures that the expected price level is (approximately) constant. By (12), the actual price level remains much the same, give or take the error term ϵ. If a change in money demand or relative prices has any effect on the price level, that effect must be a relatively minor one.

[21]For more on these details, see also Dowd (1994a).

[22]The central bank need not intervene the same day as the CPI is announced, but it should intervene the same day each month. The difference between the futures and expected future spot prices involves an interest component, and if the central bank intervenes on different days in different months, the interest component varies and moves the expected spot price around unnecessarily.

contracts, and it would then leave it to private agents to come and trade with it. If there was not enough money in circulation to produce a price level consistent with the pegged futures contract price, private agents would find it worthwhile to come and obtain additional base money from it by selling it futures contracts, the money supply would rise, and so on. If there was too much money in circulation, private agents would present it for redemption by demanding futures contracts from the central bank, and so on. The central bank would have only to choose the days when it wanted to intervene to peg the futures contract prices, and the price at which it wanted to peg them. The latter, in principle, is arbitrary, but if the central bank were concerned about price stability, as it surely should be, it would presumably pick a value consistent with the value of the price level at the time the new system was introduced. The price level would in effect be tethered to a value close to its previous prevailing level, and there should no significant transitional jump in prices as the new regime was introduced.

Conclusion

Central bankers have often claimed that they would like to eliminate inflation, but lack a technically efficient means of doing so. Monetary targeting is very imperfect for all sorts of reasons. Traditional commodity standards such as the gold standard are much easier to operate, but the price level still fluctuates with changes in the relative price of gold. Economists have long suggested that a way round the latter problem would be to peg the price of a basket of commodities rather than the price of a single individual commodity, and the broader the basket, the more stable the price level would be. However, broad baskets also involve high resource costs of one kind or another, and the basket we would really like to stabilize—the basket of goods and services behind the CPI—involves such massive resource costs that we have no alternative but to dismiss it as unfeasible.

There is however a very simple solution. Instead of stabilizing the spot price of the CPI, the central bank should stabilize the price of CPI futures. The beauty of futures contracts is that their prices are closely linked to the prices of the corresponding spot commodities, but handling futures contracts involves no resource costs worth speaking of even if the resource costs of handling the corresponding spot commodities or services are high. A central bank can handle such contracts even when resource costs prevent it from handling the spot commodities or services underlying them.

References

Angell, W. (1990) "The Case for Price Stability." Mimeo. Board of Governors of the Federal Reserve System. Washington, D.C.

Bordo, M.D. (1992) "The Gold Standard, Bretton Woods, and Other Monetary Regimes: A Historical Appraisal." Paper presented to the 17th annual policy conference, Dimensions of Monetary Policy: Essays in Honor of Anatol B. Balbach. Federal Reserve Bank of St. Louis. October.

Boughton, J.M., and Branson, W.H. (1989) "Commodity Prices as a Leading Indicator of Inflation." In K. Lahiri and G. Moore (eds.) *Leading Economic Indicators*. New York: Cambridge University Press.

Brunner, K. (1984) "From the 'Upper Tail Theory of Inflation' to the 'Lower Tail Theory of Deflation'." Mimeo. Shadow Open Market Committee. New York.

Cameron, R. (1967) "Scotland, 1750–1845." In R. Cameron et al., *Banking in the Early Stages of Industrialization: A Study in Comparative Economic History*, 60–99. New York: Oxford University Press.

Cozier, B., and Selody, J. (1991) "Inflation and Macroeconomic Performance: Some Cross-Country Evidence." Mimeo. Bank of Canada.

Dowd, K. (1994a) "A Proposal to Eliminate Inflation." *Economic Journal* 104(425): 828–40.

Dowd, K. (1994b) "The Costs of Inflation and Disinflation." *Cato Journal* 14(2): 305–31.

Dowd, K. (1994c) "Deflating the Productivity Norm." *Journal of Macroeconomics* 17(4): 717–32.

Dowd, K. (1995a) "The Mechanics of Indirect Convertibility." *Journal of Money, Credit, and Banking* 27(1): 67–88.

Driffil, J.; Mizon, G.E.; and Ulph, A. (1990) "Costs of Inflation." In B.F. Friedman and F.H. Hahn (eds.) Handbook of Monetary Economics 2: 1013–66. Amsterdam: North Holland.

Friedman, M. (1986) "The Resource Cost of Irredeemable Paper Money." *Journal of Political Economy* 94(3): 642–47.

Garner, C.A. (1985) "Commodity Prices and Monetary Policy Reform." Federal Reserve Bank of Kansas City *Economic Review* 70 (February): 7–21.

Gavin, W.T., and Stockman, A.C. (1988) "The Case for Zero Inflation." Federal Reserve Bank of Cleveland *Economic Commentary*. 15 September.

Gavin, W.T., and Stockman, A.C. (1991) "Why a Rule for Stable Prices May Dominate a Rule for Zero Inflation." Federal Reserve Bank of Cleveland *Economic Review* 27(1): 2–8.

Hall, R.E. (1982) "Explorations in the Gold Standard and Related Policies for Stabilizing the Dollar." In R. E. Hall (ed.) *Inflation: Causes and Effects*, 111–22. Chicago: University of Chicago Press.

Horrigan, B.R. (1987) "The CPI Futures Market: The Inflation Hedge That Won't Grow." Federal Reserve Bank of Philadelphia *Business Review* (May–June): 3–14.

Hoskins, W.L. (1992) "Views on Monetary Policy." Paper presented to the 17th annual policy conference, Dimensions of Monetary Policy: Essays in Honor of Anatol B. Balbach. Federal Reserve Bank of St. Louis. October.

Howitt, P. (1990) "Zero Inflation as a Long-term Target for Monetary Policy." In R.G. Lipsey (ed.) *Zero Inflation: The Goal of Price Stability*, 67–108. Toronto: C.D. Howe Institute.

Jansen, D.W. (1989) "Does Inflation Uncertainty Affect Output Growth? Further Evidence." Federal Reserve Bank of St. Louis *Review* (July-August): 43–54.

Jordan, J.L. (1992) "Skepticism about the Direction of Inflation: Causes, Costs, and Cures." Federal Reserve Bank of Cleveland *Economic Commentary*. 1 December.

Marshall, A. ([1887]1926) "Preliminary Memorandum" [to the Royal Commission on the Values of Gold and Silver]. Reprinted in A. Marshall *Official Papers*, 19–31. London: Macmillan.

Phelps, E.S. (1973) "Inflation in the Theory of Public Finance." *Swedish Journal of Economics* 75(1): 67–82.

Schnadt, N., and Whittaker, J. (1993) "Inflation-Proof Currency? The Feasibility of Variable Commodity Standards." *Journal of Money, Credit, and Banking* 25(2): 214–21.

Selgin, G.A. (1990) "Monetary Equilibrium and the Productivity Norm of Price-Level Policy." *Cato Journal* 10(1): 265–87.

Selgin, G.A. (1992) "On Foot-Loose Prices and Forecast-Free Monetary Regimes." *Cato Journal* 12(1): 75–80.

Selgin, G.A. (1994) "The Case for a 'Productivity Norm': Comment on Dowd." *Journal of Macroeconomics* 17(4): 733–40.

Selody, J.G. (1990) "The Goal of Price Stability: a Review of the Issues." Bank of Canada *Technical Report No. 54*. Ottawa: Bank of Canada.

Sumner, S. (1989) "Using Futures Instrument Prices to Target Nominal Income." *Bulletin of Economic Research* 41(2): 157–62.

Woolsey, W.W., and Yeager, L.B. (1991) "Is There a Paradox of Indirect Convertibility?" Paper presented to the Durell Foundation conference, American Money and Banking: Financial Fitness for the 1990s. Scottsdale, Ariz.

Yeager, L.B. (1992) "Toward Forecast-Free Monetary Institutions." *Cato Journal* 12(1): 53–73.