Alan S. Blinder is a professor in the department of economics at Princeton University.

Commentary

Alan S. Blinder

The title of Stephen G. Cecchetti's article is "Measuring Short-Run Inflation for Central Bankers," and that is exactly the viewpoint from which I will discuss it—the viewpoint of a *former* central banker.[1] I agreed to discuss this article because of my three-plus years of experience in month-to-month inflation tracking as a central banker and, before that, as a member of the Council of Economic Advisers. The name of the game then was distinguishing the signal from the noise, which was often difficult. The key question on my mind was typically: What part of each monthly observation on inflation is durable and what part is fleeting?

Actually, what I meant by the term durable differs subtly from Cecchetti's meaning. To me, the durable part of the information in each monthly inflation report was the part that was useful in medium- and near-term inflation fore-casting. This focus is slightly different from Cecchetti's; however, the two are related. The difference will assume some importance at the end of my remarks. But first I want to heap some praise on Cecchetti's article.

A little story explains why I agreed so readily to discuss this article. One day in 1995, I was having the chat I usually had with Fed staffers on the day the producer price index (PPI) or consumer price index (CPI) release came out. We did this exercise twice a month, generally focusing on the question: What's news and what's noise in this report? Suddenly, a thought occurred to me—the kind of idea that sometimes pops into your head and turns out to be on the right track, but not quite correct, once you have an opportunity to think about it more deeply. Why are we always taking the food and energy components out, I asked rhetorically? Because we believe they are extremely volatile. But we never actually test that hypothesis statistically, nor do we apply it to the other components. Instead of giving food and energy prices a weight of zero, and all other components a weight of one, shouldn't we really to be weighting *every* component inversely to its variance, or doing something similar? Incidentally, I believe my concrete suggestion to weight inversely by variances was mistaken—for reasons I will explain shortly. But first let me finish the story.

So I suggested to the staff members around the table that we should perhaps reweight the price indexes in this way instead of looking at "core" inflation. Then one of them—I can't recall who—said, "You know, there's a paper by Steve Cecchetti that claims we really ought to be looking at the weighted median." I replied that I would like to see that paper. So what happened next? I never got the paper, and I never saw the reweighted CPI. (This was no surprise, by the way.) So, until Cecchetti's article arrived in the mail the other day, I never saw another thing on this subject. I was extremely grateful when it arrived.

Cecchetti should be applauded for writing this article. It's just the sort of work I wish academics would do more often. Unfortunately, although research like this is extremely useful to policymakers, it is a little too nitty-gritty for the tastes of most academic researchers—which is a shame. Cecchetti deals with two very concrete issues—noise and bias. To me, the two are quite different conceptually. In fact, one of my minor quibbles with him is that he tries to merge the two in ways that are perhaps not so natural. However, Cecchetti does make a good point when he observes that the bias can be, and indeed seems to be, time varying. If the bias is bouncing around over time, it will look a lot like noise.

[1] I don't believe the viewpoints of current and former central bankers are very different.

THE BIAS ISSUE

I will concentrate on the noise issue—just as Cecchetti does. But first a few words on bias, once again from the viewpoint of a central banker.

The question of bias in the CPI—or indeed in any price index—has been intellectually fascinating for a very long time and continues to be so. But from the point of view of a central banker, this intellectually engaging question is not of great practical importance when inflation is considerably above its long-run target. Why not? Because, when you know you want to reduce inflation substantially, it is not terribly important to know whether the price index overstates inflation by, say, 0.5 percent or 1.5 percent. In either case, the policy is more or less the same: Bring inflation down. But as measured inflation approaches the upper limit of credible estimates of the bias in the inflation index, the magnitude of that bias becomes a crucial, if not *the* crucial, question of monetary policy. This was clearly not the case in the 1970s. However, it is the case now.

It is important to realize that the cost of making a mistake on this matter is enormous. If you do a standard calculation using a Phillips curve and Okun's Law, it will tell you that it would cost the country about $300 billion in lost gross domestic product if the Federal Reserve were to push inflation down one percentage point "too much." Now that's a tidy sum, which will cover the salaries of Federal Reserve staff for a very, very long time—and the salaries of governors for even longer.

Therein lies my quibble with Cecchetti's analysis. Rather than treat the bias as a purely statistical problem as he does (where the bias is defined more or less as the deviation from a 36-month moving average), I would favor a two-step approach. First, do what you can to reduce the bias by fixing up the index—including the weighting, aspects of the measurement of specific prices, and a whole variety of other difficult problems. Then use the dynamic factor index (DFI) that Cecchetti recommends to extract the signal from the noise. To me, jumping right to the DFI is putting the horse before the cart—and in an overly mechanical way at that.

THE NOISE ISSUE

Let me now turn to the noise issue, because this was my principal concern as Vice-Chairman of the Federal Reserve. I think it is a principal concern of central bankers everywhere.

As Cecchetti observes, two standard approaches exist to coping with the problem of noise in monthly data. One is to use the so-called core CPI.[2] The second is time averaging over, say, 3, 6, or 12 months. The first approach tends to be the American way: We concentrate on core inflation. The second approach tends to be the European way: When the inflation rate is announced in, say, Germany or Italy, it is almost always the 12-month trailing moving average.

Cecchetti offers us four fascinating findings. First, eliminating food and energy prices from the index is actually worse than doing nothing! This was astonishing to me, because I've been dropping food and energy prices since the 1970s. It set me thinking about the criteria on which this judgment is based. I will return to this point at the very end of my remarks, for it is, in a sense, my biggest objection to Cecchetti's analysis.

Second, he finds that averaging over three months is vastly superior to using monthly numbers. This finding really warmed the cockles of my heart because, when I was in government, journalists were constantly asking me about, say, the meaning of the August PPI report. If you are a professor sitting in your university office, you can give a curt answer to such a question: It has no meaning. But if you are a member of the Council of Economic Advisers or a Governor of the Federal Reserve, you can't get away with that.[3] So I used to cite "Blinder's Law of Weighting," which says that two consecutive monthly observations are worth four times as much as one, and three consecutive months are

[2] Incidentally, those of us who are old enough to remember the invention of the core concept, and I'm one, may recall that we first started computing the CPI with food, energy, *and used car* prices omitted. This last adjustment has since disappeared, and I don't know why. I think we should bring it back for a variety of reasons, including the fact that the price of a used car is an asset price.

[3] You can, however, say, "No comment," as I frequently did.

worth eight times as much as one. It's a doubling rule; however, I never carried it beyond $2^3 = 8$. My view, in brief, is that it takes at least three consecutive monthly numbers before you have any meaningful information. So I was delighted to see Cecchetti present this particular result. I am sure it is correct!

His third interesting finding is that some limited influence estimator (like the 10 percent trimmed mean) works best by most criteria. Of course, such complicated inflation measures are not in common use at central banks, to say the least.[4] But such measures are easy to compute. So perhaps central bankers should take notice.

And finally, Cecchetti finds that at least a full percentage point movement in the underlying inflation rate—which he represents by the 36-month moving average—is required before you can declare with any confidence at all (even 25 percent confidence, which is not much) that inflation has changed. This is truly bad news for the Federal Open Market Committee and, to the extent these results can be generalized to other countries, for central banks everywhere. For I can tell you from personal experience that a full percentage point change in the inflation rate is a very big deal to central bankers in low-inflation countries. In fact, central bankers in these circles sweat blood over whether the underlying inflation rate may have risen by a quarter percentage point! That's considered really important. So this finding is extremely bad news from their perspective.

TWO SUGGESTED IMPROVEMENTS

Here I would like to offer two suggestions that are closely related. One pertains to the criterion—using the 36-month centered moving average. The other has to do with how the various components should be weighted.

Regarding the criterion, notice that a 36-month *centered* moving average treats the past and the future symmetrically.

That is natural from the viewpoint of a scholar analyzing historical data. But it is very unnatural from the viewpoint of a central banker living in real time. Historically, $t-1$ and $t+1$ look more or less the same. But if you must make a decision at time t, there is a world of difference between $t-1$ and $t+1$. The past is known; the future is not. More important, the past is the dead hand of history, but the central bank must worry about the future. That is why I said earlier that I regarded information as "durable" if it had value for near- and medium-term forecasting.

Thinking about the distinction between the past and the future gave me an idea. Why not assign different weights to each component, but not in the way I originally suggested that day at the Fed? In retrospect, that idea seems wrong. Why not weight each component, instead, by its usefulness in forecasting *future*—not past—inflation? By studying deviations from a 36-month centered moving average, Cecchetti's analysis treats the next 18 months and the past 18 months symmetrically. I am suggesting, instead, that the next 18 months matter a good deal more to central bankers than the past 18 months.

That brings me to my second suggestion—the weighting. The trimmed mean assigns a zero weight to the largest and smallest component-specific inflation rates, and a unit weight to everything else. In computing it, you pay no attention to the *name* of each component. But I want to suggest that we should perhaps pay attention to the component's name. Why? Because I imagine that some components have more ability to forecast future "underlying" inflation than others. Clearly, the weighting scheme I am now suggesting is not the same as inverse variances. Instead, it probably takes into account (a) the persistence of each component and (b) its covariance with overall inflation. From the point of view of a central banker looking to the future, I think it probably makes sense to reweight all the components of the CPI in this way—which is clearly quite different from,

[4] A prominent exception is the Bank of England, which includes one or more such indicators in its regular inflation reports.

say, computing the 10 percent trimmed mean.

A PRACTICAL ILLUSTRATION

Let me wrap up by illustrating where this alternative point of view might lead—and how it differs from Cecchetti's. At the beginning of his article, Cecchetti refers to supply shocks and exchange-rate changes as "sources of noise." That phrase practically jumped right off the page at me, for it was not at all the way I used to think about these shocks. At the Fed, I can assure you, we never thought of the exchange rate as a source of "noise." Far from being statistical marginalia, we thought of both supply shocks and exchange-rate move-ments as raising important policy issues. Consider supply shocks.

Suppose first that food and energy prices, whether or not they are "volatile," are included in the central bank's target inflation rate—the concept of inflation it really cares about. This is Cecchetti's tacit assumption, for he is always looking at deviations from the smoothed path of CPI-U. Then, if food and energy prices are persistent, we surely would *not* want to take them out of the index. We clearly would want to leave them in. This may account for the poor results Cecchetti gets when he knocks food and energy prices out of the index.

Now suppose, instead, that food and energy prices are not part of the price index the central bank is targeting. It cares instead about, say, "core" inflation. Even then, however, the central bank *might* still want to include food and energy prices in the index it monitors on a monthly basis. It all depends on whether recent values of food and energy inflation help forecast future core inflation.

The latter comes closer to my own views as a practicing central banker. Indeed, while at the Fed, I was even known to use the price of gold in forecasting inflation. That shows the depths to which you can sink if you immerse yourself too deeply in the real world!

As a central banker, I always preferred to view the inflation rate with its food and energy components removed as our basic goal. But not because these components are extremely volatile. (I must admit, however, I thought they were until Cecchetti's article arrived.) The real reason was that the prices of food (really, food at home) and energy are, for the most part, beyond the control of the central bank. The Fed cannot do much about food and energy prices—except, of course, to cause a recession deep enough to ensure that increases in these prices do not lead to overall inflation. But the central bank *can* do something about the rest of the price index—the part that comes out of the industrial *core* of the economy, so to speak.

So I believe there is a good reason why so many of the countries that now post official inflation targets exclude food and energy prices from their target inflation rates: These prices do not logically belong on their report cards.