and people are quite good in debiasing the numbers reported in the official statistics. Even so, it is important to keep in mind that the manipulation of the inflation rate in Argentina may have done harm in ways that are not studied in this paper. Manipulation could have substantial effects on inflation uncertainty with possibly detrimental welfare consequences. Furthermore, manipulation may have resulted in a substantial wealth transfer away from holders of inflation-linked bonds.

REFERENCES FOR THE NAGEL COMMENT

Carrillo, Paul E., and M. Shahe Emran. 2012. "Public Information and Inflation Expectations: Microeconometric Evidence from a Natural Experiment." *Review of Economics and Statistics* 94, no. 4: 860–77.

Drenik, Andres, and Diego J. Perez. 2016. "Price Setting under Uncertainty about Inflation." Working paper. http://www.perezdiego.org/wp-content/uploads/2016/06/Drenik_Perez_2016.pdf

Hellwig, Christian. 2005. "Heterogeneous Information and the Welfare Effects of Public Information Disclosures." Economics Online Paper no. 283. Los Angeles: University of California.

Webber, Jude. 2008. "Argentine Inflation Index under Fire." *Financial Times,* June 11.

COMMENT BY

RICARDO REIS This paper by Alberto Cavallo, Guillermo Cruces, and Ricardo Perez-Truglia provides a fascinating account of the extent to which the Argentine government manipulated inflation statistics between 2006 and 2015. The government enacted price controls, the common recipe to stop inflation that rarely works but is sure to distort relative prices and induce misallocation. More originally, the government changed the methodology used to construct price indexes, confirming an old fear among economists and statisticians that when given a range of possible estimates from alternative methods, politicians behave as if picking from a menu rather than as Bayesians facing uncertainty. This culminated with the firing of high-ranking staff members of the Argentine government's statistics bureau, the Instituto Nacional de Estadística y Censos (INDEC), going back to the old tradition of shooting the messenger when the message is not what politicians want.

The authors' figure 1 describes most of the developments in this sad story of government manipulation of statistics; it is worth reading their paper just for this figure. The authors, however, are not political novelists but top-notch economists, and true to their commitment to science, they resist the temptation to dwell on this story. Instead, they use it as a pretext to explain how people learn from statistics and form their inflation expectations. To stay focused on this goal, and to resist the siren call that comes from the authors' figure 1, I take a step further and pose the question as if it were being applied to a less interesting country: the United States.

IMITATING THE SPIRIT OF THE AUTHORS' EXPERIMENTS Imagine that I approach a crowd of economists and policymakers and ask:

> What do you think was the annual U.S. inflation rate with respect to one year ago?

Note that I am asking for a fact about the past, not a forecast for the future. Yet surely I would get a distribution of answers. Even among a very-well-informed audience, some are better informed than others. Some are more confident or optimistic, and there is even some research suggesting that gender partly determines confidence, leading to disagreement. Some would interpret the question slightly differently from others, no matter how clear I try to be. From the raw distribution of answers, I would learn only one thing: People disagree and are not perfectly informed.

Imagine now that instead I randomize among my crowd of people, dividing them into six groups. I ask the first group:

> According to official indicators published by the Bureau of Labor Statistics, the annual inflation rate with respect to one year ago was approximately 0.1 percent. What do you think was the annual U.S. inflation rate with respect to one year ago?

I would ask the second group exactly the same question, but replacing 0.1 percent with 1.4 percent. Finally, I would do the same with the third group, but now quoting a figure of 2.2 percent. What do you think the answers would be?

Perhaps my groups of survey respondents would just find the questions awkward, and repeat back to me the number that I had given them in the question. The distribution of answers across the three groups would then have three points, with exactly the same number of respondents in each.

Perhaps instead my respondents would have thought that I must be tricking them (why would I ask such a silly question anyway!?), and so would give me a different number from the one in the question. Still, my strong prior belief is that those in the first group would give lower answers than those in the second group, and lower even than those in the third group. As long as they put at least some weight on the possibility that the number that I was giving them had some credibility, it seems plausible that this would affect their estimate. And, by the way, my three numbers are not lies, but come from the Bureau of Labor Statistics' (BLS) Consumer Price Index (CPI) economic news release for January 2016: 0.1 percent, 1.4 percent, and 2.2 percent were the 12-month changes in the CPI for the Cleveland area, for all items in the nation, and for all items except food and energy.[1]

Alternatively, imagine that I ask the fourth group a different question:

> According to *other* indicators published by the *Bureau of Economic Analysis,* the annual inflation rate with respect to one year ago was approximately –2.0 percent. What do you think was the annual U.S. inflation rate with respect to one year ago?

The fifth and sixth groups would get the same question, but with the numbers 0.3 percent and 1.0 percent. Again, these are all true: The three numbers refer to the change in the deflators for nondurable goods, personal consumption expenditures, and gross domestic product.

My guess is that again the fourth group would expect lower inflation than the fifth, and even lower than the sixth. I would also venture that there would be differences in the distribution between these three groups and the previous three. My informed respondents would note that I refer in the question to these indicators as *other* rather than official, perhaps increasing their suspicion toward me. Moreover, they would know that the more commonly used measure of inflation is the CPI computed by the BLS, not the deflators computed by the Bureau of Economic Analysis, so they might regard this information as not quite as reliable as the previous one.

In essence, this is what the authors do in their surveys of Argentines. Their respondents are not as trained in economics and statistics as my hypothetical ones, even if they are more educated than the typical Argentine, and they are used to living in a country that often faces high and volatile inflation, making them more attentive to this economic indicator. To be clear, with my thought experiment, I do not want to undermine the authors' remarkable work designing and implementing these surveys, nor to undervalue how important it is to go from thought experiments to actually collecting data that may well end up challenging one's priors. My goal is instead to focus on what information was being given to the respondents and what was being asked of them, so I can proceed to discuss what we may or may not learn from it.

WHAT CAN WE LEARN FROM THE RESULTS? The first result that the authors obtain is that providing information has an effect on the answers that

1. BLS, "Consumer Price Index: January 2016," news release, February 19, 2016 (http://www.bls.gov/news.release/archives/cpi_02192016.pdf).

people give. My six groups described above would not have given the same answer if they were like the Argentines in the authors' sample. The authors read this as a triumph for the Bayesian proposition that people do not ignore valuable pieces of information, but use them to update their priors toward new posteriors.

I agree. But this is also a fairly low bar. Only if the information were absolutely and completely useless would a Bayesian ignore it. All six numbers that I provided in my hypothetical survey, and likewise the authors' six numbers in the actual survey, were not just true but also definitely informative about what inflation must be. Even in the case of the biased government statistics, the respondents to the authors' survey certainly had some information about true inflation, even if it was muddled by the government's manipulations. Moreover, one would expect that even if the information provided was indeed useless, the people receiving it in the way described in the interview might well presume that it was somewhat useful.

Moreover, even a non-Bayesian would be expected to react to this information. Endless psychological studies have shown that cues affect responses. Providing a number, even if it is arbitrary and useless, anchors future responses to questions that ask for numbers (Tversky and Kahneman 1974). Moreover, the very-well-known Hawthorne effect states that subjects of a study have their behavior affected by being aware that they are being observed. In the case of this survey, this would likely lead even a non-Bayesian to have the number that they were given in the question affect his or her answer, even if this number had no effect on their actual expectations of inflation and on their subsequent economic choices. Having an interviewer tell you that inflation is 0.1 percent makes it hard for you to reply that it is actually 10 percent, even if this is what you really think.

The second result is that the distribution of answers across the groups that were given the official statistics is different from the distribution of answers in the group given the alternative indicators. In terms of my experiment above, the distribution across the first three groups would be different from the distribution across the last three groups.

More precisely, the authors show that people's answers are consistent with the hypothesis that when receiving information from the official Argentine indicators, they subtract a constant 10 percent perceived upward bias. Thus, the distribution of answers for a group that is told that an official statistic is 20 percent is similar to the distribution for a group that is told that an alternative indicator is 10 percent. In symbols, if the distribution of answers after an unofficial statistic is revealed appears to be drawn from some distribution with mean x and variance y, then the distribution of answers after an official statistic is revealed seems to be drawn from a similar distribution, which is different only in having a mean $x - b$, where b is the bias.

These results are again persuasive, and the differences across groups can be easily inferred visually. At the same time, failing to reject the null hypothesis that people behave as if there was a constant mean bias is not the same as accepting this hypothesis about people's behavior. Consider two alternatives. First, perhaps the bias is multiplicative, so that, instead, the distribution following the official numbers has a mean of bx. Would the data reject this alternative? Second, perhaps there is no bias but rather a perception of different precision or informativeness such that the distribution after the official number has the same mean but a variance of by. The authors' data would have trouble distinguishing this alternative.

Moreover, bias is not the same as cheating. We know that the CPI measures produced by the BLS suffer from substitution bias. Since the 1996 Boskin Commission Report, a common rule of thumb in the United States has been to subtract about 1.3 percent from the CPI statistic to get closer to the true cost of living.[2] But few people see in this any form of cheating by the BLS.

The third result of this paper is that there is an asymmetry in people's responses. Because they distrust the official sources as understating inflation, people respond more to official statistics that report higher inflation than to official statistics reporting lower inflation. The argument goes that for the government to be reporting high inflation, then actual inflation must be really high, to the point where it cannot be hidden anymore.

Interestingly, however, the asymmetry is also there in the distribution of responses that people gave after being told an unofficial inflation statistic. This suggests that the source of the asymmetry is not driven by the data they are provided, but rather by the person's responses to any information. On one hand, this may be because people in Argentina have learned to distrust *any* inflation number, regardless of its source. On the other hand, it may be the result of forming forecasts while having an asymmetric loss function in their mind. Insofar as higher inflation causes real income losses, and there is diminishing marginal utility from this income, this could justify such an asymmetry.

2. Social Security Administration, "The Boskin Commission Report: Toward a More Accurate Measure of the Cost of Living—Final Report to the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index, December 4, 1996" (https://www.ssa.gov/history/reports/boskinrpt.html).

The authors' three results are solid and hard to dispute. As often happens, however, the results are open to more than one interpretation.

WHAT CAN WE CONCLUDE ABOUT LEARNING? A separate question is whether the authors' methods, survey answers, and statistical analysis allow us to reach broader conclusions about learning and data. The authors are careful not to claim these conclusions; but it is the role of the discussant to speculate about whether they do.

First, can we conclude that their survey methodology is able to isolate the effects of information on expectations? Some notation is helpful to understand the authors' method. Let person i's prior answer on what was inflation in the past 12 months be $a^{prior}(i)$. After receiving the piece of data from the interviewer, the person will have a posterior $a^{post}(i)$. People are sorted into two groups: those treated with the official inflation reports, in group T; and those in the control group who do not receive this information, in group C. The goal is to estimate information's effect on the revision of people's answers as a result of the treatment, which can be done by comparing the two sample means:

$$\sum_{i\in T}\left[a^{post}(i)-a^{prior}(i)\right]-\sum_{i\in C}\left[a^{post}(i)-a^{prior}(i)\right].$$

However, the authors did not elicit the priors, so they do not observe $a^{prior}(i)$. As a result, their statistics are instead based on

$$\sum_{i\in T}a^{post}(i)-\sum_{i\in C}a^{post}(i).$$

Clearly, this is a valid measure only as long as

$$\sum_{i\in T}a^{prior}(i)=\sum_{i\in C}a^{prior}(i).$$

The reason why we expect this to be the case is through the randomization of people into treatment and control groups. If this randomization ensured that being part of each of the two groups is not correlated with any important source of differences across people's inflation expectations, then this condition would hold. The authors' sample plausibly satisfies this condition. The only source of concern is that their sample has a larger share of women than the population, 57 percent versus 53 percent, and there is a weak suggestion in the literature that women's inflation expectations are systematically different from men's (Bryan and Venkatu 2011).

Second, can we use their method to conclude that there is a constant inflation bias in the official data that people rationally take into account when using data from official sources to form their inflation expectations? This is a significantly harder question. The authors persuasively show that one cannot reject the null hypothesis that there is a constant 10 percent inflation bias that people take into account. But the flexibility of Bayes's rule does not allow us to confidently pin down whether the bias exists, whether it is constant, or whether it is 10 percent. With only their data, but with freedom to choose people's loss function for making forecasts and freedom to choose the two distributions from which the signals on inflation are drawn, the official and the alternative one, then we could get almost any estimate of the bias. Bayes's rule is very flexible and can accommodate many different patterns of responses.

Third, can we conclude that agents are sophisticated Bayesians, rationally discounting biased data? Again, the authors convincingly show that this null hypothesis is hard to reject. In fact, their results are even stronger. They support the modern theories of inattention, according to which the disagreement that we observe is due to people not having the same information, but once people get to pay attention—for instance, because an interviewer gives them information—they rationally update their beliefs (Reis 2006).

At the same time, the data have two features that are harder to reconcile with this optimal inattentiveness. First, why would Argentines—who by many accounts are quite informed about inflation, having lived through great price volatility many times in the recent past—have such loose priors? The authors' data show that giving one single number in an interview has a large effect on people's perceptions of inflation, which must imply that they were quite uncertain about it in the first place. Second, why do perceptions of past inflation line up so closely with expectations of future inflation (as seen in the authors' figure 5)? The serial correlation of inflation is well below 1 in the Argentine data, so this extent of persistence in perceptions and forecasts will likely lead to serially correlated forecasting errors.

ARE ARGENTINES UNSOPHISTICATED AFTER ALL? Having made a case for Argentines being quite sophisticated in using official manipulated data and forming inflation expectations, the authors move in a different direction in section III. Here, they show the result from asking people outside a supermarket about the historical price changes of the goods they have just bought. Conceptually, this is a very different question from the one considered in the rest of the paper. Here, it is not inflation—the general increase in prices for a wide basket of goods—that people are being asked about, but rather the prices of the individual goods they bought minutes ago and how they compare with what people think these prices were 12 months ago.

Impressively, the authors show that even though the Argentine government had imposed strict price controls on some goods during this period, people's perceptions of how these goods' prices changed, relative to those goods whose prices were free from government meddling, were essentially the same. This form of government manipulation—here, not of statistics but of goods' prices themselves—seems to again have had little effect on the Argentine public.

However, another conclusion is striking: Remembered price changes are extraordinarily higher than actual price changes, as shown in the authors' figure 7. Although fewer than 5 percent of prices changed by more than 60 percent, people answer that more than 40 percent of prices changed by this amount or more. By this account, Argentines' answers are so far off from the facts that they seem quite remarkably unsophisticated.

CONCLUSION This paper has two goals, and thus its results have two possible takeaways. The first is that in Argentina, people do not let the manipulation of official statistics and prices fool them. Even as the government seemed to bias official statistics down or to control the price changes of individual goods, the public's perceptions of actual inflation and future inflation remained high. Government data were debiased rather than taken at face value, and branding a piece of data as "official" led the public to treat it differently right away. Reality seemed to prevail over propaganda.

The second takeaway pertains to people's sophistication in forming perceptions about inflation. Here the bag is more mixed. In some respects, Argentines seem quite sophisticated; but in others, they are remarkably biased. The authors' data and statistics provide very valuable information with which to judge models of the formation of expectations, but they are not quite decisive toward any one particular theory.

Perhaps this paper's overall lesson, especially for policymakers, is that in spite of all the studies and research showing that people are far from rational in forecasting inflation, it does not follow that policymakers can therefore easily manipulate people's views. People may not be all that rational in dealing with economic data and forecasts, but they are experienced enough not to be duped by their governments.

REFERENCES FOR THE REIS COMMENT

Bryan, Michael F., and Guhan Venkatu. 2001. "The Curiously Different Inflation Perspectives of Men and Women." Economic Commentary. Federal Reserve Bank of Cleveland.

Reis, Ricardo. 2006. "Inattentive Consumers." *Journal of Monetary Economics* 53, no. 8: 1761–800.

Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185, no. 4157: 1124–31.

GENERAL DISCUSSION Justin Wolfers opened the discussion with some "clownish" facts about unemployment and inflation beliefs in the United States. According to a recent survey, 34 percent of Americans believe that unemployment is higher today than when President Barack Obama took office, with 53 percent of Republicans believing this.[1] A different statistic, something a little closer to the sense of the paper, is that there is a faction in the United States that believes that official CPI statistics are being terribly manipulated. Subscribers to the electronic newsletter service *Shadow Government Statistics* (www.shadowstats.com) can pay $175 per year to learn what the "real" inflation rate is. "That $175 a year," he joked, "what they'll do is they'll take the CPI and add 8 points to it for you." Ironically, he noted, the price of the subscription has not changed in eight years, implying "a substantial real price cut in the price of ShadowStats."

Both of these statistics, Wolfers noted, move the focus away from the *mean* of expectations to the *distribution*, which produces very different views about the world. What the paper shows is that the mean of expectations moves in a sensible way. However, looking at the micro data on any expectations, "Most people have completely stupid expectations." The first moment is not going to be enough; one needs to know the full distribution. "It might be that the full distribution story is people move from completely clueless to completely clueless," he concluded, which is a different story than people being quite sophisticated in debiasing.

Marshall Reinsdorf described Argentina as a very decentralized country, and remarked that one of the many fascinating things for him about Argentina during the 2006–15 period studied in the paper were the differences between the various provincial inflation rates. Looking at a random sample

1. In fact, the official U.S. unemployment rate in January 2009 was 7.8 percent; in January 2016, it was 4.9 percent.