ALBERTO CAVALLO
Massachusetts Institute of Technology

GUILLERMO CRUCES
Universidad Nacional de La Plata

RICARDO PEREZ-TRUGLIA
Microsoft Research

Learning from Potentially Biased Statistics

ABSTRACT   When forming expectations, households may be influenced by perceived bias in the information they receive. In this paper, we study how individuals learn from potentially biased statistics using data from both a natural experiment and a survey experiment during a period (2007–15) when the government of Argentina was manipulating official inflation statistics. This period is interesting because attention was being given to inflation information and both official and unofficial statistics were available. Our evidence suggests that, rather than ignoring biased statistics or naively accepting them, households react in a sophisticated way, as predicted by a Bayesian learning model. We also find evidence of an asymmetric reaction to inflation signals, with expectations changing more when the inflation rate rises than when it falls. These results could also be useful for understanding the formation of inflation expectations in less extreme contexts than Argentina, such as the United States and Europe, where experts may agree that statistics are unbiased but households are not.

Household inflation expectations play a key role in models of consumption decisions and the real effects of monetary policy, yet little is known about how these expectations are formed. In recent years, a growing body of empirical literature has been providing evidence about how individuals use information to form their inflation expectations. For example, in Cavallo, Cruces, and Perez-Truglia (2014), we show that individuals learn from both inflation statistics and supermarket prices. In this paper, we use data from a period of manipulated official statistics in Argentina to study the degree of sophistication in this learning process and the role of trust in statistics.

Our findings are based on observational and experimental evidence obtained in Argentina during the recent period, from 2007 to 2015, when the government manipulated inflation statistics. This is an ideal setting, for three main reasons. First, the inflation rate fluctuated between 15 and 30 percent, which led to high inattention costs and encouraged individuals to spend time gathering and processing information about the inflation rate.[1] Second, ample evidence suggests that the official sources of inflation information, such as the Consumer Price Index (CPI), were intentionally biased.[2] And third, the lack of reliable official data during this period promoted the creation of several unofficial inflation indicators, thereby potentially allowing individuals to counteract the government's manipulation by using other data.

We start with observational data on the comovement of inflation expectations and official and unofficial inflation statistics, both before and after the intervention by the Argentine government's statistics bureau, Instituto Nacional de Estadística y Censos (INDEC), when the government started reporting official statistics that were systematically below the unofficial estimates. Household inflation expectations quickly diverged from the official inflation indicators and instead aligned with the unofficial indicators. This change suggests that consumers are not naive learners who accept official statistics as unbiased. However, this observational evidence presents two challenges. First, we do not observe the distribution of expectations in the counterfactual scenario without manipulated official statistics. Second, the evidence does not address the nature of the learning process, such as whether individuals simply ignore official statistics or use their information in a sophisticated way.

To address these limitations in the observational data, we provide a simple model of Bayesian learners with potentially biased statistics and design a survey experiment to test its predictions. This model shows that, far from ignoring official statistics, rational learners should react to changes in official statistics by "debiasing" the signal on the basis of their perceived bias

1. Because they cannot write contracts in foreign currency or indexed by inflation, households needed to constantly estimate inflation to sign rent contracts, negotiate wages, and make savings and investment decisions. Indeed, during the period we are studying, inflation statistics were frequently mentioned and discussed in the front pages of newspapers and other media outlets, and opinion polls systematically indicated that inflation was perceived as one of the most important problems in the country.

2. For a discussion of the evidence for the manipulation of statistics, see Cavallo (2013). Our paper extends the account of the main events from 2006 until December 2015, when a new government finally suspended the publication of the official CPI.

while simultaneously updating their beliefs about the size of the official bias. In other words, we predict that rational consumers will extract useful information from potentially biased information.

In December 2012, we conducted a survey experiment in Argentina to test this prediction. We provided respondents with different inflation estimates, and we measured their subsequent inflation perceptions and inflation expectations, as well as their confidence in these perceptions. By leveraging the variety of inflation indicators available at the time, we cross-randomized, in a nondeceptive way, two features of the message that was provided to subjects: the source of the inflation statistics (official and unofficial), and the inflation rate (10, 20, or 30 percent).

Our experimental evidence rejects the hypothesis that individuals ignore information from biased official statistics. Subjects reacted significantly to all signals, including official statistics. For example, compared with individuals who were told that the official inflation rate was 20 percent, individuals who were told that it was 10 percent reported lower inflation perceptions and expectations, and individuals who were told that it was 30 percent reported higher ones.

The experimental data also allow us to directly test the hypothesis that there may be sophisticated learning from potentially biased statistics. Because the official statistics were consistently 10 percentage points below the unofficial estimates, our Bayesian model predicts that individuals should react similarly to a signal that official inflation is 10 percent as they would to a signal that unofficial inflation is 20 percent, and that they should react similarly to an official rate of 20 percent as they would to an unofficial rate of 30 percent. These predictions are consistent with subjects' reactions in our experiment. That is, in an environment where there are many alternative inflation indicators and much attention is being given to the topic, individuals function as sophisticated learners who can deal with potentially biased information.

The experiment also allowed us to explore another pattern found in our analysis of the observational data: Expectations follow actual inflation more strongly when actual inflation is rising than when it is falling. Consistent with this asymmetric pattern, we find that subjects were nearly twice as reactive to new information about higher inflation as they were to information about lower inflation, even when the information came from unofficial sources. Indeed, we discuss the possibility that this asymmetric learning was generated by the introduction of manipulated statistics.

A group of studies suggests that individuals form inflation expectations using information from their own consumer experiences (Bates and

Gabor 1986; Bruine de Bruin, van der Klaauw, and Topa 2011; Coibion and Gorodnichenko 2015; Kumar and others 2015; Malmendier and Nagel 2016). In particular, individuals rely heavily on their perceptions about the prices of individual supermarket products (Cavallo, Cruces, and Perez-Truglia 2014). These findings imply that the government could try to influence inflation expectations by changing the actual prices of salient products. Indeed, in an effort to curb inflation, in 2013 the Argentine government froze the prices of a relatively large and important sample of consumer products. We show that, even though the inflation rate then fell significantly, household inflation expectations did not fall. To further explore this finding, we ran a price-elicitation survey outside a large supermarket chain in Argentina during the period of price controls. We found that even though there was a substantial difference in the actual price changes between goods that were under price controls and those that were not, consumers did not perceive such price differences.

Although the context of manipulated statistics in Argentina is an extreme case, these results can nonetheless help to explain how individuals learn from inflation data in other countries. Even in developed nations, a significant share of individuals do not trust official statistics. For instance, according to a *Eurobarometer* report by the European Commission (2010), 71 percent of respondents in Europe in 2009 felt that it was necessary to know about economic indicators, but only 44 percent stated that they tended to trust official statistics such as the growth rate, the inflation rate, and the unemployment rate.[3] Among U.S. survey respondents, 27 percent rated their trust in official statistics as 4 or lower on a scale of 1 to 10 (Curtin 2009). Analysts, commentators, and the media routinely discuss the possibility of manipulated statistics, such as those that may have been reflected in the job creation rates that were released right before the 2012 election in the United States (Norris 2014).

Data from a survey of U.S. households reported in Cavallo, Cruces, and Perez-Truglia (2014) show that 32 percent of respondents do not trust official inflation statistics. Furthermore, compared with those who trust inflation statistics, respondents who do not trust statistics have inflation expectations that are 50 percent higher on average. This evidence suggests that a lack of trust in the government may explain part of the stylized

---

3. In the 2007 wave of the survey, 69 percent of respondents felt it was necessary to know about economic indicators, and 46 percent stated they tended to trust official statistics. See European Commission (2010, pp. 35, 44).

fact that households do not fully incorporate information from inflation statistics into their perceptions and expectations (Mankiw and Reis 2002; Mankiw, Reis, and Wolfers 2004; Carroll 2003).

To the best of our knowledge, this paper is the first to study how individuals learn from manipulated statistics. More generally, the study of biased statistics goes back to the seminal contribution by Oskar Morgenstern (1963) on measurement, accuracy, and uncertainty in economics. Morgenstern's book discusses how both private companies and governments have strong incentives to manipulate information, and he applies this argument to the problems of measuring prices.[4] Recent studies use data to measure the degree of bias in official statistics, including examples of inflation in Argentina (Cavallo 2013), debt manipulation indicators in Greece (Rauch and others 2011), and alternative growth and inflation estimates in China (Nakamura, Steinsson, and Liu 2014). Tomasz Michalski and Gilles Stoltz (2013), in turn, use statistical regularities in economic indicators to suggest that countries seem to manipulate economic data systematically.

Our paper also relates to a growing body of literature on the formation of household economic expectations. In particular, it is widely recognized that identifying the formation of inflation expectations is important to understand the link between the nominal and real sides of the economy (Bernanke 2007; Coibion and Gorodnichenko 2015; Bachmann, Berg, and Sims 2015; D'Acunto, Hoang, and Weber 2016). Several studies provide evidence that inflation statistics play a significant role in driving inflation expectations, including the analysis of variation in the media's coverage of statistics (Lamla and Lein 2008; Badarinza and Buchmann 2009; Dräger 2015), quasi-experimental variation in reporting official statistics (Carrillo and Emran 2012), and information-provision experiments (Roos and Schmidt 2012; Armantier and others 2016; Cavallo, Cruces, and Perez-Truglia 2014).

Finally, this paper relates to a theoretical literature about whether central banks (or other government agencies) should commit to provide timely and accurate information about economic fundamentals. For instance, some authors argue that information disclosure enhances welfare (Hellwig 2005), whereas others argue that it can reduce welfare

4. Morgenstern (1963) also covers the difficulties of measuring the national product, and in fact Argentina's government also falsified INDEC's GDP indicator (Camacho, Dal Bianco, and Martinez-Martin 2015), for political reasons and to avoid the payment of a GDP warrant (a bond that only pays debtors if the GDP grows at a certain rate).

(Morris and Shin 2002). The majority of these studies focus on the margin of disclosing truthful information or not. We focus on a margin that has been widely overlooked: manipulating the information that is disclosed.

The paper proceeds as follows. Section I describes the period of manipulation of official statistics in Argentina and presents the observational evidence. Section II presents a simple model of Bayesian learning from manipulated statistics, as well as the design of the survey experiment and its results. In section III, we discuss the period of price controls in 2013. Section IV concludes.

I. The Manipulation of Inflation Statistics in Argentina

This section describes the main events related to the manipulation of official inflation statistics, as well as the emergence of unofficial estimates and their comovement with consumers' inflation expectations.

I.A. The Government's Intervention at INDEC

After a severe economic crisis in 2001–02, the Argentine economy started to recover in 2003, mostly due to an unprecedented increase in commodity prices. Inflation levels were relatively low at the beginning of the recovery, but they reached double digits in 2005 (12.3 percent per year).

Figure 1 provides a timeline of the most important events from 2006 to 2015. During 2006, the government imposed a series of price controls and organized public boycotts against some retailers. The government also pressured the professional staff at INDEC to make methodological changes that could lower the annual inflation rate. For example, the government asked INDEC to reveal which stores were collecting data, to introduce automatic substitutions to reduce the weight from items that had higher inflation, and to use prices from goods on price control lists even if those goods were not available for sale at the stores where the data were being collected.

In February 2007, facing a second year of inflation above 10 percent and unwilling to scale back its expansionary monetary policy, the government made the drastic decision to fire high-ranking members of the INDEC staff, including Graciela Bevacqua, the statistician in charge of the team that computed the CPI. The monthly inflation rate fell from 1.1 percent in January 2007 to 0.4 percent in February and continued falling in the subsequent months. INDEC's employees publicly disclosed what had happened in the previous months, which increased suspicions that the CPI was being manipulated. INDEC stopped publishing several disaggregated

Figure 1. Timeline of the Manipulation of Inflation Statistics in Argentina, 2006–16

Feb 2006	The secretary of interior commerce, Guillermo Moreno, tries to gain access to micro data protected by statistical confidentiality laws
Oct 2006	Moreno hires a market-research firm, Tomadato, to produce an alternative CPI
Jan 2007	The director of INDEC announces that Beatriz Paglieri, Moreno's assistant, will be visiting the institution for one month to check the last estimations
	First meeting with Paglieri and the directors of INDEC
	Paglieri decides to stop the publication of the CPI for the greater Buenos Aires metropolitan area, the CPI-GBA
	Graciela Bevacqua, the director of the Prices Department, is suspended
Feb 2007	The government officially intervenes at INDEC
	The first manipulated CPI-GBA monthly index is published
	Bevacqua is officially fired and replaced by Paglieri
	First mobilization of INDEC employees takes place (repeated every month since)
	Senators from the opposition ask a federal prosecutor to intervene
Mar 2007	The director of INDEC, Leilo Marmora, resigns
May 2007	Manuel Garrido, a federal prosecutor, says serious irregularities took place at INDEC
Jul 2007	Cynthia Pok, in charge of the National Household Survey, is fired
	First official strike by INDEC employees
Sep 2007	INDEC modifies Mendoza province's inflation rate before publication
	The calculation of the CPI-Nacional is changed
Dec 2007	Cristina Kirchner becomes the president of Argentina, succeeding her husband
Jan 2008	INDEC employees receive wage cuts
Mar 2008	Launch of www.inflacionverdadera.com, a website where alternative indicators using online prices are updated on a daily basis; the inflation rate is three times higher than official CPI estimates
May 2008	INDEC stops publishing the CPI-Nacional, an index that used price data from seven provinces
	INDEC announces new CPI weights; food becomes more important in the new index
	Some employees of INDEC are physically assaulted by government supporters at the Finance Ministry building
Nov 2010	The government announces an agreement with the International Monetary Fund for the normalization of the statistics
Feb 2011	Moreno asks private consultants to share the methodology of their CPI calculations; most of them refuse
Mar 2011	Some private consultants are fined 500,000 pesos for failing to comply with Moreno's request
May 2011	The Congress Average index (an average of private consultants' inflation rates) is born
	A judge rejects the fines imposed on private consultants
Sep 2011	Private consultants receive letters from the government threatening them with criminal prosecution if they continue to publish their inflation estimates
Feb 2012	The International Monetary Fund announces that Argentina did not improve the CPI-GBA according to the international rules
	The Economist stops publishing Argentina's official statistics and uses instead the index produced by PriceStats (a company working with the Billion Prices Project at MIT)
Feb 2013	The International Monetary Fund issues a "motion of censure" against Argentina for the bad-quality statistics
	Moreno is replaced by Augusto Costa as Secretary of Interior Commerce
Jan 2014	The CPI-GBA is replaced by a new index, called CPI-Nu; it initially shows similar monthly inflation rates to unofficial estimates, but starts to diverge once again within a few months
Apr 2014	The government announces that the official poverty index will no longer be published
Dec 2015	Mauricio Macri, a member of the opposition, becomes the new president of Argentina
	Jorge Todesca becomes INDEC's director, and Bevacqua returns as its technical director
	Todesca says that INDEC is like a "scorched earth," and suspends publication of the CPI and other price indexes
Jan 2016	Bevacqua announces that it will take eight months for INDEC to publish a new CPI
Feb 2016	Bevacqua is fired again

Sources: Various newspaper articles and other sources, compiled by the authors.

inflation series, and it announced "methodological changes" that were never publicly disclosed.

The government's intervention at INDEC had immediate negative consequences for the Argentine economy, as discussed by Eduardo Levy Yeyati and Marcos Novarro (2013). Although the government paid less in the short run on inflation-linked bonds, most of this debt was held by the government's own pension funds. The price of these bonds quickly fell, as investors internalized the manipulation. The government also paid much higher interest rates for newly issued debt.[5] Economic uncertainty increased, bank deposits fell, and capital outflows surged, which eventually, in 2011, led the government to impose foreign exchange controls. Despite the controversy and the negative effects on the economy, the manipulation of the official CPI continued until December 2015, when a new government was elected.

I.B. *Unofficial Inflation Statistics*

INDEC's unusual situation led to the creation of alternative measures of inflation, which we generally term "unofficial" inflation indicators. The main alternative indicator we use is computed by PriceStats, a private firm based in the United States that since 2007 has been using online prices from large retailers. The PriceStats index is published weekly in *The Economist.*[6] A second alternative indicator, published since 2008, is produced by the organization named Buenos Aires City, a think tank led by Graciela Bevacqua (the former head of INDEC's CPI team that was fired by the government in 2007). Buenos Aires City uses prices collected from a sample of products in the city of Buenos Aires and follows the old INDEC methodology.[7]

A third unofficial indicator is the Provincial Index, based on CPIs from nine Argentine provinces. Whereas the official national index by INDEC was historically based only on the greater Buenos Aires area, provincial statistical agencies also collected regional price data and computed their own CPIs. The federal government pressured the provinces to manipulate or stop publishing these indexes, but those provinces that were not aligned

---

5. For example, in 2008 the government paid an interest rate of 15 percent for newly issued debt sold to the government of Venezuela.

6. PriceStats is a private company connected with the Billion Prices Project, an academic initiative based at the Massachusetts Institute of Technology (MIT) that was created in 2008 by Alberto Cavallo and Roberto Rigobon to experiment with the use of online data in the production of price indexes and other macroeconomic and international research applications. For details of the Billion Prices Project, see Cavallo and Rigobon (2016).

7. For the details, see Bevacqua and Salvatore (2009).

with the federal government continued disseminating their own unadulterated data. This index is computed as a geometric, weighted mean of nine provincial CPIs for the post-2006 period, with weights computed to maximize the correlation between the provincial aggregate and the official index (based on greater Buenos Aires) during the pre-manipulation period. Finally, the Congress Average index is an average of private inflation indicators that were widely cited in the media in 2011, after the government started to fine and prosecute economists who were publishing their own unofficial inflation estimates. Some members of Congress from the opposition political parties, who were immune from prosecution, compiled and published a monthly average of "private" estimates. Other alternative indicators also were publicized. The online appendix provides a comprehensive list, with characteristic details and methodologies.[8]

Figure 2 shows the annual inflation rate for all these unofficial indicators and the official CPI. The vertical line shows the month of the intervention at INDEC, where the official and unofficial indicators immediately diverged. All unofficial indicators showed similar results, despite differences in their data sources and methodologies. On average, the inflation rate in the unofficial indicators was approximately 10 percentage points higher than that in the official data.

### I.C. Inflation Expectations and Inflation Statistics

The surge in inflation during 2006 motivated a renewed interest in the measurement of household expectations. In August 2006, the Finance Research Center (Centro de Investigación en Finanzas) at Torcuato Di Tella University began a national household survey of inflation expectations.

In figure 3, we plot the official inflation rate, our main unofficial inflation indicator (PriceStats), and the median inflation expectation from the household survey. These monthly time series allow us to study the coevolution of available inflation indicators and of inflation expectations for seven years of uninterrupted manipulation of official statistics.

Over time, household inflation expectations aligned with the unofficial inflation level. The PriceStats index was not disseminated until March 2008, but newspapers reported other unofficial estimates before then.[9] In the online

---

8. The online appendixes for this and all other papers in this volume may be found at the *Brookings Papers* web page, www.brookings.edu/bpea, under "Past Editions."

9. An earlier version of the PriceStats index started to be published in a website called www.inflacionverdadera.com in March 2008.

Figure 2. Official Inflation and Alternative Unofficial Indicators, 2004–13[a]

Percent
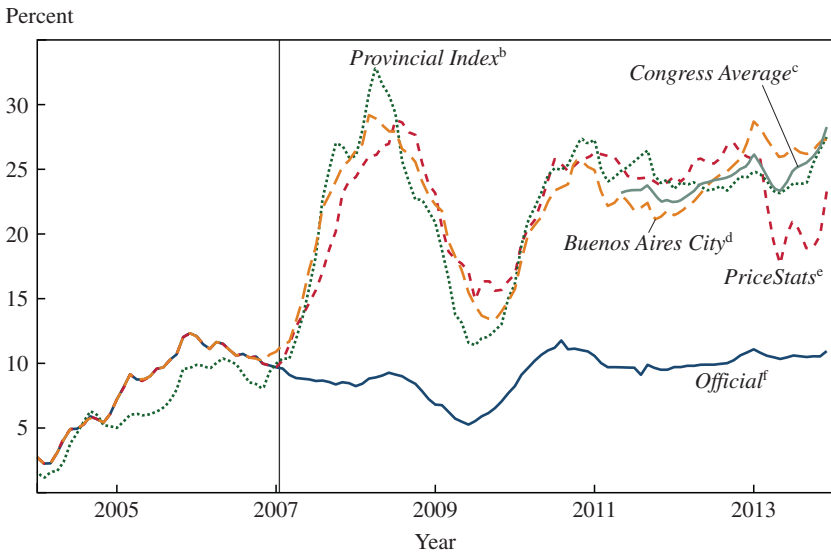
Sources: INDEC; PriceStats LLC; Cavallo and Rigobon (2016); Cosas que Pasan (CqP); Bevacqua and Salvatore (2009).

a. The vertical line represents the start of the government's intervention at INDEC in January 2007. All inflation indicators are monthly series.

b. Annual inflation rate based on a geometric, weighted mean of nine provincial statistical agencies' CPIs.

c. Annual inflation rate based on an average of several unofficial indicators compiled by Congress representatives from opposition parties.

d. Annual inflation rate compiled by the organization named Buenos Aires City from a sample of products and prices from the greater Buenos Aires metropolitan area, following INDEC's traditional methodology (the two series coincide until September 2006). According to Bevacqua and Salvatore (2009), this index uses data from the Mendoza provincial index in 2007 and from a "private consulting firm" (possibly PriceStats) in 2007. This could explain the similarities seen with the other indexes at those times.

e. Annual inflation rate compiled by PriceStats LLC based on prices from online retailers.

f. Annual inflation rate reported by INDEC based on a sample of prices from the greater Buenos Aires metropolitan area.

appendix, we plot the annual inflation rates mentioned in these newspaper articles and show that they track inflation expectations during 2007.

There is also some evidence of an asymmetric response of inflation expectations to the actual inflation rate. Two periods in particular show sticky expectations on the way down. First, from September 2008 to July 2009, when the country was experiencing the effects of the global financial crisis, the unofficial inflation rate fell by 13 percentage points, but median inflation expectations fell by only 7 percentage points. Second, from December 2012 to July 2013, due to both significant price controls and another recession, the unofficial inflation rate fell by 5 percentage points, but inflation

Figure 3. Official Statistics, Unofficial Statistics, and Inflation Expectations, 2006–13[a]
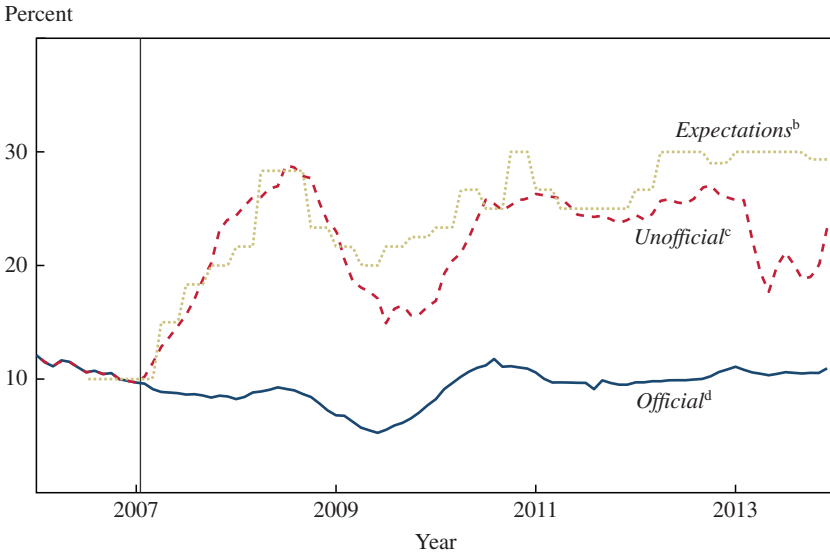
Percent

Sources: INDEC; PriceStats LLC; Cavallo and Rigobon (2016); Torcuato Di Tella University, Finance Research Center (Centro de Investigación en Finanzas), Inflation Expectations Survey (Encuesta de Expectativas de Inflación).

a. The vertical line represents the start of the government's intervention at INDEC in January 2007.

b. Quarterly averages of the monthly median of inflation expectations for the following 12 months from the Encuesta de Expectativas de Inflación. This survey collects information on the inflation expectations for the following 12 months among the general population of Argentina, based on a standard question such as, "What do you expect the annual rate of inflation will be during the next 12 months?"

c. Annual inflation rate compiled by PriceStats LLC based on prices from online retailers.

d. Annual inflation rate reported by INDEC based on a sample of prices from the greater Buenos Aires metropolitan area (monthly series).

expectations increased by 1 percentage point. We discuss this asymmetric reaction in the next section, including the possibility that statistical manipulation caused this asymmetry. The observational evidence suggests that, if anything, manipulating inflation statistics made things worse from the point of view of curbing inflation expectations.

II. Experimental Evidence

The patterns that emerge from the time series analyzed in the previous section support the hypothesis that individuals are not naive learners who accept official statistics without question. However, we cannot make causal inferences from these observational data, and it is unclear whether individuals are simply ignoring the official data or are adjusting to them in a

rational way. To address these limitations, this section develops a Bayesian learning model of inflation expectations in the presence of biased signals, and it uses experimental evidence to test some of its predictions.

II.A. A Model of Learning with Biased Statistics

For the sake of simplicity, we study the static case where the inflation rate is fixed at π_{actual} and an individual must learn about this rate of inflation indirectly from a series of signals. We also assume that price changes for each individual product in the economy are normally distributed with mean π_{actual} and variance σ^2_{actual}, and that the variance is known to the individual. Relaxing these assumptions would complicate the algebra but would not change the model's main intuition.

The individual can observe two signals based on the information about the price changes for individual products. The first signal comes from the price changes for a randomly drawn subset of N_u products, with an associated mean $\bar{u}$ and variance $\frac{1}{N_u}\sigma^2_{actual}$. This signal could be an unbiased, unofficial inflation index or could represent the information that individuals obtain by using averages of their own memories about price changes for a set of products. The second signal is the government's official statistics. We assume that the government also takes a randomly drawn subset of N_o products and computes its average price change, producing a signal with associated mean $\bar{o}$ and variance $\frac{1}{N_o}\sigma^2_{actual}$. However, the government does not report $\bar{o}$ but instead adds a bias, b_{actual}, before reporting it. In other words, the government reports $\bar{o}' = \bar{o} + b_{actual}$ instead of $\bar{o}$. Note that N_u and N_o determine the precision of the unofficial and official signals. To simplify notation, we replace $\sigma^2_u = \frac{1}{N_u}\sigma^2_{actual}$ and $\sigma^2_o = \frac{1}{N_o}\sigma^2_{actual}$.[10]

The individual has two beliefs: one about the inflation rate, π, and another about the government bias, b. We denote π_0 as the belief about the inflation rate before obtaining new information, and π_1 as the belief about the inflation rate after doing so; and b_0 and b_1 are similarly defined. The normality assumption about the distribution of price changes determines

---

10. In practice, σ^2_u and σ^2_o represent not only pure statistical errors driven by sample size but also other sources of error. For example, individuals may perceive σ^2_o to be high because they do not understand how precise these statistics are or because they do not believe that these statistics are representative of their own consumption bundle. Similarly, σ^2_u may take into account the individual's imprecision in remembering historical prices.

that the conjugate distribution for beliefs about inflation and bias is bivariate normal. For the sake of notational simplicity, we focus on the case where the prior beliefs about the inflation rate and the bias are orthogonal. As shown in the online appendix, this assumption leads to these posterior beliefs:

$$(1) \qquad \pi_1 = (1 - \omega_1 - \omega_2)\pi_0 + \omega_1\bar{u} + \omega_2(\bar{o}' - b_0)$$

$$(2) \qquad b_1 = (1 - \psi_1 - \psi_2)b_0 + \psi_1(\bar{o}' - \pi_0) + \psi_2(\bar{o}' - \bar{u}).$$

The mean posterior belief about the inflation rate, π_1, is a weighted average between the mean prior belief, π_0; the unofficial inflation rate, $\bar{u}$; and the bias-adjusted official statistics, $\bar{o}' - b_0$. The mean posterior belief about government bias, b_1, is a weighted average between the prior belief, b_0; the gap between the official statistics and the prior belief about inflation, $\bar{o}' - \pi_0$; and the gap between the official statistics and the unofficial statistics, $\bar{o}' - \bar{u}$. The parameters ω_1, ω_2, ψ_1, and ψ_2 are weights that depend on the precision of the signals and prior beliefs. Details about these weights are provided in the online appendix.

The most important prediction of this model is that a Bayesian learner is not expected to ignore biased statistics, but instead rationally adjust to the perceived bias. The following two scenarios are useful for understanding this model's intuition.

The first scenario explores how an individual who starts thinking that the government is not lying reacts to an official signal that is different from its prior. In particular, consider an individual who starts with $b_0 = 0$ and gets signals $\bar{u} = \pi_0$ (the unofficial signal equals the prior) and $\bar{o}' < \pi_0$ (the official signal is lower than the prior). The individual can attribute the low level of the official statistic to a bias, or can believe that it is driven by sampling variation. How fast would the individual learn about a bias? By making the relevant replacements in the formula given above for b_1, we get $b_1 = (\psi_1 + \psi_2)$ $(\bar{o}' - \pi_0)$. The term $\psi_1 + \psi_2$ is a set of weights that increases with the precision of both the official and unofficial signals. So, for example, if the individual perceived that there is a lot of measurement error in either of those signals, he or she would not so rapidly change his or her belief about a bias in the official data.

The second scenario explores how an individual who believes that the government is manipulating statistics reacts to the official statistics compared with the unofficial statistics. In the next sections, we study this scenario by means of a series of information experiments during the period of manipulated statistics. Consider an individual who starts out thinking

that the government biases the inflation statistics downward; that is, $b_0 < 0$. How does the individual react to official statistics compared with unofficial statistics? From the formula for π_1, it follows that, qualitatively, the individual reacts to $\bar{o}'$ in the same way as he or she reacts to $\bar{u}$, but with the exception that first it subtracts from $\bar{o}'$ the ex ante perceived bias; that is, it uses $\bar{o}' - b_0$ instead of $\bar{o}'$. So if the individual believes that the bias is $b_0 = -10$ percent, then he or she should react qualitatively to the signal $\bar{o}' = 10$ percent in the same way that he or she reacts to $\bar{u} = 20$ percent. These reactions are qualitatively the same but potentially quantitatively very different, because the weights ω_1 and ω_2 could be potentially very different. For instance, these weights would be very different if there is a large difference in precision between the unofficial and official statistics, $\dfrac{1}{\sigma_u^2}$ and $\dfrac{1}{\sigma_o^2}$. However, if these statistics are similarly precise, then we would expect a reaction that is quantitatively very similar.[11]

II.B. The Experimental Design

The survey experiment in this section is related to a group of recent studies on how individuals learn about inflation and how they form their inflation expectations (Roos and Schmidt 2012; Armantier and others 2016; Cavallo, Cruces, and Perez-Truglia 2014). We first collect background information about respondents (see the online appendix for a translation of the questionnaire). We then randomly assign subjects to different groups. The control group receives no information. The other informational treatments receive either official or unofficial statistics about inflation rates for the previous 12 months. After the information provision, we elicit subjects' inflation perceptions and expectations and measure how a particular signal about inflation affects the distribution of inflation perceptions and expectations.

The inflation perceptions correspond to a question about current inflation levels (that is, the respondent's perception of the annual inflation rate during the previous 12 months). We also include a question about the respondents' subjective assessments of their confidence in their answers, measured on a scale from 1 ("not at all confident") to 4 ("very confident").

---

11. Note that even if the precision of unofficial and official statistics were exactly the same, $\dfrac{1}{\sigma_u^2} = \dfrac{1}{\sigma_o^2}$, we would still have $\omega_1 > \omega_2$, and thus the individual would react more to $\bar{u}$ than to $\bar{o}' - b_0$. The reason is that, when doing the correction $\bar{o}' - b_0$, the individual is using b_0, which has some uncertainty of its own.

The subject's inflation expectations correspond to the expected inflation rate during the following 12 months. Argentina's economic history implies that the general public understands the meaning of the word "inflation," which is discussed routinely in the media.[12] Thus, when eliciting inflation perceptions and expectations, we state our question using the word "inflation," instead of referring to "changes of prices in general" or other indirect references to inflation that are commonly used in U.S. surveys and in other low-inflation countries.[13]

The message about inflation provided in the survey experiment has the following structure:

> According to *[SOURCE]*, the annual inflation rate with respect to a year ago was approximately *[X percent]*.

In this message, *[SOURCE]* could be "one of the official indicators published by INDEC" (that is, official statistics) or "one of the unofficial indicators published by consulting firms, analysts, and research centers" (that is, unofficial statistics). The large variety of inflation indicators allows us to cross-randomize two features of this message in a nondeceptive way: the source of the inflation statistics (official or unofficial), and the inflation rate (10, 20, or 30 percent).

For the official statistics, the first indicator produced by INDEC is the CPI, which is the most common inflation index in the world. This was the main indicator targeted for government manipulation. At the time of our experiment, the annual growth rate of the official CPI was approximately 10 percent. INDEC also computed other indicators that reflected different inflation levels. One was the GDP deflator, which is sometimes used as a measure of inflation and which closely tracked the CPI in Argentina before 2007. At the time of the experiment, the GDP deflator was close to

12. Moreover, the previous rounds of the online opinion poll, into which we built our survey experiment, used the wording in terms of inflation, as did other sources for inflation expectations, such as the Inflation Expectations Survey (Encuesta de Expectativas de Inflación) conducted by the Finance Research Center (Centro de Investigación en Finanzas) at Torcuato Di Tella University. Also, we did not provide any incentives for respondents to answer accurately, such as prizes for guessing the right figures. As shown by Armantier and others (2011) in the context of similar studies, there is a significant correlation between incentivized and nonincentivized responses on inflation expectations.

13. For instance, the University of Michigan's Survey of Consumers elicits inflation expectations by means of the following questions: "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?"—with three options, "go up," "stay the same," and "go down"—and then asks, "By about what percent do you expect prices to change, on average, during the next 12 months?" with an open numerical answer.

20 percent. The government could not allow the GDP deflator to be as low as the CPI (10 percent), because that would have implied an implausibly high real GDP growth rate (more than 15 percent). We also use a third statistic compiled by INDEC and routinely used by local economists as an inflation proxy: the rate of growth of nominal wages. At the time of our survey, this measure indicated an annual inflation rate close to 30 percent. We followed a similar strategy to exploit the variation in unofficial statistics. We chose one index published by an unofficial source that indicated an inflation rate close to 20 percent, and another index that indicated an inflation rate close to 30 percent. A third unofficial index, published by a think tank with close ties to the government, indicated an inflation rate close to 10 percent.

We emphasize that we did not deceive the experimental subjects; we conveyed information from the public discussion in Argentina at that time. We did not claim that the information provided was true or false; nor did we endorse or disavow, implicitly or explicitly, any of the sources. We merely stated that, according to a given source, the level of annual inflation was estimated to be *X* percent. In any case, because individual judgment about the information can vary depending on the source, we included a debriefing statement at the end of the survey. In this statement, we disclosed that the information about inflation that we provided was randomly selected from six possible messages, and we included a detailed source and explanation for each message. We presented the same statement to all subjects, irrespective of their assigned treatment group. Our purpose was that the subjects should leave the experiment with more information than what they had at the beginning of the experiment.

II.C. The Subject Pool and Experimental Results

The sample and survey are based on the ones used by an established public opinion research firm that carries out a quarterly online survey of adults in Argentina, which has had a stable set of questions since 2011. The experiments were conducted in December 2012, while the government was still manipulating official statistics, and almost six years after the government started to do this manipulation. We slightly modified the standard format of this public opinion survey to fit our experimental design. In particular, our survey experiment was included early in the questionnaire's flow, after which it continued with the usual set of questions about politics, politicians, and public affairs. These questions are not used to determine outcomes in our analysis, although we use some of them for descriptive purposes and to verify the balance between treatment groups. The respondents were assigned to the control group with a probability of 22.6 percent, and to each of the

Table 1. Descriptive Statistics for the Online Opinion Survey Sample and Argentina's Total Population[a]

	Share female	Mean age	Share living in greater Buenos Aires	Share with college degree[b]	Share who voted for Kirchner[c]
Authors' sample	0.570	41.1	0.677	0.607	0.242
Argentina's population	0.528	44.9	0.363	0.156	0.541

Sources: Authors' online opinion survey (see text); INDEC, Annual Survey of Urban Households (Encuesta Anual de Hogares Urbanos); the 2011 presidential election results.

a. All statistics are based on individuals ages 20 or older. The sample size for the authors' online opinion survey is 3,138.

b. Share of respondents who have completed college or another form of postsecondary education.

c. Share of respondents who reported voting for Cristina Kirchner in the 2011 presidential election.

treatment groups with a probability of 12.9 percent. The final sample on which the following analysis is based consists of all the respondents who completed the questions on inflation perceptions and inflation expectations, yielding a final sample of 3,138 observations.[14]

Table 1 presents summary statistics about the demographics of the sample, along with the corresponding indicators for the general population. This sample is not representative of the general Argentine population; though it is roughly similar in age and gender composition, our sample is substantially more educated and richer than average. Nevertheless, the qualitative results are similar if we reweight the observations to match the distribution of characteristics at the national level (not reported).

Table 2 presents descriptive statistics for all the variables used in the analysis, including pretreatment and posttreatment variables, for the control group and for each of the treatment groups. The last column reports

14. A small but nonnegligible number of individuals abandoned the questionnaire after the information treatment and the question on inflation perceptions, and before reporting their inflation expectations (105 out of 3,243, or 3.24 percent of the original sample). Although this type of attrition also occurred in previous rounds of the opinion poll (for instance, the sample had a dropout rate of 5.8 percent for the June 2012 round), in this case this might be a concern if the attrition were due to (and correlated with) the information treatments, because this could introduce biases in the experiment and complicate the interpretation of the treatment effects. For instance, government supporters who believe that inflation is low may have abandoned the experiment because they did not like to see information from unofficial sources reporting high inflation levels (the opposite situation could arise with respondents opposed to the government and with high inflation perceptions). However, this does not seem to be a concern in practice, because we cannot reject the null hypothesis of equal attrition across treatment groups (p value = 0.79).

Table 2. Average Posttreatment and Pretreatment Responses, by Treatment Group[a]

	Control	Official– 10 percent	Official– 20 percent	Official– 30 percent	Unofficial– 10 percent	Unofficial– 20 percent	Unofficial– 30 percent	p value[b]
Posttreatment								
Inflation perception, previous 12 months	28.31 (0.591)	28.55 (0.812)	33.58 (0.809)	42.10 (0.815)	26.37 (0.805)	28.89 (0.809)	34.78 (0.816)	<0.01
Confidence in inflation perception[c]	−0.0903 (0.0319)	0.119 (0.0438)	−0.0144 (0.0437)	−0.119 (0.0440)	0.0122 (0.0435)	0.0804 (0.0437)	0.0912 (0.0441)	<0.01
Inflation expectation, following 12 months	28.20 (0.595)	28.22 (0.817)	33.32 (0.814)	39.24 (0.820)	26.29 (0.810)	28.62 (0.814)	33.99 (0.821)	<0.01
Pretreatment								
Share female	0.560 (0.0181)	0.539 (0.0248)	0.563 (0.0247)	0.586 (0.0249)	0.601 (0.0246)	0.545 (0.0247)	0.608 (0.0250)	0.28
Age	41.11 (0.390)	41.06 (0.536)	40.64 (0.534)	40.79 (0.538)	40.93 (0.532)	41.37 (0.534)	41.16 (0.539)	0.97
Share with college degree[d]	0.633 (0.0178)	0.610 (0.0245)	0.650 (0.0244)	0.581 (0.0246)	0.572 (0.0243)	0.600 (0.0244)	0.578 (0.0246)	0.12
Own economic situation is better[e]	0.261 (0.0155)	0.237 (0.0213)	0.228 (0.0212)	0.203 (0.0214)	0.233 (0.0211)	0.257 (0.0212)	0.211 (0.0214)	0.26
No. of observations	750	397	400	394	404	400	393	

Source: Authors' online opinion survey (see text).

a. Each cell represents the mean of each of the row variables for the corresponding control and treatment groups in the column headers. Treatment groups are broken down by respondents' source of inflation statistics (official or unofficial), and the source's reported inflation rate (10, 20, or 30 percent). Standard errors are in parentheses.

b. Reports the *p* value of a balance test in which the null hypothesis is that the mean of each variable is equal between all seven experimental groups (the control group and the six treatment groups).

c. Represents the respondent's own confidence in his or her response to the perceptions question on a scale of 1 ("not confident at all") to 4 ("very confident").

d. Share of respondents who have completed college or another form of postsecondary education.

e. Share of respondents who reported that their current economic situation was better compared with 12 months earlier.
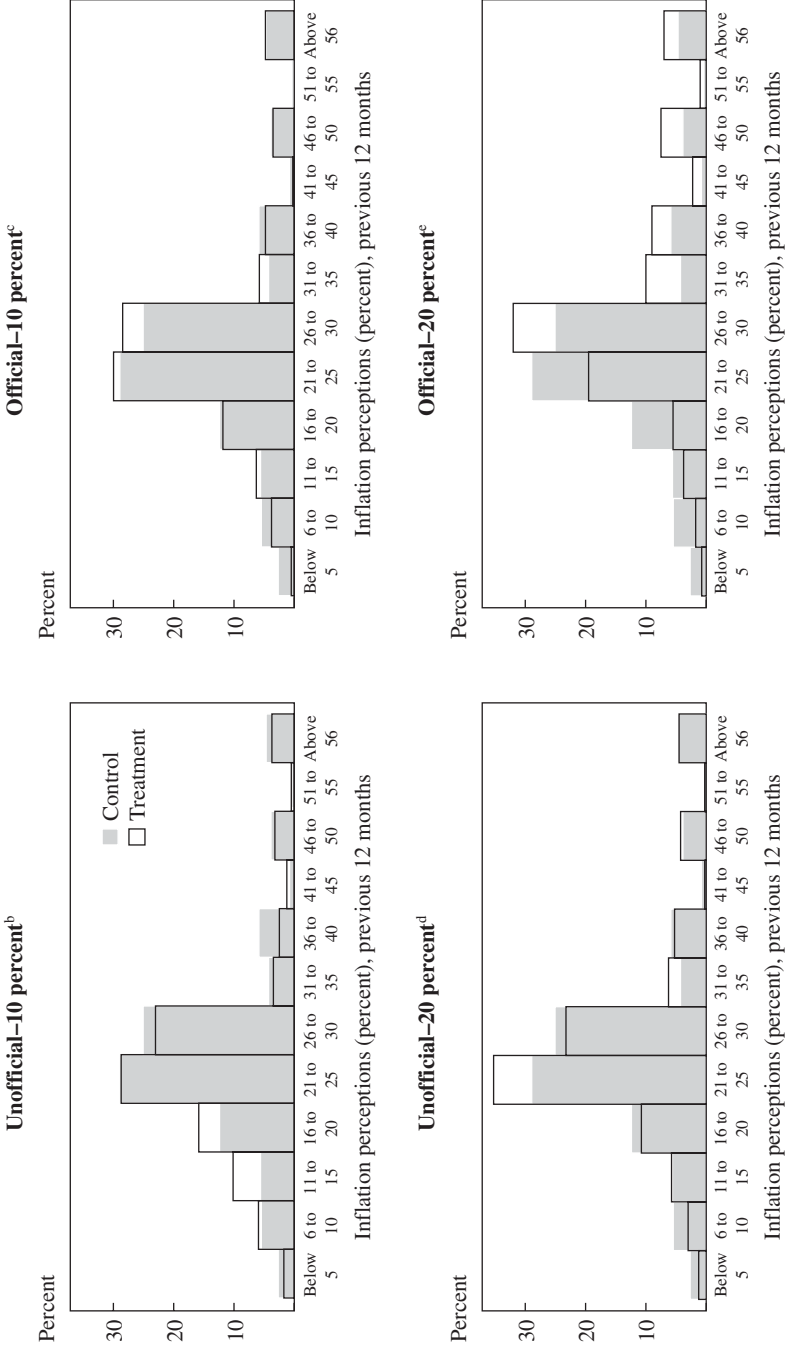
the p value of a test in which the null hypothesis is that the mean of each variable is equal in all seven experimental groups. As expected, these tests are not rejected for any of the pretreatment variables, suggesting that the randomization was indeed balanced. The top panel shows the posttreatment variables: inflation perceptions, confidence in these perceptions, and inflation expectations. We discuss this impact in more detail below. Additionally, the main experimental results are presented in two complementary ways. In figure 4 just below, we show the distribution of inflation perceptions in the control group compared with that of each of the six other informational treatments. And figure 6 below summarizes the effects of the six informational treatments on the means of various posttreatment outcomes relative to the control group.
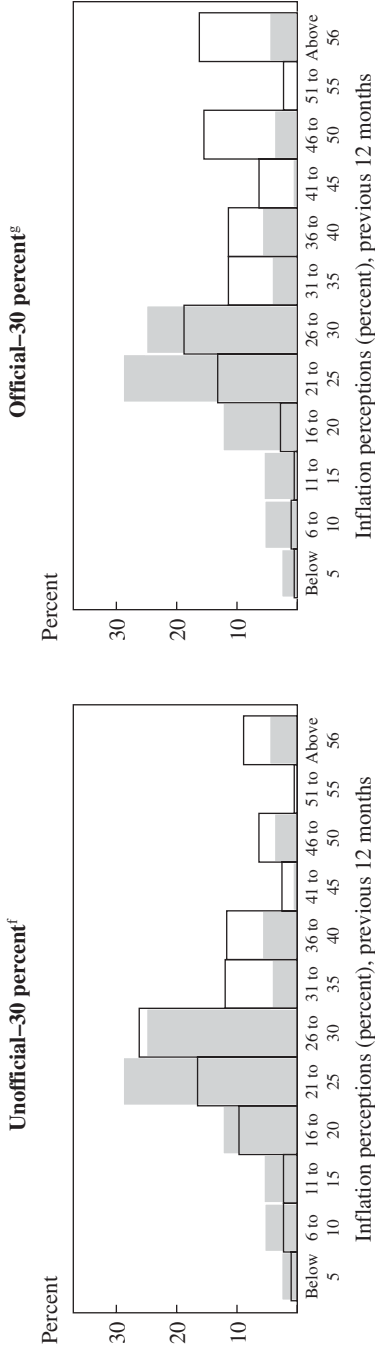
Our benchmark results in this section are based on the effects on inflation perceptions, which are directly related to the information signals provided by the experiment (past 12 months' data), but they are equivalent to those that are obtained from inflation expectations, as we discuss below.

We begin by measuring whether individuals' inflation perceptions were influenced by the messages with unofficial statistics. The left-side panels of figure 4 show the distribution of perceptions in the control group and each of the messages about the unofficial statistics. The data suggest that individuals did not ignore this information; compared with individuals who were told that inflation according to official statistics was 20 percent, individuals who were told that official statistics were lower (10 percent) reported lower inflation perceptions, and individuals who were told that official statistics indicated higher inflation (30 percent) reported higher perceptions. We conducted the Epps–Singleton (ES) two-sample test using the empirical characteristic function, a version of the Kolmogorov–Smirnov test of equality of distributions valid for discrete data (Goerg and Kaiser 2009). According to the ES test, all these pairwise differences are statistically significant at the 1 percent level. Additionally, these differences are economically significant. In sum, individuals seemed eager to learn from unofficial sources.

The first hypothesis to test is whether individuals reacted at all to the messages about official statistics. The right-side panels of figure 4 show the distribution of perceptions in the control group and each of the messages about the official statistics. In comparison with individuals who were told that inflation according to official statistics was 20 percent, individuals who were told that official statistics were lower (10 percent) reported lower inflation perceptions, and individuals who were told that official statistics indicated higher inflation (30 percent) reported higher inflation perceptions. According to the ES test, these pairwise differences in distributions

Figure 4. Histograms Comparing the Distribution of Inflation Perceptions between Treatment and Control Groups[a]

Unofficial–10 percent[b]

Official–10 percent[c]

Unofficial–20 percent[d]

Official–20 percent[e]

Unofficial–30 percent[f]

Percent

Inflation perceptions (percent), previous 12 months

Official–30 percent[g]

Percent

Inflation perceptions (percent), previous 12 months

Source: Authors' online opinion survey (see text).

a. Panel captions indicate the treatment group. Respondents in the treatment group were given nondeceptive information about inflation estimates from unofficial (left-side panels) or official (right-side panels) sources. Respondents in the control group were not given any information about inflation statistics. There were 750 observations in the control group. The histograms are censored at 5 percent and 56 percent (inclusive), but these bins represent the cumulative observations below 5 percent and above 56 percent, respectively. The Epps–Singleton p value is a measure of the equality of the two distributions (Goerg and Kaiser 2009).

b. Respondents were told that the annual inflation rate with respect to a year ago was approximately 10 percent according to an unofficial source. There were 404 observations in the treatment group. Epps–Singleton p value < 0.01.

c. Respondents were told that the annual inflation rate with respect to a year ago was approximately 10 percent according to an official source. There were 397 observations in the treatment group. Epps–Singleton p value < 0.01.

d. Respondents were told that the annual inflation rate with respect to a year ago was approximately 20 percent according to an unofficial source. There were 400 observations in the treatment group. Epps–Singleton p value = 0.06.

e. Respondents were told that the annual inflation rate with respect to a year ago was approximately 20 percent according to an official source. There were 400 observations in the treatment group. Epps–Singleton p value < 0.01.

f. Respondents were told that the annual inflation rate with respect to a year ago was approximately 30 percent according to an unofficial source. There were 393 observations in the treatment group. Epps–Singleton p value < 0.01.

g. Respondents were told that the annual inflation rate with respect to a year ago was approximately 30 percent according to an official source. There were 394 observations in the treatment group. Epps–Singleton p value < 0.01.

are statistically significant at the 1 percent level. These differences are economically significant; for instance, the mean of inflation perceptions is 28.5 percent for the official–10 percent group, 33.6 percent for the official–20 percent group, and 42.1 percent for the official–30 percent group.

The second hypothesis to test is the naive learning hypothesis, according to which households react to information on a given level of inflation from an official source in the same way that they would react to the same figure if it were coming from unofficial sources. The data strongly reject this hypothesis; the ES test indicates that the difference between the distribution of inflation perceptions across individuals given messages official–10 percent and unofficial–10 percent is significant at the 1 percent level; the same is true when comparing the distribution of perceptions for the official–20 percent and unofficial–20 percent groups, and for the official–30 percent and unofficial–30 percent groups. These differences are not only statistically but also economically significant—for instance, compared with the unofficial–10 percent, the message official–10 percent created inflation perceptions that were 2.1 percentage points higher; compared with the unofficial–20 percent, the message official–20 percent created inflation perceptions that were 4.7 percentage points higher; and compared with the unofficial–30 percent, the message official–30 percent created inflation perceptions that were 7.3 percentage points higher.

The third hypothesis to test is the rational learning hypothesis. A plausible heuristic for the period under study is that official inflation rates were systematically 10 percentage points below those from unofficial sources. On the basis of this approximation, the learning model predicts that individuals should react to information conveying an official inflation level of X percent in the same way as they would react to information from unofficial sources conveying a level of inflation of $X - 10$ percent. The results from our experiment are consistent with this hypothesis; we cannot reject the null hypothesis that the distributions of inflation perceptions are equal between individuals in the groups official–10 percent and unofficial–20 percent (ES p value = 0.91), and we cannot reject the null hypothesis that the distributions of inflation perceptions are equal between individuals with the messages official–20 percent and unofficial–30 percent (ES p value = 0.61). These differences are not only statistically insignificant; they are also economically small. For instance, the difference in the mean perceived inflation is only 0.34 percentage point between the official–10 percent and unofficial–20 percent groups, and 1.2 percentage points between the official–20 percent and unofficial–30 percent groups.

The experiment also allows us to further explore the seemingly asymmetric relationship between perceived and actual inflation suggested by the analysis of the nonexperimental time series data discussed in the previous sections. According to table 2, increasing the unofficial inflation rate shown to the subject from 10 percent to 20 percent also increased the mean perceived inflation rate by 2.52 percentage points (p value < 0.05). Instead, increasing the unofficial inflation rate shown to the subject from 20 percent to 30 percent, which is also an increase of 10 percentage points, raised the mean perceived inflation rate by 5.89 percentage points (p value < 0.01). That is, the effect of going from 20 percent to 30 percent is almost twice the size of the effect of going from 10 percent to 20 percent, and this difference is statistically significant (p value = 0.072). This evidence suggests that individuals were twice as reactive to information about higher inflation than to information about lower inflation. The results are similar for those who received signals from official sources; inflation perceptions are 5 percentage points lower in the official–10 percent group than in the official–20 percent group, and they are 8.5 percentage points lower in the official–20 percent compared with the official–30 percent group. The difference between these two effects is statistically significant (p value = 0.082).

This asymmetry in the experimental effects is consistent with the observational evidence analyzed in section I.C. In a related paper (Cavallo, Cruces, and Perez-Truglia 2014) we find no evidence of asymmetry in the reaction to information about the price changes of supermarket products, which suggests that the asymmetry is particular to inflation statistics. Furthermore, it is possible that this asymmetry was generated by the manipulation of official statistics, although we do not have experimental evidence preceding the period of manipulation to test this hypothesis directly. According to this conjecture, when individuals observe an inflation statistic indicating that inflation is below their prior belief, they suspect that such a statistic may be subject to government bias, even if it comes from a nongovernmental agency, and thus they put less weight on learning from this piece of information. On the contrary, because the government's goal is to reduce inflation perceptions, individuals do not worry that a statistic has a government-induced bias when it suggests that inflation is above their prior belief, thus explaining the asymmetry.[15]

15. Note that this conjecture predicts that the asymmetry would be reversed if the government were interested in increasing rather than reducing inflation perceptions.

Figure 5. Past Inflation Perceptions and Future Inflation Expectations
for the Control Group[a]

Inflation expectations (percent), next 12 months

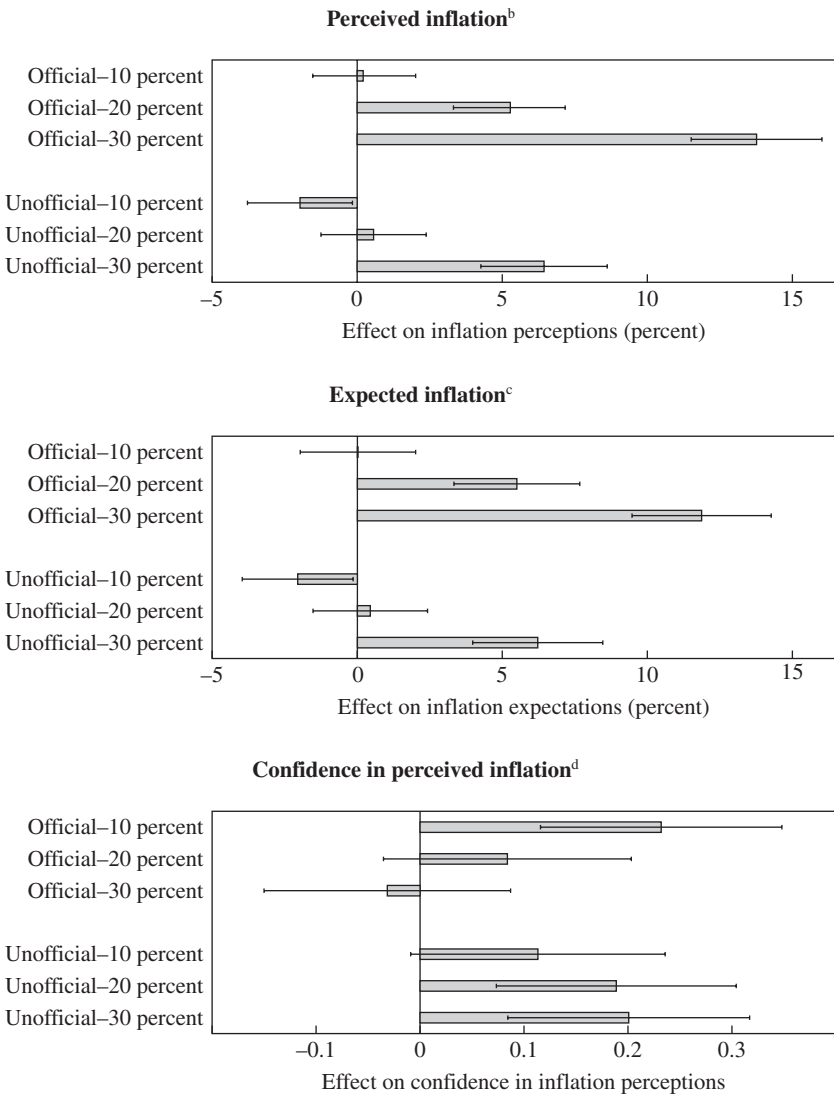Inflation perceptions (percent), previous 12 months
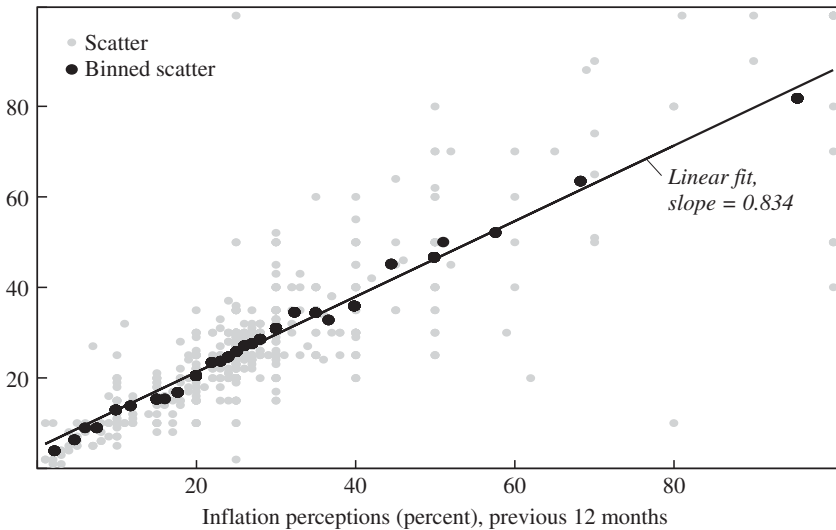
Source: Authors' online opinion survey (see text).
  a. There are 777 observations in the control group. Respondents in the control group were not given any information about inflation statistics.

Perceptions about past inflation are a key input in the formation of inflation expectations (Jonung 1981; Cavallo, Cruces, and Perez-Truglia 2014). For instance, figure 5 shows a binned scatter plot of the relationship between inflation perceptions and inflation expectations in our experimental sample, based on the subsample of respondents in the control group. This figure shows that, as expected, these two variables are strongly associated. If our information treatments affected inflation perceptions, we would expect to observe a similar effect on inflation expectations. The comparison of effects between the top and middle panels of figure 6 indicates that the effects on perceived inflation (top panel) were very similar to the effects on inflation expectations (middle panel). For instance, compared with the control group, the unofficial–30 percent message increased inflation perceptions by 6.47 percentage points and inflation expectations by 5.79 percentage points, with the difference between the two effects being close to zero and statistically insignificant.

Additionally, the bottom panel of figure 6 shows the effect of the different treatments on the subjects' reported confidence in their inflation

Figure 6. Treatment Effects on Inflation Perceptions and Expectations[a]

Perceived inflation[b]

Effect on inflation perceptions (percent)

Expected inflation[c]

Effect on inflation expectations (percent)

Confidence in perceived inflation[d]

Effect on confidence in inflation perceptions

Source: Authors' online opinion survey (see text).

a. Each bar represents the point estimate of the effect of the specific subtreatment compared with the control group. Respondents in the control group were not given any information about inflation statistics. See notes to figure 4 for sample sizes. Robust standard errors are reported.

b. Mean effect on the treatment groups' inflation perceptions for the previous 12 months compared to the control group. The mean perceived inflation rate for the control group is 28.31 percent.

c. Mean effect on the treatment groups' inflation expectations for the following 12 months compared with the control group. The mean expected inflation rate for the control group is 28.2 percent.

d. Mean effect on the treatment groups' confidence in perceived inflation compared with the control group, which received no information about inflation statistics. A respondent's confidence is reported on a scale from 1 ("not confident at all") to 4 ("very confident").

perceptions, coded from 1 ("not confident at all") to 4 ("very confident"). Bayesian learning predicts that, after observing a useful signal, individuals should usually be more confident about their posterior beliefs relative to the counterfactual with no information. The bottom panel of figure 6 indicates that, as expected, several of the informational treatments significantly increased the subjects' reported confidence in their inflation perceptions. Interestingly, the effect on reported confidence is almost twice as large for unofficial statistics (1.15) than for official statistics (0.86; p value of the difference = 0.066), suggesting that the information from unofficial sources was on average more useful.

Note that part of the reaction to the information provided in the experiment may be spurious, for example, due to numerical anchoring or experimenter-demand effect. This is an important concern because, as shown in Cavallo, Cruces, and Perez-Truglia (2014), only about half the reaction to the inflation information provided in similar survey experiments is genuine rather than spurious. In any case, according to the results from our earlier paper, it is plausible that the effects identified in this paper are quantitatively different, but qualitatively robust to spurious learning.

III. The 2013 Price Controls

Studies suggest that individuals form inflation expectations using information from their own consumer experiences, such as their memories of the prices of supermarket products (Bates and Gabor 1986; Bruine de Bruin, van der Klaauw, and Topa 2011; Cavallo, Cruces, and Perez-Truglia 2014; Coibion and Gorodnichenko 2015; Kumar and others 2015). This implies that the government could try to manipulate inflation expectations by changing the actual prices of salient products. This section discusses evidence about a 2013 effort by the government to control the prices of supermarket products.

In February 2013, the government of Argentina significantly extended its policy of price controls on retail products. These "price agreements" with big companies and large supermarket chains were temporarily applied to hundreds of products in carefully selected categories. The government targeted goods that had a significant weight in the CPI basket, and it focused on brands and retailers with large market shares. To enforce these price controls, the government publicly asked its supporters to help monitor prices. The program, which was called Precios Cuidados (Protected Prices), was widely advertised and discussed in the media. Although there were some problems with its implementation, most of the goods included

in the agreements were available for sale at the agreed-on prices. It is possible that the government hoped that by controlling the prices of key individual goods, it could influence inflation expectations. Consistent with this interpretation, the finance minister repeatedly mentioned that the price controls were meant to "provide predictability to the economy."[16]

The inflation rate did temporarily fall; according to the unofficial statistic of PriceStats shown in figure 3 above, the annual inflation rate fell from 25.8 percent in January 2013 to 17.7 percent in May 2013. This is not surprising, given that the PriceStats index draws its data mostly from large multichannel retailers (which sell both online and offline), where the government was focusing its price control efforts. Even though the inflation rate did fall, figure 3 suggests that there was no effect on inflation expectations, which remained stable, at near 30 percent.

There are several possible explanations for this lack of an effect on inflation expectations. One possibility is that people knew the effect would be temporary, so expectations about future inflation were not affected. This probably played an important role, but we do not have a way to test it. Another possibility is that even though people experienced more stable prices for some goods, this information did not affect their perceptions of prices. To test this, we ran a consumer-intercept survey at the front doors of four branches of one of the largest supermarket chains in the city of Buenos Aires, during the period of the price controls. The subject pool consisted of supermarket customers who, having just made a purchase, were invited to participate in a short survey for an academic study. Using hand-held scanners, our interviewers scanned respondents' receipts from their supermarket purchases, which contained product identifiers that were matched to a database of scraped online data from the same supermarket.[17] After their receipts were scanned, respondents were asked about the current prices of the products they had just purchased, and the corresponding prices as of 12 months before.

The top panel of figure 7 depicts the distribution of actual price changes for products with controlled prices, and for those with no controls. The figure shows that products with controlled prices did have a substantially lower inflation rate—with an average change of 1.9 percent, compared with 21.7 percent for noncontrolled products. However, the bottom panel of figure 7 shows that the program was not effective in changing individual

16. For example, see https://www.scribd.com/doc/312284129/Precios-Cuidados.
17. Prices were scraped from the websites of the supermarket by the Billion Prices Project at MIT. See Cavallo (2013) for details.

Figure 7. Actual and Remembered Price Changes for Products with Government-Controlled Prices[a]

Actual price changes

Remembered price changes

perceptions of price changes; the distribution of remembered price changes (as reported by the consumers) are very similar, and statistically indistinguishable, between controlled and noncontrolled products.[18] In any case, just as happens with the manipulation of the aggregate official index, price controls did not seem to be an effective way to influence inflation expectations, at least not in the short term.

IV. Conclusions

To understand how households learn from potentially biased statistics, we utilize data from a natural experiment and a survey experiment based on the period when the Argentine government manipulated inflation statistics. We find that consumers are sophisticated users of information. Rather than simply ignoring biased statistics or accepting them as unbiased, individuals can effectively adjust for the perceived bias using other available information. Furthermore, the publication of biased statistics may have led to an asymmetric reaction to inflation signals, even the unbiased ones, with expectations changing more when inflation rises than when it falls. The government's attempt to manipulate inflation expectations seems to have been ineffective and plausibly counterproductive.

These lessons are useful for understanding the formation of inflation expectations in less extreme contexts than Argentina, such as in the United States and Europe, where experts may believe that statistics are unbiased but the general population does not. For example, using data on a survey of U.S. households reported in Cavallo, Cruces, and Perez-Truglia (2014), we find that 32 percent of respondents did not trust the official inflation data and had inflation expectations that were significantly higher than the rest. The average inflation expectation for the group that did not trust the official statistics was 6.36 (standard error = 7.19), compared with an average of 4.22 (standard error = 4.26) in the rest of the sample (p value of the difference < 0.01). Our study suggests that the difference could be driven by the way individuals adjust for perceived biases in the official data. One policy implication is that governments should focus on providing information and make efforts to reduce the perception of a potential

---

18. It is important to note that individuals were roughly accurate in remembering the current prices of the products that they just purchased, and that—even among the noncontrolled goods—the remembered price changes had a large and systematic upward bias (results not reported).

bias. This lesson may be relevant for inflation statistics as well as other governmental statistics.

For policymakers, our results are useful for better understanding the process of belief formation in contexts of increasing inflation. In particular, as the Federal Reserve and other counties' central banks return to a more neutral monetary policy, these results, together with our other research, imply that the process whereby inflation expectations are formed may quickly change. As inflation rises, individuals devote more attention to inflation and adjust their expectations accordingly. Inflation expectations that may seem well anchored at low levels of inflation can react to information quickly and in a sophisticated way as inflation rises and becomes more important for decisionmaking processes.

Ultimately, more empirical evidence is needed to understand how inflation and other expectations are formed. Experimental evidence, in particular, can shed light on the questions raised by a large body of theoretical literature on models of expectation formation—such as adaptive, rational, and natural expectations (Fuster, Laibson, and Mendel 2010); and diagnostic expectations (Bordalo, Gennaioli, and Shleifer 2015). In particular, the evidence from Argentina suggests that perceived biases in the signals are capable of creating asymmetric responses in expectations. Future research should try to explain why this happens and to clarify the circumstances under which individuals tend to overreact or underreact to information.

Finally, in this paper, we focus on the effect of manipulating statistics on average inflation perceptions. In reality, the manipulation most likely affects other perceptions, such as uncertainty about inflation, which can have real effects on the economy. In other words, even if it was unsuccessful in reducing the average perceived and expected inflation rates, the manipulation of statistics may have increased individual subjective uncertainty about inflation. Although we think it is unlikely, it is still possible that one of the government's goals behind the manipulation was precisely to obfuscate beliefs about inflation. Thus, more research is needed to understand the effects of potentially biased statistics on subjective uncertainty.

References

Armantier, Olivier, Wändi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2011. "Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?" Staff Report no. 509. Federal Reserve Bank of New York.

Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2016. "The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment." *Review of Economics and Statistics* 98, no. 3.

Bachmann, Rüdiger, Tim O. Berg, and Eric R. Sims. 2015. "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence." *American Economic Journal: Economic Policy* 7, no. 1: 1–35.

Badarinza, Cristian, and Marco Buchmann. 2009. "Inflation Perceptions and Expectations in the Euro Area: The Role of News." Working Paper no. 1088. European Central Bank.

Bates, John M., and André Gabor. 1986. "Price Perception in Creeping Inflation: Report on an Enquiry." *Journal of Economic Psychology* 7, no. 3: 291–314.

Bernanke, Ben S. 2007. "Inflation Expectations and Inflation Forecasting." Speech given at the Monetary Economics Workshop, National Bureau of Economic Research Summer Institute, Cambridge, Mass., July 10. http://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm

Bevacqua, Graciela, and Nicolás Salvatore. 2009. "Argentina: La Reconstrucción de la Serie de Inflación Minorista—El IPC City" [Argentina: Reconstruction of the Retail Inflation Series—The City CPI]. Unpublished paper. https://www.scribd.com/doc/312284130/Buenos-Aires-City

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2015. "Diagnostic Expectations and Credit Cycles." Working paper. http://scholar.harvard.edu/shleifer/publications/diagnostic-expectations-and-credit-cycles

Bruine de Bruin, Wändi, Wilbert van der Klaauw, and Giorgio Topa. 2011. "Expectations of Inflation: The Biasing Effect of Thoughts about Specific Prices." *Journal of Economic Psychology* 32, no. 5: 834–45.

Camacho, Maximo, Marcos Dal Bianco, and Jaime Martinez-Martin. 2015. "Toward a More Reliable Picture of the Economic Activity: An Application to Argentina." *Economics Letters* 132: 129–32.

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.

Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics* 118, no. 1: 269–98.

Cavallo, Alberto. 2013. "Online and Official Price Indexes: Measuring Argentina's Inflation." *Journal of Monetary Economics* 60, no. 2: 152–65.

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2014. "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Field Experi-

ments." Working Paper no. 20576. Cambridge, Mass.: National Bureau of Economic Research.

Cavallo, Alberto, and Roberto Rigobon. 2016. "The Billion Prices Project: Using Online Prices for Measurement and Research." *Journal of Economic Perspectives* 30, no. 2: 151–78.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7, no. 1: 197–232.

Curtin, Richard. 2009. "What U.S. Consumers Know about the Economy: The Impact of Economic Crisis on Knowledge." Paper presented at 3rd World Forum on Statistics, Knowledge, and Policy, Organization for Economic Cooperation and Development, Busan, South Korea, October 27–30. http://www.oecd.org/site/progresskorea/43964049.pdf

D'Acunto, Francesco, Daniel Hoang, and Michael Weber. 2016. "Unconventional Fiscal Policy, Inflation Expectations, and Consumption Expenditure." Working paper. http://faculty.chicagobooth.edu/michael.weber/research/pdf/inflationExpectations.pdf

Dräger, Lena. 2015. "Inflation Perceptions and Expectations in Sweden: Are Media Reports the Missing Link?" *Oxford Bulletin of Economics and Statistics* 77, no. 5: 681–700.

European Commission. 2010. "Europeans' Knowledge of Economic Indicators." *Special Eurobarometer* 323, wave 72.1. Brussels: TNS Opinion & Social. http://ec.europa.eu/public_opinion/archives/ebs/ebs_323_en.pdf

Fuster, Andreas, David Laibson, and Brock Mendel. 2010. "Natural Expectations and Macroeconomic Fluctuations." *Journal of Economic Perspectives* 24, no. 4: 67–84.

Goerg, Sebastian J., and Johannes Kaiser. 2009. "Nonparametric Testing of Distributions—The Epps–Singleton Two-Sample Test Using the Empirical Characteristic Function." *Stata Journal* 9, no. 3: 454–65.

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

Jonung, Lars. 1981. "Perceived and Expected Rates of Inflation in Sweden." *American Economic Review* 71, no. 5: 961–68.

Kumar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko. 2015. "Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand." *Brookings Papers on Economic Activity,* Fall: 151–208.

Lamla, Michael J., and Sarah M. Lein. 2008. "The Role of Media for Consumers' Inflation Expectation Formation." Working Paper no. 201. Zurich: KOF Swiss Economic Institute.

Levy Yeyati, Eduardo, and Marcos Novaro. 2013. *Vamos por Todo: Las 10 Decisiones más Polémicas del Modelo* [Taking It All: The 10 Most Controversial Decisions of the Model]. Buenos Aires: Random House Mondadori.

Malmendier, Ulrike, and Stefan Nagel. 2016. "Learning from Inflation Experiences." *Quarterly Journal of Economics* 131, no. 1: 53–87.

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117, no. 4: 1295–328.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2004. "Disagreement about Inflation Expectations." *NBER Macroeconomics Annual* 18: 209–48.

Michalski, Tomasz, and Gilles Stoltz. 2013. "Do Countries Falsify Economic Data Strategically? Some Evidence That They Might." *Review of Economics and Statistics* 95, no. 2: 591–616.

Morgenstern, Oskar. 1963. *On the Accuracy of Economic Observations,* 2nd ed. Princeton University Press.

Morris, Stephen, and Hyun Song Shin. 2002. "The Social Value of Public Information." *American Economic Review* 92, no. 5: 1521–34.

Nakamura, Emi, Jón Steinsson, and Miao Liu. 2014. "Are Chinese Growth and Inflation Too Smooth? Evidence from Engel Curves." Working Paper no. 19893. Cambridge, Mass.: National Bureau of Economic Research.

Norris, Floyd. 2014. "Doubting the Economic Data? Consider the Source." *New York Times,* November 6.

Rauch, Bernhard, Max Göttsche, Gernot Brähler, and Stefan Engel. 2011. "Fact and Fiction in EU-Governmental Economic Data." *German Economic Review* 12, no. 3: 243–55.

Roos, Michael W. M., and Ulrich Schmidt. 2012. "The Importance of Time-Series Extrapolation for Macroeconomic Expectations." *German Economic Review* 13, no. 2: 196–210.

Comments and Discussion

COMMENT BY

STEFAN NAGEL
 This paper by Alberto Cavallo, Guillermo Cruces, and Ricardo Perez-Truglia is an interesting one, and it studies a question that has not received much empirical attention. For the most part, we take it for granted that individual decisionmakers in the economy have access to high-quality data on important macroeconomic variables, such as inflation rates and GDP growth rates. Although macroeconomists have paid some attention to the noise in official statistics and the effect of data revisions, we know little about the effects of intentional manipulation of official macroeconomic statistics. This paper provides valuable new evidence on this question.

For thinking about the potential consequences of data manipulation, it is useful to first reflect on the benefits associated with a government's provision of macroeconomic statistics. Because private sector agents could potentially learn from their own observations and through social learning channels, it is not entirely obvious how important government provision of macroeconomic information really is. Though some exceptions have been noted in the literature, better public information is typically thought to be welfare-improving (Hellwig 2005). Furthermore, the existing evidence suggests that private sector agents rely to a substantial extent on official statistics when forming perceptions about current macroeconomic conditions and expectations about conditions in the future. For instance, in an interesting study of a software bug that caused an error in the inflation statistics in Ecuador for several months, Paul Carrillo and Shahe Emran (2012) show that the error had a substantial effect on expectations.

The example from Ecuador is one of an accidental error, which is quite different from the intentional manipulation considered in this paper. In the case of Argentina considered here, the government's intentions and its attempts at manipulation were probably quite clear—at least to parts of

the population. The paper provides quite convincing evidence that people in Argentina were not naively misled by the government's manipulation. In terms of the posterior mean of the perceived inflation rate, individuals appeared to do quite a good job in debiasing the information provided in official inflation statistics. The effects on the posterior mean are the focus of the analysis in this paper. As I discuss in more detail below, the manipulation of statistics could potentially also have interesting and important effects on uncertainty about the inflation rate, not just on the posterior mean.

ASYMMETRIC REACTION TO INFLATION RATES In their online survey experiment, the authors present individuals with different official and unofficial inflation rates, and they elicit their inflation perceptions. Roughly speaking, subjects react to an official inflation rate of X percent in a similar way as to an inflation rate from unofficial sources of $X + 10$ percent. Thus, the difference in their reaction is roughly in line with the magnitude that the actual bias in the inflation rate reported in the Consumer Price Index (CPI) seemed to have. On average, individuals' processing of information seems to be quite well aligned with the authors' simple Bayesian model. One interesting pattern—and one that is less straightforward to understand—is the asymmetry in individuals' reaction to inflation rates at different levels. Varying the inflation rate that individuals are treated with from 10 to 20 percent (official or unofficial) produces a much smaller difference in subjects' perception of the true inflation rate than moving from 20 to 30 percent. This effect does not quite fit with the authors' normal prior–normal likelihood Bayesian model. To generate such an asymmetric reaction in the model, one probably needs to introduce some nonlinearity—for example, in the form of an upper bound of the bias. I suspect that if the bias was bounded above at zero, such an asymmetric reaction would arise within the model.

INTERPRETATION OF TREATMENT In the experiment, subjects are treated randomly with one of six inflation rates. Each of these represents a rate that was actually reported by a government agency or a private sector institution. The official inflation rates from the government agencies differ because they refer to different inflation indexes (CPI, GDP deflator, nominal wage growth). During the period in question, the differences between these rates were substantial. Because the official rates used in the treatment are actually reported ones, just for different types of indexes, the authors consider their treatment to be nondeceptive. This may not be an entirely accurate characterization, because subjects are not being told to which official inflation index the rate with which they are treated refers. Because the CPI is the index with which subjects are presumably most familiar, most subjects probably think that the rates with which they are treated are CPI

rates. In this sense, the experiment is effectively deceptive. Moreover, for the experiment to work as intended, this kind of deception is actually necessary. The authors want subjects to think that the official rates with which they are treated are CPI rates. The authors' interpretation of the findings is based on the assumption that this deception was successful. However, because one cannot be sure that the deception worked perfectly, the reliance on deception complicates the interpretation of the results. When subjects are presented with, say, a high official inflation rate, do they infer, to some extent, that this could be nominal wage growth rather than CPI inflation? And as for the perceived inflation rate that they report to the experimenters, is this now their perceived CPI inflation rate or the perceived rate for some other basket or index?

PRICE CONTROLS The authors' consumer-intercept survey shows little evidence that the government's price controls that applied to certain types of products affected people's perceived goods-specific inflation rates. Individuals' recalled price changes are similar for both controlled and noncontrolled products. The authors' conclusion is that the price controls did not affect people's inflation perceptions. Though this is a plausible interpretation of the evidence, other interpretations are also possible. In particular, while individuals may not be able to correctly recall product-specific price changes, it is still possible that price controls had an effect on their overall perceptions of inflation. Individuals' recollections of the average price changes they experienced could very well be affected by the price controls, even though they cannot quite recall correctly any product-specific price changes anymore.

The authors conclude from their evidence that "the government's attempt to manipulate inflation expectations seems to have been ineffective and plausibly counterproductive." This claim seems to be largely true, given the evidence in the paper, if one interprets it as meaning that the government's manipulation had no effect on the mean of the perceived inflation distribution. It would be incorrect, however, to conclude from this that the government's manipulation had no effect at all on the perceived inflation distribution. Furthermore, the fact that Argentina had a huge amount of outstanding inflation-linked bonds means that the manipulation may have resulted in a substantial wealth transfer away from the holders of these bonds.

UNCERTAINTY ABOUT INFLATION Although the government's manipulation had little effect on the mean of consumers' subjective posterior distribution of inflation, the manipulation could have had a substantial effect on consumers' uncertainty about inflation. The authors do not focus on uncertainty effects, and their experiments are not designed to measure effects on uncertainty, but the uncertainty channel could be an important one. As the

authors' model shows, even with a fixed bias, the presence of the bias raises the posterior uncertainty about the inflation rate. The degree of uncertainty would be further magnified if one extended the model to allow for a random component in the bias. An elevated level of inflation uncertainty could lead to adverse economic consequences. For example, with more noise in the public inflation signal, firms might put more weight on idiosyncratic signals, leading to greater price dispersion and, as a consequence, misallocation. Price dispersion indeed seems to have grown following the manipulation; Andres Drenik and Diego Perez (2016) find a 13 percent increase in price dispersion in Argentina following the manipulation of the official inflation rate. It would be interesting to study in more detail to what extent a distortion of official inflation rates raises inflation uncertainty.

INFLATION-LINKED BONDS Even without any effect on consumers' perceived inflation distribution, the policy could still be an effective one from the viewpoint of a (short-termist) government (and a nasty surprise for holders of inflation-linked bonds). At the start of the manipulation period, Argentina had about $50 billion worth of inflation-linked outstanding debt (Webber 2008). Inflation-linked bonds are supposed to protect investors' real wealth against inflation, but they only do so if the inflation rate used in the calculation of bond payments is not manipulated.

On each coupon date t, inflation-linked bonds pay a contractually fixed coupon rate times the ratio $CPI(t)/CPI(0)$, where $CPI(t)$ is the CPI level at the time of the coupon payment and $CPI(0)$ is the CPI level at the time the bond was issued. Similarly, the principal paid back to bondholders at maturity T is a contractually fixed face value times $CPI(T)/CPI(0)$. With $50 billion worth of inflation-linked outstanding debt, downward manipulation of the inflation rate by 10 percent a year saves the government $500 million in coupon interest payments and $5 billion in accrued principal each year. In present value terms, if we take a 10-year, zero-coupon bond for a back-of-the-envelope calculation, a 10 percent downward manipulation of the CPI inflation rate over the life of the bond (that is unanticipated at the time of the bond issue) would amount to a 50 percent loss for bondholders and a gain of similar size for the government.

My figure 1 presents the "real" yield of Argentina's inflation-linked bonds (constructed by Datastream as a weighted average across all outstanding maturities). In the case of Argentina, the usual calculation of a real yield no longer delivers the real yield once the government starts manipulating the inflation statistics. Instead, the "real" yield becomes

true real yield + true inflation rate − manipulated inflation rate.

Figure 1. The "Real" Yield of Argentina's Inflation-Linked Bonds, 2004–15

Percent
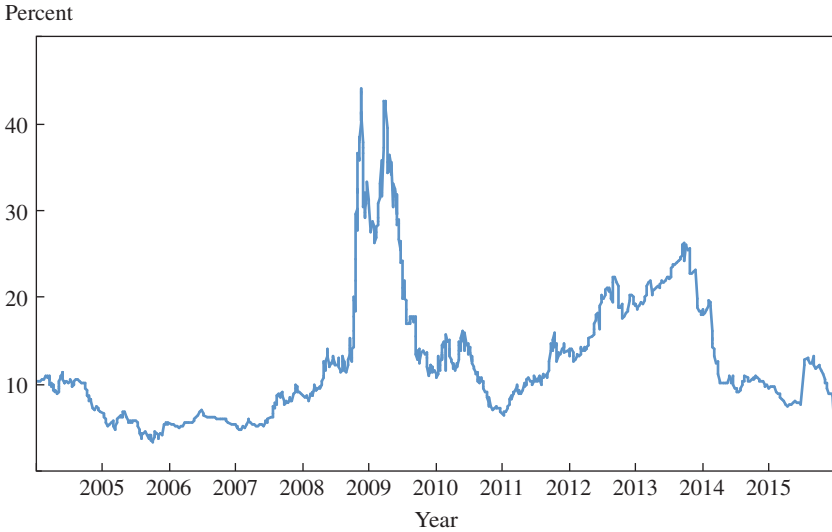
Source: Datastream, code BEMA0Y.

A downward bias in the official inflation rate thus raises the "real" yield that the usual calculation will deliver. In the extreme case where the government manipulates the CPI inflation rate to zero, the "real" yield would equal the nominal yield on a nominal bond. In this case, inflation protection has become completely ineffective and the bond trades like a nominal bond. As my figure 1 shows, the bond market is moving toward pricing inflation-linked bonds more like nominal bonds, as the "real" yield rises from about 5 percent in 2007 to close to 15 percent in 2008, consistent with the bias of about 10 percentage points in the official CPI inflation rate.[1]

Thus, from the viewpoint of the government, the manipulation of the CPI may have been highly effective in terms of its fiscal consequences, even if the policy did not succeed in affecting the mean of individuals' perceived inflation distribution.

CONCLUSIONS Overall, this paper provides useful evidence on how individuals dealt with manipulated official statistics in Argentina. Unofficial statistics are helpful as a substitute for information from official sources

1. The big jump in the "real" yield in late 2008 coincided with the announcement by Argentina's government that it was nationalizing private pension plans. This rise in yields thus likely reflected different concerns, not the manipulation of the inflation rate.

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 Analy-sis,* 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:

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

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.

of five Argentine provinces, he noted that one reported an inflation rate of 40 percent, two reported inflation rates of 30 percent, and two reported inflation rates of 10 percent. He suggested that an interesting experiment might be to compare reactions to inflation statistics across provinces, and wondered if this sort of experiment might be possible.

Wojciech Kopczuk had two questions. First, he wondered to what extent black market prices and other ways of pricing were prevalent in Argentina. Specifically, he noted that the most important black market price that is easy to observe is the exchange rate. The exchange rate is informative about other prices, so it is interesting to explore to what extent it factors into perception of the bias and how people are reacting to information. Second, he wondered to what extent people actually understood what inflation is. Fundamentally, the price level is measured relative to a basket of goods, but most people probably have different ideas about what sorts of items are included, especially in the presence of price controls. For instance, if some prices are fixed, one might not want to include them in the average.

Carol Graham pointed out that she grew up in Latin America, where hyperinflation was rampant. She noted that even very poor people seemed to be sophisticated about things like inflation and consumer baskets. For instance, workers will get their wages in whatever the home currency is and cash them in for dollars at night; they are very well aware of what the exchange rate is, in a way that most people in the United States just are not because they are not living with hyperinflation. Any kind of hyperinflation becomes a way of living, of trying to survive. She warned not to underestimate the sophistication of consumers in this story.

Joe Beaulieu made one minor point. As had been suggested, it was a very interesting idea to look at the effect of the dispersion of prices once the Argentine government had introduced its new policies toward indexation. He cautioned that at the same time, there appeared to be a sharp increase in inflation, and that there is a fairly robust fact in the literature that the two are related in all sorts of ways, both in terms of inflation rates and the actual dispersion of prices.

Alberto Cavallo noted that most people in Argentina do in fact know what inflation is. On the question of price controls, Cavallo stated that the authors looked at not only how price controls affected how Argentines remembered prices but also whether they affected expectations. They found that price controls did not significantly affect expectations. Surprisingly, they also found that price controls did not lead to shortages, the reason being that the government very quickly moved from undertaking a massive price control program to setting highly targeted price controls.

In his presentation, discussant Stefan Nagel had noted that there appears to be a puzzling asymmetry, or downward stickiness, in people's perceptions of inflation. For example, at lower levels of reported inflation—say, 10 or 20 percent—there were relatively small differences in perceived inflation. However, at higher levels of reported inflation—say, 20 or 30 percent—there seemed to be a much bigger difference in perceived inflation. Cavallo argued that the asymmetry appears to only be present in government statistics, not in the prices of goods.

Nagel had one additional comment for the authors. He wondered, when it came to paying off the holders of inflation-linked bonds, whether the Argentine government unrigged the official numbers. Cavallo responded in the negative, and noted that the only time the government unrigged them was in January 2014, when it realized that if it ever wanted to borrow again, it would need to clean up its act. The government launched a new CPI, which did not recognize any prior data; as a result, bondholders were swindled.