Learning from Potentially Biased Statistics

Household Inflation Perceptions and Expectations in Argentina*

Alberto Cavallo
MIT and NBER

Guillermo Cruces
CEDLAS-FCE-UNLP
CONICET and IZA

Ricardo Perez-Truglia†
Microsoft Research

April 5, 2016

Abstract

When forming expectations, households may be influenced by perceived bias in the information they receive. For instance, in the United States, households that do not trust inflation statistics have, on average, 50% higher inflation expectations. 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 of government manipulation of inflation statistics in Argentina (2007–2015). This period is interesting because of the attention to inflation information and the availability of official and unofficial statistics. 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 to understand 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.

JEL Classification: C93, D83, D84, E31.

Keywords: statistics, manipulation, expectations, learning, inflation, natural experiment, survey experiment.

---

*This paper was prepared for the Spring 2016 conference of the Brookings Papers on Economic Activity. We thank the editors, Janice Eberly and James Stock, for their guidance in revising this paper. We also thank Robert Barro, Raj Chetty, Luciano Cohan, David Laibson, Roberto Rigobón, N. Ruiz, Guido Sandleris, Martin Tetaz and Fernando Yu for their valuable comments, as well as those from seminar participants at Harvard University, Universidad de San Andres, MIT Sloan and Universidad Torcuato Di Tella. Julián Amendolaggine, Martín Caruso, and Maria Fazzolari provided excellent research assistance. We also thank Carolina Yellati for her collaboration in conducting the survey experiment, and Guido Sandleris and Fernando Freijedo for the CIF-UTDT household inflation expectations survey data. Funding for this experiment was generously provided by MIT Sloan and CEDLAS-FCE-UNLP. This project was reviewed and approved by the Committee on the Use of Humans as Experimental Subjects at MIT.

†Corresponding author: rtruglia@microsoft.com. Microsoft New England Research and Development (NERD) Lab, Office 12073, 1 Memorial Drive, Cambridge MA 02142.

1 Introduction

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 form. In recent years, a growing empirical literature provides 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 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.

We study how households learn, in a Bayesian sense, in an environment that has high interest in learning about inflation and several sources of inflation statistics, some of which are biased by the government. Our findings are based on observational and experimental evidence obtained during the recent period of manipulation of inflation statistics in Argentina from 2007 to 2015. This is an ideal setting for several reasons. First, the inflation rate fluctuated between 15% and 30%, which promoted 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] Third, the lack of reliable official data promoted the creation of several unofficial inflation indicators during this period, thereby potentially allowing individuals to counteract government manipulation by using other data.

We start with some observational data on the co-movement of inflation expectations and official and unofficial inflation statistics before and after the intervention of the National Statistics Institute (INDEC), when the government started reporting official statistics that were systematically below the unofficial estimates. The households' inflation expectations quickly diverged from official inflation indicators and instead aligned with 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 counter-factual 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 of the observational data, we provide a simple model of Bayesian learners with potentially biased statistics and design a survey experiment to test its predictions.

[1]Since 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, 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 of statistics manipulation, 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.

The model shows that, far from ignoring official statistics, rational learners should react to changes in official statistics by "de-biasing" the signal based on their perceived bias 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.

We conducted a large-scale survey experiment in Argentina in December 2012 to test this prediction. We provided respondents with different inflation estimates and 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 non-deceptive 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%).

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, relative to individuals who were told that the official inflation rate was 20%, individuals who were told that it was 10% reported lower inflation perceptions and expectations and individuals who were told that it was 30% reported higher inflation perceptions and expectations.

The experimental data also allows us to test directly the hypothesis of 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% as they would to a signal that unofficial inflation is 20% and that they should react similarly to an official inflation of 20% as they would to an unofficial inflation of 30%. These predictions are consistent with subjects' reactions in our experiment. That is, in an environment with many alternative inflation indicators and much attention to the topic, individuals are 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 suggest that individuals form inflation expectations using information from their own consumer experiences (Bates and Gabor, 1986; Bruine de Bruin et al., 2011; Coibion and Gorodnichenko, 2015). In particular, individuals rely heavily on their perceptions about 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, the Argentine government froze the prices of a relatively large and important sample of consumer products in 2013. We show that, even though the inflation rate fell significantly, household inflation expectations did not fall. To explore this finding further, we ran a price-elicitation survey outside a large supermarket chain in Argentina during the time of the 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 help explain how individuals learn from inflation data in other countries. Even in developed nations, a significant share of individuals do not trust official statistics. According to data from Eurobarometer (2008), 69% of respondents in Europe felt that it was necessary to know about economic indicators, but only 46% stated that they tended to trust official statistics such as the growth rate, the inflation rate, and the unemployment rate. Among U.S. survey respondents, 27% rated their trust in official statistics as 4 or lower on a 1–10 scale (Curtin, 2009). Analysts, commentators, and the media routinely discuss the possibility of manipulated statistics, such as job creation data that was released right before the 2012 election in the United States.[3]

Data from a survey of U.S. households reported in Cavallo, Cruces, and Perez-Truglia (2014) show that 32% of respondents do not trust official inflation statistics. Furthermore, relative to those who trust inflation statistics, respondents who do not trust statistics have inflation expectations that are 50% higher on average. This evidence suggests that lack of trust in the government may explain part of the stylized fact that households do not fully incorporate information from inflation statistics in their perceptions and expectations (Mankiw and Reis, 2002; Mankiw, Reis and Wolfers, 2003; Carroll, 2003).

To the best of our knowledge, our 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 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 et al., 2011), and alternative growth estimates and inflation in China (Nakamura et al., 2014). Michalski and 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 literature about 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, 2013). Several studies provide evidence that in-

[3]See Norris, Floyd. 2014. "Doubting the Economic Data? Consider the Source." The New York Times, November 6. http://www.nytimes.com/2014/11/07/business/economy/doubting-the-economic-data-consider-the-source.html.

[4]Morgenstern also covers the difficulties of measuring the national product, and in fact Argentina's government also falsified INDEC's GDP indicator (Camacho et al., 2015), for political reasons and to avoid the payment of a GDP warrant (a bond that only paid debtors if GDP grows at a certain rate).

flation statistics play a significant role in driving inflation expectations, including the analysis of variation in media coverage of statistics (Lamla and Lein, 2008; Badarinza and Buchmann, 2009; Drager, 2011), quasi-experimental variation in reporting official statistics (Carrillo and Emran, 2012), and information-provision experiments (Roos and Schmidt, 2012; Armantier et al., 2014; Cavallo, Cruces and Perez-Truglia, 2014).

Last, this paper relates to a theoretical literature about central bank transparency and 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 is welfare-enhancing (Hellwig, 2005), whereas others argue that it can reduce welfare (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 2 describes the period of manipulation of official statistics in Argentina and presents the observational evidence. Section 3 presents a simple model of Bayesian learning from manipulated statistics, as well as the design of the survey experiment and its results. In Section 4, we discuss the period of price controls in 2013. The last section concludes.

2 Manipulation of Inflation Statistics in Argentina

2.1 The Intervention of the National Statistical Institute (INDEC)

After a severe economic crisis in 2001–2002, 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 reached double digits in 2005 (12.3% 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 the National Statistics Institute (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 of items that had higher inflation, and to use prices from price-control lists even if those goods were not available for sale at the stores where the data was collected.

In February 2007, facing a second year of inflation above 10% 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% in January 2007 to 0.4% in February and continued falling in the following months. INDEC employees publicly disclosed what had happened in the previous months, which increased suspicions of a manipulation of the CPI.

INDEC stopped publishing some disaggregated inflation series and announced "methodological changes" that were never publicly disclosed.

The government intervention at INDEC had immediate negative consequences for the Argentine economy, as discused by Levy-Yeyati and Novarro (2013). Although the government paid less in the short-run on inflation-linked bonds, most of that 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, eventually leading to the government's imposition of foreign exchange controls in 2011. Despite the controversy and negative effects on the economy, the manipulation of the official CPI continued until December 2015, when a new government was elected.

2.2 Unofficial Inflation Statistics

The unusual situation with INDEC led to the creation of alternative measures of inflation that we refer to generally as "unofficial" Inflation indicators. The main alternative indicator we use is computed by PriceStats, a private firm based in the United States that uses online prices from large retailers since 2007. The PriceStats index is published weekly in The Economist.[6] A second alternative indicator, published since 2008, is produced by Buenos Aires City, a think tank lead by Graciela Bevacqua (the 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 on the Greater Buenos Aires area only, provincial statistical agencies also collected regional price data and computed their own CPIs. The federal government pressured the provinces to manipulate or stop publishing those indexes, but provinces that were not aligned 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 (based on Greater Buenos Aires) index 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 that were publishing their own unofficial inflation estimates. Some members

[5] For example, in 2008 the government payed an interest rate of 15% in dollars for newly issued debt sold to the government of Venezuela.

[6] PriceStats is a private company connected with the Billion Prices Project at MIT, an academic initiative created in 2008 by Alberto Cavallo (an author in this paper) and Roberto Rigobon to experiment with the use of online data in the production of price indices and other macro and international research applications. See Cavallo and Rigobon (2016) for more details on the Billion Prices Project.

[7] See Bevacqua and Salvatore (2009) for details.

Figure 1: Time-line of the Manipulation of Inflation Statistics in Argentina

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 the INDEC announces that 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 the INDEC
	Paglieri decides to stop the publication of CPI-GBA
	Bevacqua, the director of the Prices Departmen is suspended
Feb 2007	The government officially intervenes INDEC
	The first manipulated CPI-GBA monthly index is published
	Bevacqua is officially fired and replaced by Paglieri
	First mobilization of the 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	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 of the INDEC employees
Sep 2007	The 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
March 2008	Launch of'www.inflationverdadera.com", a website were alternative indicators using online prices are updated on a daily basis. The inflation rate is three times higher than in CPI estimates.
May 2008	INDEC stops publishing the CPI-Nacional, an index that used price data from 7 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 IMF 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 with 500.000 pesos for failing to comply with Moreno's request
May 2011	The "CPI Congreso" (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 goverment threatening them with criminal prosecutions if they continue to publish their own inflation estimates
Feb 2012	The IMF 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 IMF 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 became the director and Bevacqua returns as the technical director of INDEC
	Todesca says that INDEC is like a "scorched earth" and suspends the publication of the CPI and other price indexes
Jan 2016	Bevacqua announces that it will take 8 month for the INDEC to publish a new CPI.
Feb 2016	Bevacqua is fired again.

Notes: Compiled by the authors from newspaper articles and other sources.

of Congress from the opposition, who were immune from prosecution, compiled and published a monthly average of "private" estimates. Other alternative indicators also were publicized. The Appendix provides a comprehensive list, with characteristic details and methodologies.

Figure 2: Official Inflation and Alternative Unofficial Indicators, 2004-2013

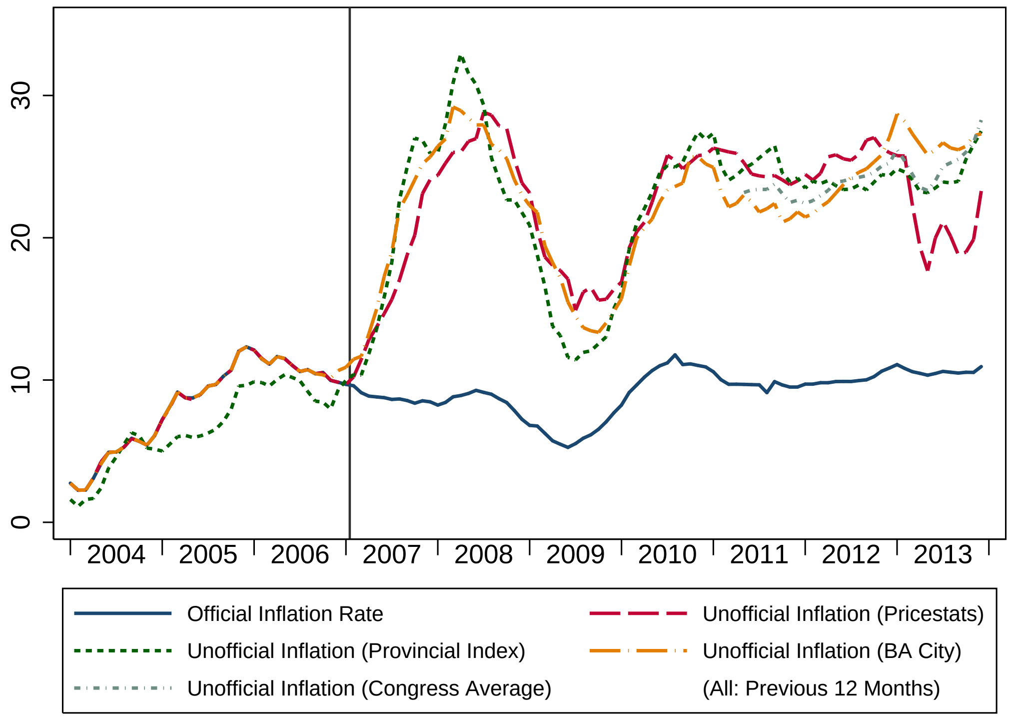

Notes: The vertical line represents the start of the intervention of the national statistical agency (INDEC) in January 2007. "Official Inflation" is the annual inflation rate reported by INDEC. The underlying Consumer Price Index is based on a sample of prices from the Greater Buenos Aires Metropolitan Area (GBA). We present in this Figure a series of unofficial inflation rate estimates: a) "Unofficial Inflation (PricesStats)" is an indicator compiled by PriceStats LLC based on prices from online retailers (derived from the MIT's Billion Prices Project; see Cavallo, 2013, for more details). b) "Unofficial Inflation (Provincial Index)" is based on a geometric average of nine provincial statistical agencies' Consumer Price Indexes (source: CqP). c) "Unofficial Inflation (BA City)" is computed by Buenos Aires City from a sample of products and prices from the GBA 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 indices at those times d) "Unofficial Rate (Congress Average)": an average of several unofficial indicators compiled by Congress representatives from opposition parties. All inflation indicators are monthly series.

Figure 2 shows the annual inflation rate for all of these unofficial indicators and the official CPI. The vertical line shows the month of the intervention at INDEC, indicating that the official and unofficial indicators diverged immediately. 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.

2.3 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 Centro de Investigación en Finanzas at Torcuato Di Tella University began a national household survey of inflation expectations that they published monthly.

Figure 3: Official Statistics, Unofficial Statistics and Inflation Expectations, 2006-2013

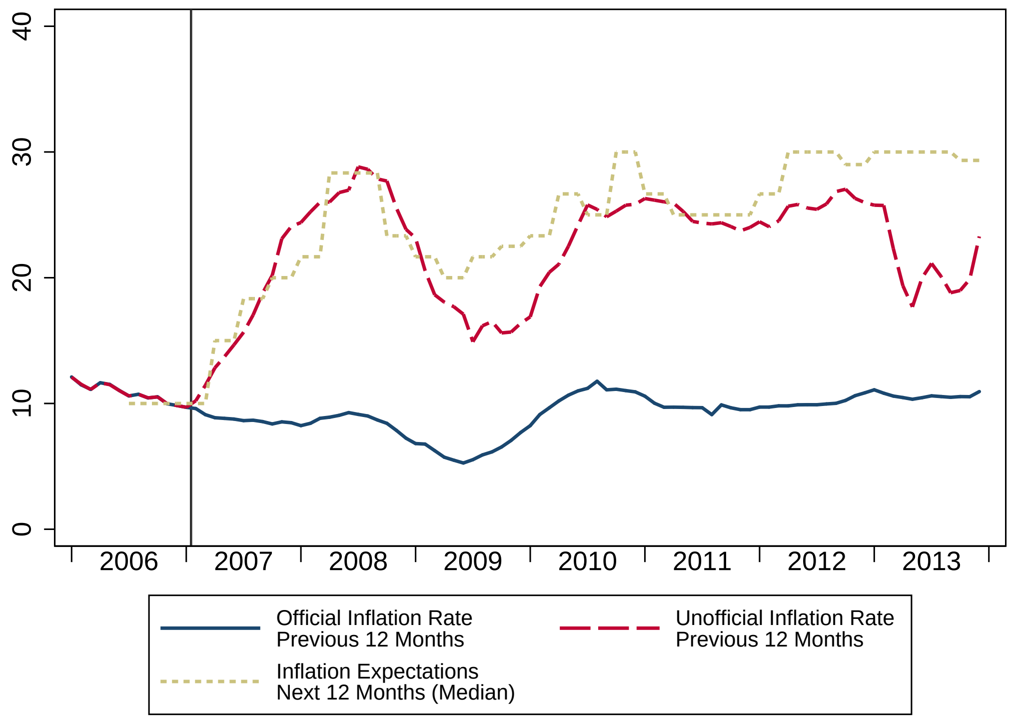

Notes: The vertical line represents the start of the intervention of the national statistical agency (INDEC) in January 2007. "Official Inflation" is the annual inflation rate reported by INDEC. The "Unofficial Inflation" indicator is computed by PriceStats (see notes in Figure 2 for more details). The median of inflation expectations for the following 12 months correspond to quarterly averages of the monthly median from the Encuesta de Expectativas de Inflación. This survey, conducted by the Centro de Investigación en Finanzas (Universidad Torcuato Di Tella) since August 2006, collects information on the inflation expectations for the following 12 months among the general population of Argentina, based on a standard question for this type of survey ("What do you expect the annual rate of inflation will be during the next 12 months?").

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 co-evolution of available inflation indicators and of inflation expectations for seven years of uninterrupted manipulation of official statistics.

Household inflation expectations aligned with the unofficial inflation level over time. The PriceStats index were not disseminated until March 2008,[8] but newspapers reported other unofficial estimates during that time. In the Appendix, we plot the annual inflation rates mentioned in these articles and show that they track inflation expectations during 2007.

There is also some evidence of an asymmetric response of 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 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.

3 Experimental Evidence

The patterns that emerge from the time series 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 this observational data, and it is unclear whether individuals are simply ignoring the official data or adjusting to it 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.

3.1 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 that 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 main intuition of the model.

The individual can observe two signals based on the information about the price changes for the 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 represent the information that individuals obtain by using

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

averages of their own memories about price changes over a set of products. The second signal is the government's official statistics. We assume 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 adds instead 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 will replace for $\sigma^2_u = \frac{1}{N_u}\sigma^2_{actual}$ and $\sigma^2_o = \frac{1}{N_o}\sigma^2_{actual}$.[9]

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 of the inflation rate prior to obtaining new information, and π_1 as the belief of the inflation rate after doing so, while b_o and b_1 are similarly defined. The normality assumption about the distribution of price changes determines that the conjugate distribution for beliefs about inflation and bias is bi-variate normal. For the sake of notational simplicity, we focus on the case where the prior beliefs for the inflation rate and the bias are orthogonal. As shown in Appendix **??**, this assumption leads to the following posterior beliefs:

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

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

The mean posterior belief for 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 for the 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 for these weights are provided in the 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 to understand the intuition of the model.

A 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 that 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 believe that is driven by sampling variation. How fast would the individual learn about a bias? By making the relevant replacements in the above formula for b_1 we get that

[9] In practice, σ^2_u and σ^2_o represents not only a pure statistical error 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.

$b_1 = (\psi_1 + \psi_2)(\bar{o}' - \pi_0)$. 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, she would not change so rapidly her belief about a bias in the official data.

A second scenario explores how an individual who believes that the government is manipulating statistics reacts to the official statistics relative to 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 that starts out thinking that the government biases the inflation statistics downwards: i.e., $b_0 < 0$. How does the individual react to official statistics relative to unofficial statistics? From the formula for π_1 it follows that, qualitatively, the individual reacts to $\bar{o}'$ in the same way as she reacts to $\bar{u}$, but with the exception that first it subtracts from $\bar{o}'$ the ex-ante perceived bias: i.e., it uses $\bar{o}' - b_0$ instead of $\bar{o}'$. So if the individual believes that the bias is $b_0 = -10\%$, then she should react qualitatively to the signal $\bar{o}' = 10\%$ in the same way than she reacts to $\bar{u} = 20\%$. 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 unofficial and official statistics, $\frac{1}{\sigma_u^2}$ and $\frac{1}{\sigma_o^2}$. However, if these precisions are similar, then we would expect a reaction that is quantitatively very similar.[10]

3.2 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 (e.g., Roos and Schmidt, 2012; Armantier et al., 2014; Cavallo, Cruces and Perez-Truglia, 2014). We first collect some background information about respondents (see Appendix C 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 (i.e., the respondent's perception of the annual inflation rate over the previous 12 months). We also include a question about the respondents' subjective assessments of their confidence in their answers, measured in a scale from 1 ("Not at all sure") to 4 ("Very sure"). The subject's inflation expectations correspond to the expected inflation rate over the following 12 months. Argentina's economic history implies that the general public understands the meaning of the word "inflation," which is

[10] Note that even if the precision of unofficial and official statistics were exactly the same, $\frac{1}{\sigma_u^2} = \frac{1}{\sigma_o^2}$, 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.

discussed routinely in the media.[11] 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.[12]

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%]*"

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

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 by government manipulation. At the time of our experiment, the annual growth rate of the official CPI was approximately 10%. 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 20%. The government could not allow the GDP deflator to be as low as the CPI (10%), because that would have implied an implausibly high real GDP growth rate (over 15%). 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%. We followed a similar strategy to exploit the variation in unofficial statistics. We chose an index published by one unofficial source that indicated an inflation rate close to 20% and another index that indicated an inflation rate close to 30%. A third unofficial index, published by a think tank with close ties to the government, indicated an inflation rate close to 10%.

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

[11]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 Encuesta de Expectativas de Inflación from the Centro de Investigación en Finanzas, Universidad Torcuato Di Tella. Also, we did not provide any incentives for respondents to answer accurately, such as prizes for guessing the right figures. As shown by Armantier et al. (2012) in the context of similar studies, there is a significant correlation between incentivized and non-incentivized responses on inflation expectations.

[12]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.

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%. In any case, because individual judgment about the information can vary depending on the source, we included a debriefing at the end of the survey. In this debriefing 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 debriefing 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.

3.3 Subject Pool and Experimental Results

The sample is based on an established public opinion research firm that carries out a quarterly online survey of adults in Argentina with an stable set of questions since 2011. The experiments were conducted on December of 2012, while the government was still manipulating official statistics, and almost 6 years after the government started with 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 as 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%, and to each of the treatment groups with a probability of 12.9%. The final sample on which the following analysis is based consists of all the respondents that completed the questions on inflation perceptions and inflation expectations, yielding a final sample of 3,138 observations.[13]

[13]A small but not negligible 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 3.23% of the original sample). While this type of attrition occurred also in previous rounds of the opinion poll (for instance, dropout of 5.8% of the sample for the June 2012 round), in this case this might be a concern if the attrition was due to (and correlated to) the information treatments, since 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 of equal attrition across treatment groups (p-value: 0.79).

Table 1: Descriptive Statistics. Opinion Poll Experiment Sample Characteristics Compared to Argentina's Population

	Share Female	Mean Age	Share Buenos Aires	Share College	Voted Kirchner
Online Experiment Sample	57%	41.1	67.7%	60.7%	24.2%
Argentina's Average (20 years old or older)	52.8%	44.9	36.3%	15.6%	54.1%

Notes: Source: First row (Online Experiment Sample): opinion poll carried out in December 2012 (N=3138). Second row (Argentine Average): computed from the Encuesta Anual de Hogares Urbanos, 2012 (INDEC), except for the "Voted C. F. Kirchner" variable, obtained from the 2011 presidential election results. All statistics are based on individuals aged 20 or older in the corresponding samples. "Share Female" is the percentage of women in each sample. "Mean Age" is average age in years. "Share Buenos Aires" represents the percentage of those living in the Greater Buenos Aires Metropolitan Area. "Share College" represents the proportion of respondents who have a college degree or completed tertiary education. "Voted Kirchner" is the percentage who reported to have voted for Cristina Fernandez de Kirchner.

Table 1 presents some summary statistics about the demographics of the sample, along with the corresponding indicators for the general population. This sample is not representative of the Argentine general population: while it is roughly similar in terms of age and gender composition, our sample is substantially more educated and richer than average. Nevertheless, the qualitative results are similar if we re-weight 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 pre- and post-treament variables, for the control group and for each of the treatment groups. The last column reports the p-value of a test in which the null 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 pre-treatment variables, suggesting that the randomization was indeed balanced. The top panel shows the post-treatment 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 we show the distribution of inflation perceptions in the control group compared to that of each of the six other informational treatments. And Figure 6 summarizes the effects of the six informational treatments on the mean of various post-treatment outcomes relative to the control group.

Table 2: Average pre- and post-treatment responses by treatment group

	Control	Official 10%	Official 20%	Official 30%	Unofficial 10%	Unofficial 20%	Unofficial 30%	P-value
Post-Treatment								
$\pi_{\mathrm{i,t}}$	28.31	28.55	33.58	42.10	26.37	28.89	34.78	<0.01
	(0.591)	(0.812)	(0.809)	(0.815)	(0.805)	(0.809)	(0.816)	
Confidence in $\pi_{\mathrm{i,t}}$	-0.0903	0.119	-0.0144	-0.119	0.0122	0.0804	0.0912	<0.01
	(0.0319)	(0.0438)	(0.0437)	(0.0440)	(0.0435)	(0.0437)	(0.0441)	
$\pi_{\mathrm{i,t+1}}$	28.20	28.22	33.32	39.24	26.29	28.62	33.99	<0.01
	(0.595)	(0.817)	(0.814)	(0.820)	(0.810)	(0.814)	(0.821)	
Pre-Treatment								
Female	0.560	0.539	0.563	0.586	0.601	0.545	0.608	0.28
	(0.0181)	(0.0248)	(0.0247)	(0.0249)	(0.0246)	(0.0247)	(0.0250)	
Age	41.11	41.06	40.64	40.79	40.93	41.37	41.16	0.97
	(0.390)	(0.536)	(0.534)	(0.538)	(0.532)	(0.534)	(0.539)	
College Degree	0.633	0.610	0.650	0.581	0.572	0.600	0.578	0.12
	(0.0178)	(0.0245)	(0.0244)	(0.0246)	(0.0243)	(0.0244)	(0.0246)	
Own Economic Situation OK	0.261	0.237	0.228	0.203	0.233	0.257	0.211	0.26
	(0.0155)	(0.0213)	(0.0212)	(0.0214)	(0.0211)	(0.0212)	(0.0214)	
Observations	750	397	400	394	404	400	393	

Notes: Each cell represents the mean of each of the row variables for the corresponding control and treatment groups (in columns). The standard errors of these means are reported in parentheses below. The last column reports the p-value of a balance test in which the null is that the mean of each variable is equal between all seven experimental groups (the control group and the six treatment groups). $\pi_{i,t}$ represents the respondent's inflation perceptions for the previous twelve months. "Confidence in $\pi_{i,t}$" represents the respondent's own confidence on her response to the perceptions question on a 1 to 4 scale ("Not confident at all" to "Very confident"). $\pi_{i,t+1}$ represents the respondent's inflation expectations for the following twelve months. "Female" indicates the proportion of women. "Age" is the average age in years. "College Degree" represents the proportion of respondents who have a college degree or completed tertiary education. "Own Economic Situation OK" is the proportion who responded that their economic situation was better with respect to one year earlier, and zero otherwise. Source: opinion poll carried out in Argentina in December 2012.

Figure 4: Histograms Comparing Distribution of Inflation Perceptions between Control and Treatment Groups

a.1 Unofficial-10% vs Control

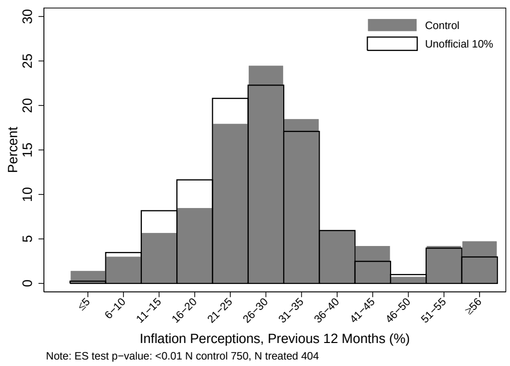

b.1 Official-10% vs. Control

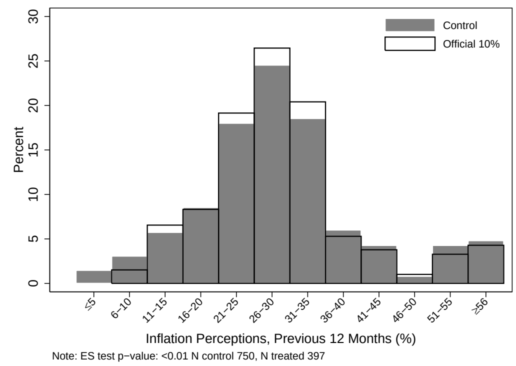

a.2 Unofficial-20% vs Control

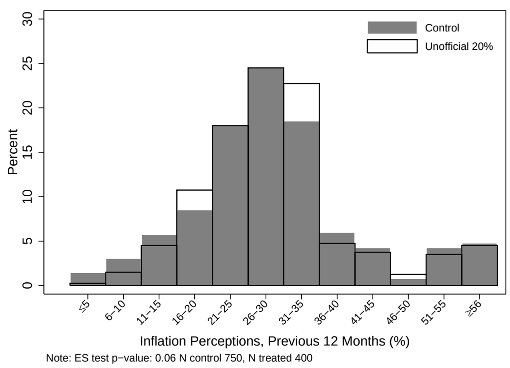

b.2 Official-20% vs. Control

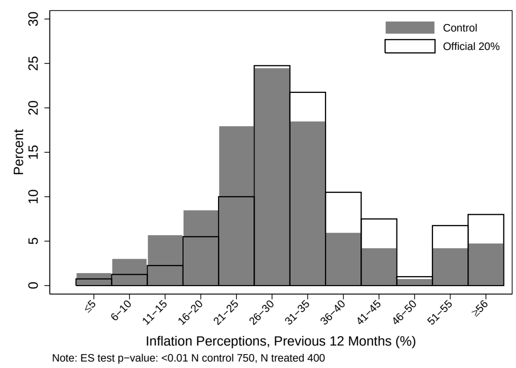

a.3 Unofficial-30% vs. Control

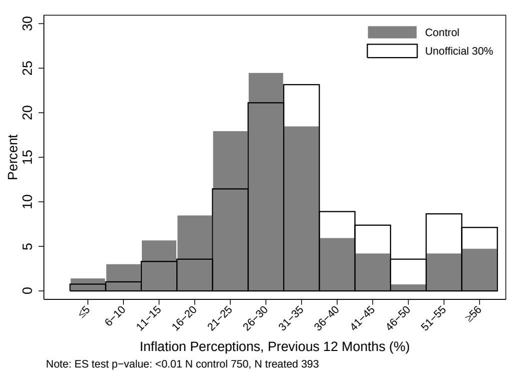

b.3 Official-30% vs. Control

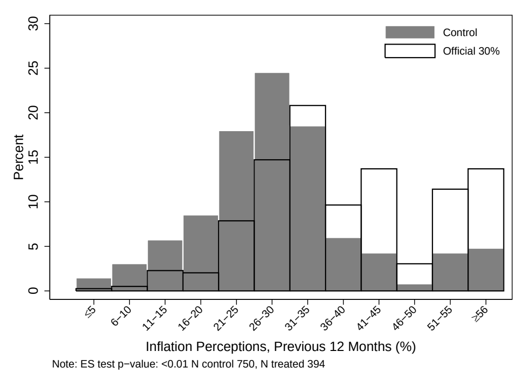

Notes: Observations: 750 in the control group, 397 in the Official-10% group, 400 in the Official-20% group, 394 in the Official-30% group, 404 in the Unofficial-10% group, 400 in the Unofficial-20% group, and 393 in the Unofficial-30% group. Respondents in the control group were not given any information about inflation statistics. Respondents in the Official-X% groups were provided non-deceptive information about inflation estimates from official sources (either 10%, 20% or 30% in the previous year). Respondents in the Unofficial-X% groups were provided non-deceptive information about inflation estimates from unofficial sources (either 10%, 20% or 30% in the previous year). ES is the Epps–Singleton characteristic function test of equality of two distributions (Goerg and Kaiser, 2009). The histograms are censored at 5% and 56% (inclusive), but these bins represent the cumulative observations below 5% and above 56% respectively. Source: opinion poll carried out in Argentina in December 2012.

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. Figure 4.a.1, 4.a.2 and 4.a.3 show the distribution of perceptions in the control group and each of the messages about the unofficial statistics. The data suggests that individual did not ignore this information: relative to individuals who were told that inflation according to official statistics was 20%, individuals who were told that official statistics was lower (10%) reported lower inflation perceptions, and individuals who were told that official statistics indicated higher inflation (30%) reported higher inflation 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 of these pairwise differences are statistically significant at the 1% 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. Figure 4.b.1, 4.b.2 and 4.b.3 show the distribution of perceptions in the control group and each of the messages about the official statistics. Relative to individuals who were told that inflation according to official statistics was 20%, individuals who were told that official statistics was lower (10%) reported lower inflation perceptions, and individuals who were told that official statistics indicated higher inflation (30%) reported higher inflation perceptions. According to the ES test, these pairwise differences in distributions are statistically significant at the 1% level. These differences are economically significant: for instance, the mean of inflation perceptions is 28.5% for the Official-10% group, 33.6% for the Official-20% group and 42.1% for the Official-30% 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 than they would to the same figure if it was coming from unofficial sources. The data strongly rejects this hypothesis: the ES test indicates that the difference between the distribution of inflation perceptions across individuals given messages Official-10% and Unofficial-10% is significant at the 1% level; the same is true when comparing distribution of perceptions for the Official-20% and Unofficial-20% groups, and for the Official-30% and Unofficial-30% groups. These differences are not only statistically significant, but also economically significant: for instance, relative to Unofficial-10%, the Official-10% message created 2.1 percentage points higher inflation perceptions; relative to Unofficial-20%, the Official-20% message created 4.7 percentage points higher inflation perceptions; and relative to Unofficial-30%, the Official-30% message created 7.3 percentage points higher inflation perceptions.

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. Based on this approximation, the learning model predicts that individuals should react to information conveying an official inflation level of X% in the same way as they would react to information from unofficial sources conveying a level of inflation of (X-10)%. The results from our experiment are consistent with this hypothesis: we cannot reject the null hypothesis that the distribution of inflation perceptions are equal between individuals in the groups Official-10% and Unofficial-20% (ES test p-value of 0.91), and we cannot reject the null hypothesis that the distribution of inflation perceptions are equal between individuals with messages Official-20% and Unofficial-30% (ES test p-value of 0.61). These differences are not only statistically insignificant, but also economically small. For instance, the difference in the mean perceived inflation is only 0.34 percentage points between the Official-10% and Unofficial-20% groups, and 1.2 percentage points between the Official-20% and Unofficial-30% groups.

The experiment also allows us to explore further the seemingly asymmetric relationship between perceived and actual inflation suggested by the analysis of the non-experimental time series data discussed in the previous sections. According to Table 2, increasing the unofficial inflation rate shown to the subject from 10% to 20% increased the mean perceived inflation by 2.51 percentage points ($p\text{-value}<0.05$). Instead, increasing the unofficial inflation rate shown to the subject from 20% to 30%, which is also a 10 percentage points increase, increased the mean perceived inflation by 5.89 percentage points ($p\text{-value}<0.01$). That is, the effect of going from 20% to 30% is almost twice the size of the effect of going from 10% to 20%, and this difference is statistically significant ($p\text{-value}=0.072$). 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% group than in the Official-20% group, and they are 8.5 percentage points lower in the Official-20% compared to the Official-30% group. The difference between these two effects is statistically significant ($p\text{-value}=0.082$).

This asymmetry in the experimental effects is consistent with the observational evidence from Section 2.3. 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 prior to 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 statistic may be subject to the government-bias, even if it comes from a non-governmental agency, and thus they put less weight when learning from that piece of information. On the contrary, because the government's goal is to reduce inflation perceptions, individuals do not worry that an statistic has a government-induced bias when it suggests that inflation is above their prior belief, thus explaining the asymmetry.[14]

[14]Note that this conjecture predicts that the asymmetry would be reversed if the government was interested in

Figure 5: Past Inflation Perceptions and Future Inflation Expectations, Binned Scatter plot

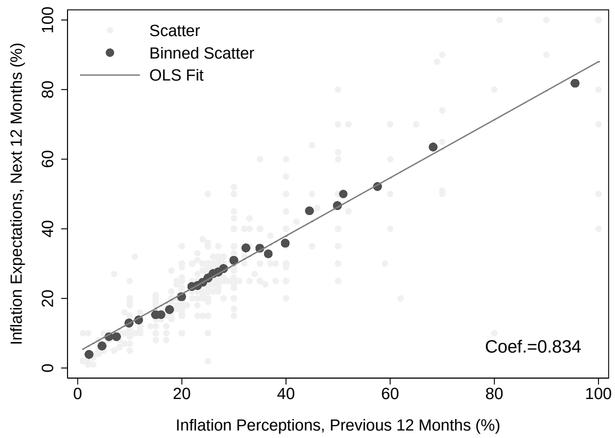

Notes: N=777. This data is for subjects in the control group, i.e., those who were not provided any information about inflation. Source: opinion poll carried out in Argentina in December 2012

Perceptions about past inflation are a key input in the formation of inflation expectations (e.g., Jonung, 1981; Cavallo, Cruces and Perez-Truglia, 2014). For instance, Figure 5 shows a binned scatterplot 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 Figures 6.a and 6.b indicates that the effects on perceived inflation (6.a) were very similar to the effects on inflation expectations (6.b). For instance, relative to the control group, the Unofficial-30% 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, Figure 6.c also shows the effect of the different treatments on the subjects' reported confidence on their inflation perceptions, coded from 1 (not sure at all) to 4 (very sure). Bayesian learning predicts that, after observing a signal that is useful, individuals should usually be more confident about their posterior beliefs relative to the counterfactual with no information. Figure 6.c indicates that, as expected, several of the informational treatments increased significantly the subjects' reported confidence on their inflation perceptions. Interestingly, the effect on reported confidence is almost twice as large for unofficial statistics than for official statistics (0.151 compared to 0.086, p-value of the difference 0.066), suggesting that the information from unofficial sources was on average more useful.

increasing rather than reducing inflation perceptions.

Figure 6: Treatment Effects on Inflation Perceptions and Expectations

a. Mean Perceived Inflation ($\pi_{i,t}$)

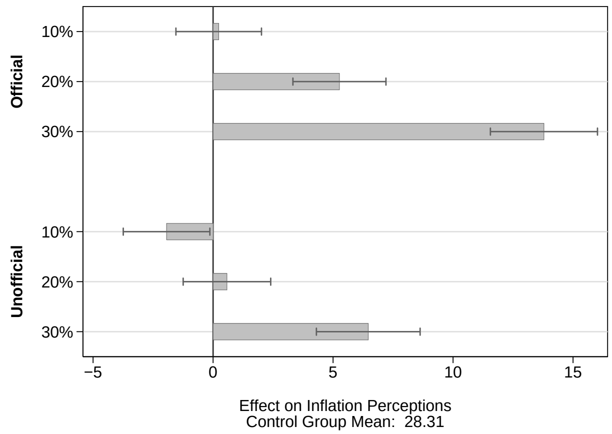

b. Mean Expected Inflation ($\pi_{i,t+1}$)

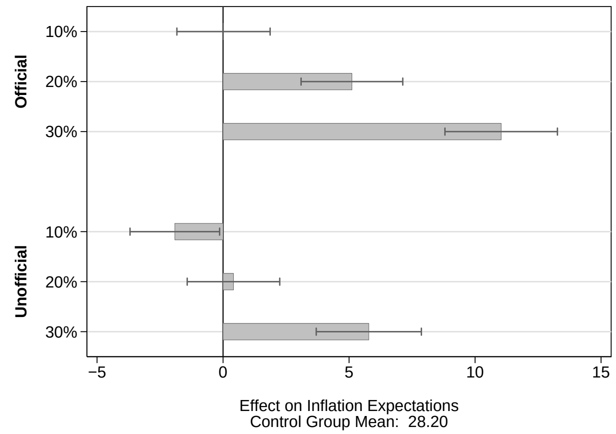

c. Mean Confidence in Perceived Inflation ($\pi_{i,t}$)

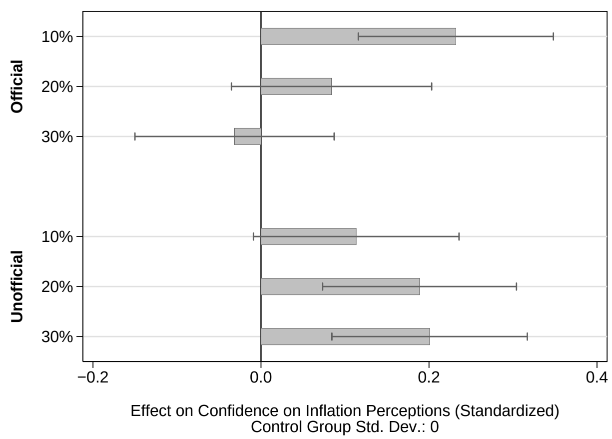

Notes: Observations: 750 in the control group, 397 in the Official-10% group, 400 in the Official-20% group, 394 in the Official-30% group, 404 in the Unofficial-10% group, 400 in the Unofficial-20% group, and 393 in the Unofficial-30% group. Each bar represents the point estimate of the effect of the specific sub-treatment compared to the control group. Robust standard errors reported. $\pi_{i,t}$ represents the respondent's inflation perceptions for the previous twelve months. "Confidence in $\pi_{i,t}$" represents the respondent's own confidence on her response to the perceptions question on a 1 to 4 scale ("Not confident at all" to "Very confident"). $\pi_{i,t+1}$ represents the respondent's inflation expectations for the following twelve months. Source: opinion poll carried out in Argentina in December 2012

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 of the reaction to inflation information provided in similar survey experiments is genuine rather than spurious. In any case, according to the results from that paper, it is plausible that the effects identified in this paper are quantitatively different, but qualitatively robust to spurious learning.

4 The Price Controls of 2013

Studies suggest that individuals form inflation expectations using information from their own consumer experiences, such as their memories of prices of supermarket products (Bates and Gabor, 1986; Bruine de Bruin et al., 2011; Cavallo, Cruces and Perez-Truglia, 2014; Coibion and Gorodnichenko, 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 an effort of the government in 2003 to control prices of supermarket products.

In February 2013, the government of Argentina significantly extended its policy of price controls on retail products. These were "price agreements" with big companies and large supermarket chains temporarily applied for hundreds of products in carefully selected categories. The government targeted goods that had a significant weight in the CPI basket, and focused on brands and retailers with large market shares. To enforce the price controls, the government publicly asked its supporters to help monitor prices. The program was called "Precios Cuidados" ("Protected Prices"). It was widely advertised and discussed in the media. While there were some problems in the implementation, most of the goods included in the agreements were available for sale at the agreed prices. It is possible that the government hoped that by controlling the prices for some 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."[15]

The inflation rate did temporarily fall: according to the unofficial statistic of PriceStats shown in Figure 3, the annual inflation rate fell from 25.8% in January 2013 to 17.7% in May 2013. This is not surprising, given that the Pricestats index draws its data mostly from large multi-channel 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 near 30%.

There are several possible explanations for the lack of effect on inflation expectations. One possibility is that people knew the effect would be temporary, so expectations about future inflation were not impacted. This probably played an important role, but we do not have a way to test it.

[15] See for example http://www.economia.gob.ar/wp-content/uploads/2014/04/07-04-20142.pdf

Another possibility is that even though people experienced more stable prices for some goods, this information did not affect perceptions of prices. To test this, we ran a consumer-intercept survey in the front door of four branches of one of the largest supermarket chains in the city of Buenos Aires, during the time 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' receipt from the supermarket purchase, which contained product identifiers that were matched to a database of scrapped online data from the same supermarket.[16] After scanning their receipts, respondents were asked about the current prices of the products they had just purchased, and the corresponding prices as of 12 months before.

Figure 7.a depicts the distribution of actual price changes for products with controlled prices, and for those with no controls. This figure shows that products with controlled prices did have a substantially lower inflation rate – average change of 1.9% compared to 21.7% for non-controlled products. However, Figure 7.b shows that the program was not effective in changing individual perceptions of price changes: the distribution of remembered price changes (as reported by the consumers) are very similar, and statistically indistinguishable, between controlled and non-controlled products.[17] In any case, just as it happens with the manipulation of the aggregate official index, price controls did not seem to be an effective way to influence inflation expectations, or at least not in the short term.

[16]Prices were scraped from the websites of the supermarket by the Billion Prices Project at MIT. See Cavallo (2013) for more details.

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

Figure 7: Actual and Remembered Price Changes for Products with Government Controlled Product Prices, Supermarket Survey

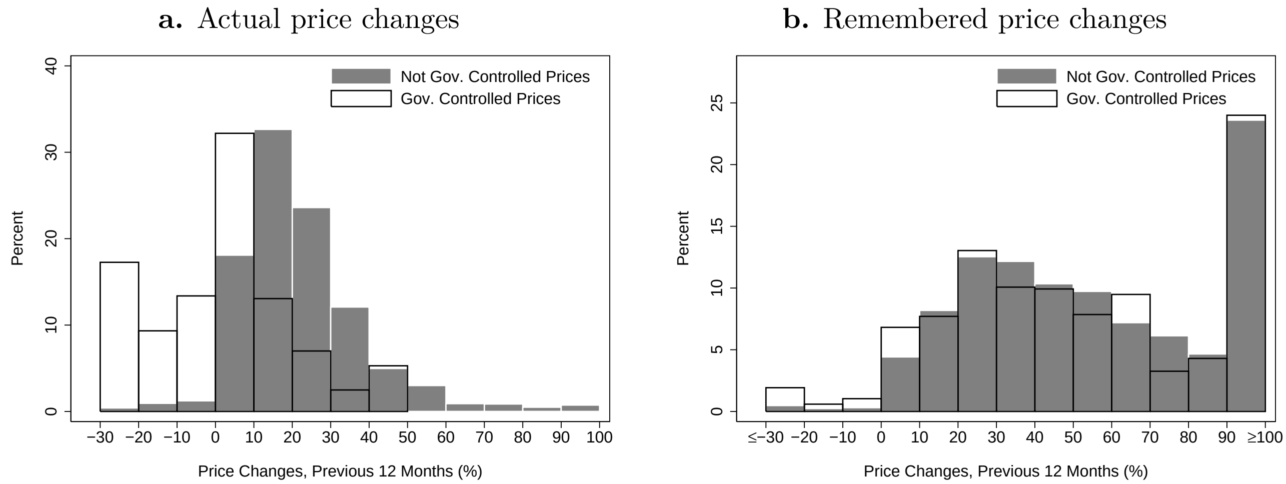

Notes: The total number of observations is 1,140. Respondents to this survey were asked about the price changes with respect to a year earlier of products they had just purchased at a supermarket (remembered price changes), and we matched those products with their current and past prices of the same products in the same stores (actual price changes; see Section 4 and Cavallo et al., 2014, for more details about the survey). Panel a. presents the distribution of actual prices changes from our database of historical supermarket prices. Panel b. represents the remembered price changes for the same products as reported by the respondents. The two figures present separate distributions for products with prices controlled by the government and those with no price controls at the time of the survey. Source: consumer intercept survey.

5 Conclusions

To understand how households learn from potentially biased statistics, we exploit data from a natural experiment and a survey experiment based on the period of government manipulation of inflation statistics in Argentina. 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% 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 (s.e. 7.19), compared to an average of 4.22 (s.e. 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 bias. This lesson may be relevant for inflation statistics as well as other governmental statistics.

For policymakers, our results are useful to better understand the process of belief formation in contexts of increasing inflation. In particular, as the Federal Reserve and other Central Banks return to a more neutral monetary policy, these results, together with our other work, imply that the process by which 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 decision-making processes.

Ultimately, more empirical evidence is needed to understand how inflation and other expectations form. Experimental evidence, in particular, can shed light on the questions raised by a large theoretical literature on models of expectation formation such as adaptive, rational, natural (Furster et al., 2010), and diagnostic expectations (Bordalo et al., 2015).. In particular, the evidence from Argentina suggests that perceived biases in the signals are capable of creating asymmetric responses in expectations. Future work should try to explain why this happens and clarify the circumstances under which individuals tend to over-react and under-react to information.

Finally, in this paper, we focus on the effect of manipulation of 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 at reducing the average perceived and expected inflation rate, the manipulation of statistics may have increased individual subjective uncertainty about the inflation rate. 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 work is needed to understand the effects of potentially biased statistics on subjective uncertainty.

References

[1] Armantier, O., Bruine de Bruin, W., Topa, G., der Klaauw, V., Wilbert, H., and Zafar, B. (2012). "Inflation expectations and behavior: Do survey respondents act on their beliefs?" Federal Reserve Bank of New York Staff Report No. 509.

[2] Armantier, O., Nelson, S., Topa, G., van der Klaauw, W. and Zafar, B. (2014). "The Price Is Right: Updating of Inflation Expectations in a Randomized Price Information Experiment," Review of Economics and Statistics, forthcoming.

[3] Atkeson, A. and Ohanian, L, (2001). "Are Phillips curves useful for forecasting inflation?," Quarterly Review, Federal Reserve Bank of Minneapolis, Winter issue, pp. 2-11.

[4] Badarinza, C. and Buchmann, M. (2009). "Inflation Perceptions and Expectations in the Euro Area: The Role of News," ECB Working Paper 1088.

[5] Bates, J. M. and Gabor, A. (1986). "Price perception in creeping inflation: Report on an enquiry," Journal of Economic Psychology, Vol. 7, pp. 291–314.

[6] Bernanke, B. (2007). "Inflation Expectations and Inflation Forecasting", Speech at the Monetary Economics Workshop of the NBER Summer Institute, Cambridge, Massachusetts, July 10, 2007. Available at: http://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm, last accessed on December 2012.

[7] Bevacqua, G. and Salvatore, N. (2009). "Argentina. La reconstrucción de la serie de inflación minorista: El IPC City." Buenos Aires City, Facultad de Ciencias Económicas, Universidad de Buenos Aires.

[8] Bishop, C. (2006), "Pattern recognition and machine learning." New York: Springer.

[9] Blanchflower, D. and Mac Coille, C. (2009). "The formation of inflation expectations: an empirical analysis for the UK," Paper presented at Banco do Brasil X1 Annual Inflation Targeting Seminar, Rio de Janeiro.

[10] Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2015). "Diagnostic Expectations and Credit Cycles," Working Paper.

[11] Branch, W. (2004). "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations," Economic Journal, Royal Economic Society, vol. 114(497), pages 592-621, 07.

[12] Bruine de Bruin, W., van der Klaauw, W., Downs, J., Fischhoff, B., Topa, G. and Armantier, O. (2010). “The effect of question wording on reported expectations and perceptions of inflation,” Staff Reports 443, Federal Reserve Bank of New York.

[13] Bruine de Bruin, W., van der Klaauw, W., Topa, G. (2011), “Expectations of inflation: The biasing effect of thoughts about specific prices,” Journal of Economic Psychology, Volume 32, Issue 5, Pages 834-845.

[14] Burke, M. and Manz, M. (2011). “Economic literacy and inflation expectations: evidence from a laboratory experiment,” Public Policy Discussion Paper 11-8, Federal Reserve Bank of Boston.

[15] Camacho, M., Dal Bianco, M. and Martinez-Martin, J. (2015). “Toward a more reliable picture of the economic activity: An application to Argentina.” Economics Letters 132129-132.

[16] Carrillo, P.E., and Shahe Emran, M. (2012). “Public information and inflation expectations: Microeconometric evidence from a natural experiment,” Review of Economics and Statistics, Vol. 94 (4), pp. 860-877.

[17] Carroll, C. (2003). “Macroeconomic Expectations of Households and Professional Forecasters,” Quarterly Journal of Economics, Vol. 118 (1).

[18] Cavallo, A.; Cruces, G. and Perez-Truglia, R. (2014), “Inflation Expectations, Learning and Supermarket Prices: Evidence from Field Experiments,” NBER Working Paper 20576.

[19] Cavallo, A. (2013). “Online and Official Price Indexes: Measuring Argentina’s Inflation”, Journal of Monetary Economics. Volume 60, Issue 2, pp. 152-165.

[20] Cavallo, A. and Rigobon, R. (2016). “The Billion Prices Project,” Journal of Economic Perspectives. Forthcoming.

[21] Coibion, O. and Gorodnichenko, Y. (2015). “Is The Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation,” American Economic Journal: Macroeconomics, Vol. 7 (1), pp. 197-232.

[22] Curtin, R. (2009), “What U.S. Consumers Know About The Economy: The Impact Of Economic Crisis On Knowledge.” Presented at the 3rd OECD World Forum On “Statistics, Knowledge And Policy,” Busan, Korea - 27-30 October 2009.

[23] Eurobarometer (2008). “Europeans’ knowledge of economic indicators,” Special Eurobarometer 323, Wave 67.2 – TNS Opinion & Social, European Commission, Brussels.

[24] Fuster, Andreas, David Laibson, and Brock Mendel (2010). “Natural Expectations and Macroeconomic Fluctuations,” Journal of Economic Perspectives, 24(4): 67-84.

[25] Goerg, S. and Kaiser, J. (2009). “Nonparametric testing of distributions—the Epps–Singleton two sample test using the empirical characteristic function,” Stata Journal, Vol. 9(3), pp. 454-465.

[26] Hellwig, C. (2005). “Heterogeneous Information and the Benefits of Public Information Disclosures,” UCLA Economics Online Papers No. 283.

[27] Jonung, L. (1981). “Perceived and expected rates of inflation in Sweden,” The American Economic Review, Vol. 71 (5), pp. 961-968.

[28] Lamla, M.J. and Lein, S.M. (2008). “The Role of Media for Consumers’ Inflation Expectation Formation.” KOF Swiss Economic Institute, Working Paper No. 201.

[29] Lee, D. (2009). “Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects.” Review of Economic Studies 76(3).

[30] Yeyati, Eduardo Levy and Novaro, Marcos (2013). “Vamos por todo: Las diez decisiones más polémicas del modelo.” SUDAMERICANA.

[31] Malmendier, U. and Nagel, S. (2013). “Learning from Inflation Experiences,” Working Paper, Berkeley.

[32] Mankiw, N.G. and Reis, R. (2002). “Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve,” Quarterly Journal of Economics, Vol. 117 (4), pp. 1295-1328.

[33] Mankiw, N.G., Reis, R. and Wolfers, J. (2003). “Disagreement About Inflation Expectations.” In NBER Macroeconomics Annual 2003, ed. by M. Gertler, and K. Rogoff.

[34] Manski, C. (2004), “Measuring expectations,” Econometrica, Vol. 72 (5), pp. 1329–1376.

[35] Michalski, T. and Stoltz, G. (2013). “Do Countries Falsify Economic Data Strategically? Some Evidence That They Might.” Review of Economics and Statistics, Vol. 95, No. 2, Pages 591-616.

[36] Morgenstern, Oskar von (1963). “On the Accuracy of Economic Observations.” Princeton, NJ: Princeton University Press.

[37] Morris, S. and Shin, H.S. (2002). “Social Value of Public Information,” The American Economic Review, Vol. 92 (5), pp. 1521-1534.

[38] Norris, F. (2014). “Doubting the Economic Data? Consider the Source.” New York Times, November 6, High & Low Finance Op-Ed Column.

[39] Paredes, J., Pérez, J. and Pérez Quirós, J. (2015). "Fiscal targets. A guide to forecasters?", CEPR Discussion Paper 10553.

[40] Rauch, B., Göttsche, M., Brähler, G. and Engel, S. (2011). "Fact and Fiction in EU-Governmental Economic Data." German Economic Review, 12: 243–255

[41] Roos, M.W.M. and Schmidt, U. (2012). "The Importance of Time-Series Extrapolation for Macroeconomic Expectations," German Economic Review, Vol. 13(2), pp. 196–210.