

TELL ME SOMETHING I DON'T ALREADY KNOW: LEARNING IN LOW AND HIGH-INFLATION SETTINGS

Michael Weber
U. Chicago and NBER

Bernardo Candia
UC Berkeley

Hassan Afrouzi
Columbia University
and NBER

Tiziano Ropele
Bank of Italy

Rodrigo Lluberas
Universidad ORT
Uruguay

Serafin Frache
Universidad de
Montevideo

Brent Meyer
Federal Reserve Bank
of Atlanta

Saten Kumar
Auckland University of
Technology

Yuriy Gorodnichenko
UC Berkeley and
NBER

Dimitris Georgarakos
European Central Bank
and CEPR

Olivier Coibion
UT Austin and NBER

Geoff Kenny
European Central Bank

Jorge Ponce
Central Bank of Uruguay

Abstract

Using randomized control trials (RCTs) applied over time in different countries, we study whether the economic environment affects how agents learn from new information. We show that as inflation rose in advanced economies, both households and firms became more attentive and informed about publicly available news about inflation, leading them to respond less to exogenously provided information about inflation and monetary policy. We also study the effects of RCTs in countries where inflation has been consistently high (Uruguay) and low (New Zealand) as well as what happens when the same agents are repeatedly provided information in both low- and high-inflation environments (Italy). Our results broadly support models in which inattention is an endogenous outcome that depends on the economic environment.

JEL: E3, E4, E5

Keywords: Inattention, RCTs, inflation expectations.

Acknowledgement: We thank the Fama-Miller Center and the Initiative on Global Markets, both at the University of Chicago Booth School of Business for financial support for conducting the surveys as well as the NSF (SES #1919307). We thank Antar Diallo and Justus Meyer for research assistance and Shannon Hazlett at NielsenIQ for her assistance with the collection of the PanelViews Survey. We are grateful to conference/seminar participants at many institutions for helpful comments. The views presented here do not represent those of the European Central Bank, the Bank of Italy, the Central Bank of Uruguay, or the Federal Reserve Bank of Atlanta. Ordering of author names is randomized. We thank Scott Baker for sharing data on media coverage of inflation. The randomized control trial is registered at the AER RCT Registry (#AEARCTR-0011948). The following datasets are used in the analysis: BIE, ECB-CES, ECB Data Portal, EEE, Gallup News Service, IPC, Nielsen, OECD, SCE, SPF, SIGE, MSC.

“Tell me and I forget. Teach me and I remember. Involve me and I learn.” B. Franklin

I Introduction

The environment in which we live shapes our behavior and beliefs. Those who grew up during the Great Depression, for example, tend to be more wary of taking on financial risk (Malmendier and Nagel 2011). Those who lived through hyperinflations are similarly scarred by the experience and are less likely to invest in risky assets (Fajardo and Dantas 2018). While the effects of historical episodes on behavior can be studied ex-post, it is more challenging – but of paramount importance for policy making – to study how the beliefs of individuals evolve in real time. In this paper, we study how a changing inflation environment alters the learning process of individuals.

To characterize how learning evolves with the economic environment, we bring together a wide range of randomized control trials (RCTs) across countries and time in which some individuals were provided with publicly available information about inflation, such as the most recent inflation rate or the central bank’s target. The extent to which individuals adjust their economic expectations in response to this information tells us about their learning process and prior knowledge about inflation. In a nutshell, when economic agents place a lot of weight on the provided information, this indicates that it is new to them, a sign of having previously been inattentive to publicly available information about inflation. When individuals are already informed about such news, the information provided should have little effect on their beliefs. Thus, the strength of the response of expectations to exogenously provided information speaks directly to the inattentiveness of individuals to such news.

We show that as inflation has increased to historically high levels in the past few years, households and firms in the U.S. and euro area have become less responsive to information treatments involving information about inflation. According to our theoretical framework, four channels could explain this time variation in treatment effects: changing uncertainty about inflation, changing persistence of inflation, changing trust in inflation statistics or monetary policy, or changing prior knowledge of publicly available information. We provide new evidence that the latter explanation (i.e. changing knowledge about publicly available information about inflation) provides the best explanation for the empirical patterns that we document. As the inflation environment has changed, so too has the degree of inattention of individuals to publicly available news about inflation. Two underlying mechanisms could explain this change, according

to our model. If high inflation is perceived as more volatile inflation, then agents will become more willing to pay the fixed cost of seeking out inflation statistics. Alternatively, if the cost of acquiring this information declines with higher inflation, e.g. because news reporting about inflation becomes more prevalent when inflation is high, then again more respondents would be aware of this information prior to participating in the survey. Our results therefore complement other recent studies that have examined the changing degree of inattention as inflation rises (e.g., Bracha and Tang forthcoming, Korenok, Munro and Chen 2023, Pfäuti 2023).

Assessing changes in the degree of inattention across different inflation regimes is empirically challenging. In a changing environment, economic agents are subject to idiosyncratic and aggregate shocks that affect them differently due to their heterogeneous characteristics. As a result, economic agents' time-varying unobserved characteristics (e.g., economic sentiment, risk aversion) correlate with prevailing conditions and are likely to confound the inference on their attention to inflation. Our key innovation relative to existing studies is that we rely on a sequence of RCTs to assess how inattention changes across economic environments. By design, the random allocation of subjects (and their unobserved characteristics) between treatment and control groups ensures that the role of attention can be consistently estimated at each given point in time and allows us to obtain reliable comparisons across inflation regimes.

To this end, we construct a unique collection of many such RCTs fielded in nationally representative surveys of households and firms for different countries and periods to speak directly to the changing degree of attention. Our first setting for doing so is a sequence of RCTs applied to surveys of U.S. households participating in the Nielsen Homescan Panel, starting in 2018Q2, when inflation was close to 2%, and continuing through much of 2021 to 2023, the period in which U.S. inflation rose sharply. We show that as inflation rose, survey participants responded significantly less to exogenously provided information about inflation, consistent with them becoming more informed. The change in the effect is particularly strong for treatments involving recent inflation rates, indicating that households have been paying much more attention to inflation dynamics, and is smaller for treatments involving the Federal Reserve's inflation target, indicating that learning about monetary policy has been more limited. Using five different RCTs implemented first in the Netherlands (in 2018Q2) and then in the euro area using the European Central Bank's (ECB) Consumer Expectations Survey (CES) from 2021 to 2023, we

similarly find that European households' response to information about inflation fell sharply as the inflation rate increased. Finally, using two RCTs conducted in the Atlanta Fed's Business Inflation Expectations survey in 2019 and 2023, we again document a decline in the responsiveness of U.S. firms to exogenously provided information as the inflation rate increased.

Why necessarily attribute this time variation in treatment effects to a different inflation environment? First, we provide evidence based on the ECB's CES that 60% of households surveyed in 2023M1 reported that they were paying more attention to inflation when inflation was high than they had previously. Furthermore, households that report being attentive to inflation have expectations and perceptions of inflation that are much closer to actual levels of inflation and generally respond significantly less to information treatments than do households that report paying little attention to inflation. Second, we use four RCTs from firms in Uruguay to study the effects of repeated information treatments in an environment where annual inflation has consistently been high (approximately 8%) during the 2018-2023 period. We show that Uruguayan firms' short-term inflation expectations did *not* respond to information treatments about recent inflation or the central bank's inflation target in 2018, 2019 and 2023, in line with the notion that agents in higher inflation environments consistently choose to pay more attention to inflation. Third, we use four RCTs applied to firms in New Zealand from 2014 to 2019, when inflation was consistently low. We find for this setting that all information treatments had large and powerful effects on the expectations of these firms, in agreement with the notion that agents in low inflation environments consistently choose to pay little attention to inflation. Fourth, using repeated quarterly RCTs applied to a panel of firms in Italy over a decade, we show that, again, the magnitude of the estimated effects of information treatments fell as the inflation rate rose. Finally, pooling all RCTs across countries and time, we find a clear negative relationship between the level of inflation and treatment effects.

Our paper builds on a growing literature that applies RCTs in macroeconomics to study how new information shapes expectations and how these expectations subsequently affect economic decisions. Much of this literature has focused on inflation expectations (e.g., Armantier et al. 2016) as we do here, but others have applied similar techniques to study expectations of housing prices (Armona, Fuster and Zafar 2019, Chopra, Roth and Wohlfart 2023), income expectations (D'Acunto et al. 2020), the state of the business cycle (Roth and Wohlfart 2020), asset prices (Beutel and Weber 2022), monetary policy (Coibion et al. 2023a), economic

uncertainty (Coibion et al. 2022, Kumar et al. 2023), and other topics. These studies typically focus on a single RCT to generate exogenous variation in the beliefs of treated individuals relative to an untreated control group, potentially raising concerns about external validity if a similar RCT were to be implemented in a different context. Relative to these studies, our main contribution is to consider a large number of comparable RCTs applied to households and firms and in different countries, periods and economic environments. As a result, we shed more light on the state-dependence of inattention to inflation. Our results therefore inform policymakers on how anchored inflation expectations are and how powerful policy communication can be.

Our paper is also closely related to recent work studying the time variation in inattention paid by individuals to economic conditions. Coibion and Gorodnichenko (2015) estimated time variation in information rigidities of professional forecasters, showing that information rigidities went up during the Great Moderation. Goldstein (2022) finds that inattention falls after large shocks. Bracha and Tang (forthcoming) focus on inattention by U.S. households to inflation, as measured by people saying “I don’t know” when asked about current inflation levels, and show that this metric historically declines when inflation is higher.¹ Korenok, Munro and Chen (2023) show that, across many countries, Google searches for “inflation” rise with the level of inflation whenever inflation exceeds a threshold around 4%. Pfäuti (2023) estimates how strongly inflation expectations of households and professionals in the U.S. respond to past forecast errors and shows that higher inflation periods are associated with larger responses to past errors, consistent with changing inattention. Other papers document that inattention to broader macroeconomic conditions is procyclical (An, Abo-Zaid and Shen 2023, Song and Stern 2023, Flynn and Sastry 2023 and Link et al. 2023b). Relative to these papers, we use the response of expectations to exogenously provided information in RCTs to measure inattention across countries and environments. Our RCT-based findings complement these other papers by illustrating the endogenous nature of inattention.

Finally, our paper builds most closely on the path-breaking work of Cavallo, Cruces and Perez-Truglia (2017). They compare a treatment providing information about recent inflation to college graduates and supermarket shoppers in Argentina, where inflation was over 20% at the time of the survey, and to crowd workers on Amazon Mechanical Turk in the U.S., where inflation

¹ In related work, Binder (2017) documents that one can use rounding of reported inflation forecasts to measure knowledge and uncertainty about inflation.

was about 2%. They document a striking difference in how strongly respondents in the two countries react to the public information about inflation: Argentine individuals placed far less weight on the provided information and more weight on their priors than U.S. individuals, consistent with people living in a high-inflation environment being more attentive to inflation.² Like them, we compare the effects of RCTs in low- and high-inflation environments to characterize how the level of inflation affects how attentive individuals are. Yet, due to the much larger number of RCTs available to us, we can address some limitations associated with this prior work. For example, because there are many differences between Argentina and the U.S., one cannot necessarily attribute the difference in the effects of the information treatments estimated at a given point in time to the level of inflation. In contrast, because we study the changing effects of RCTs *within a country* over time, we can more precisely identify the role of the inflation environment in driving inattention. Furthermore, we can do so for both households and firms in nationally representative samples. In addition, we use a theoretical model to discipline our empirical analysis and distinguish among possible mechanisms. Overall, our results strongly support the view of Cavallo, Cruces and Perez-Truglia (2017) that the inflation environment has first-order effects on how attentive individuals are to inflation developments.

The paper is organized as follows. Section II describes the randomized provision of information and how the results of RCTs speak to the inattention and optimal information choice of economic agents through the lens of a theoretical model. Section III presents empirical evidence for U.S. households, euro area households, as well as U.S. firms and examines the underlying theoretical mechanisms that could be at work. Section IV considers additional evidence from firms in Uruguay, firms in New Zealand, and firms in Italy. Section V presents results pooled across all RCTs, while Section VI concludes.

II Inattention, Information Treatments and the Economic Environment

When processing information is costly to agents, either because of the opportunity or cognitive costs involved, they will naturally make decisions about how much attention to allocate to

² A related result is in Link et al. (2023a) who rely instead on cross-sectional variation in inattention within a country. They study the effects of an information provision experiment in Germany that was applied to both households and firms. They show first that firms are overall better informed about recent conditions than households. They then find that firms respond less to the provided information than households, again consistent with the notion that more informed agents are less responsive to new information.

different areas that may affect them. The macroeconomic environment is one such domain. When economic conditions are volatile or risky, agents may choose to pay more attention to their economic environment than during normal times.

2.1 Existing Evidence of Time-Varying Inattention

To what extent do we see variation in inattention as economic conditions change? Bracha and Tang (forthcoming) study this question for U.S. households participating in the University of Michigan's Survey of Consumers (MSC). Using the phrasing of the inflation expectations question, Bracha and Tang (forthcoming) note that one can identify the fraction of households that anticipate constant inflation but do not know the current inflation rate. The latter can be interpreted as one measure of inattention, and they show that this measure of inattention is greater when U.S. inflation is lower. A closely related measure of inattention is to compare households' reported perceived inflation rates with actual inflation rates, the idea being that attentive households would have better knowledge of recent inflation than inattentive households. In Figure 1, we plot the perceived inflation rates of U.S. households (measured using the Nielsen survey described in Section 3.1) against actual inflation (Panel A) as well as that of euro area households (Panel B) using the CES (described in Section 3.2). In both cases, we see that households significantly overestimated inflation when inflation rates were low but average perceptions got very close to actual inflation once inflation started rising. Korenok, Munro and Cheng (2023) use the intensity of Google searches about inflation to measure how attentive households are to inflation and find that, in many countries, attentiveness increases with the level of inflation once inflation exceeds a threshold. Pfäuti (2023) studies how strongly expectations of households and professionals in the U.S. respond to past forecast errors, a measure of inattention derived from theoretical models. He finds that higher inflation periods are associated with larger responses to past forecast errors. Coibion and Gorodnichenko (2015) show that the predictability of forecast errors stemming from ex-ante forecast revisions provides another metric of how attentive agents are. They find that U.S. professional forecasters' attentiveness declined during the Great Moderation. Goldstein (2022) uses a similar approach to study time variation in inattentiveness of professional forecasters in Israel. Borraz, Orlik and Zacheo (2023) emphasize that firms in Uruguay have consistently been well informed about inflation during a period in which inflation was consistently high. Focusing on inattention to broader economic conditions, recent

papers have documented the countercyclicality of attention (An, Abo-Zaid and Shen 2023, Song and Stern 2023, Flynn and Sastry 2023 and Link et al. 2023b).

In Figure 2, we provide additional evidence in the same spirit but from households in the euro area showing that their attentiveness to inflation has increased as the level of inflation in the euro area has risen. In the 2023M1 wave of the CES, households were asked how attentive they were to inflation prior to any information treatments. As shown in Panel A, only about 20% of households reported that they paid no attention or little attention to inflation, indicating that most households were paying at least some attention to inflation. Households were also asked whether they were paying more or less attention to inflation compared to 12 months prior, when inflation was lower. As shown in Panel B, over 60% of households answered that they were paying more attention to inflation, consistent with inattention varying with the level of inflation. Furthermore, as shown in Panel C, inattention is not innocuous: those households who reported paying more attention to inflation tended to have forecasts closer to recent inflation levels (8.6% in January 2023). However, more attention does not seem to translate into more confidence: Panel D shows that uncertainty in inflation forecasts does not vary systematically with attention.

2.2 Measuring Inattention through Information Treatments

While the accuracy of the perceived level of recent inflation is a natural measure of inattention, it should be viewed as only suggestive because inattention is self-reported and causality toward forward-looking beliefs cannot be established. Furthermore, it does not tell us how much, or even whether, new information would change expectations, which is of direct interest for policymaking and communication. Instead, our aim is to measure the attentiveness of economic agents through their responsiveness to exogenously provided information about inflation and monetary policy. In this approach, survey respondents are assigned either to a control group that receives no information or to a treatment group that is provided with publicly available information (e.g., Armantier et al. 2016, Cavallo, Cruces and Perez-Truglia 2017, and Coibion, Gorodnichenko and Kumar 2018). The effect of the treatment on beliefs can then be evaluated through the following regression specification of posterior beliefs on prior beliefs:

$$posterior_i = \alpha + \beta \times prior_i + \delta \times \mathbb{I}_i + \gamma \times \mathbb{I}_i \times prior_i + error_i \quad (1)$$

where \mathbb{I}_i is an indicator variable equal to one if agent i is in the treatment group and thus receives a signal. In principle, one should expect $\alpha = 0$, $\beta = 1$, and $\gamma \in [-1,0]$. Figure 3 shows a visual

representation of one such experiment on inflation expectations of U.S. households participating in the Nielsen Homescan Panel (we provide more details on this survey in Section 3.1; see also Coibion, Gorodnichenko and Weber 2022). All participants are first asked for their inflation expectations using a distributional question (assign probabilities to pre-specified bins of possible future inflation rates)³ and then are assigned to either a control group that does not receive any additional information or one of several treatment groups which receive information. The three treatments in Figure 3 reflect being informed about recent inflation, the Fed’s inflation target, or the FOMC’s inflation forecast.⁴ Finally, all respondents are asked to provide their inflation expectations again, this time through a point forecast. In equation (1), the coefficient β represents the relationship between prior and posterior beliefs of the control group. As said above, one would expect the slope coefficient to be one. However, since priors and posteriors are measured using two different questions, it is not uncommon for the estimated slope to differ from one and in this case the estimated slope is 0.85 and statistically different from one.⁵

Learning by households in this context is best captured by γ which measures the change in the slope of the relationship between priors and posteriors for the treated groups. If the provided information has no effect on beliefs, γ will be equal to zero and the slope linking priors and posteriors will be the same as for the control group. However, a negative γ indicates that the treatment group is placing less weight on their priors and more weight on the new information. When $\beta + \gamma = 0$, households are placing all the weight on the provided signal in forming their posteriors and none on their prior beliefs. The fraction of β that is being offset by γ is therefore the key metric that allows us to assess how household beliefs change when presented with new information. In Figure 3, it is immediately clear that the slope for each treatment group is much flatter than for the control group. In each case, the slope coefficient is approximately 0.2, indicating that households are placing a lot of weight on the newly provided information and very

³ When we compute implied means and standard deviations, we use mid-points of the bins. For the top bin (inflation will be greater than 12%) we use 14% as the mid-point. For the bottom bin (deflation will be greater than 12%), we -14% as the mid-point.

⁴ Because high inflation is often associated with volatile inflation, one should also have treatments that target the second moment of households’ beliefs (see e.g. Kumar et al. 2023) to separate level vs. volatility effects of inflation. Unfortunately, we do not have such treatments in our sample. As a result, we estimate the “total” effect of changes in the level and volatility of inflation. We hope that future work can address this limitation of our analysis.

⁵ RCTs often use two different question formulations to measure priors and posteriors because asking survey participants to answer the exact same question multiple times in the same survey can lead to increased panelist attrition rates and raises the concern of survey demand effects (see Haaland et al. 2023).

little on their priors when forming their posterior beliefs. However, because the slope coefficient for the control group is less than one, we cannot directly interpret the estimated γ as capturing how household beliefs change when presented with the new information. Furthermore, as we discuss later, some experiments measure posteriors in subsequent waves rather than immediately after the information provision.⁶ In this case, β can be less than one as information decays over time.⁷ Hence, one needs to normalize $\hat{\gamma}$ by the estimated slope of the control group to recover the effective weight on priors. As a result, we will focus on $\hat{\gamma}/\hat{\beta}$ (i.e., the scaled change in slope) as the most informative metric of how inattentive agents are, that is, how much flatter the relationship between priors and posteriors is for the treatment group *relative to the control group*.

Our empirical strategy consists of studying how these information treatment effects vary across different inflation environments. This approach builds explicitly on (i) Armantier et al. (2016) in considering settings in which some randomly selected survey participants are provided with information about inflation or monetary policy and comparing their posterior expectations to those of a control group which were not provided with such information; (ii) Cavallo, Cruces and Perez-Truglia (2017) in comparing the effects of these RCTs across countries to assess the role that the inflation environment plays in explaining how informed economic agents are about recent inflation dynamics; and (iii) Coibion, Gorodnichenko and Kumar (2018) in using the weight on the prior to measure the sensitivity to signals about inflation. Unlike these studies, however, we can do these comparisons across a number of different countries and agents as well as within a country over time, which allows us to effectively control for country-specific fixed effects and more precisely identify the role of inflation in determining how informed economic agents are. Table 1 summarizes the countries and surveys that we will rely on for this purpose.

2.3 Theoretical Predictions for Information Treatment Effects

Before turning to the empirical results, we first consider what theory predicts about the size of our estimated treatment effects in different inflation environments. To build intuition and preserve

⁶ Although some variation in RCT design across surveys exists, the design is generally fixed within a survey and thus we can compare results over time.

⁷ For example, consider forecasting x_{t+1} that follows an AR(1) process $x_t = \rho x_{t-1} + e_t$ with $\rho \in (0,1)$. If posterior beliefs are measured one period later, the slope coefficient on the prior for the control group is $\beta = \rho < 1$ rather than $\beta = 1$.

tractability, we examine a two-period framework with a continuum of households, where time is indexed by $t \in \{0,1\}$.⁸ Proofs of all Propositions are included in Weber et al. (2024).

A Consumption Choice Problem with Arbitrary Information Sets. Suppose each household receives a nominal income of W at time 0 and can spend it on consumption across the two periods. Households know the price of the consumption good C at $t = 0$, normalized at $P_0 = 1$, and can save some of their nominal income in cash, denoted by M , to purchase with it at period 1 at nominal price P_1 . We assume that inflation at $t = 1$ is given by

$$\pi_1 = \rho\pi + u, \quad \pi \sim N(0, \sigma_\pi^2), \quad u \sim N(0, \sigma_u^2), \quad \sigma_\pi^2 > 0, \sigma_u^2 > 0$$

where π is the component of $\pi_1 \equiv P_1/P_0 - 1$ that is drawn by nature in period 0, and u is an unanticipated shock to π_1 that is drawn by nature in period 1. Therefore, $\rho > 0$ captures the persistence of inflation from period 0 to period 1. The household receives utility of $u(C_0) + u(C_1)$, where $u(C) = \frac{C^{1-\psi}-1}{1-\psi}$, with $\psi > 0$ being the coefficient of relative risk aversion, and C_0, C_1 are consumption levels in periods 0 and 1. Finally, for simplicity we assume that households perfectly observe π_1 at the beginning of period 1—so that C_1 is measurable in π_1 —but are rationally inattentive in period 0 and optimally inform themselves about π_1 in a sense that we will make precise below. For now, we note that given an arbitrary information set S_i , the problem of household i at time 0 is

$$V(S_i) \equiv \max_{C_{i,0}, C_{i,1}(\pi_1)} \mathbb{E} \left[u(C_{i,0}) + u(C_{i,1}(\pi_1)) \mid S_i \right] \text{ s. t. } C_{i,0} + M_i \leq W, C_{i,1}(\pi_1) \leq \frac{M_i}{1 + \pi_1}$$

where $V(S_i)$ captures the value of the information set S_i in terms of the quality of the household's consumption choice. We can then prove the following proposition:

Proposition 1. *Let \mathbb{S} denote the information set of an agent that perfectly observes all available information at time 0, including π . Then, a quadratic approximation to a household's ex-ante consumption equivalent losses from imperfect information set $S_i \subset \mathbb{S}$, around a non-stochastic point with $C_0 = C_0^*$ and $\pi_1 = 0$, is given by*

$$\mathbb{E}_0 \left[\frac{V(S_i) - V(\mathbb{S})}{u'(C_0^*)C_0^*} \right] \approx -\frac{B\rho^2}{2} \mathbb{E}_0[(\mathbb{E}[\pi|S_i] - \pi)^2] = -\frac{B\rho^2}{2} \text{Var}(\pi|S_i) \quad (2)$$

where \mathbb{E}_0 is the expectation operator at time 0 before π is realized and $B \equiv \frac{\psi(1-\psi^{-1})^2}{2}$.

⁸ For an infinite horizon dynamic model with inattention, we refer the reader to Mackowiak and Wiederholt (2024).

Equation (2) shows that the household's ex-ante losses from imperfect information S_i are proportional to the variance of the predictable part of inflation π conditional on this information set. The reason that uncertainty about π leads to ex-ante expected losses is that the household's optimal consumption-saving decision depends on inflation in period 1, which determines the return on their cash-holdings. Since π_1 depends on π only through the persistence parameter ρ , the benefit of learning about π should obviously increase with this persistence, captured by ρ^2 in $B\rho^2/2$. Furthermore, the parameter B , which depends on the degree of risk aversion, determines the curvature of the household's utility and thus also contributes to the value of information.⁹

Treatment Effects Under Arbitrary Information Sets. In mapping the survey to the model, we assume that households participate in the surveys after they have acquired some information about the predictable component of π_1 , denoted by π above. Hence, the control group consists of households that have already acquired *some* information about inflation. Thus, before specifying the information acquisition problem of households in period 0 and characterizing their optimal information set, it is useful to map our setup to the structure of the survey and derive how *arbitrary* information sets at time 0 will shape the treatment effects that we identify.

Formally, suppose household i observes a subset of signals S_i from a set of available Gaussian signals about π at time 0, denoted by \mathbb{S} , and then rationally forms her posterior belief given the joint (Gaussian) distribution of π and S_i , with conditional mean of π being:

$$\pi_i \equiv \mathbb{E}[\pi_1 | S_i] = \text{Cov}(\vec{S}_i', \pi_1) \text{Var}(\vec{S}_i)^{-1} \vec{S}_i$$

where π_i is the expectation of household i about inflation in period 1 conditional on S_i , and $\vec{S}_i \equiv \text{vec}(S_i)$ is the vectorized version of the information set S_i . For now we continue to think of S_i as being an arbitrary finite set of Gaussian signals about π .

At the treatment stage in the survey, a researcher picks a signal $S_p = \pi + \nu_p \in \mathbb{S}$, $\nu_p \perp \pi$, $\nu_p \sim N(0, \sigma_{\nu,p}^2)$, $\sigma_{\nu,p}^2 > 0$, about π and provides it to a random sample of the agents who form the treatment group, which we denote with T . We assume all agents in T perfectly observe S_p and

⁹ We note that B is non-monotonic in ψ as it controls both risk aversion and intertemporal substitution: at $\psi \rightarrow 1$, income and substitution effects fully offset each other, and the household's decision is independent of the realization of π_1 , captured by the fact that information, in this case, has no value at $B = 0$. However, as ψ deviates from 1, information gains value ($B > 0$) as the household's optimal consumption depends on π_1 , with either of income or substitution effects dominating depending on whether $\psi > 1$ or < 1 .

update their beliefs based on Bayes' law. While v_p is independent of π , we also assume that it can be correlated with agents' signals in any S_i only through S_p , i.e., $v_p \perp (\pi, S_i \setminus \{S_p\})$. Note that if S_p is a component of S_i , i.e. the agent has already seen S_p in the pre-treatment stage, it follows that the posterior after the treatment should be the same as the pre-treatment belief: $\tilde{\pi}_i \equiv \mathbb{E}[\pi_1 | S_i, S_p] = \mathbb{E}[\pi_1 | S_i] = \pi_i$. Intuitively, in this case, the agent has not observed any new information and their belief should not move due to the treatment. Alternatively, if $S_p \notin S_i$, since \mathbb{S} only contains Gaussian signals about π , the implied posterior belief for treated individuals is:

$$\tilde{\pi}_i \equiv \mathbb{E}[\pi_1 | S_i, S_p] = \pi_i + \frac{\text{Cov}(S_p, \pi_1 | \vec{S}_i)}{\text{Var}(S_p | \vec{S}_i)} (S_p - \mathbb{E}[S_p | \vec{S}_i])$$

With this result, we obtain the following proposition that maps the model to the empirical specification in equation (1).

Proposition 2. *Consider any information set $S_i \subseteq \mathbb{S}$ at time 0 and suppose $v_p \perp (\pi, S_i \setminus \{S_p\})$. Then, a treated household i 's post- and pre-treatment beliefs are related according to:*

$$\underbrace{\tilde{\pi}_i}_{\text{post.}} = \underbrace{1}_{\beta} \times \underbrace{\pi_i}_{\text{prior}} + \underbrace{\frac{\rho \text{Var}(\pi | \vec{S}_i)}{\text{Var}(\pi | \vec{S}_i) + \sigma_{v,p}^2}}_{\delta} \mathbf{1}_{\{S_p \notin S_i\}} \times S_p \times \mathbb{I}_i - \underbrace{\frac{\text{Var}(\pi | \vec{S}_i)}{\text{Var}(\pi | \vec{S}_i) + \sigma_{v,p}^2}}_{\gamma} \times \mathbf{1}_{\{S_p \notin S_i\}} \times \pi_i \times \mathbb{I}_i$$

where \mathbb{I}_i is the indicator that i is treated with S_p . Furthermore, comparing this equation with the empirical specification in Equation (1), the scaled treatment effect γ/β is given by:

$$\frac{\gamma}{\beta} = - \underbrace{\frac{\text{Var}(\pi | \vec{S}_i)}{\text{Var}(\pi | \vec{S}_i) + \sigma_{v,p}^2}}_{\text{Kalman gain of } S_p \text{ conditional on } \vec{S}_i} \times \underbrace{\mathbf{1}_{\{S_p \notin S_i\}}}_{\text{control for } S_p \in S_i} \leq 0 \quad (3)$$

This result shows the two main margins of adjustment in treatment effects. One is the extensive margin, consisting of the share of people who already know the information being provided, i.e. whether or not S_p is already in the agent's information set S_i . The other is the intensive margin, which captures how strongly the uninformed respond to the information treatment. The strength of the intensive margin, operating through the Kalman gain, depends on (1) the prior uncertainty of the agents entering the survey ($\text{Var}(\pi | \vec{S}_i)$) and (2) the perceived noise in the provided treatment ($\sigma_{v,p}^2$). If changes in the economic environment affect either the Kalman gain or the likelihood that agents are already aware of the provided treatment, then treatment effects will vary.

Treatment Effect Under Optimal Information Sets. To make further progress in understanding what determines the agents’ uncertainty as well as who is already informed, we need to focus on agents’ incentives to acquire information. To this end, we present a simple model with rational inattention that disciplines the joint distribution of π , S_i , and S_p and makes predictions for how γ/β depends on the underlying incentives of the agents at the pre-treatment stage.

Intuitively, rational inattention models hinge on the idea that while agents have access to arbitrarily accurate information, they may consciously choose not to use some of it due to cognitive costs of processing such information into beliefs about inflation. This is different from S_p , which in our experiments stands for information about inflation that has already been processed in the sense described above, and thus is not subject to such cognitive costs. Put simply, agents can decide to pay a fixed cost to research official statistics—like searching on the web, acquiring professional forecasts of inflation or watching inflation-related news—or they can rely on their own information processing, e.g. use price samples from their personal experiences and use cognitive resources to convert those prices into an inflation statistic, or do both. This broader framework nests classic rational inattention models when the fixed cost of accessing official statistics becomes infinitely high. To operationalize this insight, we assume that agents prior to participating in the survey behave according to a standard rational inattention model with the additional element that they also have the option to observe S_p by paying a fixed cost $\phi \geq 0$.

With the ex-ante loss function described by Proposition 1, we assume that the cost of processing information is linear in the reduction in entropy between the prior and posterior distributions, as measured by mutual information, where the constant of proportionality, denoted by $\omega > 0$, captures the cost of processing each unit of information.¹⁰ This setting translates to a

¹⁰ This assumption is a standard specification for the cost of processing information in the rational inattention literature, but in our model, it is also desirable for an additional reason as it implies that the optimal posterior uncertainty is independent of the prior uncertainty of households about π_1 . This implication is important because, as we discuss later, we do not observe changes in the posterior uncertainty of survey participants across different inflation regimes. It is also consistent with Panel D of Figure 2, which shows that inflation uncertainty does not vary with how attentive households report being to inflation. This empirical result is consistent with a linear cost of mutual information, but not convex transformations of it. In fact, if we were to assume that the cost of information is convex in mutual information (or posterior variance) we would conclude that both posterior uncertainty and the treatment effect should increase with inflation volatility.

problem where the household decides if they want to pay the fixed cost and observe S_p and how much more information they want to process. The formal problem for choosing the optimal S_i is:

$$\min \left\{ \phi + \min_{S_p \subseteq S_i \subseteq \mathbb{S}} \left\{ \frac{1}{2} B \rho^2 \text{Var}(\pi | \vec{S}_i) + \omega I(\vec{S}_i; \pi | S_p) \right\}, \min_{S_i \subseteq \mathbb{S}} \left\{ \frac{1}{2} B \rho^2 \text{Var}(\pi | \vec{S}_i) + \omega I(\vec{S}_i; \pi) \right\} \right\}.$$

Here, the first min operator captures the decision whether or not to acquire S_p : the first argument states the rational inattention problem of the agent conditional on observing S_p and the second argument captures the rational inattention problem without directly observing S_p .

This problem nests the conventional rational inattention problem as $\phi \rightarrow \infty$.¹¹ Here, $I(X; Y) = \frac{1}{2} \ln \left(\frac{\text{Var}(X)}{\text{Var}(X|Y)} \right)$ is the mutual information between Gaussian random variables X and Y .

Finally, we assume ω is small enough that agents always process *some* information on their own—i.e., they are never in a corner solution in which their information set is empty or just S_p (e.g., agents always have some sample of prices in their information set that they use for forecasting inflation). Formally, as we show in the proof of the following proposition, the necessary and sufficient condition for this assumption to hold is $\frac{\omega}{B \rho^2} < \text{Var}(\pi | S_p)$.

Proposition 3. *Suppose $\frac{\omega}{B \rho^2} < \text{Var}(\pi | S_p)$. Then,*

1. *It is optimal for the agent to always process enough information prior to taking the survey so that their subjective uncertainty about π is set to the cost-benefit ratio $\omega/(B \rho^2)$, regardless of whether the agent chooses to observe S_p or not:*

$$\text{Var}(\pi | \vec{S}_i) = \frac{\omega}{B \rho^2}.$$

2. *The household pays the fixed cost and acquires S_p if and only if the fixed cost of observing the pre-processed signal S_p is smaller than the cognitive cost of processing the amount of information revealed by S_p about π*

$$S_p \in S_i \Leftrightarrow \phi \leq \omega I(S_p, \pi) = \frac{\omega}{2} \ln \left(1 + \frac{\sigma_\pi^2}{\sigma_{v,p}^2} \right).$$

¹¹ This broader specification is of interest to us because, in a conventional rational inattention problem, agents have no incentive to pay attention to official statistics like S_p since official statistics are weakly noisier signals about inflation than π itself. Hence, if agents can process arbitrarily precise information about π and S_p at the same cognitive cost, learning directly about inflation is always more advantageous than learning about it through the signal S_p . In such a case, one can then show that agents will never directly pay attention to S_p . Taking into account that official statistics are pre-processed makes them attractive to agents despite their inherently noisier nature.

The first part of Proposition 3 shows that the optimal uncertainty of the household about π only depends on ω, ρ and B and is independent of the other parameters of the model, including the prior volatility of inflation, σ_π^2 , and the variance of the noise in the public signal, $\sigma_{v,p}^2$. Furthermore, the fact that the optimal subjective uncertainty is independent of the decision to observe S_p is particularly interesting because it shows that observing official statistics operates only on a substitution margin, as it does not affect the final subjective uncertainty of agents once they have processed their own information because the cost of attention is separable in agents' uncertainty about inflation prior to processing information.

If agents' incentives are such that they acquire the official statistic S_p on their own prior to taking the survey, which happens when the condition in Part 2 of Proposition 3 holds, then providing the official statistics to agents in the treatment group during the survey is a redundant task that should have no effect on their beliefs. If S_p is not observed by agents prior to taking the survey, then providing the treatment group with S_p during the survey should affect their beliefs relative to the beliefs of agents in the control group. We can see this by substituting the optimal subjective uncertainty, $\text{Var}(\pi|\vec{S}_i) = \frac{\omega}{B\rho^2}$, in Equation (3), yielding:

$$\frac{\gamma}{\beta}\Big|_{i \in T} = \begin{cases} -\frac{\omega}{\omega + B\rho^2\sigma_{v,p}^2} & S_p \notin S_i \\ 0 & S_p \in S_i \end{cases} \quad (4)$$

This expression provides the precise magnitude of the treatment effect when $S_p \notin S_i$, i.e. along the intensive margin: once the households update their beliefs, they put a positive weight on the treatment signal which delivers the negative γ/β ratio. If high inflation periods are such that pre-processed signals are more informative about inflation ($I(S_p, \pi) \uparrow$) or the cost of acquiring them is lower ($\phi \downarrow$), so much that the condition in Part 2 of Proposition 3 holds, then agents would already have S_p in their information set and treating them with S_p during the survey would have no effect.

We conclude this section with the following comparative statics Proposition that formalizes the channels through which a changing inflation environment can alter estimated treatment effects.

Proposition 4. *Suppose $\frac{\omega}{B\rho^2} < \text{Var}(\pi|S_p)$. Then within this parameter region:*

1. If $\phi > \omega I(S_p, \pi)$, then the size of the treatment effect, $|\gamma/\beta|$, strictly increases with the cost of processing information ω . It also strictly decreases with the inflation persistence, ρ , the sensitivity of the loss function captured by B , and the variance of the noise in the public signal $\sigma_{v,p}^2$.
2. Starting from $\phi > \omega I(S_p, \pi)$, if the cost ϕ falls sufficiently, or alternatively inflation volatility relative to the variance of the noise in public signal, $\sigma_\pi^2/\sigma_{v,p}^2$, increases sufficiently so that $\phi < \omega I(S_p, \pi)$, then the size of the treatment effect strictly decreases.

In short, there are several mechanisms through which we may see a change in the estimated treatment effect in a higher inflation environment. If higher inflation raises the value of information (B) or the perceived persistence of the inflation process (ρ), then households would reduce their posterior variance, which would in turn reduce their sensitivity to treatments about e.g. current inflation. These effects would operate solely through the intensive margin of information acquisition. Also operating along the intensive margin, an increase in the perceived noise in the treatment ($\sigma_{v,p}^2$), coming from less trust in public statistics, would reduce the weight that individuals assign to the new information but without changing the posterior uncertainty. If higher inflation reduced the cost of processing information (ω), this would lead to a reduction in the posterior variance of households and therefore reduce treatment effects via the intensive margin. However, a reduction in ω could lead to a reduction in the share of people who acquire S_p , which would *raise* treatment effects via the extensive margin and would therefore make the overall change in the treatment effect ambiguous. Finally, operating solely through the extensive margin of the share of people knowing the treatment information going into the survey, the treatment effect could fall if inflation volatility (σ_π^2) increases with inflation or the cost of acquiring the public signal ϕ decreases enough so that $\phi < \omega I(S_p, \pi)$ holds for more respondents.

III Time-Varying Inflation and the Changing Effects of Information Treatments

In this section, we focus on RCTs applied to households and firms in the U.S. and the euro area where we have the largest sample sizes and can compare within-country estimates in low- and high-inflation regimes. Our analysis uses information treatments that provide three types of information: *i*) past inflation (π_t); *ii*) the inflation target (π^*); *iii*) the inflation forecast from the central bank

$(F_t^{CB} \pi_{t+h})$.¹² These treatments should be relevant for inflation expectations and maximize the coverage across countries and time. We report these treatments in Supplemental Appendix (henceforth SA, Appendix Figure A.9; see Weber et al. 2024).

3.1 U.S. Households

The Nielsen Homescan panel consists of approximately 80,000 nationally representative households that regularly scan their purchases and participate in occasional surveys run by Nielsen (see, e.g., D’Acunto et al. 2021). These surveys typically achieve response rates of around 20-25%, yielding survey sample sizes of 15,000-20,000 on average. Prior to the information treatments, all households are asked about their inflation expectations through a distribution question in which they assign probabilities to a range of possible inflation outcomes, following the question design from the Federal Reserve Bank of New York’s Survey of Consumer Expectations (SCE). From this question, we construct an implied mean forecast of inflation that represents the prior belief of the household. Following the information treatments, all respondents (including the control group) are asked to provide a point forecast for inflation over the next 12 months, which measures the posterior belief.

To assess how and whether inattention among U.S. households has changed over time, we rely on the fact that similar RCTs as the one in 2018Q2 described in Section 2.2 were also applied in subsequent survey waves. For example, in 2019Q1, another RCT was done in which only the information treatment with the recent inflation rate was applied. Then, three more RCTs were run in 2021, another two were done in 2022, and three more in 2023. Most of these included all three information treatments. We plot the resulting estimates of the scaled treatment effect γ/β for each wave and treatment separately in Panel A of Figure 4, along with the time series of U.S. inflation and the average inflation expectations of households participating in the Nielsen surveys.¹³ A clear pattern arises: the treatment effects remain very large (in fact even larger compared to 2018) in 2019 but fall (in absolute value) as inflation rises starting in 2021. For example, the scaled treatment effects from providing the most recent inflation rate go from around -0.75 in 2018 to -

¹² If the forecast from the central bank was not available and not used in the treatment, we use the inflation forecast from a survey of professional forecasters (SPF). The sensitivity to provided information may vary with the credibility of the information source. Thus, whether inflation forecasts come from a central bank or a survey of professional forecasters can matter. In practice, inflation forecasts from these two sources are very similar in our sample.

¹³ We present all unscaled estimates of γ_j in the Appendix. These are qualitatively the same as the scaled estimates but generally present even stronger evidence of time-variation in inattention linked to the level of inflation.

0.25 in late 2021 and early 2022, before increasing slightly in absolute value in late 2022 as inflation started to decline. While there is some sampling variation depending on the specific treatment and survey wave, the results point toward a clear pattern of declining treatment effects when inflation rises. Given that the change in the magnitude of the estimated effect is strongest for treatments involving recent inflation rather than the FOMC target or forecast, this finding suggests that households have become much more informed about recent inflation dynamics but only somewhat more knowledgeable about the Federal Reserve's inflation target.

One might worry that treatment effects may reflect a desire on the part of survey participants to please the surveyors by reporting forecasts close to the provided information (survey demand effects), without real learning taking place. There are three considerations against this view. First, there is no a priori reason to expect survey demand effects to change over time given that the RCTs are implemented in a consistent manner across survey waves and therefore cannot readily explain the time variation in treatment effects that we document. Second, demand effects are weaker in online surveys (De Quidt et al. 2018), the mode for most surveys in our data. Third, one way to address this concern is to examine the persistence of treatment effects. For example, since the Nielsen survey of households is implemented quarterly, one can consider treatment effects after three months rather than immediately after the treatment is provided to households. There is little reason to believe that survey demand effects would persist beyond the current survey that implements the RCT, so this setting provides a natural check against this alternative explanation. We do so by estimating the same specification as before but using posterior beliefs measured using the subsequent quarterly survey. We report results for scaled treatment effects in Panel B of Figure 4. While the treatment effects are smaller overall in absolute value after three months than they were contemporaneously, especially when using the inflation target or the inflation forecasts of the central bank, the same time series variation obtains: treatment effects decline in absolute value as inflation rises, converging to around zero when inflation reaches its peak. Survey demand effects are unlikely to explain this time variation.¹⁴

These results are robust to a number of reasonable variations. For example, if we focus on the unscaled size of treatment effects instead of the scaled version, the estimates are essentially

¹⁴ Treatment effects after three months are naturally smaller given that treated agents partially forget the provided information but also due to them receiving additional signals in between two survey waves.

unaffected, both in terms of instantaneous treatment effects as well as treatment effects after three months (SA Appendix Figure A.1). Another possibility is that agents learn about inflation as they participate in the survey repeatedly, as emphasized in Kim and Binder (2023). In general, the RCT set-up should be robust to this concern as survey participants with different tenures are equally present in the control and treatment groups and some panel refreshment typically takes place in online surveys. In any case, when we restrict our attention to households who have not participated in the last wave or in the last two waves, we find the same patterns (SA Appendix Figure A.2) although the precision of estimates decreases due to smaller sample sizes.¹⁵ We find similar results when we explicitly control for the number of waves in which survey respondents have participated. Nor is this pattern driven by only a subset of survey participants. When we split samples by age (SA Appendix Figure A.11), political party (SA Appendix Figure A.12), education (SA Appendix Figure A.13) or gender (SA Appendix Figure A.14), we do not find any clear differences in the time variation in treatment effects along any of these metrics. In short, these results confirm the findings of Cavallo, Cruces and Perez-Truglia (2017) that inflation treatment effects are much smaller when inflation is high and agents are attentive, but using multiple RCTs within the same country.

3.2 Examining Underlying Mechanisms

Our theoretical model points toward three possible mechanisms for this time-variation in treatment effects: changes in the uncertainty of respondents going into the survey, changes in the perceived reliability of the treatment information, or changes in the share of people who already knew the information provided in the treatment. We consider each of these in turn.

Inflation Uncertainty. One of the direct determinants of the size of the Kalman gain is the uncertainty of respondents prior to receiving the information treatment. To explain smaller treatment effects over time, we would need respondents to have become more confident in their inflation forecasts as the inflation rate increased. In our model, this could arise if with higher inflation, either the cost of processing information, ω , declines (e.g., perhaps larger price changes are more cognitively discernable when inflation is higher) or the sensitivity of losses to inflation

¹⁵ As we discuss below, we find in other surveys that responsiveness to treatments stays stable within a given inflation environment (e.g., New Zealand and Uruguay) when we use fresh draws of respondents for different survey waves.

uncertainty, captured by B , increases. Both these changes reduce the cost-to-benefit ratio of acquiring information about the predictable component of inflation, and according to Proposition 4, reduce the size of the treatment effect, consistent with our findings. Intuitively, both these changes motivate agents to acquire and process more information on their own. As a result, all households become more informed relative to low inflation periods. Importantly, this prediction goes beyond our assumption of the linear mutual information cost: any cost that allows agents to partially respond to their incentives—e.g., costs that are convex in mutual information or posterior variance—would imply that a lower cost-to-benefit ratio should translate into lower posterior variance. A key implication of this channel is therefore that uncertainty in inflation forecasts should decline when inflation rises. Panel A of Figure 5 suggests that this prediction is not supported by the data: uncertainty in inflation forecasts has been flat or, if anything, weakly increased since the start of the recent inflation spurt. Panel D of Figure 2 also indicates that inflation uncertainty is unrelated to how attentive to inflation households report being.¹⁶

One could also hypothesize an alternative but closely related mechanism that the perceived persistence of inflation has changed. Because our surveys collect not only expectations but also perceptions of inflation, we can regress expectations on perceptions wave by wave and examine whether the regression coefficient covaries with inflation. We find (Panel C of Figure 5) that the perceived persistence of inflation is increasing in the level of inflation.¹⁷ Hence, a change in perceived persistence could potentially explain a change in treatment effects over time. Indeed, Proposition 4 indicates that an increase in perceived persistence would be expected to deliver a decline in estimated treatment effects for treatments involving current inflation, consistent with results in Figure 4.

¹⁶ A related explanation could be that economic agents pay more attention to economic news during turmoil periods such as the COVID-19 pandemic which may confound our analysis of responses to information treatments in high- vs. low-inflation environments. Unfortunately, we do not have RCTs during COVID-19 with low inflation for the Nielsen data. However, our results based on other surveys suggest that, in terms of attention to inflation, the marginal contribution of COVID-19 per se seems small. For example, Italy's SIGE survey indicates that 2020Q1 may have more attention (γ/β rises) but in the 2020Q2 wave attention goes back to pre-COVID levels. We also find little if any change in Uruguay which had high inflation during COVID and outside the COVID period.

¹⁷ Using much longer time series for inflation forecasts at multiple horizons from the Survey of Professional Forecasters and the Michigan Survey of Consumers, we find that this pattern holds more generally (SA Appendix Figure A.10). This evidence also rules out an alternative potential effect of persistence where signals about past inflation, for example, are less useful for predicting future inflation and hence treatment effects should decrease with inflation.

However, a change in perceived persistence also implies differential predictions for how treatment effects would be affected depending on whether the provided information is about current or future inflation. To see this, recall that in the model, an increase in ρ is predicted to reduce posterior uncertainty about current inflation: $Var(\pi|S_i) = \omega/(B\rho^2)$. As a result, higher persistence, like a lower ω or a higher B , would lead agents to become more informed about current inflation and therefore respond less to information treatments about recent inflation. But unlike with a lower ω or a higher B , posterior uncertainty about *future* inflation would not be expected to decline, since $Var(\pi_1|S_i) = \rho^2 Var(\pi|S_i) + \sigma_u^2 = \omega/B + \sigma_u^2$. Thus, we cannot rule out the persistence channel using Panel A of Figure 5, since the constant uncertainty about future inflation is not inconsistent with a perceived change in inflation persistence. However, an implication of this constant uncertainty about future inflation is that treatment effects coming from information about future inflation over the same horizon should not vary over time. To see this, consider a treatment with the FIRE forecast for inflation in period 1: $S_p^f = E^f[\pi_1] + v_p^f = \rho\pi + v_p^f$ where $v_p^f \sim N(0, \sigma_{v,p}^f)$ and compare it with our baseline treatment with last years' inflation $S_p = \pi + v_p$. We can then see that when households do not pay the fixed cost of acquiring these signals independently, $\{S_p, S_p^f\} \not\subseteq S_i$, the treatment effect for S_p should decline with ρ (Proposition 4), but the treatment effect for S_p^f is given by $(\gamma/\beta)^f = -\frac{\omega}{\omega + B(\sigma_{v,p}^f)^2}$ which does not depend on ρ .¹⁸ Intuitively, an increase in ρ makes agents become more informed about recent inflation, reducing the effect of treating them with last year's inflation S_p , but it also increases the signal to noise ratio of S_p^f , which neutralizes the first effect and leaves the treatment effect of S_p^f unchanged. However, Panel A of Figure 4 makes clear that we observe a decline in estimated treatment effects coming from information treatments involving forecasts of future inflation or the Fed's inflation target as well as treatments involving recent inflation. Hence, the rise in perceived ρ cannot explain all of the time variation in treatment effects that we observe.

Credibility of Treatment Information. The second general mechanism which could explain time variation in treatment effects is if official statistics are perceived as less credible/informative

¹⁸ To see this, one can rescale $S_p^f = \pi + v_p^f/\rho$, which mathematically maps its treatment to our baseline treatment in the model with $S_p = \pi + v_p$, where $\sigma_{v,p}^2 = (\sigma_{v,p}^f)^2/\rho^2$.

about future inflation in high inflationary periods, i.e. $\sigma_{v,p}^2$ increases with inflation. In our model, this reduces the Kalman gain and therefore changes estimated treatment effects through the intensive margin. While we do not have direct measures of credibility/informativeness, we can utilize proxies to evaluate this mechanism. Specifically, we assess this channel by examining whether trust in the Federal Reserve and other government institutions has changed over time. Panel B of Figure 5 shows that, according to Gallup surveys, the level of trust for not only the Federal Reserve but also other government institutions has been generally declining since the early 2000s with a bump-up in trust during the pandemic and reversal to the trend after the pandemic subsided.¹⁹ The level of trust for the Fed chair was similar in 2014 and 2023. Thus, it seems unlikely that changes in credibility can account for our empirical results.

Share of Informed Respondents. Finally, through the lens of our model, the decrease in the estimated treatment effect during high inflation periods could also come from the extensive margin, namely an increase in the share of individuals who are already informed about the information provided in the treatments. This could stem from a fall in ϕ , the cost of accessing pre-processed signals about inflation, or an increase in σ_{π}^2 , which increases $I(S_p, \pi)$, the informativeness of such signals about inflation. Panel D of Figure 5 shows that not only did households search more intensively for information about inflation during the inflation surge (see Korenok, Munro and Cheng 2023), but the media also supplied more inflation-related information. Furthermore, we note that, when inflation rose, households searched more intensively for inflation forecasts, which is consistent with the information in treatments (signal S_p in our model) already being in households' priors. These two mechanisms are both consistent with the decline in treatment effects with no change in inflation uncertainty. In short, better awareness about publicly available inflation-related news in a high-inflation environment appears to be the most promising explanation for the decrease in the power of our information interventions during the inflation spike, although we cannot separate between demand (more volatile inflation leading people to search for information about inflation) and supply (more news about inflation makes it easier for people to learn about inflation).

¹⁹ We find similar results when we use surveys about trust in institutions from Pew Charitable Trust. Interestingly, trust in the European Central Bank, European Commission and European Parliament plunged during the government debt crisis in 2014 but has been recovering since. This dynamic also does not support the notion that changes in trust can explain the variation in estimated treatment effects that we observe.

However, to the extent that one would assign the higher coverage of inflation partly to an increase in σ_π^2 , then our findings have implications for the type of information cost functions that might be governing the agents' information acquisition. In particular, the fact that we do not observe any changes in households' uncertainty after they have acquired some information strongly favors the linear cost of attention in terms of mutual information because any curvature in the cost function would suggest that posterior uncertainty should increase (for convex cost) or decrease (with non-convex costs such as fixed cost of attention models) with households' a priori uncertainty.

While both an increase in σ_π^2 or a decrease in ϕ increase attention of households to public news and decrease the treatment effect, they have different implications for the total amount of attention paid to inflation. Recall that total attention of household i to inflation in our model is given by the mutual information between inflation and their information set S_i :

$$I(\pi, S_i) = \frac{1}{2} \ln(\text{Var}(\pi)/\text{Var}(\pi|S_i)) = \frac{1}{2} \ln(\sigma_\pi^2) - \frac{1}{2} \ln\left(\frac{\omega}{B\rho^2}\right).$$

A key observation here is that this mutual information does depend on σ_π^2 but not on ϕ . Thus, while our evidence suggests that households' attention to public news has increased with higher inflation either due to an increase in σ_π^2 or a decline in ϕ , their total attention is only weakly increasing in the sense that it might have either increased or remained the same depending on which channel is at work. To see why, we note that unlike usual unidimensional rational inattention problems, our model has a non-trivial margin of substitution between observing pre-processed public news and processing information by sampling from the distribution of prices. All else equal, a decline in ϕ operates only through the substitution margin by increasing the relative cost of processing information versus observing pre-processed signals. But it does not affect total attention as households simply reduce their own information processing to keep their posterior uncertainty unchanged. However, an increase in σ_π^2 has two separate effects on attention: first, it increases the value of the public signal as $I(S_p, \pi)$ goes up and accordingly triggers the substitution margin where households substitute away from processing information and towards acquiring public signals. Second, it increases total attention as households would now need to process more information to attain the same level of posterior uncertainty, denoted by $\frac{\omega}{B\rho^2}$.

3.3 Euro Area Households

To complement the findings for U.S. households, we utilize a series of RCTs applied to the ECB's CES. The CES was established in 2020 and originally included France, Germany, Spain, Italy, Belgium, and the Netherlands, while starting in 2022 the survey was also piloted in five additional countries (Austria, Finland, Greece, Ireland, and Portugal). More detailed information about the survey is provided in ECB (2021) and Georgarakos and Kenny (2022). The CES can use occasional ad hoc modules to run RCTs to study how various information interventions affect the beliefs of households in the euro area. We focus on RCTs implemented in 2021Q4, 2022Q1, 2022Q2 and 2022Q4, all of which included at least one information treatment about inflation to a randomized subset of participants. In the CES we measure prior beliefs of households using one-year ahead inflation point forecasts reported before any information treatment. After information treatments, households provide a point forecast for year-ahead inflation, which serves as our measure of posterior beliefs.²⁰ Each RCT also includes a control group that is not provided with any information.

To assess the effects of information treatments on euro area households, we apply the same empirical specifications as for the Nielsen survey, using both the instantaneous change in forecasts within the survey as well as the inflation forecasts three months later. Panel A of Figure 6 plots the resulting estimates of scaled instantaneous treatment effects whereas Panel B of Figure 6 plots treatment effects after three months. In 2021Q4, inflation in the euro area was already around 5%, so initial instantaneous scaled treatment effects are small, around -0.2. As the inflation rate rose further to around 10% in 2022, we see that the treatment effects become even smaller, even insignificantly different from zero in the final available RCT in 2022Q4 (when inflation stood at 8.6%). Hence, we can observe the same decline in instantaneous treatment effects in the CES as was visible in the Nielsen survey of U.S. households, albeit over a shorter time sample. Treatment effects after 3 months are consistently estimated to be close to zero throughout the sample. Again, the results are broadly similar across information treatments.

One clear feature of the above experiments implemented in the CES is that by the time they began, inflation was already relatively high and in the news, so treatment effects were small to start with and it is difficult to identify time variation in these effects within this limited time frame. We

²⁰ Only the most recent RCT (2022Q4) uses a distributional question after treatments to measure posterior beliefs. In this case, we compare these posterior beliefs to respondents' prior beliefs using information from a corresponding distributional question asked before treatments.

consider two independent strategies to address this limitation. First, we include an additional comparable RCT that was run in the Netherlands before the inflation run-up on the Dutch National Bank’s household survey (DHS). Second, we provide cross-sectional evidence from the CES that confirms that households that report paying a lot of attention to inflation respond significantly less to information treatments than those that report paying little attention.

The Dutch RCT, which was run in 2018Q2, used a nearly indistinguishable survey design from the CES in which the treated households were informed about the most recent inflation rate in the Netherlands (see Coibion et al. 2023 for a detailed description). The survey was smaller in size (about 2,000 respondents), but it was large enough to obtain reasonably precise estimates. A follow-up wave was implemented three months later.²¹ We include results in Panels A and B of Figure 6. In each case, we find much larger treatment effects in 2018 than those we obtain later in the CES sample, providing more evidence that as the inflation rate increased in the euro area, information treatment effects became smaller as households became more attentive to inflation.

Another approach that we can use to verify the role played by attention is to exploit the fact that, in a recent ad hoc module of the CES, some households explicitly report being more informed about inflation than others. Specifically, we split respondents in the 2022Q4 wave into two groups: low-attention and high-attention (53% and 47% of the sample, respectively) based on self-reported attention to inflation. We then estimate the instantaneous treatment effect for each group separately and report the results in Table 2 (columns 1 and 2). For the high-attention group, we find no treatment effect, either in terms of the slope or the intercept. For the low-attention group on the other hand, we identify a negative scaled slope effect and a positive intercept. Hence, there is a clear difference in how the two groups respond. Those who are attentive place no weight on the provided information, likely because they already know the prevailing inflation rate, whereas those who are less attentive to inflation update their beliefs when presented with information about recent inflation.

3.4 U.S. Firms

Finding comparable evidence for firms is inherently challenging: there are far fewer large representative surveys of firms in which RCTs are allowed or feasible compared to household surveys. One exception is the Federal Reserve Bank of Atlanta’s Business Inflation Expectations

²¹ Dutch respondents in the CES have inflation expectations comparable to households in other euro area countries (ECB 2021). Dutch inflation is highly correlated with inflation in the euro area ($\rho=0.96$) for the 2015-2023 period.

survey (BIE). The BIE is a monthly survey of firms in the 6th District of the Federal Reserve System. The industry composition of the survey roughly conforms to the industrial mix of the United States, so that it can be viewed as broadly representative. Each month, around 300 firms are surveyed. More details about this survey are provided in Bryan, Meyer and Parker (2015) and Meyer and Sheng (2022). Note that this sample is much smaller than household surveys, making it more difficult to implement RCTs with strong statistical power.

The Atlanta Fed implemented two such RCTs in January of 2019 and February of 2023. In each case, a randomly selected subset of firms was provided with the most recent inflation rate. Prior to the information provision, all firms had been asked about what they thought the inflation rate had been over the previous twelve months, which we use as the prior. After the treatment, all firms were asked to provide a point forecast for aggregate inflation in the U.S. over the next 12 months, which serves as our measure of the posterior. Thus, we can estimate the instantaneous effect of information treatments on firms' expectations in a manner directly analogous to that used for households. We report estimates of the scaled treatment coefficient in Figure 7. In 2019, when inflation was low, the estimated weight on priors for treated firms was 73 percent smaller than for the control group. By 2023, this coefficient had declined to 52 percent smaller than the control group, suggesting that firms' attention to inflation also increased as the inflation rate rose. However, given the small samples, we cannot reject the null of equality across the two survey waves, although we can strongly reject this null when we use the unscaled treatment effects (SA Appendix Figure A.4). At the same time, Meyer and Sheng (2022) document a pattern of increased attention to inflation in a high inflation environment among firms in this district. Specifically, the share of firms indicating that inflation has at least a "moderate" influence of business decision-making rose from below half of the panel in January 2015 (when overall inflation was low) to nearly 2/3 of the panel in May 2022 (when the 12-month growth rate in the CPI was 8.6 percent). Schwartzman and Waddell (2024) find that more U.S. firms in the 5th District of the Federal Reserve System reported that they were paying close attention to inflation as the U.S. inflation rate rose in 2022. Hence, despite the statistical ambiguity in the regression estimates, the combined body of evidence is consistent with the notion that inattention to inflation among U.S. firms has likely declined as inflation has risen.

IV Additional Evidence from Other Settings

RCTs in the U.S. Nielsen survey, euro-area CES, and Atlanta Fed's BIE survey all allow us to compare information treatments before and during the recent global rise in inflation. In this Section, we consider other settings that also speak to this question, albeit each from a different angle. First, we consider the case of Uruguay, which experienced relatively high inflation in the past two decades. Second, we consider firms in New Zealand over a six-year period during which inflation was consistently low. Third, we consider the case of firms in Italy, some of which were repeatedly provided with information about inflation since 2012 while others were not, thereby providing another laboratory to study how information treatments may have changed over time.

4.1 Uruguay: Information treatments in a consistently high-inflation environment

We plot inflation dynamics in Uruguay since 2017 in Figure 8: inflation averaged around 8% over this period and never fell below 5%. This inflation level has been sustained since the mid-2000s and is somewhat above the central bank's inflation target range.²² Interestingly, only a mild increase in inflation occurred between 2021 and 2023 in Uruguay, and it has proven to be transitory. Thus, unlike the U.S. or the euro area, Uruguay can be characterized as having experienced consistently high inflation (by the standards of advanced economies) over the entire time period.

The National Institute of Statistics (INE) of Uruguay, on behalf of the Central Bank of Uruguay, runs a monthly representative survey of firms. The survey is relatively large, with around 550 firms participating per month, and quantitative in nature. It includes questions on inflation and cost expectations of firms, among other topics. The survey is described in more detail in Frache and Lluberias (2019) and Borraz and Mello (2020). We focus on four RCTs which were implemented in 2018M3, 2018M6, 2019M6 and 2023M3. In each survey wave, a randomly selected subset of firms was provided with the inflation rate over the last 12 months or the central bank's inflation target, while other firms were not provided with information. Prior to the information treatments, all firms were asked to provide a point forecast for what they expected inflation to be over the next 12 months. Because no comparable question was asked immediately after the treatments, we use firms' inflation expectations in the next month as the posterior.

²² This target range has fluctuated over time, both in terms of level and spread of the range. The target range was 3%-7% between July 2013 and September 2022, and it has been 3%-6% since September 2022.

We estimate the same empirical specification as before to measure the treatment effects of information about inflation on firms' inflation forecasts and report results in Figure 8. The scaled treatment effects on short-term inflation expectations are consistently close to zero in magnitude and never statistically different from zero or each other. In other words, we find no change in inattention of firms in Uruguay. Throughout the sample, they appear to be well-informed about inflation and monetary policy so that, when provided with information about either inflation or the central bank's target range, they do not change their forecasts. This "zero effect" of inflation information treatments is precisely what one would expect from agents living in a high-inflation environment: they are constantly attentive to and already informed about inflation and monetary policy.

4.2 New Zealand: Information treatments in a consistently low-inflation environment

The case of Uruguay is unique in that it covers multiple RCTs over the course of many years in a high-inflation environment. What happens over the course of many years in a low-inflation environment? We consider this case using repeated RCTs of firms that were implemented in New Zealand from 2014 to 2019, a time period during which inflation never exceeded 2.5% and occurred after more than two decades of low and stable inflation since New Zealand adopted its 2% inflation target in 1990. Unlike previously considered settings, the RCTs in New Zealand were not implemented in the context of a regular ongoing survey. Instead, they were implemented individually at different times. Prior inflation expectations were measured using a distributional question while posteriors were measured using a point forecast for inflation over the next 12 months. The first two RCTs in New Zealand (2014Q4 and 2016Q2) were part of a sequence of surveys described in Coibion, Gorodnichenko and Kumar (2018). In 2014Q4, around 1,600 firms were randomly assigned to either a control group or one of three treatment groups. The latter received either the most recent inflation rate, the central bank's inflation target, or professional forecasts of one-year ahead inflation. Applying our same empirical specification, we find (Figure 9) that the treatments had large effects on inflation expectations, with scaled slope treatment effects ranging from -0.55 (central bank target) to -0.95 (professional forecasts).

In 2016Q2, another information treatment was applied to a new representative group of firms in New Zealand. In this case, around 2,000 firms were either randomly assigned to the control group or were provided with the central bank's inflation target. Using the same empirical specification, we estimate a slightly smaller scaled treatment effect of around -0.35,

perhaps reflecting the fact that inflation was close to the deflationary zone and may therefore have been receiving more news coverage than in 2014. Another RCT was applied to a new representative group of firms in 2018Q1, as described in more detail in Coibion et al. (2021b). In this case, 251 firms received only the past inflation treatment or were in the control group. As shown in Figure 9, the estimated scaled treatment effect in this case is -0.63, effectively indistinguishable from that estimated with the same treatment in 2014Q1, when inflation had been running at a similar level as in 2018. Finally, yet another RCT was implemented on a new group of around 1,000 New Zealand firms in 2019Q3. In this case, the information treatment consisted of a combination of the previous period's inflation rate and central bank inflation target. Hence, the treatment is not directly comparable to the previous ones. Nonetheless, the estimated scaled treatment effect is still similar as in prior waves, at -0.9. In short, over a 6-year time interval during which inflation was relatively low and stable, we find across four RCTs of firms in New Zealand what looks like systematically high levels of inattention. This evidence is consistent with New Zealand's long history of inflation targeting and low inflation.

4.3 Italy: The effect of repeatedly treating firms in low- and high-inflation environments

Finally, we consider another unique setting, that of Italy, where an RCT has been repeatedly applied for over a decade. In the Italian SIGE, some firms have been *repeatedly* provided with information about the most recent inflation rate, whereas others have not, over the course of years, thereby providing a unique setting to study how the level of inflation shapes inattention.

The SIGE is a quarterly survey of firms in which approximately 1,150 firms per quarter participate. As described in Grasso and Ropele (2018) and Coibion, Gorodnichenko and Ropele (2020), at infrequent intervals firms are randomly assigned to one of two groups. One group is asked what they expect inflation to be over the next 12 months. The other group is also asked about their inflation expectations, but after being told what the most recent inflation rate was both in Italy and in the euro area. Firms remain in their group until the next reshuffling, meaning that in between re-assignments, some firms are repeatedly provided with information while others are not. Before 2012Q3, all firms were provided with the same information about recent inflation. In 2012Q3, approximately one-third of firms were randomly assigned to the group that is not provided with any information. In 2012Q4, the firms were randomly reshuffled across the two groups and remained in them until 2017Q2, when another reshuffling took place. A final reshuffling took place in 2019Q4.

The survey only asks for inflation expectations after information is provided to firms (for those in the treatment group). As a result, we use firms' inflation expectations from the previous wave as the measure of their prior belief. Applying the same cross-sectional regression as before yields a time series of estimated $\hat{\gamma}_t/\hat{\beta}_t$. We plot this time series in Figure 10 (time series for unscaled slopes are in SA Appendix Figure A.7). While there is significant variation over time in the estimates, we note a clear increase in $\hat{\gamma}_t/\hat{\beta}_t$ from -0.45 for 2012Q3-2021Q3 when inflation is below 1% on average to -0.04 for 2021Q4-2023Q1 when inflation exceeds 5%. Hence, these results again suggest that firms became more attentive to inflation as the inflation rate increased in recent years.

V Pooled Evidence

Having considered these country-specific results in isolation, we now bring them together to assess the extent to which the level of inflation is related to how (in)attentive households and firms are to inflation. We do so by combining the results from all the RCTs of U.S. households in Nielsen, euro-area households in the CES, U.S. firms in the BIE, Uruguayan firms, and New Zealand firms. For the Italian SIGE, we pool estimates from 2012-2021 into one low-inflation estimate and estimates from 2022 into one high-inflation estimate. We then plot in Figure 11 the level of CPI inflation existing at the time of each RCT against the scaled slope treatment effect ($\hat{\gamma}/\hat{\beta}$) of each RCT. There is a striking positive correlation ($\rho = 0.6$) between the two (SA Appendix Figure A.8 plots the equivalent results for unscaled treatment effects and finds an even stronger positive correlation), consistent with inattention to inflation being more pervasive in low-inflation than high-inflation environments.

Despite the different treatment types, the different questions used to measure priors and posteriors, and the fact that we consider both households and firms, all of which should tend to attenuate any underlying correlation, we still uncover a clear positive link between inflation and inattention. When we pool estimates across countries, times, and treatments and regress $\hat{\gamma}/\hat{\beta}$ on the rate of inflation at the RCT time, we find that a one percentage point increase in the rate of inflation is associated with a 0.064 (s.e. 0.013) increase in $\hat{\gamma}/\hat{\beta}$. This fitted relationship suggests that households and firms pay very close attention when annual inflation reaches 11.5 percent (i.e., $\hat{\gamma}/\hat{\beta} \approx 0$) while the degree of inattention is high ($\hat{\gamma}/\hat{\beta} \approx -0.6$) when inflation is close to 2 percent.

VI Conclusion

When inflation is higher, households and firms pay more attention to publicly available news about inflation. Our comprehensive set of results documenting this pattern through repeated RCTs in different countries complement other recent evidence such as Cavallo, Cruces and Perez-Truglia (2017), Bracha and Tang (forthcoming), Korenok, Munro and Chen (2023) and Pfäuti (2023). Jointly, this line of research presents clear evidence, using a variety of empirical strategies, that attention to inflation is endogenous and varies with the level of inflation.

These results have broad implications. For example, when agents are more inattentive, the Phillips curve is flatter (Afrouzi and Yang 2023), forward guidance is less powerful (Kiley 2021) and the ZLB constrains monetary policy more (Pfäuti 2023). Each of these mechanisms is central to monetary policy decisions. Incorporating the systematic endogeneity of inattention should therefore be an important objective for future work in optimal policy design.

Endogeneity of inattention also matters for policy communication and management of inflation expectations. When agents are inattentive, the main challenge for policymakers who seek to affect expectations is *how* to reach households and firms. Conditional on reaching them, communication is very powerful, as found in Coibion, Gorodnichenko and Weber (2022), and can enhance central bank credibility (Ehrmann, Georgarakos and Kenny 2022). In contrast, when agents are attentive, reaching them is less of a challenge. Instead, the difficulty becomes that they are less responsive to policy communications since they are already better informed. *What* information is relayed to them therefore becomes the main challenge (Candia, Coibion and Gorodnichenko 2020; D’Acunto et al, 2020). Policymakers interested in steering expectations to better stabilize economic outcomes should consider how the economic environment shapes the way to successfully communicate with the public.

Methodologically, our results also provide support for the use of RCTs along with a call for caution. We find that similar RCTs implemented in different countries at different times but experiencing similar economic environments yield results that are broadly similar. This indicates that RCTs can be viewed as having some external validity. But the “similar economic environment” is an important caveat. As emphasized in the Lucas (1976) critique, a changing environment will lead to changing behavior on the part of economic agents. Our results provide yet more evidence for Lucas’ insight, in this case by showing that the level of inflation affects how inattentive households and firms are to macroeconomic conditions.

References

- Afrouzi, H., and C. Yang, 2021. "Dynamic Rational Inattention and the Phillips Curve," mimeo.
- An, Z., S. Abo-Zaid and X.S. Sheng, 2023. "Inattention and the impact of monetary policy," *Journal of Applied Econometrics* 38(4): 623-643.
- Armantier, O., S. Nelson, G. Topa, W. 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(3): 503-523.
- Armona, L., A. Fuster, and B. Zafar, 2019. "Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment," *Review of Economic Studies* 86(4): 1371-1410.
- Beutel, J., and M. Weber, 2022. "Beliefs and Portfolios: Causal Evidence," Chicago Booth Research Paper No.22-08.
- BIE - Business Inflation Expectations. Federal Reserve Bank of Atlanta. <https://www.atlantafed.org/research/inflationproject/bie>
- Binder, C., 2017. "Measuring uncertainty based on rounding: New method and application to inflation expectations," *Journal of Monetary Economics* 90(C): 1-12.
- BLS - U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCSL], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/CPIAUCSL>. Accessed 2024Q1.
- Borraz, F., A. Orlik, and L. Zacheo, 2023. "On Firms' Inflation Expectations, (In)Attention and Inflation Targeting: Survey Evidence for Uruguay," Manuscript.
- Borraz, F., and M. Mello, 2020. "Communication, Information and Inflation Expectations," Working Paper of the Banco Central de Uruguay No. 005-2020.
- Bracha, A., and J. Tang, 2019. "Inflation Levels and (In)Attention," Forthcoming in *Review of Economics Studies*.
- Bryan, M., B. Meyer, and N. Parker, 2015. "The Inflation Expectations of Firms: What Do They Look Like, Are They Accurate, and Do They Matter?" Federal Reserve Bank of Atlanta Working Paper 2014-27a.
- Candia, B., O. Coibion, and Y. Gorodnichenko, 2020. "Communication and the Beliefs of Agents," in *Navigating the Decade Ahead: Implications for Monetary Policy*, Economic Policy Symposium (Jackson Hole, WY) Proceedings, FRB Kansas City.
- Cavallo, A., G. Cruces, and R. Perez-Truglia. 2017. "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments." *American Economic Journal: Macroeconomics* 9(3): 1-35.

- Chopra, F., C. Roth, and J. Wohlfart, 2023. “Home Price Expectations and Spending: Evidence from a Field Experiment,” Manuscript.
- Coibion, O. and Y. Gorodnichenko, 2015. “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” *American Economic Review* 105(8), 2644-2678.
- Coibion, O., D. Georgarakos, Y. Gorodnichenko and M. Weber, 2023a. “Forward Guidance and Household Expectations,” *Journal of European Economic Association* 21(5): 2131-2171.
- Coibion, O., D. Georgarakos, Y. Gorodnichenko, and M. van Rooij, 2023. “How Does Consumption Respond to News about Inflation? Field Evidence from a Randomized Control Trial,” *AEJ: Macroeconomics* 15(3): 109-152.
- Coibion, O., D. Georgarakos, Y. Gorodnichenko, G. Kenny and M. Weber, 2022. “The Effect of Macroeconomic Uncertainty on Household Spending,” forthcoming in *American Economic Review*.
- Coibion, O., Y. Gorodnichenko, and M. Weber, 2022. “Monetary Policy Communications and their Effects on Household Inflation Expectations,” *Journal of Political Economy* 130(6): 1537-1584.
- Coibion, O., Y. Gorodnichenko, and S. Kumar, 2018. “How Do Firms Form Their Expectations? New Survey Evidence,” *American Economic Review* 108(9): 2671-2713.
- Coibion, O., Y. Gorodnichenko, and T. Ropele, 2020. “Inflation Expectations and Firm Decisions: New Causal Evidence,” *Quarterly Journal of Economics* 135(1): 165-219.
- Coibion, O., Y. Gorodnichenko, S. Kumar, and J. Ryngaert, 2021b. “Do You Know That I Know That You Know...? Higher Order Beliefs in Survey Data,” *Quarterly Journal of Economics* 136(3): 1387-1446.
- D’Acunto, F., D. Hoang, M. Paloviita, and M. Weber, 2020. “Effective Policy Communication: Targets versus Instrument,” Working Paper.
- D’Acunto, F., U. Malmendier, J. Ospina, and M. Weber, 2021. “Exposure to Grocery Prices and Inflation Expectations,” *Journal of Political Economy* 129(5): 1615-1639.
- De Quidt, J. J. Haushofer, and C. Roth, 2018. “Measuring and Bounding Experimenter Demand Effects,” *American Economic Review* 108(11): 3266-3302.
- ECB Data Portal. European Central Bank. <https://data.ecb.europa.eu/>. Data accessed 2024Q1.
- ECB-CES - Consumer Expectations Survey. European Central Bank. https://www.ecb.europa.eu/stats/ecb_surveys/consumer_exp_survey/html/index.en.html. Data accessed 2024Q1.

- EEE - Encuesta de Expectativas Empresariales. Instituto Nacional de Estadística de Uruguay. <https://www.gub.uy/instituto-nacional-estadistica/encuesta-expectativas-empresariales>. Data accessed 2024Q1.
- Ehrmann, M., D. Georgarakos, and G. Kenny, 2022. “Credibility Gains from Communicating with the Public: Evidence from the ECB’s new Monetary Policy Strategy,” ECB WP No. 2785.
- European Central Bank, 2021. “ECB Consumer Expectations Survey: an overview and first evaluation”, Occasional Paper Series, No 287, December.
- Fajardo, J., and M. Dantas, 2018. “Understanding the impact of severe hyperinflation experience on current household investment behavior,” *Journal of Behavioral and Experimental Finance* 17(March): 60-67.
- Flynn, J., and K. Sastry, 2023. “Attention Cycles,” manuscript.
- Frache, S., and R. Lluberás, 2019. “New Information and Inflation Expectations among Firms,” BIS working paper No. 781.
- Gallup News Service. 2023. Gallup Poll Social Series: Economy and Personal Finance.
- Georgarakos, D., and G. Kenny, 2022. “Household spending and fiscal support during the COVID-19 pandemic: Insights from a new consumer survey,” *Journal of Monetary Economics* 129(S): S1-S14.
- Goldstein, N., 2022. “Tracking Inattention,” forthcoming in *Journal of European Economic Association*.
- Goldstein, N., and Y. Gorodnichenko, 2022. “Expectations Formation and Forward Information,” NBER Working Paper 29711.
- Google Trends. “inflation: (United States)” “inflation forecasts”. Data accessed 2024Q1.
- Grasso, A., and T. Ropele, 2018. “Firms’ inflation expectations and investment plans,” Bank of Italy Working Paper Series 1203.
- Haaland, I., C. Roth, and J. Wohlfart, Johannes, 2023. “Designing Information Provision Experiments” *Journal of Economic Literature* 61(1): 3-40.
- Huber, P., 1964. “Robust estimation of a location parameter,” *Annals of Mathematical Statistics* 35: 73–101.
- IPC - Índice de Precios del Consumo. Instituto Nacional de Estadística de Uruguay. <https://www.gub.uy/instituto-nacional-estadistica/indice-precios-consumo>. Data accessed 2024Q1.
- Kiley, M., 2021. “Policy Paradoxes in the New Keynesian Model,” FEDS WP No. 2014-29.

- Kim, G., and C. Binder. 2023. “Learning-through-Survey in Inflation Expectations,” *American Economic Journal: Macroeconomics* 15(2): 254-78.
- Korenok, O., D. Munro, and J. Chen, 2023. “Inflation and Attention Thresholds,” forthcoming in *Review of Economics and Statistics*.
- Kumar, S., Y. Gorodnichenko, and O. Coibion, 2023. “The Effect of Macroeconomic Uncertainty on Firm Decisions,” *Econometrica* 91(4): 1297-1332.
- Link, S., A. Peichl, C. Roth, and J. Wohlfart, 2023a. “Information Frictions among Firms and Households,” *Journal of Monetary Economics* 135: 99-115.
- Link, S., A. Peichl, C. Roth, and J. Wohlfart, 2023b. “Attention to the Macroeconomy,” mimeo.
- Lucas, R.E., 1976. “Econometric Policy Evaluation: A Critique,” *Carnegie-Rochester Conference Series on Public Policy* 1: 19-46.
- Maćkowiak, B., and M. Wiederholt, 2024. “Rational Inattention during an RCT,” manuscript.
- Maćkowiak, B., F. Matějka, and M. Wiederholt. 2023. “Rational Inattention: A Review,” *Journal of Economic Literature* 61(1): 226-73.
- Malmendier, U., and S. Nagel, 2011. “Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?” *Quarterly Journal of Economics* 126(1): 373-416.
- Meyer, B., and X.S. Sheng, 2022. “Unit Cost Expectations and Uncertainty: Firms’ Perspectives on Inflation,” Manuscript.
- MSC - University of Michigan Survey of Consumers. Survey Research Center University of Michigan. <https://data.sca.isr.umich.edu/>. Data accessed 2024Q1.
- Nielsen - Nielsen Homescan Panel. Chicago Booth Kilts center. <https://www.chicagobooth.edu/research/kilts>. Various waves: 2018Q2-2023Q4.
- OECD - Organization for Economic Co-operation and Development, Consumer Price Indices (CPIs, HICPs), COICOP 1999: Consumer Price Index: Total for New Zealand [NZLCPIALLQINMEI], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/NZLCPIALLQINMEI> . Data Accessed 2024Q1.
- Pfäuti, O., 2023. “Inflation – Who Cares? Monetary Policy in Times of Low Attention,” forthcoming in *Journal of Money, Credit and Banking*.
- Roth, C., and J. Wohlfart, 2020. “How Do Expectations about the Macroeconomy Affect Personal Expectations and Behavior?” *Review of Economics and Statistics* 102: 731-748.
- SCE - Survey of Consumer Expectations. Federal Reserve Bank of New York. <https://www.newyorkfed.org/microeconomics/sce.html#/>. Data accessed 2024Q1.

- Schwartzman, F., and S. Waddell, 2024. “Inflation *Expectations* and Price-Setting Among Fifth-District Firms,” Federal Reserve Bank of Richmond *Economic Brief* No. 24-03.
- SIGE - The Bank of Italy’s Survey of Inflation and Growth Expectations. Banca d’Italia Indagine sulle aspettative di inflazione e crescita. 2022. Various waves: 2012Q3-2023Q1. <https://www.bancaditalia.it/statistiche/tematiche/indagini-famiglie-impres/aspettative-inflazione/index.html>
- Song, Wenting, and Samuel Stern, 2023. “Firm Inattention and the Efficacy of Monetary Policy: A Text-Based Approach,” manuscript.
- SPF - Survey of Professional Forecasters. Federal Reserve Bank of Philadelphia. <https://www.philadelphiafed.org/surveys-and-data/data-files>. Data accessed 2024Q1.
- Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberias, Serafin Frache, Brent Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion, Geoff Kenny, and Jorge Ponce, 2024. “Supplement to “Tell Me Something I Don’t Already Know: Learning in Low and High-Inflation Settings,”” *Econometrica Supplemental Material*.

Table 1: Overview of RCTs

Country	Agents	RCT dates	Priors	Posteriors	Information treatments
United States	Households (~20K per wave)	2018Q2, 2019Q1, 2021Q2-Q4, 2022Q3-Q4, 2023Q2-Q4	One-year ahead inflation expectations from distribution	One-year ahead inflation expectations from point forecast	<ul style="list-style-type: none"> • Inflation over the last year • FOMC inflation target • FOMC inflation forecast
Euro area	Households (~10K per wave)	2021Q4, 2022Q2-Q2, 2022Q4	One-year ahead inflation expectations from distribution	One-year ahead inflation expectations from point forecast	<ul style="list-style-type: none"> • Inflation over the last year • ECB inflation target and past inflation • Professional inflation forecast
Netherlands	Households (~2,000)	2018Q2	One-year ahead inflation expectations from distribution	One-year ahead inflation expectations from point forecast	<ul style="list-style-type: none"> • Inflation over the last year
United States	Firms (~300 per wave)	2019Q1, 2023Q1	Perceived inflation over last year	One-year ahead inflation expectations from point forecast	<ul style="list-style-type: none"> • Inflation over the last year
Uruguay	Firms (~500 per wave)	2018Q1-Q2, 2019Q2 2023Q1	One-year ahead inflation expectations from point forecast	One-year ahead inflation expectations from next wave	<ul style="list-style-type: none"> • Inflation over the last year • Central Bank of Uruguay inflation target range
New Zealand	Firms (~2,000 per wave)	2014Q4, 2016Q2, 2018Q1, 2019Q3	One-year ahead inflation expectations from distribution	One-year ahead inflation expectations from point forecast	<ul style="list-style-type: none"> • Inflation over the last year • Reserve Bank of NZ inflation target • Professional forecast of inflation • Combination
Italy	Firms (~1000 per wave)	2012Q3- 22Q4	Inflation expectations in previous quarter from point forecast	One-year ahead inflation expectations from point forecast	<ul style="list-style-type: none"> • Inflation over the last year in Italy and euro area

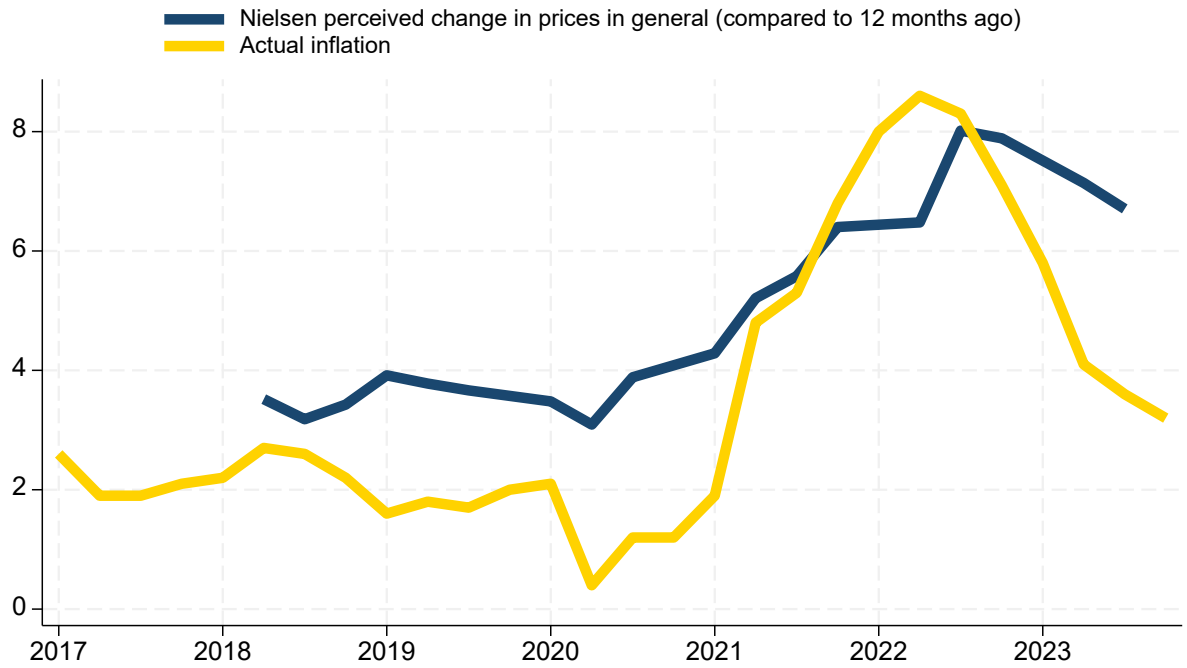
Notes: The table summarizes surveys, measurement of expectations, and information treatments used in our analysis.

Table 2: Treatment Effects for Attentive and Inattentive Households

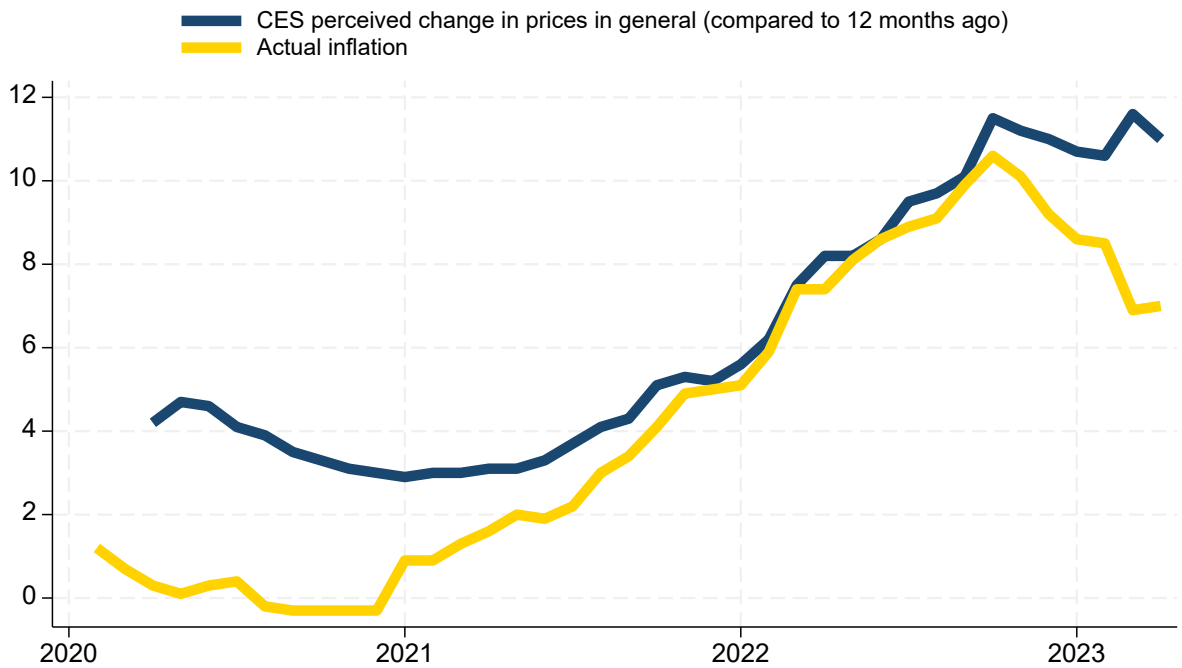
	Treatment effects		Implied moments (prior)	
	Slope (scaled)	Intercept	mean	uncertainty
	(1)	(2)	(3)	(4)
High attention to inflation	0.01 (0.08)	-0.05 (0.41)	7.15 [4.54]	2.85 [2.99]
Low attention to inflation	-0.19*** (0.06)	1.22*** (0.05)	8.32 [4.61]	2.79 [2.99]
p-value equality	0.061	0.022	<0.001	0.486

Notes: Columns (1) and (2) report estimates for γ/β (scaled slope) and δ (intercept) in specification (2) for ECB's CES based on whether respondents pay high or low attention to inflation. Columns (3) and (4) report moments of inflation expectations (prior to RCT) for the two groups. The low-attention group includes respondents who report that they pay "almost no attention", "a little attention" or "some attention" to inflation. The high-attention group includes respondents who report that they pay "much attention" or "a great deal of attention" to inflation. The estimates in columns (1) and (2) are based on the Huber (1964) robust regressions. Robust standard errors are reported in parentheses. ***, **, * indicate statistical significance at 1, 5, and 10 percent levels. Standard deviations for each moment for the two groups are reported in square parentheses in columns (3) and (4).

Figure 1: Actual Inflation and Perceived Inflation by Households
Panel A: U.S. Households



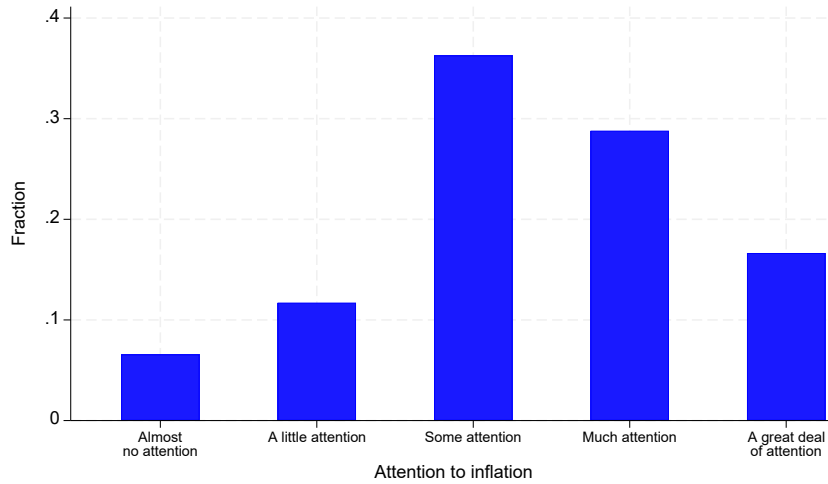
Panel B: Euro Area Households



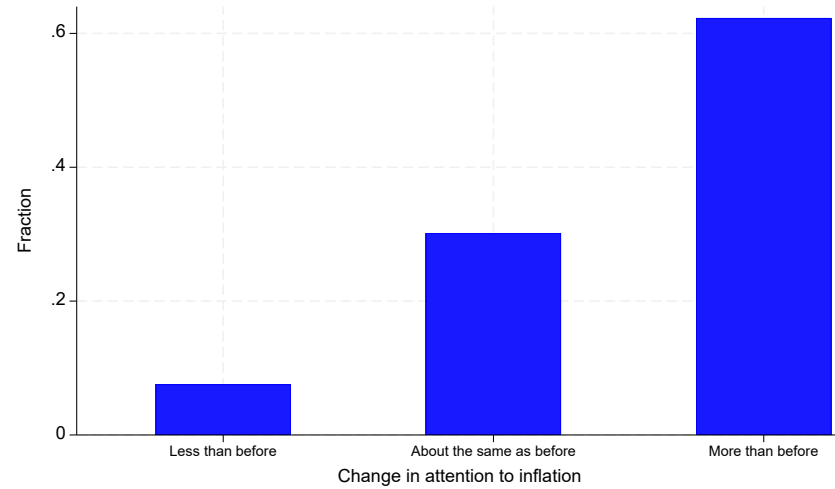
Notes: The figure shows time series of actual inflation and average perceived inflation in the US (Panel A) and the euro area (Panel B).

Figure 2: Attention to Inflation by Euro Area Households

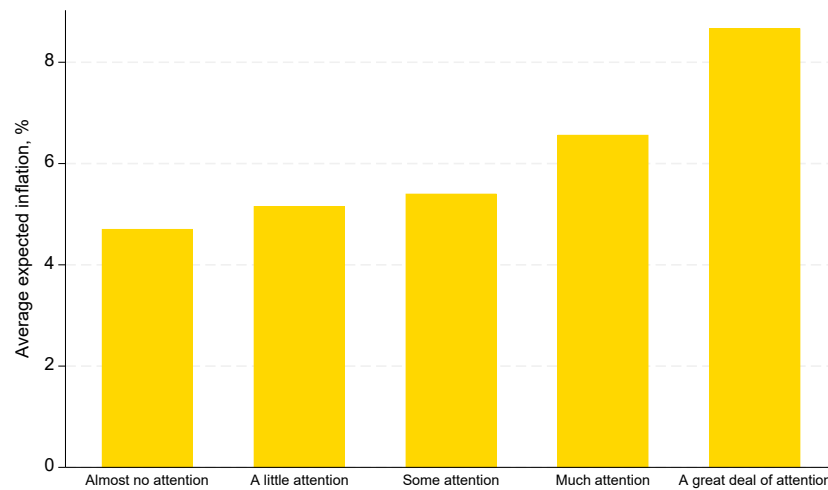
Panel A: Level of Attention to Inflation



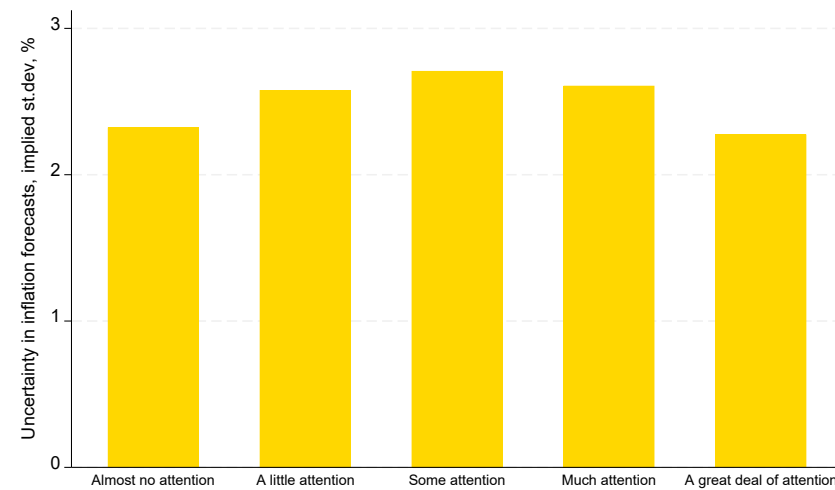
Panel B: Change in Attention to Inflation



Panel C: Inattention and Inflation Forecasts

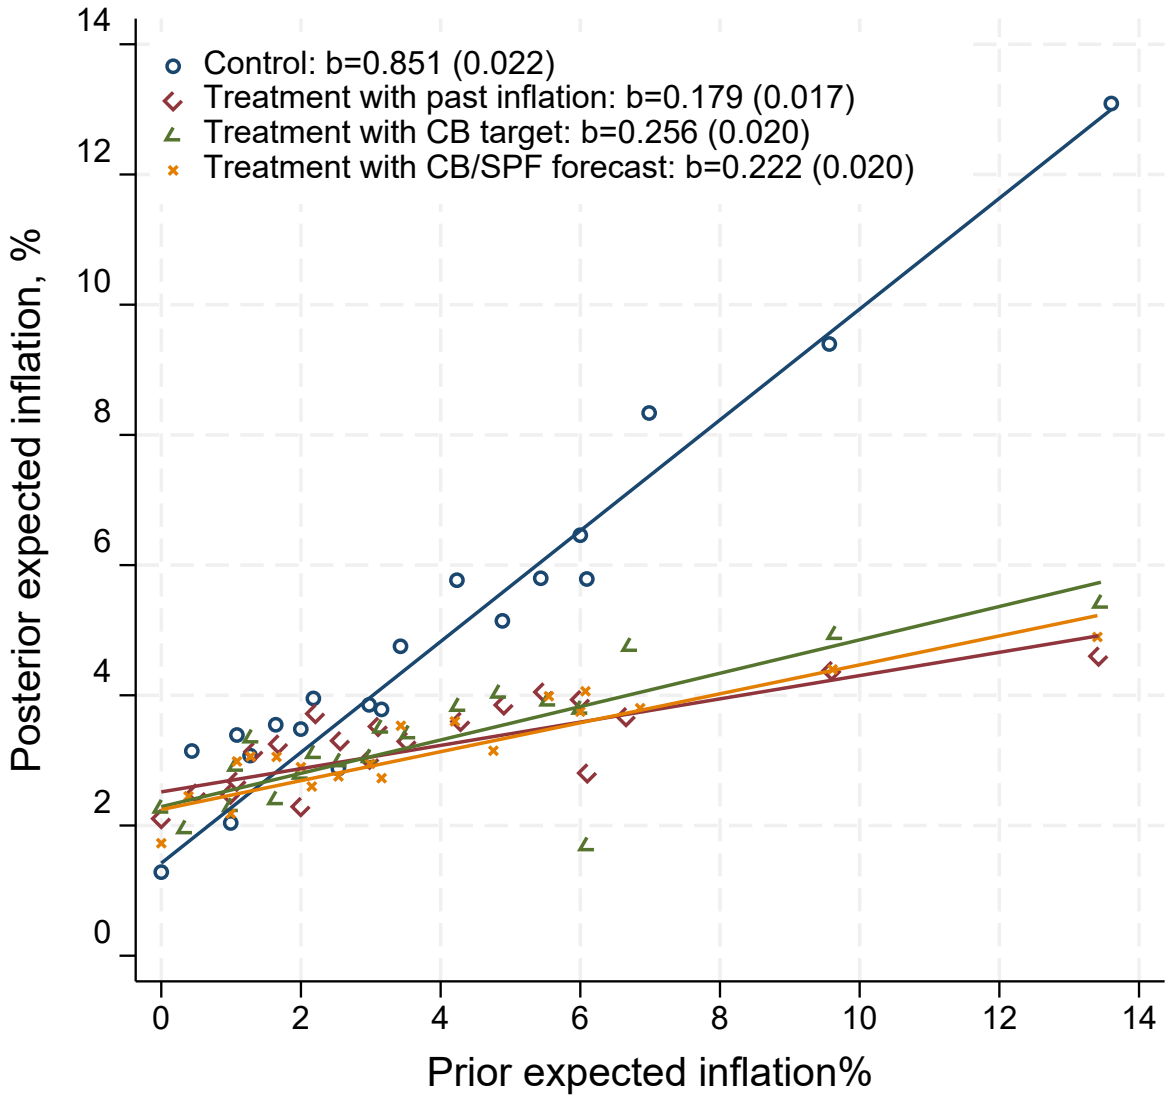


Panel D: Inattention and Uncertainty about Future Inflation



Notes: The figures report the distribution of respondents by the level (or change) of attention to inflation in the 2023M1 wave of the CES as well as their inflation forecasts and uncertainty in their inflation forecasts. Uncertainty in inflation forecasts is measured with the standard deviation of the reported subjective distribution. Subjective distributions are elicited via questions asking respondents to assign probabilities to various possible ranges (bins) of future inflation.

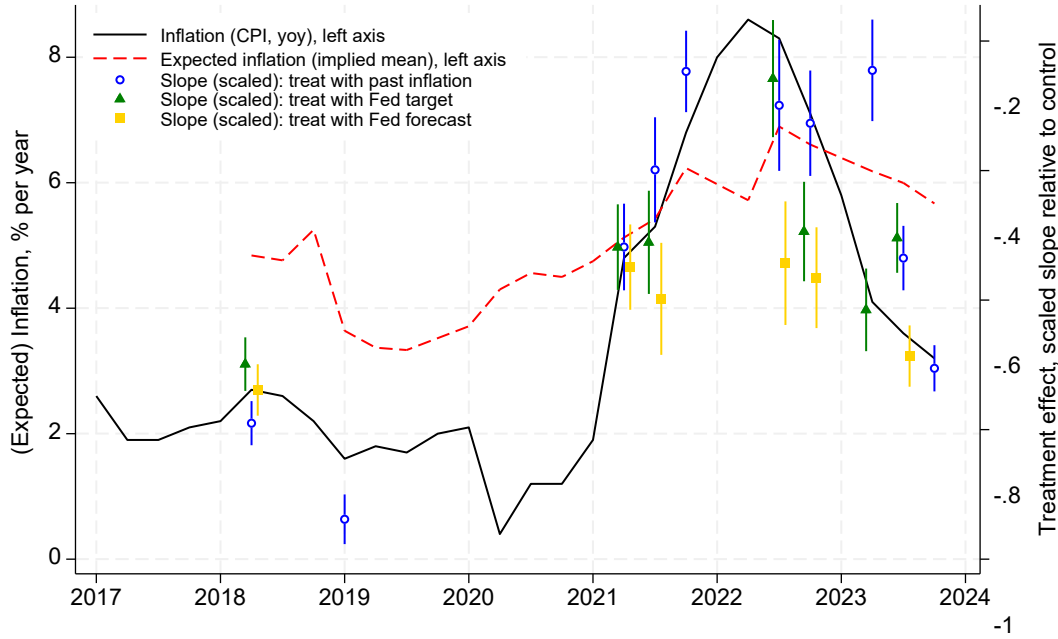
Figure 3: Priors and Posteriors of U.S. Households, 2018Q2



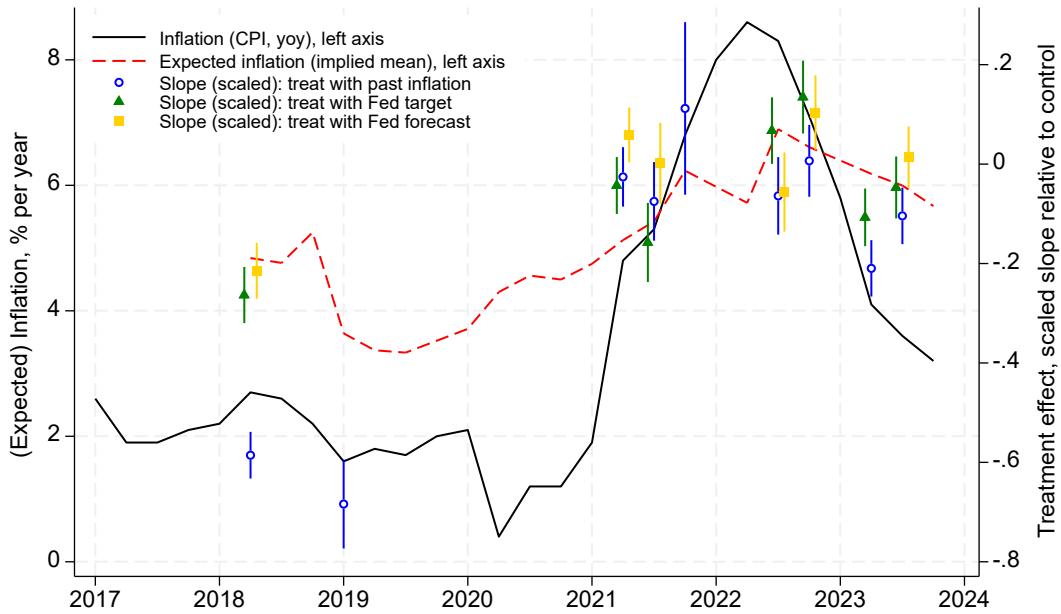
Notes: The figure plots binscatters of priors (x-axis) versus the posteriors (y-axis) of households in the control and treated groups in the Nielsen survey in 2018Q2.

Figure 4: The Changing Effects of Information Treatments on U.S. Households

Panel A: Instantaneous Treatment Effects

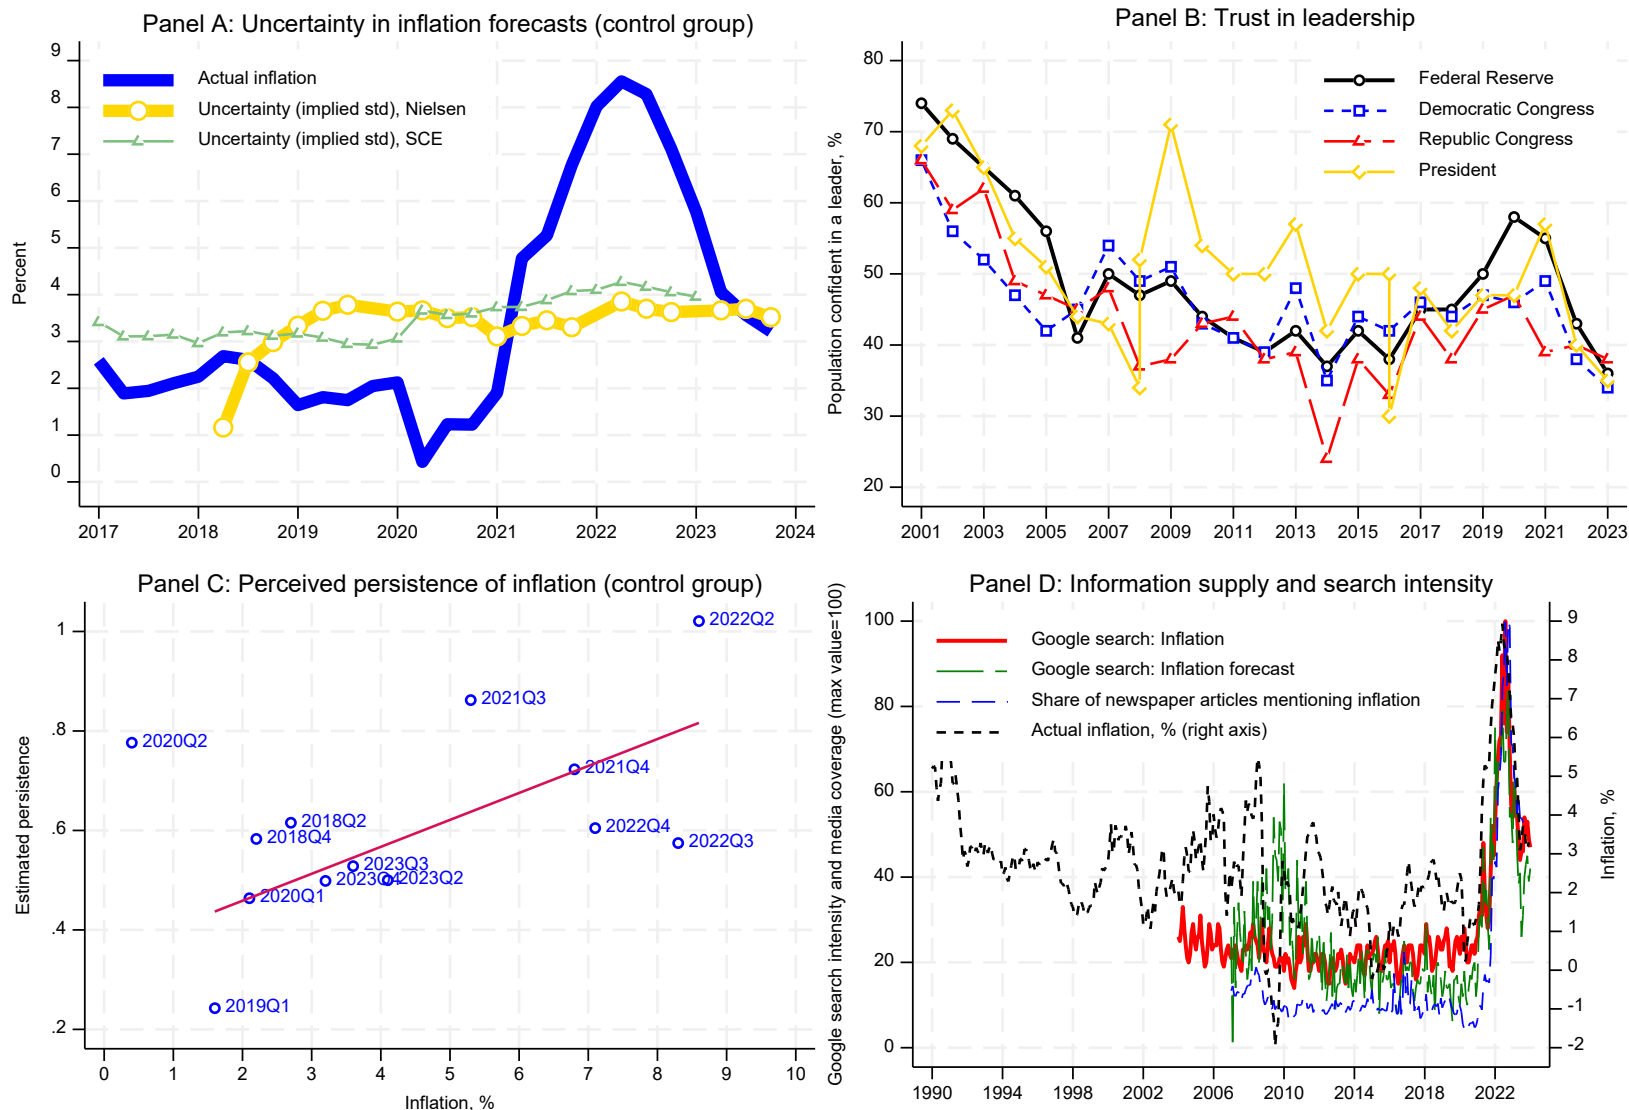


Panel B: Treatment Effects after 3 Months



Notes: Each panel shows the time series of actual inflation and average expected inflation as well as the scaled slopes (γ/β in specification (1) for Panel A and γ/β in specification (1) with posterior measured 3 months later for Panel B) for various treatments across RCTs. The whiskers show the 90% confidence intervals based on heteroskedasticity robust standard errors.

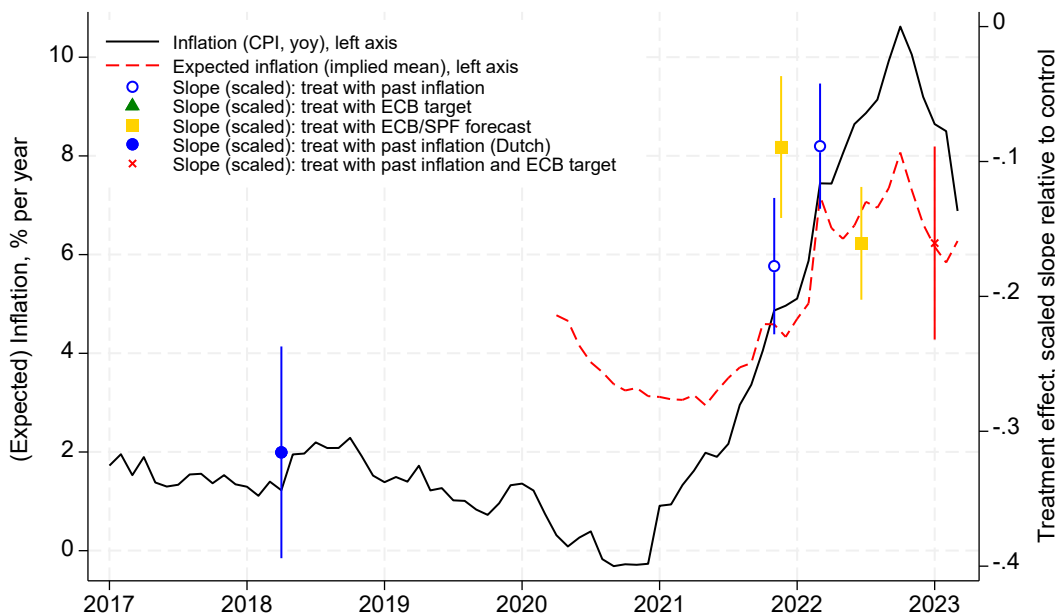
Figure 5. Examining the Channels



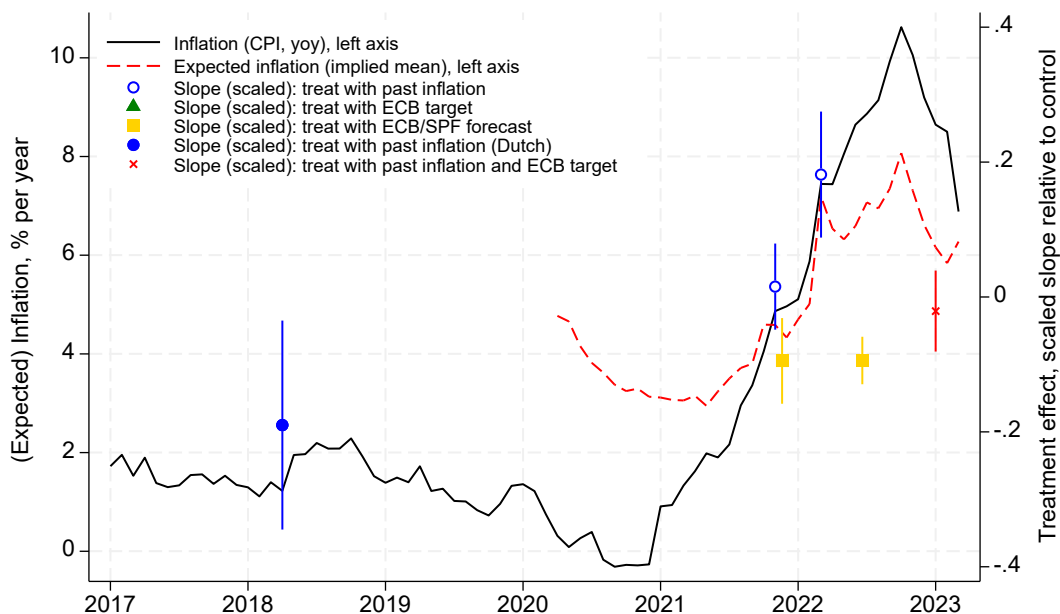
Notes: Panel A plots the time series of uncertainty (standard deviation implied by subjective probability distributions) in households’ inflation expectations in the Survey of Consumer Expectations (SCE; run by the Federal Reserve Bank of New York) and in the Nielsen Homescan Panel. Panel B plots the time series of the share of U.S. population having trust in the leader of a government institution; the data are from Gallup surveys. Panel C plots the estimated persistence of inflation (the estimated slope in the regression of one-year-ahead inflation forecast on perceived inflation over the previous 12 months) in the Nielsen Homescan Panel vs. the actual rate of inflation. Panel D plots the time series of search intensity (Google Trends) for “inflation” and “inflation forecasts”. Each search intensity is normalized so that the maximum value in the reported sample is equal to 100.

Figure 6: The Changing Effects of Information Treatments on Euro Area Households

Panel A: Instantaneous Treatment Effects

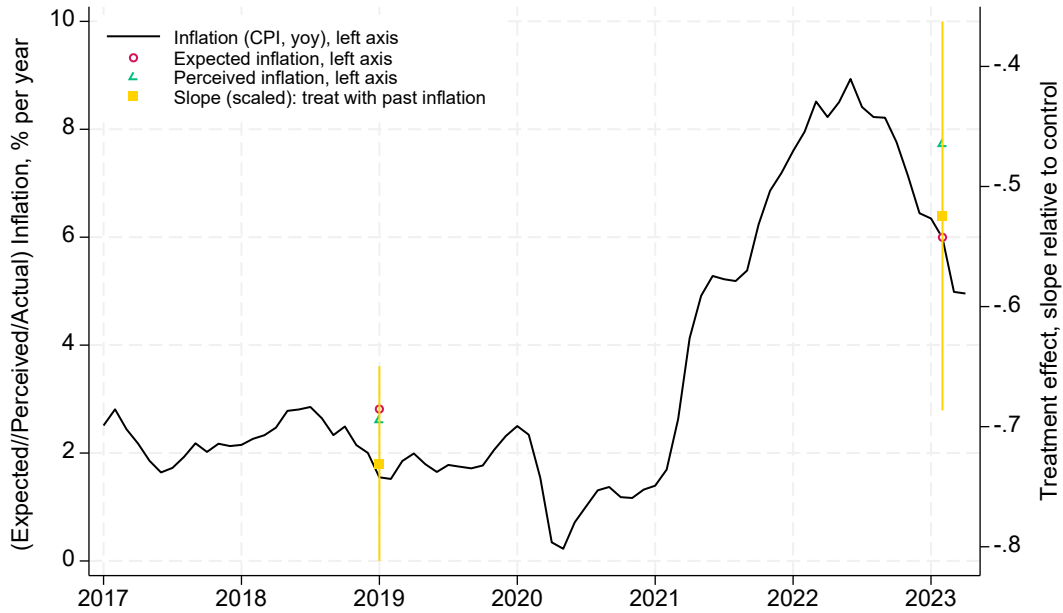


Panel B: Treatment Effects after 3 Months



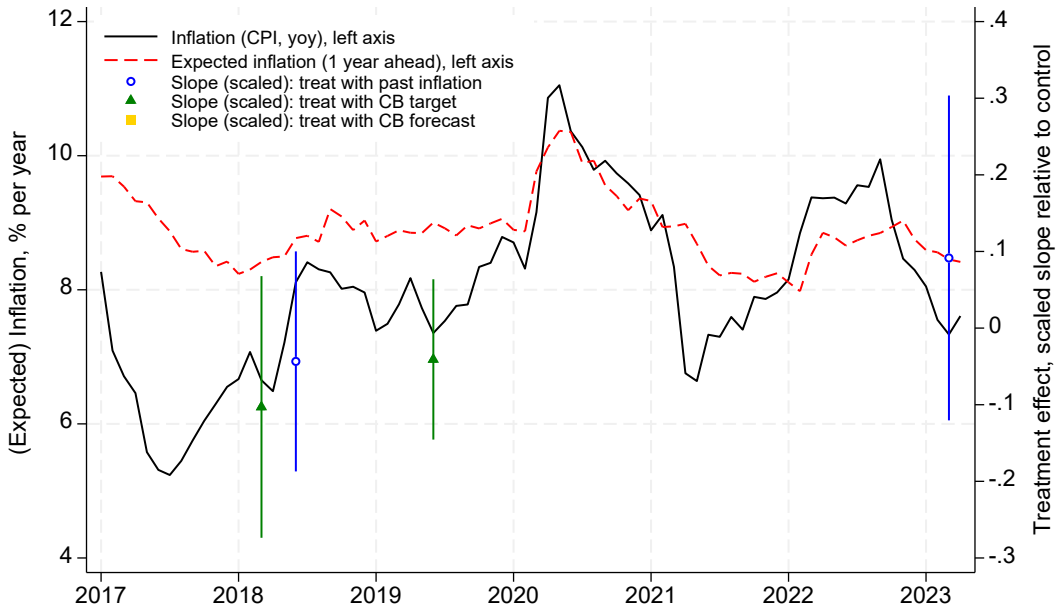
Notes: Each panel shows the time series of actual inflation and average expected inflation as well as the scaled slopes (γ/β in specification (1) for Panel A and γ/β in specification (1) with posteriors measured three months later for Panel B) for various treatments across RCTs. The whiskers show the 90% confidence intervals based on heteroskedasticity robust standard errors. The last observation in Panel B is one month after the corresponding treatment.

Figure 7: The Changing Effects of Information Treatments on U.S. Firms



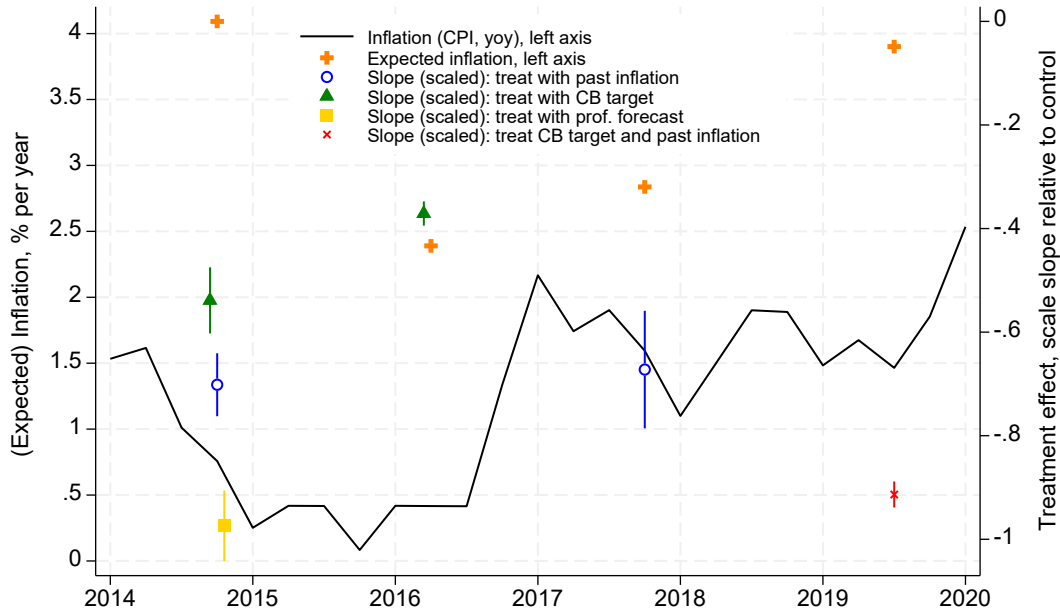
Notes: The figure shows the time series of actual inflation as well as the scaled slopes (γ/β in specification (1)) for various treatments across RCTs. The whiskers show the 90% confidence intervals based on heteroskedasticity robust standard errors. The figure also reports average expectation and perceived inflation at the time when RCTs were conducted.

Figure 8: Time Variation in Treatment Effects on Firms in Uruguay



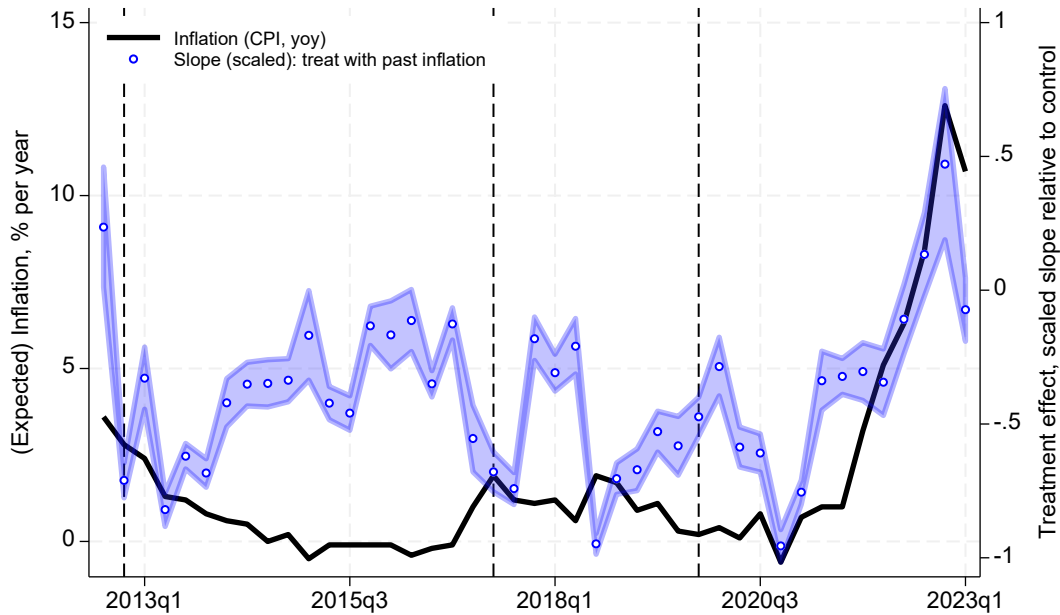
Notes: the figure shows the time series of actual inflation and average expected inflation as well as the scaled slopes (γ/β in specification (1)) for various treatments across RCTs. The whiskers show the 90% confidence intervals based on heteroskedasticity robust standard errors.

Figure 9: Time Variation in Treatment Effects on Firms in New Zealand



Notes: The figure shows the time series of actual inflation as well as the scaled slopes (γ/β in specification (1)) for various treatments across RCTs. The whiskers show the 90% confidence intervals based on heteroskedasticity robust standard errors.

Figure 10: Time Variation in Treatment Effects on Firms in Italy



Notes: The figure shows the time series of actual inflation as well as the scaled slopes (γ/β in specification (1)) for various treatments across RCTs. The shaded area shows the 90% confidence intervals based on heteroskedasticity robust standard errors. The dashed vertical lines show times when firms were randomly reshuffled into treatment and control groups.

