# The Effects of Interest Rate Increases on Consumers' Inflation Expectations: The Roles of Informedness and Compliance

Edward S. Knotek II, James Mitchell, Mathieu Pedemonte, and Taylor Shiroff*

This version: June 2025

Abstract

We study how monetary policy communications associated with increasing the federal funds rate causally affect consumers' inflation expectations in real time. In a large-scale, multi-wave randomized controlled trial (RCT), we find weak evidence that communicating these policy changes lowers consumers' medium-term inflation expectations on average. However, information differs systematically across demographic groups, in terms of ex ante informedness about monetary policy and ex post compliance with the information treatment. Monetary policy communications have a much stronger effect on the subset of consumers who had not previously heard news about monetary policy and who take sufficient time to read the treatment. Our findings show that, in an inflationary environment, these consumers expect that raising interest rates will lower inflation. More generally, our results emphasize the importance of measuring both respondents' information sets and their compliance with treatment when using RCTs in empirical macroeconomics to better understand the real-world implications of monetary policy communications.

*Keywords*: Expectations formation; Policy communication; Monetary policy; Inflation; Surveys

*JEL Codes*: E31, E52, E58

* Knotek: Federal Reserve Bank of Cleveland, Email: Edward.knotek@clev.frb.org; Mitchell: Federal Reserve Bank of Cleveland, Email: james.mitchell@clev.frb.org; Pedemonte: Inter-American Development Bank, Email: mathieu.pedemonte@gmail.com; Shiroff: Federal Reserve Bank of Cleveland, Email: taylor.shiroff@clev.frb.org.
We thank Oli Coibion, Francesco D'Acunto, Ricardo Reis, Michael Weber, and Mirko Wiederholt for helpful comments, as well as conference and seminar participants at the Federal Reserve Bank of Cleveland, the Oxford Communication Workshop, the Bank of Canada, the Cleveland Fed's Central Bank Communications: Theory and Practice conference, the Second Paris Conference on the Macroeconomics of Expectations, ISF 2024, Venice Summer Institute, and the CEBRA Annual Meeting 2024. The RCT was registered as AER RCT ID AEARCTR-0009172 (Knotek 2022). The views expressed here are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

## 1. Introduction

Monetary policymakers on the Federal Open Market Committee (FOMC) raised the federal funds rate at a relatively rapid pace during 2022 via a sequence of rate increases that began on March 16, 2022. The primary reason for this rapid tightening of policy was mounting concern that inflation was proving long-lived, with the risk of inflation expectations becoming unanchored. This led FOMC participants to talk openly and in advance of FOMC meetings about the pivot to a tighter monetary policy stance and the need for higher interest rates. While these communications were reflected in rising market-based expectations of the future federal funds rate, it is not clear whether the consumers comprising the "general public" were paying attention or, if they were, how they were reacting to these communications.[1]

This paper estimates the causal effects of communicating interest rate increases on consumers' inflation expectations using five waves of a randomized controlled trial (RCT) conducted via an online survey. We focus on communicating to consumers in very simple terms the federal funds rate increases of 2022 and assessing the impact on their inflation expectations. We do so by conducting in real time a specially designed set of RCTs immediately following the March, May, June, July, and September 2022 FOMC meetings. Each of these meetings resulted in increases to the federal funds rate target. Our RCT-based estimates of the *average* causal effect of communicating these hikes on households' inflation expectations have a wide range. We find that providing consumers with information about the latest interest rate hike reduced their expectations for inflation over the next five years, on average, between 0.17 and 2.18 percentage points, depending on the precise information transmitted.

While these treatment effects are statistically significant on average, we raise two difficulties about their interpretation that much of the wider RCT-based literature in empirical macroeconomics also faces. First, our RCTs treat consumers with information that is already publicly available and, moreover, they do so in an environment where monetary policy news was salient given high inflation readings. We provide novel evidence that consumers were indeed more informed about monetary policy in 2022 than they had been in prior years.[2] As a result, some treated individuals likely already had the treatment in their ex ante information sets, a phenomenon

[1] See https://www.cmegroup.com/insights/economic-research/2022/fed-rate-hikes-expectations-and-reality.html.
[2] This finding that consumers were more aware of monetary policy in 2022 than in prior years complements other work that has documented rising awareness of and attention paid to inflation in particular during this episode, such as Pfäuti (2024) and Bracha and Tang (2024).

also described in Weber et al. (2024) and Mackowiak and Wiederholt (2024). Second, while our RCTs explicitly communicate information to the treatment groups, it is hard to know how much of this information is really read and/or absorbed by respondents within the survey. We make the analogy to medicine and a physician prescribing a pill to a patient: the patient is only truly treated after swallowing the pill, not by having been prescribed or even given the pill.[3] Voluntary compliance with the treatment can only be confirmed ex post.[4]

Methodologically, we introduce to empirical macroeconomic studies that use RCTs the importance of ex ante informedness and ex post compliance. We propose a novel, easy, and accessible way to control for both the ex ante information set of respondents and how compliant they are with the information treatment ex post in a manner that is feasible in most online surveys. In our application, we capture informedness by asking consumers whether they had recently heard monetary news, and we measure voluntary compliance—rather than attempting to enforce compliance—through the time that respondents choose to spend reading the information treatment in the online survey. In this way, we proverbially do not force patients to take the pill, but we observe through a two-way mirror whether they choose to take the pill.

Using econometric methods familiar to microeconomists when undertaking causal inference with non-experimental data, we obtain complier average treatment (or causal) effects by upweighting (downweighting) control group respondents based on their predicted probability of complying (not complying) with the treatment and excluding treatment group non-compliers. After accounting for compliance with the treatment and controlling for consumers' ex ante informedness in our regressions, we find that monetary policy communications about increases in interest rates have a statistically significant and economically meaningful negative impact on medium-term inflation expectations for the consumers who were previously uninformed about recent policy actions and compliant with the treatments. Thus, we document a causal channel through which interest rate increases reduce inflation expectations. To the extent that the previously informed individuals in our survey swallowed the real-world monetary policy pill upon hearing monetary policy news, then we expect they too reduced their inflation expectations at that time, implying

[3] In the 1999 movie "The Matrix," when Morpheus gives Neo the choice between taking the blue pill and staying blissfully unaware, or taking the red pill and joining the movement to undermine the matrix, he watches as Neo takes the red pill and washes it down with a glass of water.

[4] An alternative interpretation is that there is a key intensive margin to treatments that can vary across respondents and may not be completely random. Fuster et al. (2022) find that people with less uncertain prior beliefs are more likely to spend time reading the information treatment.

that monetary policy actions were helpful in (re-)anchoring inflation expectations during this episode. However, RCTs are unable to adequately capture their behaviors, because the information treatments are already part of their information set, as noted in Weber et al. (2024).[5]

Our empirical results contrast with the findings in Andre et al. (2022), who reported strong disagreement among consumers' responses to hypothetical situations involving monetary policy communications based on survey responses from a low-inflation period in 2019. Our findings, coming from a high-inflation period when consumers may respond differently to information shocks, suggest that consumers had a better understanding of the objectives of monetary policy when given the right information—they understood the basic mechanism that higher interest rates would reduce inflation. In a theoretical rational inattention model, Mackowiak and Wiederholt (2024) show that individuals have more incentive to pay attention to the macroeconomy and to comply with information treatments when inflation is high. We explicitly measure both margins and show empirically that the treatment effect, consistent with their model, is higher for the uninformed and for those consumers who pay attention during the survey.

By observing rather than enforcing compliance, we can analyze consumers' choice to comply with the treatment. We document systematic demographic differences across individuals who are both less informed about monetary policy and more compliant with the treatment, and we show that these differences are relevant when interpreting average and heterogeneous effects of information treatments in the macroeconomics RCT literature. For example, Coibion, Gorodnichenko, and Weber (2022) find that communication directly by the FOMC is more effective in moving household expectations than indirect communication via the media. They also find considerable heterogeneity across respondents, with female respondents' inflation expectations reacting more strongly to monetary policy information treatments. In our RCT, we find that women are more likely to spend a longer time reading the information treatments than men and to thus be considered as having complied with the treatment. This finding means that the gender differential in Coibion, Gorodnichenko, and Weber (2022) may simply reflect women

[5] At the same time, we do not believe that an event study around the FOMC meetings in our sample is appropriate. Increases in the federal funds rate during our sample were telegraphed in advance of the meetings—subject to some uncertainty about their size—and hence it is unclear when the "events" transpired. Furthermore, we provide ample evidence below that there are systematic differences across individuals who are and who are not informed about monetary policy. Thus, comparing individuals who had not heard about monetary policy before an FOMC meeting and those who had heard about monetary policy after the meeting likely confounds a number of factors beyond the actual information provided by the meeting—i.e., selection into informedness is not random.

paying more attention to the RCT treatment than men. We find support for this conjecture in our study: after we control for informedness and compliance, the gender gap is no longer statistically significant.

Our analysis also allows us to consider issues related to monetary policy awareness more broadly. Our findings suggest that central bank communications could be augmented to reach consumers who typically do not hear much monetary policy news. At the same time, our compliance results show that even when portions of the public are presented directly with monetary policy information, there is no guarantee that the information will be processed. Our focus on monetary policy awareness and compliance thus dovetails with the work of D'Acunto, Fuster, and Weber (2021), who find that the salience of female and minority representation on the FOMC affects how Fed information influences consumers' expectations, particularly for selected demographic groups, potentially offering a pathway to enhance real-world compliance to engage with monetary policy news.

**Related Literature**

RCTs have gained prominence in empirical macroeconomics to understand expectations formation; e.g., see Armantier et al. (2016), Binder and Rodrigue (2018), and Coibion, Gorodnichenko, and Weber (2022). Using RCTs, Haldane and McMahon (2018) and Bholat et al. (2019) find simple relatable communications by the central bank to be more effective in influencing households' expectations, thus motivating our focus on information treatments that are "short and sweet." To test whether the information treatment has additional power to affect expectations if the rationale for the policy change is also communicated, additional treatment groups in our survey are given some narrative or "vignette" (see Andre et al., 2022) around the rate increase, e.g., by explaining that the FOMC is raising rates to reduce inflation. Our paper thus revisits the question of whether it is best to communicate targets or instruments, but it does so for the US, orienting the information treatments around the actual federal funds rate decisions made by the FOMC through 2022 rather than hypotheticals. In our case, we provide simple treatments and compare that treatment to others that provide additional information.

Our results build on a body of literature that studies communication as a central policy or lever of modern central banking. Much of this work has focused on whether and how central banks can communicate to consumers and firms (e.g., Blinder et al., 2008; Blinder et al., 2024; Binder,

2017). Effective communication is important when central banks want to shape the expectations of firms and consumers (Coibion et al., 2020). This RCT-based literature has found that presenting consumers with simple statistics about current and past inflation, the FOMC's inflation target, and inflation forecasts can affect inflation expectations. In turn, Coibion et al. (2023a) show that forward guidance about policy rates can have strong effects on household expectations. It is also well documented from this literature that inflation expectations, and their determinants, vary across demographic groups. Recent studies have focused on the type of instrument for communicating (D'Acunto et al., 2020) or the type of communication method used (Gorodnichenko, Pham, and Talavera, 2023, and Pedemonte, 2024). In this paper, we emphasize the importance of measuring a respondent's information set, as stressed in Weber et al. (2024), which can have important implications for the subjective model(s) of the economy that consumers use to process information treatments, as studied in Andre et al. (2022). In addition, especially in online survey settings, measuring compliance is key, as has been stressed in other fields, notably medicine. Understanding whether there is compliance is particularly important given growing evidence that the public does not pay much attention to central bank communications; see Coibion et al. (2023b). Our paper thus contributes to the literature on the effectiveness of monetary policy communication by identifying those people most likely to be both informed and compliant if not initially informed.

While we find some evidence of stronger treatment effects when consumers are informed about not just the change in the interest rate (instrument) but also the intent of the policy change (to reduce inflation), we nevertheless find strong evidence that communicating the policy change itself is effective in lowering medium-run inflation expectations, especially for those who were previously unaware of but willing to pay attention to the communications. Our results therefore suggest that consumers did have some common understanding of the monetary policy transmission mechanism and of the source and intent of the FOMC's policy actions in 2022. Absent this specific context, the "information effect" (see Nakamura and Steinsson, 2018; Coibion, Gorodnichenko, and Weber, 2022) could in principle dominate a priori: consumers could interpret FOMC announcements of increases in the federal funds rate as signaling a stronger economy. Under this interpretation, inflation could be expected to increase, and it is that anticipated increase that is driving the funds rate higher. This finding contrasts with findings in Andre et al. (2022), where both increases and decreases in inflation are predicted by substantial numbers of consumers in response to a *hypothetical* rise in the federal funds rate. Our paper shows that in the specific and

real-world context of the *actual* 2022 rate hikes in the US, on average consumers saw the demand-side effects as dominating, and accordingly they lowered their medium-term inflation expectations.

Our paper is also broadly related to the literature on inattention. Using a New Zealand survey, Coibion, Gorodnichenko, and Kumar (2018) found that firm managers have dispersed expectations of inflation. This is attributed to their inattention to inflation, opening the door for information treatments in RCTs to be effective in causing firms to update their heterogeneous inflation expectations. But Weber et al. (2024) show that, in higher inflation environments, this need not be true, as households and firms are more informed about inflation and so adjust less to the exogenously provided inflation treatment. Our results suggest that the real-world effects of monetary policy communications on inflation expectations are likely stronger than what one would estimate based on naïve average treatment effects, because these would omit the change in expectations that had already happened prior to the RCT. Since the information in our treatments is publicly available, and moreover because the federal funds rate increases that commenced in March 2022 were largely telegraphed in the media ahead of FOMC meetings, our finding that some consumers do update their inflation expectations when treated with information about monetary policy actions is consistent with models of imperfect information (see Maćkowiak, Matějka, and Wiederholt, 2023, for a review) in which agents do not fully update their information sets or incorporate all available information into their expectations. Indeed, our surveys provide direct empirical evidence that not everyone pays attention to monetary policy news, even during a period of high inflation. But, consistent with models of rational inattention (Sims, 2003), another novel feature of our survey is that it finds that a growing proportion of consumers did follow monetary policy news amid high and rising inflation in 2021 and 2022. Peaks in attentiveness coincide with FOMC meetings. Lamla and Vinogradov (2019) reach similar conclusions on an earlier lower inflation sample, when conducting surveys shortly before and after FOMC meetings.

The structure of the paper is as follows. Section 2 sets out the design of our RCTs. Section 3 presents tests showing that the information treatments were effective in reducing inflation expectations for the *average* consumer. Because the information treatments are conveying public information, in Section 4 we distinguish between consumers who were previously informed or uninformed about monetary policy actions. Section 5 uses propensity score weighted regressions to produce estimates of the treatment effect that adjust for less-than-perfect compliance in the treatment groups, thereby acknowledging the reality that respondents in RCTs need not always

comply with the treatment. Section 6 discusses the implications of our results and shows evidence that even uninformed and compliant consumers, who do expect inflation to fall because of the rate hikes, do not consistently expect this disinflation to be accompanied by a worsening real economy. Section 7 concludes. Online appendices contain supplementary tables, figures, and results.

## 2. Data and Design of the Randomized Controlled Trials

### 2.1 Survey Background

Our sample was collected as part of a daily online survey of consumer expectations conducted by the Federal Reserve Bank of Cleveland and administered by Qualtrics Research Services (see Knotek et al., 2020; Dietrich et al., 2022; Coibion et al., 2023b). Respondents are representatively drawn from several actively managed, double-opt-in market research panels, complemented using social media (Qualtrics, 2019). The survey has been conducted daily since the onset of the COVID-19 pandemic in March 2020. While our focus will be on the survey data during 2022, given the informational treatments we embedded into the survey the day after the FOMC raised interest rates in March, May, June, July, and September 2022, we will also make use of the historical data. In the survey, consumers are asked a variety of questions about their demographic characteristics, economic expectations, and their exposure to news about monetary policy and the economy.[6] In total, our main sample consists of 33,728 responses collected across the five treatment waves in 2022, with roughly 100 responses per day and over 5,000 consumers per wave. The five waves began in March, May, June, July, and September on the day after the respective month's FOMC press statement was issued and ran through the Wednesday before the next meeting. Each wave contains a repeated cross-section of observations; respondents are not followed over time. As described below, we reweight our respondents to ensure that our sample is representative of the US population.[7]

Appendix A shows the full statements from the FOMC issued immediately after each of the five meetings. In summary, the policy actions, which form the basis of our informational treatments, were that after the March 2022 meeting the FOMC raised the target range for the

[6] A list of the main survey questions is available in Appendix H.

[7] Haaland, Roth, and Wohlfart (2023) discuss online surveys and their consistency with surveys conducted using more traditional modes. As suggested in that work, our survey respondents receive monetary compensation for their time (the amount received is proprietary); we use ReCAPTCHA scores to reduce the incidence of bots; and we exclude respondents who spent extremely short (<1$^{st}$ percentile) or long (>99$^{th}$ percentile) times on the survey.

federal funds rate by ¼ percentage point, to a range of ¼ to ½ percent. In May, they raised the target funds rate by ½ percentage point, to a range of ¾ to 1 percent, and they also implemented the plan to start reducing the size of the Fed's balance sheet. In June, the target funds rate was raised by ¾ percentage point, to a range of 1½ to 1¾ percent; in July, the target funds rate was raised by ¾ percentage point, to a range of 2¼ to 2½ percent; and at the September meeting, the FOMC raised the target funds rate by ¾ percentage point, to a range of 3 to 3¼ percent.

To enable a post-treatment versus pre-treatment comparison, we first directly measure inflation expectations for the next 12 months. To this end, we ask "Over the next 12 months, do you think that there will be inflation or deflation? Please give your best guess." Depending on the answer, the respondent proceeds to a follow-up question asking for a numerical estimate: "What do you expect the rate of inflation (deflation) to be over the next 12 months?"[8]

We next ask respondents about their knowledge of interest rates and monetary policy. In particular, we asked whether they thought that interest rates in general had changed recently and whether they heard any news related to monetary policy in the last week. As we discuss in detail in Section 3, we find that most respondents were aware that broad interest rates had gone up recently, but a majority had not heard news about monetary policy.

At this stage, respondents are randomly assigned to a control group or one of the treatment groups. The control group does not receive any additional information. The treatment groups are provided with an information treatment. For the majority of respondents who reported that they had not heard news about monetary policy, the treatment is presumably news to them; Coibion et al. (2023b) show that, at earlier points within this same survey, respondents were generally inattentive to recent monetary policy decisions. The total number of treatments in each wave was four in March; two in May, June, and July; and three in September. The exact content of the information treatment varied with each wave, as the information changed between waves based on the most recent monetary policy announcement. Nonetheless, the objective for each treatment is to provide information about recent FOMC policy decisions.

Treatment 1 in each wave provides a simple description of the FOMC's most recent action:

**Treatment 1**: "*On [date of most recent FOMC press release], the Federal Open Market Committee (FOMC) raised its primary policy interest rate (the federal funds rate) by [fraction] of a percentage point, to a target range of [lower bound] to [upper bound]*

[8] We ensure that inflation responses are non-negative and deflation responses are non-positive.

*percent. The FOMC also said that it would [begin/continue] to reduce the size of its balance sheet.*"

The numbers provided in Treatment 1 were updated in each wave. The wording and length of this treatment, and of those that follow, are deliberately chosen to be to-the-point and factual and use FOMC language to mimic the sort of statement that might be communicated via a social media "tweet." Importantly, this treatment only contains information about the policy action. It does not provide any suggestive information about either the effect of the action on inflation or the motivation for the action. Andre et al. (2022) showed that, in a period of low inflation, consumers disagree about the inflationary effects of monetary policy. As this finding might change in higher inflation environments, we add additional treatments to see whether explaining the reasons for the policy change or the policy objective itself delivers stronger effects on consumers' expectations.

Treatment 2a provides the same information about *what* the most recent monetary policy action was, but it also provides information (a narrative or "vignette") about the *intent* of the action. This treatment was included in the March through July waves:

**Treatment 2a**: Treatment 1 plus: *"These actions were part of an effort to help bring inflation back down toward its objective."*

In September, we instead included information about the longer-term policy objective:

**Treatment 2b**: Treatment 1 plus: *"Federal Reserve Chair Jerome Powell said, 'The FOMC is strongly resolved to bring inflation down to 2 percent and we will keep at it until the job is done.'"*

The first wave in March also included two additional treatments that further expanded on the intent of the policy action as provided in Treatment 2:

**Treatment 3**: Treatment 1 plus *"The FOMC said that, 'with appropriate firming in the stance of monetary policy, the Committee expects inflation to return to its 2 percent objective and the labor market to remain strong.'"*

**Treatment 4**: Treatment 2 plus *"[The FOMC] anticipated that ongoing increases in the target range will be appropriate."*

Finally, a "placebo" treatment was included in September. This involved giving consumers information that had numerical content identical to the other two treatments above but which should be irrelevant to forming inflation expectations. The objective is to control for whether the

provision of this information is changing consumers' answers, or if giving a numerical anchor is changing them as well. The placebo describes population growth:

> **Placebo**: *"From 2015 to 2021, the population in the United States grew in a range of 3 to 3-1/4 percent.*"

After randomly receiving one of the information treatments or being allocated to the control group that did not receive any information treatment, respondents were asked for their (posterior) inflation expectations. To avoid possible attrition or survey fatigue from asking consumers the same question twice, this time we asked for their expectations over the next five years, as in Coibion et al. (2023b). As with the first question on inflation expectations, we first asked whether the respondent expected inflation or deflation: "Over the next 5 years, do you think there will be inflation or deflation on average?" This was followed up with a question asking for numerical input: "What do you expect the rate of inflation (deflation) to be over the next 12 months? Please give your best guess."

### 2.2. Weighting, Randomization, and Data Quality

Our survey sample comprises a large, representative, and reasonably high-quality sample of the US adult population. As in other survey research, including the Federal Reserve Bank of New York's Survey of Consumer Expectations (SCE) and Roth and Wohlfart (2020), our sample is somewhat over-educated relative to the adult population, and, unlike these other studies, our sample is also somewhat younger. These caveats aside, our sample is generally representative of the over-18 US population (see Appendix Table B1). Nevertheless, we use raking to calculate survey weights to bring our sample in line with target population proportions sourced from the 2019 American Community Survey across gender, income, age, education, race, ethnicity, and region. These weights are calculated separately for each treatment period, so that each treatment period's sample is a representative sample. While we make use of these weights in much of our analysis, omitting them has little impact on our results.[9] We do not use demographic controls in any of our regression models unless otherwise specified.

[9] These weights are not used in the propensity score weighted regressions in Section 5. Rather than reweighting to match population distributions, these regressions reweight the control group to match distributions within the treatment group compliers/non-compliers.

Our sample is well-randomized and balanced, with only minor differences between the demographic make-up of any of the treatment groups and the control groups (see Appendix Table B2). A daily quota is used to ensure that the appropriate quantity of responses is collected even if respondents drop out of the survey before completing it. Respondents who fail the survey's ReCAPTCHA check or who are otherwise flagged by Qualtrics as likely bots or spammers do not count toward the quota, nor are they included in our data set.

We drop or otherwise alter as few responses from our sample as possible. The only respondents who are outright removed from our sample are those whose total survey completion times are either too short or too long. We drop respondents who took less than 6 minutes (n=574 consumers), approximately the 1st percentile, as well as those who took more than an hour to complete the survey (n=220), since the quality of responses from these respondents is typically poor, owing to rushing through the survey or simply forgetting about it. The median respondent took 14.4 minutes to complete the survey, and just under 90% completed it in less than 25 minutes. For the July and September survey waves, a timer was enabled to capture the time that respondents spent reading the screen with an information treatment, if relevant. Importantly, respondents choose how long to spend on the information treatment page in our survey; we do not force them to spend a specific amount of time on this page.

To deal with outliers, we winsorize responses at the 2nd and 98th percentiles for all point expectations and use Huber-robust regressions. Twelve percent of our sample reported that they were expecting deflation over the next year; among these respondents, 60% anticipated deflation between zero and 10%.

The raw median (across consumers and over time) response to the question on prior (that is, pre-treatment) year-ahead inflation expectations was 8%, with the 25th percentile at 4% and 75th percentile at 20%. The median posterior (that is, post-treatment) five-year average inflation expectation was 5%, with the 25th percentile at 2% and 75th percentile at 15%. After applying Huber weights, the median, 25th percentile, and 75th percentile prior expectation are 7%, 4%, and 10%, respectively; and 5%, 2%, and 8% for the posterior.

## 3. Treatment Effects of Different Communication Tools

To estimate the average treatment effect, $\beta_j$, for treatments $j=1,\ldots,J$, we run the following regression:

$$\pi_{i,t}^{5y} - \pi_{i,t}^{1y} = \alpha + \sum_{j=1}^{J} \beta_j \times I\left(1\ if\ treatment_j = 1\right) + \varepsilon_{it}, \qquad (1)$$

where $\pi_{i,t}^{5y}$ is the posterior, five-year inflation expectation for individual $i$ in wave $t$; $\pi_{i,t}^{1y}$ is the prior, 12-month inflation expectation; $J$ is the number of treatments within a wave; and $I\left(1\ if\ treatment_j = 1\right)$ is a 0-1 dummy that takes a value of 1 if respondent $i$ received treatment $j$. If treatment $j$ is effective in changing the posterior inflation expectation relative to the prior inflation expectation, then $\beta_j$ will be different from zero, implying that the treatment induces a different response on average for individuals who receive the treatment relative to the control group. Note that $\beta_j$ measures the average treatment effect; changes to the distribution of responses that do not affect the average will not be captured in our regression. A value of $\beta_j < 0$ signifies that treatment $j$ lowers respondents' inflation expectations relative to their prior on average. To help filter outlier responses, we apply Huber weights obtained from a similar regression of five-year inflation expectations on one-year expectations, treatment indicators, and their interactions.

Table 1 shows the estimation results for each wave separately, as well as a pooled version that includes a wave fixed effect and allows us to talk about general effects over the full sample. In that column, the wave fixed effects are important because they control for the common information that the treated and control groups had at the time.

Table 1 reports a negative average treatment effect for each treatment, implying that the *average* respondent reacts to the information treatment by reducing their inflation expectations. This negative effect is present even when consumers only receive information about the new federal funds rate (Treatment 1). This suggests that consumers may have some understanding of the mechanism behind monetary policy actions. This is particularly relevant for Treatment 1, as it does not include information about either the policy objective or inflation. The effect for Treatment 2, however, seems to be smaller in magnitude. This could be explained by an "information effect" (see Nakamura and Steinsson, 2018) that can confound communications about interest rates: when consumers are told the aim of the FOMC's actions is to reduce inflation, there is implicitly an acknowledgment of an inflationary problem, which may reduce the size of the treatment effect.[10]

[10] In Appendix C, we complement Table 1 by showing that the treatment effects are larger for those consumers with higher prior expectations for inflation.

Table 1: Posterior Minus Prior on Treatments

| | (1) March | (2) May | (3) June | (4) July | (5) Sept | (6) Pooled |
|---|---|---|---|---|---|---|
| Treatment 1 | -1.90*** | -0.13 | -0.38* | -1.18*** | -0.35 | -1.55*** |
| | (0.33) | (0.19) | (0.21) | (0.28) | (0.32) | (0.13) |
| Treatment 2 | -0.25 | -0.17 | -0.45** | -0.84*** | -1.59*** | -0.35*** |
| | (0.21) | (0.18) | (0.21) | (0.29) | (0.36) | (0.09) |
| Treatment 3 | -0.50** | | | | | -2.08*** |
| | (0.21) | | | | | (0.30) |
| Treatment 4 | -0.68*** | | | | | -0.61*** |
| | (0.21) | | | | | (0.19) |
| Placebo | | | | | 1.09*** | 0.36** |
| | | | | | (0.27) | (0.18) |
| Observations | 7961 | 5212 | 4899 | 5996 | 5876 | 29884 |

Columns 1-5 contain no controls or fixed effects. Column 6 shows a pooled regression with wave fixed effects. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

There are some effects in Table 1 that are hard to rationalize. For example, the placebo treatment seems to increase inflation expectations. We re-evaluate these results below, once we have accounted for consumers' ex ante informedness about monetary policy and their ex post compliance with the information treatment.

## 4. Informed and Uninformed Respondents: Expanding the Reach of Monetary Policy

Table 1 showed that the information treatment was effective in reducing the inflation expectations of the *average* respondent. But since this information treatment is conveying public information at the time of the experiment, in principle our treatment should only affect respondents who are uninformed about recent monetary policy decisions. Informed respondents should already know the information communicated in the treatment and hence should be unaffected by it. In this section, we explore whether there are differences between the respondents who receive the treatment and are likely uninformed (the "local average treatment effect") relative to the effect on all of the treated (the average treatment effect).

We start by identifying the uninformed group and exploring its demographic, behavioral, and socio-economic characteristics. We do so by exploiting the fact that our survey includes a pre-treatment question that asks respondents if mortgage rates had changed recently (and, if so, how) and whether they had heard news about monetary policy. As our treatments provide information about monetary policy decisions and changes in interest rates, these questions help identify which respondents were likely to have already known the information in the treatment. We find that

82.4% of respondents who had heard news about monetary policy were also aware that interest rates had changed recently. Only 47.3% of respondents who had not heard monetary policy news were aware of such an increase, and only 30.2% correctly described the size of the change, less than half of the share among those who had heard monetary policy news. Since the question on monetary policy news is indeed correlated with informedness about interest rates, we use whether respondents had heard news as our indicator of informedness.[11] Figure 1 shows the evolution of the answers to this question over time.

Figure 1 reveals that the share of respondents who indicate that they have heard news about monetary policy has been rising since 2020. While only between 25% and 30% of respondents heard news about monetary policy in October 2020, around 45% of respondents had heard news of monetary policy in October 2022. This is consistent with the rise in inflation and the consequent increased public discussion of monetary policy through 2022. In addition, we also see from Figure 1 that respondents are more likely to have heard news about monetary policy immediately after FOMC meetings, as represented by the vertical dashed lines in Figure 1.[12]

We should expect smaller treatment effects for respondents who have already heard news about monetary policy. To test this hypothesis, we re-run the regressions in Table 1 but only on that subset of respondents who reported that they had not heard news about monetary policy. Note that random assignment to a treatment group is not correlated with having heard news about monetary policy by design (see Appendix Table B2).[13]

[11] Using the question on interest rates also introduces additional complications. For example, since the possible initial responses were "No," "Not Sure," and "Yes," we would need to consider the cases of truly uninformed (i.e., "No" respondents) and little-informed (i.e., "Not sure" respondents), under the tenuous assumption that these labels are accurate. Using these as alternative indicators of informedness does not significantly change our results.
[12] Using a stratified random sample of the US public two days before and two days after the FOMC press conference, Lamla and Vinogradov (2019) also find that monetary policy announcements lead to an increase in the proportion of people who have heard monetary policy news.
[13] Appendix Table E2 presents estimates showing how the probability that respondents had heard news about monetary policy varies with demographic factors and as a function of the time since the most recent FOMC meeting.

Figure 1: Percentage of Respondents Hearing News about Monetary Policy

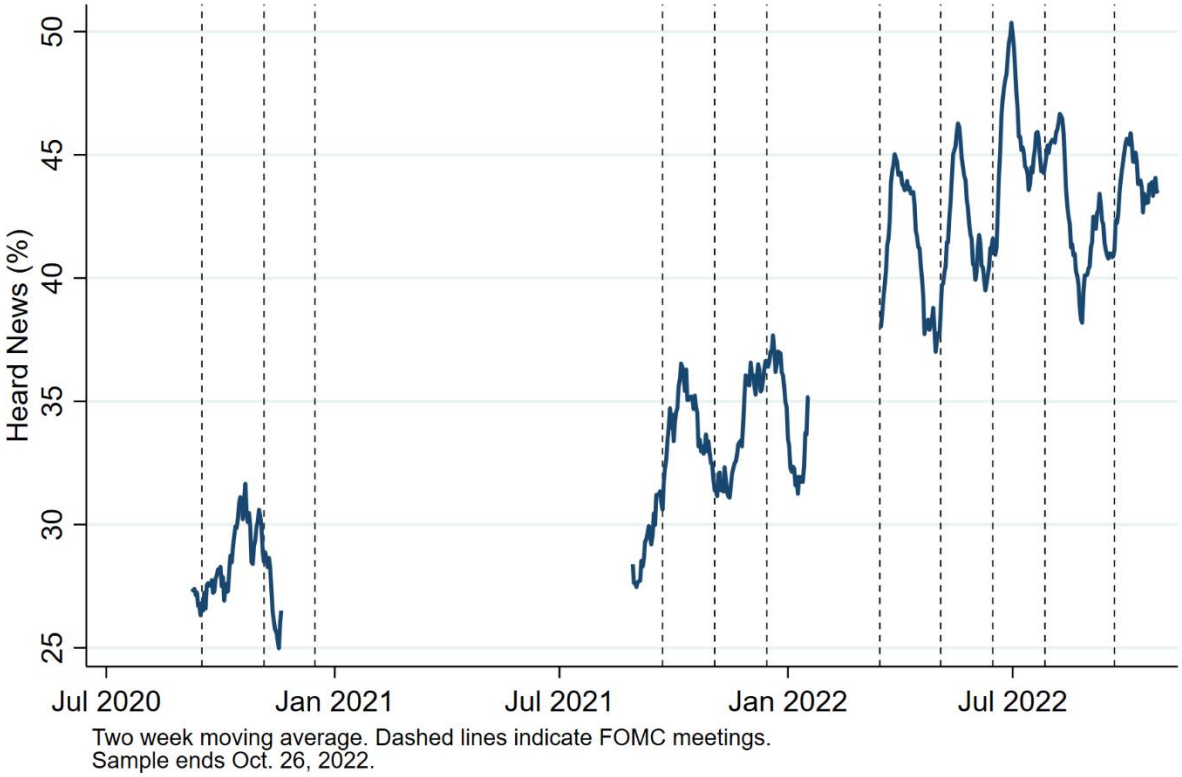


Table 2 confirms that the estimated treatment effects for respondents who have *not* heard news are indeed often larger than those reported in Table 1, likely because for these respondents the information contained in the treatment is more informative than it is for those who have already heard some monetary policy news. In particular, in July and September we observe stronger treatment effects in Table 2 than in Table 1.

This differential effect speaks to a general limitation of RCTs. The treatment effect of communicating information depends on how "informative" the information provided actually is. Respondents who are already fully informed have no need to update their priors post-treatment. In the next section, we propose a way of distinguishing between the informational content of the treatment using a measure of whether the respondent complied with the treatment.

Table 2: Posterior Minus Prior on Treatments, Heard News = No

| | (1) March | (2) May | (3) June | (4) July | (5) Sept | (6) Pooled |
|---|---|---|---|---|---|---|
| Treatment 1 | -1.29** | -0.19 | -1.34* | -2.03*** | -0.93** | -1.36*** |
| | (0.62) | (0.59) | (0.70) | (0.45) | (0.47) | (0.26) |
| Treatment 2 | 0.21 | -0.01 | 0.73 | -1.30*** | -2.83*** | -0.72*** |
| | (0.46) | (0.56) | (0.58) | (0.43) | (0.61) | (0.23) |
| Treatment 3 | -1.25** | | | | | -1.37*** |
| | (0.55) | | | | | (0.50) |
| Treatment 4 | -0.70 | | | | | -0.84* |
| | (0.49) | | | | | (0.44) |
| Placebo | | | | | -0.62 | 0.01 |
| | | | | | (0.45) | (0.41) |
| Observations | 4522 | 3075 | 2669 | 3503 | 3339 | 17110 |

Columns 1-5 contain no controls or fixed effects. Column 6 shows a pooled regression with wave fixed effects. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

## 5. Who Reads the Treatment? Measuring Compliance

Assigning a respondent to a treatment group does not guarantee that the respondent complies with—and thus actually receives—the full extent of the information contained within the treatment. A parallel can be made with medical RCTs, in which patients may be randomly assigned the treatment, but they may not comply, for example, by not swallowing the pill prescribed by the physician. In our case, there could be many reasons for non-compliance. Consumers might be inattentive within the survey itself: they could be distracted, for example, and continue to the next question without having processed the treatment. Alternatively, they might not be willing to pay attention to the treatment at all, and they simply skip to the next question as fast as possible without digesting the information. While demographic characteristics and question-based assessments of reading and/or numerical literacy can help control for varying levels of attentiveness and understanding of the treatment, they cannot assess who reads—and hence complies with—the treatment and who does not. Because non-compliance may be self-selected, it is "nonignorable," since it undermines the random allocation into the treatment group required for unbiased estimation of the local average treatment effect (the complier average causal effect); see Imbens and Angrist (1994) and Angrist, Imbens, and Rubin (1996). Therefore, the estimates in Tables 1 and 2 offer unbiased estimates of the effect of assignment, the so-called "intention-to-treat," not of the treatment itself. They reflect the efficacy of both the treatment and the compliance.

In this section, we propose a novel tool to measure whether respondents read (that is, comply with) the treatment, which we then use to estimate the average treatment effect free from any confounding effects of non-compliance. This is facilitated by the fact that in the July and September waves we measured how much time respondents chose to spend on the treatment page.[14] We use this measure as a proxy for whether respondents read and processed the treatment; we then separate the sample into attentive and inattentive respondents. Because the treatments vary in length, content, and across waves, separate rules are calculated for each treatment in each wave. We use a rule whereby a respondent is considered to have "read the treatment" if they took at least half of the average amount of time spent by respondents assigned to their treatment in their wave to read the treatment. This cutoff is 4.1 and 5.4 seconds for Treatments 1 and 2 in July, respectively, and 4.3, 5.9, and 2.5 seconds for Treatment 1, Treatment 2, and the placebo in September, respectively. Our key results are not sensitive to using different cutoffs.[15]

While we cannot know from reading time alone whether a respondent truly read, processed, and correctly understood the information in the treatment, we can reasonably assume that those with reading times below these cutoffs did not do so. Even though our treatments provide succinct, "tweet-style" snippets of information, it is difficult to imagine that a respondent could read the entirety of the treatment in fewer than 4 or 5 seconds. Our measure of compliance, therefore, is a conservative one. There may very well be respondents we consider compliant who in fact were not compliant, but very few respondents whom we could be incorrectly considering non-compliant, given our reading cutoff times.

Table 3 shows how demographic and other characteristics correlate with attentiveness to the treatment. We see that women tend to pay more attention to the treatment, even conditional on not having already heard news. In addition, respondents who are white; non-Hispanic; older; more educated; and have greater numerical literacy are more likely to read the treatment. In total, we can explain a good proportion of compliance based on measurable characteristics that are not idiosyncratic (e.g., how tired a respondent is). This result is useful because it allows us to

[14] This type of information is easy to implement in online surveys. Fuster et al. (2022) use a similar counter to track the relevance of information. They measure time spent reporting the posterior. They find a positive relationship between uncertain priors and the time spent reporting the posterior.

[15] We discuss alternative cutoffs below. In general, our "half-the-average" rule requires a treatment reading time slightly below the median. For a plot of the empirical distribution of reading times, see Appendix Figure D1.

characterize potential candidates who are likely to be attentive in the control group, and therefore to see if the effect of the treatment is stronger for consumers who pay attention.

Table 3: Likelihood of Reading the Treatment

| | Full Sample | | | | No News Only | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) Logit Odds Ratio | | (2) OLS | | (3) Logit Odds Ratio | | (4) OLS | |
| Male | 0.81*** | (0.05) | -0.04*** | (0.01) | 0.81** | (0.07) | -0.04** | (0.02) |
| Nonwhite | 0.64*** | (0.05) | -0.09*** | (0.02) | 0.62*** | (0.07) | -0.10*** | (0.02) |
| Hispanic | 0.64*** | (0.07) | -0.09*** | (0.02) | 0.65*** | (0.09) | -0.09*** | (0.03) |
| Primary Shopper | 0.90 | (0.10) | -0.02 | (0.02) | 1.10 | (0.13) | 0.02 | (0.02) |
| Numerical Literacy | 1.77*** | (0.16) | 0.11*** | (0.02) | 1.43*** | (0.17) | 0.07*** | (0.02) |
| Heard News | 0.93 | (0.06) | -0.02 | (0.01) | | | | |
| Age: | | | | | | | | |
| 36-50 | 2.11*** | (0.16) | 0.17*** | (0.02) | 2.22*** | (0.21) | 0.18*** | (0.02) |
| 51-65 | 5.24*** | (0.46) | 0.38*** | (0.02) | 4.79*** | (0.53) | 0.36*** | (0.02) |
| 66+ | 11.96*** | (1.26) | 0.51*** | (0.02) | 9.19*** | (1.37) | 0.47*** | (0.02) |
| Income: | | | | | | | | |
| $35,000-$49,999 | 1.17* | (0.11) | 0.03* | (0.02) | 1.22* | (0.13) | 0.04* | (0.02) |
| $50,000-$99,999 | 0.91 | (0.07) | -0.02 | (0.02) | 0.93 | (0.09) | -0.02 | (0.02) |
| $100,000 or more | 0.75*** | (0.08) | -0.05*** | (0.02) | 0.62*** | (0.09) | -0.09*** | (0.03) |
| Education: | | | | | | | | |
| Some College | 1.62*** | (0.13) | 0.09*** | (0.02) | 1.44*** | (0.14) | 0.07*** | (0.02) |
| Bachelor's Degree | 1.64*** | (0.15) | 0.10*** | (0.02) | 1.83*** | (0.21) | 0.12*** | (0.02) |
| Advanced Degree | 1.57*** | (0.18) | 0.09*** | (0.02) | 1.94*** | (0.32) | 0.13*** | (0.03) |
| Political Party: | | | | | | | | |
| Democrat | 0.93 | (0.07) | -0.02 | (0.01) | 0.87 | (0.08) | -0.03 | (0.02) |
| Republican | 1.08 | (0.09) | 0.02 | (0.02) | 1.03 | (0.11) | 0.01 | (0.02) |
| Constant | 0.46*** | (0.06) | 0.32*** | (0.02) | 0.42*** | (0.06) | 0.31*** | (0.03) |
| Observations | 9427 | | 9427 | | 5350 | | 5350 | |

Columns 1 and 3 report results from a logit model predicting compliance as a function of the listed variables as odds ratios, while columns 2 and 4 report results from OLS regressions of an otherwise identical model. All columns include a treatment wave fixed effect. Robust standard errors are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

We find a close relationship between the predicted probabilities of compliance and the Huber weights used in Tables 1 and 2: respondents with Huber weights close to 1 are predicted to be much more likely to read the treatment than those downweighted in the Huber-robust regressions (see Appendix Figure D2). In other words, those who are predicted to be more likely to comply with their assigned treatment are far less likely to provide outlier responses, and vice-versa. Accounting for compliance therefore has an added benefit in that doing so downweights

outlier responses on the basis of the respondents' behavior instead of using measures derived from the statistical properties of all (or a subset of) responses, as in the case of Huber weights.

While having heard news does not predict assignment to a treatment group, reading the treatment does; respondents could only pass the reading time cutoff if they were given a treatment to read. This problem is compounded by the predictability of compliance demonstrated in Table 3: selection into compliance or non-compliance is predictable based on respondent characteristics. Therefore, simply excluding non-compliers from the sample and rerunning the regressions in Tables 1 and 2 would no longer leave us with a randomly assigned treatment group. This implies that the average treatment effect estimates provided in Tables 1 and 2 are likely to understate the local average treatment effect for compliers, because those estimates do not account for the partial endogeneity of treatment selection; that is, respondents must be randomly assigned a treatment group to be treated, but they may or may not comply with being treated in a non-random fashion. This underestimation is understood by noting (e.g., Imbens and Angrist, 1994; Angrist, Imbens, and Rubin, 1996) that, under the exclusion restriction that treatment does not affect compliance, the local average treatment effect is the ratio of the estimated intent-to-treat effect (as shown in Tables 1 and 2) and the estimated proportion of compliers (as modeled in Table 3).

Given the evidence from Table 3 that we have covariates that explain compliance, we use these to calculate propensity scores that, in turn, are used to re-estimate the treatment effect. Following Jo and Stuart (2009), we use a two-step process to estimate complier and non-complier treatment effects. In the first step, similar to Follmann (2000), we estimate propensity scores ($\hat{p}$) for respondents in the control group of a given survey wave by fitting a logit model of compliance within the treatment group using the covariates seen in Table 3. In the case of additional sample restrictions (e.g., excluding respondents who heard news about monetary policy), the logit model is estimated on a subsample with the same restrictions, and fitted values are only calculated for respondents in the control group belonging to the same group (e.g., who also did not hear news about monetary policy). Because of randomization, the covariates used in the treatment group should also explain compliance in the control group, for whom we cannot directly measure compliance because they did not receive a treatment and thus were not timed.

In the second step, we return to the regression model in Tables 1 and 2 but now use the estimated propensity scores to reweight respondents. To estimate the average causal effect among compliers (CACE), respondents in the control group are assigned weights $\hat{p}_i/(1-\hat{p}_i)$, while

treatment group compliers and non-compliers are assigned weights equal to one and zero, respectively (thus equivalently excluding treatment group non-compliers). To estimate the average causal effect among non-compliers (NACE), respondents in the control group are assigned weights $(1-\hat{p}_i)/\hat{p}_i$, while treatment group compliers and non-compliers are assigned weights of zero and one, respectively (the inverse of the CACE weighting regime).[16] Intuitively, these regressions estimate complier (non-complier) treatment effects by giving the most likely and least likely compliers (non-compliers) within the control group the largest and smallest weights, respectively, effectively reweighting the control group as a whole to match the characteristics of the treatment group compliers (non-compliers).

Table 4 reports CACE and NACE estimates obtained via the two-step procedure outlined above, for the July and September waves and broken down by prior news exposure. To account for the estimation uncertainty from using estimated propensity scores in the second-stage weighted regression, we bootstrap each step of the two-step process. Table 4 reports the means and standard deviations of the distribution of each coefficient's draws, with p-scores calculated using these means and standard deviations.

Table 4 reveals large differences between compliers and non-compliers in both the size and statistical significance of the treatment effects. Among compliers, we find considerably larger treatment effects for the sample as a whole in column (1) and compliers without prior knowledge of recent monetary policy news in column (3). For the compliers for whom the treatment represented true news about monetary policy, we see that informing them of increases in the federal funds rate causally reduced their medium-term inflation expectations by an economically and statistically significant amount, as shown in column (3). For this group, the placebo treatment has no effect. We also see that compliers who had already heard news about monetary policy had no significant treatment effects in column (2), consistent with the news having already been in their information sets. For non-compliers, on the other hand, we find no significant treatment effects except for the placebo—which, given their non-compliance, might not be surprising.

Comparing the size of the treatment effects on the compliers who had not previously heard news across the different waves, it is noteworthy that Treatment 1 is more powerful than Treatment

[16] The reason that treatment group compliers and non-compliers get weights of 1 and 0, respectively (and vice-versa when estimating the NACE), rather than predicted scores, is that we know with certainty whether they passed the reading cutoff time or not. Using predicted scores over the observed compliance would discard this information.

2 in July, but the reverse is true in September. This may be because the September version of Treatment 2, which supplements the policy action with some discussion of its intent, includes a direct quote from Chair Powell. Previous research has also found that monetary policy communications resonate more when reinforced with a direct quote (e.g., Hoffmann, Moench, and Schultefrankenfeld, 2023). However, as discussed for Table 1, it is not clear that monetary policy communications are more effective when the reasons for the policy action are explained, since Treatment 3 in Table 1 is ineffective in reducing consumers' inflation expectations. Certainly, the words used to communicate the rationale for the policy seem to matter.

Table 4: Treatment Effects Via Propensity Score Weighted Regressions

| | Compliers | | | Non-compliers | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Heard News = | All | Yes | No | All | Yes | No |
| *July* | | | | | | |
| Treatment 1 | -1.76*** | 0.20 | -3.20*** | -0.92 | -1.12 | -0.93 |
| | (0.67) | (0.88) | (0.98) | (1.22) | (2.09) | (1.49) |
| Treatment 2 | -2.00*** | -1.20 | -2.68*** | -0.07 | 0.50 | -0.69 |
| | (0.69) | (0.86) | (1.03) | (1.24) | (2.16) | (1.47) |
| Observations | 4674 | 1994 | 2680 | 4284 | 1833 | 2451 |
| *September* | | | | | | |
| Treatment 1 | -1.79** | 0.04 | -3.24*** | 0.64 | 0.31 | 2.13 |
| | (0.79) | (1.16) | (1.07) | (1.39) | (2.25) | (1.65) |
| Treatment 2 | -2.57*** | -1.17 | -3.57*** | -0.68 | -2.11 | 1.72 |
| | (0.81) | (1.20) | (1.10) | (1.41) | (2.21) | (1.80) |
| Placebo | -0.04 | 0.35 | -0.42 | -0.99 | -4.48* | 3.22* |
| | (0.78) | (1.02) | (1.17) | (1.46) | (2.39) | (1.71) |
| Observations | 4339 | 1900 | 2439 | 3982 | 1834 | 2148 |

One hundred sets of fitted propensity scores are generated by bootstrapping a logit regression predicting compliance among the pooled treatment groups within each treatment period 100 times for each combination of prior information and compliance type (i.e., for each column within each panel). Then, each of the 100 sets of propensity scores from the first stage is used in a bootstrapped weighted OLS regression with another 100 repetitions. Reported above are the means and (in parentheses) standard deviations of the distribution of the resulting 10,000 "draws" of each coefficient. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

These results are robust to alternative cutoff parameters for compliance. Appendix Figure D3 shows that as we increase the reading time cutoff threshold used to classify compliance, the resulting treatment effects for the uninformed and compliant (the analog to the coefficients in

column 3 in Table 4) grow steadily larger in absolute terms to a point, consistent with our interpretation that these individuals are, indeed, spending more time reading the treatment.[17]

In summary, our results confirm that identifying compliant and uninformed consumers is important when understanding the heterogeneity that lies beneath average treatment effects, as initially reported in Table 1.

We can go further in understanding the importance of distinguishing compliant and uninformed consumers from the non-compliant and informed by drawing on our results that compliant and uninformed consumers can be systematically characterized by their demographic characteristics. Coibion, Gorodnichenko, and Weber (2022) find that gender differences are important; their RCTs suggest that female respondents are more responsive to information treatments in terms of updating their inflation expectations. However, we find (see Appendix Table E2) that women are both more likely to have read (and hence comply with) the treatment and to report being less informed about monetary policy news, on average, and so their stronger reaction to informational treatments could simply follow from the fact that compliant and less informed respondents have higher average treatment effects, and have nothing to do with gender per se.

To explore this further, Table 5 reports gender differences in treatment effects in our sample for the July and September waves. For simplicity, we pool the information treatments to focus on gender differentials. Abstracting from informedness and compliance, column (1) replicates the Coibion, Gorodnichenko, and Weber (2022) findings: females exhibit a greater (absolute) response to the treatments than males. This result still holds in the September wave when we limit the sample to uninformed individuals without controlling for compliance in column (3). After we account for respondent compliance, we find that these gender differences disappear in columns (2) and (4): the coefficient on Treated x Male is no longer statistically significant.

[17] Table G5 in the Appendix shows that these differences in the estimated coefficients in Table 4 between consumers that heard and did not hear news are statistically significant. The effect on the informed is close to zero, especially for Treatment 1. In addition, Table G6 shows that consumers who spend more time reading the treatment tend to have stronger treatment effects than consumers who spend less time reading the treatment, even conditional on compliance.

Table 5: Gender Differences in Treatment Effects

| Heard News = | All | | No | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Model = | Baseline | Compliers | Baseline | Compliers |
| *July* | | | | |
| Treated | -2.17*** | -2.64*** | -1.72*** | -3.41*** |
| | (0.31) | (0.97) | (0.47) | (1.24) |
| Treated x Male | 1.29*** | 1.66 | 0.42 | 1.22 |
| | (0.39) | (1.14) | (0.72) | (1.58) |
| Male | 0.39 | 1.16 | 0.93* | 0.67 |
| | (0.27) | (0.93) | (0.54) | (1.25) |
| Observations | 6018 | 4674 | 3498 | 2680 |
| *September* | | | | |
| Treated | -2.48*** | -1.74* | -2.97*** | -2.73** |
| | (0.36) | (0.92) | (0.68) | (1.16) |
| Treated x Male | 1.59*** | -0.92 | 2.16** | -2.01 |
| | (0.48) | (1.24) | (0.95) | (1.72) |
| Placebo | 0.14 | -0.13 | -0.94 | -0.03 |
| | (0.34) | (1.24) | (0.67) | (1.61) |
| Placebo x Male | 0.67 | -0.09 | 0.70 | -1.37 |
| | (0.46) | (1.55) | (0.88) | (2.22) |
| Male | 0.18 | 2.60*** | 0.70 | 2.34** |
| | (0.34) | (0.73) | (0.64) | (1.10) |
| Observations | 5899 | 4339 | 3343 | 2439 |

Treatments are pooled within each wave, with the placebo in September kept separate. "Baseline" columns are a variation of equation (1), the model used in Tables 1 and 2, while "Compliers" columns are a variation of the propensity score procedure outlined in Section 5 and implemented in Table 4. "Compliers" columns show the results for compliers. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In summary, our results highlight that in order to meaningfully compare information treatment effects across demographic groups it is important to measure who is paying attention, in order to capture the underlying treatment effect. Much attention has been paid in the now-considerable empirical macroeconomics RCT literature to the effectiveness of treatments (e.g., Coibion, Gorodnichenko, and Weber, 2022; D'Acunto et al., 2020; Weber et al., 2024). However, there has been little discussion given to whether the treatments under consideration are, in fact, effective simply because respondents are paying attention to the information provided, or because they were previously uninformed and the information provided is news to them. Differences in average treatment effects across demographic groups can reflect differences in information sets rather than responses to the treatments per se. One of this paper's contributions, therefore, is to

propose a simple way to measure informedness and compliance and to show that, in certain contexts, apparent demographic differences can be explained by informedness and compliance.

## 6. Discussion

Our results highlight the importance of controlling for respondents' information sets when running and interpreting RCTs in empirical macro settings, especially when considering differences in responses across demographic groups that could be proxies for informedness and compliance.

Taken at face value, our results find that monetary policy communications can have a large impact on consumers' inflation expectations. The RCTs find large negative treatment effects for consumers who had not previously heard news about monetary policy and who read the treatment. But the RCTs have little to say about consumers who "swallowed the real-world monetary policy pill"—that is, for consumers who were paying close attention to monetary policy news all along—because the information provided to them was presumably uninformative. To the extent that those consumers responded in a similar fashion when they first heard the news that policymakers would be raising interest rates to combat inflation, FOMC policy actions in 2022 should have worked to directly reduce their inflation expectations as well.[18] Consistent with this belief, on average we find that informed respondents had lower inflation expectations during our sample period than uninformed respondents.[19]

Our results serve as a check on the external validity of several different but related approaches to capturing how consumers respond to monetary policy. Using hypothetical "vignettes," Andre et al. (2022) show that surveys reveal considerable disagreement across consumers in terms of how they update their expectations after an unexpected monetary policy shock. The authors attribute this to people having different subjective models of the economy, rather than to people having different information sets about the state of the economy. In contrast, having differentiated people's information sets, we find that uninformed and compliant consumers do understand the policy objective, as communicating the—for them, unanticipated or unknown—

[18] Because many of the interest rate increases in 2022 were telegraphed in advance, the realization of some of this news would have occurred prior to our study.

[19] However, this relationship need not hold, especially at all points in time. Outside of the laboratory of an RCT, many different information shocks could have been buffeting the informed and uninformed respondents that could explain the time-varying relationship between the series. See Appendix Figure E1.

policy change directly reduces their inflation expectations. In the real-life context of the high inflation of 2022, these consumers appear to use a common subjective model of the economy, at least to the extent that they reduce their medium-term inflation expectations once they are informed that the FOMC has recently increased interest rates. Adding the narrative or vignette for why rates are rising, as administered in our RCTs by comparing Treatment 2 with Treatment 1, does not reduce consumers' inflation expectations further. Consumers appear to already know why the FOMC is raising the federal funds rate. As Weber et al. (2024) discuss, consumers may well react differently to information shocks when the level of inflation is lower. This may help explain why our results differ from those in Andre et al. (2022), who ran their experiments in 2019, a period when inflation was low and stable. Because of the rising inflation through 2022, consumers may well have learned or had a general sense that the intent of the increases in the federal funds rate was to lower inflation. Therefore, even absent communication about the intent of the policy action, consumers still understood the FOMC's intended monetary policy transmission mechanism.

On the surface, our results also appear to contrast with some of those in Roth, Wiederholt, and Wolfhart (2023). Presenting alternative hypothetical scenarios to consumers about what might happen to the funds rate at the March 2022 and September 2022 FOMC meetings, they find that a higher hypothetical funds rate was associated with a higher expected path for inflation.[20] In our surveys, we find that communicating ex post about an actual increase in the funds rate resulted in lower medium-term inflation expectations, on average but especially for the compliant and previously uninformed. While there are a number of differences between the exercises in terms of approach and key details, we highlight the role of information. In Roth, Wiederholt, and Wolfhart (2023), it is not clear which of the hypothetical scenarios is more closely aligned with respondents' priors. Based on the situation at the time—with high inflation and telegraphed plans for higher policy rates—presenting a baseline hypothetical scenario of an unchanged funds rate at the upcoming meeting was likely a surprise to relatively informed consumers. This, in turn, could have led them to reassess and to lower their inflation expectations as they reverse engineered why the funds rate would not rise as they thought it would. Our focus on the treatment effects from

[20] By contrast, in a survey conducted in 2021 that used different hypothetical scenarios, they found a negative relationship between the expected path for the fed funds rate and inflation, suggesting a mechanism qualitatively consistent with our findings.

providing information to the previously uninformed is meant to guard against such reverse causality.

To probe more deeply at the mechanisms—that is, the underlying subjective models—through which consumers thought that increases in the policy rate would reduce inflation, we summarize results testing whether consumers update their prior expectations for GDP growth and personal income in light of the information treatments as they do their inflation expectations. This additional exercise is possible given that our survey asks consumers questions about these two variables similar to those it asks about inflation.[21] In summary, as shown in the tables in Appendix F, we find that even when focusing on the previously uninformed and compliant respondents, consumers' posterior expectations for GDP growth and their personal income on average are unaffected by the treatments. In other words, even uninformed and compliant consumers do not consistently associate an increase in interest rates that induces a decline in inflation with either a worsening real macroeconomy or deteriorating personal economic prospects. This implies that, if consumers think the slope of the Phillips curve is flat, they do not expect the increases in the federal funds rate to move the economy along a flat Phillips curve. Instead, monetary policy shifts inflation expectations directly, similar to the shifts seen in Hazell et al. (2022), who estimate a flat Phillips curve slope using state-level data. However, our results admittedly cannot rule out the possibility that consumers alternatively believe that there is a very steep Phillips curve, and hence there is little cost to disinflation, or that there is some mixture of these types of consumers among whom our analysis cannot differentiate. While future work should aim to disentangle these two possibilities, our results suggest that consumers think the Fed's disinflationary efforts will have little cost in terms of lost output.

While our treatments and analysis focus on the communication of monetary policy decisions, consumers may be more aware of other, longer-term interest rates—which potentially incorporate forward-looking expectations of monetary policy—because these rates are more relevant for the real-world monetary transmission mechanism for borrowers and savers. In Appendix G, we re-estimate the treatment effects in Table 4 while controlling for two additional measures of broader awareness about interest rates that our survey asks about pre-treatment:

[21] These questions were asked in exactly the same manner as the questions eliciting inflation expectations: a question prior to the treatment first asked whether the respondent expected GDP/personal income to increase or decrease over the next year, and a follow-up posterior question after the treatment asked for the expected percent change (see Section 2 or Appendix H).

whether the respondent reported that the interest rates that people pay to borrow money in general had changed recently; and whether the respondent indicated that borrowing rates in general had gone up by more than 1 percentage point recently.[22] As in Table 4, we continue to find negative and statistically significant treatment effects for all respondents who complied with the treatment and for those respondents who complied and were uninformed about monetary policy. We find no evidence that prior awareness about borrowing rates attenuates the treatment effects, as seen by the statistically insignificant interaction terms. Thus, our results are not dependent on providing information to consumers who were completely inattentive to the economic environment; the monetary policy treatments also moved the inflation expectations of those who already had some sense of what was happening to interest rates. These results provide suggestive evidence that it is the combination of communicating about the *actor* (that is, the FOMC) and the *action* (that is, the decision to raise the policy rate) in our treatments that has delivered the large negative impact on inflation expectations, and not a change in interest rates alone. Future work can refine the role of these two margins further.

Our results highlight the importance of controlling for both informedness and compliance so that RCTs have a higher degree of external validity and can better help economists understand the real-world effects of monetary policy communications. Informedness is not the same for everybody (see Appendix Table E2), with some agents accruing information about policy ahead of others in the "real-world" rather than in the "laboratory" of the RCT. Informed consumers update their inflation expectations outside of the RCT. This means that the small and statistically insignificant treatments effects we find for them (Table 4, columns (2) and (5)) do not mean that the monetary policy communication has been ineffective, just that this group of consumers may have already updated their expectations. The uninformed do have an understanding of the main mechanism of monetary policy; so new, targeted communications to reach this group in the real-world should also be expected to weaken their responsiveness to treatments within the RCT.

[22] Appendix G documents that there is a high correlation between informedness about monetary policy and awareness of increases in borrowing rates more generally, but not a perfect correlation between the two different concepts, which forms the basis for these regression results.

## 7. Conclusion

This paper reports and analyzes results from a specially designed multi-wave RCT to test whether and how communications about actual increases in the federal funds rate in 2022 causally affected consumers' inflation expectations. We find that simple communications about increases in the federal funds rate reduced consumers' medium-term inflation expectations, most notably for those who were previously unaware of but willing to pay attention to (that is, in our language, "compliant" with) the information communicated in the RCT. Our results thus highlight the importance of measuring the ex ante informedness and the ex post compliance of respondents in applied macroeconomics RCTs.

Our results therefore provide suggestive evidence that the FOMC's policy actions of 2022 likely helped to directly (re-)anchor medium-term inflation expectations and contribute to the disinflation process for some consumers. In turn, the fact that real-time and real-life monetary policy communications of the policy action alone—even absent any information on its intent—are found to lower longer-term inflation expectations suggests that consumers did have some common understanding of the monetary policy transmission mechanism and of the FOMC's intention that the rate increases should drive inflation down.

Our finding that communicating the monetary policy changes of 2022 had small effects on inflation expectations on average, but larger effects on the previously uninformed and compliant, reinforces the growing consensus in the literature (e.g., Coibion and Gorodnichenko, 2015; Andre et al., 2022) that there is considerable heterogeneity across consumers in terms of how they form and update their inflation expectations. During the high inflation of 2022, we find that not everyone was paying attention to monetary policy, nor was everyone equally attentive to the information treatment administered in our five RCTs, even when presented with very short, tweet-like monetary policy communications designed to be easy to read. By identifying groups of consumers who tend to report being less informed about monetary policy news and to be more likely to pay attention to news when it is shared with them, and by controlling for those groups in our regressions, we find evidence that there is scope to increase the impact of monetary policy communications by targeting specific groups of the general public, notably women. More generally, our results suggest that it is important when interpreting the heterogeneous treatment effects commonly found in RCTs in macroeconomics to unpack both the compliant from the non-compliant and those for whom the informational treatment is news from those for whom the

information is already known. Otherwise, as we illustrate, it is possible to misclassify the drivers of the heterogeneity and misdiagnose the best ways of increasing the effectiveness of monetary policy communications.

**References**


Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart. 2022. "Subjective Models of the Macroeconomy: Evidence from Experts and Representative Samples." *Review of Economic Studies* 89 (6): 2958-91. https://doi.org/10.1093/restud/rdac008.

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91: 444–55. https://doi.org/10.1080/01621459.1996.10476902.

Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2016. "The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment." *Review of Economics and Statistics* 98 (3): 503–23. https://doi.org/10.1162/REST_a_00499.

Bholat, David, Nida Broughton, Janna Ter Meer, and Eryk Walczak. 2019. "Enhancing Central Bank Communications Using Simple and Relatable Information." *Journal of Monetary Economics* 108: 1-15. https://doi.org/10.1016/j.jmoneco.2019.08.007.

Binder, Carola. 2017. "Fed Speak on Main Street: Central Bank Communication and Household Expectations." *Journal of Macroeconomics* 52 (C): 238-51. https://doi.org/10.1016/j.jmacro.2017.05.003.

Binder, Carola and Alex Rodrigue. 2018. "Household Informedness and Long-Run Inflation Expectations: Experimental Evidence." *Southern Economic Journal* 85: 580-98. https://doi.org/10.1002/soej.12306.

Blinder, Alan. S., Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. 2008. "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence." *Journal of Economic Literature* 46 (4): 910–45. https://doi.org/10.1257/jel.46.4.910.

Blinder, Alan S., Michael Ehrmann, Jakob de Haan, and David-Jan Jansen. 2024. "Central Bank Communication with the General Public: Promise or False Hope?" *Journal of Economic Literature* 62(2): 425-57. https://doi.org/10.1257/jel.20231683.

Bracha, Anat and Jenny Tang. 2024. “Inflation Levels and (In)Attention.” *Review of Economic Studies*. https:// doi.org/10.1093/restud/rdae063.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts.” *American Economic Review* 105 (8): 2644-78. https://doi.org/10.1257/aer.20110306.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. 2018. “How Do Firms Form Their Expectations? New Survey Evidence.” *American Economic Review* 108 (9): 2671–2713. https://doi.org/10.1257/aer.20151299.

Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Mathieu Pedemonte. 2020. “Inflation expectations as a policy tool?” *Journal of International Economics*, 124: 103297. https://doi.org/10.1016/j.jinteco.2020.103297.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Michael Weber. 2023a. “Forward Guidance and Household Expectations.” *Journal of the European Economic Association*. 21 (5): 2131-71. https://doi.org/10.1093/jeea/jvad003.

Coibion, Olivier, Yuriy Gorodnichenko, Edward S. Knotek II, and Raphael Schoenle. 2023b. “Average Inflation Targeting and Household Expectations.” *Journal of Political Economy Macroeconomics* 1 (2): 403-46. https://doi.org/10.1086/722962.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2022. “Monetary Policy Communications and Their Effects on Household Inflation Expectations.” *Journal of Political Economy*, 130 (6): 1537-84. https://doi.org/10.1086/718982.

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber. 2020. “Effective Policy Communication: Targets versus Instruments.” Working Paper 2020-148. University of Chicago Becker Friedman Institute for Economics. https://doi.org/10.2139/ssrn.3712658.

D’Acunto, Francesco, Andreas Fuster, and Michael Weber. 2021. “Diverse Policy Committees Can Reach Underrepresented Groups.” Working Paper 29275. National Bureau of Economic Research. https://doi.org/10.3386/w29275.

Dietrich, Alexander M., Edward S. Knotek II, Kristian Ove R. Myrseth, Robert W. Rich, Raphael S. Schoenle, and Michael Weber. 2022. “Greater Than the Sum of the Parts:

Aggregate vs. Aggregated Inflation Expectations." Federal Reserve Bank of Cleveland, Working Paper No. 22-20. https://doi.org/10.26509/frbc-wp-202220.

Follmann, Dean A. 2000. "On the Effect of Treatment Among Would-Be Treatment Compliers: An Analysis of the Multiple Risk Factor Intervention Trial." *Journal of the American Statistical Association* 95: 1101–09. https://doi.org/10.1080/01621459.2000.10474306.

Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar. 2022. "Expectations with Endogenous Information Acquisition: An Experimental Investigation." *The Review of Economics and Statistics* 104 (5): 1059–1078. https://doi.org/10.1162/rest_a_00994

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera. 2023. "The Voice of Monetary Policy." *American Economic Review* 113 (2): 548-84. https://doi.org/10.1257/aer.20220129.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart. 2023. "Designing Information Provision Experiments." *Journal of Economic Literature* 61 (1): 3-40. https://doi.org/10.1257/jel.20211658.

Haldane, Andrew and Michael McMahon. 2018. "Central Bank Communications and the General Public." *AEA Papers and Proceedings* 108: 578-83. https://doi.org/10.1257/pandp.20181082.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson. 2022. "The Slope of the Phillips Curve: Evidence from U.S. States." *The Quarterly Journal of Economics* 137 (3): 1299–1344. https://doi.org/10.1093/qje/qjac010.

Hoffmann, Mathias, Emanuel Moench, and Guido Schultefrankenfeld. 2023. "A KISS for Central Bank Communication in Times of High Inflation." *mimeo*

Imbens, Guido W. and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62 (2): 467–75. https://doi.org/10.2307/2951620.

Jo, Booil and Elizabeth A. Stuart. 2009. "On the Use of Propensity Scores in Principal Causal Effect Estimation." *Statistics in Medicine* 28 (23): 2857-75. https://doi.org/10.1002/sim.3669.

Knotek, Edward S., II, Raphael S. Schoenle, Alexander M. Dietrich, Keith Kuester, Gernot J. Müller, Kristian Ove R. Myrseth, and Michael Weber. 2020. "Consumers and COVID-19: A

Real-Time Survey." *Economic Commentary* 2020-08. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-ec-202008.

Lamla, Michael J. and Dmitri V. Vinogradov. 2019. "Central Bank Announcements: Big News for Little People?" *Journal of Monetary Economics* 108: 21-38. https://doi.org/10.1016/j.jmoneco.2019.08.014.

Maćkowiak, Bartosz and Mirko Wiederholt. 2024. "Rational Inattention during an RCT." https://cepr.org/publications/dp18875.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. 2023. "Rational Inattention: A Review." *Journal of Economic Literature* 61 (1): 226-73. https://doi.org/10.1257/jel.20211524.

Nakamura, Emi and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3): 1283–1330. https://doi.org/10.1093/qje/qjy004.

Pedemonte, Mathieu. 2024. "Fireside Chats: Communication and Consumers' Expectations in the Great Depression." *Review of Economics and Statistics* 106(4): 997–1011. https://doi.org/10.1162/rest_a_01200.

Pfäuti, Oliver. 2024. "The Inflation Attention Threshold and Inflation Surges." Unpublished manuscript.

Qualtrics. 2019. "ESOMAR 28: 28 Questions to Help Buyers of Online Samples." April 2019, Technical Report. https://www.iup.edu/arl/files/qualtrics/esomar.pdf.

Roth, Christopher and Johannes Wohlfart. 2020. "How Do Expectations about the Macroeconomy Affect Personal Expectations and Behavior?" *Review of Economics and Statistics* 102 (4): 731-48. https://doi.org/10.1162/rest_a_00867.

Roth, Christopher, Mirko Wiederholt, and Johannes Wohlfart. 2023. "The Effects of Monetary Policy: Theory with Measured Expectations." CESifo Working Paper No. 10216. Available at SSRN: https://doi.org/10.2139/ssrn.4329475.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665-90. https://doi.org/10.1016/S0304-3932(03)00029-1.

Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent H. Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion, Geoff Kenny, and Jorge Ponce. 2024. “Tell Me Something I Don’t Already Know: Learning in Low and High-Inflation Settings.” Working Paper 31485. National Bureau of Economic Research. https://doi.org/10.3386/w31485.