THE CAUSAL EFFECTS OF INFLATION EXPECTATIONS ON HOUSEHOLDS' BELIEFS AND ACTIONS

Olivier Coibion
UT Austin and NBER

Yuriy Gorodnichenko
UC Berkeley and NBER

This draft: October 9, 2024

Abstract: We discuss how Randomized Control Trials (RCTs) can be used to study the causal effects of inflation expectations on the decisions of households. RCTs create exogenous variation in the inflation expectations of survey participants. When linked to either external information on their actions or subsequent survey waves that measure their ex-post decisions, this can provide clear causal identification of expectations on decisions. We review recent evidence using this strategy and discuss potential challenges associated with this approach.

JEL: E3, E4

Keywords: inflation expectations, surveys, consumption.

1. Introduction

If everyone in the economy expects higher inflation in the future, how will this affect their decisions? The answer, as illustrated in Werning (2022), is not as simple as looking at aggregate structural relationships like the New Keynesian Phillips curve in the case of prices or the Euler equation in the case of consumption. In these relationships, expectations already embody a wide range of restrictions that that make the thought exercise ill-defined. Instead, Werning (2022) suggests that we need to start answering the question by first focusing on how exogenous changes in inflation expectations at the individual level affect their decisions holding everything else constant. In this paper, we review new papers that tackle this question directly in the case of households by utilizing RCT methods to create exogenous variation in inflation expectations of individual households.

Identifying the passthrough of expectations into decisions is not just relevant for policymakers: it is at the heart of every macroeconomic model in which most decisions are forward looking and depend on beliefs about the future. Despite this, it has been challenging for researchers to actually quantify the extent to which expectations affect decisions as posited by theoretical models. This reflects the fact that taking up this challenge requires measures of people's beliefs as well as a source of exogenous variation in those beliefs. While the increasing availability of survey data has made the former less of a constraint, it is not enough to measure expectations and decisions. A positive correlation between consumption and inflation expectations could reflect the intertemporal channel through which inflation expectations are predicted to affect consumer spending, but causation could also run in the opposite direction if, when consumers are spending more, they think others will do the same and the increase in demand will cause prices to rise more rapidly. The only way to break through this identification challenge is to find variation in inflation expectations that is unrelated to these other channels.

Randomized control trials (RCTs) provide one way to do so. The idea is to provide information to randomly selected survey participants that leads them to revise their inflation expectations relative to other households who do not receive this information. But because the choice of who to provide the information to is randomized, the selection of those who revise their beliefs is as well. Therefore, if those individuals who changed their beliefs due to this treatment also tend to change their subsequent spending in a clear way relative to the untreated, one can label the change in spending as being caused by the change in inflation expectations. In this paper, we review recent papers that implement this approach to study how households' inflation expectations affect their decisions and discuss potential challenges that arise in the implementation of this RCT approach.

One paper that first used this strategy in the context of inflation expectations is Coibion, Gorodnichenko and Weber (2022). Using information about inflation that was provided to randomly selected U.S. households, these authors were successful in creating large exogenous changes in the inflation expectations of randomly selected participants and link those changes to their subsequent spending decisions. They found that, consistent with the intertemporal substitution channel, higher inflation expectations led households to persistently and significantly raise their spending on non-durables in subsequent months. However, the effects on durables goods purchases seemed to go in the opposite direction, a feature they speculated had to do with how households changed other expectations along with inflation.

A subsequent and recent paper by Georgarakos et al. (2024) tackles this possibility by separating changes in inflation expectations from changes in inflation uncertainty. Because high inflation tends to be more volatile inflation, it would be natural for households who expect higher inflation to also be more uncertain about future inflation, a feature that Georgarakos et al. (2024) confirm for European households. Higher uncertainty could induce households to postpone large durable goods purchases, as in Bloom (2009), so higher inflation expectations, to the extent that they are associated with higher inflation uncertainty, could have ambiguous overall effects on durable goods purchases. Through the use of multiple information treatments that generate separate variation in inflation expectations and inflation uncertainty, Georgarakos et al. (2024) show that this is precisely what happens: higher inflation uncertainty (holding expectations constant) tends to reduce the likelihood of subsequent durable goods purchases whereas higher inflation expectations (holding uncertainty constant) tends to increase it. The combined effect is dominated by the uncertainty channel, which explains why Coibion, Gorodnichenko and Weber (2022) found negative overall effects of inflation expectations on durable goods purchases.

In addition to consumer spending, Georgarakos et al. (2024) are able to assess how inflation expectations and uncertainty affect other margins of adjustment available to households. They report three key findings. First, higher inflation uncertainty leads households to adjust their financial portfolios away from (illiquid) retirement accounts and toward (very liquid) checking and savings accounts, while higher inflation expectations lead to the opposite behavior. Second, when households become more uncertain about future inflation, they expect to search more actively for work in coming months, whereas higher inflation expectations lead to reductions in job search intensity. Third, these changes in job search behavior are reflected in labor market outcomes, especially in the case of inflation uncertainty. As households become more uncertain about future inflation and search for work more actively, their probability of being unemployed declines in subsequent months as does their probability of being in part-time work, even as their likelihood of being out of the labor force declines. In short, higher inflation uncertainty leads to more workers (both employed and unemployed) moving into full-time employment.

Together, these results provide decisive causal evidence that consumers' beliefs about future inflation shape their decisions along several different margins. We view this as providing support for the use of RCTs in studying how expectations affect decisions. However, this approach faces a number of challenges to its widespread use. We review these challenges in the paper and discuss how they can potentially be overcome. Doing so would allow macroeconomists to answer many new questions and potentially validate and quantify some of the most fundamental mechanisms at work in our models.

Furthermore, these results speak to important monetary policy discussions, particularly about the use of communication to shape the expectations of households and firms. One potential implication is that policy communications, even if they are successful in changing expectations such as about inflation, may have effects that differ from those that might be expected from theoretical models if they induce indirect effects on decisions, such as by changing uncertainty. Similarly, policy pronouncements that emphasize the "data dependence" of future policy may help avoid tying the hands of policymakers but they may have negative effects on economic activity by increasing uncertainty about future inflation outcomes.

The paper is organized as follows. Section 2 describes the RCT approach and why it can address fundamental identification issues involving expectations and decisions of economic agents. Section 3 discusses recent evidence from RCTs on the effects of inflation expectations on households' decisions. Section 4 considers potential challenges facing RCTs in the future and potential limits to how they can be used, as well as some possible solutions to these challenges. Section 5 concludes by discussing some policy implications.

2. The RCT Strategy to Identifying the Causal Effects of Expectations on Decisions

How do we identify the effect on inflation expectations on the decisions of economic agents? In this section, we review some of the main challenges as well as how RCTs can help break through this identification challenge.

2.1 Challenges to Identification

Expectations of the future are omnipresent in macroeconomic models and the decisions that are embodied in them. Yet characterizing how expectations actually affect the decisions of agents is challenging. There are two major issues. First, the measurement of expectations. Second, the endogeneity of expectations.

The most direct approach to measuring expectations is to directly ask economic agents what they expect about the future. This is now commonly done in surveys of households, firms, and financial market participants. But in practice, many issues arise. First, surveys need to be representative and not based on convenience samples. This has proven to be a major challenge for surveys of firms, where getting randomly selected top executives of large firms to participate is often difficult. Second, the specifics of how expectations are measured can matter for the resulting measures of beliefs. Whether questions are formulated as point forecasts or distributions or whether questions are formulated in terms of prices or inflation, for example, can affect the answers provided by respondents. See De Bruin et al. (2011) and Armantier et al. (2013) for an extensive discussion of these points. Different surveys often take different approaches, occasionally leading to diverging patterns for measured expectations.

The second major challenge in identifying how expectations affect decisions is endogeneity: beliefs are not formed in a vacuum and are themselves a function of the economic environment. Suppose for example that we observe in the data that households who tend to expect higher inflation in the future also expect to consume more in the next month. One potential reason for this correlation could be the intertemporal substitution motive of the consumption Euler equation: if prices are expected to rise more rapidly in the future, households have an incentive to purchase more before those price increases take place. But causality could also run in the opposite direction. If households expect to increase their consumption (e.g. because their income is rising), then they might expect other households to do the same which, by increasing aggregate demand, could raise inflation in the future. So it is not enough to simply be able to measure the expectations of agents, one also needs to be able to identify exogenous variation in those beliefs to assess how changes in expectations pass through into decisions. While finding this type of variation in historical data has been very challenging, RCTs can provide a clear source of exogenous variation that allows for causal identification.

2.2 The RCT Approach

In principle, the RCT strategy is simple and follows three steps, presented visually in Figure 1. First, one needs to measure the initial expectations of economic agents through a survey. These provide a measure of the "prior" beliefs of all survey participants. Second, a piece of information is provided to a randomly selected subset of survey participants. This is the "information treatment" step, which serves to create the exogenous variation in beliefs of some agents. It is exogenous since participants are randomly assigned to either the control group (that receives no information) or the treatment group (that receives the information), making the subsequent change in beliefs orthogonal to all of the characteristics of participants. Haaland et al. (2023) provide an extensive review of how information treatments can be applied in surveys. The third step is to measure the ex-post beliefs and decisions of respondents. Measuring ex-post beliefs helps the researcher identify how well the information treatment worked in terms of creating exogenous variation in expectations of agents. Measuring ex-post decisions is essential to being able to characterize whether those changes in beliefs then impacted the economic decisions of the survey participants. Decisions can be measured using external sources of information on the economic actions of the agents (when such data can be matched to the survey) or through self-reported decisions of the respondents measured through subsequent waves. In the absence of such measures, one can also use planned decisions measured in the same survey as the information treatment but asked after the treatment is applied. Roth and Wohlfart (2020) provide an excellent illustration of how this can be done in the context of how beliefs about the probability of a future recession affect planned household spending.

Conceptually, this approach is therefore very similar to the use of quasi-experimental variation in applied work (see Angrist and Pischke 2009). But whereas the latter relies on historical episodes which provide variation that is similar to an experiment, the RCT approach directly *creates* the exogeneity in beliefs via the information provided to randomly selected participants in the survey. As we describe in the next section, RCTs can sometimes serve as very powerful sources of variation in beliefs. However, in practice there are many challenges that can arise with this empirical strategy, and we will discuss some of these challenges and how they can potentially be addressed.

3. Recent Causal Evidence on the Effects of Inflation Expectations on Household Decisions

In this section, we review two papers that applied this RCT strategy to study how inflation expectations of households affect their economic decisions.

3.1 Inflation Expectations and Consumer Spending

The first paper is Coibion, Gorodnichenko and Weber (2022, CGW henceforth). This paper uses households that participate in the Nielsen Homescan panel, which requires them to track their individual retail purchases. This data is commonly used to study household spending decisions using the scanner data. However, Nielsen also allows researchers to implement surveys of households, which can then be linked to their subsequent spending decisions, and the surveys can include randomized information treatments, making this an ideal setting to utilize the RCT approach. Furthermore, because the size of the survey is very large (~15,000-25,000 respondents per survey wave), it allows for large numbers of participants in each treatment arm.

In 2018, CGW ran a survey of Nielsen households which applied the strategy described above. First, survey participants were asked to assign probabilities to different possible inflation outcomes, providing a measure of their initial inflation expectations. Then, survey participants were randomly assigned to either the control group or one of multiple information treatment groups. Three of these treatments are particularly relevant for us. One treatment group was provided with the most recent inflation statistic (close to 2%). A second treatment group was provided with the most recent SPF or FOMC 12-month ahead inflation forecast (again close to 2%). The third group was provided with the Federal Reserve's inflation target of 2%. Subsequently, survey participants were asked to provide a point forecast for inflation over the next 12 months, which provided a measure of their "posterior" beliefs.

The information treatments proved to be quite powerful in changing the inflation expectations of survey participants. One way to see this is visually in Figure 1. This figure presents a binscatter plot of households' prior inflation expectations against their posterior inflation expectations, broken down by each treatment arm. Consider first the control group. For this group, posteriors and priors are closely linked, and the slope of the regression line linking the two is close to one. This is what one would expect since these participants were not provided with any information, so their posterior beliefs should be the same as their prior beliefs. However, because two different question formulations were used to measure prior and posterior beliefs (to minimize survey fatigue), the slope of the line is somewhat less than one from attenuation bias.

For the three treatment groups on the other hand, the regression line linking posteriors and priors is very flat: respondents' posterior beliefs moved strongly in the direction of the provided information and away from their prior beliefs. This indicates that they placed a lot of weight on the new information. Had they chosen to dismiss the information, their posteriors would be close to their priors, as was the case for the control group. They could have dismissed the information if, for example, they already knew it or if they viewed it as non-credible or irrelevant. But instead, they responded very strongly to the information, leading to strong revisions in expectations toward the provided signal. Thus, the information treatment can be viewed as having been very successful in generating exogenous variation in the inflation expectations of survey participants.

To quantify the effect of inflation expectations on household spending, CGW ran the following regression:

$$\log(spend)_{i,t+h} = \beta E_i^{post}\pi + \gamma E_i^{prior}\pi + \kappa \log(spend)_{it} + Controls_{it} + error_{i,t+h}$$

where ex-post spending is measured using Nielsen scanner-level reported spending of households on non-durable goods, $E_i^{post}\pi$ is the posterior inflation expectations of households and $E_i^{prior}\pi$ is the prior inflation expectations. To identify the exogenous variation in inflation expectations, CGW use an IV strategy in which the first stage is given by:

$$E_i^{post}\pi = a + \sum_j b_j \times Treat_{i,j} + \sum_j \gamma_j \times Treat_{i,j} \times E_i^{pre}\pi + \psi \times E_i^{pre}\pi + error$$

where $Treat_{i,j}$ is an indicator variable equal to one if household i belongs to treatment group j. Since this specification includes the interaction of the treatment indicator with households' prior inflation expectations, it is effectively reproducing the regressions presented visually in Figure 1.

We present the results of this regression from CGW in Table 1 for spending levels 3 months and 6 months after the information treatment. Note first that the F-statistics for the first stage are well above 100, indicating again that the information treatments provided a powerful source of exogenous variation in inflation expectations. Second, the estimated coefficient on inflation expectations is close to one, indicating that a one percentage point increase in inflation expectations is followed by a close to one percent increase in consumption over the next 3-6 months by U.S. households. Hence, this presents clear evidence of a large and persistent causal relationship running from the inflation expectations of households to their spending decisions.

CGW document two other notable results. First, when they use self-reported spending outcomes from subsequent survey waves rather than the Nielsen measures of spending, the estimates are much noisier, albeit still significantly different from zero. This is because self-reported spending measures incorporate recall error, rounding, and other properties of self-reported data that introduce noise. As a result, it is more difficult to precisely estimate the effects of expectations on decisions with these self-reported spending measures. This illustrates the importance of being able to match the surveys with external sources of information on decisions, when possible.

Second, CGW find that when inflation expectations of households rise, they became less likely to engage in purchases of large durable goods in subsequent periods, a finding at odds with the intertemporal channel typically emphasized for inflation expectations. CGW speculate that this could reflect the fact that other expectations of households could be changing as well. For example, since higher inflation tends to be more volatile inflation, it could be that respondents who raise their inflation expectations also tend to increase their uncertainty about inflation. Because uncertainty tends to lead to "wait and see" effects on durable goods purchases (see e.g. Coibion et al. 2024), it could be that the response of durable goods purchases confounds these two different expectations channels. Because CGW did not measure inflation uncertainty after the treatments, they could not assess whether this channel was indeed at work. In a similar experiment applied to households in the Netherlands, Coibion et al. (2023) similarly find a negative effect of inflation expectations on consumer durables purchases. The next paper we discuss tackles this question directly and tries to separate out the effects of inflation expectations and uncertainty on household decisions.

3.2 Inflation Expectations, Uncertainty and Household Decisions

It is a well-known empirical pattern that high inflation is volatile inflation. So changes in inflation expectations could well be associated with changes in inflation uncertainty. A recent paper by

Georgarakos et al. (2024, GGCK henceforth) addresses this joint dynamic of inflation expectations and uncertainty directly and tries to separately identify their effects on the decisions of households.

To do so, GCCK use the European Central Bank's Consumer Expectations Survey (CES), a monthly survey of around 19,000 households in 11 European countries. In September 2023, GGCK added a special set of questions to the regular CES survey. One question asked respondents about what the lowest and highest inflation rates they considered likely were over the next 12 months, which provided an initial measure of inflation uncertainty as well as a measure of their average inflation expectation. Following CGW, they then implemented an information treatment, which was followed by another question that measured posterior inflation expectations and uncertainty. Specifically, following Altig et al. (2022), respondents were asked to provide three scenarios for inflation (low, medium and high) then assign probabilities to each. With these questions, GGCK could assess the effects of information treatments on both inflation expectations and uncertainty. As shown in Figure 2, GGCK confirmed the intuition that respondents who expected higher inflation also tended to be more uncertain about their inflation forecast, with a strong positive correlation between the two being visible.

The information treatments were designed to generate different relative variation in the first and second moments of respondents' inflation expectations. To do so, GGCK randomly assigned participants either to a control group or one of three treatment groups. One treatment group was told about the average inflation forecast of professional forecasters. The second treatment group was instead told about the difference in inflation expectations between the most optimistic and the most pessimistic professional forecasters. The third group was provided with both pieces of information. The objective, following Coibion et al. (2024), was to generate separate variation in the first and second moments of households' beliefs about inflation, which would be necessary to identify their separate effects on household decisions.

Through follow-up waves of the survey, GGCK are able to assess how inflation expectations and uncertainty affect household decisions along a number of different margins. Their first set of results focus on household durable goods purchases, following CGW. The results for how inflation expectations and inflation uncertainty affect ex-post durable goods purchases from GGCK are shown in Table 2, with the estimation done in a similar manner as CGW, but now including both inflation expectations and uncertainty as RHS variables and instrumenting for both using the information treatments. Panel A shows the main results of these regressions for different categories of durable goods purchases, one month after the information treatment. They find a clear pattern: high inflation uncertainty is followed by a lower probability of households purchasing durable goods whereas higher inflation expectations are followed by a greater probability. The latter is the opposite of the finding in CGW and suggests that the strong positive correlation between expectations and uncertainty was responsible for the negative relationship between expectations and durable goods spending that they had found. Indeed, when GGCK only include inflation expectations as a right-hand side variable (Panel B of Table 2), they similarly find that inflation expectations appear to affect durable goods purchases negatively. But from Panel A, it is clear that this is because there are two effects are work. First, the direct effect of higher inflation expectations is to move durable goods purchases forward in time. Second, the indirect effect is that higher inflation expectations lead to more uncertainty about inflation, which tend to reduce durable goods purchases. The results in Panel B indicate that the indirect effect is stronger than the direct effect.

Panel C focuses on the importance of the RCT for identification purposes. Specifically, it presents estimates of the same specification as in Panel A but using OLS instead of IV, in other words without taking advantage of the exogenous variation created by the treatments. In this case, we see that the estimated coefficients are small and generally insignificant. One cannot clearly identify the effect of either inflation expectations or uncertainty on the durable goods purchases of households. This illustrates the key role played by the information treatment in generating the exogenous variation in beliefs that is necessary to identify their effects on decisions. In short, without the RCT, there is no identification.

With respect to non-durable goods purchases, GGCK do not have access to external measures of spending as was the case for CGW and instead must rely on self-reported measures of ex-post spending. They find that the latter are too noisy to yield precise estimates of the effects of either inflation expectations or uncertainty on household purchases of non-durables and services. Given that one would expect that both the direct and indirect effects of inflation expectations were at work in the estimates of CGW for non-durables as well as durables, this suggests that the estimated effects found by CGW for non-durables should be thought of as a lower bound for the direct effect of inflation expectations on this type of spending.

In addition to spending decisions, GGCK were able to quantify other decisions taken by survey participants in the months following the information treatment, allowing them to characterize the effects of inflation expectations on different margins of adjustment available to households. One such margin is the composition of their financial portfolio. GGCK rely on two outcomes. One is a hypothetical question asking respondents how they would allocate an income windfall across different types of financial assets, immediately after the information treatment. The second outcome comes from the fact that respondents were asked to provide a decomposition of their financial wealth in terms of different asset types two months after the treatment.

Table 3 therefore presents the estimated effects of inflation expectations and inflation uncertainty on the share of assets that households would either allocate a financial windfall to (Panel A) or on their actual financial portfolio (Panel B). In both cases, the main result is that higher inflation uncertainty leads to a reallocation of the portfolio away from retirement assets (which are highly illiquid) and toward checking/savings accounts (which are highly liquid), whereas higher inflation expectations lead to the opposite pattern. Hence, both inflation expectations and uncertainty lead to clear and economically large portfolio reallocations by households that take place fairly rapidly.

Another margin of adjustment available to households is their labor supply and job search decisions. Because the CES asks questions about how intensively respondents search for work and track their employment status across waves, GGCK are also able to assess the effects of inflation expectations and uncertainty on these decisions in outcomes. The effects on job search are presented in Table 4 for different horizons. The main finding is that higher uncertainty about inflation leads respondents to report that they plan to search more actively for work, whereas higher inflation expectations do the reverse. Consistent with this search behavior, unemployed respondents who become more uncertain about inflation raise the probability they will have a job in 3 months whereas unemployed respondents with higher inflation expectations view this outcome as less likely. For working respondents, those who become more uncertain about inflation view it as more likely that they will be searching for a new job over the next 3 months but not because they are more likely to be fired, which indicates that it will be a conscious decision to search harder while on the job, whereas the reverse is true when inflation expectations rise.

Does this changing search behavior affect employment outcomes? Table 5 presents results for how inflation expectations and uncertainty affect subsequent employment outcomes of respondents. GGCK find that when respondents become more uncertain about inflation, they become less likely to be either unemployed or working part-time in subsequent months, but not through movements out of the labor force. Consistent with higher job search, the results instead suggest that high uncertainty about inflation induces workers to take on more work or move into full-time work. In contrast, there is little clear effect of higher inflation expectations on subsequent job outcomes.

Together, these results indicate that inflation expectations matter for the decisions of households, but that it is important to separate out the direct effects of these expectations from their indirect effects that operate through inflation uncertainty. Note that for policymaking, the total effect may be the more relevant statistic: if communication raises inflation expectations and uncertainty jointly, then the total effect will provide a good approximation to the likely effects on households' decisions. But for modeling purposes or whenever the channels can be separated via targeted communication, knowing the difference between the direct and indirect effects will be important. For example, a policy communication which raises inflation expectations while simultaneously reducing inflation uncertainty would be predicted to have a much larger effect on household decisions than one that did not control for the endogenous response of uncertainty to inflation expectations.

4. Other Challenges to RCT Identification

The potential need to separately identify the direct and indirect effect of a change in expectations arising from an RCT is one issue that can make the inference from this type of approach challenging. But it is not the only one. In this section, we briefly review a number of other challenges that can arise, as well as potential solutions to them.

Measurement of Expectations and Survey Implementation

As already discussed, surveys are subject to a wide range of issues, from ensuring their representativeness to how questions are asked, in what order, etc., all of which can matter for the results. Given this, every concern that applies to surveys will be relevant for the RCT strategy that relies on surveys to measure expectations. However, because RCTs entail a comparison between treatment and control groups, any issue that affects the two groups equally (e.g. a bias induced by a question formulation) will effectively be differenced out. To the extent that one is also examining the change in expectations of one individual over time, this will also take out individual fixed effects. The result is that the RCT, by effectively delivering a difference-in-difference setting, mitigates some of the issues that are unavoidable in surveys. As a result, while some of the concerns commonly associated with surveys are still relevant, many are much less relevant in this context.

Where Can You Run an RCT?

Another concern is that there are few settings in which this strategy can be implemented. Many of the existing surveys (e.g. Michigan Survey of Consumers) do not allow researchers to introduce information experiments in the survey. And the cost of running one's own survey may be prohibitive for many researchers. If there are no settings in which a researcher can implement the method, what good is it to have this new tool? This is a legitimate concern and expanding the range of settings where researchers can run RCTs would help improve the robustness of this type of estimate and ensure more external validity. It is worth noting that the number of surveys in which RCTs can be implemented, in conjunction with survey research teams, is expanding. Furthermore, while some surveys do not allow for information treatments, they do allow for the addition of hypothetical questions which can target the same kind of question, as we discuss further below. Third, online surveys are another affordable option for researchers interested in the beliefs of ordinary households.

Successful Information Treatments

A necessary condition for the strategy to correctly identify how beliefs affect decisions is that the first stage must be successful: the information treatment must create sufficient exogenous variation in expectations. The ones described in CGW and GGCK worked very well, but this will not always be the case. Information treatments that involve beliefs that are first-order concern for agents will tend to be less successful, since agents will be more informed about these topics in the first place. Indeed, even the same information treatments may have different effects over time or in different places depending on the incentives agents have to be informed about the topic. For example, while CGW find large treatment effects on the inflation expectations of households in 2018, the same information treatments had much smaller effects in 2022, when inflation was high (Weber et al. 2024), reflecting the fact that U.S. households became more informed about inflation during the inflation spike. So researchers must take care to choose information treatments and expectations that are likely to be responsive, which may limit the scope of topics that can be studied with this approach.

Measurement of Outcomes

As already illustrated in CGW, how outcomes are measured can be important. The ideal scenario is for surveys to be matched to external sources of information such as administrative data that provide precise measures of the decisions taken by agents, as is possible in the Nielsen Homescan data. But this is not always feasible. There are two options that can be used when external data is unavailable. One is to rely on subsequent survey waves and to ask respondents about their decisions in those subsequent periods. CGW show, for example, that self-reported measures of spending line up closely with actual spending levels on average. However, it can take a large number of observations for this average to obtain. This is because self-reported data will include rounding by respondents, recall error, and additional forms of measurement error. As a result, it may be more difficult to establish clear causal links between expectations and decisions than would be the case with better data on outcomes. One potential solution is to use surveys with large numbers of respondents, which can help improve identification. A second solution is to focus on outcomes that are less susceptible to measurement and recall error, for example the extensive margin of whether a durable good was purchased as opposed to the intensive margin of how much was spent on durable goods. Third, there is scope for matching more surveys to external sources of information, as in Caplin et al. (2023).

External Validity

One concern that always arises with RCTs is whether the results would generalize to another time period or another setting, i.e. the external validity of the study. This is an important concern which can ultimately only be addressed by repeating RCTs in many different contexts. The replication of studies with new experiments is therefore particularly important for this line of work. But it is also important to recall that, to the extent that estimated parameters are often reduced form coefficients rather than structural parameters, one should not necessarily expect them to be the same in different contexts. Understanding the mapping from estimated coefficients in RCTs to underlying structural relationships is important to identifying which results would be expected to hold in different situations and which would be expected to vary. This is another reason why it can be useful to separately identify direct and indirect effects of changes in expectations, since their combined effects are unlikely to map easily into structural parameters whereas direct effects may have a more natural structural interpretation. Consistent with this, Candia et al. (2024) argues that the total response of durable consumption to inflation expectations should vary depending on the inflation environment and provides evidence that this is indeed the case.

Alternatives to RCTs

Despite the growing availability of surveys that allow for RCTs, it remains true that access to these surveys is limited and the cost of running one's own survey can be prohibitive. Ideally, one would like to have complementary approaches that are more easily accessible. One such approach is the use of hypothetical scenarios, as discussed in detail in Colarieti et al. (2024). The idea is to ask survey respondents what they would do if they held different beliefs. In principle, this can provide a direct estimate of the passthrough coefficient that would be estimated through an RCT. But unlike having to do the full RCT, a hypothetical question does not require different treatment arms (and therefore can be done in smaller surveys), does not require external information on decisions or follow-up waves (so it can be implemented in a single wave), and can be phrased in such a way as to either allow for or exclude indirect effects of expectations. This approach has been used extensively in the New York Fed's SCE, which does not allow for information treatments but can incorporate occasional additional questions that are framed as hypotheticals (e.g. Armantier et al. 2022). As reviewed in Colarieti et al. (2024), available comparisons of RCTs and hypotheticals is supportive of the notion that they tend to yield similar results. However, it remains to be established how generally this is the case and what circumstances might cause the two to yield different results. Despite this, we view the use of hypotheticals as very complementary to RCTs since they aim to estimate similar passthroughs but can be implemented in a much more tractable and accessible manner.

Partial versus General Equilibrium

A final concern worth emphasizing is that coefficients from RCTs like the ones described here at best identify partial equilibrium responses of agents to a change in beliefs, which may differ significantly from general equilibrium outcomes. The latter cannot be assessed directly from RCTs. Instead, moving from partial equilibrium estimates to general equilibrium predictions requires a model and a mapping from the RCT estimates into such a model. One example of this is provided in Ropele et al. (2024), who use RCT estimates of how inflation expectations affect firms' employment and investment decisions to study the general equilibrium costs of misallocation induced by inflation. But this is more the exception than the rule: there would be much value added from seeing more research on how to reconcile RCT estimates of expectations passthrough into decisions with theoretical models that allow us to speak to general equilibrium outcomes.

5. Conclusion and Policy Implications

With RCTs providing clear exogenous variation in the inflation expectations of households, recent research has been able to characterize in unprecedented detail the different ways in which these expectations affect households' decisions. In parallel, other papers have used similar strategies to study how inflation expectations affect the decisions of firms, while another line of research has utilized the same type of method to quantify the passthrough of different expectations into the decisions of firms and households. For macroeconomists, this line of research is providing new evidence on just how expectations affect the decisions of agents, a question which has long been outside the scope of clear causal empirical tests.

For policymakers, the results should also be of interest. Since policy communication often aim to directly affect the beliefs of economic agents, understanding how these beliefs pass through into decisions is important. In this respect, the total effects estimated in papers like Coibion, Gorodnichenko and Weber (2022) speak directly to how communications that change inflation expectations are likely to affect decisions. But understanding direct and indirect effects may be important for policymakers as well. For example, communication that can change expectations without changing uncertainty or vice versa may have more powerful economic consequences when first and second moments tend to have offsetting effects, as found in Coibion et al. (2024). Indeed, communication that actively uses both dimensions can be even more powerful, e.g. by simultaneously raising inflation expectations and reducing inflation uncertainty. Because of the strength of the unconditional correlation between first and second moments, doing so may not be easy and may require different communication styles, but the economic potential of this type of more targeted communication should make the notion worth considering.

References:

Altig, David, Jose Barrero, Nick Bloom, Steven Davis, Brent Meyer, and Nicholas Parker. 2022. "Surveying Business Uncertainty," *Journal of Econometrics* 231(1): 282-303.

Angrist, Joshua D. and Jorn-Steffen Pischke, 2009. Mostly Harmless Econometrics, Princeton University Press, Princeton, NJ.

Armantier, Olivier, Wändi Bruine de Bruin, Simon Potter, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar, 2013. "Measuring Inflation Expectations," *Annual Review of Economics* 5(1): 273-301.

Armantier, Olivier, Argia Sbordone, Giorgio Topa, Wilbert van der Klaauw, and John C. Williams, 2022. "A New Approach to Assess Inflation Expectations Anchoring using Strategic Surveys," *Journal of Monetary Economics* vol. 129.

Binetti, Alberto, Francesco Nuzzi, and Stefanie Stantcheva, 2024. "People's Understanding of Inflation," forthcoming in *Journal of Monetary Economics*.

Bloom, N. 2009. "The Impact of Uncertainty Shocks," *Econometrica* 77(3): 623-685.

Bruine De Bruin, Wändi, Charles Manski, Giorgio Topa, and Wilbert van der Klaauw, 2011. "Measuring consumer uncertainty about future inflation," *Journal of Applied Econometrics* 26(3): 454-478.

Candia, Bernardo, 2024. "Inflation Expectations and Household Spending: Different Patterns in Low and High-Inflation Settings," Manuscript.

Caplin, Andrew, Victoria Gregory, Eungik Lee, Soren Leth-Petersen, and Johan Saeverud, 2023. "Subjective Earnings Risk," NBER Working Paper 31019.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, Geoff Kenny, and Michael Weber. 2024. "The Effect of Macroeconomic Uncertainty on Household Spending," *American Economic Review* 114(3): 645-77.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Maarten van Rooij. 2023. "How Does Consumption Respond to News about Inflation? Evidence from a Randomized Control Trial," *American Economic Journals: Macroeconomics* 15(3): 109-152.

Coibion, Olivier, Yuriy Gorodnichenko, Michael Weber, 2022. "Monetary Policy Communications and their Effects on Household Inflation Expectations," *Journal of Political Economy* 130(6): 1537-1584.

Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva, 2024. "The How and Why of Household Reactions to Income Shocks," NBER Working Paper 32191.

Georgarakos, Dimitris, Yuriy Gorodnichenko, Olivier Coibion and Geoff Kenny, 2024. "The Causal Effects of Inflation Uncertainty on Households' Beliefs and Decisions," NBER Working Paper 33014.

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

Ropele, Tiziano, Yuriy Gorodnichenko and Olivier Coibion, 2024. "Inflation Expectations and Misallocation of Resources: Evidence from Italy," *AER Insights* 6: 246-261.

Roth, Christopher and Johannes Wohlfart, 2020. "How Do Expectations about the Macroeconomy Affect Personal Expectations and Behavior?" *The Review of Economics and Statistics*, 102(4): 731-748.

Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent 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," forthcoming in *Econometrica*.

Werning, Ivan, 2022. "Expectations and the Rate of Inflation," manuscript.

Table 1. The Effects of Inflation Expectations on Household Non-Durable Goods Spending

Dep. var. is indicated in the title of the panel	Actual spending, horizon, month	
	3 months	6 months
	(1)	(2)
Effect on Total Spending, scanner		
Posterior inflation expectations	0.950***	0.864**
	(0.286)	(0.336)
Observations	13,170	13,132
1st stage F-stat	134.8	128.1

Notes: The table is taken from Coibion, Gorodnichenko and Weber (2022) and reports the effect of inflation expectations on household spending on non-durables measured in Nielsen Homescan panel using the IV strategy described in section 3.1.

Table 2. The Effect of Inflation Expectations and Uncertainty on Purchases of Durable Goods,

	Dependent variable: indicator variable is a durable good is purchased.					
	Home	Durable	Car	Holiday package	Luxury items	Other
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. One month after treatment, IV.						
Posterior mean	0.421	4.812***	0.483	1.934	0.539*	0.451
	(0.268)	(1.369)	(0.315)	(1.578)	(0.283)	(0.863)
100×log(Posterior uncertainty)	-0.025**	-0.230***	-0.024*	-0.091	-0.021**	-0.055*
	(0.010)	(0.057)	(0.013)	(0.065)	(0.011)	(0.034)
Observations	11,514	11,506	11,502	11,512	11,519	11,483
R-squared	0.002	-0.041	-0.001	0.100	0.022	0.036
1st stage F-stat (mean)	118.4	113.8	117.6	114.8	118	112.7
1st stage F-stat (uncert)	100.5	99.29	99.10	100.7	101.9	101.2
KP Wald test	10.63	9.532	10.34	10.51	10.48	10.19
Panel B. One month after treatment, IV.						
Posterior mean	-0.305***	-1.695***	-0.325***	-1.158**	-0.208***	-1.452***
	(0.066)	(0.400)	(0.078)	(0.501)	(0.071)	(0.267)
Observations	8,658	8,652	8,645	8,653	8,662	8,646
R-squared	0.01	0.04	0.01	0.11	0.03	0.02
1st stage F-stat (mean)	208.3	200.1	206.6	212.8	207.3	202.6
Panel C. One month after treatment, OLS						
Posterior mean	0.077	-0.014	0.120	0.246	0.217***	0.685***
	(0.081)	(0.332)	(0.083)	(0.273)	(0.082)	(0.215)
100×log(Posterior uncertainty)	-0.131	3.383**	0.140	0.086	-0.465	-0.790
	(0.430)	(1.645)	(0.312)	(1.351)	(0.339)	(1.074)
Observations	2,654	2,638	2,644	2,629	2,647	2,634
R-squared	0.011	0.080	0.012	0.105	0.100	0.085

Notes: This table is taken from Georgarakos et al. (2024). The table reports estimated coefficients on posterior beliefs about inflation. Panel C includes the control group and the specification does not include pre-treatment beliefs. The dependent variables takes values 0 (no purchase) and 100 (a purchase is made). Heteroskedasticity robust standard errors are reported in parentheses. ***, **, * denote statistical significance at 1, 5 and 10 percent levels.

Table 3. The Effect of Inflation Expectations and Uncertainty on Portfolio Allocations.

	Portfolio shares							
	Cash	Curr./Saving account	Stocks	Mutual funds	Retirement account	Bonds	Crypto assets	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Scenario-based, immediately after treatment								
Posterior mean	0.525	-2.346	-0.876	1.144	1.039*	-0.233	0.157	0.442
	(1.068)	(1.642)	(0.665)	(0.712)	(0.595)	(0.548)	(0.195)	(0.518)
100×log(Posterior uncertainty)	-0.045	0.173***	0.025	-0.049	-0.065***	0.001	-0.009	-0.025
	(0.042)	(0.066)	(0.029)	(0.030)	(0.025)	(0.023)	(0.009)	(0.021)
Observations	13,601	13,601	13,601	13,601	13,601	13,601	13,601	13,601
R-squared	0.10	0.05	0.05	0.08	0.02	0.11	0.02	0.05
1st stage F-stat (mean)	143.9	143.9	143.9	143.9	143.9	143.9	143.9	143.9
1st stage F-stat (uncertainty)	122.5	122.5	122.5	122.5	122.5	122.5	122.5	122.5
KP Wald	12.78	12.78	12.78	12.78	12.78	12.78	12.78	12.78
Panel B. Actual, two months after treatment								
Posterior mean	-0.325	-4.894***	1.026*	0.589	1.833*	0.678**	0.016	1.612**
	(0.398)	(1.723)	(0.526)	(0.537)	(1.073)	(0.330)	(0.063)	(0.741)
100×log(Posterior uncertainty)	0.010	0.233***	-0.053**	-0.025	-0.076*	-0.036***	-0.003	-0.065**
	(0.015)	(0.071)	(0.022)	(0.023)	(0.044)	(0.013)	(0.003)	(0.030)
Observations	9,121	9,121	9,121	9,121	9,121	9,121	9,121	9,121
R-squared	0.07	0.02	0.05	-0.04	-0.05	-0.11	0.01	0.06
1st stage F-stat (mean)	101.1	101.1	101.1	101.1	101.1	101.1	101.1	101.1
1st stage F-stat (uncertainty)	91.79	91.79	91.79	91.79	91.79	91.79	91.79	91.79
KP Wald	11.30	11.30	11.30	11.30	11.30	11.30	11.30	11.30

Notes: This table is taken from Georgarakos et al. (2024). The table reports estimated coefficients on posterior beliefs about inflation. Heteroskedasticity robust standard errors are reported in parentheses. ***, **, * denote statistical significance at 1, 5 and 10 percent levels.

Table 4. The Effect of Inflation Expectations on Job Search Expectations.

	Job search intensity (# of job application)	Subj. prob. of finding a job in 3 months	Subj. prob. of losing a job in 3 months	Subj. prob. of looking for a job in 3 months
	(1)	(2)	(3)	(4)
Panel A. One month after treatment				
Posterior mean	-1.149***	-10.240**	-0.274	-1.808**
	(0.415)	(4.892)	(0.931)	(0.744)
100×log(Posterior uncertainty)	0.056***	0.365**	0.016	0.053*
	(0.017)	(0.171)	(0.034)	(0.030)
Observations	1,411	461	7,597	7,251
R-squared	-0.07	-0.07	0.03	0.03
1st stage F-stat (mean)	11.03	2.383	70.18	75.18
1st stage F-stat (uncertainty)	10.14	3.878	65.76	69.30
KP Wald	1.887	1.232	5.896	9.996
Panel B. Four months after treatment				
Posterior mean	-0.268	38.871**	0.246	-0.225
	(0.272)	(19.038)	(0.716)	(0.686)
100×log(Posterior uncertainty)	-0.007	-1.702*	-0.014	-0.007
	(0.013)	(0.915)	(0.027)	(0.027)
Observations	848	274	5,810	5,632
R-squared	0.07	-4.99	0.05	0.05
1st stage F-stat (mean)	12.54	3.101	58.56	59.61
1st stage F-stat (uncertainty)	11.52	3.184	55.05	51.87
KP Wald	3.379	0.238	6.412	7.460

Notes: This table is taken from Georgarakos et al. (2024). The table reports estimated coefficients on posterior beliefs about inflation. Heteroskedasticity robust standard errors are reported in parentheses. ***, **, * denote statistical significance at 1, 5 and 10 percent levels.

Table 5. The Effects of Inflation Expectations and Uncertainty on Employment Status.

	Employed (any)	Employed (full-time)	Employed (part-time)	Unemployed	Other (out of labor force, laid-off, etc.)
	(1)	(2)	(3)	(4)	(5)
Panel A. One month after treatment					
Posterior mean	0.649	-0.981	1.402	0.503	-1.175
	(1.873)	(1.967)	(1.098)	(0.693)	(1.826)
100×log(Posterior uncertainty)	0.031	0.116	-0.076*	-0.065**	0.032
	(0.073)	(0.077)	(0.044)	(0.026)	(0.071)
Observations	11,426	11,426	11,426	11,426	11,426
R-squared	0.37	0.32	0.04	0.03	0.40
1st stage F-stat (mean)	112.4	112.4	112.4	112.4	112.4
1st stage F-stat (uncertainty)	101.5	101.5	101.5	101.5	101.5
KP Wald	11.43	11.43	11.43	11.43	11.43
Panel B. Four months after treatment					
Posterior mean	-0.259	-2.327	2.173*	0.822	-0.716
	(1.886)	(2.026)	(1.201)	(0.565)	(1.854)
100×log(Posterior uncertainty)	0.044	0.161**	-0.121**	-0.071***	0.026
	(0.076)	(0.082)	(0.049)	(0.022)	(0.075)
Observations	8,666	8,666	8,666	8,666	8,666
R-squared	0.41	0.35	0.01	0.02	0.43
1st stage F-stat (mean)	96.75	96.75	96.75	96.75	96.75
1st stage F-stat (uncertainty)	85.54	85.54	85.54	85.54	85.54
KP Wald	8.570	8.570	8.570	8.570	8.570

Notes: This table is taken from Georgarakos et al. (2024). The table reports estimated coefficients on posterior beliefs about inflation. Employment status is measures as an indicator variable equal to one if in a given status and zero otherwise. Category “other” includes laid-off workers. Heteroskedasticity robust standard errors are reported in parentheses. ***, **, * denote statistical significance at 1, 5 and 10 percent levels.

Figure 1. Effect of Information Treatments on Expectations in Coibion, Gorodnichenko and Weber (2022)

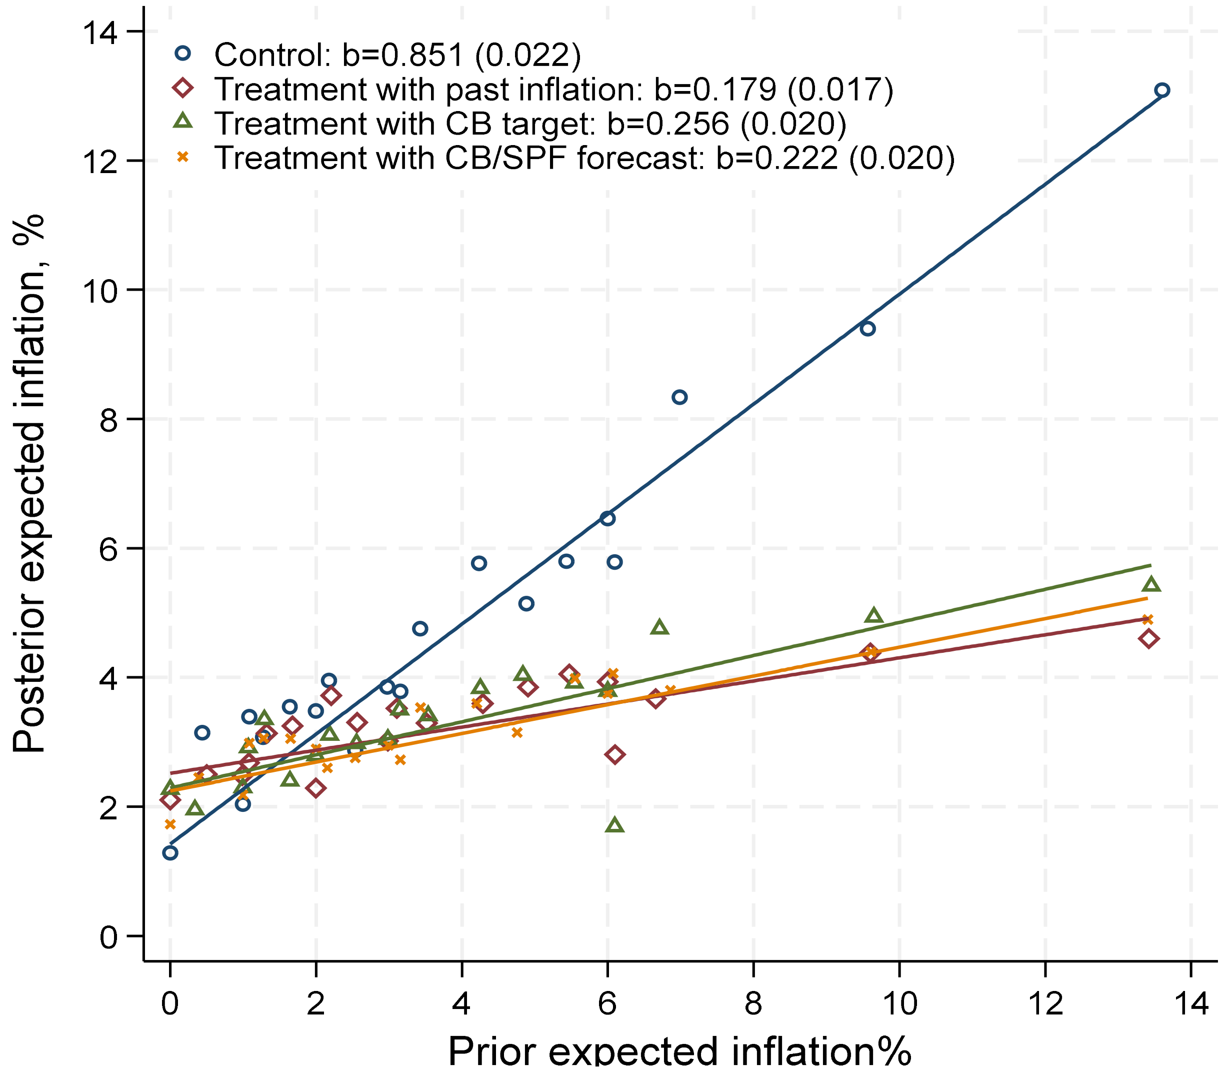

Notes: This figure presents results from Coibion, Gorodnichenko and Weber (2022). It is a binscatter plot of households' prior inflation expectations against their posterior inflation expectations for different groups of respondents based on which treatment arm they were randomly assigned to. The "control" group was not provided with any information. Other participants in the figure were either provided with the most recent inflation statistic at the time (close to 2%), the most recent SPF or FOMC 12-month ahead inflation forecast (close to 2%) or the Federal Reserve's inflation target (of 2%).

Figure 2. The Correlation of Inflation Expectations and Inflation Uncertainty

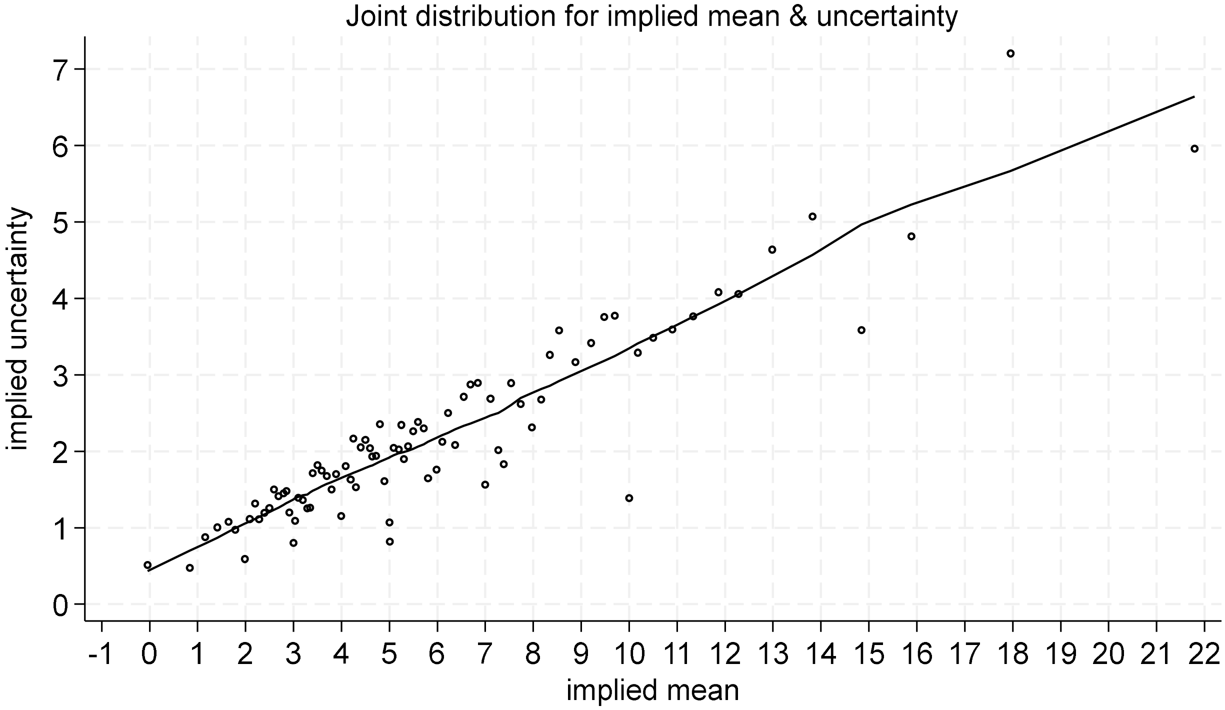

Notes: This figure is taken from Georgarakos et al. (2024) and plots the correlation between the inflation expectations and inflation uncertainty of European households in the ECB's CES in September 2023.