Household Disagreement about Expected Inflation*

Salomé Fofana LSE

Paula Patzelt LSE

Ricardo Reis LSE

March 2024

Abstract

We survey the main facts that have emerged from research on disagreement between households on what they expect inflation to be. We document them using figures and correlations that capture: the statistical regularities on the observable drivers of disagreement, the measurement of residual disagreement, the usefulness of disagreement to forecast inflation, the response of disagreement to shocks, the disagreement between households and professionals, and the relation between disagreement, risk, and uncertainty.

*Contacts: s.fofana@lse.ac.uk, p.h.patzelt@lse.ac.uk, r.a.reis@lse.ac.uk. This work was supported by the UK Research and Innovation grant number EP/Eo25039/1. First draft: February 2024.

1 Introduction

During the heyday of Keynesian macroeconomics, expectations were not the focus of attention. Models built on Hicks' IS-LM framework were static, so expectations of the future were at best an exogenous shock to the system. They were called changes in sentiment, or animal spirits, and were treated as outside unmodelled forces, akin to how economists today still sometimes refer to technology or oil shocks. An illustration are the speeches of Arthur Burns, chair of the Federal Reserve between 1970 and 1978 and before that a distinguished scholar of business cycles, in which he blames the initial rise of inflation in the 1970s on the psychology of markets and businessmen.

The rise of new classical economics, with dynamics and forward-looking behaviour, put expectations at the center of models. At the same time, the assumption of full-information rational expectations meant that expectations were fully determined by the outcomes within the model. Expectations went from being a mysterious outside force to being solely a mirror of the outcomes that we could measure in the aggregate data. Prescott (1977) argued that measuring expectations was as futile as measuring utility: it was best to focus on measuring actions and equilibrium aggregates.

Against this background, economists in the early 1980s started having access to surveys of households with numerical answers to expected inflation, namely from the Michigan Survey of Consumers. The initial empirical work was, understandably, focussed on testing for the validity of rational expectations. One early test looked at whether expectations were unbiased, so that forecast errors were on average zero. For households, this was then, and still is, easy to reject. Whether on average or individually, across time periods or across countries, expectations are biased. A second wave of work tested for the orthogonality condition that is behind many statistical estimators: forecast errors should be uncorrelated with information available at the time of the survey. One particularly salient information is the forecast error that was just realised. A powerful test of rationality is to see if forecast errors are serially uncorrelated. Again, this has been decisively rejected, over and over again.

By the turn of the century, work moved beyond clearing this low bar of rejecting rational expectations. Researchers extracted facts about household expectations to guide the building and testing of alternative models of expectations. One branch kept the focus on average forecast errors, describing how they slowly adjust to exogenous shocks or to forecast revisions. Coibion and Gorodnichenko (2015) was a key paper in that literature. Another branch, kickstarted by Mankiw, Reis and Wolfers (2004), instead looked at the

second moment in the cross-sectional distribution of expectations. By itself, the existence of this disagreement rejects full information rational expectations where, as Thomas Sargent often emphasised, there is a "communism of beliefs" shared by every agent in the model. More usefully, the properties of this disagreement and how it evolves with inflation, responds to shocks, and varies over time have fuelled the growth of models where there is incomplete information and no common knowledge.

This article lays down some of the lessons from this empirical work after twenty years. We do so by presenting some of the main facts about disagreement in the data, using original figures and tables. These build on the robust findings across many articles, but we do not provide a survey of these papers and the debates between them, nor a connection between data and theory. Instead, we will present main facts in the simplest way we were able to do, with plots and correlations.[1]

Section 2 describes a set of fixed observable characteristics that account for some disagreement. Section 3 describes measures of residualised disagreement driven by unobservables. Section 4 presents the facts on how that disagreement is driven by inflation, or helps forecast future inflation, while section 5 links disagreement to exogenous shocks. Our focus is entirely on disagreement between households, but section 6 describes its link to disagreement among professionals and to market price measures of inflation uncertainty. Section 7 concludes.

2 Disagreement driven by observable individual characteristics

We use the three most popular sources of households inflation expectations data, all at a monthly frequency, and focusing on the 1-year ahead horizon. The first is the Michigan Survey of Consumers (MSC) of US households. Its main virtue is that it has run continuously since 1978:1, while its main flaw is that is covers only 400 to 1,500 households per wave, with recent counts at the lower end of this range. This is a small cross-section over which to calculate second moments. The second is the ECB Consumer Expectations Survey (CES), which only starts in 2020:4, but has 9,000 to 22,000 respondents per wave. In between them is the Survey of Consumer Expectations (SCE) of US households run by the Federal Reserve Bank of New York, which starts in 2013:6, and covers approximately

[1] There are many good recent surveys of empirical work on survey expectations, for instance D'Acunto, Malmendier and Weber (2023) and Coibion, Gorodnichenko and Kamdar (2018).

1,300 respondents. Throughout we use samples that end some time in 2023.[2]

2.1 The facts

Figure 1 shows the coefficients of a panel regression of expected inflation on dummy variables for social, economic, and demographic characteristics. This amounts to a comparison of the average expected inflation over time across each of these groups.[3] We do the analysis for each survey, as well as for their common sub-samples, in order to judge whether the facts are robust across time, region, and survey. For the MSC, we also separate the period before 1995 when, arguably, the monetary and inflation regime of the US was different, with higher, more volatile, and more persistent inflation outcomes.

The first characteristic is gender. An often-repeated claim is that that women expect higher inflation than men, although it is hard to trace it to particular papers and actual statistics. Figure 1 confirms this is indeed a robust fact, across time, region, and survey.

The second characteristic is income, which we split into three categories for the US—less than $50,000 annual household income, between $50,000 and $100,000, and above $100,000—and into five quantiles for the EA. The general pattern is that the higher is income, the lower is expected inflation. This is true not just between the richest and the poorest but quite consistently throughout the income distribution.

Marital status is a third robust characteristic associated with differences in expected inflation. While not always statistically significant, still across all times, regions, and surveys, those married usually expected higher inflation.

Education and age are more subtle. Starting with the first, we measure it for the US by separating respondents with a college degree, or some college attendance, relative to those without any college. Since 1995, holders of a college diploma consistently report lower expected inflation than those without college, and those with some college in between. The reverse was true between 1978 and 1994 in the US. In Europe, the separation is between going to college, having a high school degree, or less. It is also the case that those who went to college expect lower inflation than those who just have a high school degree. But, not even finishing high school leads to the lowest inflation expectations.

Turning to age, until 1994, those less than 40 years' old reported significantly higher

[2]CES data is available until 2023:12, MSC data until 2023:7, and SCE data until 2023:4.

[3]One might worry that the composition of the survey according to these characteristics has changed over time in a way that could correlate with changes in the differences in expected inflation across groups. It has not, as can be confirmed by including time fixed effects in the regression, which leads to very similar estimates of the coefficients.

expected inflation. Since then, this pattern reversed with the younger reporting lower expected inflation. In turn, those older than 60 expected significantly less inflation in 1978-94, but after 1994 their expectations are statistically indistinguishable from those aged 40 to 60.

Figure 1: Socio-demographic-economic determinants of inflation expectations

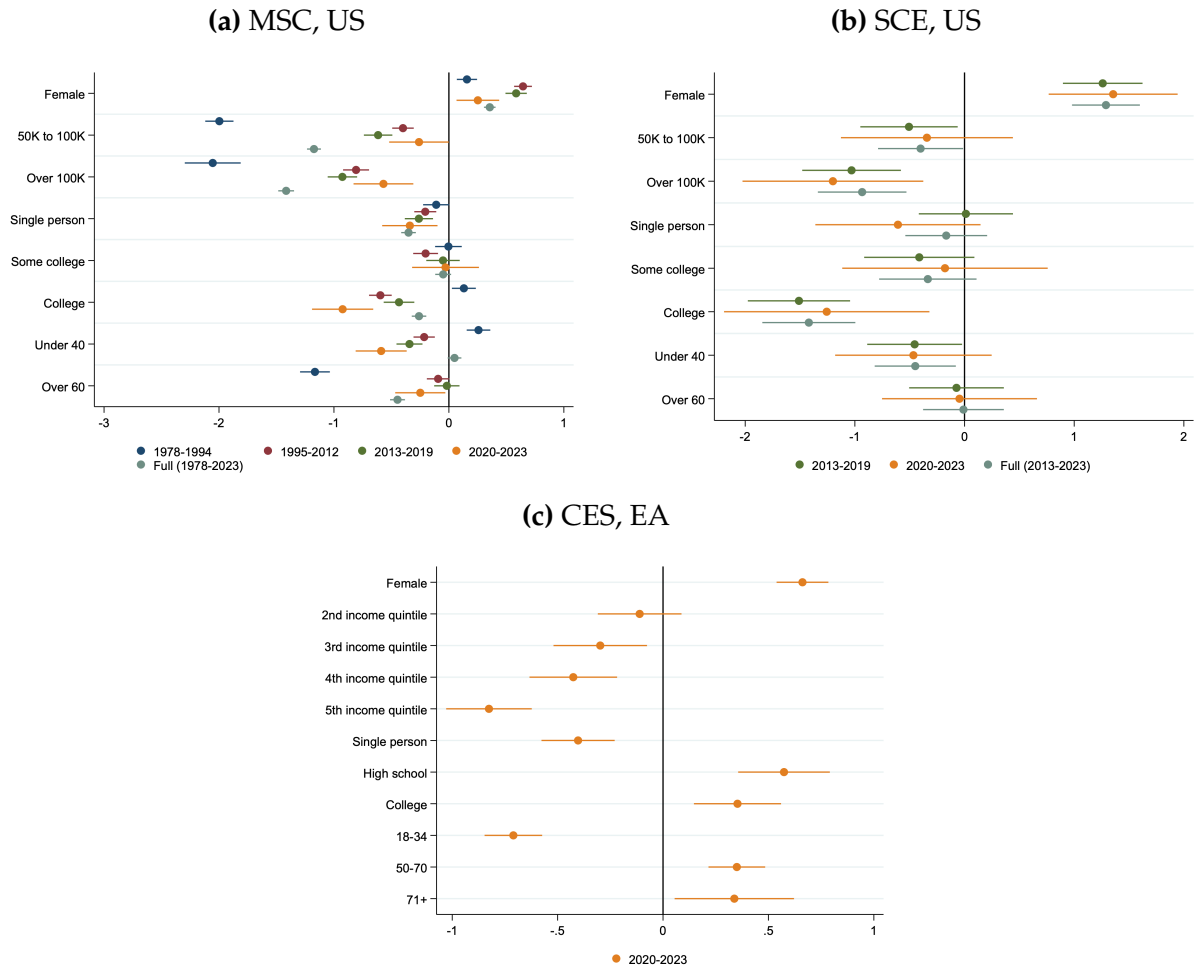

Note: Coefficient estimates and 95% confidence bands from a regression of expected inflation on socio-demographic characteristics by survey and subsample period.

Figure 2 digs deeper into this variation over time, by using the MSC's longer sample to estimate the education and age coefficients separately at each year through a series of cross-sectional regressions. For education, this shows that the positive effects of education on expected inflation we found for the pre-1995 sample are driven solely by the first four

years of the sample, 1978 to 1981. It is also the case that, for these four years alone, even the coefficients on gender (not reported), which are among the most robust at all other dates and regions, are of the opposite sign. This makes us skeptical of these years at the start of the survey. Without them, there is strong evidence for a negative, and relatively stable, relation between education and expected inflation. This relation has become even stronger in the last two years of the sample.

Figure 2: The time-varying influence of education and age on inflation expectations

(a) Education: college (left) and some college (right) versus high school

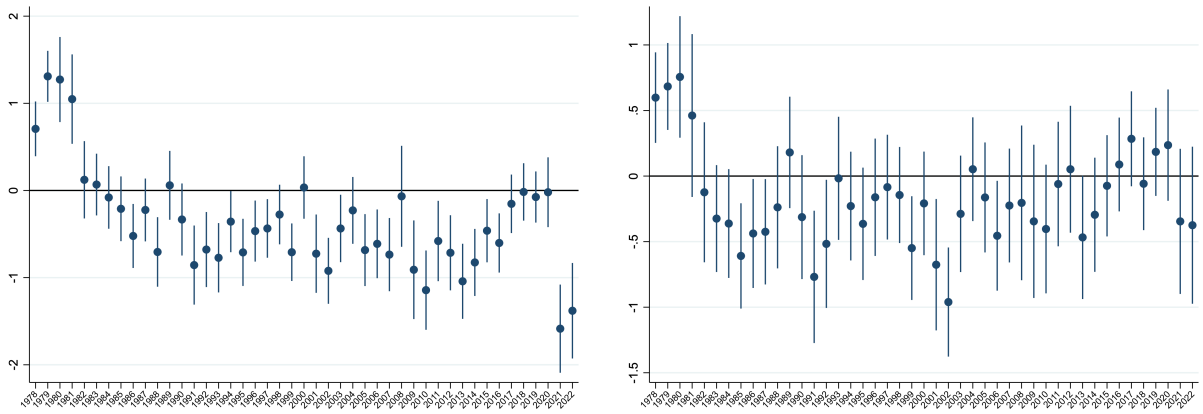

(b) Age: under 40 (left) and over 60 (right) versus 40–60

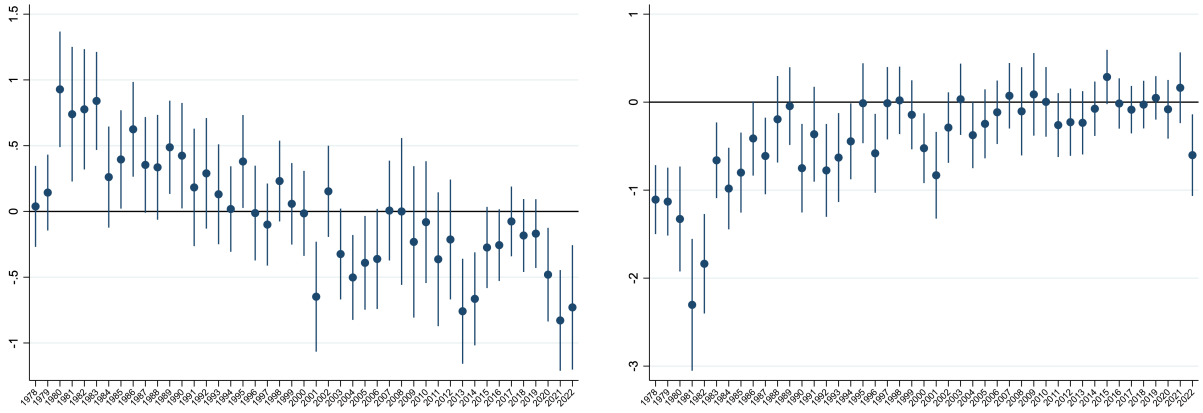

Note: Coefficient estimates and 95% confidence bands from regressions of inflation expectations on socio-demographic characteristics for the MSC survey for each separate year.

For age, there is a clear trend over time in the estimates until the end of the 1990s across all age groups, with older age increasingly being associated with expecting less inflation.

The effect for the first two decades of this century has been approximately zero. But then, and consistent across all surveys and regions, in 2021 and 22 as inflation exploded, the younger have started expecting lower inflation.

A concern with all survey answers, especially pressing when studying disagreement, is what to do with extreme observations. We have censored all answers that expected inflation above 20% or below -20%. Investigating the robustness of the results to somewhat higher or lower threshold is best practice in this literature. A closely related, but often ignored, concern is the large share of respondents who do not answer or respond "don't know" to this question, which is mostly between 5% and 15%, averaging to 9% since the start of the MSC. This is not a random sample of the population, since DK/NA answers on expected inflation are correlated with the same observable characteristics that drive disagreement.

To illustrate this issue, we use the Fofana (2024) imputation procedure to create pseudo-answers on expected inflation using the respondents' answers to all the other questions, excluding the questions on socio-demographic characteristics that we evaluated above. There are 14 such questions asked consistently since 1978, and the procedure essentially uses the predicted values from regressions for a sample of respondents that provide expected inflation in order to guess what the answer would have been for those who did not. Doing so, the imputed missing DK/NA responses for both men and women lead to a higher average expected inflation, but more so for women. The average disagreement between the men and women increases from 0.94% to 1.01%. This suggests that DK/NA answers lead to an understatement of the extent of disagreement in the data.

2.2 Interpretation

A long literature has investigated the reasons behind this predictable disagreement that makes a married, poor, older, less-educated woman expect higher inflation than her counterparts. One line of work has emphasised that people are exposed to different baskets. When they answer what they expect inflation to be, people are referring not to the CPI, or to some common aggregate, but rather to their individual price basket. For instance, if women's basket differs from men's, the different price changes in those baskets might explain their disagreement. Because these basket differences are observable, this is one of the most tested explanations for observable disagreement. It explains some of it, but far from all.

A related hypothesis is that people differ in the history of price changes that they

have gone through. To give one example, a young and an old person have experienced different inflation, and this may shape the way they make forecasts about the future. When it comes to differences in age, this has been a powerful explainer of disagreement, but it has been comparatively less explored for the other characteristics.

Forecasts reflect the information of the person making them. Disagreement about the future tends to be correlated with disagreement on perceptions of the present. Measuring that individual information is not easy, but studies have tried to measure the sources of information that people have, for instance through their social networks, to explain the differences in observed disagreement.

Information is an input into a mental model to produce a forecast. That model might reflect different views of how the economy works, as well as different attitudes towards optimism or an aversion to being wrong. Because the models that people use to make inferences have been studied in many other contexts and associated with observed socio-demographic characteristic, they can be linked to the differences in expected inflation in a disciplined way that is consistent across different forms of behaviour. For instance, higher risk and loss aversion by married people across experiments is consistent with them expecting higher inflation on average. More recent research has shown that this model is affected by political partisanship, as disagreement about expected inflation is correlated with party affiliation.

Finally, insofar as forecasting is one of many tasks that people undertake with varying success, disagreement could be explained by observable differences in financial literacy or broader measures of quantitative skills. Some work has shown this association may be there in the data.

3 Measuring disagreement

The R^2 of the regressions of expected inflation on observable characteristics were uniformly low. That is, after taking into account economic, social, and demographic characteristics, people still disagree substantially on what they expect that inflation will be. From now onwards, we will use as data the residualised expected inflation from the full-sample regression on these observable characteristics.

3.1 The facts

Figure 3 plots the standard deviation of expected inflation across households for the three surveys on the left panel, and the skewness coefficient on the right panel. Disagreement, across measures, samples, regions, and surveys is always large, well above the horizontal axis.

Disagreement also varies over time. Disagreement measured by a second moment, as in the left panel, was higher in the early 1980s, when inflation was higher, but also spiked up during the great financial crisis when inflation was unusually low. It rose again markedly during the 2021-22 inflation disaster. Disagreement measured by a third moment, shown in the right panel, behaves differently. It is significantly noisier varying from month to month, and it has a downward trend since the 2010s. Like the second moments, there are clear fluctuations at business-cycle frequencies.

Figure 3: Disagreement over time

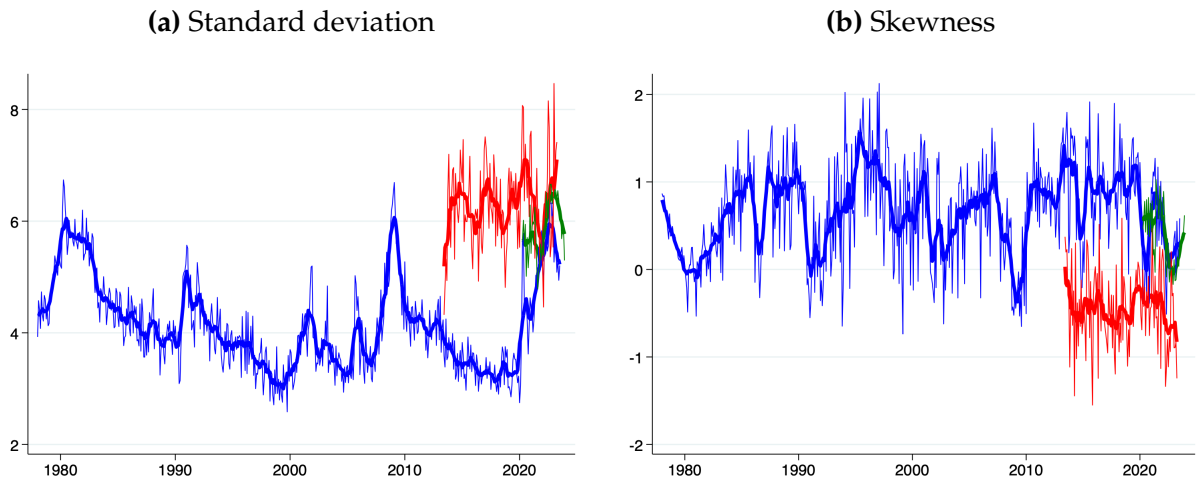

Note: Each series corresponds to a survey, with the MSC in blue, the SCE in red, and the CES in green. Thin lines show the monthly series, thick lines show a 9-month centred moving average.

Are these two statistical measures the right way to capture disagreement? The top panel of table 1 shows the correlations for each of the three surveys, with their different periods and regions, across three alternative measures of the second moment of the distribution of expectations: a trimmed standard deviation over a truncated sample that excludes the top and bottom 10%, one half of the difference between the 16th and 84th percentiles (sometimes referred to as a quasi-standard deviation) and the interquartile

range. For all three surveys, they are positive and usually quite high. Results are similar across the sub-samples in time as in figure 1. This suggests that, for most applications, it will not make much of a difference which measure is used. Also included in the survey are the correlations with the first moment, measured by the mean. These are lower, and can vary considerably over sub-sample.

The bottom panel of Table 1 shows similar correlation coefficients for the higher moments, starting with the standard Fisher coefficient of skewness, and reporting also mode skewness (which is three times the non parametric skew) and the kurtosis. The noise that we saw in the figure also shows up as lower correlations than among the second moments. Picking which one to use is a potentially larger concern. The table also shows the correlation between these higher-order moments and measures of the first and second moments in the distribution, respectively the mean and the standard deviation. They are not stable across surveys. Not shown, they also change considerably over time periods within one sample.

Table 1: Alternative measures of (residualised) disagreement

(a) Correlations with the standard deviation

	MSC	SCE	CES
Standard deviation	1.00	1.00	1.00
Trimmed standard deviation	0.49***	0.23***	0.32**
Quasi-standard deviation	0.88***	0.62***	0.81***
Inter-quartile range	0.83***	0.32***	0.81***
Mean	0.43***	-0.13	0.53***

(b) Correlations with skewness

	MSC	SCE	CES
Skewness	1.00	1.00	1.00
Mode skewness	-0.01	0.24*	0.49**
Kurtosis	0.19***	-0.37***	0.60***
Standard deviation	-0.54***	-0.49***	-0.91***
Mean	0.11**	0.26**	-0.58***

3.2 Interpretation

Why would there be disagreement beyond observable characteristics? An old tradition, dating back to the famous Lucas islands' model, thinks of people as receiving individual

signals with which they form expectations. Disagreement in expectations would then reflect the distribution of signals across people. From this starting point, the literature has moved in complementary directions.

One seeming puzzle is that information about current inflation, or about official forecasts from reputed organisations, takes seconds to get with an internet connection. One would expect disagreement to be minor, or at least to have a strong downward trend over the last thirty years with the spread of information technologies. The data strongly contradicts this.

The response of economic theorists for more than twenty years has been to emphasise inattention as a foundation for disagreement. People could get information that would make disagreement tiny, but they choose not do so, because the benefits do not exceed the costs. Relative to the individual risk they face, aggregate risk from inflation in developed countries has been small. While the costs of acquiring information have gotten smaller, the costs of absorbing it and processing have not, in the sense of building sufficient statistics and plugging them into optimal policy functions.

Models of inattention justify not just incomplete information but also its dispersion. If each person is choosing a few bits of information to absorb through a channel with limited capacity, then there will be individual noise in that channel generating the disagreement. Moreover, people update their information at different times. They face slightly different costs of obtaining different types of news. They have some natural complementarities between paying attention to inflation and their other individual actions. All of these can justify the large residualised disagreement that we observe.

One feature that data from countries with richer inflation experiences than the US or the EA has shown is that changes in the monetary policy regime and in inflation outcomes can have a permanent impact on disagreement. Namely, adopting an inflation target and giving independence to the central bank to pursue it tend to have an appreciable downward effect on disagreement. Whether his happens because they also lower the average and the volatility of inflation, or because they lower the costs of getting information about inflation, is harder to disentangle.

The other feature of disagreement in the figure is that it varies over time. In the Lucas model, disagreement was a constant, and it is constant as well in some inattention models. Given the data, time-varying disagreement should be a feature of all models of expectations, especially the ones used to understand business cycles.

Perhaps the variation is driven by cycles in attention. Or, perhaps it is driven by the

same shocks that drive the business cycle, also affecting the forecast errors that agents make. Either way, the study of business cycles as a study of the mechanism through which shocks amplify and propagate has already benefitted greatly from looking at disagreement as a measure of some of these channels, and much more can be done.

4 The information from disagreement about inflation

A regression of realised inflation on expected inflation for the MSC, SCE and CES surveys gives coefficients of 1.7, 0.8, and 2.5, and R^2's of 0.51, 0.33, and 0.74, respectively. That is, the average expected inflation from surveys is typically a strong statistical predictor of inflation, but the coefficient can be far from one. This leads to predictable forecast errors.

4.1 The facts

Table 2 uses measures of disagreement to improve on those inflation forecasts. Namely, it regresses the forecast errors that result from using the first moment alone, on the second and higher-order moments of expectations. In the US sample covered by the SCE, disagreement does not improve forecasts. But over either a longer sample in the US, as well as over the more recent years of the inflation disaster, disagreement through second and higher moments helps to explain (and so lower) the forecast errors.

Table 3 continues this study of the forecasting power of disagreement by estimating instead quantile regressions of realised inflation on moments of the distributions of expected inflation. Having at least one measure of disagreement in expectations often improves our ability to predict episodes of low, moderate, or high inflation.

Table 4 moves one step farther to investigate whether disagreement is a signal of inflation disasters. We identify in the US sample periods when in the past 12 months inflation increased or decreased by at least 3%. We then regress a measure of disagreement on these dummy variables for fast-changing inflation. The data very strongly indicates that that inflation disasters are times of disagreement about inflation.

Overall, ignoring disagreement leaves out much significant information in the surveys about what inflation will be.

Table 2: Disagreement and forecast errors over inflation

	MSC		SCE		CES	
	(1)	(2)	(3)	(4)	(5)	(6)
Standard deviation	0.499***	0.186	0.424	0.129	-3.422***	-0.211
	(0.167)	(0.150)	(0.291)	(0.482)	(1.098)	(1.760)
Skewness	-0.570***	-0.644***	-0.047	-0.636	1.350	3.599**
	(0.178)	(0.181)	(0.451)	(0.835)	(1.734)	(1.583)
Kurtosis		-0.332***		-0.226		1.392*
		(0.087)		(0.241)		(0.709)
Constant	2.106***	5.462***	0.235	3.238	24.882***	-1.519
	(0.670)	(0.954)	(1.682)	(3.987)	(7.031)	(13.531)
Observations	535	535	107	107	34	34
R^2	0.075	0.126	0.021	0.028	0.452	0.532

Note: Estimates from time-series regressions of the forecast error of the mean expected inflation on different measures of disagreement. Newey-West standard errors with one lag in parentheses.

Table 3: Disagreement and inflation quantiles

	MSC			SCE			CES		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Quantile	0.25	0.50	0.75	0.25	0.50	0.75	0.25	0.50	0.75
Mean	0.096	1.058***	1.756***	0.031	0.430**	1.359***	1.160***	1.084**	-0.002
	(0.105)	(0.165)	(0.138)	(0.134)	(0.206)	(0.346)	(0.388)	(0.512)	(0.477)
Standard deviation	0.247***	0.505***	0.408***	-0.344*	0.294*	0.485	-3.996*	-3.110	-2.891
	(0.080)	(0.100)	(0.156)	(0.196)	(0.160)	(0.426)	(2.163)	(2.261)	(1.894)
Skewness	-0.293***	-0.640***	-0.659***	-0.197	-0.181	-0.288	0.284	3.571	0.161
	(0.111)	(0.167)	(0.210)	(0.315)	(0.475)	(0.561)	(2.624)	(3.017)	(2.565)
Constant	1.214***	2.116***	4.084***	3.327***	0.362	0.971	26.892*	22.081	24.540**
	(0.374)	(0.516)	(0.718)	(1.135)	(0.802)	(2.812)	(13.524)	(14.128)	(11.448)
Observations	535	535	535	107	107	107	34	34	34

Note: Estimates from time-series quantile regressions of actual inflation over the following 12 months on moments of inflation expectations for these 12 months. For each survey, the three columns show results for quantiles 0.25, 0.50, and 0.75 in this order. Newey-West standard errors in parentheses.

4.2 Interpretation

There is a large and varied literature writing models for inflation that is driven by incomplete and disperse information. In fact, one may even interpret models of sticky prices as

Table 4: Disagreement and inflation disasters

	(1)	(2)
Disaster in previous year	1.118***	0.426***
	(0.126)	(0.085)
Disaster in forecasted year	0.413***	1.344***
	(0.081)	(0.109)
Constant	3.884***	3.854***
	(0.045)	(0.038)
Observations	545	545
R^2	0.335	0.485

Note: Estimates of a time-series regression of the standard deviation of residualised expected inflation on two dummy variables indicating whether an inflation disaster occurred in the year leading up to the forecast, or a disaster occurred in the forecasted year. An inflation disaster is defined as at least one instance when year-on-year inflation was 3pp higher (column 1), or 3pp lower than 12 months prior (column 2). Newey-West standard errors with one lag in parentheses.

particular versions of models with sticky information: people infrequently update their information, and then are further constrained to choose a price plan that consists of a single fixed number. Price dispersion, which is the driver of the Phillips curve and of the welfare costs of inflation, is in part a reflection of disagreement about expectations.

One difficulty with this work becoming more widely used has been the absence of a single parsimonious model for expectations that replaces the widespread use of rational expectations. There have been some proposals for such a model, reflecting a relative agreement on what its properties should be.[4] If one of these or a new alternative takes off, linking disagreement to other macroeconomic variables beyond inflation should be a fruitful avenue for work.

5 Disagreement after shocks

Dynamic responses of endogenous macroeconomic variables to identified shocks are useful guides to models.

[4]For instance, Angeletos, Huo and Sastry (2020), Reis (2020), or Bordalo, Gennaioli and Shleifer (2022).

5.1 The facts

Figure 4 shows the impulse response of disagreement measured by the standard deviation in the MSC, estimated through a local projection on four commonly-used shocks. The first is a monetary policy shock that comes from high-frequency responses of interest-rate futures around dates of policy meetings and speeches of policymakers, calculated by Bauer and Swanson (2023). The second is a supply shock, calculated as changes in the prices of futures contracts on oil prices following OPEC production target announcements, provided by Känzig (2021). The third is instead a demand shock, coming from news about additional federal military spending, from Ramey and Zubairy (2018). The fourth and final is a catch-all shock, from Angeletos, Collard and Dellas (2020), that starting from a set of the main 9 macroeconomic aggregates, orthogonalizes them according to their contribution to variance at business-cycle frequencies. We use the shock that maximizes business-cycle volatility of private consumption. We square each of the shocks, to see what happens to disagreement after a shock, whether it is positive or negative. The figure also reports 66% confidence bands calculated with Newey-West standard errors and, depending on the shock series, it is calculated over quarters or months.

Across all shocks, a shock tends to increase disagreement. The effect is imprecise, and often not significant after 6 months, but it is quite robust across measures of disagreement, samples, or alternatives to these four shocks (not reported). We do not report (but calculated) what happens in response to the shock itself, not squared. Most of those responses are insignificant. This is consistent with theories of disagreement, since both positive and negative shocks will raise disagreement if agents are slowly or imperfectly learning. On average their effect would be zero, as we find in the data.

Taking one step further, we look at large shocks. Oil prices in the US, measured using the West Texas Intermediate spot price, had their largest jumps since 1970 on the following four time periods: 1973:12 to 1974:1 (134%), 1979:4 to 1980:4 (148%), 2007:1 to 2007:11 (73%) and 2020:4 to 2022:6 (592%). Figure 5 shows the distribution of inflation expectations around each of these periods. Each of these shocks comes with an increase in mass on the right of the distribution, which would translate to a rise in the standard deviation and a rise in skewness. For each of them as well, this initial increase in the right-ward mass is temporary and is followed by a shift in the middle mass of the distribution to the right.

Figure 4: The response of disagreement to squared shocks

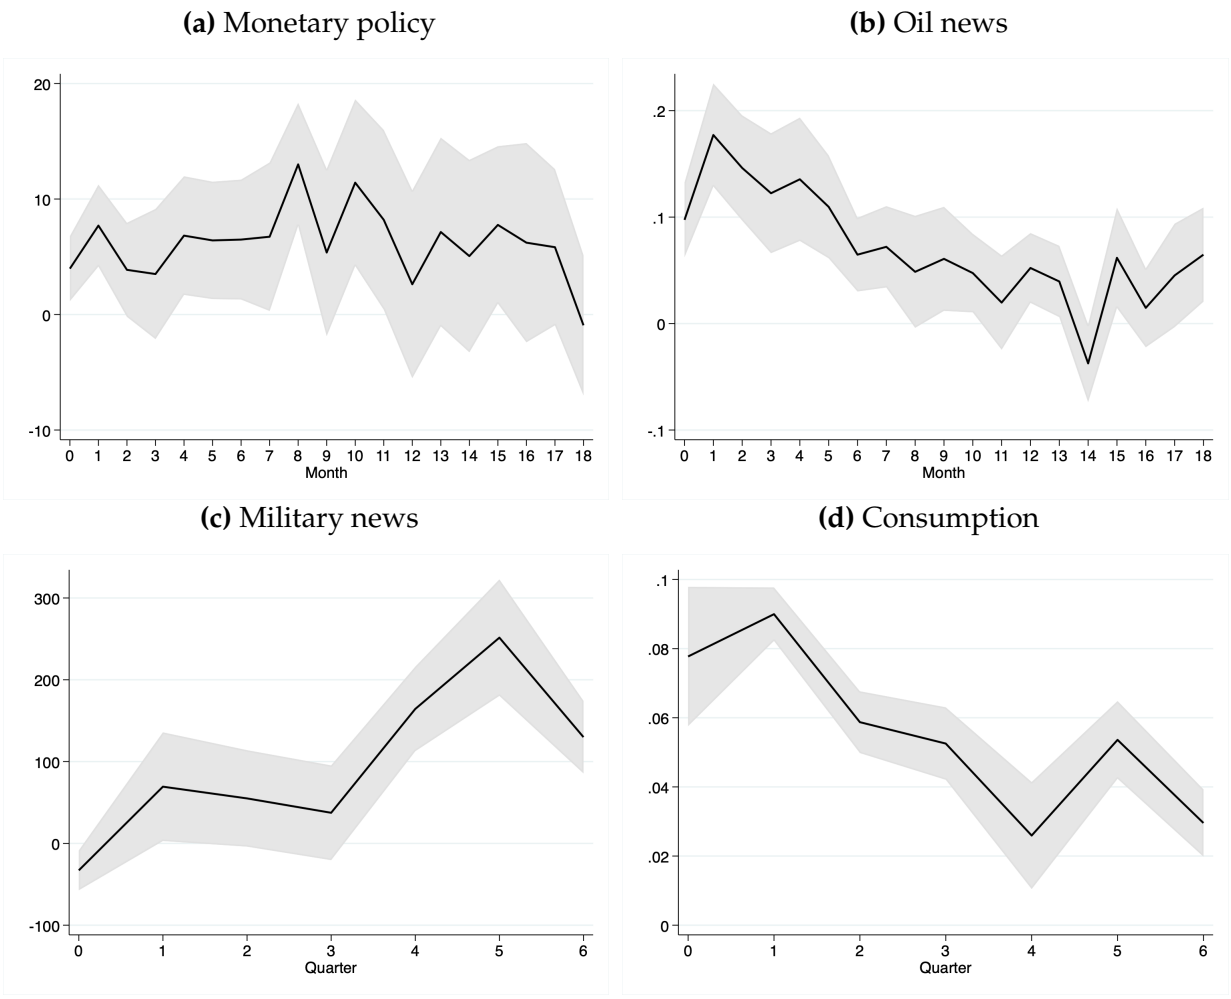

Note: Household disagreement is measured by the residualized standard deviation of MSC 1-year ahead inflation expectations. The series of monetary policy surprises comes from Bauer and Swanson (2023), using FOMC announcements only, and orthogonalized. Oil news shocks are from Känzig (2021) using OPEC announcements. Military news shocks are from Ramey and Zubairy (2018) treated by a Gordon-Krenn procedure. Consumption shocks are from Angeletos, Collard and Dellas (2020). 66% confidence bands are displayed, using Newey–West standard errors.

5.2 Other shocks

With a focus on expectations, it is not only fundamental shocks matter for the macroeconomy. Also, news about these fundamentals are important. In parallel with the study of expectations, a literature on news shocks has tried to model the changes in the information sets of economic agents. Data on expectations and disagreement are central to

Figure 5: Disagreement around large oil shocks

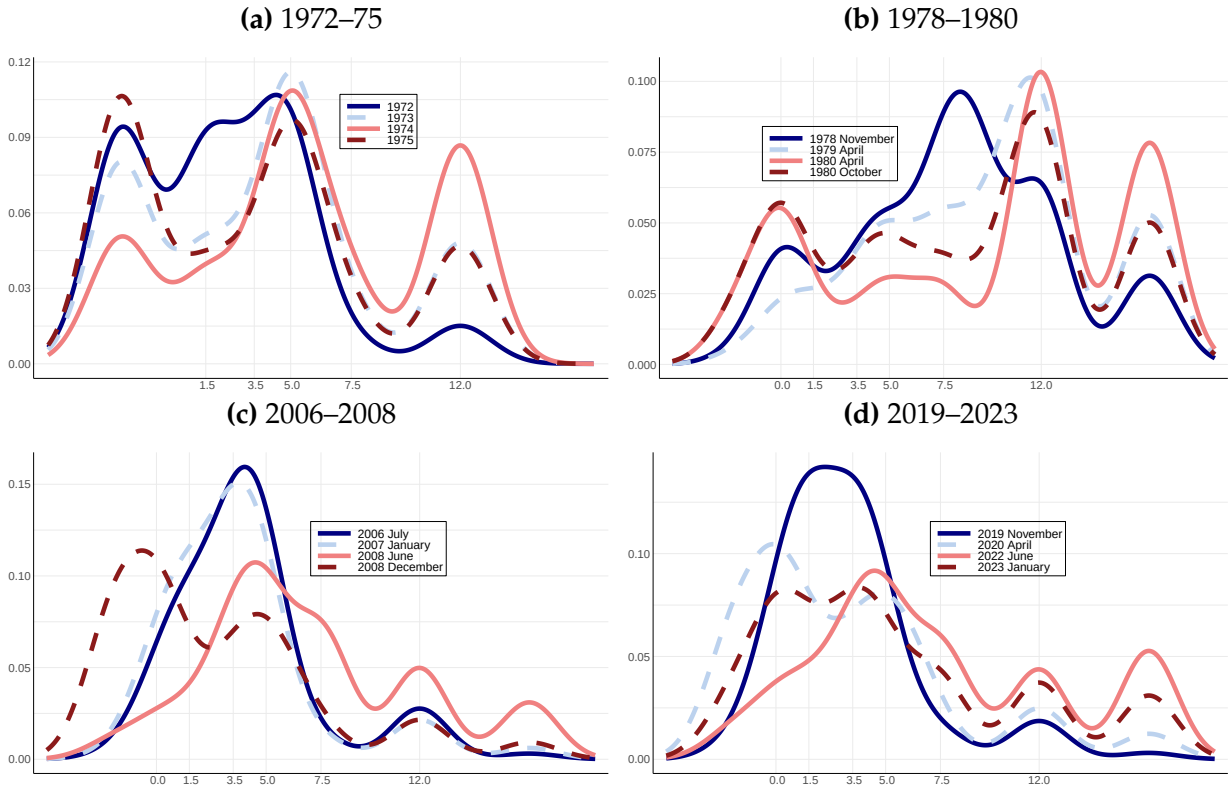

Note: Distribution of 1-year ahead inflation expectations in the MSC around oil shocks. The data for 1972–75 in subfigure (a) is from Reis (2021).

discipline those models.

Communication by central banks provides information about both monetary policy shocks as well as the central bank's assessment of the economy. A rich literature has examined the impact of that communication on disagreement. It is often hard to identify communication shocks that are both exogenous and large enough. For the few that we have, they reduce disagreement, unlike all the other shocks in figure 4.

In turn, data on the disagreement responses to communication teaches us about the effect of different types of communication. An important finding is that communicating with the public is quite different from communicating with markets or professionals. Further, a typical policy message will create disagreement between these groups, a topic we turn to in the next section.

Before doing so, one last use and interpretation of the results on disagreement around large shocks in figure 5 is in relation to credibility. The large movements in skewness

first, and standard deviation afterwards, are arguably associated with the response of monetary policy to these shocks. The data on longer-horizon inflation expectations and on risk premia in inflation contracts confirms this.

Perhaps the most striking episode of the usefulness of disagreement data is that plots like those in figure 5 provided an early warning sign in 2021 of the inflation disaster that followed. The apparent small movements in the first moment of the distribution hid large movements in the higher moments. Disagreement data alerted policymakers to not look through the supply shocks, but respond aggressively to re-anchor expectations.

6 Disagreement among households versus professionals and market prices

Our focus in this article is on how households disagree with each other.

Over the past decade, there have been many new surveys of executives within firms (see Candia, Coibion and Gorodnichenko, 2023, for a survey). While households are responsible for consumption and savings decisions, firms choose production and investment, so their expectations are potentially just as important. There are not as many datasets or as many research articles studying firm expectations as there is for households. Still, the existing ones all point in the same direction: both average expected inflation and disagreement about expectations of firms has the same statistical patterns as that of households. Firms also disagree significantly, over and beyond what some observable drivers of these differences are, and the disagreement varies over time and in response to shocks.

A much larger literature has studied the predictions of professional forecasters. Disagreement between professionals takes on a different dimension. Because their careers depend in part on the quality of their forecasts, and since they interact closely with each other, there is a strong strategic dimension to how and why they disagree. Moreover, they also interact closely with policymakers, so official communication strategies become an important suppressor of disagreement.

6.1 Households versus professionals via the Carroll regressions

A smaller literature has looked at the extent to which household and professionals disagree. Influential early work by Carroll (2003) found evidence that households followed

professionals in their response to news. After a shock or a large movement in inflation, disagreement between the two would rise, and then gradually fall as households catch up to professionals.

Table 5 revisits the Carroll (2003) evidence, using the MSC to take advantage of its longer sample, as well as the Survey of Professional Forecasters (SPF). We consider both the full sample as well as three sub-samples, and both the mean as well as the median to capture the first moment of expected inflation. The top panel shows Granger causality tests, while the bottom panel shows estimates of the Carroll (2003) regression. The original finding turns out to not be robust. There is no decisive evidence that professionals Granger cause households, or that the latter follow the former.[5]

Table 6 presents a version of the previous table but now using measures of disagreement, the standard deviation and skewness. When it comes to disagreement, households actually seem to Granger cause professionals. That is, when disagreement between households rises, then professionals start disagreeing afterwards. Looking at the regressions, for either measure, the disagreement between professionals almost never predicts disagreement between households.

6.2 Household expectations versus market risk and uncertainty

Figure 6 shows instead the time series of disagreement against two measures of inflation risk. The first is the standard deviation of 10-year ahead inflation options, from Hilscher, Raviv and Reis (2022). The second is the inflation risk premium estimated in a yield curve model and reported by the FRB Cleveland (Haubrich, Pennacchi and Ritchken, 2011). The correlation of disagreement with these measures of risk is very small, respectively, 0.14 and 0.07. Disagreement is not risk, at least in this statistical sense.

Returning to household surveys, the SCE asks households to report the uncertainty of their forecasts. The correlation between the time-series of the average uncertainty across households, and their disagreement is either 0.22 or 0.25, depending on which proxy of uncertainty—respectively the variance or the interquartile range of individual density forecasts—is chosen. Many articles have tried and failed to find a robust association between disagreement and uncertainty.

[5]For Carroll's original sample from 1981:q3 to 2000:q2, we get similar results as his, with a statistically significant coefficient of 0.39.

Table 5: Household expectations do not follow professionals

(a) Granger causality tests

		Full	1981-1994	1995-2012	2013-2023
SPF causes MSC	Mean	0.540	0.084	0.601	0.169
	Median	0.145	0.164	0.670	0.040
MSC causes SPF	Mean	0.039	0.082	0.050	0.006
	Median	0.206	0.046	0.094	0.000

(b) Carroll regressions

	Full sample		1978-1994		1995-2012		2013-2020	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SPF	0.080	0.021	0.462***	0.308**	0.232	0.140	0.344	0.428
	(0.068)	(0.059)	(0.088)	(0.123)	(0.214)	(0.144)	(0.466)	(0.355)
L.MSC	0.565***	0.336***	0.092	-0.047	0.512***	0.457***	0.889***	0.397***
	(0.090)	(0.091)	(0.168)	(0.155)	(0.111)	(0.122)	(0.149)	(0.132)
L.CPI	0.023	0.134***	-0.021	0.128*	-0.012	0.020	-0.081	0.094
	(0.051)	(0.041)	(0.051)	(0.074)	(0.123)	(0.061)	(0.121)	(0.084)
Constant	1.310***	1.621***	1.626***	1.646***	1.206**	1.227***	-0.154	0.756
	(0.284)	(0.234)	(0.437)	(0.502)	(0.493)	(0.438)	(0.985)	(0.662)
Observations	167	167	53	53	72	72	42	42
R^2	0.432	0.418	0.501	0.447	0.282	0.244	0.752	0.735

Note: Panel (a) reports p-values of Granger causality tests between MSC and SPF forecasts, using their respective mean or median, by subsample. Panel (b) shows results for the Carroll (2003) regression: $MSC_t(\pi_{t,t+4}) = \alpha_0 + \alpha_1 SPF_t(\pi_{t,t+4}) + \alpha_2 MSC_{t-1}(\pi_{t-1,t+3}) + \alpha_3 \pi_{t-5,t-1} + \epsilon_t$ where t is a quarter, due to the frequency of the SPF, and MSC_t and SPF_t refer to the household and professional mean expected inflation, respectively. The columns show estimates by subsample, for each of them first using the MSC and SPF mean forecasts, and second using the respective median forecasts. Newey-West standard errors with one lag in parentheses.

6.3 Looking ahead

Taken together, this evidence suggests that the class of models that see professionals as leading households are perhaps not so promising. Rather, it is more promising to see household disagreement, and disagreement between professionals and households, as two separate features of the data. Few articles have explored how to use the contrast of markets versus the people to learn more about the sources of disagreement and their effects, but this is a promising route for future research.

Table 6: Household disagreement is unrelated to professional disagreement

(a) Granger causality tests

		Full	1981-1994	1995-2012	2013-2023
SPF causes MSC	Standard deviation	0.030	0.049	0.372	0.016
	Kurtosis	0.893	0.391	0.312	0.429
MSC causes SPF	Standard deviation	0.041	0.982	0.308	0.000
	Kurtosis	0.250	0.107	0.398	0.015

(b) Carroll regressions

	Full sample		1978-1994		1995-2012		2013-2020	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SPF disagreement	0.278*	0.035	0.251	0.059	2.125***	0.150	0.454	-0.059
	(0.150)	(0.044)	(0.220)	(0.084)	(0.698)	(0.106)	(0.449)	(0.050)
L.MSC disagreement	0.617***	0.288***	-0.416***	0.116	-0.272	0.064	0.348*	-0.009
	(0.079)	(0.083)	(0.133)	(0.133)	(0.172)	(0.174)	(0.179)	(0.101)
L.CPI	0.070***	-0.043	0.227***	0.241***	0.172*	0.164	0.126	0.300***
	(0.025)	(0.027)	(0.056)	(0.034)	(0.096)	(0.104)	(0.092)	(0.041)
Constant	1.172***	0.607***	4.569***	2.279***	2.935***	2.499***	1.884***	2.469***
	(0.258)	(0.138)	(0.528)	(0.186)	(0.702)	(0.203)	(0.554)	(0.170)
Observations	167	167	53	53	72	72	42	42
R^2	0.630	0.107	0.407	0.395	0.262	0.128	0.569	0.697

Note: Panel (a) reports p-values of Granger causality tests between MSC and SPF disagreement, using standard deviation and kurtosis as respective measures, by subsample. Panel (b) reports results for Carroll regressions based on MSC and SPF disagreement. For each subsample, we first show results using the standard deviation, and second using skewness as a measure of disagreement. Newey-West standard errors with one lag in parentheses.

7 Conclusion

We collected some of the main facts about disagreement. These were established over two decades of much research spread over many articles, and we presented them here in terms of simple figures or correlations. There are, of course, many other statistics that the literature, or our own statistical work for this article, have found to not be particularly robust. Moreover, it is the nature of scientific progress that maybe in ten years' time, some of these facts have been reevaluated and abandoned, or that some others have replaced them. Finally, there are a few, more subtle facts, that have emerged in the context of

Figure 6: Disagreement versus risk premium

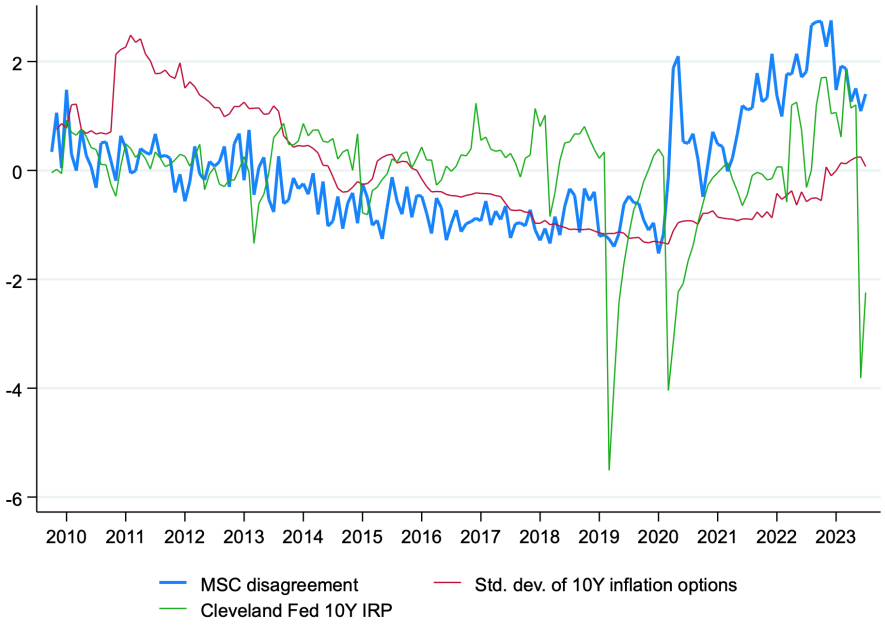

Note: Disagreement measured as the standard deviation of residualized 1-year inflation expectations in the MSC. Risk measured as either the standard deviation of expected inflation in 10-year ahead inflation options, or the inflation risk premium over 10 years in the Cleveland Fed model. To make them visually comparable, the three series are normalised by centering them around their average over the time period displayed, and scaling them using their standard deviation over the period.

testing for particular features of some models of disagreement that we failed to capture in a simple statistical relation.

With this in mind, we leave the list of facts here:

- On average, women, the poorer, married, and the less educated expect higher inflation.
- Until the mid to the end of the 1990s, the older expected lower inflation but less so every year. Between 2000 and 2020 there was no age effect, but since 2021, the old have started expecting higher inflation.
- The extent of disagreement is likely understated in the data because of don't know / not available answers.
- We can measure residual disagreement with one statistic of the second moment of

the distribution of answers, and one statistic for third (or higher) moments. Two moments are important, but which particular pair of measures is used seems less important.

- Disagreement measures carry information that is not in statistics for the first moment of inflation expectations. They significantly improve the forecasting performance of surveys for realised inflation.
- Measures of disagreement are especially informative at the tails of inflation realisations, and around large inflation disasters.
- Disagreement is systematically affected by identified shocks. Second moments rise for 1-2 quarters after either positive or negative shocks.
- The expectations of professionals are not informative about disagreement with or by households.
- Disagreement does not have a clear stable relation with risk in markets for inflation, or with self-reported uncertainty by households.

References

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2020. "Business-Cycle Anatomy." *American Economic Review*, 110(10): 3030–70.

Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry. 2020. "Imperfect Macroeconomic Expectations: Evidence and Theory." *NBER Macroeconomics Annual*, 35: 1–86.

Bauer, Michael D., and Eric T. Swanson. 2023. "An Alternative Explanation for the "Fed Information Effect"." *American Economic Review*, 113(3): 664–700.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2022. "Overreaction and Diagnostic Expectations in Macroeconomics." *Journal of Economic Perspectives*, 36(3): 223–244.

Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko. 2023. "Chapter 11 - The macroeconomic expectations of firms." In *Handbook of Economic Expectations*. , ed. Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, 321–353. Academic Press.

Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *The Quarterly Journal of Economics*, 118(1): 269–298.

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.

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar. 2018. "The Formation of Expectations, Inflation, and the Phillips Curve." *Journal of Economic Literature*, 56(4): 1447–91.

D'Acunto, Francesco, Ulrike Malmendier, and Michael Weber. 2023. "Chapter 5 - What do the data tell us about inflation expectations?" In *Handbook of Economic Expectations*. , ed. Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, 133–161. Academic Press.

Fofana, Salomé. 2024. "Empty vessels make the most noise: "don't know" answers in household expectations surveys." LSE manuscript.

Haubrich, Joseph G., George Pennacchi, and Peter Ritchken. 2011. "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps." *FRB Cleveland, Working Paper No. 11-07*.

Hilscher, Jens, Alon Raviv, and Ricardo Reis. 2022. "How Likely Is An Inflation Disaster?" *CEPR Discussion Paper 17224*.

Känzig, Diego R. 2021. "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements." *American Economic Review*, 111(4): 1092–1125.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2004. "Disagreement about Inflation Expectations." *NBER Macroeconomics Annual*, 18: 209–248.

Prescott, Edward C. 1977. "Should control theory be used for economic stabilization?" *Carnegie-Rochester Conference Series on Public Policy*, 7: 13–38.

Ramey, Valerie A., and Sarah Zubairy. 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data." *Journal of Political Economy*, 126(2): 850–901.

Reis, Ricardo. 2020. "The People versus the Markets: A Parsimonious Model of Inflation Expectations." *CFM discussion paper 2033*.

Reis, Ricardo. 2021. "Losing the Inflation Anchor." *Brookings Papers on Economic Activity*, 307–361.