The People versus the Markets: A Parsimonious Model of Inflation Expectations *

Ricardo Reis

LSE

October 2020

Abstract

Expected long-run inflation is sometimes inferred using market prices, other times using surveys. The discrepancy between the two measures has large business-cycle fluctuations, is systematically correlated with monetary policies, and is mostly driven by disagreement across and within groups of people. A parsimonious structural model of dispersed expectations and financial markets for inflation risk can make sense of the data, and provides estimates of the underlying expected inflation anchor. Applied to US data, the estimates suggest that inflation became gradually, but steadily, unanchored through the 2010s. A model of monetary policy suggests that a central bank that focuses on either measure of expected inflation will have less anchored inflation, partly because of their endogeneity with respect to policy itself.

*First draft: March 2019. Contact: r.a.reis@lse.ac.uk. I am grateful to comments from Gianluca Benigno, Rodrigo Guimaraes, Peter Kondor, Ian Martin, Harald Uhlig, and to seminar participants at Chicago Booth, FRB Cleveland, Oxford University, University of Essex, the 2019 EABCN conference in Warsaw, the 50th Konstanz seminar, the 2019 MacCalm conference, and the 2019 Rebuilding Macroeconomics conference, and to Adrien Couturier, Jose Alberto Fernandes, Maxi Guennewig, and Chao He for research assistance. This project has received funding from the European Union's Horizon 2020 research and innovation programme, INFL, under grant number No. GA: 682288.

1 Introduction

Expectations are the bedrock of dynamic economic models. Among them, inflation expectations attract special attention. Inflation is the most commonly asked variable to forecast in surveys, and the central questions in monetary economics—What explains inflation? How can central banks control it? What is the trade-off between inflation and real activity?—depend crucially on what private agents expect inflation will be. Most academic attention, however, has fallen on short-term (usually one-year ahead) inflation expectations by households and firms, while every central banker instead repeats that her focus is on anchoring long-term expectations, and these are usually measured using market prices.

This difference matters. At the end of 2010, with the global financial crisis behind but the Eurozone sovereign debt crisis raging, the median short-term EZ inflation expectation was 1.5%, the long-term median was 1.9%, and the market long-term forecast was 1.65%.[1] Four years later, the survey long-term forecast was almost the same, at 1.8%, but the market one had plunged to 0.71%. Policymakers focussed on the latter to justify the introduction of quantitative easing in the EZ, while stating that the former was the crucial one for success.[2] Since the start of 2019, long-term market-based measures have been persistently falling, in both the US and the EZ, while long-term survey-based measures have barely changed. Are inflation expectations anchored or not? Are extra unconventional monetary policies justified?

This paper makes three contributions. First, it proposes a new object for study: the business-cycle dynamics of the discrepancy between market and survey measures of long-run inflation expectations; for short, the *discrepancy*. This variable combines three characteristics. First, it concentrates on long-horizon inflation expectations, averaged over 5 or 10 years out. Second, it measures their fluctuations at a business-cycle frequency. Third, it focuses on why two legitimate ways to measure expectations, surveys of people, or prices of assets, give very different estimates. These three characteristics make this variable quite useful for economists trying to understand how expectations, in

[1] The survey measures are from the Survey of Professional Forecasters, at the 1 and 5 year horizon respectively, and the market measure is from inflation swaps at a 5-year horizon

[2] Reporting on the ECB's president, Mario Draghi's, speech at Jackson Hole which pointed to the 5-year-5-year forward market expectation fo inflation, the Financial Times wrote "Mr Draghi had highlighted the inflation swap rate ... never before August's Jackson Hole speech had a president of the ECB made such a clear link between its behavior and policy action." But in 2018, in Sintra, reflecting on the unconventional policy measures, Draghi showed the 5-year survey expected inflation to conclude success because:"What is key is that inflation expectations remain well anchored".

general, are formed. The focus on the long-horizon is useful from the perspective of behavioral theories of beliefs since differences in (i) when the expectations are measured, (ii) the timing of data releases on present inflation, (iii) anchoring effects on surveys, and (iv) asymmetric payoffs in financial prices, can all have a large impact on 1-year ahead expectations, and complicate the study of expectations alone, but are less relevant for long-run expectations. The focus on business-cycle frequencies raises macroeconomic questions, as opposed to the financial questions in low and high frequency studies of inflation risk premium and liquidity. Finally, the focus on the difference between market prices and surveys creates a challenge that typical papers on expectations do not meet, since it disciplines the model of subjective beliefs by not just the observed survey responses but also by the equilibrium prices that result.

The second novelty is an empirical, structural, *parsimonious model of expectations* for both people and markets that is flexible enough that it can be used as a measurement tool. An active literature over the last two decades has explored different behavioral theories of expectations and different models of market prices, testing some of their main qualitative predictions with survey and other data. This paper attempts to take the next step by writing a parsimonious model of expectations that builds on some of the leading theories, while being flexible. The model takes as input the cross-sectional moments from surveys, namely its first three moments, as well as the observed asset prices It can fit the discrepancy exactly, while providing insights into what drives it. It also yields a measure of the underlying rational expectation of inflation that is the anchor for inflation dynamics. Crucial for the estimates turns out to be the third moment of the distribution of expectations across people, the skewness.

The third and final part of the paper integrates the model of expectations into a simple (and standard) general-equilibrium model of inflation with interest-rate rules. This leads to three lessons. First, when the discrepancy is lower, this contributes to lower inflation. Second, to minimize the variance of inflation, monetary policy cannot over-rely on either survey or market measures of expected inflation. Third, policy itself affects the relative informativeness of the different measures.

Outline. Section 2 presents data to measure the discrepancy and establishes three facts about it. First, it has a large business-cycle component, robust across different measures. Second, it is systematically correlated with the state of monetary policy, and is significantly affected by policy shocks and policy regimes. Third, the discrepancy can be decomposed into three terms. The first is compensation for inflation risk required by a fictitious representative agent. The second is a difference between the subjective expectations of the public and that of market traders. The third is the difference between the subjective expectation of the marginal and the average trader. US data suggests that the two last terms drive most of the discrepancy. The section concludes by clarifying that the discrepancy is related, but is not the same, as an inflation risk premium.

Section 3 presents a parsimonious model of how people form expectations that tries to capture some of the main insights from the literature. The emphasis is on capturing the cross-sectional moments in the distribution of expectations, not their time-series dynamics. It provides estimates of its key moments for the Michigan survey of households.

Section 4 inserts these agents in an equilibrium model of dispersed information and financial markets, in the spirit of the classic work of Grossman and Stiglitz (1976). In the model, financial participants can observe market prices, which allows them to form more accurate forecasts than households. However, prices are contaminated by "noise" in the form a supply shock that can be interpreted as resulting from noise traders, animal spirits, or market fads. The novel focus is again on the cross-section of expectations across traders, and especially on the difference between the marginal and the average trader.

Section 5 applies the model to US data. The model takes as inputs the cross-sectional moments of surveys of expectations for households (and sometimes traders), as well as the observed asset prices. It fits the discrepancy perfectly, and produces two outputs. The first is the decomposition of the discrepancy between the two disagreement terms. This suggests that changes in the skewness of expectations has played an important role in accounting for the discrepancy. The second is an estimate of the fundamental rational-expectation best forecast of 5-year ahead inflation. According to the estimates, inflation has been significantly less anchored than surveys would suggest. In particular, expectations persistently drifted below 2% in the mid-2010s.

Section 6 inserts the model of expectations and financial markets into a model of monetary policy and inflation. The discrepancy affects monetary policy through two channels. First, because it provides a signal of fundamental expected inflation. Second, because it can transmit noise from financial markets into inflation volatility. Minimizing the variance of inflation requires adjusting with different weights to the two measures of expected inflation to reflect not just their relative informativeness, but especially the equilibrium feedback between this informativeness and the policy rule itself.

Finally, section 7 concludes with an application to the Euro-area.

Link to the literature. This paper contributes to four literatures in monetary economics.

First, a long literature has studied inflation expectations in economics. Most of that literature focuses on inflation over the next year though.[3] The focus of this paper is on long-run inflation expectations, and the measures used are of inflation on average over the next 5 or 10 years.[4] Another significant literature has focussed on disagreement on inflation expectations. Some focus on the second moment within surveys, others on disagreement across surveys, and yet others on difference in expected inflation across financial prices.[5] This paper instead focuses on disagreement between people and markets.[6]

The empirical model of household expectations combines elements from the literatures on imperfect information, overconfidence, learning, and sticky information, and builds more closely on Angeletos, Huo and Sastry (2020).[7] The model of financial markets builds more closely on Albagli, Hellwig and Tsyvinski (2013).[8] While these papers use the model to isolate a theoretical channel or to make predictions, here the model is a generalization that can be used as a measurement tool. As a result, the model is flexible enough to fit the expectations data exactly so that it can be used as a filter to detect the underlying fundamentals. Finally, the general-equilibrium model of inflation and monetary policy is a classic new Keynesian model, but set in continuous time building more closely on Reis (2019).[9]

2 A new variable of interest: the discrepancy

Let π_t, T denote the change in the log of the price level between dates t and T. The discrepancy ϕ_t is defined as:

$$\phi_t = \mathbb{E}_t^*(\pi_{t,T}) - \mathbb{E}_t^p(\pi_{t,T}). \tag{1}$$

[3]For recent surveys, see Coibion, Gorodnichenko and Kamdar (2018) applied to the Phillips curve, Clements (2019) applied to forecasting, and Kose et al. (2019) applied to emerging markets.

[4]Other papers that likewise focus on long-term expectations, include Beechey, Johannsen and Levin (2011), Garcia and Werner (2018),Eusepi et al. (2019), Moessner and Takáts (2020), and Yetman (2020).

[5]For an examples of each, see Mankiw, Reis and Wolfers (2004) and Andrade and Le Bihan (2013), citeCarroll and Coibion and Gorodnichenko (2012), and Fleckenstein, Longstaff and Lustig (2014) and Andreasen, Christensen and Riddell (2017), respectively.

[6]See also Ang, Bekaert and Wei (2007), citeFaustWright but from the perspective of a horse-race between the two, as opposed to explaining their difference.

[7]See, respectively, Woodford (2003*a*) and Sims (2003),Bordalo et al. (2020) and Guo and Wachter (2019), Malmendier and Nagel (2015) and Eusepi and Preston (2018), and Mankiw and Reis (2002) and Coibion and Gorodnichenko (2015).

[8]It is a version of the model in Grossman and Stiglitz (1980) surveyed in Vives (2008) or Veldkamp (2011).

[9]The classic analysis is in Woodford (2003*b*) and the focus on monetary policy responding to expected inflation is as in Clarida, Galí and Gertler (2000).

The first term, $\mathbb{E}_t^*(.)$, is the expectation implicit in asset prices at date t, sometimes also referred to as: the risk-adjusted expectation, the expectation under the risk-neutral measure, or the break-even inflation. The second term, $\mathbb{E}_t^p(.)$, is the subjective expectation by the public as reflected in answers to surveys.

Measuring the discrepancy requires defining the country to which it refers to, the frequency of t, the horizon T, the asset price to extract the market measure, and the survey to measure the people's beliefs. Since the focus on this paper is on long-term expectation, T should be at least 5 years out, and given the focus on business-cycle movements, the frequency should be at least quarterly.

This still leaves many possible measures. For its baseline, the paper uses data for the United States, at a monthly frequency, and for a 5-year horizon. Data for market prices comes from inflation swap contracts, which are available daily since June of 2007. The swap market is quite liquid, and is heavily used by pension funds to hedge long-run inflation risk. Data for expectations by the public is from the Michigan survey of households, starting in 1978, taken as the median over the around 500 responses every month that answer:"By about what percent per year do you expect prices to go up/down on the average, during the next 5 to 10 years?". The baseline series is demeaned and starts in 2010.

2.1 Fact 1: the discrepancy has significant business-cycle fluctuations

The baseline series is plotted in figure 1. It moves around significantly across quarters with a standard deviation of 0.30%. For comparison, the standard deviation of actual annual inflation during that period is 0.37%, and these are long-horizon expectations.

Figure 2 confirms this by plotting in the top panel the spectral density of the discrepancy, as well as inflation's. In the grey box are the usual business cycle frequencies of 6-32 quarters. The discrepancy series has a significant amount of power in this frequency: 49% of its variation is accounted by the business cycle. In this, it resembles actual inflation.

The middle panel investigates different series used to construct the discrepancy, plotting these series and their correlation with the baseline. Starting with the horizon, inflation swaps are also available for a 10-year horizon, so the alternative uses that for the market measure. A different market price comes from the difference between the yield in CPI inflation-indexed bonds (TIPS) and the yield in Treasury bonds.[10] A third alternative

[10] These are available for 5 or 10 year horizons, which are also the more liquid maturities. Reliable prices from liquid markets come from Gurkaynak, Swanson and Wright (2005) and start in 1999. A main caveat

Figure 1: The discrepancy over time

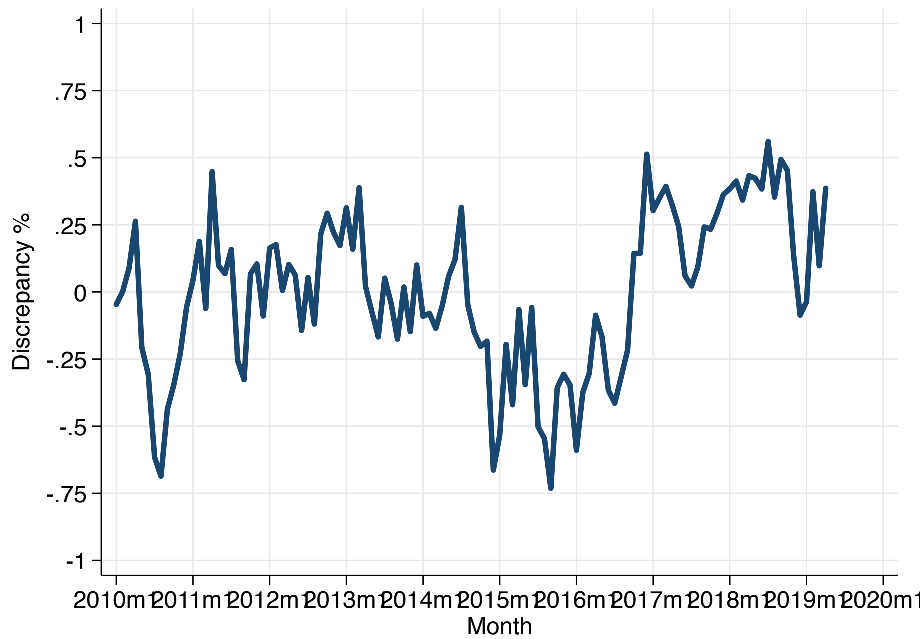

comes from the choice of series to measure the people's expectations. In the baseline, I took the median across respondents, whereas now I use the mean.[11] Still focusing on the people's expectation, the fourth series uses data from the Survey of Professional Forecasters, which asks about CPI inflation over the next 5 and over the next 10 years, in two separate questions. The survey is quarterly, but a monthly measure of its central tendency was calculated by the FRB Philadelphia.[12] The figure shows that all of these series are highly correlated with our baseline, and at a business-cycle frequency the coherence between the two series is very high. Where they differ significantly is in their level and their high-frequency variation.

The previous panel plotted the data since August of 2007, when the series for swaps are available. There is a clear large outlier at the end of 2008. The peak of the financial crisis significantly disrupted markets, and this especially affected the price of inflation-

of this measure is that for a significant part of the sample inflation was near zero, yet TIPS are not indexed to negative inflation, so that changes in this measure reflect changes in the option value embedded in this payoff, rather than changes in expected inflation.

[11] The median is a more reliable measure of central tendencies because a significant fraction of respondents provide absurdly large or small answers.

[12] This series includes business people and market participants: section 5 explores that distinction.

Figure 2: The business-cycle variation of the discrepancy

(a) Spectral density excluding frequency zero

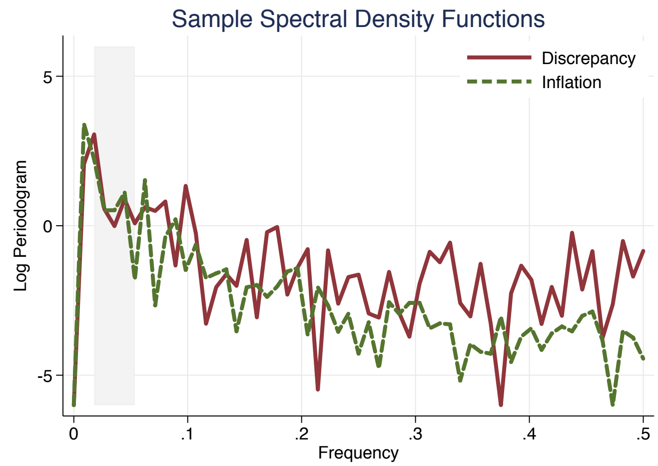

(b) Alternative measures

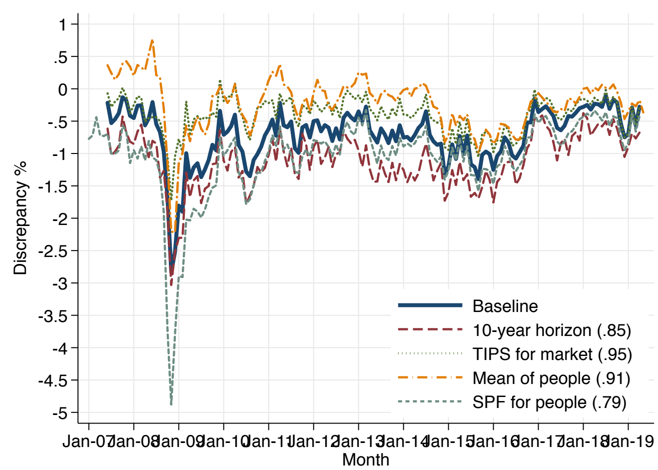

(c) Longer sample

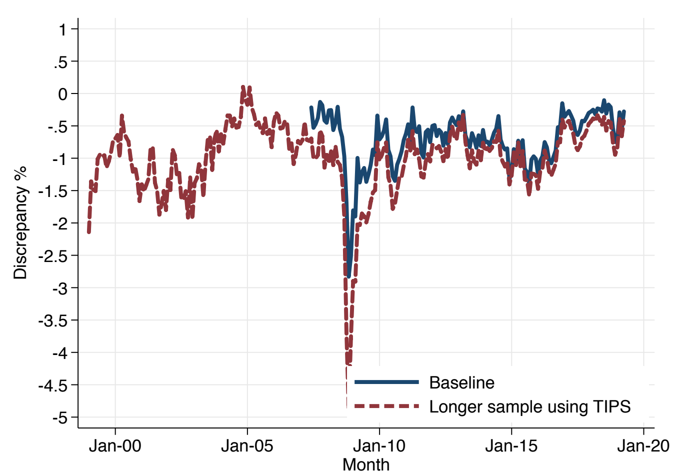

indexed bonds. Insofar as the liquidity problems in that market, as well as the Fed's differential response in the nominal bonds versus the indexed-bonds market, distort the measures of expected inflation, in the baseline series, I start the sample in 2010, when these markets seemed to be operating normally again.To ensure 2008 was indeed exceptional, and to try to maximize the length of the sample, the bottom panel plots a longer alternative series, that uses TIPS for the market measures, since the data is available since 1999. The movements in the discrepancy between 2008 and 2009 are indeed exceptional. Moreover, the movements of the discrepancy before 2007 resemble those in our baseline sample. The main difference is that the series is slightly more volatile, as so was inflation.

More generally, the fluctuations in figure 1 are persistent, not just the results of high-frequency movements. Market liquidity frictions are important for daily movements, but they do not seem to interfere with the monthly or quarterly movements in the series. The correlation across series from different markets are very high at business-cycle frequencies even if they are much lower at a daily frequency.

2.2 Fact 2: the discrepancy is systematically related to monetary policy

The top panel of table 1 shows estimates of a regression of the discrepancy on a summary indicator of the stance of monetary policy: the 2-year yield. Tight monetary policy is strongly negatively correlated with a fall in the discrepancy. This is consistent with market expected inflation falling by more than the people's expectation. The regression controls for inflation and its squared change.

The bottom panel instead regresses the discrepancy on a shock to monetary policy, constructed form movements in the Federal Funds forward market, from Nakamura and Steinsson (2018). Tighter unexpected changes in policy likewise have a negative impact on the discrepancy.

Figure 1 extend the analysis across countries. For the Eurozone, there are no indexed bonds at the eurozone level but there are liquid inflation swaps, both at the 5 and the 10 year horizons, with daily data since 2004. For the people's expectations, an imperfect measure comes from the ECB's survey of professional forecasters, which asks respondents every quarter since 2004 for their inflation expectations over the long run. The figure shows the discrepancy using the 5-year swaps and the median of respondents to be closer to the US numbers. The United Kingdom has active and liquid markets for both inflation swaps and indexed government bonds at both the 5 and the 10 year horizon, the former

Table 1: The proximate determinants of the discrepancy

	Determinants	Policy shocks
	(1)	(2)
2-year yield	0.149***	
	(0.0273)	
Inflation	0.177***	
	(0.0233)	
Squared change inflation	-0.200	
	(0.159)	
Monetary shocks		-6.717*
		(3.884)
Observations	111	43
R-squared	0.512	0.068

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: xxx.

since 2007 and the latter since 1997. For the people's expectations, I use the median of the Bank of England's 3-year ahead Survey of Economic Forecasters, which is quarterly and collected since 1998. Finally, for Japan, there are market prices for 5-year and 10-year indexed bonds and swap contracts since 2009. Subjective expectations comes from Consensus economics, which has quarterly data on 10-year and 5-year ahead inflation expectations since 1989.[13]

All series are in figure 3 together with a quarterly version of our baseline series. While it is well-known that actual inflation is quite correlated across these countries, the correlation of the discrepancies is negative (for the EA) or low (for Japan). Policies in the US and the EZ were different during this time, and even more different in the other countries. Movements in the Japanese series match some of the changes in the policy regime there, from the introduction of qualitative and quantitative easing to yield curve control. In the UK, the end of sample shows a rising discrepancy following the Brexit referendum. In short, different policy regimes across these countries are correlated with different behaviors of the discrepancy.

[13] A caveat of all of these subjective measures is that they mix market participants and industry observers, unlike our baseline series from the Michigan survey, which surveys households directly. Section 5 discusses this further.

Figure 3: The discrepancy across countries

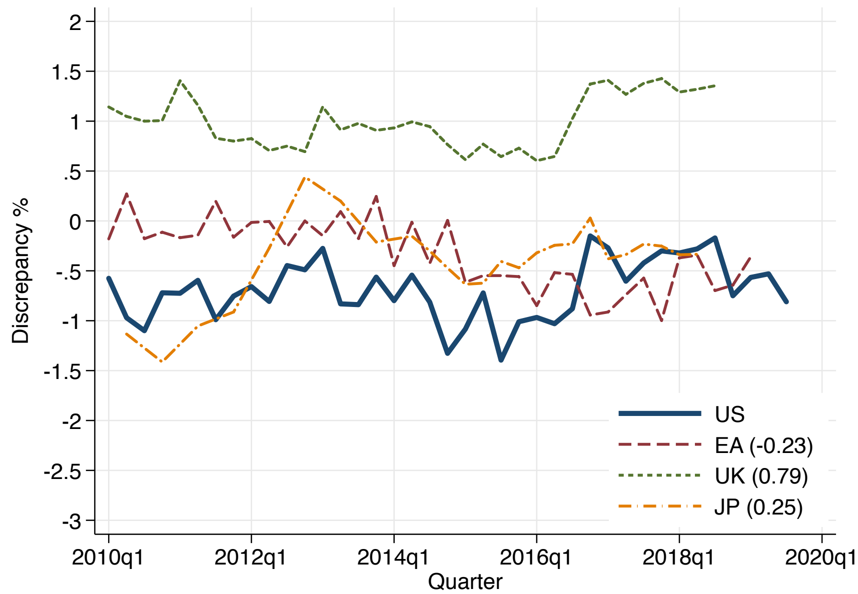

2.3 Fact 3: a decomposition points to disagreement

Define $\mathbb{E}_t^b(.)$ and $\mathbb{E}_t^m(.)$ as the subjective expectations of bond-traders, respectively of the average and of the marginal trader. These can be different from each other, and from the subjective expectations of the households answering surveys $\mathbb{E}_t^p(.)$, because each individual forms her own expectations using her information and her beliefs. One can then decompose the discrepancy into three terms:

$$\phi_t = \underbrace{\mathbb{E}_t^b(\pi_{t,T}) - \mathbb{E}_t^p(\pi_{t,T})}_{\text{disagreement across}} + \underbrace{\mathbb{E}_t^m(\pi_{t,T}) - \mathbb{E}_t^b(\pi_{t,T})}_{\text{disagreement within}} + \underbrace{\mathbb{E}_t^*(\pi_{t,T}) - \mathbb{E}_t^m(\pi_{t,T})}_{\text{risk compensation}} \tag{2}$$

The first term captures disagreement across types of agents, namely market traders and the public. We might expect the first group to be better informed about inflation, since it is part of their job, although they may also be more susceptible to fads, conformity biases, and short-term thinking. The second term is also about disagreement, but now between the marginal trader in the market, whose views prices reflect, and the average trader. Heterogeneity of views in markets is what causes trade in the first place. At the same time, imperfections in financial markets arising from liquidity shocks, changes in risk-taking capacity across agents, and shifts in heterogeneous levels of confidence together with short sales constraints, would show up here. Finally, the third term is the pure compensation for inflation risk that this individual marginal trader will require if she is risk averse.

There are empirical proxies for the first and the third term. Starting with the former, the FRB New York surveys about 50 financial market dealers eight times per year on their expected inflation over the next 5 years.[14] These are dealers in Treasury markets trading inflation risk, therefore matching closely the $\mathbb{E}_t^b(.)$ concept.[15] Turning to the latter, the arguments in MARTIN and WAGNER (2019) for equity risk applied to the market for inflation risk, suggest that the risk-neutral variance of inflation measures compensation for inflation risk. This can be measured using options for inflation as in Hilscher, Raviv and Reis (2014).

Figure **??** plots the discrepancy and these two measures, while table **??** regresses the discrepancy onto the two components. Both are statistically significant at the 5% level. But, while the disagreement between the beliefs of traders and households can account for a significant part of the time-series variation in the discrepancy, risk compensation instead quantitatively accounts for very little.

This regression treats the within disagreement term as a residual. It is hard to believe that this omitted term is orthogonal to the other two. At the same time, observing who is the marginal trader is a hopeless task. Instead, the next section uses a model of financial markets with heterogeneous beliefs that can provide estimates of the within-disagreement term as function of observables.

2.4 Inflation risk premia

Before proceeding, it is worth clarifying how risk premia fit into the discussion. Let $\pi_{t,T}^e$ denote the *rational expectation* of inflation. Then, the inflation risk premium is defined as:

$$\text{IRP}_t \equiv \mathbb{E}_t^*(\pi_{t,T}) - \pi_{t,T}^e = \underbrace{\phi_t}_{\text{discrepancy}} + \underbrace{\mathbb{E}_t^p(\pi_{t,T}) - \pi_{t,T}^e}_{\text{irrationality}} \tag{3}$$

[14]The survey also includes quartiles of the distribution, and an expectation of inflation starting in 5 years' time for the next 5 years. However, these extra series add little additional information.

[15]An inferior alternative is the Blue Chip survey, which surveys mostly financial market participants, but across many different financial marktes.

Figure 4: The decomposition of the US discrepancy

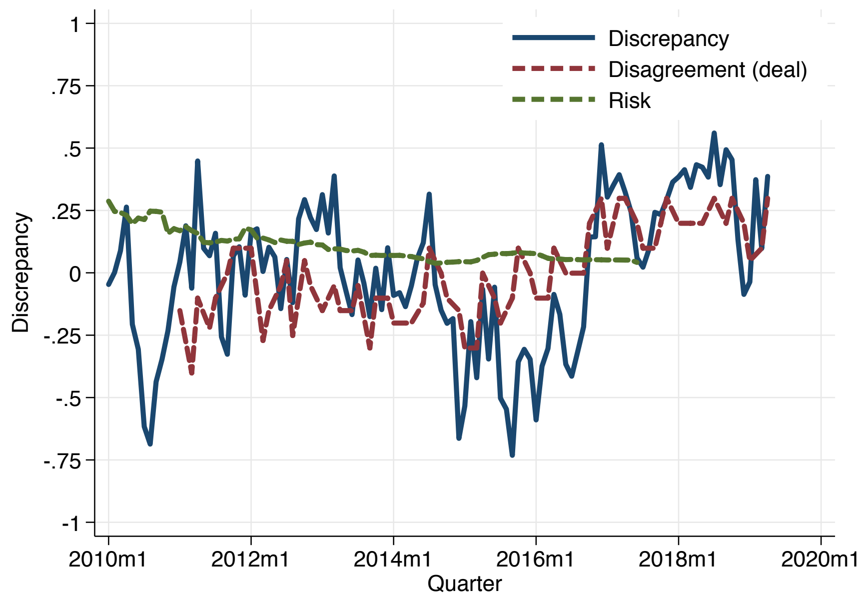

As the equality shows, this is not the same as the discrepancy. Only if there are homogenous rational expectations would they be the same, in which case the inflation risk premium would be equal to the risk compensation in the decomposition in figure **??**. Not only has the hypothesis of rational expectation been strongly rejected many times using our Michigan survey data (and many other datasets), but the pure risk compensation accounted for a small part of the variation in the discrepancy.[16]

Figure **??** is instead consistent with the many models of financial risk that rely on disagreement between marginal and average trader as opposed to individual-level risk aversion. Departing from the benchmark of Milgrom and Stokey (1982), a voluminous literature has considered imperfect information, behavioral biases, liquidity shocks, short sales constraints, among many others. Section 5 will derive estimates of $\pi^e_{t,T}$ with which one could test those models.

[16]See also Gürkaynak, Sack and Wright (2010).

3 A parsimonious model of the people's expectations

The anchor of people's model of expectation is the rational expectations fundamental $\pi^e_{t,T}$. The only assumption made on it is that it is unbiased and leads to serially uncorrelated forecast errors: $\mathbb{E}_t(\pi_{t,T}) = \pi^e_{t,T}$, and $\mathbb{E}_t(\pi^e_{t,T}(\pi_{t,T} - \pi^e_{t,T})) = 0$ if $\mathbb{E}_t(.)$ is a statistically optimal expectation operator. People do not know what this is. Indexing a household by i, and dropping the time subscript (t, T) that would otherwise appear everywhere, its individual expectation is denoted by v^h.

3.1 Four behavioral features of expectations

The first property of household expectations is *incomplete information*. Starting from some common prior π^*, each household receives a idiosyncratic noisy signal drawn from a distribution centered at π^e but with variance σ^2. This signal induces a dispersion of expectations across people as a result of different signals being drawn. At the same time, because households know their signals are noisy, this imperfect information leads average inflation forecasts across households to under-react to news, as a large literature has found (e.g. Coibion and Gorodnichenko, 2015). I impose a simplifying assumption: that the distribution of signals is normally distributed.[17]

The second property of expectations is *over-confidence*. Households behave as if their signals are more precise than what they really are. Therefore, if the response of their individual expectation to a signal is given by θ, this can be high even when $1/\sigma^2$ is low, and could even be above 1, something a rational imperfectly informed agent would never choose to do. This matches the also extensive literature that individuals, in the cross-section, over-react to signals (e.g. Bordalo et al., 2020). The simplifying assumption in modeling this is to assume that the relation is (at least approximately) linear.[18]

Third, households *learn from experience*. Namely, they suffer from a bias z^c in their beliefs, that arises from past experiences that left a scar. Empirically, a growing literature has found that experiences of inflation, especially in younger ages, account for a significant share of the disagreement observed across age cohorts (e.g. Malmendier and Nagel, 2015). The simplification here is to assume that this bias is linear in the age of the cohort,

[17] This would be optimal if households had a quadratic objective function, the prior was normal, and they suffered from rational inattention citepSimsJME03.

[18] If the overconfidence shows up as the perceived variance of signals being smaller than the actual variance, then with a normal signal, the linearity follows from the properties of the conditional normal distribution.

so if $c = 0, 1, 2...$ denotes the cohort, then $z^c = c\pi^z$ for some constant π^z. This matches the fact that in the US, inflation has trended down since the 1980s, so that as one moves to older cohorts, so c rises, the bias is higher.[19]

Finally, cohorts update their expectation bias according to *sticky information*. In particular, in the population, a small share of household of any given cohort every year updates its information and eliminates the bias it had from the past. This leads to a slow dissemination of information, and to disagreement that evolves endogenously with the shocks to inflation, as observed in the data (e.g. Mankiw, Reis and Wolfers, 2004). I make the usual simplification in this literature that this process evolves according to a memoryless Poisson process, so that at any date in time there is a fraction $\lambda(1-\lambda)^c$ of people that have a bias according to cohort c.[20]

These four properties capture the main features of a significant share of the research on modeling expectations over the last two decades. Together with the simplifying assumptions introduced along the way, they give rise to the following parsimonious empirical model of expectations, where the horizon is T for all of them, but the time subscript t is re-introduced:

$$v_t^h = c_t \pi_t^z + \pi_t^* + \theta_t (e_t^i + \pi_t^e - \pi_t^*) \tag{4}$$

$$e_t^h | \pi_t^e \sim N(0, \sigma_t^2) \;\text{ and }\; c_t \sim Exp(\lambda_t) \tag{5}$$

The model has four parameters capturing the strength of the four behavioral mechanisms described above: σ_t^2 on how disperse and imperfect is information, θ_t on how over-confident people are, π_t^z on how large are the scars of past high inflation, and λ_t on the stickiness of information on updating biases. Conditional on the two unobservables, the prior π_t^* and the actual fundamental π_t^e, the model predicts that the observable individual expectation v_t^i has a distribution $F_t(.)$ that is an exponentially-modified Gaussian. Its first three moments are non-zero. With data on the average, the standard deviation (or interquartile range), and the skewness of inflation expectations in the household survey, the model identifies three parameters: $\sigma_t^2, \theta_t, \lambda_t / \pi_t^z$.

[19] If inflation followed a random walk with white noise, that happened to have mostly negative permanent shocks in the last 40 years, then a least squares learning formula that is infrequently updated would generate this process endogenously.

[20] If agents did not update XXX

3.2 Application to the Michigan survey data

Figure 5 fits this model to the data on answers to long-term inflation expectations in the Michigan survey since 1990. The top panel shows the shape of $F_t(.)$ when using the average over time of the mean, standard deviation, and skewness in the data. The distribution looks like a normal distribution, but it has a fatter tail on the right reflecting the presence of the bias due to experience.

The middle panel shows the three moments in the data over time. As has been noted before, the average long-run inflation expectation fell throughout the 1990s, before stabilizing around 3%. Some interpret the fall as the success of expectations anchoring, while others note that the fact that it stabilizes a full percentage point above the target fo the Federal Reserve as a failure. Less appreciated is that the standard deviation of expectations was roughly constant throughout, but then started falling around 2014. At the same time, the skewness also fell, even if only slightly, after it had risen in the 2008-10 period.

The bottom panel shows the implied σ_t^2 and λ_t/π_t^z to match these moments using the model. According to the model, after a spike in 2008-10, there is a visible trend downwards in the dispersion of information that, by the end of the decade, is persistently at e very low level. At the same time, the bias arising from learning from experience has fallen significantly from 2014 onwards, as either the size of that bias π_t^z has fallen, or the frequency with which agents update and reduce it λ_t has risen.[21]

With these parameters varying over time, the model fits the first three moments of the expectations data exactly. At the same time, there are three over-identification tests of whether the model fits the overall data well.

The first is that the model imposes that the two series on the bottom panel of the figure have to be positive. In the expectations data, this turn out to be so at every single data. This was not guaranteed, since if the skewness was much higher at any one date, the implied σ_t^2 at that date would have been negative. The fact that it never is supports the setup of the model.[22]

Second, the model predicts that kurtosis (and higher-order moments) should be zero. In the data, on average that is approximately the case over the last thirty years. Focussing on the first three moments, as the model does, is a good statistical approximation.

Third, the model predicts that the following combination of data statistics should be

[21] The model does not identify the remaining parameter, θ_t, independently of the fundamental π_t^e.

[22] As a validation, in the Survey of Professional Forecasters, this is not the case.

equal to the average prior over time:

$$\mu_t \equiv Mean_t - StDev_t(0.5Skew_t)^{1/3} \quad \Rightarrow \quad \lim_{T\to\infty} \frac{\sum_t \mu_t}{T} = \pi^* \tag{6}$$

Even tough the average of $Mean_t$ is 3.59%, the left-hand side of this expression is 2.26%. Post 2010, the average long-run inflation expectation in the Michigan survey is 3.11%, but according to the model π^* will have been only 1.92%. That is, without any free parameters, the model can make sense of the high survey expectations reported in the surveys as being consistent with a plausible value for the underlying long-run average of implied expected inflation.

To be clear, the fit of this parsimonious empirical model does not by itself provide a test of each of the four behavioral assumptions and associated simplifications. What it shows is that the model is a good measurement tool with which to filter the noisy expectations data and that is broadly consistent with the major insights from micro-founded models.

4 A parsimonious model of financial market expectations

Consider a continuum of traders indexed by i in the unit interval. They are drawn from the population of households, so they enter markets with a prior expectation of inflation that is a draw v^i from the posterior $F(.)$ distribution (again omitting time subscripts). Traders have superior information through because they observe an extra piece of information: the market price q of a nominal bond that next period gives 1 nominal unit next period. Prices are not fully revealing but they reveal some information on what the fundamental is, which is captured by the distribution $g(q|\pi^e)$. Being Bayesian, the traders' posterior for inflation therefore is:

$$p(\pi^e|v^i, q) \propto \frac{g(q|\pi^e)f(\pi^e|v^i)}{\int g(q|\pi^e)f(\pi^e|v^i)d\pi^e} \tag{7}$$

4.1 Individual behavior

The goal of each trader is to maximize the expected discounted profits by choosing bond-holdings b^i:

$$\max \int \left[m(\pi)e^{-\pi} - q\right] b^i p(\pi^e|v^i, q)d\pi^e \tag{8}$$

Figure 5: The model of households expectations and the data

(a) Representative distribution $F(.)$

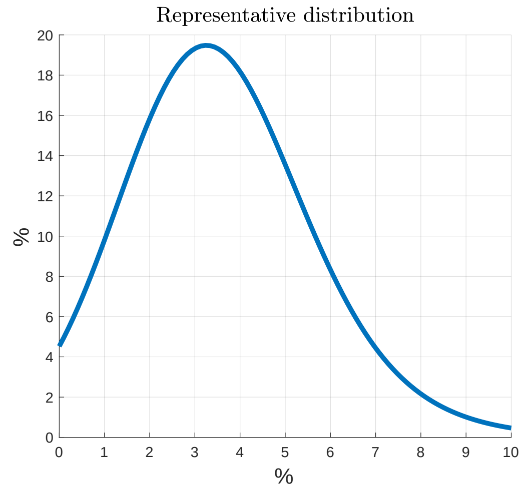

(b) Three first moments of Michigan household survey

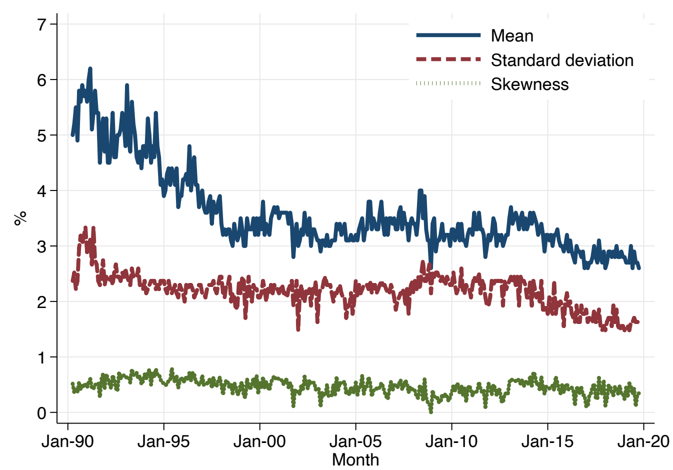

(c) Parameters over time

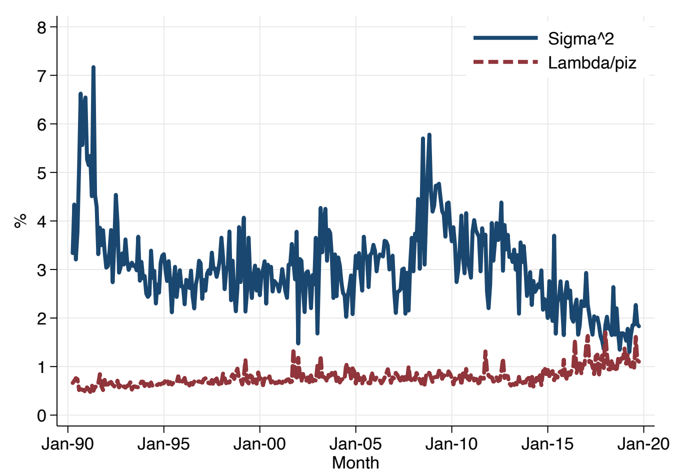

where $m(\pi)$ is the stochastic discount factor, that depends on inflation insofar as there is aversion to inflation risk. I impose two simplifying assumptions on this problem. First, each trader has some positive wealth w^i, and cannot borrow or short the bond, so that $b_i \in [0, w_i]$. Second, the stochastic discount factor is common across all traders, given to them by the head household to whom they return their profits at the end. Since each of them is infinitesimal, then the payoff from investing in the bond is: $y(\pi^e) \equiv \mathbb{E}\left[m(\pi)e^{-\pi}|\pi^e\right]$ which does not depend on each individual trader. While there is risk aversion captured in the curvature of $y(\pi^e)$, each trader individually behave as if she was risk neutral.

These two assumptions, together with the fact that the exponentially-modified Gaussian distribution $F(.)$ has a monotone likelihood ratio, imply a simple solution to the problem. As long as the posterior after observing the price is not degenerate, then traders that expect high inflation at the start, $v^i > v^*$ do not want to hold the nominal bond $b^i = 0$; traders that expect low inflation $v^i < v^*$ invest all their wealth in the bond $b^i = w^i$, and the marginal trader who is just indifferent defines the threshold v^* as:

$$\int y(\pi^e)p(\pi^e|v^*, q)d\pi^e = q \tag{9}$$

4.2 Market clearing and noise

Given a total supply of bonds B, the standard market clearing condition is $\int_0^1 b^i di = B$. Given the threshold strategy for investment just derived, then letting w be average wealth across agents, the market clearing condition becomes the simple condition:

$$F(v^*|\pi^e) = B/w \tag{10}$$

If wealth and the supply of bonds were known, then the market price would perfectly reveal to every trader what the actual fundamental expected inflation π^e is. Their posteriors $p(\pi^e|v^i, q)$ would all be identical and degenerate (Grossman and Stiglitz, 1980), and they would be indifferent between trading the bond or not (Milgrom and Stokey, 1982). A long literature has broken this efficiency of markets by arguing for different behavioral assumptions on how traders fail to solve this inference properly, different market structure assumptions that lead to noise traders and time-varying frictions that map into shocks to w, and shocks to the supply of bonds or to their liquidity mapping into B. A parsimonious model of these imperfections is to assume that $\omega \equiv B/W$ is a random variable that the traders cannot observe. It is then a source of noise contaminating asset prices.

The range of this variable ω is the unit interval. I make the simplifying assumption that this follows a symmetric Beta distribution with parameter β. When β approaches 1, the distribution approaches the uniform, while when it grows to infinity, it becomes concentrated in a point mass at $1/2$. For any finite value, it implies that market prices will reflect both the fundamental as well as this noise, matching the high volatility observed in the prices of inflation swaps.

4.3 Equilibrium

A perfect Bayesian equilibrium in this financial market is a price function $q(\pi^e, \omega)$ such that the beliefs if each trader follow Bayes rule in equation (7), their investment choices satisfy the optimal threshold strategy in equation (9), and the bond market clears in equation (10). Following Albagli, Hellwig and Tsyvinski (2013) though, inspecting the three equations that determine equilibrium, it is clear that the threshold v^* is a sufficient statistic for the pair (π^e, ω) to determine prices. Letting $q = Q(v^*)$, note that only v^* appears in the belief function in equation (7) and in the investment threshold function in equation (9). As for the market clearing equation (10), one can invert the $F(.)$ distribution function to obtain: $v^* = \pi^e + F^{-1}(\omega)$, using the fact that the signals in the survey distribution are centered around the fundamental π^e. Given the Beta distribution for ω, this equation produces a distribution function for the threshold that is consistent with markets clearing: $G(v^*|\pi^e)$.

The solution of the model is therefore the solution of the equation:[23]

$$Q(v^*) = \frac{\int y(\pi^e) g(v^* - \pi^e) f(v^* - \pi^e) d\pi^e}{\int g(v^* - \pi^e) f(v^* - \pi^e) d\pi^e}. \tag{11}$$

Given a solution for $Q(.)$, recovering the solution for the original model follows from: $q(\pi^e, \omega) = Q(\pi^e + F^{-1}(\omega))$.

The model and its solution have a few properties that make it useful to understand market-implied inflation expectations as they relate to survey expectations. First, an increase in the prior mean π^* shifts the predicted asset price q one-to-one. Therefore the prior, which is a free parameter of the model, anchors the results in a transparent way. The other free parameter is β, which determines how informative asset prices are, again transparently accommodating the views of different researchers.

[23] There was a slight abuse of notation in using $g(q|\pi^e)$ in equation (7) and then replacing it $g(v^*|\pi^e)$.

Second, the asset price is monotonic in both the underlying fundamental π^e and noise ω. Therefore, higher expected inflation according to markets may either reflect signal or noise. It is easy to show that when ω approaches its lower (upper) limit of 0 (1), then the threshold v^* approaches minus (plus) infinity, and likewise for the price. Therefore, for any observed market price, the model can always make sense of it in terms of a realization of noise, even if very unlikely. Like the model of expectations, this model of markets, can always fit the data.

Third, the model has a natural mapping to the empirical objects in section 2. For starters, the market expectation of inflation is $\mathbb{E}^*(\pi^e) = 1/q(\pi^e, \sigma) - (1+r)$ where r is a measure of the safe real rate. Next, replacing the payoff function with a linear approximation: $y(\pi^e) = \bar{y} + \pi^e$, then $\bar{y}$ is the *risk compensation* under homogeneous rational expectations. Adjusted for risk, the asset price is the expectation of inflation according to the marginal trader, or $\mathbb{E}^m(.) = \mathbb{E}(.|v^*, q)$. In turn, letting the median of the $G(v^*|\pi^e)$ distribution be v^{med}, then the median in a survey of traders would be:

$$\mathbb{E}^b(\pi) = \frac{\int \pi^e g(v^* - \pi^e) f(v^{med} - \pi^e) d\pi^e}{\int g(v^* - \pi^e) f(v^{med} - \pi^e) d\pi^e}. \tag{12}$$

The *disagreement within* term is then the disagreement between the trader with prior expectation (before observing market prices) v^* and the trader with prior expectation v^{med}. Finally, the *disagreement across* is due to traders observing prices, while households do not, and so the former having a posterior distribution $P(.)$ while the latter have a distribution $F(.)$ for inflation. Therefore, the model fits perfectly into the decomposition of the discrepancy in section 2.3. Because the expectation of the marginal trader is unobservable with data alone, and v^* is an equilibrium object, the model is able to complete the decomposition.

Finally, and combining all of these properties, adding financial markets allows the parsimonious empirical model to now also fit market-based expectations of inflation, and to explain the discrepancy in terms of its decomposition into underlying terms. At each date, t, the parameters of the model are π^*_t, anchoring the fundamental, and β on the volatility of the noise. The data are the three moments of the survey distribution of households $F_t(.)$, the median expectation of traders, and the market-implied expectation in the price q_t. They identify the five unknowns, three from the behavioral model of expectations $\theta_t, \sigma^2_t, \lambda_t/\pi^z_t$, one from the model of financial markets ω_t, and finally the fundamental expected inflation anchor π^e_t. The model therefore provides a measurement tool—a computational filter—through which to measure fundamental expected inflation, and thus to assess whether long-run inflation expectations are truly anchored.

5 Measuring US long-run inflation expectations

There are many expected inflations in the model. Markets, people, and traders all deviate from the fundamental expected inflation π^e, but they are all anchored by it. Moreover, in a general-equilibrium model (like the next section will illustrate) this fundamental is the equilibrium object that may be anchored or not depending on policy. Therefore, answering the question in the introduction—Are inflation expectations anchored?—boils down to estimating π^e. At the same time, estimating the marginal trader v^* provides for the decomposition of disagreement within and across that is driving the discrepancy.

5.1 The model's mechanics

The model has only two parameters. The prior for inflation expectations is set at a constant $\pi^*_t = 2\%$. This was the announced target of the Federal Reserve, and assuming a prior on it, if anything biases the results towards finding anchored expectations at this target. The appendix reports an alternative where instead one year's prior is taken to be the previous year's estimate, that is: $\pi^*_t = \pi^e_{t-1}$. The estimates turn out to be quite similar. As for β_t I set it equal to 2, so the prior on market noise is reasonably flat. Setting it instead at 1.5 or 4 gave almost indistinguishable results.

Figure 6 shows how the model's predictions for expected inflation and the marginal trader vary with changes in the two sources of disagreement, by varying the market expected inflation and the survey expected inflation, while keeping the trader's average belief fixed. As disagreement across becomes more negative, so as people expect higher inflation while trader's expectation is unchanged, the model predicts that fundamental expected inflation must be higher. In the other direction, if we observe trader's expectation falling relative to the people's, the model will signal a fall in fundamental expected inflation.

Further, as disagreement within becomes more negative, so the asset price falls further below the average trade's belief and the discrepancy rises, then the bottom panel shows that the marginal trader moves towards the left tail away from 0.5. The model partly interprets this as a result of noise, so the fall in expected inflation is muted but still posi-

Figure 6: The model at work

(a) The model's predicted expected inflation

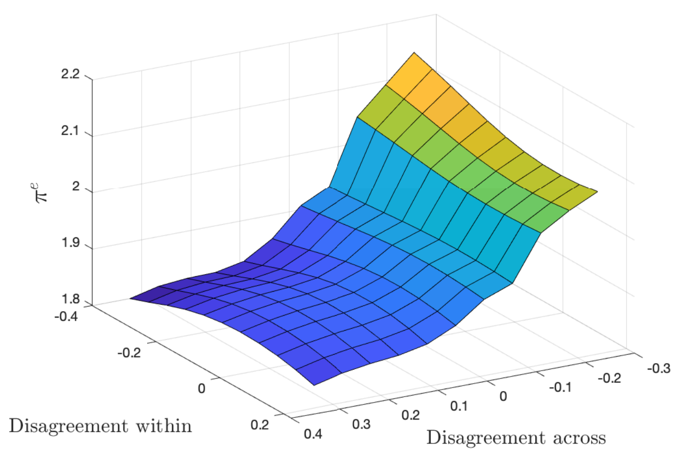

(b) The model's predicted marginal trader

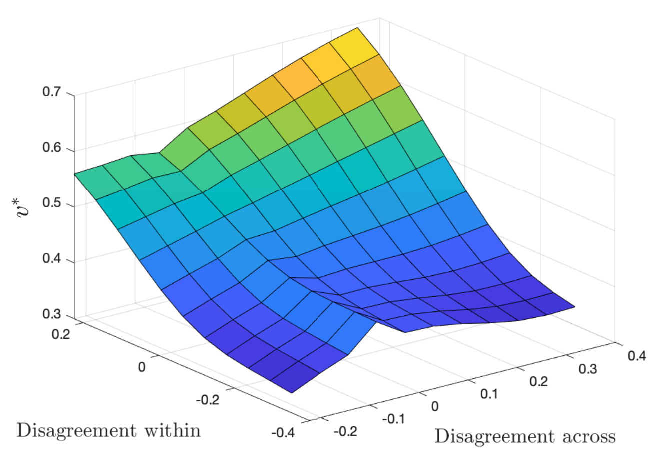

tive. If instead disagreement within is positive, while disagreement across was negative, that is if the average household expects higher inflation than the average trader, but the market-prices are also above those of traders, the model interprets this partly as a result of noise pushing prices up, but still increases its estimate of expected inflation.

Quantitatively, ceteris paribus, higher people's expectations have a stronger impact on estimates of the fundamental than higher market expectations because of the role of market noise.

5.2 Estimates post-2011

The baseline discrepancy series starts in 2011, and uses swap contracts to measure market-implied expectations, the implied variance in options contracts to measure risk compensation, the Michigan survey to measure the first three moments of the distribution of household's expectations, and finally the FRB New York survey of dealers to measure the expectations of the median trader. The latter series is not available every month, since the survey takes place only 8 times per year. Moreover, as figure 5 showed, the survey data is quite noisy month-to-month. Therefore, I average each of the data input within each of the nine years of the sample, and solve the model ten times, once per year.[24]

The top panel of figure **??** shows the implied fundamental expected inflation. This is the unique solution to the model, that is the only series consistent with the the data from the people, the markets and the traders. It shows a clear downward trend from 2014 onwards. Starting from 2%, long-run expected inflation has fallen to around 1.8% by 2019.

The middle panel shows who is the marginal trader at each date in time. Strikingly, between 2013 and 2016, the marginal trader goes from being roughly the same as the average trader to being the trader on the 20th percentile left tail. Two factors seem to be driving this result. First, the skewness of the survey increases during this time. Second, the market-implied expectation falls. The model interprets these facts as there being more pessimists in the population so that they start dominating the market.

The bottom panel shows the decomposition between the two sources of disagreement, which confirms this interpretation. It is the disagreement within market traders

[24]The appendix reports the results of solving the model for every month that the data is available. The annual average of the estimates is almost identical to the estimates using annually averaged data. Moreover, the month-to-month fluctuations are erratic and driven by the jumps on the moments of the survey distribution.

Figure 7: Model estimates for baseline US sample

(a) Fundamental long-run inflation expectations

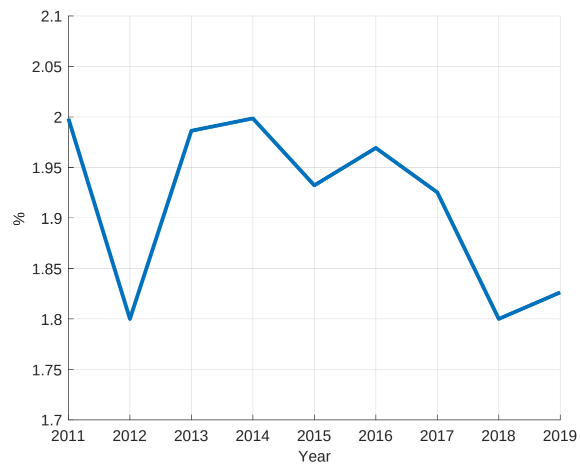

(b) The percentile of the marginal trader

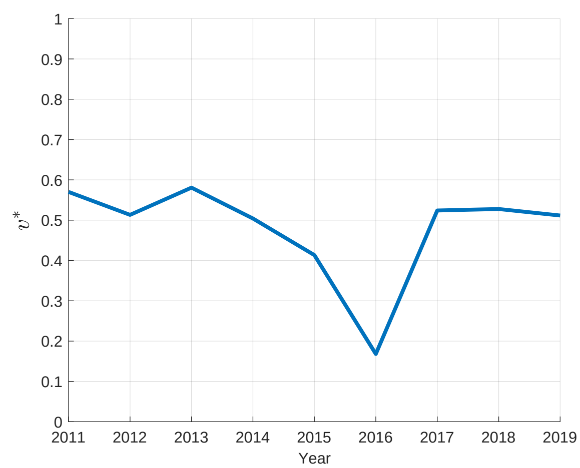

(c) The decomposition of the discrepancy over time

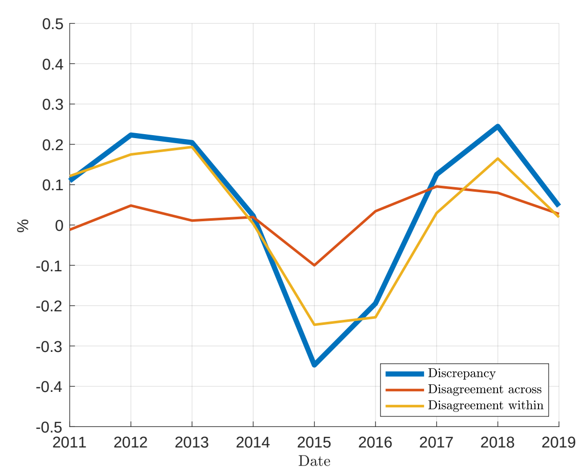

Figure 8: Estimates of expected long-run US inflation since 2000

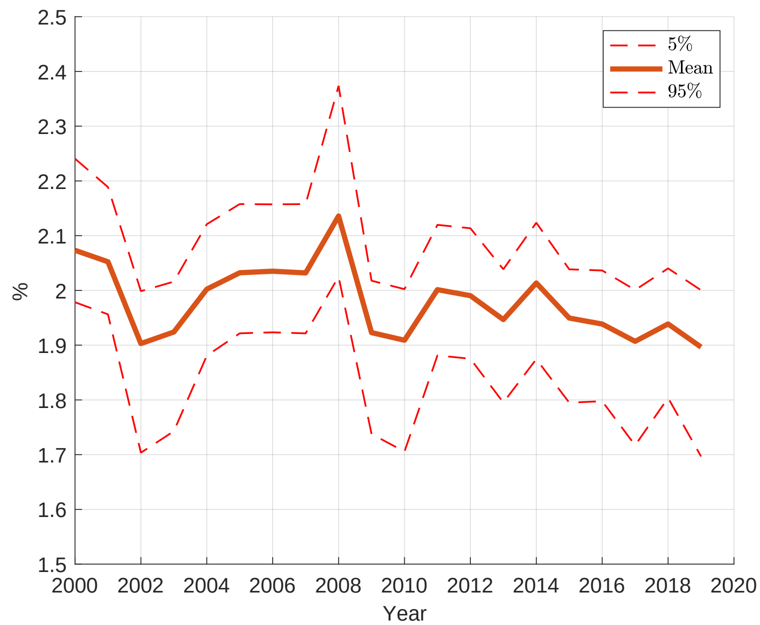

that drove the discrepancy down in 2013-16, with the marginal trader moving further away from the median. This reverts from 2017 onwards, but by then there is a fall in the mean of the survey, which keeps expected inflation down even as the discrepancy closes.

5.3 Estimates post-2000

Before 2011, there is no reliable data for the expectation of traders. There is however data from the survey of households, and from market prices using TIPS since the start of 2000. With this missing data, the model no longer pins down a unique fundamental expected inflation at every date. Instead, at each date, there is a pair (π_t^e, ω_t) that is consistent with the data. There is now an inference problem.

I solve it by taking a Bayesian approach. Like the agents in the model, a Beta distribution for ω provides a prior for it. Also like the agents in the model, but removed of their behavioral biases, I choose a normal prior for π_t^e centered at π_t^* and with a standard deviation of 0.1. The model has one curious property, however. The parameter determining the sensitivity of household expectations to their signal is:

$$\theta_t = \frac{\mu_t - \pi_t^*}{\pi_t^e - \pi_t^*} \tag{13}$$

Since $\theta_t > 0$, in the dates where the data implies that $\mu_t > \pi_t^*$, then the model immediately infers that $\pi_t^e > \pi_t^*$, and vice-versa when the sign switch. Therefore, the effective prior is a truncated normal, either to the left or to the right of π_t^*. As survey expectations move around, and μ_t hovers around 2%, the prior moves considerably and this would move most of the estimates. To deal with this issue, instead of fixing $\pi_t^* = 2\%$ I set a hyper-prior for it, following a normal distribution with standard deviation 0.1%. This way, at every date, different values of π_t^* are considered while the data fixes μ_t. The prior switches but the hyper-prior averages around these switches.

Figure 8 shows the mean and two percentiles of the posterior distribution for expected inflation. There is still a fall after 2000, although the new prior imposing no change makes it less dramatic. In the new sample, there is an interesting steady increase from 2002 to 2008, a period where arguably monetary policy was loose, at least relative to a Taylor rule. The model suggests that the common view that long-run inflation expectations have been essentially constant at 2% for the last 20 years is misleading. While the time-series changes have not been extreme, they are still clearly visible.

5.4 The euro-area

Looking for more data across regions, as opposed to over time, figure 9 shows the evolution of the discrepancy (top panel) and of the distribution of expectations in the Survey of Professional Forecasters for the Euro-area. The SPF is not the best sample for the model, since its forecasters are considerably more informed than the households but, at the same time, many of them are not traders that we could map into our model. Therefore, there isn't the data to estimate the model as there was for the US. At the same time, as the top panel shows, the discrepancy has been particularly large and persistent in the Euro-area, with market expectations stable around 1% since 2014, while the SPF expectations stable at 1.9%.

Informed by the lessons form the model, the bottom panel provides hints at what is going on. The distribution of expectations shows a noticeable increase in variance between 2010 and 2014. More striking is the emergence of a clear negative skew. This suggests that the marginal trader in Euro-area inflation contracts was shifting noticeably to the left

Figure 9: The discrepancy and expectations in the Euro-area

(a) The people/traders mis in the SPF, and the markets

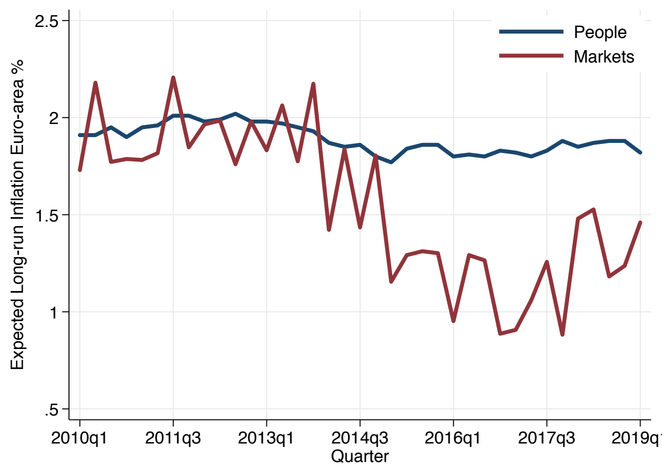

(b) The distribution of people/traders in the SPF

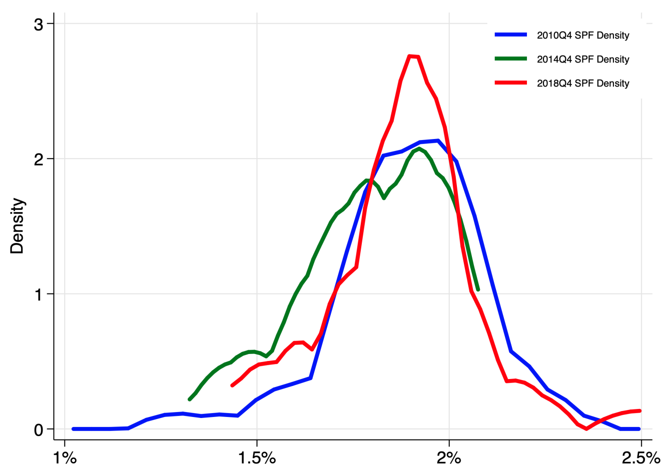

around the same time as it did in the US. However, unlike the US, in the following years, the mean of the distribution fell (as did its variance). This suggests that disagreement within in a first stage, followed by disagreement across in the second stage, contributed to keep market expectations low, and likely to a lower fundamental expected inflation in the Euroarea than in the US.

6 Monetary policy and inflation

This section embeds the partial-equilibrium model of financial markets into a simple general-equilibrium model of inflation determination. The goal is to endogenize the fundamental π_t^e, so as to understand the interaction between it, the discrepancy, and monetary policy. There are many other interactions to study in future research, but the model here is kept as simple as possible.

In terms of notation, time t is continuous, p_t is the price level, g_t is the expected growth rate of output, and Z_t is a vector of all the shocks hitting the economy, which are all Ito processes. Inflation is given by:

$$\frac{dp_t}{p_t} = \pi_t^e dt + \alpha' dZ_t \tag{14}$$

where π_t^e is the expected inflation that is te focus of this paper, and α is the vector with the sensitivity of inflation to each shock. Both are equilibrium objects, which the model will solve for.

6.1 Financial markets, households, and the transmission of the discrepancy

A central bank sets the nominal interest rate i_t^{CB} that prevails in financial markets. Because this is reflected in the price of bonds, then the shadow real interest rate out of financial markets is: $r_t = i_t^{CB} - \mathbb{E}^m(dp_t/dt)$, where crucially it is the expectation of inflation from financial prices that is relevant.[25]

[25] A simplifying assumption is to leave in the background the link from the long-run inflation expectations discussed earlier to these expectations of short-run inflation. It would not be hard to extend the model to have bonds traded at different maturities and make this link. Since the model is affine, a standard affine term structure would follow leading to similar conclusions (Reis, 2019).

The households, within a period, live in an island isolated from the traders, and without observing their nominal prices or expectations. They take instead as given the shadow real interest rate offered by the financial markets, and use their expectations of inflation to offer savings contracts to consumers in nominal terms at a rate: $i_t = r_t + \mathbb{E}^p(dp_t/dt)$. This implies that there is a wedge between the nominal interest rate set by the central bank, and the one offered for savings, and that wedge is the discrepancy: $i_t = i_t^{CB} - \phi_t$.

Recall that monetary policy will affect ϕ_t through its effects on expected inflation. Therefore, ϕ_t is standing in more generally as a proxy for the frictions in the transmission of monetary policy, from policy rates to the ons used for intertemporal decisions. There are surely other shocks and variables affecting the wedge, but given this paper's focus on the discrepancy, that is the one used here.

Finally, there is a representative agent making consumption decisions. At the end of the period, households and traders return to this agent, and by aggregating their information, she has rational expectations over inflation. As a result, the Euler equation is given by:

$$g_t = \ln(\beta) + i_t - \pi_t^e + \alpha'\alpha \tag{15}$$

so that the expected growth rate of consumption is equal to minus the subjective discount rate $(\ln(\beta))$, plus the shadow real rate $(i_t - \pi_t^e)$, plus the inflation risk premium $(\alpha'\alpha)$.

It only takes some ingenuity to make different assumptions on how the three agents are partially or fully segmented in their markets and information. In the end, they would all lead to a version of the transmission of monetary policy to real decisions of the form:

$$g_t = \ln(\beta) + i_t^{CB} + \alpha'\alpha - \pi_t^e - \theta\phi_t, \tag{16}$$

for different values of θ. In the model laid out above, $\theta = 1$. In an extreme, all the expectations formed by markets and people are irrelevant, as savings choices are made with rational expectations, so $\theta = 0$. In between, different values of $\theta \in [0,1]$ correspond to different weights that the expectations of different agents have on ultimate savings decisions. I will consider these different cases in the analysis that follows.

6.2 Monetary policy and its effects

The central bank sets interest rates according to a rule:

$$di_t^{CB} = -\rho(i_t^{CB} - i^*)dt + \eta\left(\frac{dp_t}{dt} - \pi^*\right) + \gamma d\phi_t \tag{17}$$

The central bank has a long-run target for inflation, which coincides with the prior for inflation that agents use, and which is taken here to be constant. Keeping inflation and its expectations near to this target is the goal. Consistent with this target, nominal interest rates in the long-run are equal to $i^* = \pi^* + \ln(\beta) + \alpha'\alpha$. The central bank though smoothes interest rate adjustments towards that target, at a rate ρ, following a continuous-time partial adjustment process.

The two key policy parameters of interest to this paper are η and γ. The first describes how aggressive monetary policy is in response to inflation. It is well understood that η has to be sufficiently high to ensure that inflation expectations are anchored in the sense of delivering a determinate equilibrium for inflation. The question is whether that threshold—the Taylor principle–change sin the presence of the discrepancy. The second parameter describes whether the central bank responds to the discrepancy itself. Central banker speeches leaves no doubt that they respond to measures of inflation expectations, and the influence of the fundamental rational expectation of inflation is already present through η. The γ allows us to study the consequences of responding differentially to the different measures of expectations.

Monetary policy does not respond to output. This is a consequence of the final : the classical dichotomy will hods, so output growth will be exogenous to the model. This is consistent with the focus of the paper on long-run inflation expectations. In the long run, arguably, prices are flexible and the Phillips curve is vertical, justifying this assumption. More practically, introducing firms and price stickiness would require introducing the expectations of firms distinct from those of households and markets, and other discrepancies, which are best left for future work.[26] Output growth follow a continuous-time autoregressive process:

$$dg_t = -\kappa_g(g_t - g^*)dt + \sigma_g dz_t^g \tag{18}$$

with innovations dz_t^g. These are one fundamental source of shocks in the economy. They are typically referred to as shocks to the natural rate of interest, and I will follow this

[26] On the discrepancy between household and firm expectations, see Coibion, Gorodnichenko and Kumar (2018).

convention.

6.3 The discrepancy and equilibrium

Finally, the model of information of the previous sections translates into a function: $\phi_t = \Phi_t(\pi_t^e, \omega_t)$. Empirically, the last section effectively characterized this function for the U.S. In this section, to preserve the linearity of the model, I approximate this functions log-linearly with:

$$\phi_t = \alpha_\pi(\pi_t^e - \pi^*) + \alpha_\omega \hat{\omega}_t \tag{19}$$

On the US data, the empirical estimates of the two elasticities are: $\alpha_\pi = 0.29$ and $\alpha_\omega = 0.87$.

Finally, I make a modest change to a previous assumption: the deviation of the noise shocks from its mean of 1/2, $\hat{\omega}_t$ will follow an autoregressive process with normal innovations:

$$d\omega_t = -\kappa_\omega \hat{\omega}_t dt + \sigma_\omega dz_t^\omega \tag{20}$$

as opposed to the Beta distribution it followed before. This is to preserve the linear-Ito structure of the problem.

All combined, an equilibrium of this monetary economy is a solution for inflation in (14) in terms of fundamental expected inflation π_t^e and the volatility of inflation in response to economic shocks α_g and noise shocks α_ω, subject to the Euler-Fisher equation in (16), the monetary policy rule in (17), and the discrepancy equilibrium law of motion , as well as the exogenous law of motion for the two shocks in equations (18) and (20).

The discrepancy plays two roles in the economy. First, it affects monetary policy through the parameter γ, and it affects its transmission to inflation through the parameter θ. If both of these are zero, then we are back at the conventional analysis of inflation with an interest-rate rule. Second, it is affected and affects the equilibrium expected inflation through the parameter α_π, and it introduces financial shocks as drivers of inflation through the parameter α_ω. Again, if both are zero, then the discrepancy would be zero, and we would get a conventional analysis.

6.4 Basics of anchoring: determinacy

A minimal definition of anchoring inflation is to make sure that the model of its dynamics has a determinate equilibrium. The appendix proves the following result:

Proposition 1. *Inflation is determinate as long as:*

$$\eta > \rho + \rho\theta\alpha_\pi \quad \textit{and} \quad \alpha_\pi(\gamma - \theta) < 1 \tag{21}$$

If the discrepancy was exogenous, driven by noise, unrelated to inflation, then $\alpha_\pi = 0$ and the condition reduces to the standard Taylor principle: $\eta > \rho$. The response of interest rates to inflation has to be large enough to offset interest-rate smoothing. The discrete-time equivalent with no smoothing as stated by Taylor is $e^\eta > 1$. This would also be the case in the first condition if the discrepancy did not affect the transmission of policy rates to the savings decisions so $\theta = 0$.

The presence of the discrepancy requires monetary policy to be unambiguously more aggressive in response to inflation. The reason is that when expected inflation rises, markets update by more than people, so the discrepancy rises. But, as a result of, market interest rates fall, which pushes inflation up. This endogenous mechanism works against the response of policy, which can now therefore be stronger. The second condition puts an upper bound on the policy response to the discrepancy γ. If that is too high, then an upwards-deviation of expected inflation will come with a lower expected change in the discrepancy looking forward, and thus interest rates fall.

6.5 The volatility of expected inflation

The appendix proves the following result:

Proposition 2. *Expected inflation equals:*

$$\pi^e = \pi^* + \frac{(\rho - \kappa_g)(g_t - g^*)}{\eta - \rho + \rho\theta\alpha_\pi + \kappa_g(1 - \alpha_\pi(\gamma - \theta))} + \frac{\alpha_\omega[(\rho - \kappa_\omega)\theta + \gamma\kappa_\omega]\hat{\omega}_t}{\eta - \rho + \rho\theta\alpha_\pi + \kappa_\omega(1 - \alpha_\pi(\gamma - \theta))} \tag{22}$$

If inflation is determinate, then the conditions in proposition 1 ensure that the denominator in both fractions is strictly positive. The first term therefore reflects the standard result that an increase in the natural rate of interest (g_t) will tend to raise expected inflation as long as interest rates persist relatively little relative to the shock (so ρ is high relative to κ_g). More interesting is the new term on financial shocks. When there is a financial shock, market expectations of inflation rise, while the people are unchanged. The discrepancy rises, so for a given policy rate, the private economy behaves as if the real interest rate was lower, where this effect is controlled by the size of θ. This pushes expected inflation up. At the same time, a financial shock leads to a tightening of monetary policy, where the size of this policy response is controlled by the size of γ. This lowers actual inflation, but raises expected inflation in the transition as is standard.

If policy wants to keep inflation expectations anchored, it may want to reduce the volatility due to financial shocks. This would require setting $\gamma = -\theta(\rho/\kappa_\omega - 1)$. That is, insofar as a positive discrepancy leads to together financial conditions, monetary policy wants to lower interest rates to loosen them back. However, focusing now on the effects of the natural rate shocks lowering γ leads to a higher volatility in response to those shocks. This is because now, when expected inflation rises, the observed increase in the discrepancy will be met by the monetary authority with looser policy. Lowering the variance of expected inflation de to natural rate shocks calls instead for a positive γ. Depending on the relative variance of the two shocks, the policymaker will choose. The two forces for the discrepancy are at play and point in opposite directions. Ignoring the discrepancy $\gamma = 0$ will in general not be optimal.

6.6 Who is right: people vs. markets?

Who forecasts better: the people or the markets? In the model, because market traders are better informed, their expectations are more responsive to the fundamental expected inflation. At the same time, they are contaminated by the noise that comes from the financial shocks. Monetary policy makes the comparison trickier, since a higher responsiveness to the discrepancy makes actual inflation respond more to natural-rate shocks, but less to financial shocks. An extreme example makes this clearer. Imagine monetary poicy is so aggressive that inflation is always exactly on target. Fully uninformed households would be perfect forecasters, even as markets would fail due to the financial shocks

The appendix proves the more general result.

Proposition 3. *If monetary policy responds to the discrepancy is large enough $\gamma > \gamma^*$, then financial market forecast of inflation have a higher instantaneous mean squared error than the people's forecast*

The fact that surveys seems to sometimes forecast US inflation better than markets has been seen as puzzling (Ang, Bekaert and Wei, 2007). After all, market participants should have superior information. The model in this paper, by having the discrepancy be endogenous, and especially depend on monetary policy, can rationalize it. If monetary policy responds to markets to keep inflation under control, then markets can do worse.

7 Conclusion

How are expectations of macroeconomic variables formed, and how should policy adapt to it? This paper added a new perspective on this classic question. It proposed a parsimonious model of subjective expectations that is flexible enough to fit the US survey data, and a parsimonious model of financial markets that is flexible enough to fit the US market data. The model can therefore be used to uncover what is the underlying fundamental expected inflation, and what accounts for the discrepancy between market and survey expectations that we observe. When integrated in a model of monetary policy, it was used to understand why inflation expectations may be better or worse anchored.

The paper reached a few conclusions. First, that the discrepancy in long-run inflation expectations has large business-cycle fluctuations, it is systematically related to monetary policy, and it is driven by disagreement across groups in the population as well as disagreement between the average and the marginal market traders. Second, that a combination of imperfect information, over-confidence, learning from experience, and sticky information, can explain the three first moments of the cross-sectional survey data on household long-run inflation expectations. Variations in skewness are important, even though they have been so far mostly ignored. Third, the marginal trader in the US data was significantly bearish on inflation in 2015 and 2016, and this played large role in explaining the large negative discrepancy that arose during this time. Fourth, disagreement within has roughly been as important as disagreement across in explaining the discrepancy in the last decade. Fifth, long-run expected inflation has been trending down in the US since 2014 and is now around 1.8% (with the incomplete Euroarea data suggesting it may be significantly lower there). Sixth, according to the model, determinacy of inflation requires monetary policy to respond more aggressively to inflation. Seventh, monetary policy must trade off financial shocks versus natural rate shocks when choosing how much to respond to the discrepancy, and in doing so it may result in either markets or people being the better forecaster.

References

Albagli, Elias, Christian Hellwig, and Aleh Tsyvinski. 2013. "A Theory of Asset Prices based on Heterogeneous Information." C.E.P.R. Discussion Papers CEPR Discussion Papers 9291.

Andrade, Philippe, and Hervé Le Bihan. 2013. "Inattentive professional forecasters." *Journal of Monetary Economics*, 60(8): 967 – 982.

Andreasen, Martin M., Jens H. E. Christensen, and Simon Riddell. 2017. "The TIPS Liquidity Premium." Federal Reserve Bank of San Francisco Working Paper Series 2017-11.

Ang, Andrew, Geert Bekaert, and Min Wei. 2007. "Do macro variables, asset markets, or surveys forecast inflation better?" *Journal of Monetary Economics*, 54(4): 1163 – 1212.

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

Beechey, Meredith J., Benjamin K. Johannsen, and Andrew T. Levin. 2011. "Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area Than in the United States?" *American Economic Journal: Macroeconomics*, 3(2): 104–29.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review*, 110(9): 2748–2782.

Clarida, Richard, Jordi Galí, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *The Quarterly Journal of Economics*, 115(1): 147–180.

Clements, Michael P. 2019. *Macroeocnomic Survey Expectations. Palgrave Texts in Econometrics*, Palgrave MacMillan.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120(1): 116–159.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review*, 105(8): 2644–2678.

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.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. 2018. "How Do Firms Form Their Expectations? New Survey Evidence." *American Economic Review*, 108(9): 2671–2713.

Eusepi, Stefano, and Bruce Preston. 2018. "The Science of Monetary Policy: An Imperfect Knowledge Perspective." *Journal of Economic Literature*, 56(1): 3–59.

Eusepi, Stefano, Emanuel Moench, Bruce Preston, and Carlos Viana de Carvalho. 2019. "Anchored Inflation Expectations." *CEPR Discussion Paper 13900*.

Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig. 2014. "The TIPS-Treasury Bond Puzzle." *Journal of Finance*, 69(5): 2151–2197.

Garcia, Juan Angel, and Sebastian Werner. 2018. "Inflation News and Euro Area Inflation Expectations." International Monetary Fund IMF Working Papers 18/167.

Grossman, Sanford J, and Joseph E Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review*, 70(3): 393–408.

Guo, Hongye, and Jessica A. Wachter. 2019. ""Superstitious" Investors." National Bureau of Economic Research, Inc NBER Working Papers 25603.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2010. "The TIPS Yield Curve and Inflation Compensation." *American Economic Journal: Macroeconomics*, 2(1): 70–92.

Hilscher, Jens, Alon Raviv, and Ricardo Reis. 2014. "Inflating Away the Public Debt? An Empirical Assessment." *NBER Working Paper 20339*.

Kose, Ayhan, Hideaki Matsuoka, Ugo Panizza, and Dana Vorisek. 2019. "Inflation Expectations: Review and Evidence." *CEPR discussion paper 13601*.

Malmendier, Ulrike, and Stefan Nagel. 2015. "Learning from Inflation Experiences." *Quarterly Journal of Economics*, forthcoming.

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics*, 117(4): 1295–1328.

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

MARTIN, IAN W. R., and CHRISTIAN WAGNER. 2019. "What Is the Expected Return on a Stock?" *The Journal of Finance*, 74(4): 1887–1929.

Milgrom, Paul, and Nancy Stokey. 1982. "Information, trade and common knowledge." *Journal of Economic Theory*, 26(1): 17–27.

Moessner, Richhild, and Előd Takáts. 2020. "How well-anchored are long-term inflation expectations?" *BIS Working Papers 869*.

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

Reis, Ricardo. 2019. "Central Banks Going Long." In *Monetary Policy and Financial Stability: Transmission Mechanisms and Policy Implications*. , ed. Alvaro Aguirre, Markus Brunnermeier and Diego Saravia, 43–81. Central Bank of Chile.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics*, 50(3): 665–690.

Veldkamp, Laura L. 2011. *Information Choice in Macroeconomics and Finance.* Princeton University Press.

Vives, Xavier. 2008. *Information and Learning in Markets: The impact of market microstructure.* Princeton University Press.

Woodford, Michael. 2003*a*. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S.* , ed. Joseph Stiglitz Phillipe Aghion, Romain Frydman and Michael Woodford, 25–58. Princeton University Press.

Woodford, Michael. 2003*b*. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton, NJ:Princeton University Press.

Yetman, James. 2020. "The pass-through from short-horizon to long-horizon inflation expectations." *BIS Papers 111*.