Four mistakes in the use of measures of expected inflation

By Ricardo Reis*

It is often a good idea to simplify and to focus. To learn something, when there are multiple measures of that something, we often try to find the best measure, and discard all others. For most problems, the best measure is not publicly available. The focus should be on some average (an index or aggregate, to be more precise) over the different measures. The discussion turns then to what should the weights on each of the series given the question at hand. A common strategy is to argue that the weight should rise with some characteristic, observe that one series is much better in this dimension, and set its weight to 100%. Superficially, this seems sensible, or even optimal, given limited attention. But it can lead to severe mistakes.

Over the last two decades, there has been a remarkable amount of progress in understanding how people form expectations of inflation and how these affect inflation outcomes. Much of this progress has been made through measurement, using both surveys of expectations and models of prices in financial markets (Weber et al., 2022). This work was partly validated in 2021-22, as these measures provided valuable early signals that an inflation surge was on the way (Reis, 2022). At the same time, because there are now several series for expected inflation for the major economies, too often researchers and policymakers will pick one to focus on arguing it is the right one, or the best one. In this note I work through four common arguments made for these choices. I find all of them to be weak and each to lead to mistakes in using expectations data to understand what drives inflation and to guide monetary policy.

* LSE, Houghton Street, London WC2A 2AE, United Kingdom (e-mail: r.a.reis@lse.ac.uk). I am grateful to a UKRI grant for financial support. The references and the full model used in this article are available in its CEPR discussion paper version.

I. Mistake 1: Focus on firms' expectations because firms set prices

Until recently, the only surveys that spanned a few decades and included a few hundred respondents were those of households, like the famous Michigan survey for the US. They were sometimes dismissed using the following argument: Firms choose prices in the economy. It is their expectations that matter for how prices are set, and therefore for what inflation will be. In a modern Phillips curve equation, it is firms' expectations that appear on the right-hand side of the equation as a key driver of inflation. Households take prices as given and their expectations do not appear in that equation. So, the argument goes, they are not relevant, and so household surveys are of little use to understand inflation.

Looking forward, researchers have recently produced new surveys of firms' expectations (Candia, Coibion and Gorodnichenko, 2022). As of now, it seems that the estimates they produce do not track inflation so well, and firm managers share with households many of their biases and inefficiencies when making forecasts. A new version of this argument is that, since theory says that these are the right measures of expectations that should be tightly linked to prices, then surveys are not useful at all.

The appeal to the Phillips curve in this argument is superficially plausible, but also weak, or downright wrong. In a large class of models of nominal rigidities there is a partial-equilibrium relation in the goods market, derived from monopolistic firms maximizing real profits, given demand, and subject to nominal rigidities:

$$\pi = \pi^f + rmc \tag{1}$$

where π is inflation, rmc are real marginal costs, and π^f are *firms' expectations of inflation*, all as log-linear deviations from a

steady state. Intuitively, firms want to raise their prices relative to the prices that they expect their competitors are setting when the cost of producing an extra good is higher.

However, in almost any plausible general-equilibrium environments, rmc depends on other agents' expectations as well. In fact, in some cases, it even depends on firm expectations with the opposite sign. Take an extreme case where the only variable input is capital and it is lent to the firm at a nominal interest rate i^l. Then, real marginal costs are $rmc = i^l - \pi^f$. Firm expectations cancel out in the equation. Intuitively, facing a fixed nominal interest rate, if the firm expects higher inflation, it expects that in real terms the cost of hiring inputs is lower, so it can lower its price. This exactly offsets the initial desire to raise prices.

What does the lending rate i^l depend on? It is set by financial institutions, for which the marginal cost of funds depends on the interbank market rate targeted by the central bank i. But, as long as there are financial frictions that require the financial institutions to use some of their net worth, the required return on net worth affects the lending rate as well. Taking that real return to be constant for simplicity, then:

$$i^l = i + \gamma(\pi^m - i), \tag{2}$$

with the parameter γ capturing the extent of the financial frictions.

For a given real return on net worth, higher *financial markets' expectations of inflation* π^m raise the interest rate that is charged to the firm. If markets start expecting higher inflation, they will raise the interest rates they charge on loans, which raises the financial costs of firms, leading them to raise prices. This is a general-equilibrium effect, from combining the goods market with the loan market.

Labor is also a variable input in production, and letting α be the capital share in production, real marginal costs are instead:

$$rmc = \alpha(i^l - \pi^f) + (1-\alpha)(w - \pi), \tag{3}$$

where w is the nominal wage. It is set by workers. The more labor is used and output produced, the more they must be paid for the disutility of working, with an elasticity of θ. If unions and workers have some bargaining power and set wages subject to nominal rigidities, they will also have to form some expectations of inflation. In equations, if the *workers/unions' expectations of inflation* are π^w, then:

$$w = \pi^w + \theta y. \tag{4}$$

If workers expect higher inflation, they ask for higher nominal wages. If inflation has not itself yet changed, this raises the real wages that firms pay. They respond by raising their prices, which causes inflation. Again, it is general equilibrium, now working from the labor market to the goods market, that make higher inflation expectations elsewhere in the economy drive the firms to raise their prices.

To conclude, superficially it is firms that are setting prices and they respond to their expectations. But they respond as well to the costs they face. Those costs depend on the expectations of inflation of workers and financial markets. In economic equilibrium, choices depend on other's actions, and a priori any of the beliefs could be more or less important for the decisions that are made.

In practice, these influences are not negligible. The expectations of financial markets translate quickly to the financial conditions facing all agents, so they have a fast and powerful impact on inflation. Monetary policy relies on this transmission channel, which seems to be strong. In turn, when expectations of wages move away from the central bank's target they are hard to re-anchor, and can start wage-price spirals. This is, arguably, the major concern about inflation at the start of 2023. For a central bank today, firms' expectations may very well be the least important of the three, even if they are the ones that set the prices.

II. Mistake 2: focus on the big players as their choices drive aggregates

Large firms, unions, or banks have a large weight in the averages of production, labor, and credit that are behind inflation

outcomes. This is especially so in financial markets, as private credit in most countries is concentrated on a handful of banks. Another common argument is that a larger weight should be put on surveys of large firms, especially in the financial sector, like the Blue Chip survey in the US.

An immediate objection is that market prices reflect the actions of the marginal agent, not the average over agents. The lending rare i^l is set at the margin where demand and supply for credit meet. The bank that is just indifferent between lending or not can just as well be small. With financial frictions, the quantity of credit can matter independently of its price. But how much, and whether large lenders have an outsized role, varies considerably across models. In practice, measures of expected inflation from market prices differ systematically from the survey measures of bankers or dealers in bonds and swaps.

Further, consider what determines expected inflation. There are many well-developed models in the literature of how people form their beliefs. At one extreme, if they have rational expectations and perfect foresight, expected inflation equals actual inflation. At the other extreme, expectations are animal spirits. A reduced-form way to capture the in-between is to say that

$$\pi^f = (1-\lambda^f)\pi + \lambda^f \hat{\pi}^f, \tag{5}$$

for firms where $\hat{\pi}^f$ are the exogenous spirits, and λ^f is a parameter between zero and one. The same applies to workers and financial markets with λ^w and λ^m respectively.

The closer the λ's are to zero, the less useful it is to measure expectations through expensive surveys or sophisticated techniques. The measures are just mirrors of what is going on in reality, and researchers are better off measuring outcomes and fundamental shocks. Plausibly, large firms with chief economists will have a small λ^f. Therefore their $\hat{\pi}^f$ spirits will not be so important on aggregate outcomes. The players may be large, and their choices drive outcomes, but the *autonomous* changes in their expectations that could bring a shock to inflation are small, and drive little of the variation that we see in the data.

This can be seen mathematically by combining all the equations presented so far to get the actual Phillips curve for the economy, the structural relation that links real activity to inflation as a result of general equilibrium across markets:

$$\pi = \pi^e + \kappa y + \xi(i - \pi) \tag{6}$$

The coefficients κ and ξ are formulae of all the other parameters and presented in the accompanying discussion paper. More interesting, expected inflation π^e is a weighted average that sums to one of the expectations of firms, markets, and workers:

$$\pi^e = \frac{\alpha\gamma\lambda^m\hat{\pi}^m + (1-\alpha)(\lambda^w\hat{\pi}^w - \lambda^f\hat{\pi}^f)}{\alpha\gamma\lambda^m + (1-\alpha)(\lambda^w - \lambda^f)} \tag{7}$$

Each agents' expectation has a larger weight on π^e if their $\lambda's$ are larger.

Again, in practice this is not negligible. Both in the EA and the US, surveys of chief economists in firms are almost always quite close to the central bank's internal forecast. When inflation is close to target, they do not add much information. When the central bank's model got it wrong in US history—the rise of inflation in the late 1960s, its fall in the early 1980s, and the new rise in 2021-22—the professional forecasters were just as wrong, and instead it was household expectations that seemed to provide impetus for the dynamics of inflation and their survey measures were the ones that contained useful signals (Reis, 2021).

III. Mistake 3: focus on the measures with smaller forecast errors

Some people do not care about what drives inflation, but are only interested in forecasting it. They would then argue that the economic arguments in the precious sections should be ignored. Rather, they would compare the forecasting performance of different measures of expected inflation, and focus on the one which does best according to a criteria like mean squared error. The answer in many countrres and

many decades is a survey of professional forecasters. A more brusque version of this argument discards household expectations entirely because, since they are biased and have persistent forecast errors, their mean squared forecast errors are large.

Even from a statistical perspective, this argument is weak. To start, if the goal is forecasting performance alone, then the measure of expected inflation with lowest mean squared forecast error in most advanced economies is the forecast published by the central bank. Since these forecasts often include data from other measures of expected inflation, not much is learned from this exercise.

Moreover, inflation in history has gone through regimes. Surveys of professional forecasts do well within regimes, but not at times of regime change. A careful evaluation of forecast performance is tricky and the samples for which we have measures of expectations are not quite long enough to reach definitive conclusions.

Finally, focussing on forecast performance confuses concept with measurement. Surveys might be poor but can be improved through better design. Expectations may be biased, but theories of those biases can be used to de-biase them. A direct measure of expected inflation from a survey may seem far off from reality. But putting this measure into a careful model that links it to the relevant concept may well be very informative.

Turning bank to economics, models are mostly used, not for unconditional forecasting, but rather for forecasting what will happen conditional on a shock. To close the model developed so far, start by adding an equation for aggregate demand:

$$y = -\omega(i^l - \pi) + \sigma(\pi^c - \pi) \tag{8}$$

The first term captures the fall in current spending (or rise in savings) when returns are higher. For a fixed nominal interest rate in financial markets, the second term captures the force that higher *consumers' expected inflation* π^c leads them to want to spend more today before prices rise. With a single representative agent, $\omega = \sigma$, but in richer models these terms separate as σ includes not just intertemporal substitution but also the response of consumption to cureent and expected future income.

Close the model with a standard rule for monetary policy: $i = \phi\pi + \phi_y y + \varepsilon$ with policy parameters ϕ and $\phi_y > 0$ and policy shocks ε. Then, focusing only on the response of inflation to a shock to consumer expectations (the same could be done with respect to the other agents's expectations) gives:

$$\frac{\partial \pi}{\partial \hat{\pi}^c} = \frac{\sigma}{\rho + \sigma} \tag{9}$$

where ρ is a composite parameter that is larger than one if inflation is determinate. The message is clear: shocks to expectations of consumers matter more for outcomes if σ is higher.

In the model, this parameter determines the transmission *from expectations to actions*. A similar conclusion applies to the other expectation shocks with respect to the parameters that capture how much their actions respond to their expectations. In general, this transmission is the key parameter to decide how much attention to pay to surveys of consumers, not whether the measures are statistically accurate or not.

IV. Mistake 4: focus on the expectations that policy can move

When policymakers change monetary policy, or give a speech, financial market expectations of inflation move within minutes. Household expectations, instead, rarely move at all with policies or communications. In fact, many people, including those running firms, usually cannot state what is the goal or mandate of the central bank, or who is currently its head. It is then natural for policymakers to focus on financial market expectations, and to devote their energy to managing those (Haldane, Macaulay and McMahon, 2021).

Of course, financial market prices do not just react to news. They often over-react to them, as well as to noise unrelated to fundamentals. Moreover, market prices re-

flect both expected inflation and risk premia. Removing the latter is hard and imperfect. A policymaker that responds to every movement in market expectations of inflation may end up propagating shocks to risk attitudes.

In the simple model of this paper, the responsiveness of market expectations π^m to a policy shock ε is captured by λ^m. All else equal, algebra shows that a low λ^m raises $\partial\pi/\partial\varepsilon$. But at the same time it lowers $\partial y/\partial\varepsilon$. That is, and perhaps unsurprisingly, more responsive expectations make the Phillips curve steeper. Conversely, very sluggish household expectations, captured by a high λ^w or λ^c, make the curve flatter. This changes the trade-offs that policymakers face in stabilizing both inflation and output. But it does not mean that one form of expectations are more important to understand or predict the effects of policy.

Treating the λ's as fixed parameters is a useful approximation when inflation is in a stable regime. But across regimes, economists have long known that the responsiveness of expectations to policy is endogenous to policy and the steepness of the Phillips curve changes. With multiple λ's across different agents, these changes will be different as well. If financial markets were already very responsive to news over two decades of low and stable inflation, then there is little room for change when inflation becomes high and volatile. Instead, since households were so unresponsive when inflation was low and stable, there is plenty of room for them to start paying more attention now. When this happens, inflation becomes more volatile in response to the shocks as well.

The experience from countries that go through prolonged periods of high and volatile inflation shows that this happens and it is large. A major task of a central bank in an inflation disaster is to re-anchor expectations. This can in part be understood as trying to convince agents to become inattentive again.

To conclude, when inflation is low and stable, the sluggishness of household expectations is not a reason to ignore them, but rather what gives the central bank power to affect output. When inflation becomes high and volatile, these expectations move, and policy managing to quiet them back into a stupor is part of what brings inflation down.

V. Conclusion

Across the four mistakes, I argued that the expectations of firms, large banks, professionals, or financial markets do not have any claim to being more useful than the expectations of households in order to understand inflation outcomes or to guide monetary policy. While household expectations are the ones that are more often dismissed, it would be just as mistaken to conclude from this article that one should only focus on household expectations and downweight the expectations of other agents. The simple, perhaps obvious, but often forgotten, conclusion is that one needs models to extract as much signal as possible from different measures and to combine them to provide the best guide.

REFERENCES

Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko. 2022. "The Macroeconomic Expectations of Firms." In *Handbook of Economic Expectations*. Elsevier: North Holland.

Haldane, Andrew, Alistair Macaulay, and Michael McMahon. 2021. "The Three E's of Central-bank Communication with the Public." In *Independence, Credibility, and Communication of Central Banking*. Central Bank of Chile.

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

Reis, Ricardo. 2022. "Expected inflation in the euro area: measurement and policy repsonses." *ECB Sintra Forum*.

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion. 2022. "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications." *Journal of Economic Perspectives*, 36(3): 157–84.