Oxford Review of Economic Policy, 2025, 41, 426–451
https://doi.org/10.1093/oxrep/graf028
Article

OXFORD

The straw that breaks the camel's back: inferential expectations and sudden belief changes

Timo Henckel,[*] Gordon D. Menzies,[**] and Daniel J. Zizzo[***]

[*]Australian National University and CAMA, Australia, e-mail: timo.henckel@anu.edu.au
[**]University of Technology and CAMA, Australia, e-mail: gordon.menzies@uts.edu.au
[***]The University of Queensland and CAMA, Australia, email: d.zizzo@uq.edu.au

We are grateful to seminar participants at the Australian National University, at the International Conference on Industrial Engineering and Informatization (IEIS), Beijing Jiaotong University, Beijing, and at the Feng Chia-ASEAN International Conference on Business and Management (ICBM), Feng Chia University, Taichung City, for helpful comments. Special thanks go to David Vines for his generous support, encouragement, insight, and feedback, and to Alison Gomm for her editorial assistance.

Abstract

We propose a theory of abrupt shifts in beliefs—termed inferential expectations (IE)—to explain significant structural changes observed in economies. According to this theory, agents conduct hypothesis tests on economic fundamentals, holding firmly to their initial beliefs until the weight of contrarian evidence reaches a critical threshold, prompting a sudden and seemingly disproportionate revision of expectations. We explain the intuition behind this theory, refer to its microfoundations, and list examples of possible applications. We also simulate a simple 3-equation macro model to provide qualitative insight into the recent inflation surge and the policy dilemma it presents.

Keywords: belief updating, expectations, anchoring, central bank credibility, inflation.

JEL codes: C34, C91, D03, D84, E03

I. Introduction

Economies rarely follow a smooth trajectory; instead, they often undergo sudden, substantial shifts that compel economic agents to rapidly adapt their behaviours, reassess their expectations, and reconsider their understanding of economic mechanisms. These structural changes can be triggered by external forces, such as a shifting climate, or arise from transformative developments within the economy itself, like the advent of a new general-purpose technology that boosts global productivity.

Structural shifts introduce significant challenges as they disrupt established relationships and assumptions that underpin economic models, decision-making, and policy frameworks. For the agents themselves—firms, households, banks, and so on—these shifts create uncertainty about future economic conditions, complicating investment, consumption, and saving decisions. For instance, a shift from a manufacturing-based economy to a service-oriented one can destabilize job markets, affecting household income and consumption patterns. For economic analysts, structural shifts make the estimation of structural relationships, like the association between inflation and unemployment or productivity and wages, far more complex, as traditional models may fail to account for new dynamics. Policymakers in turn face heightened difficulties in responding effectively, as past policy responses may no longer yield desired outcomes. Thus, structural shifts not only alter the economic landscape but also demand a rethinking of theoretical models and policy frameworks to navigate new uncertainties.

In this paper we argue that abrupt shifts in beliefs are one important source, and manifestation, of structural changes within an economy. We advance a theory of such abrupt belief changes—called *inferential expectations* (Menzies and Zizzo, 2009)—which formalizes how individuals, when confronted with uncertain or evolving economic signals, suddenly adjust their expectations in ways that to an outsider may look surprising and perhaps unwarranted. We explain theoretically how inferential expectations (IE) can be individually rational and provide experimental and empirical evidence, demonstrating that IE offers a valuable lens through which to understand how individual-level shifts in beliefs can contribute to broader structural changes.

The paper proceeds as follows: section II introduces the concept of inferential expectations, situating it within the broader literature on expectations formation, presenting supporting empirical evidence, and outlining potential applications, including the de-anchoring of inflation expectations. Section III develops a 3-equation New Keynesian model incorporating IE, which is used to simulate the post-Covid inflation surge observed in many advanced economies. The analysis is contrasted with recent studies of the same episode. Section IV summarizes the broader merits of inferential expectations as a modelling framework, and section V concludes.

II. Inferential expectations

(i) Intuition

The theory of inferential expectations offers a surprisingly powerful, yet simple, way to model sudden changes in expectations following a gradual build-up of evidence. The modelling metaphor used is a statistical hypothesis test. As a reminder, every hypothesis test has a variable about which one is forming expectations (the *cognitive target*), a maintained belief (a *null hypothesis* or H_0), and possibly an alternative belief (an *alternative hypothesis* or H_1), evidence that is used to calculate a *test statistic*, a significance level α (usually 1, 5, or 10%) which defines a *rejection region*, and a sampling distribution of the test statistic under the null. In Figure 1, agents start off believing that some imperfectly observed quantity equals a. If this looks a little unlikely, based on a statistic b, it is not enough to overturn their belief. However, if the statistic reaches c, the agent will reject a and conclude that the imperfectly observed quantity exceeds a. The significance level α (the black area) is set by, or is a characteristic of, the person doing the test. A lower α implies a higher c—more evidence is needed to reject H_0.

In economic models with inferential expectations, agents act as if they are doing a classical hypothesis test using information about fundamentals from any number of sources. Initially, an agent has an opinion about the state of the world. It could be something simple such as inflation will be 2.5% next quarter or something complex such as the financial system is not in any danger from sub-prime defaults. Then they receive signals about inflation or the financial sector, but do not change beliefs with every new observation. Instead, they hold onto their belief about the state of the world (H_0) even in the face of contrarian evidence. When enough evidence cumulates against the null that it appears significant (when the statistic breaches the rejection region), the old belief is no longer tenable, and the agent abandons their old belief to form a new belief. This is the straw that breaks the camel's back. The likelihood that the agent will update their belief depends on the significance level α, which can be interpreted as the degree of belief conservatism, or belief stickiness. If α is low, only large amounts of evidence against H_0 will appear significant and lead the agent to reject the null and revise their belief. If α is high, they will not require much contrarian evidence to reject H_0. In the limiting case where $\alpha = 1$, the agent exhibits no belief stickiness at all and updates their belief with each new piece of information, no matter how 'big' or 'small' the information is. In Cohen *et al*. (2019), which we explain in more detail in section II(iii), we show that inferential expectations are individually rational.

Such behaviour, where a belief regime can resist contrarian evidence for some time, before the evidence becomes overwhelming and forces a revision, is beautifully captured by the admission made by former Federal Reserve Bank Chairman Alan Greenspan on the CBS *60 Minutes* programme on 16 September 2007: 'While I was aware a lot of these [sub-prime] practices were

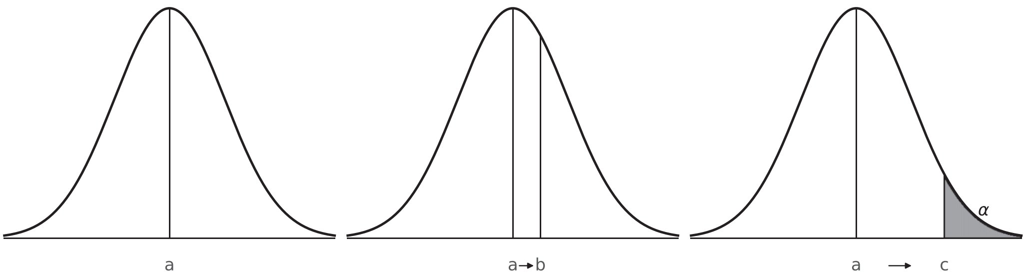

Figure 1. A hypothesis test.

going on, I had no notion of how *significant* they had become until very late. I didn't really get it until very late in 2005 and 2006.' (our italics)

Thus, IE beliefs may appear to be detached from slow-moving fundamentals but only until the last bit of information nudges the test statistic into the rejection region (at point c in Figure 1).

Menzies and Zizzo (2009) introduced this idea in a version of the famous Dornbusch (1976) 'overshooting' model, where a representative agent is endowed with inferential expectations. The representative agent forms expectations about the future exchange rate, based on observations of the money supply. In this model, the cognitive target is the money supply, the test statistic is the duration of a monetary expansion to date, and the null and alternative hypotheses are beliefs that the expansion is temporary or permanent, respectively. It is further assumed that the alternative hypothesis is the rational expectations (RE) solution, making RE a special case of IE (with $\alpha = 1$), since any null is rejected whenever α takes the value of unity. When a monetary expansion goes on for 'long enough', agents conclude that it is permanent and jump to the RE solution (Figure 2).

The model has the very attractive feature that it can explain the downward bias in tests for uncovered interest parity, appealing to the observed interest-rate forecast errors (documented in Mankiw and Miron (1986) and Gourinchas and Tornell (2004)).

Much economic analysis boils down to updating parameters when receiving sequential information. For a large class of situations, RE is well-represented by Bayes' rule which has a rational link between the persuasiveness of evidence and the extent of belief change. For example, imagine switching from one belief H_0 at time t to H_1 at time $t+1$ after receiving information I_{t+1}. A decisive belief change from H_0 to H_1 could be described as holding H_1 with a small probability ε prior to the arrival of the new information and then holding it with a high probability $1-\varepsilon$ afterwards. Thus, according to Bayes' rule:

$$Pr\,(H_1 \mid I_{t+1}) = \frac{\varepsilon\,\Pr\,(I_{t+1} \mid H_1)}{\varepsilon\,\Pr\,(I_{t+1} \mid H_1) + (1-\varepsilon)\,\Pr\,(I_{t+1} \mid H_0)} = 1-\varepsilon$$

This implies the information must be much more likely under H_1 than under H_0:

$$\Rightarrow \frac{\Pr\,(I_{t+1} \mid H_1)}{\Pr\,(I_{t+1} \mid H_0)} = \left(\frac{1-\varepsilon}{\varepsilon}\right)^2$$

Once we depart from RE, however, the rational link between the persuasiveness of evidence and the extent of belief change can be broken, leading to 'regimes' in which beliefs can resist change from contrary evidence, possibly for extended periods. When the belief change does occur, it may *appear* unjustifiably large when the full informational history is not accounted for.

(ii) Empirical support for inferential expectations

Marked belief changes are often directly observable in the data, such as when traders come to believe a currency needs to be fundamentally revalued, or when price-setters lose confidence in a central bank's inflation target. At these moments, the economy or society can change dramatically, even if the information is underwhelming—the veritable 'straw that breaks the camel's back'—most prominent examples of which are observed during crises (Figure 3).[1]

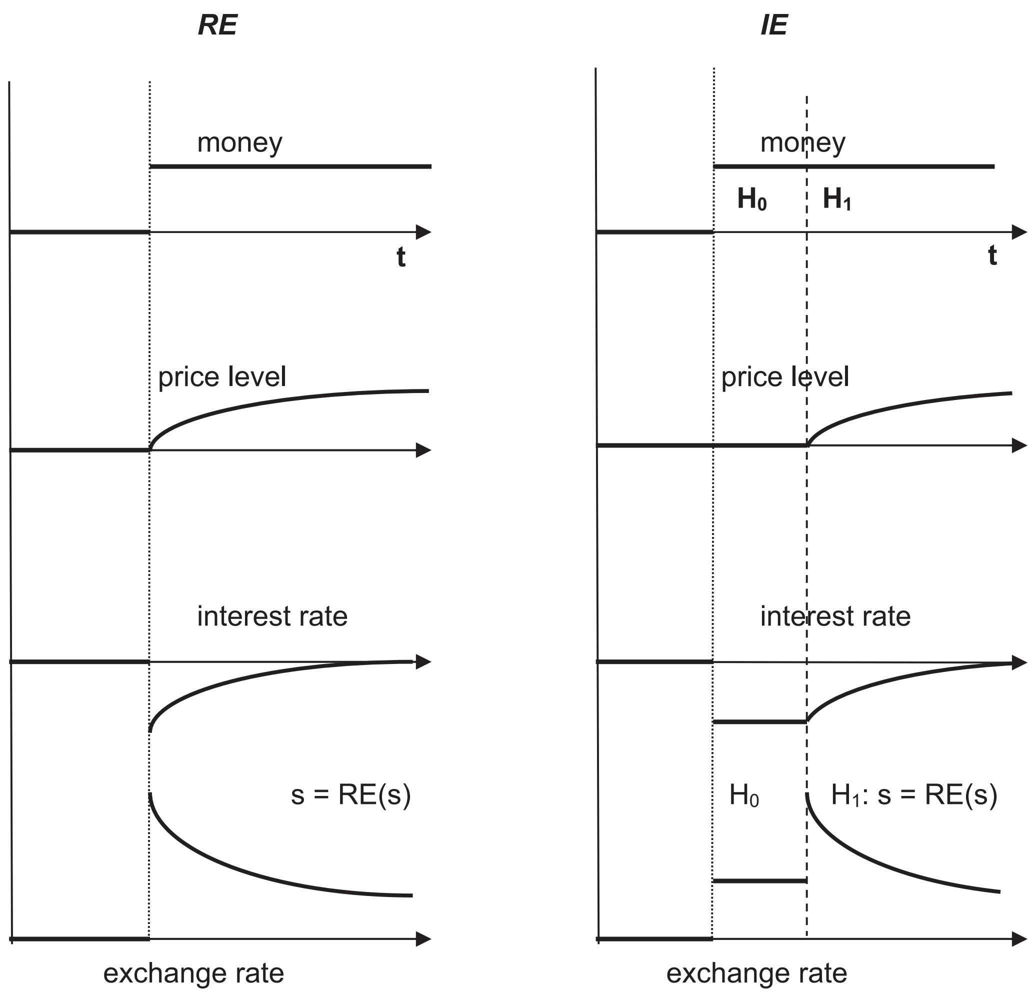

Figure 2. The Dornbusch (1976) 'overshooting' model with rational expectations (RE) and inferential expectations (IE) (under H_1 left and right panels identical).
Notes: The left-hand panel shows the standard dynamics from Dornbusch (1976) for a monetary expansion. The right-hand panel show the IE dynamics. Prior to H_0 being overturned, agents in the foreign exchange market ignore future (negative) interest rate differentials, and the depreciated (lower) long-run exchange rate. When H_0 is discarded for H_1 the exchange rate jump depreciates to the RE solution.

Time series charts relating fundamentals to beliefs, like the ones in Figure 3, are suggestive of inferential expectations. The psychology literature also provides examples in which a hypothesis test serves as a model of human inference. (See, for example, Adler and Rips (2008), Janis and Mann (1977), and Nisbett and Ross (1980)).

More recently, we provided direct empirical support for the central feature of inferential expectations, namely that agents do not revise their beliefs in response to every new piece of evidence, except in the limiting case where $\alpha = 1$. Henckel *et al.* (2022) report the findings of a classic 'poker chip bag' laboratory experiment involving two urns, each containing a distinct distribution of white and orange balls. One urn is selected randomly, and subjects receive

[1] A corollary of such behaviour is that minor fluctuations in market prices are not necessarily indicative of belief changes. We model this for an experiment extending Henckel *et al.* (2022), currently still work-in-progress, which considers the interaction of IE agents in a market setting.

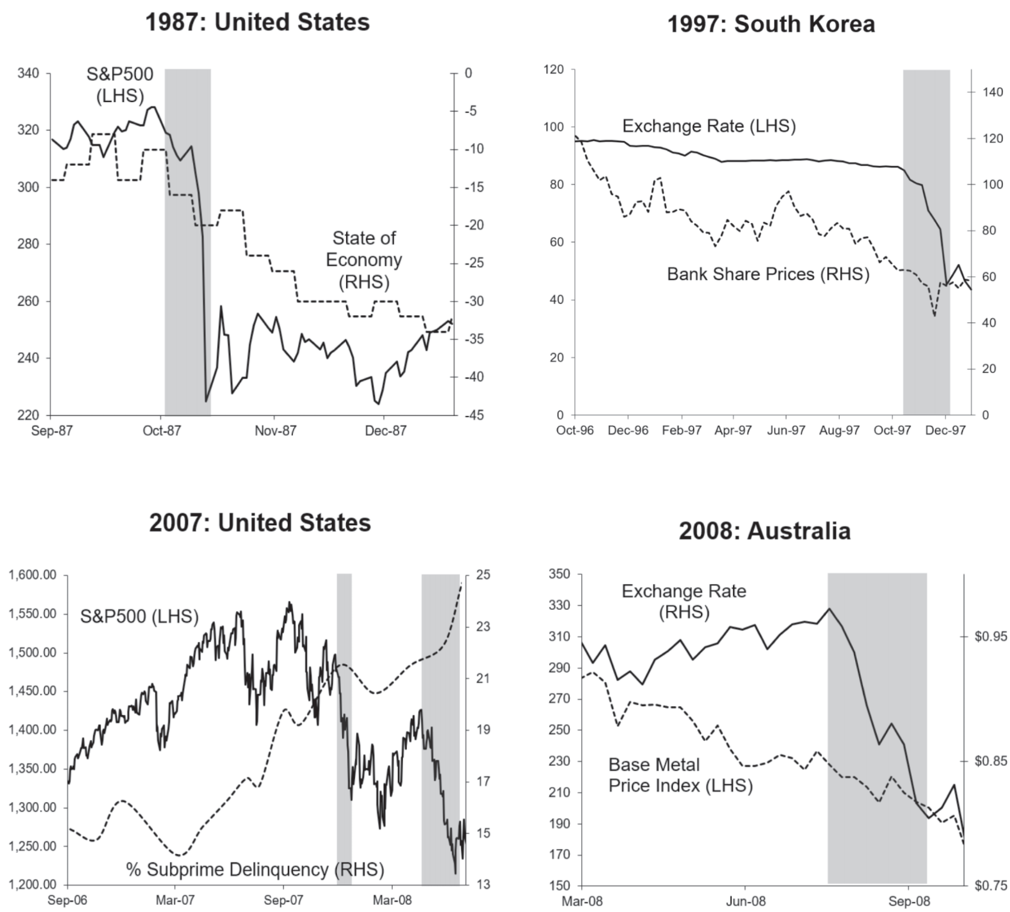

Figure 3. Structural belief changes in three decades of crises.
Notes: 1987 US: S&P500 Index and ABC News State of the Economy Index from Yahoo! Finance and Datastream. 1997 Korea: USD Korean Won index (1 August 1996 = 100) and Korean Bank Institution Stock Price Index from Datastream. 2007 US: S&P500 Index from Yahoo! Finance and Datastream and delinquency rates for sub-prime loans originated between January 2004 and August 2008 excluding home equity loans. Sub-prime loans are either identified as such by the servicer or have an original FICO score of less than 620; data from LPS Analytics. 2008 Australia: Westpac Base metals sub-index and USD exchange rate from Datastream.

signals about its contents by means of random ball draws with replacement from the chosen urn. Participants were asked to state the probability they attached to the urn being either the majority-white-ball urn or the majority-orange-ball urn. (Subjects were monetarily incentivized to provide the correct answer.) Under rational expectations (RE), the probability assigned to each urn should evolve according to Bayes' rule, with each newly drawn ball prompting an update. This would have been the best answer for subjects to give. However, subjects exhibited belief conservatism—a key implication of $\alpha < 1$ under IE—by failing to update their beliefs as frequently as predicted by RE. Even allowing for rounding,[2] participants failed to revise their probability estimates over 30% of the time, significantly more often than Bayes' rule (and RE) would suggest.

Moreover, Henckel *et al.* (2022) found evidence of quasi-Bayesian updating following rejection of the null hypothesis. In their double hurdle framework, subjects first decide whether to update according to a hypothesis test of size α. If they do, they revise their beliefs using a modified

[2] Rounding was defined as a failure to switch the stated probability guess if the RE Bayesian probability changed by less than 1%.

Updating taxonomy with results (rational expectations corresponds to $\alpha = \beta = 1$)

Hurdle 2: Extent of update / Hurdle 1: Decision to update		Interpretation of Quasi-Bayesian β in *Posterior* = (*LikelihoodRatio*)$^{\beta}$ *Prior*			
		Irrational $\beta < 0$	Underuse of information $0 \leq \beta < 0.9$	Bayes $\beta \approx 1^{*}$	Overuse of information $\beta > 1.1$
Inferential expectations α is attention towards evidence	$\alpha_{med} \approx 0 (< 0.01)$ (inattentive)	0%	4%	3%	7%
	$0.01 < \alpha_{med} < 1$ (partly attentive)	3%	22%	4%	12%
	$\alpha_{med} = 1$ (fully attentive)	2%	28%	3%	11%

$*$ The value of one is taken to be $0.9 \leq \beta \leq 1.1$.

Figure 4. Heterogeneous agents' responses to information (Henckel *et al.*, 2022).

version of Bayes' rule, in which the likelihood ratio is raised to the power β. The RE benchmark corresponds to the special case where $\alpha = 1$ and ($\beta = 1$). Figure 4 (Table 5 in Henckel *et al.* (2022, p. 54)) summarizes their results, indicating that subjects display both belief conservatism ($\alpha < 1$) as well as attenuated responsiveness when updating ($\beta < 1$). These findings suggest that, while rational expectations may remain a useful benchmark when the null is rejected, models with heterogeneous agents likely require a richer set of belief-updating behaviours to capture observed deviations.

(iii) Microfoundations of inferential expectations

Inferential expectations should be interpreted as an 'as if assumption', as it is impossible to directly observe human subjects conducting hypothesis tests in response to new information. However, there is good reason to believe that such hypothesis-testing behaviour, which has also been posited in the psychology literature, is an innate feature of human decision-making.

In the philosophy of science, Mayo (1996) and Mayo and Spanos (2006) argue that almost no one, and especially not a conventional scientist, adopts a Bayesian stance of holding multiple theories at once with probabilities attached to each of them. Instead, they argue (in the vein of Kuhn (1970)) that a theory is held, and tenaciously so, until a cumulation of contrary evidence amasses. This is of course the methodology of a hypothesis test, albeit writ large for entire frameworks, or paradigms.

At a more individual decision-theoretic level, Cohen *et al.* (2019) provide a micro-foundation for two-sided hypothesis tests. In their model, agents conduct inference in response to a sequence of signals, subject to a switching cost whenever beliefs change. They show that the optimal solution is a band of inaction similar to a confidence interval (and therefore to a hypothesis test).[3] The exact nature of the switching costs is irrelevant; for example, they may be menu costs, transaction costs, or cognitive (effort) costs associated with attention and observation or the consultation of experts. Carroll (2003) is an example of the latter where relatively inattentive agents occasionally

[3] Any rule that denotes significance outside of a $100(1-\alpha)$ % confidence interval is equivalent to a two-sided hypothesis test of size α %.

consult 'experts'. If experts are modelled as fully rational agents, RE is attained at these infrequent moments of consultation, a feature common (though not necessary) in IE models.

(iv) Comparison to other theories of expectations formation

Early macroeconomic models assumed adaptive expectations, introduced by Cagan (1956) and Friedman (1957), to capture how agents forecast in an uncertain environment. Under this framework, expectations about future values of economic variables, such as inflation or interest rates, are based on past observations. The core assumption is that individuals adjust their expectations gradually in response to past forecast errors rather than incorporating all available information instantaneously. Formally, if $E_t[X_{t+1}]$ represents the expected value of a variable X at time $t+1$, conditional on information at time t, agents update their expectations according to the equation

$$E_t\left[X_{t+1}\right] = E_{t-1}X_t + \lambda\left(X_t - E_{t-1}X_t\right),$$

where $0 < \lambda \leq 1$ is the adjustment coefficient determining the weight given to new information. This backward-looking mechanism implies that expectations adjust slowly and can persistently lag behind actual economic conditions, leading to systematic forecast errors.[4] Adaptive expectations played a key role in early Phillips curve models, underpinning the notion of a short-run trade-off between inflation and unemployment that policymakers could exploit. However, this approach was later challenged by the rational expectations revolution, which argued that agents form expectations by efficiently utilizing all available information, not merely historical data, leading to unbiased forecasts on average.

While adaptive expectations imply that expectations are slow to adjust—in other words, they exhibit some stickiness—they do adjust incrementally with each new observation of the underlying variable and, crucially, discrete jumps are ruled out.

In recent years, a variety of theories have emerged to explain deviations from the rational expectations hypothesis, reflecting a growing consensus that real-world agents form and update expectations in ways that are more complex, context-dependent, and subject to cognitive and informational constraints. Two frameworks in particular have made inroads, rational inattention models and diagnostic expectations. There are some similarities between IE and these frameworks, especially with the theory of rational inattention, but there are some important differences that are worth highlighting here.

The theory of rational inattention, developed by Sims (2003) and subsequent contributors, assumes that agents face constraints on their ability to process information due to cognitive or technological limitations. Because acquiring and processing information is costly, agents optimally allocate their attention to the most relevant variables, deliberately ignoring less critical signals. Rational inattention provides a formal framework for understanding how information frictions can lead to sluggish or incomplete adjustments in expectations, as agents rationally filter out certain details to conserve their cognitive resources.

In contrast, IE does not presuppose such deliberate attention-allocation constraints but instead focuses on how agents use available information to infer the likely state of the world. While rational inattention treats information-processing as a constrained optimization problem, IE emphasizes the cognitive process of interpreting information signals within the context of pre-existing beliefs and the broader economic environment. For example, under IE, an agent might place disproportionate weight on a new signal not because they are inattentive to other information, but because it is the sum of all signals which in their totality challenge the agent's prior belief, thus precipitating a discrete change.[5]

[4] A common specification involves setting $\lambda = 1$ so that $E_t[X_{t+1}] = X_t$, that is, the error-adjustment term disappears and the expectation of X in the next period is simply equal to the observed value of X in the current period.

[5] A simple comparison between rational inattention models and IE is difficult because the rational inattention framework has itself evolved. For example, IE assumes that information is received continuously but only processed partially, due, say, to cognitive limitations. The delays in the agents' responses are directly related to the information's cumulative perceived salience, the latter being determined by the significance level used for the hypothesis test. A key implication of IE is that expectations are state-dependent, whereas in the early vintage models of rational inattention *à la* Mankiw and Reis (2002, 2003), also known as sticky information models, expectations are time-dependent. State-dependence, apart from being intuitively more appealing, enables the modeller to address a much richer menu of issues such as anticipation effects of future events and endogenously generated regime shifts. Rational inattention models based on Shannon

This distinction has important implications for understanding structural shifts. Rational inattention may struggle to account for abrupt changes in beliefs, as these would require a sudden reallocation of attention that rational inattention frameworks typically model as gradual. By contrast, as explained earlier, IE naturally allows for abrupt shifts in beliefs.

Diagnostic expectations, as proposed by Bordalo *et al.* (2018, 2022), posit that agents overreact to salient or recent information, leading to expectations that are excessively sensitive to short-term trends. This behaviour is rooted in cognitive biases, particularly the representativeness heuristic, which causes agents to overweight the likelihood of outcomes that seem typical or consistent with recent observations. Diagnostic expectations have been applied to explain phenomena like bubbles, booms, and busts, where agents' overreaction to temporary signals amplifies economic fluctuations.

IE differs from diagnostic expectations in several key respects. While diagnostic expectations focus on systematic cognitive biases, IE emphasizes a more deliberate reasoning process, where agents weigh new evidence against their prior beliefs. IE does not generate a consistent overreaction to new information; instead, it tends to generate a protracted period of underreaction, followed by a sudden overreaction. For instance, in the case of inflation expectations, diagnostic models might predict an exaggerated response to a single period of high inflation, whereas an IE agent might not respond at all to a single period of high inflation.

Moreover, diagnostic expectations tend to presume a common behavioural bias across agents, leading to coordinated errors that amplify macroeconomic dynamics. In contrast, IE can, in theory, accommodate heterogeneity, as different agents may draw distinct inferences based on their (possibly differing) degrees of belief conservatism, prior beliefs, and idiosyncratic signals.

(v) Examples of applications

Belief in climate change

Public belief in climate change, particularly in its anthropogenic causes, has varied widely across countries and demographic groups. While scientific consensus on human-induced climate change is strong, many individuals' beliefs are shaped more by observable events and personal experiences than by abstract data or expert opinion. IE can serve as a useful metaphor to think about how individuals update their beliefs about climate change in response to extreme weather events, such as hurricanes, bushfires, and heatwaves. This dynamic provides a vivid case where individual-level belief adjustments can aggregate into a structural shift in societal attitudes, policy demands, and economic behaviour.

Extreme climate events often serve as pivotal data points that individuals interpret to update their beliefs. IE suggests that people use these events to infer causality and likelihood, particularly when prior beliefs were uncertain or weakly held. For instance, people who personally experience extreme events, like flooding or droughts, are more likely to attribute these events to climate change. And coverage of climate disasters influences belief updates, even among those not directly affected. Research has found that individuals who experience a string of extreme weather events are more likely to believe in climate change. For example, in their meta-study, Howe *et al.* (2019) show that US residents affected by record-breaking heatwaves exhibited higher acceptance of climate change. Following Hurricane Sandy in 2012, public concern about climate change spiked in the US, particularly in affected regions. This pattern of belief updates has been observed globally after high-profile disasters (see also Thomas-Walters *et al.*, 2024).

IE would explain these shifts as follows: individuals observe extreme climate events or their media representations. Based on the cumulative salience and emotional impact of these events, they update their prior beliefs about climate change and its causes. As a critical mass of individuals infer new beliefs, public attitudes shift, influencing policy debates, market dynamics (e.g. growth in green energy demand), and global climate negotiations.

information theory (e.g. Sims) are state-dependent and thus closer in spirit to IE but, as Sims (2011) readily admits, pose an immense technical challenge [p. 178]. One of the key strengths of IE is its intuitive appeal and technical simplicity.

Balance-of-payments crises

History has shown how fickle global capital flows are and how costly sudden adjustment can be. Large capital inflows, resulting from investors' search for yield, depress risk premia, flood the economy with liquidity, and lead to an appreciation of the exchange rate or active reserve management by the monetary authority. However, at a whim, investors may head for the exit, looking for a safe haven and forcing the country to quickly adapt to the reversal of capital flows. This is usually a disorderly and painful process.

Understanding how and when capital flows reverse is therefore central to understanding the causes and consequences of international financial crises and to improving policymakers' alleviation and management of these disruptive episodes.

The academic literature on the causes of international capital account crises broadly proposes two competing theories (see Chui and Gai (2005) for a concise summary). One view, so-called first-generation currency crises models, emphasizes the importance of a country's weak fundamentals as the trigger of a crisis. A sudden capital outflow is a fully rational decision by investors; there is no 'panic' or belief-driven exit, and a crisis is easily justified based on the underlying fundamentals (Krugman, 1979; Flood and Garber, 1984).

Second-generation currency models (e.g. Obstfeld, 1996; Sachs *et al.*, 1996) assume speculative attacks are largely self-fulfilling and therefore random, with sound fundamentals being neither a necessary nor sufficient condition for averting a crisis. An oft-cited example is the 1997 Asian financial crisis which was sudden and unforeseen.

An alternative perspective on the Asian Crisis, which naturally leads to IE, stresses agents' sudden realization about the extent of the fiscal authorities' contingent liabilities. In the narrative of Corbett and Vines (1999), Burnside *et al.* (2001), and Irwin and Vines (2003), agents' demand for Asian currencies collapsed due to a perceived decline in the creditworthiness of their banking sectors and a corresponding liability for the fiscal authorities. IE is able to build an intellectual bridge between first- and second-generation currency models: expectations are ultimately linked to fundamentals, but belief conservatism allows a protracted disconnect between expectations and fundamentals, making crises hard to predict.

Furthermore, an IE balance-of-payments model naturally accommodates the econometric challenge associated with low probability, tail risk events. For example, the well-known peso-problem, that is, the small probability of a large change contained in an asset price, can be parameterized using appropriate estimates of investors' belief conservatism.

Rise of awareness of the Covid-19 pandemic

The rapid emergence of Covid-19 in early 2020 provides another example of how inferential expectations can explain abrupt shifts in individual and collective belief formation. At the outset, information about the virus was scarce, and many individuals and policymakers downplayed its risks, either due to lack of direct exposure or because past outbreaks, such as SARS and MERS, had limited global impact. However, as cases and fatalities began to rise in multiple countries, individuals updated their beliefs based on accumulating data, anecdotal evidence, and government responses.

IE suggests that people's perceptions of risk evolved not through continuous information absorption but through discrete inferential jumps, occasionally triggered by highly salient events such as overwhelmed hospitals, lockdowns, or personal experiences with illness. Early in the pandemic, many individuals revised their beliefs suddenly upon encountering reports of exponential case growth or witnessing government-imposed restrictions, even if they had previously been sceptical. Studies, such as those by Fetzer *et al.* (2021) and Aksoy *et al.* (2022), document how belief formation during the pandemic was strongly influenced by exposure to local outbreaks, media framing, and government communication strategies.

The inferential expectations framework is particularly useful for explaining why some populations exhibited delayed or reluctant belief shifts, especially in cases where prior expectations were strong or when conflicting signals (e.g. misinformation, policy inconsistencies) led to hesitancy in updating beliefs. The IE approach can be extended to model belief formation during pandemics

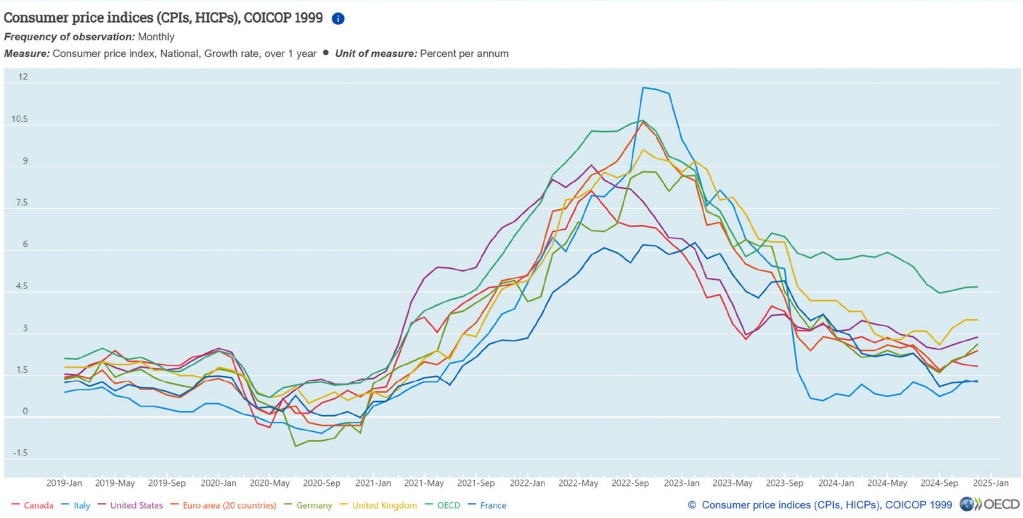

Figure 5. Inflation rates in major developed economies.

in general, helping to explain why risk perceptions fluctuate dramatically and why some groups remain resistant to updates despite mounting evidence.

De-anchoring of inflation expectations

The recent post-Covid inflation episode provides an interesting example of the need to avoid structural adjustments in beliefs. Standard macroeconomic models assume that three factors determine inflation:

$$\text{inflation} = \text{inflation expectations} + \text{demand conditions} + \text{supply conditions}$$

The recent surge in inflation (see Figure 5) was attributed both to excess demand (because of expansionary monetary and fiscal policy) and adverse supply shocks (gummed up supply chains and high energy prices).

The critical problem for central banks has been how quickly to return inflation to its desired level, lest an extended period of high inflation raises the private sector's inflation expectations. The dangers of an overly slow return can be represented by the conditions under which IE predicts a structural change in beliefs—abandoning belief in the inflation target.

Suppose that naïve agents form expectations of inflation using a simple arithmetic average and conduct an IE hypothesis test to determine whether this is consistent with the target. In any hypothesis test like this, the null will be abandoned if the following test statistic is too large. Let the hypothesized value of the inflation rate (the cognitive target) be denoted by π^{H_0}, the average of the observed inflation rates by $\bar{\pi}$, the standard deviation by σ_π, and the number of observations by n. Then the corresponding z-statistic will be given by,[6]

$$z = \frac{\bar{\pi} - \pi^{H_0}}{\sigma_\pi / \sqrt{n}}. \tag{1}$$

A key policy insight of IE is that a structural change in beliefs can occur in (at least) two ways. The z-statistic can be large if a single large inflation gap raises $\bar{\pi} - \pi^{H_0}$. However, many small departures from target will, if sustained for sufficiently long, drive up z via the $1/\sqrt{n}$ term in the denominator, even if $\bar{\pi} - \pi^{H_0}$ is small. Small inflation deviations matter *if sustained for a sufficiently long time* which is the nub of the policy dilemma faced by central banks wanting to get back to target.

[6] For simplicity, we are employing large sample tests, using the population standard deviation. One could easily assume small sample sizes instead and use a t-statistic, with the sample standard deviation.

What happens to the inflation process if credibility is lost? Henckel *et al.* (2019) explored this econometrically. They posited an IE Phillips curve of the following form:

$$\pi = \beta \pi^{H_0} + (1-\beta) E[\pi] + \sum_i \gamma_i X_i + \eta. \tag{2}$$

Aggregate inflation π depends on the inferential expectations of price-setters, a fraction β of whom fail to reject H_0 and a fraction $1-\beta$ of whom reject H_0 and work out the rational expectation of inflation, based on (1). The pronumeral π^{H_0} is the value inflation takes under the null and $E[\pi]$ is the mathematical expectation of π, the rational expectations solution. The X_i are other explanators, variables like the output gap, which are zero in the steady-state. The estimation error η is assumed to be i.i.d. They simplified the above equation by noting that, when agents adopt rational expectations, they can use (2) to solve for $[\pi]$:

$$E[\pi] = \beta \pi^{H_0} + (1-\beta) E[\pi] + \sum_i \gamma_i X_i \quad \Rightarrow \quad E[\pi] = \pi^{H_0} + \sum_i \frac{\gamma_i}{\beta} X_i, \tag{3}$$

which can be substituted into (2) to obtain an IE Phillips curve:

$$\pi = \pi^{H_0} + \sum_i \frac{\gamma_i}{\beta} X_i + \eta. \tag{4}$$

Equation (4) is noteworthy because the marginal impacts of the X_is depend upon the proportion of agents that are anchored onto H_0. The more they abandon the inflation target (the lower β) the more they must rely on other variables for their expectations, which magnifies their marginal impacts. Thus, a testable implication of the IE Phillips curve is that all the coefficients on the X_is will rise by the same proportion when agents abandon the null. This was in fact what Henckel *et al.* (2019) found in the Euro area in the aftermath of the European sovereign debt crisis. It is a striking feature of many crises that previously unimportant variables can become more important, and (4) is a new (and testable) model of why this might be so.[7]

The novelty and significance of the above point is worth emphasizing: the β-term interacts with the coefficients of the explanatory variables. Different βs thus generate different slope estimates for the explanators, making the coefficients state-dependent. In the extreme, where β falls from one and approaches zero, the parameter acts like an interactive dummy variable which switches the γ_i coefficients from standard effects to unboundedly large ones. Crucially, this dummy is endogenous, *not* exogenous.

Numerous other studies have also found that the slopes of Phillips curves worldwide have steepened post-pandemic, almost as though the Phillips curve has been awakened out of its decade-long slumber. For example, Hobijn *et al.* (2023) compare the Phillips curves pre-pandemic to the post-pandemic recovery phase for 29 countries and identify significant shifts in their slopes, summarized in Figure 6.

As the above discussion makes clear, in an IE world, the Phillips curve coefficients change whenever β, the proportion of agents who adhere to their null beliefs, changes. Sudden, possibly large shifts in beliefs, will correspondingly be reflected in large changes in the relationship between the dependent and independent variables and thus in large shifts in the estimated coefficients.

III. An inferential expectations model of the post-Covid inflation surge

The previous section on the de-anchoring of inflation expectations was primarily illustrative and, in the process, highlighted important properties of econometric models involving inferential expectations. In this section, we build on the previous section by presenting a simple, bare-bones 3-equation macroeconomic model to explain the implications of inferential expectations (IE) and,

[7] The IE model thus adopts the last three recommendations of the 'Rebuilding Macroeconomic Theory Programme' (Vines and Wills, 2018): limiting the reliance on rational expectations; assuming the existence of heterogeneous agents; and having more realistic microfoundations (like Carroll (2003)). The project, described in Vol. 34, No. 1 of this journal, is a summary of the written and verbal contributions by the issue's many authors. A comment that the econometric explanatory power of variables can intensify during a crisis, as we have modelled in (3) above, is attributed to Olivier Blanchard.

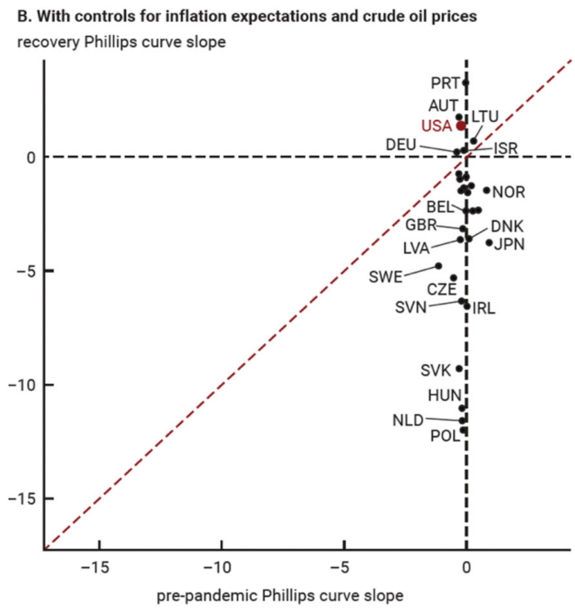

Figure 6. Phillips curve slopes for 29 countries: pre-pandemic and recovery periods.
Notes: The three-letter country codes from the International Organization for Standardization (ISO) are available online. Data points are Phillips curve slopes for countries in the pre-pandemic period (2013: Q1–2019: Q4) and the recovery period (2021: Q1–2022: Q2). The red dashed line is the 45-degree line. Control variables for inflation expectations and (Brent) crude oil prices are included in the estimation.
Source: Hobijn *et al.* (2023).

more to the point, provide a useful way to think about the recent surge, and subsequent decline, of inflation post-Covid-19.

A central debate in both policy-making circles and academic economics concerns the appropriate response of central banks to inflationary supply shocks. The post-pandemic period is widely characterized as an adverse supply shock, driven by disrupted global supply chains and exacerbated by the sharp rise in energy prices following Russia's invasion of Ukraine. The question of how forcefully central banks should react to such a shock remains contentious. The two extreme points of view were forcefully argued by Joseph Stiglitz and Paul Beaudry, respectively, in a seminar at a conference on post-pandemic macroeconomics held on 22 and 23 March 2024 in the School of International and Public Affairs at Columbia University.

Stiglitz argued that the inflationary surge was primarily transitory, driven by supply-side distortions that would self-correct without necessitating an aggressive monetary policy response. According to this perspective, excessive tightening in response to a short-lived inflationary shock could needlessly stifle economic recovery and employment growth.

Beaudry took a more cautionary view, emphasizing the risks of allowing even a temporary inflationary surge to take hold. The primary concern is that inflation shocks, while potentially transient, may still be large and persistent enough to de-anchor inflation expectations. Once expectations become unmoored from the central bank's target, inflation dynamics can become self-perpetuating, making disinflation significantly more costly. This raises a critical policy trade-off: should central banks prioritize immediate price stability, even at the risk of recession, or should they tolerate short-term inflation to allow the economy to normalize organically?

Inferential expectations provide a handy vehicle to shed light on this problem, by considering how the anchoring and de-anchoring of inflation expectations affect the macroeconomy during an inflationary surge and how they subsequently adjust as inflation declines. These dynamics have direct implications for monetary policy, particularly regarding the speed and intensity of interest rate adjustments in response to inflationary pressures.

To examine these dynamics, we propose a minimalistic 3-equation macroeconomic model, employing numerical simulations to explore the *qualitative* behaviour of inflation expectations under a couple of different scenarios. Our objective is not to provide a fully calibrated empirical model but rather to outline a theoretical foundation that can inform future empirical research.

Consider a standard expectations-augmented Phillips curve of the form:

$$\pi_t = \pi_t^e + \theta x_t + u_t, \tag{5}$$

where π is inflation, π^e is expected inflation, x denotes a vector of demand conditions, as a deviation from a non-stochastic steady-state, and u is an autoregressive error term:

$$u_t = \rho u_{t-1} + \eta_t, \quad \rho \in (0, 1),$$

where ρ is the autoregressive parameter and η an i.i.d random process. Henceforth, we reduce x to a simple scalar, namely the output gap x.

Following Fuhrer and Moore (1995), we assume that inflation expectations are a convex combination of adaptive (backward-looking) expectations and forward-looking expectations:

$$\pi_t^e = (1-\gamma)\pi_t^{AE} + \gamma \pi_t^{FE}, \tag{6}$$

where the superscript AE denotes adaptive expectations and FE forward-looking expectations, while the parameter γ is the fraction of agents with forward-looking expectations. In principle, one could include a fixed component, such as an explicit inflation target π^T, in the π^{AE} term, or introduce multiple lags, or allow for partial adjustment (as in section II(iv)). However, for simplicity, we assume that adaptive expectations take the simple autoregressive form:

$$\pi_t^{AE} = \pi_{t-1}. \tag{7}$$

Typically, in standard expectations-augmented Phillips curve models, the forward-looking term π^{FE} is assumed to be given by rational expectations (RE), whereby agents form expectations using all available information in an unbiased manner. Instead, we introduce inferential expectations (IE) as the mechanism by which forward-looking agents form their inflation expectations, albeit with a possible delay:

$$\pi_t^{FE} = \pi_t^{IE}. \tag{8}$$

(i) Inferential expectations

Agents begin with a null hypothesis about inflation that aligns with the central bank's official inflation target:

$$H_0 \ : \ \bar{\pi} = \pi^T, \tag{9}$$

where $\bar{\pi}$ is the (sample) mean of the inflation rate. Thus, agents initially expect inflation to remain close to the target set by the monetary authority. However, rather than taking this expectation as fixed or predetermined, they engage in a hypothesis-testing process to assess whether the history of observed inflation rates is consistent with this assumption. In other words, agents continuously evaluate whether the inflation target remains a valid predictor of actual inflation outcomes.

To formalize this process, let the agents construct a standard hypothesis test, where the test statistic is based on the deviation of the sample mean of realized inflation, $\bar{\pi}_t$, from the official target. Specifically, assume that agents employ a z-statistic, given by

$$z_t = \frac{\bar{\pi}_t - \pi^T}{\sigma_\pi/\sqrt{n}}, \tag{10}$$

where σ_π, as before, denotes the standard deviation of inflation, and n denotes the number of observations used in the hypothesis test. The denominator captures the precision of the observed inflation deviations, increasing as the number of observations grows.

Agents apply classical statistical inference: the null hypothesis (H_0) is rejected if the test statistic exceeds a critical threshold, meaning that the observed inflation outcomes have deviated significantly from the official target. This rejection can occur under two key conditions:

(i) *Magnitude Effect*: if, for a given σ_π and n, the absolute gap between realized mean inflation $\bar{\pi}_t$ and the target π^T becomes too large, the z-statistic will exceed the critical threshold;
(ii) *Persistence Effect*: even if individual deviations are small, a prolonged sequence of deviations from the target increases n, raising the test statistic through the $1/\sqrt{n}$ term. Over time, if inflation persistently fails to return to target, the cumulative evidence may be sufficient to reject H_0.

The threshold value of the z-statistic, at which the null hypothesis is marginally rejected for a two-sided hypothesis test, is denoted by $z^*_{\alpha/2}$, where α is the significance level (equivalently, the test size).[8] Thus H_0 is maintained if $z \leq |z^*_{\alpha/2}|$, and it is rejected otherwise.[9]

If H_0 is rejected, then we assume that IE agents switch to the alternative hypothesis, H_1, namely that inflation expectations are rational:

$$H_1 \quad : \quad \pi_t = E\left[\pi_t\right], \tag{11}$$

where $E[\cdot]$ denotes the mathematical expectations operator, conditional on $t-1$ information.

(ii) Solving for inflation

The purpose of the 3-equation model is to provide a theoretical structure for thinking about the post-pandemic inflation surge, in particular the difficulty this posed for the world's central banks. Before simulating the economy, identifying the basic mechanics, and solving for inflation in equilibrium, helps to build some intuition.

Assume the economy starts in a steady-state with $x = 0$, $\pi = \pi^T$, and $i = r^n + \pi^T$. Schematically, we can divide the inflation surge into two phases:

(i) *UP-Phase*: A supply shock (equivalently, a Phillips curve shock), e.g. a positive innovation in η, pushes up π and, in subsequent periods, also π^e due to the presence of *AE*-agents. Depending on how long inflation deviates from target and how large the inflation gap is, *IE*-agents will maintain or reject the null hypothesis. If they maintain it, they continue to assume that $\bar{\pi}_t = \pi^T$, in other words, their expectations remain well-anchored. If they reject H_0, they switch to H_1 instead. In other words, their expectations have become unanchored, and inflation will rise further because now π^e jumps up the moment *IE*-agents switch from H_0 to H_1.
(ii) *DOWN-Phase*: This is the adjustment phase, when the inflation innovation has passed and the central bank is responding to the inflation gap by raising the interest rate to bring inflation back to target faster. Clearly, the higher are inflation expectations, the more difficult the inflation-busting job for the central bank becomes. Alternatively, *ceteris paribus*, if *IE*-agents' expectations remain anchored (they never rejected H_0 in the *UP-Phase*), the task of bringing inflation back to target is easier than if *IE*-agents' expectations had become unanchored. (By 'easier' we mean less costly in the sense of a lower sacrifice ratio.) This provides a rationale for central banks to be relatively aggressive early on.

When *IE*-agents maintain the null hypothesis, the PC becomes

$$\pi_t = \pi_t^e + \theta x_t + u_t = (1-\gamma)\pi_{t-1} + \gamma\pi^T + \theta x_t + u_t. \tag{12}$$

When *IE*-agents reject the null hypothesis and switch to RE, the PC becomes

$$\pi_t = \pi_t^e + \theta x_t + u_t = (1-\gamma)\pi_{t-1} + \gamma E\left[\pi_t\right] + \theta x_t + u_t. \tag{13}$$

[8] In many contexts it might be more natural to use a one-sided hypothesis test. For our purposes this does not matter.
[9] Though methodologically different, our analysis shares some similarity with Carvalho *et al.* (2023) who, in a New Keynesian model with statistical learning, study how a chronic overshooting of inflation targets leads to unanchoring of long-term expectations.

Solve for $E[\pi]$ by running the expectations operator through the entire equation above and collecting terms:

$$E\left[\pi_t\right] = \pi_{t-1} + \frac{\theta}{1-\gamma}x_t. \tag{14}$$

To obtain an expression for equilibrium inflation, substitute this back into the PC above:

$$\begin{aligned}\pi_t &= (1-\gamma)\,\pi_{t-1} + \gamma E\left[\pi_t\right] + \theta x_t + u_t \\ &= (1-\gamma)\,\pi_{t-1} + \gamma\left(\pi_{t-1} + \frac{\theta}{1-\gamma}x_t\right) + \theta x_t + u_t \\ &= \pi_{t-1} + \frac{\theta}{1-\gamma}x_t + u_t. \end{aligned} \tag{15}$$

Once H_0 is rejected and H_1 'adopted', one must take a stand on how inflation expectations are subsequently formed. One possibility assumes that expectations are fully rational forever after. Alternatively, one can assume that agents form a new null hypothesis and continue to hold onto this belief until, once again, after repeated hypothesis testing, this new null is rejected in favour of yet another alternative hypothesis. This assumption has the advantage of (*a*) being more consistent with the general notion that people are belief-conservative; (*b*) preserves the symmetry of agents holding onto a null both on the way up as well as on the way down; and (*c*) allows a much clearer analysis, and discussion, of the problem of the central bank *regaining* credibility on the way down. The disadvantage is that it is not obvious which rate of inflation the IE agents should latch onto and then, once again, how the (new) alternative hypothesis should be defined. This is an issue that needs to be settled comprehensively through further empirical analysis; however, the results in Henckel *et al.* (2022) provide a starting point.

For expositional purposes, we proceed by assuming that agents have inferential expectations also on the way down and that the new null is defined by the inflation rate at time $t = \tau$, the period in which IE agents reject the prior null, H_0, and solve for the rational expectation of inflation instead. In particular, the new null (in the *DOWN-Phase*) is given by

$$H_0^{'} \;:\; \bar{\pi} = \pi^{\tau} > \pi^{T}, \tag{9$'$}$$

where

$$\pi^{\tau} = \pi_{\tau-1} + \frac{\theta}{1-\gamma}x_{\tau} + u_{\tau}, \tag{16}$$

is the solution to the inflation rate at time $t = \tau$.

(iii) The remainder of the model

Two additional equations are necessary to complete the 3-equation macro model. First, we have the IS schedule,

$$x_t = E_t\left[x_{t+1}\right] - \psi\left(i_t - \pi_t^{e} - r^{n}\right), \quad \psi > 0, \tag{17}$$

where, as a reminder, x is the output gap, i the nominal interest rate, r^n the natural real rate of interest, and ψ a parameter.

Second, we have a Taylor rule describing how the policy rate is set as a function of the inflation gap and the output gap:

$$i_t = r^{n} + \pi^{T} + \varphi_{\pi}\left(\pi_t - \pi^{T}\right) + \varphi_x x_t, \quad \varphi_{\pi}, \varphi_x > 0, \tag{18}$$

where φ_{π} and φ_x are parameters, the weights attached to the inflation and output stabilization objectives, respectively.

(iv) Simulation results

To illustrate how inferential expectations (IE) operate within the 3-equation model, we adopt parameter values commonly used in New Keynesian frameworks. (See, for example, Galí (2015).)

Specifically, we set the natural real interest rate $r^n = 2$, the inflation target $\pi^T = 2$, the slope of the Phillips curve $\theta = 2$, and the Taylor rule coefficients $\varphi_\pi = 1.5$ and $\varphi_x = 0.5$.

The sensitivity of the IS equation to the policy rate is captured by $\psi = 0.15$, and the autocorrelation parameter ρ governing the shock process is set to 0.8. For the IE-specific parameters, we choose values that seem plausible: we set $\gamma = 0.2$ (the fraction of IE agents), $\sigma_\pi = 2$ (the standard deviation of inflation signals), and $\alpha = 0.05$ (the significance level in the underlying hypothesis test).

The simulation begins in a non-stochastic steady state where $\pi = \pi^e = \pi^T$, $r = r^n = 2$, and the output gap $x = 0$. An unanticipated shock, denoted η, occurs in period 10, causing inflation to deviate from the target. Figures 7a–c show the impulse response functions (IRFs) for inflation, which is our primary focus, for the output gap and for the nominal interest rate.

In this baseline scenario, shown in solid blue, IE agents observe the rise in inflation after the shock but continue to hold onto their initial null hypothesis $\bar{\pi} = \pi^T$ until a sufficiently large or persistent deviation from π^T accumulates in the data. In period 15, they reject the original null H_0 (based on the test formalized in section IV(i)) and update their expectations to a higher inflation rate. This discrete jump in expectations feeds quickly into realized inflation, which reaches another peak a couple of periods later, as the central bank steadily tightens monetary policy to bring inflation back down.

The rejection of the original null in period 15 produces a visible kink in the inflation path. Notably, IE agents reject the null *despite inflation already declining and monetary tightening under way*. This occurs because the z-statistic IE-agents use for their hypothesis test continues to increase, driven by both the *magnitude* and *persistence effects*, explained in section III(i). The policy implication is as novel as it is remarkable: even as headline inflation falls, underlying inflationary pressure from IE agents—unseen in the official data—continues to mount, only to be released as a renewed surge in inflation when IE-agents' inflation expectations become unanchored. Or, put differently, declining headline inflation need not signal relief: latent IE-based pressures, missed by official data, can build and erupt into a fresh inflation spike when IE agents' expectations lose their anchor.

When IE agents reject the original null on the way up, they switch to the rational-expectations solution, compute the implied inflation rate at that moment, and treat this new rate as their updated null hypothesis, H_0'. In the baseline simulation, that revised null is approximately 6.6%. From that point onward, they remain anchored at 6.6% until incoming inflation data persuade them otherwise. This can produce a plateau on the downward leg of the inflation path, because agents remain stuck at their updated null if realized inflation does not deviate enough from that level. When the cumulated evidence is eventually sufficient, in period 40 in this simulation, they reject the 6.6% null and revert to expecting inflation equal to the official target.

The IRF thus exhibits a second kink, marking the abrupt re-anchoring process. Inflation resumes its decline, accelerating somewhat after period 40, even overshooting the target level π^T, since IE agents have reverted to believing that inflation should converge back to the official target. Had they not rejected the null as early as period 40, inflation would have settled above target for an extended period. Eventually, as the sample size grows, IE-agents would reject the null hypothesis and inflation would return to the official target level, but in practical terms inflation could remain elevated for a considerable period.

A useful counterfactual in the same figure plots an alternative path (the dashed orange line) where IE agents never reject the null hypothesis of $\bar{\pi} = \pi^T$. In this scenario, inflation remains much more subdued. Central bank credibility, in this counterfactual, is never lost—IE agents continuously believe in the target, so the monetary authority faces a less formidable battle against inflation. In contrast, when IE-agents do reject the target and adopt a higher reference level, the resulting jump in expectations pushes inflation up more forcefully, so the policy tightening required to rein in inflation becomes more pronounced.

The central takeaway from this analysis is that inflation will only return to target once expectations re-anchor, which may take considerable time and evidence. Even as inflation declines following its peak, it can remain persistently above the official target if agents continue to

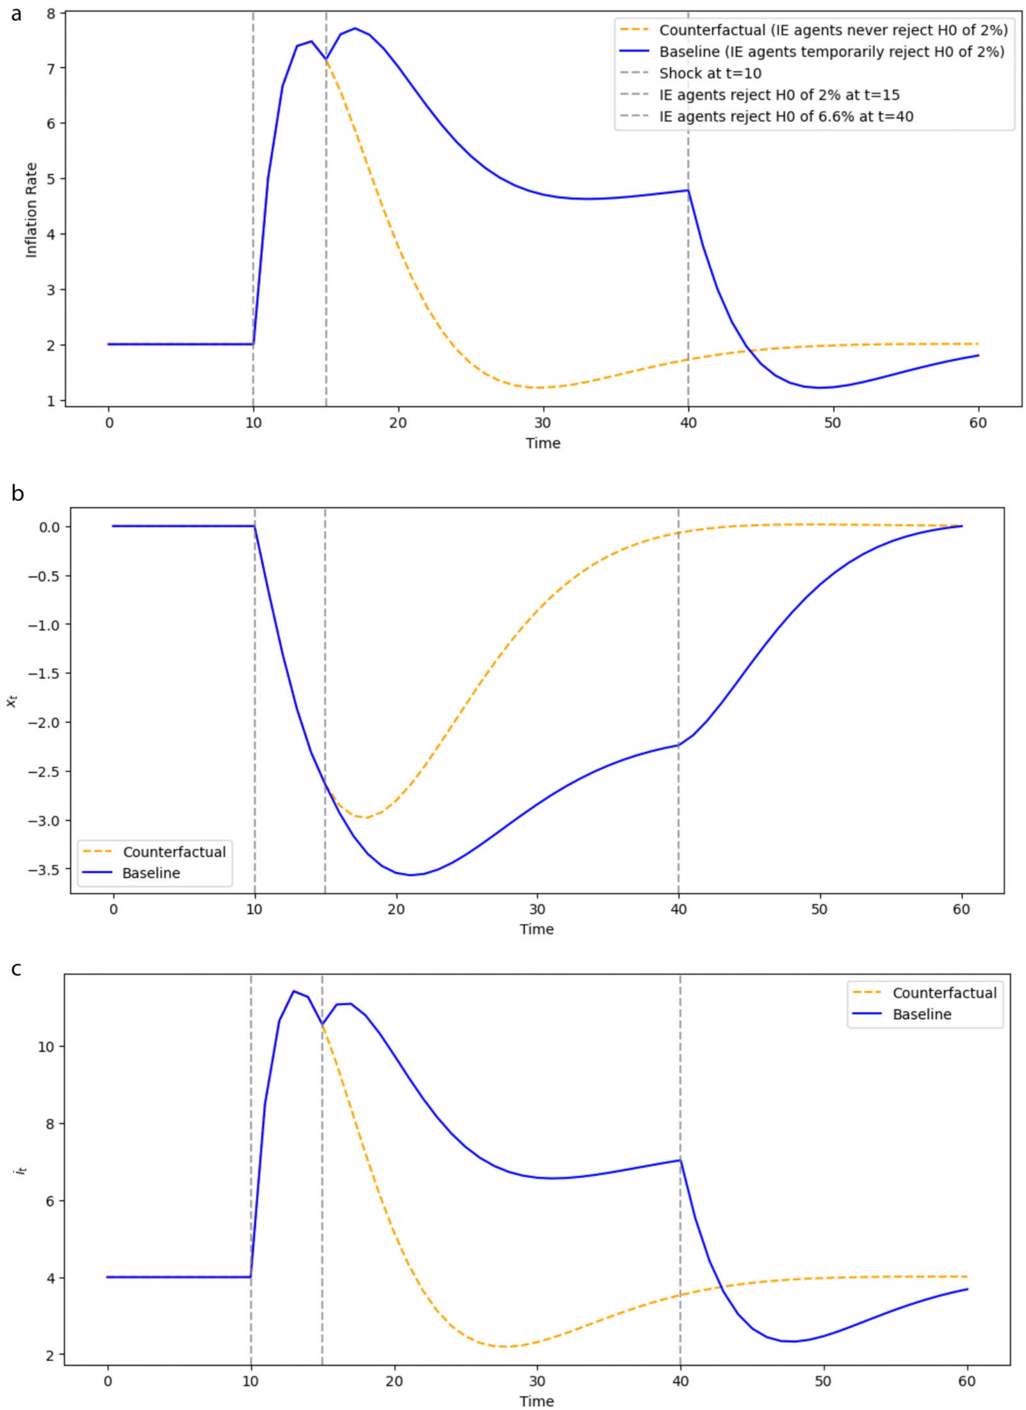

Figure 7. (a) Inflation, (b) Output gap, and (c) Nominal interest rate.

treat a higher rate as their reference point. This creates the risk of premature policy complacency: headline inflation may fall, but the underlying expectations that drive future inflation dynamics may remain misaligned. The disinflation process, therefore, is not complete until IE-agents decisively reject the elevated null and restore credibility to the central bank's target. Until that point, inflation may plateau at an undesirably high level, demanding sustained and potentially aggressive policy intervention. Observers have noted such a possibility in recent data, especially in the United States, where inflation initially came down from post-pandemic highs but seemed to stabilize at a rate above 2%. Baumann *et al.* (2024, 2025) make a parallel argument for Europe, remarking that although euro area inflation declined from its peak, there were signs that inflation was not fully anchored back to target levels, thereby complicating the last mile of disinflation.

The relationship between adaptive expectations and IE is valuable to understand how quickly inflation adjusts. If γ (the fraction of IE agents) is small, most individuals are backward-looking and do not strongly pull the inflation path back to the target, so realized inflation may rise quickly and prompt an early rejection of the null for IE agents. With relatively few IE agents, the switch in expectations, when it does happen, exerts only limited additional upward pressure on inflation. Conversely, if γ is large, the strong influence of IE agents keeps inflation closer to π^T initially, delaying any definitive rejection of the target on the way up. By the time the evidence against π^T becomes overwhelming, a large fraction of agents revise their beliefs, but that revision is smaller in magnitude because inflation was never allowed to deviate too far from the target in the first place. A similar pattern applies on the way down. In both phases, intermediate values of γ can produce more interesting, and potentially more volatile, dynamics, because the moderate balance between IE and AE agents ensures there is neither a quick jump in expectations nor a firm anchor.

To further explore the role of IE, we consider an alternative scenario where IE agents, after having rejected the old null, H_0, switch to rational expectations and continue to update their expectations using the rational-expectations solution each period until inflation has fallen sufficiently close to π^T. Only then do they switch back to the official target. This scenario is depicted in Figure 8a–c. Here, though difficult to see with the naked eye, there is a mild kink around period 25, when IE agents stop relying on the rational-expectations calculation and re-anchor to π^T. The orange line again displays the counterfactual where the IE null is never rejected; as in the baseline, that path is much more benign. Notably, inflation in this second scenario comes down more rapidly than in the baseline, avoiding the possibility of settling at a level above the target. This underscores how desirable it is, from a policy perspective, to prevent IE agents from anchoring to an alternative rate in the first place, or at least to ensure they re-anchor to π^T relatively quickly.

Naturally, these impulse response functions vary with the choice of parameter values. Different calibrations for α, γ, σ_π, and other model parameters can yield IRFs with distinct shapes or speeds of adjustment. A thorough sensitivity analysis would explore how each parameter affects the timing of rejections, the magnitude of expectation swings, and the persistence of inflation. For brevity, we do not exhaust those possibilities here as we are only aiming to provide a qualitative analysis. Instead, the main takeaway is that IE can generate plausible episodes where inflation temporarily settles at a level above the official target, forcing policy-makers to take account of the risk that expectations may remain stuck unless there is clear evidence of a robust disinflation.

Our model is a firm nod to Rowthorn's (1977) pioneering insight that inflation goes largely unnoticed until it reaches a threshold, at which point agents revise and act on inflation expectations.[10] Our model also relates to Bernanke and Blanchard (2023, 2024) and Michl and Rowthorn (2024), each of whom decomposes the expectations formation process into two components. In their study of the recent inflation surge in the US, Bernanke and Blanchard (2023) model expectations as comprising a fast-moving, quasi-rational component and a slower, backward-looking one anchored near the inflation target. This set-up allows for a discussion of the degree to which inflation expectations appear anchored or unanchored. For instance, if actual inflation rises sharply but agents treat it as transitory, then long-run inflation expectations remain stable, leading to

[10] We are grateful to David Vines for drawing our attention to Rowthorn's work and for the many valuable discussions we have had about it. To be fair, Rowthorn's model is much richer than this one, focusing on the distributional conflict between wages and profits. In his model, there may be a threshold value of the inflation rate at which workers switch from ignoring price inflation (relying on nominal representation) to incorporating it fully (using a real representation). Akerlof *et al.* (2000) and Blanchard (2018) make similar arguments. See Skott (2023) for a concise discussion.

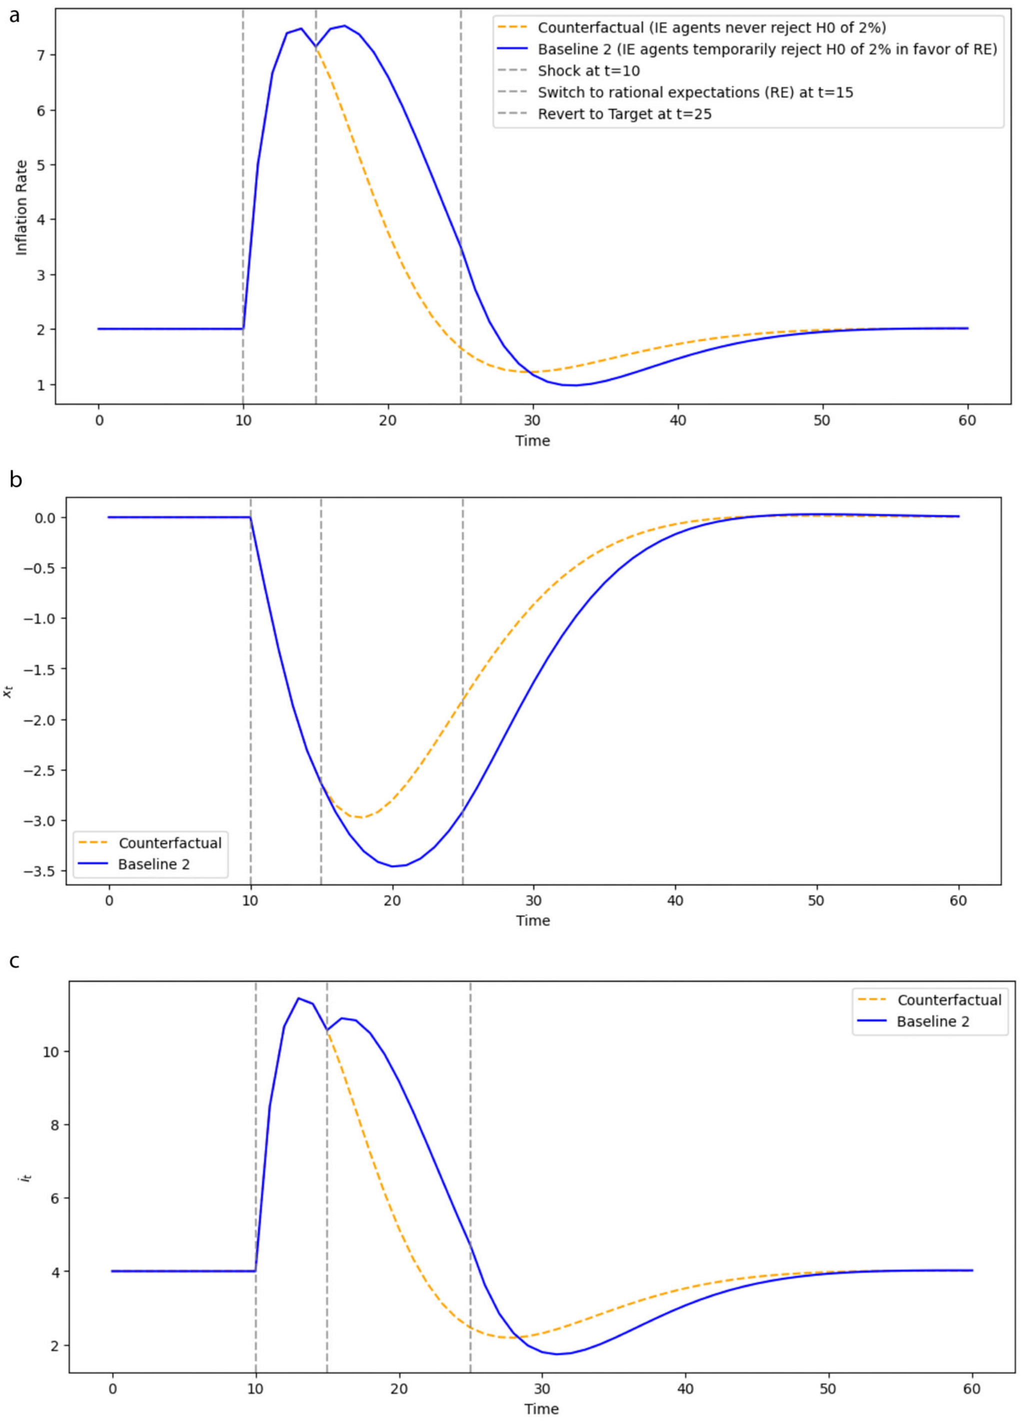

Figure 8. (a) Inflation, (b) Output gap, and (c) Nominal interest rate.

minimal second-round wage or price feedback. Conversely, if the shocks persist, or if the public perceives a change in policy stance, long-run expectations may drift upward, though Bernanke and Blanchard's empirical estimates suggest that, thus far, long-run expectations in the United States and many other advanced economies have not become dramatically unanchored (Bernanke and Blanchard, 2024).

Michl and Rowthorn (2024), by contrast, assume one group of agents is permanently anchored to the target, while another forms rational, forward-looking expectations. Our framework is closer to the latter but introduces a key extension: we allow the target-anchored group to revise their beliefs and switch to rational expectations when confronted with sufficient evidence. At the same time, unlike Michl and Rowthorn but akin to Bernanke and Blanchard, the remaining agents in our model form expectations in a backward-looking manner rather than rationally.

Like Bernanke and Blanchard, we highlight the outsized role that changes in expectations can play in driving inflation dynamics, yet our IE mechanism differs in its potential for sudden, non-linear belief revisions triggered by small signals. Bernanke and Blanchard's estimated equations do allow for some catch-up when price surprises occur, but they do not include a tipping-point dynamic in which a minor piece of evidence prompts agents to overhaul their entire inflation outlook. Instead, their model suggests a relatively smooth updating process, especially for the long-run component of inflation expectations.

(v) Implications for policymakers

A central welfare consideration under inferential expectations (IE) is the potentially steep cost of losing the inflation expectations anchor. Once IE agents reject their null hypothesis, $H_0 \ : \ \bar{\pi} = \pi^T$, and switch to a higher reference level, realized inflation rises further, compelling the central bank to respond with sharper monetary tightening. This dynamic lies at the heart of the policy debate between Beaudry and Stiglitz, referenced at the beginning of section III. Beaudry advocated for a pre-emptive tightening to prevent inflation expectations from becoming unanchored, thereby avoiding a costlier disinflation later. In contrast, Stiglitz maintained that expectations would remain anchored, allowing inflation to return to target without requiring forceful intervention, implying that the central bank could afford to adopt a more accommodating policy stance.

In standard New Keynesian set-ups, the social loss function is often approximated by a discounted quadratic form,

$$L_t = \sum_{k=0}^{\infty} \delta^k \left[\left(\pi_{t+k} - \pi^T \right)^2 + \lambda x_{t+k}^2 \right], \tag{19}$$

where π_t is the inflation rate, π^T the central bank's inflation target (as before), x_t the output gap (as before), $0 < \delta < 1$ a discount factor, and $\lambda > 0$ the relative weight on output-gap stabilization. Furthermore, assume that the central bank follows a standard Taylor rule, given by

$$i_t = r^n + \pi^T + \varphi_\pi \left(\pi_t - \pi^T \right) + \varphi_x x_t, \tag{20}$$

where $\varphi_\pi > 1$ and $\varphi_x > 0$. Deviations of inflation from π^T lead to a higher $(\pi_t - \pi^T)^2$ term. According to (20), to counteract the high inflation rate, the central bank must raise the interest rate, which worsens the output gap x_t and further amplifies the λx_t^2 term. This creates a costly trade-off between inflation and a negative output gap (Clarida *et al.*, 1999; Woodford, 2003).

Now suppose a policymaker seeks to choose the coefficient φ_π in the Taylor rule (20) to minimize the loss defined by (19). If the policymaker assumes that IE agents' inflation expectations remain anchored, they may choose a relatively low value of φ_π. Although this allows the inflation shock to persist for some time, it avoids large movements in the output gap x, thereby limiting the overall welfare loss.

However, if this assumption is incorrect and the inflation shock turns out to be large and/or persistent enough to cause IE agents to abandon their anchor, then the policy choice becomes significantly more costly. Once expectations de-anchor, bringing inflation back to target requires much larger and more prolonged movements in x, resulting in considerably higher welfare losses.

Anticipating this risk, a prudent policymaker may prefer to adopt a higher value of φ_π in (20), even if, at first glance, such a strong response appears excessive. This dynamic echoes findings in the broader New Keynesian literature that credibility and anchored expectations can be so valuable as to justify an initially stricter policy stance (e.g. Orphanides and Williams, 2007; Galí *et al.*, 2015; Gobbi *et al.*, 2019; Gáti, 2023; Christoffel and Farkas, 2025).

The trade-off between short-term and long-term costs hinges on how aggressively the central bank must respond in the near term to avert the unanchoring event. If the required rate increases are only moderate, then the net present value of forgoing a large future inflation spiral—and the associated stronger tightening later—can be net positive. However, if the shock is so large that staying below the IE rejection threshold requires an extremely severe immediate recession, policymakers might opt instead for a relatively standard Taylor response and bear the cost of de-anchoring. Quantitatively determining this threshold depends on structural parameters like γ (the fraction of IE agents), α (the significance level in their hypothesis test), and σ_π (the standard deviation of inflation innovations), which are all likely to be unknown to the central bank and very difficult to measure. Exploring such threshold effects within a linear-quadratic framework provides fruitful avenues for future research, further illuminating the interplay between credibility, the aggressiveness of policy intervention, and overall welfare outcomes (Clarida *et al.*, 1999; Orphanides and Williams, 2007). Our current research, which aims to estimate market level α's, is one step in this direction.

A more detailed comparison with the recent contributions by Beaudry *et al.* (2023, 2024) is instructive. In their 2023 paper, the authors examine the central bank's dilemma of choosing between looking through temporary supply shocks and aggressively controlling inflation expectations. Inflation expectations in their framework are explicitly forward-looking and endogenous: agents form expectations based on their perception of how the central bank responds to supply shocks. When supply shocks are sufficiently large or persistent, even temporary shocks can destabilize expectations if the central bank's response is perceived as inadequate. Agents' beliefs about future inflation hinge critically on their *interpretation* of the central bank's response, making expectations potentially sensitive to relatively modest policy missteps. Hence, Beaudry *et al.* (2023) underscore that central bank credibility and proactive response strategies are crucial to maintaining stable inflation expectations in the face of supply disruptions.

Beaudry *et al.* (2024) adopt a more structural approach to modelling the expectation formation process. They highlight the significant role of short-term inflation expectations influenced by broad-based supply shocks, particularly emphasizing bounded rationality and imperfect information. Agents do not possess fully rational expectations; rather, they rely on noisy signals about the economy, using disaggregated sector-level inflation data to form their expectations. In particular, agents interpret widespread sectoral inflation as a signal of a common inflationary component, termed broad-based supply shocks. When supply shocks affect multiple sectors simultaneously, agents infer (possibly mistakenly) that the economy-wide inflation rate has increased permanently, leading to elevated and persistent inflation expectations. This bounded rationality introduces persistent effects from what would otherwise be transient shocks, amplifying inflationary pressures through expectational channels. Therefore, even moderate but widespread supply shocks can significantly and persistently elevate inflation expectations, complicating central bank policy decisions.

What is unique to IE, in contrast to Bernanke and Blanchard's emphasis on shock magnitude and persistence mentioned earlier, and Beaudry *et al.*'s stress on broad-based shock perception, is the ability of even relatively small, single shocks to precipitate disproportionate inflationary episodes if they trigger abrupt belief revisions. The critical factor in our simulations is the inferential tipping point rather than the absolute scale or breadth of the shocks themselves. This distinctive feature of IE provides a complementary perspective, highlighting that inflation dynamics can become nonlinear and disproportionately sensitive to threshold-crossing signals.

This, in turn, has important policy implications. Bernanke and Blanchard (2023, 2024) advocate cautious but decisive monetary responses, leaning against the wind, to mitigate risks from persistent supply shocks and labour market overheating. They underscore that transparent communication and credible monetary policy are paramount to preventing long-run expectation de-anchoring. Similarly, Beaudry *et al.* (2023, 2024) highlight the critical central bank dilemma of managing expectations in the face of ambiguous supply shocks. If central banks 'look through' shocks deemed transitory, there is a risk of allowing inflation expectations to become unanchored if shocks persist longer than expected. Consequently, both Bernanke and Blanchard and Beaudry *et al*. argue for vigilant monitoring of expectations and cautious policy calibration.

Our results reinforce the importance of credible and responsive monetary policy, while the inferential expectations (IE) framework introduces an additional layer of caution. Even minor policy delays or misjudgements—by failing to recognize that signals might cross critical inferential thresholds—can trigger abrupt belief shifts and significant inflation volatility. This strengthens the case for more proactive and pre-emptive central bank communication to keep expectations anchored. It may also justify a brief, front-loaded policy response strong enough to avoid breaching these thresholds, thereby preventing more severe disruptions later. For analysts and policymakers, understanding the mechanisms that drive belief revisions is as important as identifying the magnitude and persistence of the shocks themselves.

IV. Discussion

This paper advances inferential expectations (IE) as a theoretically coherent and empirically supported mechanism for explaining sudden and significant shifts in economic beliefs. Economies rarely evolve smoothly; instead, they periodically experience abrupt and substantial changes. One possible avenue is through abrupt changes in expectations, reflecting the cumulative effect of evidence gathered by agents who conduct hypothesis tests. Under IE, agents retain their initial beliefs until accumulating information surpasses a critical threshold, triggering a sudden adjustment in expectations. The theoretical foundations provided here, complemented by empirical analyses and laboratory experiments, underscore IE's ability to account for diverse phenomena, including currency crises, inflation expectation de-anchoring, and the evolution of public beliefs about climate change and pandemics.

We list here four features of IE that make the theory attractive, especially in macroeconomic models where expectations are an important element:

(i) *IE allows for temporary departures of inflation from their fundamentals*
Because agents do not respond to each new piece of information, IE beliefs appear to be temporarily detached from the fundamentals but only until the last bit of information nudges the test statistic into the rejection region. With a high degree of belief conservatism this disconnect between expectations and fundamentals may last a considerable time. Ultimately, however, reality does catch up with agents and they are forced to revise their beliefs to consider the most current circumstances.

(ii) *IE can provide a simple metric of rationality*
In benchmark versions of the IE model, agents who reject H_0 and believe H_1 adopt RE. Therefore, RE becomes a special case of IE if agents are unconcerned about mistakenly changing their beliefs ($\alpha = 1$). This has three advantages: (i) expectations and the structure of the model cannot be permanently separated; (ii) in at least one case ($\alpha = 1$) an IE model inherits the stability properties of an RE model; and (iii) α becomes a metric for rationality: $\alpha = 0$ implies agents are unresponsive to evidence and $\alpha = 1$ implies they 'get it' immediately. In many contexts, the α can be estimated by observing the time it takes for an agent to 'get it'.

(iii) *IE is relatively tractable*
Assuming IE in economic models is relatively straightforward as most economists are already familiar with classical statistical hypothesis tests and therefore do not need to learn new tools. As explained in footnote 3, alternatives such as diagnostic expectations and rational inattention models based on Shannon information theory are technically quite demanding. There are two main challenges associated with IE: (i) the choice of hypotheses, test statistic and parameter values for which there may not always exist obvious empirical counterparts, and (ii) the solution to forward-looking structural models with IE may be

difficult to obtain given the discontinuity associated with the possible rejection of a null hypothesis.

(iv) *IE makes 'information policy' explicit*
A recurring concern in monetary policy is that central banks possess only limited direct instruments—notably, the short-term policy rate—to influence inflation and output. Yet in practice, central banks routinely deploy a broader arsenal, particularly through *information policy*: the strategic management of communication and disclosure. This includes forward guidance, publication of forecasts, and regulatory requirements for financial institutions to report specific economic indicators (Woodford, 2003, 2005). These tools aim not merely to inform, but to shape expectations and steer belief formation in financial markets and the broader public.
The IE framework provides a natural formalization of information policy, since it explicitly models how agents revise their beliefs in response to observed data. It highlights the importance of four key informational dimensions: the *volume* of data (sample size), the *quality* or reliability of each data point (its variance), the agent's *willingness to revise beliefs* (significance level α), and the *correlation structure* of the signals. From this perspective, the effectiveness of communication depends not only on how much information is made available, but on its statistical informativeness—e.g. whether aggregate inflation measures are stable enough to serve as credible signals, and whether the information being received is independent across time and sources.
Consequently, central banks cannot assume that more information automatically stabilizes expectations. If the data stream is noisy, erratic, or redundant, it may fail to trigger belief revisions when needed, or worse, provoke abrupt and destabilizing shifts. The IE framework thus implies that the success of information policy hinges on providing agents with timely, low-variance, and appropriately diverse signals. By shaping the informational environment in which expectations are formed, central banks can exert influence beyond the traditional interest rate channel, reinforcing the case for communication as an essential tool of macroeconomic stabilization.

IE exhibits some additional features that may be useful in a wider context:

(v) *IE can capture decision-making in information-poor environments*
In most economic models, fully rational agents use the distributions of underlying shocks to guess the state of the economy, in a process known as Bayesian updating. A classical statistical hypothesis test, however, requires less information, reducing the need to specify the model's stochastic properties.

(vi) *IE as a meta-theory*
IE receives support from scientific methodology. Kuhn (1970) observed that scientists will hold onto an old paradigm until the accumulation of anomalies appears significant; analogously, agents will hold onto H_0 in a hypothesis test until the accumulation of contrary evidence appears significant. More recently, Mayo (1996) and Mayo and Spanos (2006) criticize the reluctance of scientists to describe their own belief changes in terms of hypothesis testing, despite its widespread use. The key question here is that if we, as (social) scientists, habitually use hypothesis testing to help form our own beliefs—for example, by using econometric evidence—why wouldn't we assume that the agents in our models use something akin to hypothesis testing?

IE can be applied in wide and varied contexts, for example, to analyse central bank credibility and central bank responses (Henckel *et al.*, 2011, 2019) or to study the use of standards of proofs in legal decision-making (Lyons *et al.*, 2012). Several promising directions exist for future research. One natural extension would involve the integration of inferential expectations into larger, fully specified dynamic macroeconomic models, providing a richer environment to further explore policy design and implications. Additionally, studying heterogeneity among IE agents could offer more nuanced insights into market dynamics and belief coordination.

V. Conclusion

This paper has articulated the concept of inferential expectations (IE) as a micro-founded account of belief revision under uncertainty and has demonstrated why it is a useful ingredient of any economic analysis of structural change. Within the IE framework agents behave as if they conduct sequential hypothesis tests: they accumulate evidence, maintain their incumbent narrative so long as incoming signals remain within a 'confidence band', and switch abruptly once the test statistic breaches a critical threshold. Such threshold behaviour reconciles the empirical coexistence of slow-moving fundamentals with sudden breaks in macro time series—'the straw that breaks the camel's back'. By endogenizing the trigger point at which collective beliefs flip, IE provides a mechanism through which apparently discrete structural shifts can emerge endogenously from continuous, gradual forces, thereby complementing the more ad-hoc regime-switching assumptions often imposed in DSGE or Markov-switching models.

Applying IE to the formation of inflation expectations, the paper shows that welfare losses could rise disproportionately once a persistent string of price surprises crosses the rejection threshold, de-anchoring expectations from the central bank's target. Optimal policy may thus warrant an early, state-contingent departure from a mechanical Taylor rule, tightening just enough, and early enough, to keep the empirical test statistic below the critical value. This result squares with recent evidence that the post-pandemic inflation surge was first propagated by supply-side shocks but became persistent only when expectations began to adjust, magnifying the Phillips-curve trade-off (Bernanke and Blanchard, 2023; Beaudry *et al.*, 2024). By embedding an expectations channel that is both abrupt and endogenous, the IE framework helps explain why policy strategies aimed at managing expectations should account for the cumulative nature of evidence and the nonlinear dynamics inherent in belief updating.

References

Adler, J. E., and Rips, L. J. (eds) (2008), *Reasoning: Studies of Human Inference and Its Foundations*, Cambridge, Cambridge University Press. https://doi.org/10.1017/CBO9780511814273

Akerlof, G. A., Dickens, W. T., and Perry, G. L. (2000), 'Near-rational Wage and Price Setting and the Long-run Phillips Curve', *Brookings Papers on Economic Activity*, 2000, 1–44. https://doi.org/10.1353/eca.2000.0001

Aksoy, C., Eichengreen, B., and Saka, O. (2022), 'COVID-19 and Trust Among the Young', *Finance & Development Magazine*, International Monetary Fund, June.

Baumann, U., Ferrando, A., Georgarakos, D., Gorodnichenko, Y., and Reinelt, T. (2024), 'Safe to Update Inflation Expectations?', NBER Working Paper 32504, May.

— (2025), 'Do Inflation Expectations Become More Anchored During a Disinflation Episode? Evidence from Euro Area Firms', NBER Working Paper No. 33440, February.

Beaudry, P., Carter, T. J., and Lahiri, A. (2023), 'The Central Bank's Dilemma: Look Through Supply Shocks or Control Inflation Expectations?', NBER Working Paper No. 31741, September.

Beaudry, P., Hou, C., and Portier, F. (2024), 'The Dominant Role of Expectations and Broad-Based Supply Shocks in Driving Inflation', NBER Working Paper 32322, April.

Bernanke, B. S., and Blanchard, O. J. (2023), 'What Caused the US Pandemic-Era Inflation?', NBER Working Paper 31417, June.

Bernanke, B. S., and Blanchard, O. (2024), 'An Analysis of Pandemic-era Inflation in 11 Economies', *Peterson Institute for International Economics* WP24–11, May.

Blanchard, O. J. (2018), 'Should We Reject the Natural Rate Hypothesis?' *Journal of Economic Perspectives*, 32(1), 97–120. https://doi.org/10.1257/jep.32.1.97

Bordalo, P., Gennaioli, N., and Shleifer, A. (2018), 'Diagnostic Expectations and Credit Cycles', *The Journal of Finance*, 73(1), 199–227. https://doi.org/10.1111/jofi.12586

— (2022), 'Overreaction and Diagnostic Expectations in Macroeconomics', *Journal of Economic Perspectives*, 36(3), 223–44. https://doi.org/10.1257/jep.36.3.223

Burnside, C., Eichenbaum, M., and Rebelo, S. (2001), 'Prospective Deficits and the Asian Currency Crisis', *Journal of Political Economy*, 109(6), 1155–97. https://doi.org/10.1086/323271

Cagan, P. (1956), 'The Monetary Dynamics of Inflation', in M. Friedman(ed.), *Studies in the Quantity Theory of Money*, Chicago, IL, University of Chicago Press, 25–117.

Carroll, C. D. (2003), 'Macroeconomic Expectations of Households and Professional Forecasters', *The Quarterly Journal of Economics*, 118, 269–98. https://doi.org/10.1162/00335530360535207

Carvalho, C., Eusepi, S., Moench, E., and Preston, B. (2023), 'Anchored Inflation Expectations', *American Economic Journal: Macroeconomics*, 15(1), 1–47.

Christoffel, K., and Mátyás, F. (2025), 'Managing the Risks of Inflation Expectation De-anchoring', ECB Working Paper No. 3082.

Chui, M., and Gai, P. (2005), *Private Sector Involvement and International Financial Crises: An Analytical Perspective*, Oxford: Oxford University Press. https://doi.org/10.1093/0199267758.001.0001

Clarida, R., Galí, J., and Gertler, M. (1999), 'The Science of Monetary Policy: A New Keynesian Perspective', *Journal of Economic Literature*, 37(4), 1661–707. https://doi.org/10.1257/jel.37.4.1661

Cohen, S., Henckel, T., Menzies, G., Muhle-Karbe, J., and Zizzo, D. (2019), 'Switching Cost Models as Hypothesis Tests', *Economics Letters*, 175(February), 32–5. https://doi.org/10.1016/j.econlet.2018.11.014

Corbett, J., and Vines, D. (1999), 'Asian Currency and Financial Crises: Lessons from Vulnerability, Crisis, and Collapse', *The World Economy*, 22(2), 155–77. https://doi.org/10.1111/1467-9701.00198

Dornbusch, R. (1976), 'Expectations and Exchange Rate Dynamics', *Journal of Political Economy*, 84, 1161–76. https://doi.org/10.1086/260506

Fetzer, T., Hensel, L., Hermle, J., and Roth, C. (2021), 'Coronavirus Perceptions and Economic Anxiety', *The Review of Economics and Statistics*, 103(5), December, 968–78. https://doi.org/10.1162/rest_a_00946

Flood, R. P., and Garber, P. M. (1984), 'Collapsing Exchange-rate Regimes: Some Linear Examples', *Journal of International Economics*, 17(1–2), 1–13. https://doi.org/10.1016/0022-1996(84)90002-3

Friedman, M. (1957), *Theory of the Consumption Function*, Princeton, NJ, Princeton University Press. https://doi.org/10.1515/9780691188485

Fuhrer, J., and Moore, G. (1995), 'Inflation Persistence', *The Quarterly Journal of Economics*, 110(1), 127–59. https://doi.org/10.2307/2118513

Galí, J. (2015), *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications,* 2nd edn, Princeton, NJ, Princeton University Press.

Gáti, L. (2023), 'Monetary Policy and Anchored Expectations—An Endogenous Gain Learning Model', *Journal of Monetary Economics*, 140, S37.

Gobbi, L., Mazzocchi, R., and Tamborini, R. (2019), 'Monetary Policy, De-anchoring of Inflation Expectations, and the "New Normal" ', *Journal of Macroeconomics*, 61, 103070.

Gourinchas, P., and Tornell, A. (2004), 'Exchange Rate Puzzles and Distorted Beliefs', *Journal of International Economics*, 64, 303–33. https://doi.org/10.1016/j.jinteco.2003.11.002

Henckel, T., Menzies, G., Moffatt, P., and Zizzo, D. (2019), 'Three Dimensions of Central Bank Credibility and Inferential Expectations: The Euro Zone', *Journal of Macroeconomics*, 60(June), 294–308. https://doi.org/10.1016/j.jmacro.2019.01.012

— (2022), 'Belief Adjustment: A Double Hurdle Model and Experimental Evidence', *Experimental Economics*, 25, 26–67. https://doi.org/10.1007/s10683-021-09701-2

Henckel, T., Menzies, G. D., Prokhovnik, N., and Zizzo, D. J. (2011), 'Barro–Gordon Revisited: Reputational Equilibria with Inferential Expectations', *Economics Letters* 112(2), 144–47. https://doi.org/10.1016/j.econlet.2011.03.031

Hobijn, B., Miles, R. A., and Royal, J. (2023), 'The Recent Steepening of Phillips Curves', *ChicagoFedLetter*, No. 475, January.

Howe, P. D., Marlon, J. R., Mildenberger, M., and Shield, B. S. (2019), 'How Will Climate Change Shape Climate Opinion?' *Environmental Research Letters*, 14, 113001. https://doi.org/10.1088/1748-9326/ab466a

Irwin, G., and Vines, D. (2003), 'Government Guarantees, Investment, and Vulnerability to Financial Crises', *Review of International Economics*, 11(5), 860–74. https://doi.org/10.1046/j.1467-9396.2003.00422.x

Janis, I. L., and Mann, L. (1977), *Decision Making: A Psychological Analysis of Conflict, Choice, and Commitment*, New York, The Free Press.

Krugman, P. (1979), 'A Model of Balance-of-Payments Crises', *Journal of Money, Credit and Banking*, 11(3), 311–25. https://doi.org/10.2307/1991793

Kuhn, T. S. (1970), *The Structure of Scientific Revolutions,* 2nd edn, Chicago, IL, University of Chicago Press.

Lyons, B., Menzies, G. D., and Zizzo, D. J. (2012), 'Conflicting Evidence and Decisions by Agency Professionals: An Experimental Test in the Context of Merger Regulation', *Theory and Decision*, 73(3), 465–99. https://doi.org/10.1007/s11238-011-9258-3

Mankiw, N. G., and Miron, J. A. (1986), 'The Changing Behavior of the Term Structure of Interest Rates', *The Quarterly Journal of Economics*, 101, 211–28. https://doi.org/10.2307/1891113

Mankiw, N. G., and Reis, R. (2002), 'Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve', *The Quarterly Journal of Economics*, 117(4), 1295–328. https://doi.org/10.1162/003355302320935034

— (2003), 'Sticky Information: A Model of Monetary Non-neutrality and Structural Slumps', in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford (eds), *Knowledge, Information and Expectations in Modern Macroeconomics*, Princeton, NJ, Princeton University Press, 64–86. https://doi.org/10.1515/9780691223933

Mayo, D. G. (1996), *Error and the Growth of Experimental Knowledge*, Chicago, IL, University of Chicago Press. https://doi.org/10.7208/chicago/9780226511993.001.0001

Mayo, D. G., and Spanos, A. (2006), 'Severe Testing as a Basic Concept in a Neyman–Pearson Philosophy of Induction', *The British Journal for the Philosophy of Science*, 57(June), 323–57. https://doi.org/10.1093/bjps/axl003

Menzies, G. D., and Zizzo, D. J. (2009), 'Inferential Expectations', *The Berkeley Electronic Journal of Macroeconomics*, 9(1), 1–25.

Michl, T. R., and Rowthorn, R. (2024), 'Optimal Inflation Targeting with Anchoring', *Review of Keynesian Economics*, 12(2), 220–38. https://doi.org/10.4337/roke.2024.02.06

Nisbett, R., and Ross, L. (1980), *Human Inference: Strategies and Shortcomings of Social Judgment*, London, Prentice–Hall.

Obstfeld, M. (1996), 'Models of Currency Crises with Self-fulfilling Features', *European Economic Review*, 40(3–5), 1037–47. https://doi.org/10.1016/0014-2921(95)00111-5

Orphanides, A., and Williams, J. C. (2007), 'Robust Monetary Policy with Imperfect Knowledge', *Journal of Monetary Economics*, 54(5), 1406–35. https://doi.org/10.1016/j.jmoneco.2007.06.005

Rowthorn, R. E. (1977), 'Conflict, Inflation and Money', *Cambridge Journal of Economics*, 1, 215–39.

Sachs, J., Tornell, A., and Velasco, A. (1996), 'The Mexican Peso Crisis: Sudden Death or Death Foretold?', *Journal of International Economics*, 41(3–4), 265–83. https://doi.org/10.1016/S0022-1996(96)01437-7

Sims, C. (2003), 'Implications of Rational Inattention', *Journal of Monetary Economics*, 50(April), 665–90. https://doi.org/10.1016/S0304-3932(03)00029-1

Sims, C. A. (2011), 'Rational Inattention and Monetary Economics', *Handbook of Monetary Economics*, Vol. 3A, Amsterdam: Elsevier.

Skott, P. (2023), *Structuralist and Behavioral Macroeconomics*, Cambridge, Cambridge University Press. https://doi.org/10.1017/9781009367349

Thomas-Walters, L., Goldberg, M. H., Lee, S., Lyde, A., Rosenthal, S., and Leiserowitz, A. (2024), 'Communicating the Links between Climate Change and Heat Waves with the Climate Shift Index', *Weather, Climate, and Society*, 16(3), 511–20.

Vines, D., and Wills, S. (2018), 'The Rebuilding Macroeconomic Theory Project: An Analytical Assessment', *Oxford Review of Economic Policy*, 34(1–2), 1–42. https://doi.org/10.1093/oxrep/grx062

Woodford, M. (2003), *Interest & Prices: Foundations of a Theory of Monetary Policy*, Princeton, NJ, Princeton University Press.

— (2005), 'Central Bank Communication and Policy Effectiveness', NBER Working Paper 11898, December.