Discussion Papers in Economics

THE CAUSAL EFFECTS OF HETEROGENEOUS EXPECTATION FORMATION IN GENERAL EQUILIBRIUM

By

Alistair Macaulay
(University of Surrey)

DP 03/26

School of Economics
University of Surrey
Guildford
Surrey GU2 7XH, UK
Telephone +44 (0)1483 689380
Facsimile +44 (0)1483 689548
Web https://www.surrey.ac.uk/school-economics
ISSN: 1749-5075

The Causal Effects of Heterogeneous Expectation Formation in General Equilibrium

Alistair Macaulay*
University of Surrey

February 13, 2026

Abstract

This paper studies the causal effect of heterogeneously-formed expectations on aggregate economic outcomes. To do this, I develop a 'temporary-information equilibrium' approach that characterizes the causal impact of expectation changes in a broad class of macroeconomic models using simple cross-sectional statistics on how agents form expectations. This approach uncovers a novel 'narrative heterogeneity channel' through which heterogeneities in expectation formation can affect the transmission of both expectational and fundamental shocks. This channel arises in a wide range of commonly-used models of expectation formation. In two applications to data on firm and household expectations, I find that narrative heterogeneity accounts for a large but time-varying share of the causal effects of expectations on macroeconomic outcomes.

JEL codes: D83, D84, E31, E71

*School of Economics, University of Surrey. Email: a.macaulay@surrey.ac.uk. I thank Martin Ellison, Ángelo Gutiérrez-Daza (discussant), Cosmin Ilut, Sebastian Link, Michael McMahon, Pascal Meichtry (discussant), Oliver Pfäuti, Carlo Pizzinelli, Wenting Song, Mirko Wiederholt, and participants at numerous seminars and conferences for valuable feedback. I am particularly grateful to Timo Reinelt for providing me with extra results from the experiment in Baumann et al. (2024). I thank the John Fell OUP Research Fund for financial support. Alexa Kaminski and Chenchuan Shi provided excellent research assistance. This paper subsumes material circulated previously under the titles "Shock Transmission and the Sources of Heterogeneous Expectations" and "Heterogeneous information, subjective model beliefs, and the time-varying transmission of shocks".

1 Introduction

Expectations are fundamental to our understanding of macroeconomics. At least since Lucas (1972) they have been at the core of macroeconomic modeling, and more recently a vast empirical literature has demonstrated their importance for the decisions of firms, households, and investors.[1] Yet despite the seemingly obvious centrality of expectations, it is notoriously difficult to say precisely how a given change in expectations would affect aggregate outcomes.

The reason for this uncertainty is that in the equilibrium conditions that define the behavior of macroeconomic models, expectations are endogenous. Take, for example, the linearized New Keynesian Phillips Curve (NKPC), as presented in e.g. Galí (2008):

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t \tag{1}$$

This shows that inflation π_t and expected future inflation $E_t\pi_{t+1}$ are related to one another, but it *does not* reveal anything about the causal effect of inflation expectations. As emphasized recently by Werning (2022), the β coefficient combines the causal effect of $E_t\pi_{t+1}$ on π_t, the causal effect of π_t on $E_t\pi_{t+1}$, and the fact that both $E_t\pi_{t+1}$ and π_t react simultaneously to other shocks.

The best approach to date has been to study models in *temporary equilibrium* (Grandmont, 1977; Woodford, 2013; Werning, 2022). However, this approach necessarily abstracts from one of the central findings of recent empirical literature: that expectations are formed heterogeneously. Between otherwise similar decision-makers there are large differences in information, and in how people interpret any information they have.[2] How, if at all, does this heterogeneity affect the causal effects of aggregate expectation shifts, and macroeconomic dynamics?

In this paper, I develop a new method that can reveal the causal effect of expectations in general equilibrium while allowing for heterogeneous expectation formation. Applying the method to a general class of linear macroeconomic models, I uncover a novel '*narrative heterogeneity channel*', through which heterogeneity influences both the causal effects of expectations and the transmission of fundamental shocks. I show that many popular models of expectation formation implicitly generate such a channel whenever sufficient heterogeneity is allowed for. I then measure the strength of the narrative heterogeneity channel, and demonstrate the use of my approach to studying expectations in general equilibrium models, in two applications.

In the first application, I combine the Calvo model underpinning equation (1) with the information-treatment randomized controlled trial (RCT) results of Baumann et al. (2024) to

[1] See e.g. the extensive reviews in Bachmann et al. (2023), and the discussion of related literature below.

[2] See for example Link et al. (2023) for information, Andre et al. (2022) for subjective models used to interpret information, and Pfajfar and Santoro (2010), Beutel and Weber (2024), and Macaulay and Moberly (2025) for both.

measure the causal effects of firm inflation expectations on realized inflation. In doing so, I demonstrate an auxiliary benefit of the '*temporary-information equilibrium*' method I develop, that it can be used to map partial-equilibrium information-treatment RCT results into general equilibrium insights, offering a solution to the 'missing intercept' problem in these studies.

In the second application, I use unique features of a UK survey to provide a direct measurement of the narrative heterogeneity channel over time, in the context of household inflation expectations. I then combine this measurement with a New Keynesian model to measure the causal effects of household inflation expectations on inflation, output, and interest rates. In both applications, I find that the narrative heterogeneity channel is substantial. In the household application in particular, the narrative heterogeneity channel amplifies the average causal effect of household inflation expectations by $\approx 15\%$, and dampens the time-variation in that effect by a similar amount.

Causal effects of expectations. Temporary equilibrium analysis starts by solving the decision problems of agents in a model, but rather than imposing rational expectations, the expectations of all variables are held fixed. The researcher then solves for equilibrium as a function of those fixed and now exogenous expectations. Finally, they compute comparative statics of endogenous variables of interest with respect to each expectation. This reveals the causal effect of an aggregate shift in expectations of a particular variable, holding all other expectations constant.

This abstracts from two plausible features of expectation formation. First, as all expectations are shifted equally, there can be no heterogeneity in any aspect of expectations. Second, there is no general equilibrium feedback from endogenous outcomes back to expectations. When expectations shift and affect realized outcomes, agents observe those outcomes and update their expectations accordingly, potentially amplifying or dampening the initial effect.

To obtain the full general equilibrium effects of a given expectation, we therefore need to allow expectations to respond endogenously and heterogeneously to aggregate conditions. How can we reconcile this with the fact that this very endogeneity is what makes it hard to uncover causal effects in the first place?

I propose a solution, "*temporary-information equilibrium*," inspired by the information-treatment RCTs used to estimate causal effects of expectations on individual beliefs and actions in the empirical literature (see Haaland et al., 2023, for a review). In those experiments, researchers exogenously vary one part of the information available to participants, and track how that affects expectations of the shocked variable, expectations of other variables ("cross-learning"), and actions. Instrumenting the shocked expectation with the exogenously varied information, they identify the causal effect of that expectation, despite not controlling all aspects of expectation formation.

As with temporary equilibrium, I start from agent decision problems. However, rather than holding expectations fixed, I model them as composed of information (a noisy signal) and a residual, which depends on expectations of other variables. While expectations are therefore endogenous, the signal noise is exogenous, and shocking this noise creates exogenous variation in a given expectation. Since all effects of the information shift come through the expectation of interest, we can use this to study the causal effect of changes in that expectation, while still allowing cross-learning to other expectations and responses to general equilibrium effects. Moreover, we can allow for rich heterogeneity in expectation formation, and compute the causal effects of a shock causing aggregate expectations to move a given amount. This mirrors how information-treatment RCTs use exogenous variation in only part of participant information sets to assess the causal effects of interest.

This temporary-information equilibrium approach is the key methodological contribution of the paper. It nests standard temporary equilibrium as a special case, but also allows for both general-equilibrium feedback through endogenous signals and expectations heterogeneity.

The Calvo model. For clarity, I begin by demonstrating temporary-information equilibrium in the Calvo model of firm price-setting. This is one way to derive the NKPC (1), and its temporary equilibrium has been extensively analyzed by Werning (2022). The *temporary-information equilibrium* nests Werning's results, while also allowing for heterogeneity in expectation formation.

This turns out to be of first-order importance. The temporary-information equilibrium reveals a previously overlooked channel of shock transmission, which I term the '*narrative heterogeneity channel*,' and which qualitatively alters features of the passthrough from inflation expectations to realized inflation derived in Werning (2022). The novel channel arises whenever the precision of firm signals is correlated with the subjective models they use for cross-learning. Intuitively, if the firms who receive precise signals on current inflation are also those who extrapolate most strongly from current to future inflation, then a given change in inflation will have a greater effect on aggregate pricing than if all firms had the same average information precision and cross-learning.[3]

This narrative heterogeneity channel affects both the partial equilibrium effects of expectations, and the feedback into other expectations in general equilibrium. Moreover, I show that a wide range of common theories of expectation formation naturally imply narrative heterogeneity effects whenever heterogeneity is permitted in both information precision and the cross-learning used to interpret that information. This includes models of rational inattention (Maćkowiak et al., 2023), learning (Evans and McGough, 2020), level-k thinking (Farhi and Werning, 2019) and others.

Having developed these results and procedures in the simple Calvo model, I then go on to show

[3] This echoes the redistribution channel in the broader heterogeneous-agent literature (Auclert, 2019), where shocks are amplified when exposure to shocks correlate with responsiveness to those shocks.

that the same approach can be applied to a broad class of linear general equilibrium models. Many of the same insights apply, notably the potentially critical role of narrative heterogeneity effects in determining partial and general equilibrium effects of expectations.

Application 1: the effect of firm inflation expectations. Typically, researchers are unable to formally map the estimates from information-provision RCTs into macroeconomic insights, because of a "missing intercept" problem: RCTs compare treated and untreated individuals, and therefore cannot reveal the general equilibrium effects of aggregate expectation changes. Moreover, since those general equilibrium effects likely also depend on the behavior of expectations, there are no available equivalence results allowing the intercept to be obtained from alternative time-series evidence, as is the case for fiscal transfers (Wolf, 2023).

Temporary-information equilibrium offers a solution. The conceptual link between the approach and information-provision RCTs means that the estimates commonly computed in those experiments correspond closely to the key statistics in the temporary-information equilibrium of standard macroeconomic models. I demonstrate this using the RCT on European firm inflation expectations in Baumann et al. (2024), showing how the temporary-information equilibrium of the Calvo model can map their estimates into general equilibrium insights for inflation dynamics.

I find that the Baumann et al. (2024) estimates imply that general equilibrium effects of one-year-ahead inflation expectations are larger than would be implied by rational expectations, but that the differences are dampened by the narrative heterogeneity channel, and the magnitudes depend strongly on the degree of price stickiness. If the average price duration is 5 quarters, the general equilibrium effect of inflation expectations is just 6% larger than that implied by rational expectations. If, however, the average price duration is 2 years, the difference rises to 21%. Without the narrative heterogeneity channel, these numbers would be 7% and 24% respectively.

In this application the narrative heterogeneity channel is inferred by combining the results of regressions in the RCT that are informative about different statistics in the temporary-information equilibrium of the Calvo model. In the second application, I turn to a context in which the correlation of interest, between information and subjective models, can be observed directly.

Application 2: the effect of household inflation expectations. Unique features of a UK household survey directly reveal the joint distribution of household inflation information and the subjective models used to interpret it. I therefore measure the narrative heterogeneity channel here without needing to infer it from a specific model.

I document two key patterns in this data. First, households who believe inflation makes little difference to the strength of the economy use less information about inflation than households with other subjective models. There is therefore a systematic relationship between information and subjective models, implying that the narrative heterogeneity channel will operate.

Second, information that inflation is high is associated with more negative subjective models of the effects of inflation, both in the cross-section and over time. A greater proportion of households report that inflation makes the economy weaker in periods with high realized inflation, and within a period, those who believe inflation is currently higher are more likely to hold such negative subjective models. The joint distribution of information and subjective models therefore varies over time, and is systematically related to the state of the economy.

Combining these survey results with the temporary-information equilibrium of a New Keynesian model, I find that increased household inflation expectations cause output to fall. When inflation expectations rise, the majority of households revise their expectations for future output down, and thus reduce consumption. However, this causal effect is time-varying. When realized inflation rises, average cross-learning from inflation to output becomes more negative, and so the causal effects of expectations are stronger. The model implies that the causal effects of household inflation expectations was 30% larger in 2008 Q3, when annualized inflation was 4.8%, than in 2015 Q1 when it was 0.1%. The narrative heterogeneity channel amplifies the causal effect in steady state by 16%, and dampens its time-variation by a similar amount. The dampening of volatility occurs because, when inflation becomes high, almost all households agree that inflation reduces output, and so there is little heterogeneity in subjective models to drive a narrative heterogeneity channel. When inflation falls, heterogeneity increases, driving up the narrative heterogeneity effect.

Related literature. Principally, this paper contributes to the vast literature on the role of expectations in macroeconomic dynamics. While many papers examine how particular models of expectation formation affect the transmission of fundamental shocks (Maćkowiak and Wiederholt, 2015; Eusepi and Preston, 2018; Bianchi et al., 2024, among many others), the causal effects of changes in expectations themselves are much less clear. Recent progress on this issue has involved abstracting from cross-sectional heterogeneity in expectations and expectation formation (e.g. Werning, 2022; Ascari et al., 2023), despite their empirical prominence (e.g. Mankiw et al., 2004; Andrade et al., 2016; Link et al., 2025, among many others). I extend this literature by isolating the causal effects of *heterogeneously-formed* expectations on aggregate economic dynamics.

The approach uncovers a novel *narrative heterogeneity channel* through which heterogeneity in expectation formation affects aggregate dynamics, adding to the mechanisms identified in existing literature (e.g. Branch and Evans, 2006; Andrade et al., 2019, and many others). Conceptually, it relates to channels in the heterogeneous-household literature generated by correlations between shock exposure and the marginal propensity to consume (Auclert, 2019; Bilbiie, 2019).[4] I show that the relevant cross-sectional correlations for the narrative heterogeneity channel arise naturally in many common models of expectation formation, whenever heterogeneity is permitted in both information and subjective models. This complements empirical work documenting these correlations in survey data (Beutel and Weber, 2024; Macaulay and Moberly, 2025).

Methodologically, this paper is most closely related to Werning (2022), who studies the causal effect of inflation expectations in price-setting models in temporary equilibrium (Grandmont, 1977). My temporary-information equilibrium nests this, and extends it to settings with rich heterogeneity in expectation formation, as well as general equilibrium feedback from endogenous variables to expectations.[5] The method is also conceptually related to the large empirical literature on information-provision RCTs (reviewed in Fuster and Zafar, 2023; Haaland et al., 2023).

Translating the results of information-provision RCTs into insights for macroeconomic dynamics is complicated by the 'missing intercept' problem that affects all cross-sectional empirical work in macroeconomics (Wolf, 2023; Donaldson, 2025; Matthes et al., 2025).[6] Because the expectation statistics that discipline the temporary-information equilibrium are closely related to the objects estimated in these RCTs, the method provides a way to map RCT results into general-equilibrium insights, through the lens of standard macroeconomic models. I illustrate this way of addressing the missing intercept problem using the RCT on firm expectations in Baumann et al. (2024). The results speak to the growing literature on firm inflation expectations (see Candia et al., 2022, for a review).

Finally, my second application relates to the broader literature on household inflation expectations (see reviews in D'Acunto et al., 2023; Milani, 2023). My contribution is to assess the causal effect of household inflation expectations on aggregate output, uncovering substantial time variation, driven in large part by the narrative heterogeneity channel.

Outline. In Section 2, I introduce the temporary-information equilibrium approach, and the narrative heterogeneity channel, in the simple context of the Calvo model of firm price-setting. Section 3 generalizes to a broad class of linear macroeconomic models. In Section 4 I apply this to the RCT in Baumann et al. (2024), and in Sections 5-6 I apply it to household inflation expectations, using UK survey data and a New Keynesian model. Section 7 concludes.

[4] Nord (2022) and Guerreiro (2023) study a further related channel driven by cross-sectional correlations between the sensitivity of expectations to a shock, and the consumption responses to expectations.

[5] Allowing for the reaction of expectations to endogenous variables also connects this method to the literature on learning from prices (Lucas, 1972; Grossman and Stiglitz, 1980; Gaballo, 2018; Chahrour and Gaballo, 2021).

[6] A small number of papers use such RCTs to discipline macroeconomic models (e.g. Hoffmann et al., 2022; Hajdini et al., 2023; Andre et al., 2025). Related work disciplines models with non-experimental moments of expectations data (e.g. Leombroni et al., 2020; Ludwig et al., 2024). While this avoids the missing intercept problem, the data used cannot reveal how expectations change in response to shocks or policy changes (Piazzesi and Schneider, 2016).

2 Calvo model

2.1 Setup

We start, as in Werning (2022), with the log-linearized first order condition of a price-setting firm:

$$p_{it}^* - P_{t-1} = (1-\beta\lambda)\sum_{s=0}^{\infty}(\beta\lambda)^s E_{it}\left(\pi_{t+s} + mc_{t+s}\right) + \mu \tag{2}$$

where p_{it}^* is the desired reset price of firm i in period t; P_{t-1} is the aggregate price level in period $t-1$; $\pi_{t+s} \equiv P_{t+s} - P_{t-1}$ is cumulative inflation between periods $t-1$ and $t+s$; and mc_{t+s} is real marginal costs in period $t+s$. The parameters β and λ are the discount factor and the probability a firm is unable to reset their price in a given period respectively, and both are $\in [0,1]$. μ is the desired steady state markup, and the operator E_{it} denotes the expectation of firm i in period t.

In New Keynesian models, the standard approach is to impose rational expectations, and then derive the familiar New Keynesian Phillips Curve, in which the coefficient on $\bar{E}_t\pi_{t+1}$ is β (e.g. Galí, 2008). However, as discussed extensively in Werning (2022), this coefficient does not reveal the causal effects of inflation expectations, because the use of rational expectations means that expected inflation is endogenous. High inflation expectations simultaneously cause and are caused by high current inflation, making it impossible to disentangle the strength of each individual mechanism.

Temporary equilibrium. Instead of imposing rational expectations, Grandmont (1977) and others recommend holding expectations fixed, then aggregating and solving for equilibrium in terms of expectations. In the Calvo model this is straightforward:[7]

$$\pi_t = (1-\lambda)(p_{it}^* - P_{t-1}) = (1-\lambda)(1-\beta\lambda)\left(\sum_{s=0}^{\infty}(\beta\lambda)^s \bar{E}_t\left(\pi_{t+s} + mc_{t+s}\right) + \mu\right) \tag{3}$$

where $\bar{E}_t(\cdot) = \int_i E_{it}(\cdot)di$ denotes aggregate expectation across firms.

With this, we can compute comparative statics with respect to any aggregate expectation. Since we have not imposed rational expectations, expectations are exogenous in these equations, and the comparative statics can be interpreted as the causal effects of changes in expectations. Specifically,

$$\frac{d\pi_t}{d\bar{E}_t(\pi_t)} = (1-\lambda)(1-\beta\lambda) \tag{4}$$

[7] Note we do not need to include a general equilibrium process determining mc_{t+s}, because it is *expected* mc_{t+s} that enter the firm first order condition, and these (like all other expectations) are held constant in temporary equilibrium.

gives the causal effect of a change in aggregate expectations of current inflation, *holding all other expectations constant*. This is a direct consequence of removing all endogeneity from expectations: expectations of one variable cannot therefore react to expectations of another.

In many settings this is more extreme than is desirable. This Calvo model is a good example, because π_{t+s} is defined as the cumulative inflation between periods $t-1$ and $t+s$. Increasing $\bar{E}_t(\pi_t)$ but leaving all longer-horizon $\bar{E}_t(\pi_{t+s})$ unchanged, as in equation (4), therefore entails a decrease in expected inflation between periods t and $t+1$ to exactly offset the expected increase in π_t. The majority of evidence on inflation expectations of firms (as well as households) suggests that there is no such reversion of price-level expectations after shocks to short-run expected inflation (Weber et al., 2022), so this may not be the most informative way to study inflation expectations.

Extended temporary equilibrium. The simple way to avoid this problem is to impose some structure to link expectations that we think should co-move after solving for equilibrium (3), but before computing comparative statics. This is the approach taken in Werning (2022), who assumes

$$\bar{E}_t(\pi_{t+s}) = \pi^e(1+s) \tag{5}$$

which corresponds to the assumption that per-period expected inflation is constant at π^e. While each expectation in equation (3) is no longer fully exogenous (as inflation expectations at each horizon are endogenous to each other), the common factor π^e is still exogenous. The comparative static

$$\frac{d\pi_t}{d\pi^e} = \frac{1-\lambda}{1-\beta\lambda} \tag{6}$$

therefore gives the causal effect of a change in aggregate expectations of per-period inflation, assuming that per-period inflation is expected to be equal at all horizons. Werning (2022) gives an extensive analysis of this result. In principle, the researcher can also specify relationships between expected inflation and expected marginal costs, and recompute the causal effects.

This is a useful advance on studying the equilibrium New Keynesian Phillips Curve, and on the standard temporary equilibrium approach. However, it cannot address how heterogeneity in expectation formation affects shock transmission, how general equilibrium feedback alters the results, or what role expectations play in the transmission of fundamental shocks. To answer these questions, I propose an alternative approach, inspired by the information-treatment experiments in the empirical study of expectations. This "temporary-information equilibrium" nests the extended temporary equilibrium approach presented here, but allows for general equilibrium and heterogeneous expectation formation, while preserving the causal interpretation of the results.

Temporary-information equilibrium. When assessing the causal impact of a particular expectation, information-treatment RCTs do not vary that expectation directly. Rather, they exogenously vary the *information* participants have on that variable, and that information is passed through into the expectation of interest. There are also commonly effects on expectations of other variables, which are known in that literature as 'cross-learning' (Haaland et al., 2023). Taking this empirical procedure as our inspiration, we assume that the expectation of any variable z_j is composed of two parts: a signal $s_{it}(z_j)$ and a residual $R_{it}(z_j)$.

$$E_{it}(z_j) = (1 - \tau_{ij})R_{it}(z_j) + \tau_{ij}s_{it}(z_j) \tag{7}$$

This follows the form of a standard signal extraction problem, including that used by Maćkowiak and Wiederholt (2024) to model information acquisition during information-treatment RCTs. Under that signal extraction interpretation, $R_{it}(z_j)$ reflects the firm's prior belief, and τ_{ij} reflects the signal to noise ratio.

Rather than exogenously varying the expectation as in temporary equilibrium approaches, we will analyze the effects of exogenous variation in the signal, allowing the residual component to (potentially) respond endogenously. Specifically, the residual is the expectation firm i would set if they received no information on that variable. The only way to form such an expectation is by using information and beliefs about *other* variables. Since the model here is linear, we impose

$$R_{it}(z_j) = \sum_{z_k \in \boldsymbol{z}: z_k \neq z_j} M_{ijk}E_{it}(z_k) \tag{8}$$

where $\boldsymbol{z}$ is the set of all expectations of interest in the model. Stacking equation (7) over all expectations in $\boldsymbol{z}$, substituting in equation (8), and rearranging gives:

$$E_{it}(\boldsymbol{z}) = \boldsymbol{\chi}_i\boldsymbol{\tau}_i\boldsymbol{s}_{it} \tag{9}$$

where $\boldsymbol{\tau}_i$ is a square diagonal matrix, with jth diagonal element equal to τ_{ij}; $\boldsymbol{s}_{it}$ is a column vector stacking $s_{it}(z_j)$ over z_j, and the square matrix $\boldsymbol{\chi}_i$ is defined as

$$\boldsymbol{\chi}_i \equiv (I - (I - \boldsymbol{\tau}_i)\boldsymbol{M}_i)^{-1} \tag{10}$$

where $\boldsymbol{M}_i$ has zeros on the diagonal, and (j, k)th element equal to M_{ijk} in equation (8).

For this section, we will restrict attention to cases in which $\boldsymbol{M}_i$ is upper-triangular. Conceptually, this means we only consider cases where the subjective model of each agent can be expressed as a recursive system of equations, as discussed in detail in the literature on Directed Acyclic Graphs (Spiegler, 2020).[8] We relax this assumption in Section 3.

With this assumption, $\boldsymbol{\chi}_i$ is upper-triangular, with diagonal elements equal to 1. Each element of $\boldsymbol{\tau}_i$ therefore represents the degree to which a firm updates their expectations of a particular variable in response to direct information about that variable. In many applications this will be proportional to the precision of the agent's information. Off-diagonal elements of $\boldsymbol{\chi}_i$ capture cross-learning, or equivalently the links across expectations imposed in extended temporary equilibrium analysis.

The cross-learning matrix $\boldsymbol{\chi}_i$ captures both direct and indirect links between expectations of different variables: after a signal on short-run inflation, a firm may update expectations of long-run inflation partly because they believe inflation is persistent. They may also update their expectations of short-run marginal costs, which in turn affects their expectations of long-run marginal costs, which in turn affects their expectations of long-run inflation. It will be convenient from here to work directly with $\boldsymbol{\chi}_i$, rather than considering each of these possible links individually.[9]

Substituting equation (9) into equation (3) yields

$$\pi_t = (1-\lambda)(1-\beta\lambda)\left(\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\boldsymbol{\chi}_{i\pi_{t+s}:}\boldsymbol{\tau}_i \boldsymbol{s}_{it} + \boldsymbol{\chi}_{imc_{t+s}:}\boldsymbol{\tau}_i \boldsymbol{s}_{it}\right) di + \mu\right) \tag{11}$$

where $\boldsymbol{\chi}_{i\pi_{t+s}:}$ and $\boldsymbol{\chi}_{imc_{t+s}:}$ are rows of $\boldsymbol{\chi}_i$ corresponding to $E_{it}(\pi_{t+s})$ and $E_{it}(mc_{t+s})$ respectively.

2.2 The causal effect of inflation expectations.

I now use the temporary-information equilibrium (11) to study the causal effect of changes in expected π_t and in expected π_{t+1}. The former is the effect studied in Werning (2022), while the latter more closely corresponds to the questions in the empirical information-treatment literature.

Partial equilibrium. To begin with, consider an exogenous shock to the signals firms receive about current inflation $s_{it}(\pi_t)$, holding constant the signals firms receive about all other variables. All proofs for this section and Section 3 are in Appendix A.

[8] As an example, this means that if the firm uses $E_{it}\pi_t$ to update $R_{it}mc_t$, they cannot simultaneously use $E_{it}mc_t$ to update $R_{it}\pi_t$. $E_{it}mc_t$ may still affect $R_{it}\pi_{t+s}$ for $s > 0$, as these are treated as separate elements of $\boldsymbol{z}$ from π_t.

[9] This has a parallel in the literature on production networks (Carvalho and Tahbaz-Salehi, 2019). The direct links in $\boldsymbol{M}_i$ are analogous to the elements of the input-output matrix, and $\boldsymbol{\chi}_i$ is the corresponding Leontief inverse.

Proposition 1 *The causal effect of a shock ε_{π_t} that increases $s_{it}(\pi_t)$ equally for all firms is*

$$\frac{1}{\bar{\tau}_{\pi_t}}\frac{d\pi_t}{d\varepsilon_{\pi_t}} = (1-\lambda)(1-\beta\lambda)\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\chi_{i\pi_{t+s}\pi_t}\frac{\tau_{i\pi_t}}{\bar{\tau}_{\pi_t}} + \chi_{imc_{t+s}\pi_t}\frac{\tau_{i\pi_t}}{\bar{\tau}_{\pi_t}}\right) di \tag{12}$$

where $\chi_{i\pi_{t+s}\pi_t}$ and $\chi_{imc_{t+s}\pi_t}$ refer to the elements of $\boldsymbol{\chi}_i$ that correspond to the cross-learning from a signal about π_t to expectations of π_{t+s} and mc_{t+s} respectively. Scaling the derivative by the average response to π_t signals, $\bar{\tau}_{\pi_t} \equiv \int_i \tau_{i\pi_t} di$, implies that this effect comes through a unit increase in $\bar{E}_t\pi_t$.

As in temporary equilibrium analysis, this can be interpreted as the causal effect of inflation expectations. Firm signals about π_t have been varied exogenously, and the only way those signals affect the firm decision problem is through $E_{it}(\pi_t)$. This is therefore the theoretical equivalent of information-treatment RCTs using exogenous variation in information as an instrument for expectations, and thus determining the effects of expectations on other (individual-level) outcomes. In-keeping with this IV-like interpretation, we scale the effect such that the result reflects the causal effect of a unit increase in the relevant aggregate expectation, which in this case is $\bar{E}_t(\pi_t)$.

Given this interpretation, it is unsurprising that this straightforwardly nests the two standard approaches discussed above. Setting $\tau_{i\pi_t} = 1$ for all firms i recovers the case in which one expectation is shocked and all others are held constant, as in standard temporary equilibrium analysis (equation (4)). If instead $\tau_{i\pi_t} < 1$, and $\chi_{imc_{t+s}\pi_t} = 0$, $\chi_{i\pi_{t+s}\pi_t} = 1 + s$ for all i, then we return to the extended temporary equilibrium in equation (6), in which a rise in expectations of π_t increases expected per-period inflation one-for-one at all horizons, but has no effect on marginal cost expectations. By selecting different values for $\tau_{i\pi_t}$ and the elements of $\boldsymbol{\chi}_i$, we can also solve for the partial-equilibrium effect of current inflation expectations under a range of common assumptions on expectation formation, which have different implications for responses to information and the linkages between expectations across horizons and variables. We return to this point in Section 2.4.

Similarly, the causal effect of π_{t+1} expectations is given by Proposition 2.

Proposition 2 *The causal effect of a shock $\varepsilon_{\pi_{t+1}}$ that increases $s_{it}(\pi_{t+1})$ equally for all firms is*

$$\frac{1}{\bar{\tau}_{\pi_{t+1}}}\frac{d\pi_t}{d\varepsilon_{\pi_{t+1}}} = (1-\lambda)(1-\beta\lambda)\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\chi_{i\pi_{t+s}\pi_{t+1}}\frac{\tau_{i\pi_{t+1}}}{\bar{\tau}_{\pi_{t+1}}} + \chi_{imc_{t+s}\pi_{t+1}}\frac{\tau_{i\pi_{t+1}}}{\bar{\tau}_{\pi_{t+1}}}\right) di \tag{13}$$

where the scaling here implies that the effect comes through a unit increase in $\bar{E}_t\pi_{t+1}$.

The narrative heterogeneity channel. A key novel insight to notice from equations (12) and (13) concerns the impact of heterogeneity in expectation formation. There is substantial evidence of heterogeneity across households and firms in the precision of information (e.g. Link et al., 2023), and more directly in the responsiveness to signals (Balla-Elliott, 2025). There is also substantial heterogeneity in how agents update from one expectation to another (e.g. Andre et al., 2022; Macaulay and Moberly, 2025). Holding signals, rather than expectations, fixed allows us to study how this heterogeneity affects the causal effect of expectations. Specifically, using the definition of a covariance, we can rewrite equation (12) as

$$\frac{1}{\bar{\tau}_{\pi_t}}\frac{d\pi_t}{d\varepsilon_{\pi_t}} = (1-\lambda)(1-\beta\lambda)\sum_{s=0}^{\infty}(\beta\lambda)^s\left(\bar{\chi}_{\pi_{t+s}\pi_t} + \bar{\chi}_{mc_{t+s}\pi_t}\right)$$
$$+ (1-\lambda)(1-\beta\lambda)\sum_{s=0}^{\infty}(\beta\lambda)^s\left(\frac{1}{\bar{\tau}_{\pi_t}}Cov_I[\chi_{i\pi_{t+s}\pi_t}, \tau_{i\pi_t}] + \frac{1}{\bar{\tau}_{\pi_t}}Cov_I[\chi_{imc_{t+s}\pi_t}, \tau_{i\pi_t}]\right) \quad (14)$$

where $\bar{\chi}_{z\pi_t}$ is the average of $\chi_{iz\pi_t}$ across firms for each variable z, and $Cov_I[\cdot,\cdot]$ denotes a cross-sectional covariance. Equation (13) admits an equivalent decomposition for the effects of $\varepsilon_{\pi_{t+1}}$.

The first summation captures representative-agent effects: the causal effect of expected current inflation depends on the average cross-learning to expectations of marginal costs and longer-run inflation across firms. However, the second summation captures the effects of heterogeneity. If the firms with the most precise information on π_t (i.e. the highest $\tau_{i\pi_t}$) are also the firms who extrapolate most strongly from π_t to other variables (high $\chi_{i\pi_{t+s}\pi_t}$ and $\chi_{imc_{t+s}\pi_t}$), information is concentrated among those who will extrapolate and react most strongly to it. That amplifies the effect of an aggregate signal relative to what would be observed if all firms formed expectations in the same way.

We will refer to this broad category of effects as the 'narrative heterogeneity channel,' reflecting the idea that a narrative combines information with a subjective model to interpret it (Gibbons and Prusak, 2020). These effects were not visible in existing temporary equilibrium analysis, because such exercises implicitly remove all heterogeneity in expectations. Once we allow for heterogeneity, then mathematically the narrative heterogeneity terms are a trivial consequence of Jensen's inequality. In Section 2.4 we will argue that *economically* they are not trivial at all. A wide range of popular models of expectation formation imply potentially large and time-varying narrative heterogeneity effects, with important qualitative and quantitative consequences.

General equilibrium. All of the exercises so far study the effects of a given exogenous change in the signals firms receive about current or one-period-ahead inflation. This captures the causal effect of a given expectation rise, holding all other signals constant. However, holding signals constant means abstracting from second-round effects: when expectations of π_{t+1} rise, that has effects on current realized π_t (equation (13)), *which in turn* should affect the information firms receive about π_t, which has its own effects. To assess the full passthrough from inflation expectations to realized inflation in general equilibrium, this feedback should be taken into account.

The temporary-information approach can easily accommodate this. As is standard in models with noisy information (e.g. Coibion and Gorodnichenko, 2015), suppose that firm signals are equal to the variable of interest plus some additively separable noise, which here I allow to consist of a public component and an idiosyncratic component

$$s_{it}(\pi_{t+s}) = \pi_{t+s} + \varepsilon_{\pi t+s} + \tilde{\varepsilon}_{i\pi t+s} \tag{15}$$

where $\varepsilon_{\pi t+s}$ and $\tilde{\varepsilon}_{i\pi t+s}$ are both mean zero, and $\tilde{\varepsilon}_{i\pi t+s}$ is additionally i.i.d. across firms.

This functional form assumption preserves the causal interpretation of the exercises above: exogenous variation in $\varepsilon_{\pi t+s}$ still only affects firm decisions through the relevant inflation expectation, so can be used to assess the causal effects of that expectation on aggregate outcomes. However, that effect can now also include an endogenous response of signals to the shock, through the dependence of the signal on realized π_{t+s}.

For the causal effects of $\bar{E}_t(\pi_t)$, this general equilibrium effect turns out to be irrelevant. The partial equilibrium effect in equation (12) increases π_t, which now causes a second round of changes in aggregate π_t expectations, with associated consequences for realized inflation. However, all of those general equilibrium effects operate through $\bar{E}_t(\pi_t)$. As a result, when we scale the causal effect to ensure a unit impact on $\bar{E}_t(\pi_t)$, we get back exactly to the partial equilibrium effect. I show this formally in Appendix A.

However, for the causal effects of $\bar{E}_t(\pi_{t+1})$, the general equilibrium feedback is not irrelevant. From here, it will be useful to introduce the following notation: let $\Omega^{PE}(\pi_{t+s})$ denote the partial-equilibrium effect of a change in $s_{it}(\pi_{t+s})$, scaled to ensure a unit impact on $\bar{E}_t(\pi_{t+s})$ as in Propositions 1 and 2. Let $\Omega^{GE}(\pi_{t+s})$ denote the equivalent when we allow for general equilibrium feedback through the endogenous part of signals.

Proposition 3 *The causal effect of $\bar{E}_t(\pi_{t+1})$ on realized π_t, after accounting for general equilibrium effects through current inflation signals, is given by*

$$\Omega^{GE}(\pi_{t+1}) = K \cdot \frac{d\pi_t}{d\varepsilon_{\pi_{t+1}}} = \frac{\Omega^{PE}(\pi_{t+1})}{1 - \bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t) + \Omega^{PE}(\pi_{t+1}) \cdot (\bar{\chi}_{\pi_{t+1}\pi_t}\bar{\tau}_{\pi_t} + Cov_I[\chi_{i\pi_{t+1}\pi_t}, \tau_{i\pi_t}])} \tag{16}$$

where K is a scaling coefficient defined in Appendix A which ensures the impact on $\bar{E}_t(\pi_{t+1})$ in *general equilibrium is equal to 1.*

This reveals two general-equilibrium forces at work after a shock to $\bar{E}_t(\pi_{t+1})$. First, the partial-equilibrium effect of the shock on realized π_t feeds back into signals $s_{it}(\pi_t)$. This change in signals in turn affects π_t, with strength governed by $\Omega^P E(\pi_t)$. Second, the resulting increase in $E_{it}(\pi_t)$ causes firms to further update $E_{it}(\pi_{t+1})$ through cross-learning. Since we scale to ensure a unit increase in $\bar{E}_t(\pi_{t+1})$, this dampens the aggregate effect. Overall, general equilibrium effects amplify the partial-equilibrium effect of the shock if the first force dominates:

$$\Omega^{PE}(\pi_t) > \Omega^{PE}(\pi_{t+1}) \cdot \left(\bar{\chi}_{\pi_{t+1}\pi_t} + \frac{1}{\bar{\tau}_{\pi_t}} Cov_I[\chi_{i\pi_{t+1}\pi_t}, \tau_{i\pi_t}] \right) \tag{17}$$

Note that even if the desired model experiment is one in which all firms receive an identical signal about π_{t+1}, and so there is no narrative heterogeneity term in the direct partial-equilibrium effects of the shock, heterogeneity will still affect dynamics through these general equilibrium feedback channels. If realized π_t also affects real marginal costs, which have not been modeled here, that would further affect these dynamics. The temporary-information equilibrium can be used in the same way.

2.3 The role of expectations in the transmission of fundamental shocks

So far, we have been considering the causal effects of a shock or change in expectations. However, the temporary-information equilibrium apparatus is also useful in studying the related question of how expectations affect the transmission of *fundamental* shocks.

To see why, go back to equation (11), and allow for mc_t to be perfectly observed, with no cross-learning from it to any other expectations. In this setting we have

$$\pi_t = (1-\lambda)(1-\beta\lambda) \left[\sum_{s=0}^{\infty} (\beta\lambda)^s (\bar{\chi}_{\pi_{t+s},\pi_t} \bar{\tau}_{\pi_t} + Cov_I[\chi_{i,\pi_{t+s},\pi_t}, \tau_{i,\pi_t}]) \right] (\pi_t + \varepsilon_{s\pi t}) \\ + (1-\lambda)(1-\beta\lambda) mc_t + a \tag{18}$$

where a collects the steady-state markup μ and all terms involving signals of π_{t+s} and mc_{t+s} for $s \geq 1$. Consider an i.i.d. shock to mc_t. Since the Calvo model is purely forward-looking, all variables return to steady state in period $t+1$, and we can treat a as fixed. The aggregate effects of an i.i.d. mc_t shock are therefore given by

$$\frac{d\pi_t}{dmc_t} = (1-\lambda)(1-\beta\lambda)\left(\left[\sum_{s=0}^{\infty}(\beta\lambda)^s(\bar{\chi}_{\pi_{t+s},\pi_t}\bar{\tau}_{\pi_t} + Cov_I[\chi_{i,\pi_{t+s},\pi_t},\tau_{i,\pi_t}])\right]\frac{d\pi_t}{dmc_t}+1\right) \quad (19)$$

which rearranges to

$$\frac{d\pi_t}{dmc_t} = \frac{(1-\lambda)(1-\beta\lambda)}{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)} \quad (20)$$

This helps us to see the role of endogenous expectations in the transmission of the fundamental mc_t shock, as all expectational effects are contained in the $\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)$ term in the denominator. Formally, we can compute a counterfactual shock transmission in which there is no response of expectations by setting $\Omega^{PE}(\pi_t) = 0$. We can then compare this counterfactual to equation (20) to compute the share of marginal cost shock transmission which can be attributed to expectations.

Proposition 4 *The fraction of the transmission of an i.i.d. marginal cost shock to realized inflation that is due to inflation expectations is*

$$\frac{\frac{d\pi_t}{dmc_t} - \left.\frac{d\pi_t}{dmc_t}\right|_{\sim E}}{\frac{d\pi_t}{dmc_t}} = \bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t) \quad (21)$$

where $\left.\frac{d\pi_t}{dmc_t}\right|_{\sim E}$ *denotes the transmission in a counterfactual with* $\Omega^{PE}(\pi_t) = 0$.

That is, in this environment the share of the transmission of a fundamental marginal cost shock that is due to inflation expectations is equal to the partial equilibrium effect of an exogenous shock to those inflation expectations, multiplied by the average responsiveness of firms to inflation signals. The same channels determining the causal effects of an information shock also determine the share of other shocks' transmission due to expectations.

This is a particularly simple fundamental shock to analyze. However, the approach continues to work for more complicated interactions with multiple variables, or for alternative counterfactuals.[10] The key point is that studying the shock in the temporary-information equilibrium, rather than in a rational-expectations equilibrium or after imposing any other expectation-formation assumptions, allows us to separate the consequences of expectational shifts from their causes, without abstracting from heterogeneity in expectation formation.

[10]For example, if firms are assumed to perfectly observe current inflation, we could allow beliefs about π_t to adjust in the counterfactual, but shut down any transmission to expectations of future π_{t+s}. Rather than simply setting $\Omega^{PE}(\pi_t)$ to 0, the χ_{i,π_{t+s},π_t} and τ_{i,π_t} terms that make up $\Omega^{PE}(\pi_t)$ would just need to be modified accordingly.

2.4 The Narrative Heterogeneity Channel

The analysis above identifies a covariance between firms' direct response to signals, and their cross-learning, as a potential factor in the effects of expectations in general equilibrium. We now consider this 'narrative heterogeneity channel' in more detail.

An example. We begin with a simple example to give intuition for how the narrative heterogeneity channel operates, and how it can change quantitative and qualitative conclusions about the effects of expectations. Suppose that there is no cross-learning from inflation to marginal costs ($\chi_{imc_{t+s}\pi_t} = 0$), and $\chi_{i\pi_{t+s}\pi_t}$ and $\tau_{i\pi_t}$ are each split equally between two values

$$\chi_{i\pi_{t+s}\pi_t} = \{1, 1+2s\}, \quad \tau_{i\pi_t} = \{0, 1\} \tag{22}$$

These have been chosen such that $\bar{\chi}_{\pi_{t+s}\pi_t} = 1 + s$. If $\chi_{i\pi_{t+s}\pi_t}$ and $\tau_{i\pi_t}$ are independent, the narrative heterogeneity terms equal 0, and the causal effect of current inflation expectations in partial equilibrium is as in equation (6) and in Werning (2022) Proposition 1. However, if they are correlated, this causal effect can change dramatically. Proposition 5 considers the limit as $\beta \to 1$, which is a useful benchmark to gain intuition (and which is also used in Werning, 2022).

Proposition 5 *If $\tau_{i\pi_t}$ and $\chi_{i\pi_{t+s}\pi_t}$ are split equally between the values in equation (22), the causal effect of expected inflation on realized inflation as $\beta \to 1$ is given by*

$$\frac{1}{\bar{\tau}_{\pi_t}} \frac{d\pi_t}{d\varepsilon_{\pi_t}} = 1 + \lambda(2q - 1) \tag{23}$$

where $q \equiv \Pr\left(\chi_{i\pi_{t+s}\pi_t} = 1 + 2s \middle| \tau_{i\pi_t} = 1\right)$.

Werning (2022) Proposition 1 states that the transmission from expectations to inflation becomes 1 under no-discounting, irrespective of price stickiness. Equation (23) shows this is no longer true when expectations are formed heterogeneously. Unless information precision is exactly independent of subjective models (i.e. $q = \frac{1}{2}$), transmission depends on λ, even though the representative-agent terms do not.

The intuition is that $\chi_{i\pi_{t+s}\pi_t} = 1 + s$ is a special case that balances two opposing forces. As prices get stickier, each price-setting firm expects to wait longer before they can change prices again, which increases their responsiveness to expectations of inflation in future periods. At the same time, stickier prices mean fewer firms are resetting prices, decreasing the effect of expectations. When $\chi_{i\pi_{t+s}\pi_t} = 1 + s$ for all firms, these forces exactly cancel out, leaving the causal effect of expectations independent of price stickiness.

Heterogeneity breaks this balance, by affecting the first of these forces. When price stickiness increases, firms with high $\chi_{i\pi_{t+s}\pi_t}$ increase their responsiveness to expected inflation a great deal, because the extra length of a typical fixed-price spell encompasses periods of greater and greater inflation. By the same logic in reverse, firms with a low $\chi_{i\pi_{t+s}\pi_t}$ change their behavior less after an increase in stickiness. If information is concentrated among the former group, they dominate the passthrough from aggregate expectations to inflation, so that passthrough increases in stickiness.[11] If information is instead concentrated among the latter group (with low $\chi_{i\pi_{t+s}\pi_t}$), that group dominates, and the passthrough decreases in stickiness.

Common models of expectations generate narrative heterogeneity. Previous literature has typically relaxed either full information or rational expectations, but not both simultaneously (see related literature discussion above). That allows for heterogeneity in either $\boldsymbol{\chi}_i$ or $\boldsymbol{\tau}_i$, and not both, which forces the narrative heterogeneity terms to be 0. However, if heterogeneity is allowed in both components of expectation formation, many standard models of expectations naturally generate narrative heterogeneity effects.

In Appendix B, I consider the simple case in which the only possible signals are on π_t, where there is no public noise in those signals, and where both π_t and $\tilde{\varepsilon}_{i\pi_t}$ are normally distributed.[12] I then introduce a range of common assumptions on expectation formation, and show that they each imply systematic correlations between direct responses to signals ($\tau_{i\pi_t}$) and cross-learning ($\chi_{i\pi_{t+s}\pi_t}, \chi_{imc_{t+s}\pi_t}$), whenever heterogeneity is allowed in both. Since these models of expectation formation are well-known, I leave the details and derivations to Appendix B, and focus here on the intuition for why the narrative heterogeneity channel arises under each set of assumptions.

Rational Inattention (Maćkowiak et al., 2023): if firms can pay a cost to increase the precision of their signals, then $\tau_{i\pi_t}$ is endogenous to the firm's expected use of information. Firms who believe there is a stronger pass-through from current inflation to future inflation and marginal costs find information on π_t to be more valuable, and so process more of it. If there is heterogeneity in cross-learning, high $\tau_{i\pi_t}$ is therefore correlated with large magnitudes of $\chi_{i\pi_{t+s}\pi_t}$ and $\chi_{imc_{t+s}\pi_t}$.

For further intuition, suppose all firms have $\chi_{i\pi_{t+1}\pi_t} \geq 0$. Rational inattention implies $Cov_I[\chi_{i\pi_{t+1}\pi_t}, \tau_{i\pi_t}] > 0$, because firms with larger $\chi_{i\pi_{t+1}\pi_t}$ condition prices more strongly on signals of π_t, so they pay for more precision in those signals. This amplifies the causal effect of

[11] Macaulay and Moberly (2025) find evidence that among households greater Kalman gains on inflation signals are correlated with greater perceived inflation persistence, consistent with a positive narrative heterogeneity effect here and thus an amplified causal effect of current inflation expectations.

[12] The exception is level-k thinking, where for tractability I allow for signals on mc_t rather than π_t.

aggregate inflation expectations, as information is concentrated in those who respond most strongly to it (Proposition 5). This in turn amplifies the transmission of fundamental shocks (equation (20)).

Learning (Evans and McGough, 2020): if firms update their perceived laws of motion for inflation and marginal costs after observing their signals, then $\boldsymbol{\chi}_i$ is endogenous to information precision. Firms who observe more precise signals update their subjective models of the economy more strongly after each observation than those with less precise information. If there is heterogeneity in information precision, this therefore leads to systematic correlations between elements of $\boldsymbol{\tau}_i$ and $\boldsymbol{\chi}_i$. Notably, the narrative heterogeneity channel generated in this case may vary over time as firms continually update their perceived laws of motion as they receive new signals.

Sticky information (Mankiw and Reis, 2002): if firms can pay a cost to observe information more frequently (as in e.g. Reis, 2006, and others), then subjective models influence this information choice, with a similar logic to the rational inattention model.

Level-k thinking (Farhi and Werning, 2019): again, a narrative heterogeneity channel arises whenever firms can pay a cost to engage in a further round of reasoning, as in Alaoui and Penta (2016, 2022). In this case, firms with more precise information have more to gain from engaging in costly reasoning about how to use that information. Reasoning causes firms to update their cross-learning parameters, so the elements of $\boldsymbol{\chi}_i$ differ between firms with different $\boldsymbol{\tau}_i$.

Information delegation to media outlets (Nimark and Pitschner, 2019): firms with larger $\chi_{i\pi_{t+s}\pi_t}$ and $\chi_{imc_{t+s}\pi_t}$ place a greater value on information on π_t. Competitive media outlets specialize, and firms with large values of these cross-learning parameters select into reading those which report π_t more often.

In this case, the resulting narrative heterogeneity channel is state-dependent. When inflation is close to steady state, almost no outlets find inflation sufficiently newsworthy for them to report it, and most firms have $\tau_{i\pi_t} = 0$. Similarly, when inflation is very far from steady state, it is highly newsworthy, and almost all outlets do report it, leading most firms to have $\tau_{i\pi_t} = 1$. It is only in the intermediate region where some outlets report inflation, and others don't, that the precision of firm information becomes heterogeneous, generating a narrative heterogeneity channel.

Selective memory recall (Bordalo et al., 2023): if signals on current economic conditions act as cues for selective memory recall (as in Gennaioli et al., 2024), then firms with different information recall different past inflation experiences. Firms with precise signals (high τ_{π_t}) recall experiences clustered around the realized π_t, while those with less precise signals recall a wider set of experiences. If these recalled experiences are used to inform cross-learning, for example through perceptions of inflation persistence, then this implies a correlation between $\boldsymbol{\tau}_i$ and $\boldsymbol{\chi}_i$, and thus a narrative heterogeneity channel

Data generated through production (Farboodi and Veldkamp, 2021): firm subjective models affect their response to a given shock. That in turn affects how much data each firm has to learn from in subsequent periods, again generating a correlation between subjective models and information.

Heterogeneous price indices: for households, inflation rates vary substantially across the population (Kaplan and Schulhofer-Wohl, 2017; Kiss and Strasser, 2024). If households observe their own idiosyncratic price levels (heterogeneous information), and form expectations using the dynamics of those idiosyncratic price levels (heterogeneous cross-learning), then we could see a narrative heterogeneity effect even if all agents have rational expectations. I take this idea to the Calvo model by assuming firms belong to sectors which differ in the persistence of their inflation processes. Greater inflation persistence implies a weaker response to inflation signals but a stronger extrapolation from short to long-horizon inflation expectations, generating a negative covariance between $\tau_{i\pi_t}$ and $\chi_{i\pi_{t+s}\pi_t}$, even though all cross-learning is consistent with rational expectations.

Overall, the narrative heterogeneity channel is a robust feature of many models of expectation formation, either because subjective models influence information acquisition, or because information influences the formation of subjective models. With multiple frictions combined, it is possible for causation to run in both directions. Indeed, the final case (heterogeneous price indices) shows that the mechanism can arise even under rational expectations, if there is heterogeneity in the object to be forecast. Allowing for heterogeneity in both components of expectation formation therefore generates dynamics that are not present if heterogeneity is restricted to either information or subjective models alone. Indeed, the correlation between information and subjective models may provide a novel way to test competing models of expectation formation.

3 General model

We now show how to apply the temporary-information equilibrium approach to a general class of linear models in discrete time, and show that the narrative heterogeneity channel is a general feature of the causal effects of expectations.

Decision rules and expectations. There is a mass 1 of agents, indexed i. We start our analysis with a general linear policy function for agent actions, of the form

$$\boldsymbol{x}_{it} = A_i \boldsymbol{x}_{it-1} + \sum_{s=0}^{\infty} \left(B_{iws} E_{it} \boldsymbol{w}_{t+s} + B_{ixs} E_{it} \boldsymbol{x}_{t+s} \right) \tag{24}$$

where $\boldsymbol{x}_{it}$ is a $N_x \times 1$ vector of choice variables for agent i in period t; $\boldsymbol{w}_t$ is a $N_w \times 1$ vector of state variables; and $\boldsymbol{x}_t \equiv \int_i \boldsymbol{x}_{it} di$ denotes aggregate actions across agents. Each agent i takes $\boldsymbol{w}_t$ and $\boldsymbol{x}_t$ as given. A_i B_{iws}, and B_{ixs} are matrices of coefficients, which may vary across agents. This nests Equation (2) in the Calvo model above, as well as the equations determining agent choices in a range of other models.

As in Section 2.1, define $\boldsymbol{z}_t$ as the vector of all variables about which the agent will form expectations. Similarly, let B_i be the matrix that collects all B_{iws}, B_{ixs} coefficients.

$$\boldsymbol{z}_t \equiv \begin{bmatrix} \boldsymbol{x}_t \\ \boldsymbol{w}_t \\ \boldsymbol{x}_{t+1} \\ \boldsymbol{w}_{t+1} \\ \vdots \end{bmatrix}, \quad B_i \equiv \begin{bmatrix} B_{ix0}, B_{iw0}, B_{ix1}, B_{iw1}, \ldots \end{bmatrix} \tag{25}$$

Following the same steps as in Section 2.1, we can express the expectations of $\boldsymbol{z}_t$ as in equation (9). In Section 2, we restricted attention to upper-triangular $\boldsymbol{M}_i$ matrices. This implied that $\boldsymbol{\chi}_i$ had a unit diagonal. In this more general model, we relax that assumption. Doing this implies that $\boldsymbol{\tau}_i$ can no longer be interpreted as the extent of expectations updating after a direct signal.[13] However, it will be helpful for the exercises below to re-scale the $\boldsymbol{\chi}_i$ and $\boldsymbol{\tau}_i$ such that this original interpretation is restored. As in the first stage of an information-treatment RCT, we would like $\boldsymbol{\tau}_i$ to reflect the updating from signals to expectations of the variables being signalled.

To do that, define a new square matrix $\boldsymbol{D}_i$, where the diagonal elements are the diagonal elements of $\boldsymbol{\chi}_i$, but all off-diagonal elements are equal to 0. Then define $\hat{\boldsymbol{\chi}}_i = \boldsymbol{\chi}_i \boldsymbol{D}_i^{-1}$ and $\hat{\boldsymbol{\tau}}_i = \boldsymbol{D}_i \boldsymbol{\tau}_i$, so

$$E_{it}(\boldsymbol{z}) = \hat{\boldsymbol{\chi}}_i \hat{\boldsymbol{\tau}}_i \boldsymbol{s}_{it} \tag{26}$$

[13] For example, in the Calvo model a signal about π_{t+1} has a direct effect on $E_{it}\pi_{t+1}$, but also may affect $E_{it}mc_{t+1}$. If $\boldsymbol{M}_i$ is not upper-triangular, this change in $E_{it}mc_{t+1}$ may cause a further change in $E_{it}\pi_{t+1}$. $\tau_{i\pi_{t+1}}$ only captures the first of these effects, which is no longer the full update after the signal.

This is of the form in equation (9), but $\hat{\boldsymbol{\chi}}_i$ is now guaranteed to have a unit diagonal, such that $\boldsymbol{\tau}_i$ captures the updating from direct signals. Substituting this into equation (24) yields

$$\boldsymbol{x}_{it} = A_i \boldsymbol{x}_{it-1} + B_i \hat{\boldsymbol{\chi}}_i \hat{\boldsymbol{\tau}}_i \boldsymbol{s}_{it} \tag{27}$$

General equilibrium. The state variables $\boldsymbol{w}_t$ are such that

$$C\boldsymbol{x}_t + D\boldsymbol{w}_t + F\boldsymbol{\xi}_t = 0 \tag{28}$$

where $\boldsymbol{\xi}_t$ is a $N_\xi \times 1$ vector of exogenous shocks, and C, D, F are matrices of coefficients.

In the Calvo model in Section 2, the only element of $\boldsymbol{w}_t$ was mc_t, and this was exogenous to firm actions (i.e. pricing). Thus the constraint was, in effect, that $mc_t + \xi_{mc,t} = 0$. Equation (28) nests this, but also nests a wide range of other equilibrium conditions, which could consist of technological constraints, market clearing conditions, or the decision rules of other categories of agents whose expectations are not the focus of the analysis.

Temporary-information equilibrium. For given distributions of $\hat{\boldsymbol{\tau}}_i, \hat{\boldsymbol{\chi}}_i$, and given exogenous shocks $\boldsymbol{\xi}_t$, the temporary-information equilibrium is a sequence of actions $\boldsymbol{x}_{it}$, endogenous states $\boldsymbol{w}_t$, and signals $\boldsymbol{s}_{it}$ such that

1. $\boldsymbol{x}_{it}$ is chosen according to the policy function (27), and aggregate actions are $\boldsymbol{x}_t \equiv \int_i \boldsymbol{x}_{it} di$;
2. $\boldsymbol{w}_t$ satisfies the general equilibrium condition (28);
3. a specified process links signals $\boldsymbol{s}_{it}$ with states $\boldsymbol{w}_t$ and aggregate actions $\boldsymbol{x}_t$.

In principle, many different processes could be specified for item (3). For simplicity, however, I restrict attention here to the form used in Section 2.2:

$$s_{it}(z_{jt}) = z_{jt} + \varepsilon_{jt} + \tilde{\varepsilon}_{ijt} \tag{29}$$

where $\varepsilon_{jt}, \tilde{\varepsilon}_{ijt}$ are mean-zero white noise. The public noise ε_{jt} is common across agents i, the idiosyncratic noise $\tilde{\varepsilon}_{ijt}$ is independent across agents. From here, all results follow as in Section 2.

Causal effects of expectations in partial equilibrium. The causal effect of $\bar{E}_t(z_{jt})$ on an element of the aggregate choice vector $\boldsymbol{x}_t$ in partial equilibrium is given by Proposition 6. Note that partial equilibrium here means that all signals aside from $s_{it}(z_{jt})$ are held constant: the change in aggregate actions does not affect signals, and nor does any change in the state variables $\boldsymbol{w}_t$ through the general equilibrium condition (28).

Proposition 6 *The causal effect of a shock ε_{jt} in partial equilibrium is*

$$\Omega_k^{PE}(z_{jt}) = \frac{1}{\bar{\hat{\tau}}_j} \cdot \frac{dx_{kt}}{d\varepsilon_{jt}} = \sum_{m=0}^{N_z} \left(\bar{B}_{km}\bar{\hat{\chi}}_{mj} + Cov_I \left[B_{ikm}, \hat{\chi}_{imj} \frac{\hat{\tau}_{ij}}{\bar{\hat{\tau}}_j} \right] + \bar{B}_{km} Cov_I \left[\hat{\chi}_{imj}, \frac{\hat{\tau}_{ij}}{\bar{\hat{\tau}}_j} \right] \right) \quad (30)$$

where B_{ikm} denotes the k,mth element of B_i, and $\bar{B}_{km} \equiv \int_i B_{ikm} di$. The scaling by $\bar{\tau}_j^{-1}$ ensures that the information shock is such that $\bar{E}_t(z_{jt})$ increases by 1.

This consists of three sets of effects. The first term inside the summation is the *representative agent channel*: the effects of the average policy function coefficients and average cross-learning across agents. This summarizes all shock transmission channels in models with a representative agent, and indeed in many models with heterogeneity, where average expectation formation is sufficient to capture shock responses to first order.[14]

The second term is the *response heterogeneity channel*. Since $\hat{\chi}_{imj}\hat{\tau}_{ij}$ gives the response of $E_{it}(z_{mt})$ to a signal on z_{jt}, this reflects that shocks will be amplified if the agents whose expectations react the most to the shock are the agents whose actions are most sensitive to those expectations (i.e. who have the largest B_{ikm}). Macaulay and Moberly (2025) provide evidence of one such correlation, between the behavior of inflation expectations and liquidity constraints among German households. This channel, for other expectations, is also behind the novel dynamics in Broer et al. (2020), Macaulay (2021), and Guerreiro (2023).

Finally, the third term is the *narrative heterogeneity channel*, as explored for the Calvo model in Section 2.4. Heterogeneous expectations can generate a channel of aggregate shock transmission even if every agent has the same policy function, if information ($\hat{\tau}_{ij}$) is correlated with cross-learning ($\hat{\chi}_{imj}$) across agents. Cross-learning determines an agent's response to a given piece of information, so if information is concentrated among agents with particular non-representative $\hat{\boldsymbol{\chi}}_i$, that distorts the aggregate response away from the representative agent effect.

Setting $\hat{\tau}_{ij} = 1$ for all i, j recovers the effect in standard temporary equilibrium, in which $E_{it}(z_{jt})$ shifts equally for all agents, but no other expectations change. This abstracts from all effects of heterogeneity through the response heterogeneity and narrative heterogeneity channels. All that matters is how, on average, expectations of z_{jt} enter the policy function for x_{kt}.

$$\Omega_k^{PE}(z_{jt}|\hat{\tau}_{ij} = 1 \forall i, j) = \bar{B}_{kj} \quad (31)$$

[14] In Andrade et al. (2019), for example, households differ in their interpretation of forward guidance announcements, but the aggregate effects of an announcement depend only on the average over this mix of beliefs. Heterogeneity might affect the average information or cross-learning behavior sustained in equilibrium, but unless one of the other two terms in equation (30) is non-zero, those averages alone drive aggregate shock transmission.

Setting $\hat{\tau}_{ij} = \bar{\hat{\tau}}_j < 1$ and $\hat{\chi}_{ix_{jt+s}x_{jt}} = 1 + s$ for all i, j recovers the effect in extended temporary equilibrium, as in Werning (2022). Again, this abstracts from all effects of heterogeneity, though now a broader range of policy function coefficients determine the aggregate response, as signals on x_{kt} also impact expectations of x_{kt+1}, x_{kt+2}, etc.

$$\Omega_k^{PE}(z_{jt}|\hat{\tau}_{ij} = \bar{\hat{\tau}}_j < 1, \hat{\chi}_{ix_{jt+s}x_{jt}} = 1 + s, \forall i, j) = \sum_{s=0}^{\infty} \bar{B}_{kjs}(1 + s) \tag{32}$$

where $\bar{B}_{kjs}$ denotes the average coefficient on z_{jt+s} in the row of the policy function (27) corresponding to x_{ikt}.

Causal effects of expectations in general equilibrium. The causal effect of $\bar{E}_t(z_{jt})$ on an element of the aggregate choice vector $\boldsymbol{x}_t$ in general equilibrium is given by Proposition 7.

Proposition 7 *The effect of an exogenous shift in information due to ε_{jt} on aggregate action x_{kt} is*

$$\frac{dx_{kt}}{d\varepsilon_{jt}} = \left[\boldsymbol{\Lambda}^{GE}\boldsymbol{\Omega}^{PE}\right]_{k,j} \bar{\hat{\tau}}_j \tag{33}$$

where $\boldsymbol{\Omega}^{PE}$ is a $N_x \times N_z$ matrix collecting the partial-equilibrium effects of each expectation $\bar{E}_t(z_{jt})$ on each aggregate action x_{kt}. $\boldsymbol{\Lambda}^{GE}$ is a $N_x \times N_x$ matrix capturing general-equilibrium effects, defined in Appendix A.

The causal effect of $\bar{E}_t(z_{jt})$ is obtained by scaling $\frac{dx_{kt}}{d\varepsilon_{jt}}$ such that the response of $\bar{E}_t(z_{jt})$ is equal to 1, which implies

$$\Omega_k^{GE}(z_{jt}) = \frac{\left[\boldsymbol{\Lambda}^{GE}\boldsymbol{\Omega}^{PE}\right]_{k,j} \bar{\hat{\tau}}_j}{\bar{\hat{\tau}}_j + \Gamma_j} \tag{34}$$

where Γ_j, which captures general-equilibrium effects of the information shock ε_{jt} on $\bar{E}_t(z_{jt})$, is defined in Appendix A.

Importantly, the GE feedback matrix $\boldsymbol{\Lambda}^{GE}$ will typically not be diagonal. This means that when computing the effects of $\bar{E}_t(z_{jt})$ on a specific aggregate action x_{kt}, the effects of that expectation on *all* aggregate actions and state variables will affect the result. If $\bar{E}_t(z_{jt})$ affects other variables aside from x_{kt}, this will cause changes in agent signals about those other variables, leading to further changes in x_{kt} beyond the partial-equilibrium effect. This generalizes the result in Proposition 3, where shocks to $\bar{E}_t(\pi_{t+1})$ in part propagated by influencing π_t, which in turn affected $\bar{E}_t(\pi_t)$.

To showcase more clearly how this differs from the partial-equilibrium case, and from the simple Calvo model of Section 2, Corollary 1 in Appendix A shows $\boldsymbol{\Lambda}^{GE}$ in the special case where agents only receive signals about current realized $\boldsymbol{x}_t$ and $\boldsymbol{w}_t$ (i.e. $\tau_{ij} = 0$ for all z_{jt} realized after period t). In this case, $\boldsymbol{\Lambda}^{GE}$ simplifies substantially, highlighting that general equilibrium feedback from exogenous variation in expectations can come through two distinct channels. First, as in the Calvo model, the shock causes a change in aggregate actions $\boldsymbol{x}_t$, which feeds back into the information agents receive. Second, this change in $\boldsymbol{x}_t$ also affects the state variables $\boldsymbol{w}_t$ through the equilibrium condition (28), which *also* affects agent information. This was not present in Section 2 because we abstracted there from all responses of the state (marginal costs) to firm actions.

The other key insight from this is that all three channels identified in partial equilibrium also affect the strength of this general equilibrium feedback. When exogenous variation in a given expectation affects aggregate actions and states, this affects the information agents receive about those variables, changing their expectations again. The effects of that second round of expectation updating also depend on the representative-agent, response heterogeneity, and narrative heterogeneity channels identified in Proposition 6.

The role of expectations in the transmission of fundamental shocks. The causal effect of an exogenous shock $\boldsymbol{\xi}_t$ on aggregate choices is given by Proposition 8, which follows Section 2.3.

Proposition 8 *The causal effect of a fundamental shock $\boldsymbol{\xi}_t$ on aggregate choices is*

$$\frac{d\boldsymbol{x}_t}{d\boldsymbol{\xi}_t} = -\boldsymbol{\Lambda}^{GE}\boldsymbol{\Omega}^{PE}_{w0}\bar{\hat{\boldsymbol{\tau}}}_{w0}D^{-1}F \tag{35}$$

where the matrix $\boldsymbol{\Omega}^{PE}_{w0}$ collects the partial-equilibrium effects of each element of expectations about current state variables ($\bar{E}_t(\boldsymbol{w}_t)$), and $\bar{\hat{\boldsymbol{\tau}}}_{w0}$ is a diagonal matrix collecting the corresponding elements of average $\boldsymbol{\tau}_i$.

Intuitively, the general equilibrium condition (28) implies that $-D^{-1}F$ is the response of state variables $\boldsymbol{w}_t$ to a $\boldsymbol{\xi}_t$ shock, in the absence of any $\boldsymbol{x}_t$ response. $\boldsymbol{\Omega}^{PE}_{w0}\bar{\hat{\boldsymbol{\tau}}}_{w0}$ then gives the partial-equilibrium effect of that $\boldsymbol{w}_t$ change on $\boldsymbol{x}_t$, through changes in expectations of $\boldsymbol{w}_t$. Finally, $\boldsymbol{\Lambda}^{GE}$ accounts for general equilibrium feedback through both $\boldsymbol{w}_t$ and $\boldsymbol{x}_t$.

Next, construct a counterfactual transmission in which there is no response of $E_{it}(z_{jt})$ for some set of variables defined by $j \in J$. To do this, replace $\hat{\tau}_{ij} = \hat{\chi}_{ijk} = 0$ for all i, k, and for all $j \in J$, then re-compute $\boldsymbol{\Lambda}^{GE}, \boldsymbol{\Omega}^{PE}_{w0}, \hat{\boldsymbol{\tau}}_{w0}$. The share of shock transmission due to expectations of z_{jt} for $j \in J$ is then given by $\left(\frac{d\boldsymbol{x}_t}{d\boldsymbol{\xi}_t} - \frac{d\boldsymbol{x}_t}{d\boldsymbol{\xi}_t}\Big|_{\sim E_{j\in J}}\right) \cdot \left(\frac{d\boldsymbol{x}_t}{d\boldsymbol{\xi}_t}\right)^{-1}$.

4 Mapping RCT estimates into general equilibrium

Since the temporary-information equilibrium method is conceptually linked to information-treatment RCTs, the expectation statistics in the model framework correspond closely to parameters estimated in those experiments. The temporary-information equilibrium approach can therefore be used to translate RCT estimates into general equilibrium insights, overcoming the typical missing intercept problem in the empirical literature. This helps us uncover whether the departures from full-information rational-expectations observed in a given RCT are large or small, in terms of their macroeconomic implications. We can also separate the implied causal effects of expectations into their representative-agent and narrative-heterogeneity components.

I focus here on applying the method to the recent RCT conducted on European firms by Baumann et al. (2024), as this lends itself naturally to interpretation through the Calvo model already studied in Section 2. However, the methods Baumann et al. (2024) employ are common to a broad range of survey experiments, so the same procedure could be applied in other experiments, using the temporary-information equilibrium of other models as appropriate to the context.

The missing intercept in information-treatment RCTs. RCTs like Baumann et al. (2024) use information treatments to exogenously vary the expectations of some survey participants, while keeping others as a control group. This cross-sectional variation enables researchers to identify the causal effect of expectations on individual beliefs and actions. While this has advanced our understanding of the micro-level effects of expectations, it is often hard to determine whether the results are of a magnitude that is meaningful for macroeconomic outcomes. This is because the individual-level effect cannot simply be aggregated to give a total macroeconomic effect of a given expectation change. If *all* inflation expectations change, for example, then aggregate realized inflation may change, which in turn may affect firm expectations and decisions. Interest rates may also adjust, as well as the labor market, and other variables, each of which have their own consequences. These general equilibrium effects form the 'missing intercept' between the estimates from an information-provision RCT and macroeconomic insights.

This issue is present in all empirical work using cross-sectional variation to identify partial-equilibrium elasticities (see e.g. discussions in Nakamura and Steinsson, 2018; Guren et al., 2020; Wolf, 2023, among many others). In information-provision experiments, however, the problem is particularly acute, for two reasons. First, the experiment typically occurs only once, or a small number of times. There is therefore limited scope for time-series approaches that could complement the cross-sectional identification (as in e.g. Donaldson, 2025; Matthes et al., 2025).

Second, the very features of expectations that drive the partial-equilibrium RCT results also affect the missing general equilibrium feedback.[15] In studies of fiscal policy (Nakamura and Steinsson, 2014) researchers can use estimates from other literatures to quantify the general-equilibrium intercept, thanks to equivalence results in a range of macroeconomic models (Wolf, 2023). The same approach cannot be taken to study the aggregate effects of changes in expectations, because the features of expectations revealed in the partial-equilibrium RCT estimates *also* affect that intercept. If, for example, firms over-extrapolate from current to future inflation, that will affect both the RCT estimates and firm responses to subsequent general equilibrium changes in aggregate variables.

The temporary-information equilibrium approach developed here provides a way through this issue: the results in Section 3 can be applied to interpret the results from information-treatment RCTs in any model fitting the form introduced there.

4.1 Extracting key expectation statistics.

Baumann et al. (2024) elicit prior beliefs about inflation in 1, 3, and 5 years time. They then treat a random subsample of participants with information about inflation. They study two treatments: one treatment group is shown information on current inflation, and another is shown information on one-year-ahead inflation (from SPF forecasts).

Total effects. As standard in the RCT literature, Baumann et al. (2024) first estimate the regression:

$$E_{it}^{post}\pi_{t+h} = a_0 + a_1 \cdot E_{it}^{prior}\pi_{t+h} + \sum_j b_j \cdot E_{it}^{prior}\pi_{t+h} \cdot D_{ij} + \sum_j c_j \cdot D_{ij} + \Gamma_{it} + \text{ error}_i \quad (36)$$

where D_{ij} is an indicator equal to 1 if firm i is in treatment group j, and Γ_{it} is a vector of controls. The coefficients b_j is $-1\times$ the average degree of updating from the observed signal in treatment group j. For the group shown a signal about π_{t+1}, the b_j estimate when $E_{it}\pi_{t+1}$ is the dependent variable is therefore equal to $-\bar{\tau}_{\pi_{t+1}}$ in our model.

For other combinations of treatment and dependent variable, the signal is no longer precisely about the dependent variable. The coefficient therefore represents the average of $-\frac{\partial E_{it}(z_{jt})}{\partial s_{it}(z_{kt})} = -\chi_{ijk}\tau_{ik}$ for treatment variable k and outcome variable j across the sample. From Table 5 in Baumann et al. (2024) we therefore obtain $\int_i \chi_{ijk}\tau_{ik} di = \bar{\chi}_{jk}\bar{\tau}_k + Cov_I[\chi_{ijk}, \tau_{ik}]$, i.e. a sum of representative agent and narrative heterogeneity effects.

If all that is desired is an estimate of the total causal effects of inflation expectations on aggregate

[15] While not an issue in RCTs, this also suggests that any attempt to use heterogeneous exposure to aggregate expectations shocks to identify the causal effect of expectations will suffer by construction from the issue of heterogeneous exposure to general equilibrium effects discussed in Donaldson (2025).

inflation, these statistics would be sufficient. In Propositions 1-3, those causal effects are derived in terms of precisely the same average effects.

Decomposition. The second regression Baumann et al. (2024) run is:

$$E_{it}^{post} X_i = \gamma_0 + \gamma_1 \cdot E_{it}^{post} \pi_{t+h} + \gamma_2 \cdot E_{it}^{prior} \pi_{t+h} + \Gamma_{it} + \text{ error}_i \tag{37}$$

for a variety of expected dependent variables $E_{it}^{post} X_i$. $E_{it}^{post} \pi_{t+h}$ is instrumented with variation from the information treatment. Specifically, equation (36) is used as a first stage.

Baumann et al. (2024) estimate several versions of equation (37). The relevant specifications here are those in which a single treatment group is compared to the control group who received no information. In these cases, γ_1 captures the average effect of an exogenous change in expected π_{t^*} on expectations of the dependent variable, where $\pi_{t^*} \in \{\pi_t, \pi_{t+1}\}$ depending on which treatment group is being studied. This is exactly the definition of $\bar{\chi}_{\pi_{t^*},X}$ in the temporary-information equilibrium of the Calvo model. Combining this with decompositions such as equation (14) (or more generally equation (30)), we obtain an estimate of the representative-agent component of the causal effects of inflation expectations on inflation, which we can combine with the total effect described above to infer the remaining narrative-heterogeneity component.

4.2 The causal effect of expected future inflation on current inflation

I now compute the causal effect of one-year-ahead inflation expectations on realized current inflation. I focus only on the effects that operate through the updating of inflation expectations at various horizons, abstracting from any cross-learning from π_{t+1} to marginal cost expectations.

This represents the best showcase for the temporary-information equilibrium approach, as all of the required expectation statistics are reported directly by Baumann et al. (2024) or are easily calculable in their data.[16] Formally, I abstract from marginal-cost cross-learning by setting $\chi_{imc_{t+s}\pi_t} = \chi_{imc_{t+s}\pi_{t+1}} = 0$ (relaxed in Appendix C). I then calibrate the remaining elements of $\boldsymbol{\chi}_i$ and $\boldsymbol{\tau}_i$ to the estimated coefficients in equations (36) and (37), assuming one period is one year (full details in Appendix C). The only remaining parameters are then β and λ. I set the former to 0.96, and plot the results for varying values of price stickiness λ.

[16]The only statistics not reported by Baumann et al. (2024) are the estimates of γ_1 in equation (37) with posterior inflation expectations at medium- and long-term horizons as the dependent variable. While these are not required for calculating the total causal effect of expectations, they allow the decomposition into representative-agent and narrative-heterogeneity components. I am indebted to Timo Reinelt for kindly providing me with these estimates.

Benchmarks. To interpret the RCT-implied causal effects of inflation expectations, we need benchmarks with which to compare them. I consider two benchmarks: perfect foresight and an approximation to limited-information rational expectations. For perfect foresight, we straightforwardly set $\tau_{ij} = 1$ for all variables z_j, which in turn implies $\boldsymbol{\chi}_i$ becomes the identify matrix (equation (10)).

Rational expectations is harder to specify, because we do not know the true joint dynamics of inflation and marginal costs in the Euro Area. I provide a simple approximation here by regressing annual Euro Area HICP inflation on a 12-month lag of itself. I then assume that $\bar{\tau}_{\pi_{t+1}}$ is as in the RCT-based calibration, but set cross-learning so that all firms extrapolate to future inflation using the degree of persistence implied by that AR(1) regression. Full details are given in Appendix C.

Finally, note that the survey does not provide an estimate of $\bar{\tau}_{\pi_t}$, the average $E_{it}(\pi_t)$ response to signals about π_t. This will matter for the evaluation of general equilibrium effects (see Proposition 3 and the associated discussion). In the RCT calibration and in the benchmarks I assume $\bar{\tau}_{\pi_t} = 1$, as if all firms are perfectly well-informed about current inflation.

The total causal effect of expectations. Figure 1 plots the causal effect of one-year-ahead inflation expectations on current inflation as price stickiness λ varies. Panel (a) plots the partial equilibrium effect $\Omega^{PE}(\pi_{t+1})$, while panel (b) plots the general equilibrium effect $\Omega^{GE}(\pi_{t+1})$, incorporating the endogenous response of π_t signals.

Figure 1: Partial and general equilibrium effects of a shock to $s_{\pi_{t+1}}$, comparing RCT estimates to benchmarks

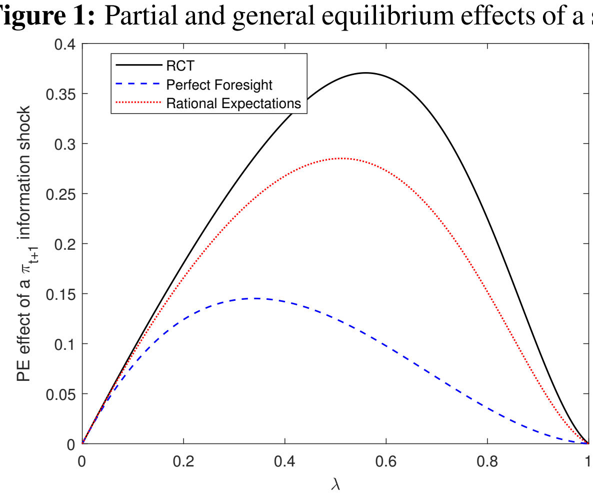

(a) Partial Equilibrium

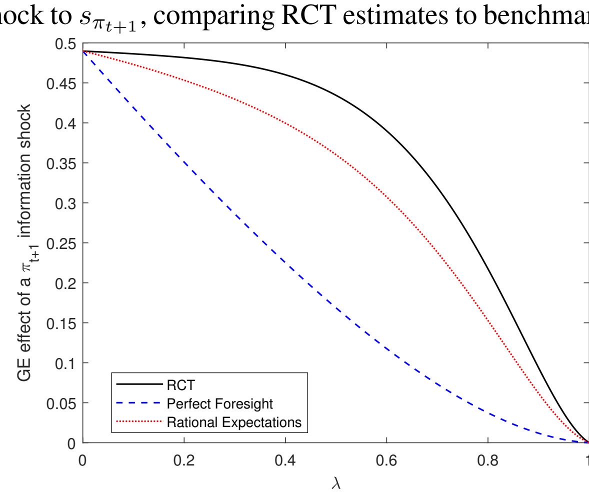

(b) General Equilibrium

Note: Figure plots the partial- and general-equilibrium causal effects of an increase in $\bar{E}_t\pi_{t+1}$ in the Calvo model, following the formulae in Propositions 2 and 3, and a calibration derived from the estimates in Baumann et al. (2024) (detail in Appendix C).

The key message from these figures is that the macroeconomic significance of the RCT results depends critically on the degree of price stickiness. With an average price duration of 5 quarters ($\lambda = 0.2$ since the calibration is annual), the partial equilibrium effect implied by the RCT estimates is 9% larger than under rational expectations, and the general equilibrium effect is 6% larger. However, if prices are sufficiently sticky that the average price duration is two years ($\lambda = 0.5$),[17] the RCT estimates imply causal effects of expectations that are 28% and 21% larger than rational expectations in partial and general equilibrium respectively. The differences to the perfect foresight benchmark are universally larger, but the same strong dependence on price stickiness is present. With $\lambda = 0.2$, the RCT implies a 46% larger partial equilibrium effect, and a 37% larger general equilibrium effect. With $\lambda = 0.5$, these figures jump to 191% and 158%.

The reason for this dependence is that the RCT estimates depart most strongly from rational expectations at long horizons, where cross-learning from π_t and π_{t+1} information is found to be much stronger than the benchmark would predict.[18] This means that a given change in one-period-ahead expectations is associated with greater updating of long-horizon expectations, and thus a larger inflation response. The difference is even more pronounced when comparing to perfect foresight, where there is no such cross-learning at all.

At very flexible prices (small λ), this does not have a large quantitative effect, because price durations are expected to be short and so long-horizon expectations are relatively unimportant for price-setting. As prices get stickier, long-horizon expectations become more important to current price-setting, and the differences between the RCT-implied model and the benchmarks grow.

Furthermore, price stickiness also determines the importance of the general-equilibrium feedback, i.e. the 'missing intercept'. At $\lambda = 0.2$, the general equilibrium effect of expected inflation is more than $2.5\times$ larger than the partial equilibrium effect. At $\lambda = 0.5$, it is only 20% larger. The other key determinants are the RCT results themselves, because the general-equilibrium effect of an increase in realized π_t depends on how firms incorporate that signal into their expectations of inflation at all horizons, which is a key object studied in the RCT. In Appendix C I show this concretely by plotting the causal effect of expectations using the RCT calibration for the partial-equilibrium effect, but taking the general-equilibrium feedback from either the perfect foresight or rational expectations benchmark. In sharp contrast to the correct results in Figure 1b, this exercise finds that the general-equilibrium effect of inflation expectations is initially *increasing* in price stickiness, with either of these alternative general-equilibrium assumptions.

Appendix C also contains a number of other results extending this analysis. In particular, I apply the methods described in Section 2.3 to study the share of transmission of a marginal-cost

[17]This is the calibration used in Werning (2022) footnote 8).

[18]To get the rational expectations benchmark to deliver the same cross-learning from π_{t+1} to π_{t+5} as implied by the RCT, the annual persistence of inflation would need to be 0.82, more than double its estimated value of 0.39.

news shock that is due to inflation expectations, and find that even the sign of the departure from rational expectations implied by the RCT results varies with price stickiness. I also compute alternative benchmarks using two popular models that depart from rational expectations: diagnostic expectations (Bordalo et al., 2020) and inattentiveness (Gabaix, 2020). In each case, adjusting the benchmark simply involves replacing the $\boldsymbol{\chi}_i$ and $\boldsymbol{\tau}_i$ parameters with those implied by the researcher's chosen theory of expectations.

The narrative heterogeneity channel. Figure 2 decomposes the partial-equilibrium effect of expected inflation on realized inflation into its representative-agent and narrative-heterogeneity components. The causal effect of expectations is stronger when we consider the representative-agent effects in isolation. This implies that the narrative heterogeneity effect is negative. In other words, the firms who process the most information about π_{t+1} extrapolate the least to inflation at longer horizons. As a proportion of the total effect, the magnitude of the narrative heterogeneity channel is monotonically increasing in λ. At $\lambda = 0.5$, it reduces the causal effect of expected inflation by 6%, narrowing the gap between the RCT results and the rational expectations benchmark.

Figure 2: Partial equilibrium effects of a shock to $s_{\pi_{t+1}}$, decomposition

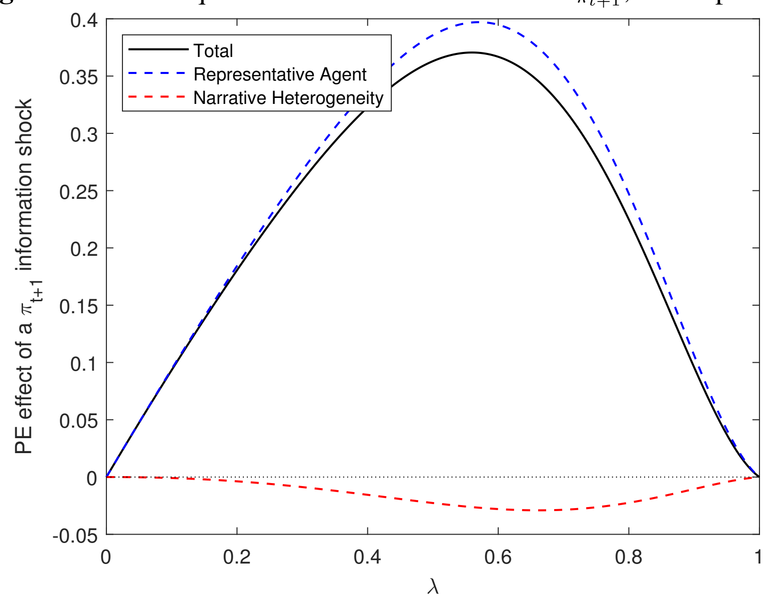

Note: Figure plots (in black) the partial-equilibrium causal effects of an increase in $\bar{E}_t\pi_{t+1}$ in the Calvo model, following the formula in Proposition 2, and a calibration derived from the estimates in Baumann et al. (2024) (detail in Appendix C). The blue and red lines are the decomposition of this effect into representative-agent and narrative-heterogeneity channels, following Proposition 6.

This negative narrative heterogeneity channel may seem initially surprising, given that simple models of information acquisition would predict that information is more valuable if the underlying process is believed to be persistent. However, many other forces could generate such a negative correlation between information and cross-learning between inflation horizons. In particular, the RCT results imply that firms on average extrapolate strongly from π_{t+1} signals to longer horizons. If this overextrapolation relative to the rational expectations benchmark is a reasoning error on the part of firms, then the firms who are better-informed may use that information to update their beliefs about inflation persistence towards the truth, thus leading to a negative correlation between information precision and cross-learning.

This sharp decomposition is only possible when considering the causal effects of expectations in partial equilibrium. In general equilibrium, the total causal effect of expectations is a nonlinear combination of representative-agent and narrative-heterogeneity channels (e.g. equation (16)). The difference between the total general-equilibrium effect and the counterfactual if only representative-agent channels were present is therefore indicative of the influence of heterogeneous expectation formation, but does not have the same interpretation as a specific covariance. Moreover, as I assume here that firms observe current inflation precisely ($\tau_{\pi_t} = 1$), there are no further narrative heterogeneity effects in the missing intercept beyond those shown in the partial-equilibrium result in Figure 2. Mechanically, therefore, the dampening of the partial-equilibrium effect due to narrative heterogeneity also dampens the general-equilibrium effect. This is shown in full in Appendix C.

5 Survey evidence on narrative heterogeneity

In Section 4, I measured the narrative heterogeneity channel implicitly, backing it out from a decomposition of the Calvo model in Section 2. In this section I measure it directly, in the context of household beliefs around inflation. I document three empirical results that reveal a strong and time-varying narrative heterogeneity channel in this context. These results are used to inform the model in Section 6, from which I observe the aggregate consequences of this instance of the narrative heterogeneity channel.

5.1 Data

Most empirical papers on information frictions or subjective models use data on realized expectations. Expectations combine both information and subjective models (Sections 2 and 3), and so cannot be used to identify the narrative heterogeneity channel. To overcome this, I use data from the Bank of England Inflation Attitudes Survey (IAS), which contains several unique questions which enable me to measure these components of expectation formation separately – and thus to measure their joint distribution.

The IAS is a quarterly survey of a repeated cross-section of UK households, run since 2001 (annual until 2003). After weighting, the sample is representative of the UK adult population. I use the individual-level response data from 2001-2019, omitting the quarters conducted after the onset of the Covid-19 pandemic, as the implementation of the survey had to be changed substantially at this time (see Bank of England, 2020).

The first survey question I use asks households about their subjective model of the relationship between inflation and the 'strength of the economy'.

Question 1 *If prices started to rise faster than they are now, do you think Britain's economy would end up stronger, or weaker, or would it make little difference?*

This differs from standard questions on expected future economic outcomes because it does not invoke the use of information about the state of the world. The answers are informative about cross-learning only (χ_{ijk} in the models above).[19] In the analysis below, I will refer to a respondent answering that inflation would make the economy stronger/little difference/weaker as having a positive/neutral/negative subjective model of inflation respectively.

There are two possible interpretations of this question. Households may view it as asking about the causal effects of inflation on the economy (as in the model of Spiegler, 2021). Alternatively, they could see it as asking about the most likely source of a rise in inflation, if they believe supply- and demand-driven inflation is associated with different real outcomes (as in e.g. Kamdar and Ray, 2024). For the purposes of this section, this distinction does not matter. χ_{ijk} is defined above as the degree to which households update their expectations of variable j when their expectation of variable k changes, irrespective of why that link exists. In this case, it is the updating of expectations about the strength of the economy when expected inflation rises. The sign of this updating is captured by the question, whether it occurs because of a perceived direct causal link from inflation to the real economy, or a belief about the type of shocks hitting the economy.

The next set of questions concern the information households use to form their inflation expectations, without asking what those expectations are. This allows us to learn about household information precision (encoded in τ_{ik}) without contamination from cross-learning (χ_{ijk}).

Question 2a *What were the most important factors in getting to your expectation for how prices in the shops would change over the next 12 months?*

Please select up to 4 [from a list reproduced in Appendix D.1].

[19] This is similar to how Andre et al. (2022) interpret the answers to the hypothetical vignettes in their surveys.

Question 2a was only asked in 2016Q1, but very similar questions were asked at other times. In each, the respondent is asked about the information sources they used to arrive at their expected inflation, or that led them to change that expectation over the previous year. In each, some of the options that respondents can choose refer to information specifically on inflation, while others refer to information on other variables, which can only be used to forecast inflation if the household engages in some cross-learning. For each such question I construct a dummy variable equal to 1 if the respondent reports using direct information about inflation, and equal to 0 if they do not. Full details of these questions, and the options representing direct information, are in Appendix D.1. Combining these dummy variables gives an indicator for if the respondent used direct information on inflation in forming their expectations, that is whether $\tau_{i\pi_t} > 0$. This indicator is observed for 8 separate quarters between 2009Q1-2019Q1. I confirm below that the key results of this section do not vary substantially with the changes in question wording over these periods.

In Appendix D.2 I confirm that these measures of information and subjective models correlate with planned household consumption, in ways that are consistent with the measures picking up the desired elements of household beliefs. A further possible test of the information indicator would ask if households who obtain direct information about inflation make more accurate forecasts. However, if beliefs about the level of inflation affect subjective models, that may in turn change the incentives to acquire further information, complicating the predicted correlation between information and forecast accuracy.[20] Indeed, recent evidence from Link et al. (2025) suggests that greater household information acquisition is often associated with larger forecast errors about inflation.

The other questions used in this section are standard, asking households to give point estimates for current and one-year-ahead annual inflation. For each question, responses consist of a 1 percentage-point bin between -5% and +10%, or end ranges $\leq -5\%, \geq 10\%$. For the exercises in Section 5.4, I code perceptions and expectations at the midpoint of the selected bin, with the lowest and highest bins coded as -5.5% and 10.5% respectively.

5.2 Information and subjective models in the cross-section

With this data, I first explore the cross-sectional relationship between the distributions of information and subjective models, which is the source of any narrative heterogeneity channel. Table 1 shows the estimated average marginal effects from a probit regression of the information indicator defined in Section 5.1 on the respondent's subjective model of inflation, represented by their answer to Question 1. This gives us our first empirical result.

[20] In a previous version of this paper (Macaulay, 2022) I explicitly model this process and show that the data line up qualitatively with the model.

Table 1: Information correlates with subjective models

	(1)	(2)
end up stronger	-0.0102	-0.00827
	(0.0191)	(0.0192)
make little difference	-0.0356***	-0.0315**
	(0.0128)	(0.0129)
dont know	-0.0627***	-0.0605***
	(0.0172)	(0.0172)
Controls	None	All
Time FE	Yes	Yes
Observations	8270	8270

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports average marginal effects from estimating a probit regression of the information indicator on the responses to Question 1. The information indicator equals 1 if the household reports using a direct source of information about inflation when forming their expectations, as defined in Appendix D.1. The omitted category is the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS. All regressions include quarter fixed-effects. The controls in column 2 are are gender, age, class, employment status, income, education, region, and home-ownership status. Age, class, income and education are all reported in bands, and included as categorical variables.

Empirical Result 1 *Households who believe inflation makes no difference to the economy acquire less information about inflation on average than households who believe inflation does affect the economy (in either direction).*

The probability of using direct inflation information is 3-3.5 percentage points lower for those with a neutral model of the effects of inflation than those who believe inflation weakens the economy. Over the whole population 23% of respondents use direct inflation information, so this difference is non-trivial. More important than the magnitude, however, is that this shows a systematic cross-sectional relationship between information and subjective models, indicating a role for the narrative heterogeneity channel. Quantifying the effect requires a model, as developed in Section 6 below.

The information indicator is composed of answers to several slightly different questions, with small wording and definition changes across survey waves. In Appendix D.3 I repeat the regressions of Table 1 on subsets of the questions, and find the results are robust. As some respondents do not answer the unique survey questions used here, I also account for the concern that there may be selection bias in whose answers are observed, using a selection model as in Heckman (1979). Again, the results are robust.

Result 1 is not consistent with models with exogenous information, as there would be no reason for information to be systematically correlated with household subjective models. It is, however, consistent with models of endogenous information acquisition, as explored in Section 2.4. Intuitively, the value of inflation information is lower for households who believe inflation makes little difference to other variables that matter for their decisions.[21]

5.3 Subjective models over time

Figure 3 shows the proportions answering Question 1 with each subjective model of inflation over time ('don't know' omitted for figure clarity).

Figure 3: Proportions giving each answer to Question 1: "If prices started to rise faster than they are now, do you think Britain's economy would end up stronger, or weaker, or would it make little difference?"

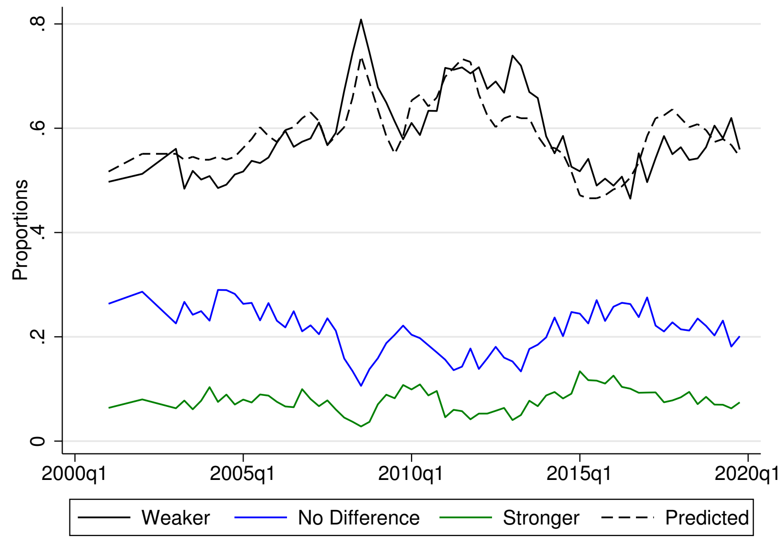

Note: Proportions shown are calculated using the survey weights provided in the IAS. Proportion answering 'Don't know' is omitted for figure clarity. The dashed line is the predicted values from regressing the proportion reporting that inflation makes the economy weaker on annual CPI inflation: $\widehat{\Pr(\text{weaker})}_t = 0.057 \times \text{CPI inflation}_t + 0.466$. The coefficient on inflation is significant at the 1% level.

The majority of households answer that inflation would make the economy weaker in all quarters, in keeping with the findings in Shiller (1997), Andre et al. (2022), and Kamdar and Ray (2024). The second-largest group report that inflation makes no difference. Combined with Empirical Result 1, this suggests that the covariance of information on inflation and cross-learning from inflation to the strength of the economy is *negative*. If households consume more when they believe the economy is strong, the narrative heterogeneity channel will therefore reduce the consumption response to inflationary shocks in partial equilibrium.

The relatively long time series of the IAS also allows us to see that the distribution of answers varies substantially over time, which gives us our second empirical result.

[21] See e.g. Roth and Wohlfart (2020) for evidence that households perceive the state of the aggregate economy as relevant to their own decisions.

Empirical Result 2 *A greater proportion of households believe inflation weakens the economy when realized inflation is high.*

The correlation between annual CPI inflation and the proportion of respondents with negative models of inflation is extremely high, at 0.8.[22] The dashed line in Figure 3 plots the predicted values from regressing this proportion on CPI inflation, showing that this correlation is strong across the whole sample. Tests in Appendix D.4 show that the correlation is robust to the addition of various macroeconomic controls, which themselves explain far less of the variation in the distribution of responses than realized inflation. The correlations are also robust to using inflation measures split by various household characteristics. Finally, the proportions giving all other answers are also shown to be significantly negatively correlated with current inflation.

A key advantage of the temporary-information equilibrium approach is that we do not need to take a stand on the source of this time-series variation in order to assess how it affects macroeconomic dynamics. In Section 6 below, I take the evolving distribution of subjective models from this data as given, and use that as an input into the quantification of the causal effects of inflation expectations.

5.4 Inflation perceptions, expectations, and subjective models

Finally, I compare perceived and expected inflation across households with different subjective model beliefs. Figure 4 shows the time series of mean perceived and expected inflation by qualitative subjective model of inflation, which gives us our third empirical result.

Empirical Result 3 *Households who believe inflation weakens the economy on average perceive higher current inflation, and expect higher future inflation, than those who believe inflation makes no difference to the economy. They, in turn, perceive and expect higher inflation than those with positive subjective models of inflation.*

The differences are large: Table 2 shows that even after controlling for the full set of available household characteristics, those with a negative model of inflation perceive that inflation has been 54 basis points higher than those with a neutral model, and 70 basis points higher than those with a positive model.[23] The gaps are similarly large and strongly significant for expected inflation.[24]

[22] Although I do not use later data due to the structural break induced by the methodology change in 2020, this pattern is robust through the 2021-22 inflation surge. Using data from 2020Q1-2024Q3, the equivalent correlation is 0.65.

[23] If anything, these gaps are larger in the more recent waves of the survey omitted from the main analysis. At the peak of the UK's inflation in 2022Q4, the mean difference in inflation perceptions between those with negative and neutral models of inflation was 374 basis points. For 1-year ahead inflation expectations it was 280 basis points.

[24] Dräger et al. (2022) similarly find for German households that inflation expectations are higher among those reporting that they would prefer inflation to be lower.

Figure 4: Inflation perceptions and expectations by subjective model.

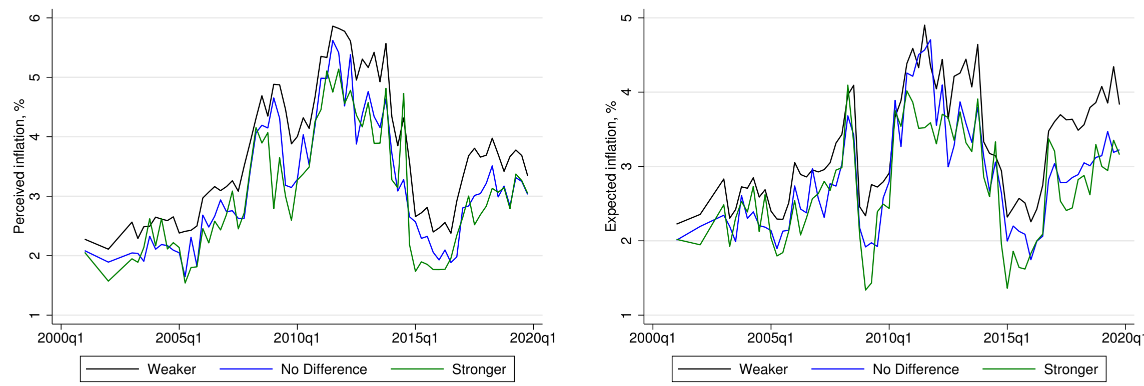

(a) Perception, past 12 months: $\mathbb{E}_t \pi_{t,t-12}$

(b) Expectation, next 12 months: $\mathbb{E}_t \pi_{t+12,t}$

Note: Perceived inflation refers to beliefs about what inflation has been over the past 12 months, and expected inflation refers to expectations for the next 12 months. Averages for each variable are calculated using the survey weights provided in the IAS. Average perceptions and expectations among respondents answering 'Don't know' to the subjective model question (Question 1) are omitted for figure clarity.

Appendix D.5 shows that these results are not driven by selection bias from missing observations for inflation perceptions and expectations.

This is not driven by the households using different kinds of information to arrive at their perceptions and expectations: Table 1 shows that the households with positive subjective models use similar information sources to those with negative models. It is, however, consistent with information about high inflation causing households to update their subjective models towards more negative views. Although the exercises here do not identify the direction of causation, such a mechanism can parsimoniously account for Results 2 and 3. Within a period, those who receive signals that inflation is high shift to more negative subjective models, and when realized inflation rises more households receive such signals.

6 Evaluating the causal effects of household inflation expectations

Having documented direct evidence of a time-varying narrative heterogeneity channel in household beliefs around inflation, I now interpret that evidence through the lens of a New Keynesian model.

Table 2: Perceived and expected inflation are higher for those with more negative subjective models.

	(1) Perceived inflation	(2) Expected inflation
end up stronger	-0.696***	-0.565***
	(0.0371)	(0.0353)
make little difference	-0.543***	-0.466***
	(0.0226)	(0.0207)
dont know	-0.462***	-0.413***
	(0.0315)	(0.0294)
Controls	Yes	Yes
Time FE	Yes	Yes
R-squared	0.179	0.113
Observations	85803	85201

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the results of regressing perceived and expected inflation on respondent subjective models (responses to Question 1) . The omitted category is the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS.

6.1 Households

There is a representative household who chooses consumption C_t, labor supply N_t, and real one-period bond holdings B_t to maximize

$$\bar{E}_t U_t = \bar{E}_t \sum_{s=t}^{\infty} \beta^s \frac{(C_s)^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \varphi \frac{N_s^{1+\psi}}{1+\psi} + \xi \frac{B_t^{1-\eta}}{1-\eta} \tag{38}$$

subject to:

$$C_t + B_t = \frac{R_{t-1}}{\Pi_t} B_{t-1} + W_t N_t + D_t - T_t \tag{39}$$

where R_t is the gross nominal interest rate, Π_t is the per-period rate of inflation,[25], W_t is the real wage, D_t is profits transferred from firms, and T_t are lump-sum taxes.

There are two differences to the household problem in standard New Keynesian models. First, I include bonds in the utility function, as in Hagedorn (2018) and Michaillat and Saez (2019). This allows the model to generate plausible marginal propensities to consume out of current and expected income (Kaplan and Violante, 2018; Auclert et al., 2024), which is important here as the key channel operates through the effect of inflation on income expectations and thus demand.

[25] Note the change from sections 2 and 4, where π_{t+s} was the cumulative inflation between period $t-1$ and $t+s$. I make the change here to be consistent with standard New Keynesian models.

The second departure from standard models is the expectations operator $\bar{E}_t$. Rather than having the household form expectations directly, the model features a continuum of 'agents' within the household. Each agent i forms their own idiosyncratic expectations of all variables, using potentially heterogeneous signals and cross-learning, and the household aggregates these before deciding C_t, N_t. This simplification allows the model to isolate the role of heterogeneous expectations through household actions, without generating further effects from heterogeneous expectations causing heterogeneities in wealth.[26]

In the notation of Section 3, the agent actions $\boldsymbol{x}_{it}$ are to choose a value for the expectation of each variable in the system. The household consumption choice, after log-linearization, then depends on the aggregate of these $\boldsymbol{x}_{it}$ through the following equation.

Proposition 9 *Normalize steady state output and bond holdings to $Y = B = 1$. The log-linear consumption function is*

$$c_t = \mu b_{t-1} + \zeta R\Lambda(r_{t-1} - \bar{E}_t\pi_t) + \Lambda \sum_{s=0}^{\infty} \Lambda^s \left(\zeta \bar{E}_t y_{t+s} - (\sigma - \zeta R\Lambda)(\bar{E}_t r_{t+s} - \bar{E}_t \pi_{t+s+1})\right) \quad (40)$$

where lower-case letters denote log-deviations of corresponding capitalized variables from their steady states, y_t is output, R is the steady state nominal interest rate, and μ, ζ, Λ are coefficients defined in Appendix E.1.

Proof. Appendix E.1. ■

This fits into the form of equation (28). Equation (40) is similar to the standard form in models that do not impose rational expectations (e.g. Eusepi and Preston, 2018), though notably the responses to future income and real interest rates differ from the standard form due to the bonds in utility function. Labor supply is then determined by a standard intratemporal labor supply condition, presented (along with all other model equations) in Appendix E.2.

6.2 Firms, policy, and market clearing

As the focus of this model is the behavior of households, I keep the rest of the model simple. Full details and derivations are in Appendix E.2.

[26] This is analogous to how McKay and Wolf (2023) and Debortoli and Galí (2025) in the HANK literature use labor-aggregating unions to abstract from effects of wealth on labor supply, to focus on core demand effects. See Broer et al. (2021) and Macaulay and Moberly (2025) for evidence that the correlation between wealth and inflation expectations is small, especially outside the hand-to-mouth region.

Monopolistically competitive intermediate goods producers set prices subject to quadratic adjustment costs, and produce using labor as the only input to a linear production function. They supply a perfectly competitive final goods producer, who combines the intermediate goods varieties with a CES production function. Log-linearizing the solution to the intermediate goods firm pricing problem about the zero-inflation steady state yields the Phillips Curve. As in e.g. Bilbiie (2024),[27] I make the simplification that firms ignore future inflation, and thus the Phillips curve is static, in order to focus on the core demand mechanisms of interest here:

$$\pi_t = \kappa y_t + \nu_{\pi t} \tag{41}$$

where $\nu_{\pi t}$ is an AR(1) markup shock with Gaussian innovations.

Monetary policy is set according to a standard Taylor rule:

$$r_t = \phi_\pi \pi_t + \phi_y y_t + \nu_{rt} \tag{42}$$

where ν_{rt} is an AR(1) monetary policy shock with Gaussian innovations.

The fiscal authority issues a constant supply of real bonds $B = 1$, and levies lump-sum taxes to pay the interest on those bonds. Goods market clearing implies

$$y_t = c_t \tag{43}$$

6.3 Agent expectations.

In principle, the temporary-information equilibrium approach can be applied for a broad range of assumptions on agent information and cross-learning. However, in this case I choose a particularly simple structure, that allows me to isolate the role of the correlations documented in Section 5 without introducing a large number of free parameters in expectation formation, for which we have no reliable disciplinary information. Documenting other aspects of household information and cross-learning that are abstracted from here would be a fruitful avenue for future research.

Information. Agents observe realized y_t and r_t. However, they only observe a noisy signal about π_t, as is common in this literature (e.g. Coibion and Gorodnichenko, 2015). That signal takes the

[27] As shown in Bilbiie (2024), this can be microfounded by assuming firms pay a Rotemberg adjustment cost relative to last-period's aggregate price level, rather than their own past prices.

form used in equations (15) and (29), which in this context is

$$s_{it}(\pi_t) = \pi_t + \varepsilon_{\pi t} + \tilde{\varepsilon}_{i\pi t} \tag{44}$$

where $\varepsilon_{\pi t}$ and $\tilde{\varepsilon}_{i\pi t}$ are both mean zero Gaussian variables.

Subjective models. Households know the true monetary policy rule (42). They also know the functional form of the relationship between output and inflation (41), but may hold incorrect and heterogeneous beliefs about the coefficient. Thus their expectations of future output and inflation are related by

$$E_{it}y_{t+s} = \chi_{iy\pi}E_{it}\pi_{t+s} \tag{45}$$

where $\chi_{iy\pi}$ is the cross-learning from inflation to output, which I discipline below with the data in Section 5. Finally, agents believe that inflation follows a simple AR(1) process, so that:

$$E_{it}\pi_{t+s} = \rho E_{it}\pi_{t+s-1} \tag{46}$$

where the perceived persistence ρ is common across agents. This is a simplification of the true inflation dynamics in equilibrium, made here to minimize the number of parameters in the subjective models which would otherwise need to be calibrated. It is also consistent with evidence that households typically use simple perceived laws of motion to forecast inflation (e.g. Adam, 2007).

6.4 Temporary-Information Equilibrium

This model fits into the framework of Section 3. As explained above, agent-specific expectations are the choice variables $\boldsymbol{x}_{it}$ in this case. The state variables c_t, y_t, π_t, r_t depend on the aggregate of these expectation choices through equations (40), (41), (42) and (43), which map into equation (28), with shocks $\nu_{\pi t}$ and ν_{rt}.[28] A formal statement of the temporary-information equilibrium of the model is presented in Appendix E.3.

Using the consumption function (40), and substituting in the assumptions on agent information and subjective models described in Section 6.3, then imposing market clearing, we obtain

$$c_t = \frac{\zeta R\Lambda}{1-\Lambda\zeta}(r_{t-1} - \bar{E}_t\pi_t) - \frac{(\sigma - \zeta R\Lambda)\Lambda}{1-\Lambda\zeta}r_t + \int \tilde{\chi}_i E_{it}\pi_t di \tag{47}$$

[28] Since bond supply is constant, in equilibrium $b_t = 0\forall t$, so we do not need to consider b_t as a state variable.

where
$$\tilde{\chi}_i = \frac{\Lambda^2\rho\left((\zeta - (\sigma - \zeta R\Lambda)\phi_y)\chi_{iy\pi} + (\sigma - \zeta R\Lambda)(\rho - \phi_\pi)\right)}{(1-\Lambda\rho)(1-\Lambda\zeta)} \tag{48}$$
captures the pass-through from perceived inflation to c_t. The full derivation is in Appendix E.3.

Given the Gaussian signal structure (44) and AR(1) perceived law of motion for inflation (46), the inflation perception $E_{it}\pi_t$ is formed according to

$$E_{it}\pi_t = (1 - \tau_{i\pi_t})\rho E_{it-1}\pi_{t-1} + \tau_{i\pi_t} s_{it}(\pi_t) \tag{49}$$

where $\tau_{i\pi_t}$ is the Kalman gain of agent i. From these expressions, we can find the causal effects of inflation expectations on y_t, as described in Section 3. I do this in partial equilibrium (holding π_t, r_t, and signals fixed, except for the public noise in $s_{it}(\pi_{t+1})$) and in general equilibrium. The full algebraic expressions for each causal effect are presented in Appendix E.3.

6.5 Calibration

I set most of the parameters to standard values in the New Keynesian literature for a quarterly calibration. The curvature parameters in the utility function are set to $\sigma = \psi = 1$, and (following the suggestion of Kaplan and Violante, 2018) $\eta = 2.5$. Price stickiness is set to target an average price duration of 5 quarters. The Taylor rule parameters are set to $(\phi_\pi, \phi_y) = (1.5, 0.125)$.

Next, I choose the discount factor β and the weight on bonds in utility ξ to jointly hit two targets: a quarterly steady state interest rate of 1%,[29] and an impact marginal propensity to consume of 0.35, which is the average estimate across the empirical studies surveyed by Sokolova (2023).

To calibrate ρ, I estimate an AR(1) regression using quarterly UK CPI inflation from Q2 2001-Q4 2019, as in the survey data sample. Since at the time of writing the UK does not publish seasonally-adjusted CPI data (Dixon and Michail, 2025), I include quarter-of-the-year fixed effects in the regression. This yields a persistence parameter of $\rho = 0.44$. More sophisticated seasonal adjustments give very similar results (see Appendix E.4).

Finally, I calibrate the joint distribution of $\chi_{iy\pi}, \tau_{i\pi_t}$ to the data in Section 5. Since that data partitions information and subjective models into discrete groups, I impose that the population of

[29] Note that with bonds in the utility function β alone does not pin down the steady state R.

agents is divided into three such groups:

$$(\chi_{iy\pi}, \tau_{i\pi_t}) = \begin{cases} (\chi_+, \tau_H) & \text{with proportion } p_{+,t} \\ (0, \tau_L) & \text{with proportion } 1 - p_{+,t} - p_{-,t} \\ (\chi_-, \tau_H) & \text{with proportion } p_{-,t} \end{cases} \tag{50}$$

with $\chi_+ > 0 > \chi_-$, and $\tau_H > \tau_L$. This distribution reflects Empirical Result 1, that those with positive and negative models of inflation's effects on "the strength of the economy" (which I interpret here as output) have similar information precision, while those who believe inflation has no effect on output use less precise information. In Section 6.6, I set the proportions in each group to match the average proportions giving each subjective model answer over the sample. In Section 6.7, I calibrate them period-by-period to the proportions observed in each wave of the survey. In each case, I pool households responding to Question 1 with "make no difference" and "don't know" into the $(\chi_{iy\pi}, \tau_{i\pi_t}) = (0, \tau_L)$ group.

The survey data is not informative about the magnitude of households' perceived relationship between inflation and output. I therefore proceed by assuming that households with χ_+ know the true slope of the Phillips curve, implying $\chi_{iy\pi} = \kappa^{-1}$. I then choose χ_- such that $\tilde{\chi}_i^2$ is equal across households with positive and negative models. This is motivated by models of rational inattention: with standard quadratic objective functions, optimal information precision is proportional to the square of the elasticity of the action to the state (see the example in Appendix B). Thus if the households with χ_+ and χ_- make the same information choices, they should have the same $\tilde{\chi}_i^2$.

Those same rational inattention models typically imply that optimal information precision is 0 for agents who do not believe the state is sufficiently valuable to learn about. I assume that $\chi_{iy\pi} = 0$ falls in this range, so $\tau_L = 0$. I then calibrate τ_H to target an estimate of the average Kalman gain across households obtained from the survey data, in a procedure described in Appendix E.4.

In both exercises below I also compute the causal effects of interest in a quasi-full-information rational-expectations (FIRE) benchmark, in which $\tau_{i\pi_t} = 1$ and all households know the true slope of the Phillips curve, implying $\chi_{iy\pi} = \kappa^{-1}$. Note this is not a true FIRE calculation, as I maintain the AR(1) perceived law of motion for inflation, with ρ as calibrated above.

6.6 The causal effects of inflation expectations in steady state

The first exercise I conduct with the calibrated temporary-information equilibrium model is to find the causal effects of an inflation expectation shock that arrives when the economy is in steady state. I calibrate the steady states of p_+ and p_- to the average proportions giving positive and negative answers to question 1 in the survey data. Using these probabilities, and applying Propositions 6 and 7, I find the causal effects of inflation expectations on output. The results are displayed in Table 3.

Table 3: The causal effect of aggregate inflation perceptions $\bar{E}_t\pi_t$ on output y_t: total and decomposition.

	Total	RA channel	NH channel	FIRE
Partial equilibrium	-1.286	-1.112	-0.174	0.428
General equilibrium	-1.075	-0.930	-0.145	0.358

Note: The table reports the causal effect of aggregate inflation perceptions on output, in both partial and general equilibrium. The total effect is decomposed into the representative-agent (RA) channel and the narrative heterogeneity (NH) channel, as described in Appendix E.3. The FIRE column reports a benchmark in which $\tau_{i\pi_t} = 1$ and $\chi_{iy\pi} = \kappa^{-1}$ for all i. The calibration is described in Section 6.5 and Appendix E.4.

Greater inflation expectations therefore reduce aggregate output, contrary to the intuition popularly put forward for models with rational expectations (see e.g. Eggertsson and Woodford, 2003; Adam and Billi, 2007), and the findings in the FIRE benchmark. This is because the majority of households believe inflation weakens the economy in the survey, in line with e.g. Hajdini et al. (2023). Since all variables here are in log-deviations, a 1% increase in $\bar{E}_t\pi_t$ causes a 1.1% fall in output in general equilibrium. As well as having the opposite sign, this is just over triple the magnitude of the causal effect in the FIRE benchmark.

As observed in the IAS data, there is a negative correlation between information τ_i and χ_i, so the narrative heterogeneity channel is also negative. Quantitatively, it amplifies the steady state GE effects of inflation expectations on output by 16%. The general equilibrium feedback is also critical to the magnitudes here: since an increase in expectations is contractionary, it causes the monetary policymaker to cut nominal interest rates, which weakens the overall effect substantially.

6.7 The causal effects of inflation expectations over time

I now feed in to the model the time series of $p_{+,t}, p_{-,t}$ observed in the IAS survey. As shown in Figure 3, p_{-t} rises and p_{+t} falls when realized inflation is high.

Figure 5 shows the resulting causal effect of $\bar{E}_t\pi_t$ on y_t in general equilibrium. The solid black line is the total effect. The dashed blue line is the representative-agent channels only.

From this, we see that the causal effect of inflation expectations varies over time. When annualized inflation was 0.1% in 2015 Q1, the causal effect was -0.93. In 2009 Q2, when annualized inflation peaked at 4.8%, it was -1.21, more than 30% larger. However, the volatility of the representative-agent channel is greater than that of the overall causal effect. This is because when inflation is high, the vast majority ($> 80\%$) of households agree on $\chi_{iy\pi} = \chi_-$, and thus there is little heterogeneity in expectation formation. The representative-agent channels dominate. When realized inflation falls, subjective models become more heterogeneous, and the narrative heterogeneity

Figure 5: The causal effect of aggregate inflation perceptions $\bar{E}_t \pi_t$ on output y_t over time.

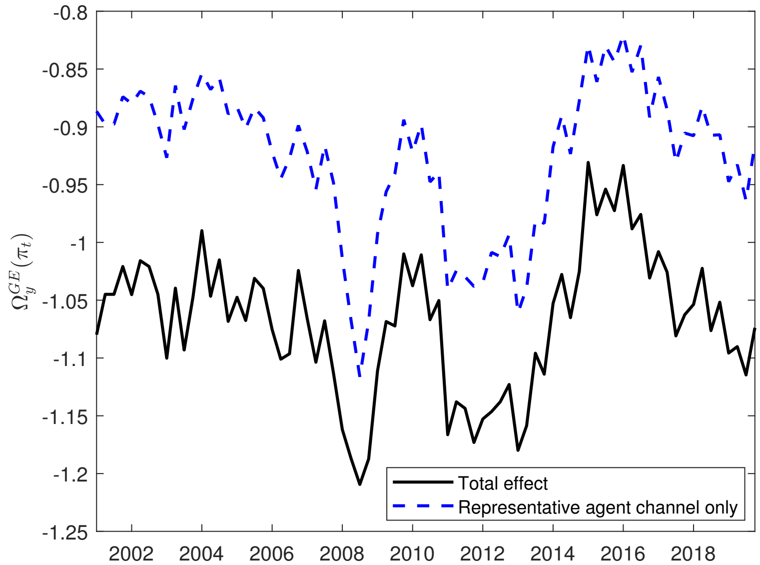

Note: Figure plots the general-equilibrium causal effect of aggregate inflation perceptions on output over time. The solid black line is the total causal effect and the dashed blue line gives the representative-agent channel only. The difference between the two lines reflects the narrative heterogeneity channel. The proportions $p_{+,t}$ and $p_{-,t}$ are calibrated period-by-period to the proportions observed in each wave of the IAS. All other parameters are calibrated as described in Section 6.5 and Appendix E.4.

channel grows. The variance of the total general-equilibrium causal effect is therefore 15% smaller than the variance of the representative-agent effect.

The narrative heterogeneity channel is therefore an important driver of the causal effect of inflation expectations. In fact, the magnitudes reported here are likely to be an underestimate of its effects, because in imposing the discrete-group structure in equation (50) I have removed much of the potential heterogeneity in $\tau_{i\pi_t}$ and $\chi_{iy\pi}$ across households, and thus limited the scope for narrative heterogeneity effects.

7 Conclusion

This paper studies the causal effects of heterogeneously-formed expectations in general equilibrium. To do this I develop a new approach, which I label the *temporary-information equilibrium*. This builds on the core insight from the recent literature on information-provision survey experiments in economics, that is that one can study the causal effect of expectations by creating exogenous variation in a part of agent information sets. I take this idea from the empirical literature and apply it to theoretical macroeconomic modeling. Specifically, this entails an extension of temporary equilibrium, where instead of taking all expectations as fixed and exogenous, we take only a noise term in agent signals as exogenous. Comparative statics with respect to this noise term yield a model equivalent of the first-stage regressions in information-provision experiments. That is, we create exogenous variation in expectations that can be used to identify their causal effects. This therefore retains the precise causal interpretation of temporary equilibrium analysis, without the requirement to hold all aspects of expectation formation fixed. This is how the method allows for cross-learning and general-equilibrium feedback into expectations, and for heterogeneity in the reception and response to information.

I show how to apply this approach to a general class of linear macroeconomic models, and to the specific cases of a Calvo pricing model, and a New Keynesian model. In doing so, I find that the causal effects of expectations, and by implication the role of expectations in the transmission of fundamental shocks, depend in general on a novel *narrative heterogeneity channel*. This operates whenever information and subjective models covary systematically across agents. Heterogeneous subjective models imply heterogeneous responses to information, so systematic patterns in the distribution of information across agents with different subjective models distort the aggregate response to shocks. Moreover, many common assumptions about expectation formation popular in current literature generate transmission channels of this kind whenever there is non-zero heterogeneity in both information and subjective models.

In the Calvo model, the presence of this channel can have large quantitative and even qualitative effects on the conclusions from temporary-equilibrium analysis that necessarily assumes no expectations heterogeneity. For the New Keynesian model, I use unique features of the Bank of England Inflation Attitudes Survey to show that household subjective models and information about inflation covary systematically with each other, implying that the narrative heterogeneity channel will indeed be present in the causal effects of household inflation expectations. Specifically, when matching the model to the empirical patterns, the narrative heterogeneity channel substantially reduces the effects of inflation expectations on output, and weakens the time-variation in that effect.

Finally, an effective use of temporary equilibrium has been to combine "theory with measured expectations" (Piazzesi and Schneider, 2016; Ludwig et al., 2024). The drawback of this approach is that the measured expectations are informative about the level of expectations, but not how they change in response to shocks, or how they affect individual behavior. Since temporary-information equilibrium is conceptually related to information-provision experiments, it can be used to combine theory with measured expectation *changes*, and expectation *transmission*. I provide one example of this, combining the temporary-information equilibrium of the Calvo model with the RCT results in Baumann et al. (2024). More broadly, this method opens up the possibility of combining information-treatment experiments with macroeconomic theory to answer a wide range of macroeconomic questions.

References

Adam, K. (2007). Experimental Evidence on the Persistence of Output and Inflation. *The Economic Journal*, 117(520):603–636.

Adam, K. and Billi, R. M. (2007). Discretionary monetary policy and the zero lower bound on nominal interest rates. *Journal of Monetary Economics*, 54(3):728–752.

Alaoui, L. and Penta, A. (2016). Endogenous Depth of Reasoning. *The Review of Economic Studies*, 83(4):1297–1333.

Alaoui, L. and Penta, A. (2022). Cost-Benefit Analysis in Reasoning. *Journal of Political Economy*, 130(4):881–925.

Andrade, P., Crump, R. K., Eusepi, S., and Moench, E. (2016). Fundamental disagreement. *Journal of Monetary Economics*, 83:106–128.

Andrade, P., Gaballo, G., Mengus, E., and Mojon, B. (2019). Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics*, 11(3):1–29.

Andre, P., Haaland, I., Roth, C., Wiederholt, M., and Wohlfart, J. (2025). Narratives about the Macroeconomy. *Review of Economic Studies*, forthcoming.

Andre, P., Pizzinelli, C., Roth, C., and Wohlfart, J. (2022). Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample. *Review of Economic Studies*.

Ascari, G., Fasani, S., Grazzini, J., and Rossi, L. (2023). Endogenous uncertainty and the macroeconomic impact of shocks to inflation expectations. *Journal of Monetary Economics*, 140:S48–S63.

Ascari, G. and Rossi, L. (2012). Trend Inflation and Firms Price-Setting: Rotemberg Versus Calvo*. *The Economic Journal*, 122(563):1115–1141.

Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review*, 109(6):2333–67.

Auclert, A., Rognlie, M., and Straub, L. (2024). The Intertemporal Keynesian Cross. *Journal of Political Economy*, 132(12):4068–4121.

Bachmann, R., Topa, G., and Van Der Klaauw, W. (2023). *Handbook of Economic Expectations*. Elsevier.

Baley, I. and Veldkamp, L. (2022). Bayesian Learning. In Bachmann, R., Topa, G., and van der Klaauw, W., editors, *Handbook of Economic Expectations*.

Balla-Elliott, D. (2025). Identifying Causal Effects in Information Provision Experiments.

Bank of England (2020). Inflation Attitudes Survey Methodology and notes - online survey.

Baumann, U., Ferrando, A., Georgarakos, D., Gorodnichenko, Y., and Reinelt, T. (2024). SAFE to Update Inflation Expectations? New Survey Evidence on Euro Area Firms. *NBER Working Paper*, 32504.

Beutel, J. and Weber, M. (2024). Beliefs and Portfolios: Causal Evidence. *Working Paper*.

Bhandari, A., Borovička, J., and Ho, P. (2024). Survey Data and Subjective Beliefs in Business Cycle Models. *The Review of Economic Studies*.

Bianchi, F., Ilut, C., and Saijo, H. (2024). Diagnostic Business Cycles. *Review of Economic Studies*, 91(1).

Bilbiie, F. O. (2019). The New Keynesian cross. *Journal of Monetary Economics*, 114:90–108.

Bilbiie, F. O. (2024). Monetary Policy and Heterogeneity: An Analytical Framework. *Review of Economic Studies*, forthcoming.

Blanchard, O. J. and Kahn, C. M. (1980). The Solution of Linear Difference Models under Rational Expectations. *Econometrica*, 48(5):1305–1311.

Bordalo, P., Conlon, J. J., Gennaioli, N., Kwon, S. Y., and Shleifer, A. (2023). Memory and Probability. *The Quarterly Journal of Economics*, 138(1):265–311.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in Macroeconomic Expectations. *American Economic Review*, 110(9):2748–2782.

Branch, W. A. and Evans, G. W. (2006). Intrinsic heterogeneity in expectation formation. *Journal of Economic Theory*, 127(1):264–295.

Broer, T., Kohlhas, A., Mitman, K., and Schlafmann, K. (2020). Heterogenous Information Choice in General Equilibrium.

Broer, T., Kohlhas, A., Mitman, K., and Schlafmann, K. (2021). Information and Wealth Heterogeneity in the Macroeconomy. *CEPR Discussion Papers*, 15934.

Bullard, J. and Suda, J. (2016). The stability of macroeconomic systems with Bayesian learners. *Journal of Economic Dynamics and Control*, 62:1–16.

Candia, B., Coibion, O., and Gorodnichenko, Y. (2022). The Macroeconomic Expectations of Firms. *Handbook of Economic Expectations*, 30042.

Carvalho, V. M. and Tahbaz-Salehi, A. (2019). Production Networks: A Primer. *Annual Review of Economics*, 11:635–663.

Chahrour, R. and Gaballo, G. (2021). Learning from House Prices: Amplification and Business Fluctuations. *The Review of Economic Studies*, 88(4):1720–1759.

Chahrour, R., Nimark, K., and Pitschner, S. (2021). Sectoral Media Focus and Aggregate Fluctuations. *American Economic Review*, 111(12):3872–3922.

Cogley, T., Primiceri, G. E., and Sargent, T. J. (2010). Inflation-Gap Persistence in the US. *American Economic Journal: Macroeconomics*, 2(1):43–69.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678.

D'Acunto, F., Malmendier, U., and Weber, M. (2023). What do the data tell us about inflation expectations? In Bachmann, R., Topa, G., and van der Klaauw, W., editors, *Handbook of Economic Expectations*, pages 133–161. Academic Press.

Dawber, J., Würz, N., Smith, P. A., Flower, T., Thomas, H., Schmid, T., and Tzavidis, N. (2022). Experimental UK Regional Consumer Price Inflation with Model-Based Expenditure Weights. *Journal of Official Statistics*, 38(1):213–237.

Debortoli, D. and Galí, J. (2025). Heterogeneity and Aggregate Fluctuations: Insights from TANK Models. *NBER Macroeconomics Annual*, 39(1):307–357.

Dixon, H. and Michail, M. G. (2025). Seasonal Adjustment of CPI and CPIH. *NIESR Policy Paper*, 45.

Donaldson, P. (2025). Cross-sectional Identification with Heterogeneous Exposure to General Equilibrium Effects. *Working Paper*.

Dräger, L., Lamla, M. J., and Pfajfar, D. (2022). The Hidden Heterogeneity of Inflation Expectations and its Implications. *CESifo Working Paper*, (9637).

Eggertsson, G. B. and Woodford, M. (2003). The Zero Bound on Interest Rates and Optimal Monetary Policy. *Brookings Papers on Economic Activity*, 1.

Eliaz, K. and Spiegler, R. (2024). News Media as Suppliers of Narratives (and Information). *Working Paper*.

Eusepi, S. and Preston, B. (2018). The science of monetary policy: An imperfect knowledge perspective. *Journal of Economic Literature*, 56(1):3–59.

Evans, G. W. and McGough, B. (2020). Adaptive Learning in Macroeconomics. *Oxford Research Encyclopedia of Economics and Finance*.

Farboodi, M. and Veldkamp, L. (2021). A Model of the Data Economy. *NBER Working Papers*, 28427.

Farhi, E. and Werning, I. (2019). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review*, 109(11):3887–3928.

Fuster, A. and Zafar, B. (2023). Survey experiments on economic expectations. In Bachmann, R., Topa, G., and van der Klaauw, W., editors, *Handbook of Economic Expectations*, pages 107–130. Academic Press.

Gabaix, X. (2020). A behavioral new Keynesian model. *American Economic Review*, 110(8):2271–2327.

Gaballo, G. (2018). Price Dispersion, Private Uncertainty, and Endogenous Nominal Rigidities. *The Review of Economic Studies*, 85(2):1070–1110.

Galí, J. (2008). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, Princeton, USA and Oxford, UK.

Gennaioli, N., Leva, M., Schoenle, R., and Shleifer, A. (2024). How Inflation Expectations De-Anchor: The Role of Selective Memory Cues. *NBER Working Papers*, 32633.

Gibbons, R. and Prusak, L. (2020). Knowledge, Stories, and Culture in Organizations. *AEA Papers and Proceedings*, 110:187–92.

Grandmont, J. M. (1977). Temporary General Equilibrium Theory. *Econometrica*, 45(3):535.

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

Guerreiro, J. (2023). Belief Disagreement and Business Cycles. *Working Paper*.

Guren, A., McKay, A., Nakamura, E., and Steinsson, J. (2020). What Do We Learn from Cross-Regional Empirical Estimates in Macroeconomics?: : Vol 35. *NBER Macroeconomics Annual*, 35.

Haaland, I., Roth, C., and Wohlfart, J. (2023). Designing Information Provision Experiments. *Journal of Economic Literature*, 61(1):3–40.

Hagedorn, M. (2018). Prices and Inflation when Government Bonds are Net Wealth. *CEPR Discussion Papers*.

Hajdini, I., Knotek, E. S., Leer, J., Pedemonte, M., Rich, R. W., and Schoenle, R. S. (2023). Low Passthrough from Inflation Expectations to Income Growth Expectations: Why People Dislike Inflation. *Federal Reserve Bank of Cleveland Working Paper*, (22-21R).

Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. *Econometrica*, 47(1):153–161.

Hoffmann, M., Moench, E., Pavlova, L., and Schultefrankenfeld, G. (2022). Would households understand average inflation targeting? *Journal of Monetary Economics*, 129:S52–S66.

Iovino, L. and Sergeyev, D. (2023). Central Bank Balance Sheet Policies without Rational Expectations. *Review of Economic Studies*, 90(6):3119–3152.

Kamdar, R. and Ray, W. (2024). Attention-Driven Sentiment and the Business Cycle. *Working Paper*.

Kanwal, R. P. (1996). *Linear Integral Equations: Theory and Technique*. Birkhäuser Boston.

Kaplan, G. and Schulhofer-Wohl, S. (2017). Inflation at the household level. *Journal of Monetary Economics*, 91:19–38.

Kaplan, G. and Violante, G. L. (2018). Microeconomic Heterogeneity and Macroeconomic Shocks. *NBER Working Papers*, 24734.

Kiss, R. and Strasser, G. (2024). Inflation Heterogeneity Across Households. *ECB Working Paper Series*, 2898.

Leombroni, M., Piazzesi, M., Rogers, C., and Schneider, M. (2020). Inflation and the Price of Real Assets. *Working Paper, Stanford University*.

Link, S., Peichl, A., Pfäuti, O., Roth, C., and Wohlfart, J. (2025). Attention to the Macroeconomy. *ECONtribute Discussion Papers*, 256.

Link, S., Peichl, A., Roth, C., and Wohlfart, J. (2023). Information Frictions Among Firms and Households. *Journal of Monetary Economics*, 135:99–115.

Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–124.

Ludwig, A., Mankart, J., Quintana, J., and Wiederholt, M. (2024). Heterogeneity in Expectations and House Price Dynamics. *SSRN Electronic Journal*.

Macaulay, A. (2021). The Attention Trap: Rational Inattention, Inequality, and Fiscal Policy. *European Economic Review*, 135:103716.

Macaulay, A. (2022). Heterogeneous Information, Subjective Model Beliefs, and the Time-Varying Transmission of Shocks. *CESifo Working Paper*, 9733.

Macaulay, A. and Moberly, J. (2025). Heterogeneity in Imperfect Inflation Expectations: Theory and Evidence from a Novel Survey. *Journal of Money, Credit and Banking*.

Macaulay, A. and Song, W. (2022). Narrative-Driven Fluctuations in Sentiment: Evidence Linking Traditional and Social Media. *University of Oxford Department of Economics Discussion Paper Series*, 973.

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2023). Rational Inattention: A Review. *Journal of Economic Literature*, 61(1):226–73.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99(3):769–803.

Maćkowiak, B. and Wiederholt, M. (2015). Business Cycle Dynamics under Rational Inattention. *Review of Economic Studies*, 82(4):1502–1532.

Maćkowiak, B. and Wiederholt, M. (2024). Rational Inattention during an RCT. *American Economic Review: Insights*.

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):1295–1328.

Mankiw, N. G., Reis, R., and Wolfers, J. (2004). Disagreement about Inflation Expectations. *NBER Macroeconomics Annual*, 18:209–270.

Matthes, C., Nagasaka, N., and Schwartzman, F. (2025). Estimating the Missing Intercept. *Working Paper*.

McKay, A. and Wolf, C. K. (2023). Optimal Policy Rules in HANK. *Working Paper*.

Michaillat, P. and Saez, E. (2019). Resolving New Keynesian Anomalies with Wealth in the Utility Function. *The Review of Economics and Statistics*, pages 1–46.

Michelacci, C. and Paciello, L. (2024). Ambiguity Aversion and Heterogeneity in Households' Beliefs. *American Economic Journal: Macroeconomics*, 16(2):95–126.

Milani, F. (2023). Chapter 18 - Expectational data in DSGE models. In Bachmann, R., Topa, G., and van der Klaauw, W., editors, *Handbook of Economic Expectations*, pages 541–567. Academic Press.

Nakamura, E. and Steinsson, J. (2014). Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *American Economic Review*, 104(3):753–792.

Nakamura, E. and Steinsson, J. (2018). Identification in Macroeconomics. *Journal of Economic Perspectives*, 32(3):59–86.

Nimark, K. P. (2014). Man-Bites-Dog Business Cycles. *American Economic Review*, 104(8):2320–67.

Nimark, K. P. and Pitschner, S. (2019). News media and delegated information choice. *Journal of Economic Theory*, 181:160–196.

Nord, L. (2022). Who Cares about Inflation? Endogenous Expectation Formation of Heterogeneous Households. *Working Paper*.

Nunes, R. and Park, D. (2020). Inflation Expectations, Interest Rates, and Consumption Behavior.

Pfajfar, D. and Santoro, E. (2010). Heterogeneity, learning and information stickiness in inflation expectations. *Journal of Economic Behavior & Organization*, 75(3):426–444.

Piazzesi, M. and Schneider, M. (2016). Chapter 19 - Housing and Macroeconomics. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 1547–1640. Elsevier.

Reis, R. (2006). Inattentive Producers. *The Review of Economic Studies*, 73(3):793–821.

Rotemberg, J. J. (1982). Sticky Prices in the United States. *Journal of Political Economy*, 90(6):1187–1211.

Roth, C. and Wohlfart, J. (2020). How do expectations about the Macroeconomy affect personal expectations and Behavior? *Review of Economics and Statistics*, 102(4):731–748.

Shiller, R. J. (1997). Why Do People Dislike Inflation? In Romer, C. D. and Romer, D. H., editors, *Reducing Inflation: Motivation and Strategy*, pages 13–70. University of Chicago Press.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Sokolova, A. (2023). Marginal propensity to consume and unemployment: A meta-analysis. *Review of Economic Dynamics*, 51:813–846.

Spiegler, R. (2020). Behavioral implications of causal misperceptions. *Annual Review of Economics*, 12:81–106.

Spiegler, R. (2021). A Simple Model of Monetary Policy under Phillips-Curve Causal Disagreements. *Working Paper*.

Weber, M., D'Acunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications. *Journal of Economic Perspectives*, 36(3):157–184.

Werning, I. (2022). Expectations and the Rate of Inflation. *NBER Working Papers*, 30260.

Wolf, C. K. (2023). The Missing Intercept: A Demand Equivalence Approach. *American Economic Review*, 113(8):2232–2269.

Woodford, M. (2013). Macroeconomic Analysis Without the Rational Expectations Hypothesis. *Annual Review of Economics*, 5:303–346.

A Calvo model proofs

Proposition 1. Differentiating equation (11) with respect to the information shock yields

$$\frac{d\pi_t}{d\varepsilon_{\pi_t}} = (1-\lambda)(1-\beta\lambda)\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\chi_{i\pi_{t+s}\pi_t}\tau_{i\pi_t} + \chi_{imc_{t+s}\pi_t}\tau_{i\pi_t}\right) di \tag{51}$$

Aggregating equation (26) across firms and then differentiating yields

$$\frac{d\bar{E}_t(\pi_t)}{d\varepsilon_{\pi_t}} = \int_i \tau_{i\pi_t} di \equiv \bar{\tau}_{\pi_t} \tag{52}$$

where we have used the fact that the diagonal elements of $\boldsymbol{\chi}_i$ are equal to 1 by construction.

Since the model is linear, scaling the shock by $\bar{\tau}_{\pi_t}^{-1}$ implies a unit change in $\bar{E}_t(\pi_t)$, which then generates equation (12).

Proposition 2. The proof follows the same steps as that for Proposition 1. Differentiating equation (11) with respect to the information shock yields

$$\frac{d\pi_t}{d\varepsilon_{\pi_{t+1}}} = (1-\lambda)(1-\beta\lambda)\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\chi_{i\pi_{t+s}\pi_{t+1}}\tau_{i\pi_{t+1}} + \chi_{imc_{t+s}\pi_{t+1}}\tau_{i\pi_{t+1}}\right) di \tag{53}$$

Aggregating equation (9) across firms and then differentiating yields

$$\frac{d\bar{E}_t(\pi_{t+1})}{d\varepsilon_{\pi_{t+1}}} = \int_i \tau_{i\pi_{t+1}} di \equiv \bar{\tau}_{\pi_{t+1}} \tag{54}$$

Since the model is linear, scaling the shock by $\bar{\tau}_{\pi_{t+1}}^{-1}$ implies a unit change in $\bar{E}_t(\pi_{t+1})$, which then generates equation (13).

GE effects of $\varepsilon_{\pi t}$. With the assumptions on signals in equation (15), equation (11) can be written

$$\pi_t = (1-\lambda)(1-\beta\lambda)\left(\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\chi_{i\pi_{t+s}\pi_t} + \chi_{imc_{t+s}\pi_t}\right)\tau_{i\pi_t} di\right)(\pi_t + \varepsilon_{\pi_t}) + a \tag{55}$$

where a collects the steady-state markup μ and all terms involving signals of mc_{t+s}, and π_{t+s} for $s \geq 1$. Differentiating with respect to $\varepsilon_{s\pi t}$ yields

$$\frac{d\pi_t}{d\varepsilon_{\pi_t}} = (1-\lambda)(1-\beta\lambda)\left(\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\chi_{i\pi_{t+s}\pi_t} + \chi_{imc_{t+s}\pi_t}\right)\tau_{i\pi_t} di\right)\left(\frac{d\pi_t}{d\varepsilon_{\pi_t}} + 1\right) \tag{56}$$

which rearranges to

$$\frac{d\pi_t}{d\varepsilon_{\pi_t}} = \frac{\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)}{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)} \tag{57}$$

where

$$\Omega^{PE}(\pi_t) \equiv (1-\lambda)(1-\beta\lambda)\left(\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\chi_{i\pi_{t+s}\pi_t} + \chi_{imc_{t+s}\pi_t}\right)\frac{\tau_{i\pi_t}}{\bar{\tau}_{\pi_t}} di\right) \tag{58}$$

denotes the (scaled) partial-equilibrium effect of π_t expectations, as in equation (12).

Aggregating equation (26) across firms and then differentiating yields

$$\begin{aligned}
\frac{d\bar{E}_t(\pi_t)}{d\varepsilon_{\pi_t}} &= \bar{\tau}_{\pi_t}\left(\frac{d\pi_t}{d\varepsilon_{\pi_t}} + 1\right) && (59)\\
&= \bar{\tau}_{\pi_t}\left(\frac{\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)}{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)} + 1\right) && (60)\\
&= \frac{\bar{\tau}_{\pi_t}}{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)} && (61)
\end{aligned}$$

where the second equality uses equation (57). To ensure the impact on $\bar{E}_t(\pi_t) = 1$, we therefore scale the shock by $\frac{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)}{\bar{\tau}_{\pi_t}}$. This implies

$$\Omega^{GE}(\pi_t) = \frac{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)}{\bar{\tau}_{\pi_t}}\frac{d\pi_t}{d\varepsilon_{\pi_t}} = \Omega^{PE}(\pi_t) \tag{62}$$

After scaling, general equilibrium effects of expected current inflation are irrelevant.

Proposition 3. The proof follows the same steps as for the GE effects of expected π_t above. Equation (11) can be written

$$\pi_t = (1-\lambda)(1-\beta\lambda)\left(\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\chi_{i\pi_{t+s}\pi_t} + \chi_{imc_{t+s}\pi_t}\right)\tau_{i\pi_t}di\right)(\pi_t + \varepsilon_{\pi_t})$$
$$+ (1-\lambda)(1-\beta\lambda)\left(\sum_{s=0}^{\infty}(\beta\lambda)^s \int_i \left(\chi_{i\pi_{t+s}\pi_{t+1}} + \chi_{imc_{t+s}\pi_{t+1}}\right)\tau_{i\pi_{t+1}}di\right)(\pi_{t+1} + \varepsilon_{\pi_{t+1}}) + a' \quad (63)$$

where a' collects the steady-state markup μ and all terms involving signals of mc_{t+s}, and π_{t+s} for $s \geq 2$. Differentiating with respect to $\varepsilon_{s\pi t+1}$ yields

$$\frac{d\pi_t}{d\varepsilon_{\pi_{t+1}}} = \bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)\frac{d\pi_t}{d\varepsilon_{\pi_{t+1}}} + \bar{\tau}_{\pi_{t+1}}\Omega^{PE}(\pi_{t+1}) \quad (64)$$

where we assume that the change in expectations does not imply any change in realized π_{t+1}. In other words, by the time period $t+1$ arrives, the firms realize that the change in $s_{it}(\pi_{t+1})$ they observed was noise, so $E_{it+1}(\pi_{t+1})$ return to steady state. Rearranging, we obtain

$$\frac{d\pi_t}{d\varepsilon_{\pi_{t+1}}} = \frac{\bar{\tau}_{\pi_{t+1}}\Omega^{PE}(\pi_{t+1})}{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)} \quad (65)$$

Aggregating equation (26) across firms and then differentiating yields

$$\frac{d\bar{E}_t(\pi_{t+1})}{d\varepsilon_{\pi_{t+1}}} = \bar{\tau}_{\pi_{t+1}} + \frac{d\pi_t}{d\varepsilon_{\pi_{t+1}}}\int_i \chi_{i\pi_{t+1}\pi_t}\tau_{i\pi_t}di \quad (66)$$
$$= \bar{\tau}_{\pi_{t+1}} + \frac{\bar{\tau}_{\pi_{t+1}}\Omega^{PE}(\pi_{t+1})}{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)}\int_i \chi_{i\pi_{t+1}\pi_t}\tau_{i\pi_t}di \quad (67)$$

This is more complicated than the case of a shock to ε_{π_t} because the general equilibrium feedback into aggregate $\bar{E}_t(\pi_{t+1})$ occurs through cross-learning from changes in realized π_t, and thus in $s_{it}(\pi_t)$. Define the constant K as

$$K \equiv \left(\bar{\tau}_{\pi_{t+1}} + \frac{\bar{\tau}_{\pi_{t+1}}\Omega^{PE}(\pi_{t+1})}{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)}\int_i \chi_{i\pi_{t+1}\pi_t}\tau_{i\pi_t}di\right)^{-1} \quad (68)$$
$$= \frac{1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)}{\bar{\tau}_{\pi_{t+1}}(1-\bar{\tau}_{\pi_t}\Omega^{PE}(\pi_t)) + \bar{\tau}_{\pi_{t+1}}\Omega^{PE}(\pi_{t+1})(\bar{\chi}_{\pi_{t+1}\pi_t}\bar{\tau}_{\pi_t} + Cov_I[\chi_{i\pi_{t+1}\pi_t}, \tau_{i\pi_t}])} \quad (69)$$

The scaled aggregate effect of the shock is therefore

$$K \cdot \frac{d\pi_t}{d\varepsilon_{\pi_{t+1}}} = \frac{\bar{\tau}_{\pi_{t+1}} \Omega^{PE}(\pi_{t+1})}{\bar{\tau}_{\pi_{t+1}}(1 - \bar{\tau}_{\pi_t} \Omega^{PE}(\pi_t)) + \bar{\tau}_{\pi_{t+1}} \Omega^{PE}(\pi_{t+1})(\bar{\chi}_{\pi_{t+1}\pi_t} \bar{\tau}_{\pi_t} + Cov_I[\chi_{i\pi_{t+1}\pi_t}, \tau_{i\pi_t}])} \tag{70}$$

which after dividing through by $\bar{\tau}_{\pi_{t+1}}$ becomes equation (16).

Proposition 4. This follows from equation (20). Using this and the definition of the counterfactual, the fraction of transmission due to expectations is

$$\frac{\frac{d\pi_t}{dmc_t} - \left.\frac{d\pi_t}{dmc_t}\right|_{\sim E}}{\frac{d\pi_t}{dmc_t}} = 1 - \frac{\left.\frac{d\pi_t}{dmc_t}\right|_{\sim E}}{\frac{d\pi_t}{dmc_t}} = 1 - (1 - \bar{\tau}_{\pi_t} \Omega^{PE}(\pi_t)) \tag{71}$$

which implies equation (21).

Proposition 5. The parameters in equation (22) have been chosen so that the representative-agent effects of inflation expectations (line 1 in equation (14)) are equal to those in equation (6). Since $\chi_{imc_{t+s}\pi_t}$ is set to 0, it only remains to find

$$\begin{aligned} Cov_I[\chi_{i\pi_{t+s}\pi_t}, \tau_{i\pi_t}] &= E_I[\chi_{i\pi_{t+s}\pi_t} \tau_{i\pi_t}] - \bar{\chi}_{\pi_{t+s}\pi_t} \bar{\tau}_{\pi_t} && (72) \\ &= E_I[\chi_{i\pi_{t+s}\pi_t} | \tau_{i\pi_t} = 1] \Pr(\tau_{i\pi_t} = 1) - \frac{1+s}{2} && (73) \\ &= \frac{((1+2s)q + 1 - q)}{2} - \frac{1+s}{2} && (74) \\ &= \frac{s}{2}(2q-1) && (75) \end{aligned}$$

where $E_I[\cdot]$ denotes a cross-sectional expectation across firms, and $q \equiv \Pr\left(\chi_{i\pi_{t+s}\pi_t} = 1+2s \middle| \tau_{i\pi_t} = 1\right)$. With this, the narrative heterogeneity effects (line 2 in equation (14)) are

$$\begin{aligned} \frac{(1-\lambda)(1-\beta\lambda)}{\frac{1}{2}} \sum_{s=0}^{\infty} (\beta\lambda)^s Cov_I[\chi_{i\pi_{t+s}\pi_t}, \tau_{i\pi_t}] &= (1-\lambda)(1-\beta\lambda)(2q-1) \sum_{s=0}^{\infty} (\beta\lambda)^s s && (76) \\ &= \frac{\beta\lambda(1-\lambda)(2q-1)}{1-\beta\lambda} && (77) \end{aligned}$$

Adding this together with the representative-agent effects from equation (6) yields

$$\frac{1}{\bar{\tau}_{\pi_t}} \frac{d\pi_t}{d\varepsilon_{\pi_t}} = \frac{1-\lambda}{1-\beta\lambda} + \frac{\beta\lambda(1-\lambda)(2q-1)}{1-\beta\lambda} = \frac{1-\lambda}{1-\beta\lambda} \cdot (1 + \beta\lambda(2q-1)) \tag{78}$$

and taking $\beta \to 1$ then generates equation (23).

Proposition 6. The proof follows the same steps as for Proposition 1. The kth row of equation (27) can be written as

$$x_{ikt} = \sum_{n=0}^{N_z} A_{ikn} x_{int-1} + \sum_{n=0}^{N_z} \sum_{m=0}^{N_z} B_{ikm} \hat{\chi}_{imn} \hat{\tau}_{in} s_{int} \tag{79}$$

where A_{ikn} denotes the k, nth element of A_i, and B_{ikm} denotes the k, mth element of B_i. $\hat{\chi}_{imn}$ and $\hat{\tau}_{in}$ are the m, nth element of $\hat{\boldsymbol{\chi}}_i$ and the nth diagonal element of $\hat{\boldsymbol{\tau}}_i$ respectively.

Aggregating across agents i and differentiating with respect to ε_{jt} we obtain

$$\frac{dx_{kt}}{d\varepsilon_{jt}} = \int_i \left(\sum_{m=0}^{N_z} B_{ikm} \hat{\chi}_{imj} \hat{\tau}_{ij} \right) di \tag{80}$$

$$= \sum_{m=0}^{N_z} \left(\bar{B}_{km} \bar{\chi}_{mj} \bar{\tau}_j + Cov_I \left[B_{ikm}, \chi_{imj} \tau_{ij} \right] + \bar{B}_{km} Cov_I \left[\chi_{imj}, \tau_{ij} \right] \right) \tag{81}$$

where the second equality uses the definition of a covariance.

Similarly, aggregating equation (26) across agents and differentiating yields

$$\frac{d\bar{E}_t(z_{jt})}{d\varepsilon_{jt}} = \int_i \tau_{ij} di \equiv \bar{\tau}_j \tag{82}$$

where we have used the fact that diagonal elements of $\hat{\boldsymbol{\chi}}_i$ are equal to 1 by construction.

Since the model is linear, scaling the shock by $\bar{\tau}_j^{-1}$ implies a unit change in $\bar{E}_t(z_{jt})$, which then implies equation (30).

Proposition 7. Aggregating equation (27) across agents, we can express the aggregate policy function as

$$\boldsymbol{x}_t = \int_i A_i \boldsymbol{x}_{it-1} di + \sum_{s=0}^{\infty} \left[\boldsymbol{\Omega}_{xs}^{PE} \bar{\hat{\boldsymbol{\tau}}}_{xs} (\boldsymbol{x}_{t+s} + \boldsymbol{\varepsilon}_{xt+s}) + \boldsymbol{\Omega}_{ws}^{PE} \bar{\hat{\boldsymbol{\tau}}}_{ws} (\boldsymbol{w}_{t+s} + \boldsymbol{\varepsilon}_{wt+s}) \right] \tag{83}$$

where $\boldsymbol{\Omega}_{xs}^{PE}$ is a $N_x \times N_x$ matrix collecting the partial-equilibrium causal effects of expectations about each element of the future aggregate choice vector $\boldsymbol{x}_{t+s}$ on each aggregate choice (i.e. each element of $\boldsymbol{x}_t$), as defined in Proposition 6. Similarly, $\boldsymbol{\Omega}_{ws}^{PE}$ is a $N_x \times N_w$ matrix collecting the partial-equilibrium causal effects of expectations about each element of the future state vector $\boldsymbol{w}_{t+s}$

on each aggregate choice. $\bar{\hat{\boldsymbol{\tau}}}_{xs}$ and $\bar{\hat{\boldsymbol{\tau}}}_{ws}$ are square diagonal matrices with diagonals equal to the average $\hat{\tau}_{ij}$ parameters across agents for signals on $\boldsymbol{x}_{t+s}$ and $\boldsymbol{w}_{t+s}$ respectively. Finally, $\boldsymbol{\varepsilon}_{xt+s}$ and $\boldsymbol{\varepsilon}_{wt+s}$ are vectors collecting the public component of noise in signals on $\boldsymbol{x}_{t+s}$ and $\boldsymbol{w}_{t+s}$. Note that the idiosyncratic noise terms in equation (29) cancel out when we aggregate across agents.

Rearranging equation (28) we obtain

$$\boldsymbol{w}_t = -D^{-1}C\boldsymbol{x}_t - D^{-1}F\boldsymbol{\xi}_t \tag{84}$$

Substituting (84) in to (83) and rearranging yields

$$\begin{aligned}\boldsymbol{x}_t = (I - \boldsymbol{\Lambda}_0)^{-1}\Bigg(&\int_i A_i\boldsymbol{x}_{it-1}di + \sum_{s=1}^{\infty}\left[\boldsymbol{\Lambda}_s\boldsymbol{x}_{t+s}\right] \\ &+ \sum_{s=0}^{\infty}\left[\boldsymbol{\Omega}_{xs}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{xs}\boldsymbol{\varepsilon}_{xt+s} + \boldsymbol{\Omega}_{ws}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{ws}(\boldsymbol{\varepsilon}_{wt+s} - D^{-1}F\boldsymbol{\xi}_{t+s})\right]\Bigg)\end{aligned} \tag{85}$$

where

$$\boldsymbol{\Lambda}_s \equiv \boldsymbol{\Omega}_{xs}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{xs} - \boldsymbol{\Omega}_{ws}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{ws}D^{-1}C \tag{86}$$

Differentiating with respect to $\boldsymbol{\varepsilon}_{xt+h}$ yields

$$\frac{d\boldsymbol{x}_t}{d\boldsymbol{\varepsilon}_{xt+h}} = (I - \boldsymbol{\Lambda}_0)^{-1}\left(\boldsymbol{\Omega}_{xh}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{xh} + \sum_{s=1}^{\infty}\left[\boldsymbol{\Lambda}_s\frac{d\boldsymbol{x}_{t+s}}{d\boldsymbol{\varepsilon}_{xt+h}}\right]\right) \tag{87}$$

Next, we turn to $\frac{d\boldsymbol{x}_{t+s}}{d\boldsymbol{\varepsilon}_{xt+h}}$. For this, note that the terms relating $\boldsymbol{x}_t$ to $\boldsymbol{x}_{it-1}$ in equation (83) can be written as

$$\int_i A_i\boldsymbol{x}_{it-1}di = \tilde{A}\boldsymbol{x}_{t-1} \tag{88}$$

where

$$\tilde{A} \equiv \int_i A_i di + Cov_I[A_i, \boldsymbol{x}_{it-1}]\boldsymbol{x}_{t-1}^{+} \tag{89}$$

and $\boldsymbol{x}_{t-1}^{+}$ is the Moore-Penrose pseudoinverse of $\boldsymbol{x}_{t-1}$, such that $\boldsymbol{x}_{t-1}^{+}\boldsymbol{x}_{t-1} = 1$.

To proceed, we will make the simplifying assumption that $Cov_I[A_i, \boldsymbol{x}_{it-1}]\boldsymbol{x}_{t-1}^{+}$ is constant over time. That is, the correlation between an agent's weight on lagged actions in their policy rule and their deviation from the aggregate actions may be non-zero, but it is unaffected by shocks. With this simplification, $\tilde{A}$ is constant.

Let the matrix R be defined by the recursion

$$R = (I - \mathbf{\Lambda}_0)^{-1} \left(\tilde{A} + \sum_{k=1}^{\infty} \mathbf{\Lambda}_k R^{k+1} \right) \tag{90}$$

Given this, we can show the following.

Lemma 1 *If the series $\sum_{k=1}^{\infty} \mathbf{\Lambda}_k R^{k+1}$ converges, and if the spectral radius of R is $\rho(R) < 1$, then*

$$\frac{d\boldsymbol{x}_{t+s}}{d\boldsymbol{\varepsilon}_{xt+h}} = R^s \frac{d\boldsymbol{x}_t}{d\boldsymbol{\varepsilon}_{xt+h}} \tag{91}$$

Proof. We prove equation (91) with a guess-and-verify approach. Rolling equation (85) forwards and differentiating, we have

$$\frac{d\boldsymbol{x}_{t+s}}{d\boldsymbol{\varepsilon}_{xt+h}} = (I - \mathbf{\Lambda}_0)^{-1} \left(\tilde{A} \frac{d\boldsymbol{x}_{t+s-1}}{d\boldsymbol{\varepsilon}_{xt+h}} + \sum_{k=1}^{\infty} \left[\mathbf{\Lambda}_k \frac{d\boldsymbol{x}_{t+s+k}}{d\boldsymbol{\varepsilon}_{xt+h}} \right] \right) \tag{92}$$

Substituting in equation (91) this becomes

$$\frac{d\boldsymbol{x}_{t+s}}{d\boldsymbol{\varepsilon}_{xt+h}} = (I - \mathbf{\Lambda}_0)^{-1} \left(\tilde{A} R^{s-1} \frac{d\boldsymbol{x}_t}{d\boldsymbol{\varepsilon}_{xt+h}} + \sum_{k=1}^{\infty} \left[\mathbf{\Lambda}_k R^{s+k} \frac{d\boldsymbol{x}_t}{d\boldsymbol{\varepsilon}_{xt+h}} \right] \right) \tag{93}$$

$$= (I - \mathbf{\Lambda}_0)^{-1} \left(\tilde{A} + \sum_{k=1}^{\infty} \left[\mathbf{\Lambda}_k R^{k+1} \right] \right) R^{s-1} \frac{d\boldsymbol{x}_t}{d\boldsymbol{\varepsilon}_{xt+h}} \tag{94}$$

which when combined with (90) verifies the guess.

The requirement that $\sum_{k=1}^{\infty} \mathbf{\Lambda}_k R^{k+1}$ converges in some consistent matrix norm ensures that R is a well-defined finite matrix. $\rho(R) < 1$ ensures uniqueness of the bounded solution in (91), ruling out explosive alternatives. We assume these conditions hold, in analogy with the determinacy conditions in rational-expectations models (existence and uniqueness of a bounded solution). ■

Substituting equation (91) into equation (87) and rearranging yields

$$\frac{d\boldsymbol{x}_t}{d\boldsymbol{\varepsilon}_{xt+h}} = \left(I - \mathbf{\Lambda}_0 - \sum_{s=1}^{\infty} \left[\mathbf{\Lambda}_s R^s \right] \right)^{-1} \mathbf{\Omega}_{xh}^{PE} \bar{\boldsymbol{\tau}}_{xh} \tag{95}$$

Applying the same steps to a shock to information about state variables yields

$$\frac{d\boldsymbol{x}_t}{d\boldsymbol{\varepsilon}_{wt+h}} = \left(I - \boldsymbol{\Lambda}_0 - \sum_{s=1}^{\infty}[\boldsymbol{\Lambda}_s R^s]\right)^{-1} \boldsymbol{\Omega}_{wh}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{wh} \tag{96}$$

Defining the $N_x \times N_z$ matrix $\boldsymbol{\Omega}^{PE} = (\boldsymbol{\Omega}_{x0}^{PE}, \boldsymbol{\Omega}_{w0}^{PE}, \boldsymbol{\Omega}_{x1}^{PE}, \boldsymbol{\Omega}_{w1}^{PE}, ...)$, we can combine these into one expression:

$$\frac{d\boldsymbol{x}_t}{d\boldsymbol{\varepsilon}_t} = \boldsymbol{\Lambda}^{GE}\boldsymbol{\Omega}^{PE}\bar{\hat{\boldsymbol{\tau}}} \tag{97}$$

where

$$\boldsymbol{\Lambda}^{GE} \equiv \left(I - \boldsymbol{\Lambda}_0 - \sum_{s=1}^{\infty}[\boldsymbol{\Lambda}_s R^s]\right)^{-1} \tag{98}$$

To extract the causal effect of the jth element of $\boldsymbol{\varepsilon}_t$ on the kth aggregate action (x_{kt}), we simply take the k, jth element of the matrix $\frac{d\boldsymbol{x}_t}{d\boldsymbol{\varepsilon}_t}$. Since $\bar{\hat{\boldsymbol{\tau}}}$ is diagonal, this implies

$$\frac{dx_{kt}}{d\varepsilon_{jt}} = \left[\boldsymbol{\Lambda}^{GE}\boldsymbol{\Omega}^{PE}\right]_{k,j}\bar{\hat{\tau}}_j \tag{99}$$

as in equation (33).

All that remains is to scale this so the impact on the relevant aggregate expectation is equal to 1. The jth row of equation (26) is

$$E_{it}(z_{jt}) = \sum_{n=0}^{N_z}\hat{\chi}_{ijn}\hat{\tau}_{in}(z_{nt} + \tilde{\varepsilon}_{int} + \varepsilon_{nt}) \tag{100}$$

Aggregating across agents and differentiating with respect to ε_{jt} we find

$$\frac{\bar{E}_t(z_{jt})}{d\varepsilon_{jt}} = \bar{\hat{\tau}}_j + \Gamma_j \tag{101}$$

where

$$\Gamma_j \equiv \sum_{n=0}^{N_z} \left[\bar{\hat{\chi}}_{jn}\bar{\hat{\tau}}_n + Cov_I[\hat{\chi}_{ijn}, \hat{\tau}_{in}]\right] \frac{dz_{nt}}{d\varepsilon_{jt}} \tag{102}$$

$$= \sum_{n=0}^{N_z} \left[\bar{\hat{\chi}}_{jn}\bar{\hat{\tau}}_n + Cov_I[\hat{\chi}_{ijn}, \hat{\tau}_{in}]\right] \cdot \left[\begin{pmatrix} I \\ -D^{-1}C \end{pmatrix} \mathbf{\Lambda}^{GE}\mathbf{\Omega}^{PE}\right]_{j,n} \bar{\hat{\tau}}_n \tag{103}$$

The second equality here comes from using equations (97) and (84) to obtain $\frac{d\boldsymbol{w}_t}{d\varepsilon_{jt}}$ in terms of $\frac{d\boldsymbol{x}_t}{d\varepsilon_{jt}}$, and using this to substitute out for $\frac{dz_{nt}}{d\varepsilon_{jt}}$.

To ensure a unit change in $\bar{E}_t(z_{jt})$, the un-scaled effect $\frac{dx_{kt}}{d\varepsilon_{jt}}$ must be multiplied by $K = (\bar{\hat{\tau}}_j + \Gamma_j)^{-1}$. Combined with equation (99) this implies equation (34).

Corollary 1 *If there are no signals about any variables realized after the current period, the general equilibrium feedback from an exogenous change in expectations of a contemporaneous variable is determined by the matrix*

$$\mathbf{\Lambda}^{GE} = \left(I - \mathbf{\Omega}_{x0}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{x0} + \mathbf{\Omega}_{w0}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{w0}D^{-1}C\right)^{-1} \tag{104}$$

where $\mathbf{\Omega}_{x0}^{PE}$ and $\mathbf{\Omega}_{w0}^{PE}$ collect the partial-equilibrium effects of each expectation $\bar{E}_t(z_{jt})$ for $z_{jt} \in \{\boldsymbol{x}_t, \boldsymbol{w}_t\}$, and $\bar{\hat{\boldsymbol{\tau}}}_{x0}, \bar{\hat{\boldsymbol{\tau}}}_{w0}$ are matrices collecting the corresponding elements of average $\boldsymbol{\tau}_i$.

Proof. If there is no forward information, then $\bar{\hat{\boldsymbol{\tau}}}_{xs}$ and $\bar{\hat{\boldsymbol{\tau}}}_{ws}$ are zero matrices for all $s \geq 1$. Equation (86) therefore implies $\mathbf{\Lambda}_s$ is a zero matrix for all $s \geq 1$. Equation (98) then becomes

$$\mathbf{\Lambda}^{GE} = (I - \mathbf{\Lambda}_0)^{-1} = \left(I - \mathbf{\Omega}_{x0}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{x0} + \mathbf{\Omega}_{w0}^{PE}\bar{\hat{\boldsymbol{\tau}}}_{w0}D^{-1}C\right)^{-1} \tag{105}$$

where the second equality uses equation (86). ■

From this, we can see that an exogenous change in $\bar{E}_t(z_{jt})$ causes two related general equilibrium feedback effects. First, as in the Calvo model in Section 2, the shock causes a change in aggregate actions $\boldsymbol{x}_t$, which feeds back into the information agents receive.

Second, this change in $\boldsymbol{x}_t$ also affects the state variables $\boldsymbol{w}_t$ through the equilibrium condition (28), which *also* affects agent information. This was not present in Section 2 because we abstracted there from all responses of the state (marginal costs) to firm actions.

Finally, note that outside of the special case in Corollary 1, there is a further category of effects, as changes in $\bar{E}_t(z_{jt})$ affect future actions and states through the lagged-action term $A_i\boldsymbol{x}_{it-1}$ in the policy function (27). If agents get signals about future aggregate actions or states, those signals will also therefore react to the change in $\bar{E}_t(z_{jt})$.

Proposition 8. Differentiating equation (85) with respect to the fundamental shock vector $\boldsymbol{\xi}_t$ yields

$$\frac{d\boldsymbol{x}_t}{d\boldsymbol{\xi}_t} = (I - \boldsymbol{\Lambda}_0)^{-1} \left(-\boldsymbol{\Omega}_{w0}^{PE} \bar{\tilde{\boldsymbol{\tau}}}_{w0} D^{-1} F + \sum_{s=1}^{\infty} \left[\boldsymbol{\Lambda}_s \frac{d\boldsymbol{x}_{t+s}}{d\boldsymbol{\xi}_t} \right] \right) \tag{106}$$

Following analogous reasoning to that in Lemma 1 we can substitute out for $\frac{d\boldsymbol{x}_{t+s}}{d\boldsymbol{\xi}_t} = R^s \frac{d\boldsymbol{x}_t}{d\boldsymbol{\xi}_t}$, where R is defined in (90). Rearranging the resulting expression implies equation (35).

B Narrative heterogeneity in common models of expectation formation

I focus here on the model in Section 2 under the simplifying assumptions that $\tau_{ij} = 0$ for all variables except π_t, and $\varepsilon_{\pi_t} = 0$, $\tilde{\varepsilon}_{i\pi_t} \sim N(0, \sigma_i^2)$, $\pi_t \sim N(0, \sigma_\pi^2)$. The reset price of firm i is then given by

$$p_{it}^* - P_{t-1} = (1 - \beta\lambda) \sum_{s=0}^{\infty} (\beta\lambda)^s [\chi_{i\pi_{t+s}\pi_t} + \chi_{imc_{t+s}\pi_t}] \tau_{i\pi_t} (\pi_t + \tilde{\varepsilon}_{i\pi_t}) \tag{107}$$

It will be helpful here to define $\tilde{\chi}_i$ as the sum of cross-learning parameters as they enter the reset price first order condition:

$$\tilde{\chi}_i \equiv \sum_{s=0}^{\infty} (\beta\lambda)^s [\chi_{i\pi_{t+s}\pi_t} + \chi_{imc_{t+s}\pi_t}] \tag{108}$$

I also assume that firms incorporate their signals into expectations of π_t optimally. Standard Bayesian updating therefore implies

$$\tau_{i\pi_t} = \frac{\sigma_\pi^2}{\sigma_\pi^2 + \sigma_i^2} \tag{109}$$

Rational inattention. Following Sims (2003), models of rational inattention assume that agents choose the precision of their signals, subject to a cost (see Maćkowiak et al., 2023, for a review).

To study this in our Calvo model, we make two assumptions. First, a resetting firm i's expected payoff is proportional to $-E_{it-1}(p_{it}^* - p_{it}^{FI})^2$, where p_{it}^{FI} is the reset price that they would choose if they precisely observed π_t.[30] Second, before the signal $s_{it}(\pi_t)$ is observed, each firm can choose the variance of $\tilde{\varepsilon}_{i\pi_t}$, subject to a cost that is proportional to the expected entropy reduction between their prior and the resulting posterior beliefs about π_t.[31]

Since there is a one-to-one mapping between σ_i^2 and $\tau_{i\pi_t}$ (equation (109)), we can formulate the information-choice problem as

$$\arg\max_{\tau_{i\pi_t}\in[0,1]} -E_{it-1}(p_{it}^* - p_{it}^{FI})^2 - \mu \log\left(\frac{1}{1-\tau_{i\pi_t}}\right) \tag{110}$$

where $\mu > 0$ is the marginal cost of information.

Substituting out for p_{it}^* and p_{it}^{FI} using (107) and evaluating the expectations, the problem becomes

$$\arg\max_{\tau_{i\pi_t}\in[0,1]} -(1-\beta\lambda)\tilde{\chi}_i^2\left((\tau_{i\pi_t}-1)^2\sigma_\pi^2 + \tau_{i\pi_t}^2\sigma_i^2\right) - \mu \log\left(\frac{1}{1-\tau_{i\pi_t}}\right) \tag{111}$$

which after substituting out for σ_i^2 using (109) and differentiating gives the first order condition:

$$\tau_{i\pi_t} = \max\left\{0, 1 - \frac{\mu}{(1-\beta\lambda)\sigma_{\pi_t}^2\tilde{\chi}_i^2}\right\} \tag{112}$$

This clearly generates correlations between $\tau_{i\pi_t}$ and $\chi_{i\pi_{t+s}\pi_t}, \chi_{imc_{t+s}\pi_t}$, as long as some firms choose to process strictly positive amounts of information ($\tau_{i\pi_t} > 0$). Intuitively, firms who believe there is a stronger pass-through from current inflation to future inflation and to marginal costs (larger $\tilde{\chi}_i^2 = [\sum_{s=0}^{\infty}(\beta\lambda)^s[\chi_{i\pi_{t+s}\pi_t} + \chi_{imc_{t+s}\pi_t}]]^2$) find information more valuable, so process more of it – implying a $\tau_{i\pi_t}$ closer to 1.

Learning. Models of rational inattention endogenize $\boldsymbol{\tau}_i$. In contrast, models of learning make subjective models endogenous: agents estimate the coefficients of aggregate laws of motion from their observations (Evans and McGough, 2020). This also generates a narrative heterogeneity effect, because agents learning from systematically different information estimate different laws of motion.

To incorporate this in a simple way into the Calvo framework used here, assume that all firms have perceived laws of motion that specify there is no link between π_t and future π_{t+s} or mc_{t+s} for

[30] An objective of this form can be obtained by taking a log-quadratic approximation to the profit function (Maćkowiak and Wiederholt, 2009).

[31] Since the firm's objective function is quadratic, and priors are Gaussian, the signal structure in equation (15) without public noise is in fact optimal (Maćkowiak and Wiederholt, 2009).

$s \geq 1$, that $\pi_t \sim N(\mu_\pi, \sigma_\pi^2)$, and that current mc_t is perfectly correlated with π_t:

$$\pi_t = \gamma mc_t \tag{113}$$

These assumptions imply that $\tilde{\chi}_i$ is equal to the firm's belief about the parameter γ.

Furthermore, suppose that at the end of period t (i.e. after resetting firms have chosen p_{it}) firms observe the realized mc_t. They can therefore use the combination of mc_t and $s_{it}(\pi_t)$ to learn about the true slope γ.

There are a range of possible learning technologies present in the literature. I proceed here with a tractable application of Bayesian learning (reviewed in Baley and Veldkamp, 2022). Using their observation of mc_t and their inflation signal from period t, firm i can construct an unbiased signal about γ^{-1}:

$$\frac{s_{it}(\pi_t)}{mc_t} = \frac{1}{\gamma} + \frac{\epsilon_{i\pi_t}}{mc_t} \tag{114}$$

To keep the learning problem simple for this example, suppose that firms start period t with a Gaussian prior belief about γ^{-1}, with mean $\mu_{\gamma it}$ and variance $\sigma^2_{\gamma it}$. Learning about the inverse coefficient γ^{-1} is then straightforward: Bayesian updating implies that posterior beliefs about γ^{-1} at the end of period t are:

$$E(\gamma^{-1}|s_{it}(\pi_t), mc_t) = \mu_{\gamma it} + K_{it}\left(\frac{1}{\gamma} + \frac{\epsilon_{i\pi_t}}{mc_t} - \mu_{\gamma it}\right) \tag{115}$$

where the gain K_{it} is given by:

$$K_{it} = \frac{mc_t^2 \sigma^2_{\gamma it}}{mc_t^2 \sigma^2_{\gamma it} + \sigma^2_\pi \left(\frac{1-\tau_{i\pi_t}}{\tau_{i\pi_t}}\right)} \tag{116}$$

This gain is strictly increasing in $\tau_{i\pi_t}$. Firms with more precise information update their subjective models more strongly in response to it. At the extremes, a firm with $\tau_{i\pi_t} = 1$ is able to exactly back out the realized γ, so $K_{it} = 1$ and their posterior belief matches the true correlation between π_t and mc_t. A firm with $\tau_{i\pi_t} \to 0$ has $K_{it} \to 0$: their signal is uninformative, so they never update their subjective model. If period-t posteriors persist to influence the prior beliefs used in pricing decisions in period $t+1$, this difference in learning speeds will create a narrative heterogeneity channel.

We can see the effects of this more concretely if we impose further structure. Assume that in an initial period 0, all firms share the same priors, with mean μ_0 and variance $\sigma^2_{\gamma 0}$. Firms are split equally between $\tau_{i\pi_t} = \{0, 0.5, 1\}$. Finally, assume all firms believe γ is constant, so prior beliefs in period 1 are centered on their period-0 posteriors.

In period 0, resetting firms choose prices before observing mc_t, so use their prior beliefs on γ_0^{-1}.[32] Since there is no heterogeneity in these beliefs, the narrative heterogeneity term is equal to 0.

After mc_0 is realized, firms update their subjective models using equation (115), which yields:

$$E_i(\gamma^{-1}|s_{i0}(\pi_t), mc_0) = \begin{cases} \mu_0 & \text{if } \tau_{i\pi_t} = 0 \\ \mu_0 + \frac{mc_t^2\sigma_{\gamma 0}^2}{mc_t^2\sigma_{\gamma 0}^2+\sigma_\pi^2}\left(\frac{1}{\gamma} + \frac{\epsilon_{i\pi_t}}{mc_0} - \mu_0\right) & \text{if } \tau_{i\pi_t} = 0.5 \\ \gamma^{-1} & \text{if } \tau_{i\pi_t} = 1 \end{cases} \tag{117}$$

In period 1, resetting firms choose their price using $\tilde{\chi}_{i1} = [E_i(\gamma^{-1}|s_{i0}(\pi_t), mc_0)]^{-1}$ as their subjective model. The narrative heterogeneity channel is then:[33]

$$Cov_I[\tilde{\chi}_{i1}, \tau_{i\pi_t}] = \frac{1}{6}(\gamma - \mu_0^{-1}) \tag{118}$$

If prior beliefs in period 0 are below the true γ, then in period 1 the narrative heterogeneity channel is positive. Firms with precise information update their beliefs about γ upwards, while those with imprecise information do not, so more precise information is associated with higher $\tilde{\chi}_{i1}$. If period-0 priors are above the true γ, the narrative heterogeneity channel has the opposite sign. The only way for the narrative heterogeneity channel to disappear is if period-0 priors are centred on γ.

Sticky information. Following Mankiw and Reis (2002), firms observe information only sporadically. With probability $q_i \in (0, 1)$, firm i observes π_t precisely, which through equation (109) implies $\tau_{i\pi_t} = 1$. With probability $1 - q_i$, firm i observes no information at all ($\tau_i = 0$). As in Reis (2006) and others, firms can pay a cost to increase q_i.[34]

Applying the same steps as used to obtain equation (111), expected utility net of the increasing, convex cost of increasing q_i ($C(q_i)$) is:

$$-(1-\beta\lambda)\tilde{\chi}_i^2\sigma_\pi^2(1-q_i) - C(q_i) \tag{119}$$

[32] In all periods, I assume that firms make pricing decisions as if they were certain of their estimate of γ, following the standard 'anticipated utility' assumption in the learning literature (Bullard and Suda, 2016).

[33] Calculating this requires taking an expectation over the inverse of the normally distributed random variable $E_i(\gamma^{-1}|s_{i0}(\pi_t), mc_0, \tau_{i\pi_t} = 0.5)$. Since the expectation of the inverse of a normally distributed variable does not exist in the conventional sense, I use the corresponding Cauchy principal value (Kanwal, 1996). This term cancels out of the resulting expression, but the step is required to ensure that $Cov_I[\tilde{\chi}_{i1}, \tau_{i\pi_t}]$ is well-defined.

[34] Specifically, Reis (2006) assumes firms can pay a cost to observe information, which is equivalent to paying to increase q_i to 1.

The firm chooses q_i to maximize this objective. The optimal q_i satisfies the first order condition:

$$(1-\beta\lambda)\tilde{\chi}_i^2\sigma_\pi^2 = C'(q_i) \tag{120}$$

which implies q_i is strictly increasing in $\tilde{\chi}_i^2 = \left[\sum_{s=0}^{\infty}(\beta\lambda)^s[\chi_{i\pi_{t+s}\pi_t} + \chi_{imc_{t+s}\pi_t}]\right]^2$. The intuition is the same as the rational inattention model: information is more valuable to firms who believe that current inflation has larger effects on future inflation and marginal costs.

By the law of iterated expectations, each narrative heterogeneity covariance can be expressed as:

$$\begin{aligned} Cov(\chi_{ij\pi_t}, \tau_{i\pi_t}) &= E(E(\chi_{ij\pi_t}\tau_{i\pi_t}|q_i)) - \bar{\chi}_{j\pi_t}E(E(\tau_{i\pi_t}|q_i)) && (121)\\ &= E(\chi_{ij\pi_t}q_i) - \bar{\chi}_{j\pi_t}E(q_i) && (122)\\ &= Cov(\chi_{ij\pi_t}, q_i) && (123)\end{aligned}$$

That is, the narrative heterogeneity channel appears whenever the probability of observing information is correlated with a firm's subjective model. Equation (120) implies that this is indeed the case. In the case where $\chi_{ij\pi_t} > 0$ for all i, then $Cov_I(\chi_{ij\pi_t}, q_i) > 0$.

Information delegation to media outlets. As in Nimark (2014), Chahrour et al. (2021), and others, assume that firms receive their signals from media outlets, who perform an editorial role. Media outlet m observes π_t precisely, and chooses whether or not to report it to their readers depending on its realization.

As is common in the literature, I focus on editorial decision rules in which large shocks are deemed more newsworthy. Media outlet m reports on π_t if $|\pi_t| > \theta_m$, where $\theta_m > 0$ is an outlet-specific threshold. If outlet m reports on π_t, any firm that reads outlet m observes π_t precisely ($\tau_{i\pi_t} = 1$). Using the same quadratic objective as in the rational inattention case, the firm has a payoff of 0. If m does not report π_t, readers do not observe any information about π_t ($\tau_{i\pi_t} = 0$). Instead of π_t, the firms read about a non-economic news story, from which they gain utility x. Each firm only reads the news from one outlet. The expected utility from reading outlet m is:

$$E_{imt-1}U_{it} = \Pr(|\pi_t| > \theta_m)\cdot 0 + \Pr(|\pi_t| < \theta_m)\cdot\left(x - (1-\beta\lambda)\tilde{\chi}_i^2 Var(\pi_t \big| |\pi_t| < \theta_m)\right) \tag{124}$$

Importantly, the variance of π_t conditional on π_t not being reported is not equal to the unconditional variance σ_π^2. This is the key insight of Nimark (2014): the fact that π_t is not reported allows firms to rule out extreme realizations of π_t.

As θ_m rises, $\Pr(|\pi_t| < \theta_m)$ rises, but $Var(\pi_t \big| |\pi_t| < \theta_m)$ also rises as fewer extreme realizations can be ruled out. This means that as π_t becomes less likely to be reported, the expected utility costs of it going unreported also rise. These costs are proportional to $\tilde{\chi}_i^2$, so firms with larger cross-learning parameters prefer π_t to be reported more frequently. This is the key mechanism that generates a narrative heterogeneity effect, as firms with greater $\chi_{ij\pi_t}$ parameters demand news with a lower threshold θ_m.

The exact value of the narrative heterogeneity channel depends on the nature of the media market. For simplicity, here I follow Nimark and Pitschner (2019) and assume that there is a continuum of media outlets, one for each firm, and that each outlet maximizes the expected utility of their associated firm. The optimal θ_i for firm i is such that firms are indifferent between reported and unreported states. That is:

$$Var(\pi_t \big| |\pi_t| < \theta_i) = \frac{x}{(1-\beta\lambda)\tilde{\chi}_i^2} \tag{125}$$

Since $Var(\pi_t \big| |\pi_t| < \theta_i)$ is strictly increasing in θ_i, greater $\tilde{\chi}_i^2$ is associated with lower θ_i. For any given realization of π_t, we therefore have $\tau_{i\pi_t} = 1$ if $\tilde{\chi}_i^2$ is above a certain threshold, and $\tau_{i\pi_t} = 0$ if not. This generates a narrative heterogeneity effect, as information is concentrated among those with large cross-learning parameters, as in the rational inattention and sticky information models. Unlike those models, the extent of this covariance is now state-dependent: at π_t close to 0, almost no firms receive news about it, so there is little heterogeneity in information and the covariance is small. At very large (positive or negative) realizations of π_t, almost all firms receive the news, and again there is little heterogeneity. The narrative heterogeneity effect is therefore strongest for intermediate realizations of π_t, when news coverage is more divided.

$$Cov_I[\tilde{\chi}_i, \tau_{i\pi_t}] = [E_I(\tilde{\chi}_i|\tilde{\chi}_i^2 > \chi^*(\pi_t)) - \bar{\tilde{\chi}}] \cdot \Pr(\tilde{\chi}_i^2 > \chi^*(\pi_t)) \tag{126}$$

where $\chi^*(\pi_t)$ is the threshold such that firm i receives news on π_t if $\tilde{\chi}_i^2 > \chi^*(\pi_t)$. The results above imply that $\chi^*(\pi_t)$ is decreasing in $|\pi_t|$. For small $|\pi_t|$, the first component is close to 0. For large $|\pi_t|$, the second component is close to 0.

This analysis abstracts from two further channels through which media provision of information may affect the joint distribution of information and subjective models. First, I considered only news reporting functions which were symmetric about $\pi_t = 0$. This implied that $E_{it}\pi_t = 0$ whenever the firm's media outlet did not report π_t. If there was a bias (e.g. if negative shocks are reported more often than positive shocks, as in Chahrour et al. (2021)), then a firm observing their outlet not reporting π_t will adjust their posterior expectation of π_t, as well as their posterior variance. Two firms may form different expectations even if neither of their outlets report on π_t, if their outlets have different reporting functions. This creates further information heterogeneity between firms based on their media consumption. Second, recent work has suggested that media may also transmit subjective models as well as information (Macaulay and Song, 2022; Eliaz and Spiegler, 2024). While not considered here, it provides a further reason why the media sector may induce correlations between information and subjective models.

Selective recall. Following Bordalo et al. (2023) and Gennaioli et al. (2024), firms form beliefs about the probability of future inflation by recalling previous inflation experiences. Assume that the bank of experiences that can be recalled is a finite history of inflation realizations $\{\pi_{t-n}\}_{n=1}^{N}$. The probability of recalling a given past inflation episode is determined by its similarity to the cues that the firm is exposed to. In Gennaioli et al. (2024), the 'numerical context' cue is the current realized inflation rate. In the model developed here, firms do not observe that current inflation within the period, but rather observe the noisy signal $s_{it}(\pi_t)$. I take this noisy signal as the cue, which gives the recall probability function

$$\Pr(\text{recall } \pi_{t-n}|s_{it}(\pi_t)) = \frac{S(s_{it}(\pi_t), \pi_{t-n})}{\sum_{m=1}^{N} S(s_{it}(\pi_t), \pi_{t-m})} \tag{127}$$

where S is a similarity function, often chosen to be exponential:

$$S(s_{it}(\pi_t), \pi_{t-n}) = \exp\left(-(s_{it}(\pi_t) - \pi_{t-n})^2\right) \tag{128}$$

Firms recall a finite number of experiences, to give them a recalled sample $\mathcal{R}_i$. In Gennaioli et al. (2024), agents form expectations of future inflation by averaging over their recalled sample. However, to fit with the framework outlined above (especially equation (7)), I instead model firms using this selectively-recalled sample to form beliefs about inflation persistence $\chi_{i\pi_{t+1}\pi_t}$. In this case, the perceived inflation persistence of firm i is the estimated first-order autocorrelation of inflation over the sample of inflation episodes in $\mathcal{R}_i$.

$$\chi_{i\pi_{t+1}\pi_t} = \frac{\sum_{t-n\in\mathcal{R}_i} \pi_{t-n+1}\pi_{t-n}}{\sum_{t-n\in\mathcal{R}_i} \pi_{t-n}^2} \tag{129}$$

where I am assuming that for every recalled episode, the firm recalls the level of inflation, and the inflation one period afterwards, so they recall how persistent inflation was in that recalled event.

This naturally raises the possibility of a narrative heterogeneity channel. Firms with precise signals on π_t (high $\tau_{i\pi_t}$) recall past inflation episodes with inflation close to π_t. Firms with less precise signals (lower $\tau_{i\pi_t}$) have a greater variance of observed signals, and thus greater dispersion in their recalled samples. If the persistence of inflation varies over time or with realized inflation (as found in e.g. Cogley et al., 2010), then these firms sampling from different parts of the historical inflation distribution will form systematically different $\chi_{i\pi_{t+1}\pi_t}$, and thus will extrapolate from signals in systematically different ways.

A simple example makes this clear. Suppose that firms only recall a single past inflation experience, and the similarity function is such that they recall the most similar past experience to their numerical cue $s_{it}(\pi_t)$ with certainty. Furthermore, assume that the support of the history of available inflation rates is $\pi_{t-n} \in [\pi^-, \pi^+]$, realized π_t is also within this support, and the sample autocorrelation of inflation is increasing in π_{t-n}.

If the inflation history is dense, then the recalled inflation rate of household i is

$$\pi_{t-n_i^*} = \begin{cases} \pi^- & \text{if } s_{it}(\pi_t) \leq \pi^- \\ s_{it}(\pi_t) & \text{if } s_{it}(\pi_t) \in [\pi^-, \pi^+] \\ \pi^+ & \text{if } s_{it}(\pi_t) \geq \pi^+ \end{cases} \tag{130}$$

The average recalled inflation for firms with information precision $\tau_{i\pi_t}$ is then

$$E(\pi_{t-n_i^*}|\tau_{i\pi_t}) = \pi_t + \sigma_\pi \sqrt{\frac{1-\tau_{i\pi_t}}{\tau_{i\pi_t}}} \cdot [a\Phi(a) - b(1-\Phi(b)) + \phi(a) - \phi(b)] \tag{131}$$

where $a \equiv \frac{\pi^- - \pi_t}{\sigma_i}$, $b \equiv \frac{\pi^+ - \pi_t}{\sigma_i}$, and $\Phi(\cdot), \phi(\cdot)$ are the CDF and PDF of $N(0,1)$.

When $\pi_t = \frac{\pi^- + \pi^+}{2}$, the probabilities of hitting either boundary of the available inflation history are equal, and $E(\pi_{t-n_i^*}|\tau_{i\pi_t}) = \pi_t$ independent of $\tau_{i\pi_t}$. However, when $\pi_t > \frac{\pi^- + \pi^+}{2}$, firms are more likely to be constrained by the highest inflation in their memory database, and so $E(\pi_{t-n_i^*}|\tau_{i\pi_t}) < \pi_t$. The greater the variance of signal noise (i.e. the lower $\tau_{i\pi_t}$), the more firms are constrained by the limits of this database, and so the greater this under-estimation will be. Since we assumed the sample autocorrelation of inflation increased with the level of inflation, this means that $Cov_I[\pi_t, \chi_{i\pi_{t+1}\pi_t}] < 0$. In contrast, when $\pi_t < \frac{\pi^- + \pi^+}{2}$, the logic is reversed, and $Cov_I[\pi_t, \chi_{i\pi_{t+1}\pi_t}] > 0$. The narrative heterogeneity channel is therefore state-dependent, as realized inflation affects firm signals, and thus the distribution of recalled experiences.

Data generated through production. Farboodi and Veldkamp (2021) have recently introduced models in which firms generate information (data) through production. To model this in our framework, note that log-linearizing the firm-level demand equation used to derive the Calvo first order condition (107) implies (relative) production is a decreasing function of relative price:

$$y_{it} - Y_t = -\gamma(p_{it} - P_t) \tag{132}$$

where y_{it} and Y_t denote firm i and aggregate output respectively, and γ is the elasticity of substitution between firms' products.

If a firm has a low relative price in period t, they therefore produce more, and thus receive a more precise signal about inflation in the following period. As in the model with media, this introduces time-variation in signal-to-noise ratios. Specifically, I follow Farboodi and Veldkamp (2021) and assume that signal precision is linearly increasing in production (i.e. is linearly decreasing in the relative price), which implies firm i's signal-to-noise ratio in period t is:

$$\tau_{i\pi_t,t} = \frac{\sigma_\pi^2}{\sigma_\pi^2 + \frac{\sigma_\epsilon^2}{z(p_{it}-P_t)}}, \quad z(p_{it} - P_t) = \max\{\delta, z^* - (p_{it} - P_t)\} \tag{133}$$

where σ_ϵ^2 is constant across firms, $z^* > 0$ is a constant, and $\delta > 0$ is arbitrarily close to 0. I focus here on the tractable case where firms are myopic, so resetting firms do not account for the value of future signals when choosing prices in period t, and the pricing decision is as in equation (107). The precision of a firm's signal in period t^* is therefore:

$$\sigma_\epsilon^{-2} z(p_{it^*} - P_{t^*}) = \sigma^{-2} \cdot \max(\delta, z^* - (p_{it^*} - P_{t^*-1}) + (P_{t^*} - P_{t^*-1})) \tag{134}$$

$$= \sigma^{-2} \cdot \max(\delta, z^* + (P_{t^*} - P_{t^*-1}) - (1-\beta\lambda)\tilde{\chi}_i \tau_{i\pi_t,t^*-1} s_{it^*-1}(\pi_{t^*-1})) \tag{135}$$

If $s_{it^*-1}(\pi_{t^*-1}) > 0$, this is strictly decreasing in $\tilde{\chi}_i$. If $s_{it^*-1}(\pi_{t^*-1}) < 0$, it is strictly increasing. The signal-to-noise ratio $\tau_{i\pi_t,t^*}$ is strictly increasing in this precision.

Consider the case where $\pi_{t^*-1} > 0$, and all $\tilde{\chi}_i > 0$. The majority of resetting firms will receive $s_{it^*-1}(\pi_{t^*-1}) > 0$, and will increase prices. Firms with larger $\tilde{\chi}_i$ increase prices by more than those with small $\tilde{\chi}_i$. As a result, large $\tilde{\chi}_i$ firms have less precise information in period t^* than those with small $\tilde{\chi}_i$: there is a negative $Cov_I[\tilde{\chi}_i, \tau_{i,t^*}]$. The reverse would be true if $\pi_{t^*-1} < 0$. Overall, the narrative heterogeneity channel arises because subjective models affect pricing decisions, which in turn affects information production. As in the model with media, the narrative heterogeneity channel produced here is therefore state-dependent.

If firms were not myopic, there would also be a second channel here: firms would set prices below the within-period optimum to increase the precision of signals in the next period. As discussed extensively above, the value of precise information is larger for firms with large $\tilde{\chi}_i^2$, so this is also likely to induce a correlation between $\tilde{\chi}_i$ and information precision.

Heterogeneous price indices. Recent literature (e.g. Kaplan and Schulhofer-Wohl, 2017) has found that households in particular face highly heterogeneous inflation rates, and so some heterogeneity in survey expectations may come because households are forecasting their own idiosyncratic inflation rates. We now adapt that idea to our simple firm model, and show that it can also generate a narrative heterogeneity channel, even if all agents have rational expectations.

Suppose that each firm i belongs to a sector k, and that each sector faces its own inflation process, which is a weighted average of past sector-specific inflation and an aggregate i.i.d. shock $u_t \sim N(0, \sigma_u^2)$:

$$\pi_{kt} = \rho_k \pi_{kt-1} + (1 - \rho_s k) u_t \tag{136}$$

where $\rho_k \in (0, 1)$ is the persistence of that sector's inflation process.

Firms observe their own sector's inflation at the end of each period, and receive signals of the form:

$$s_{i\pi_{kt}} = \pi_{kt} + \epsilon_{i\pi_{kt}}, \quad \epsilon_{i\pi_{kt}} \sim N(0, \sigma_\epsilon^2) \tag{137}$$

The posterior expectation of π_{kt} after observing this signal is:

$$E_{it}\pi_{kt} = (1 - \tau_{k\pi_t})\rho_k \pi_{kt-1} + \tau_{k\pi_t}(\pi_{kt} + \epsilon_{i\pi_{kt}}) \tag{138}$$

where

$$\tau_{k\pi_t} = \frac{(1 - \rho_k)^2 \sigma_u^2}{(1 - \rho_k)^2 \sigma_u^2 + \sigma_\epsilon^2} \tag{139}$$

From this, we see that τ_k is strictly decreasing in ρ_k for all $\rho_k < 1$. This is because greater persistence in inflation means that in any given period, prior beliefs are more informative, leading to less reaction to current signals. This heterogeneity in $\tau_{k\pi_t}$ arises even though the precision of inflation signals is assumed to be equal across sectors.

Moreover, inflation persistence also affects how firms cross-learn from current to future inflation. It is straightforward to see that

$$E_{it}\pi_{kt+s} = \rho_k^s E_{it}\pi_{kt} \tag{140}$$

and thus that $\chi_{i\pi_{t+s}\pi_t} = \rho_k^s$, which is increasing in ρ_k. The covariance between $\chi_{i\pi_{t+s}\pi_t}$ and $\tau_{i\pi_t}$ is therefore negative. Firms in sectors with less persistent inflation react more to inflation information, but extrapolate less strongly from current inflation expectations to the future. Responsiveness to information is therefore concentrated among those who cross-learn the least from it, which is a narrative heterogeneity effect that weakens the aggregate causal effect of inflation expectations.

Level-k thinking. For this, we switch the setup slightly to facilitate tractability. Specifically, now assume that the only signals available to firms are on current marginal costs mc_t, not inflation. There is cross-learning from mc_t to π_t, but no cross-learning to longer horizons. That is, $\chi_{i\pi_{t+s}mc_t} = \chi_{imc_{t+s}mc_t} = 0$ for all $s > 0$. As in the other cases above, assume there is no public noise in signals. In this case, the optimal price of a resetting firm is

$$p^*_{it} - P_{t-1} = (1 - \beta\lambda)[1 + \chi_{i\pi_t mc_t}]\tau_{imc_t}(mc_t + \varepsilon_{imc_t}) \tag{141}$$

This can alternatively be written as a static beauty contest

$$p^*_{it} - P_{t-1} = (1 - \beta\lambda)[E_{it}mc_t + E_{it}\pi_t] = (1 - \beta\lambda)\left[E_{it}mc_t + (1 - \lambda)E_{it}\int(p^*_{jt} - P_{t-1})dj\right] \tag{142}$$

Following Farhi and Werning (2019) and Iovino and Sergeyev (2023), a firm that reasons to level 0 observes their signal about mc_t, and updates their expectations of the shock accordingly, but fails to update their beliefs about the prices of others. That is, they ignore general equilibrium effects of the shock. The price choice of such a level-0 firm is therefore:

$$p^{*0}_{it} - P_{t-1} = (1 - \beta\lambda)E_{it}mc_t = (1 - \beta\lambda)\tau_{imc_t}s_{it}(mc_t) \tag{143}$$

A firm that reasons to level 1 realizes that others may react to the shock as well, but do not reason all the way to general equilibrium. Rather, they assume that all other firms are level 0. In this case, the optimal reset price is:

$$p^{*1}_{it} - P_{t-1} = (1 - \beta\lambda)\left[E_{it}mc_t + (1 - \lambda)E_{it}\int(p^{*0}_{jt} - P_{t-1})dj\right] \tag{144}$$

$$= (1 - \beta\lambda)[1 + (1 - \lambda)(1 - \beta\lambda)\bar{\tau}_{mc_t}]\tau_{imc_t}s_{it}(mc_t) \tag{145}$$

For the purpose of this example, I stop at $k = 1$. Letting χ^k be the $\chi_{i\pi_t mc_t}$ of a firm engaged in k levels of reasoning, we have:

$$\chi^0 = 0, \quad \chi^1 = (1 - \lambda)(1 - \beta\lambda)\bar{\tau}_{mc_t} \tag{146}$$

Level-1 reasoners have a strictly greater $\chi_{i\pi_t mc_t}$ than level-0 reasoners. If the depth of reasoning is correlated with the precision of information τ_i, then there will be a narrative heterogeneity effect. I now show that such a correlation arises naturally if firms are able to choose their depth of reasoning, as in Alaoui and Penta (2016). In the model they develop, agents decide whether to complete an extra level of reasoning by weighing up the perceived costs and benefits. Specifically, in their model, two games are said to be cognitively equivalent if they entail the same sequence of actions as reasoning develops (k increases). To fit the firm problem into this framework, denote the choice of $\chi_{i\pi_t mc_t}$ as the firm's action, with pricing following mechanically using equation (141). χ^0 and χ^1 are independent of all firm-level variables, so framed in this way the problem facing firm i is cognitively equivalent to the one facing firm j. As a result, the costs of reasoning are the same for all firms, and the depth of reasoning only differs according to the value of doing so.

Alaoui and Penta (2016) propose a number of functions for the value of reasoning. I demonstrate the narrative heterogeneity channel for the general 'expected gain' formulation from their Example 2, which is a generalization of the form they use for their examples, and is given an axiomatic foundation in Alaoui and Penta (2022). Adapting it to the firm problem, the value of reasoning to level 1 is given by:

$$v_i^1 = \int p(\chi') \left[\mathbb{E}_{it-1} U_{i,t}(BR(\chi'), \chi') - \mathbb{E}_{it} U_{i,t-1}(\chi^0, \chi') \right] d\chi' \tag{147}$$

where $p(\chi')$ are weights such that $p(\chi') \in [0, 1]$ and $\int p(\chi')d\chi' = 1$, $U_{i,t}(\chi_i, \chi')$ is the payoff to choosing χ_i when all other firms choose χ', and $BR(\chi')$ is the χ_i that maximizes $\mathbb{E}_{it-1}U_{i,t}(\chi_i, \chi')$ for a given χ'. That is, the value of reasoning is the expected utility gain from switching from χ^0 to the best-response χ_i, weighted by the firm's perceived probability that reasoning will reveal $\chi_j = \chi'$. Reasoning is undertaken before the observation of $s_{i,t}$, so the expectations are indexed $t - 1$.

If $\chi_j = \chi'$ for all firms $j \neq i$, the optimal price is (from equation (142)):

$$p_{it}^* - P_{t-1} = (1 - \beta\lambda) \left[1 + (1 - \lambda)(1 - \beta\lambda)(1 + \chi')\bar{\tau}_{mc_t}\right] \tau_{imc_t} s_{it}(mc_t) \tag{148}$$

The best-response χ_i is therefore:

$$BR(\chi') = (1 - \lambda)(1 - \beta\lambda)(1 + \chi')\bar{\tau}_{mc_t} \tag{149}$$

Following the steps used to compute indirect utility in the rational inattention case above, the value of reasoning to level 1 is:

$$v_i^1 = (1 - \beta\lambda)^4 (1 - \lambda)^2 \sigma_{mc}^2 \bar{\tau}_{mc_t}^2 \tau_{imc_t} \int p(\chi')(1 + \chi')^2 d\chi' \tag{150}$$

The value of reasoning is therefore linearly increasing in τ_{imc_t}, whatever perceptions $p(\chi')$ the firm holds. If we further assume that all firms use the same $p(\chi')$ when evaluating the benefits of reasoning (as in Alaoui and Penta, 2016), then there exists a threshold τ^* above which firms decide to reason to level 1, and below which they do not. This implies that high τ_{imc_t} is associated with high $\chi_{i\pi_t mc_t}$, so the narrative heterogeneity channel is positive. Formally:

$$Cov_I[\chi_{i\pi_t mc_t}, \tau_{imc_t}] = q\chi^0\tau^- + (1-q)\chi^1\tau^+ - \bar{\tau}(q\chi^0 + (1-q)\chi^1) \tag{151}$$

where q is the proportion with $\tau_{imc_t} > \tau^*$, and τ^-, τ^+ are the average τ_{imc_t} conditional on $\tau_{imc_t} < \tau^*$ and $\tau_{imc_t} > \tau^*$ respectively. Simplifying, this becomes:

$$Cov_I[\chi_{i\pi_t mc_t}, \tau_{imc_t}] = -q(1-q)(\tau^+ - \tau^-)(\chi^0 - \chi^1) > 0 \tag{152}$$

Intuitively, firms with more precise information on shocks have more to gain from understanding how to use that information, and so they will choose a higher k than firms with less information. Information is therefore concentrated among firms with large $\chi_{i\pi_t mc_t}$.

C Mapping RCT estimates to the Calvo model: details

RCT-implied calibration. To adapt the Baumann et al. (2024) estimates to the particular model set out in Section 2, we need to take two steps: first, interpolate between estimates at different horizons to get the total effect of a given signal on inflation at horizons $t+2$ and $t+4$, where there are no estimates available. I use a simple linear interpolation here. Second, the estimates refer to annual inflation, while the model is expressed in terms of cumulative inflation since period $t-1$. We therefore sum them up to get the relevant coefficients for aggregate expectations updating.

To see why we need to sum, notice that inflation in the survey is given as net annual inflation:

$$\pi^a_{t+s} = \frac{P_{t+s} - P_{t+s-1}}{P_{t+s-1}} \approx \log\left(\frac{P_{t+s}}{P_{t+s-1}}\right) \tag{153}$$

Cumulative log inflation (as in the model) is therefore given by:

$$\pi_{t+s} = \log\left(\frac{P_{t+s}}{P_{t-1}}\right) = \log\left(\prod_{m=0}^{s} \frac{P_{t+m}}{P_{t+m-1}}\right) \approx \sum_{m=0}^{s} \pi^a_{t+m} \tag{154}$$

Any change in cumulative inflation is therefore equal to the sum of changes of annual inflation at all periods between now and the relevant horizon. With these adjustments, the coefficients implied by Table 5 of Baumann et al. (2024) are presented in the left panel of Table 4.

Table 4: Mapping estimates from Baumann et al. (2024) to the Calvo model.

Model object	π_t signal	π_{t+1} signal	Model object	π_{t+1} signal
$E_I[\chi_{\text{signal period},\pi_t}\tau_{\text{signal period}}]$	1	N/A	$E_I[\chi_{\text{signal period},\pi_t}]$	N/A
$E_I[\chi_{\text{signal period},\pi_{t+1}}\tau_{\text{signal period}}]$	1.21	0.53	$E_I[\chi_{\text{signal period},\pi_{t+1}}]$	1
$E_I[\chi_{\text{signal period},\pi_{t+2}}\tau_{\text{signal period}}]$	1.51	0.945	$E_I[\chi_{\text{signal period},\pi_{t+2}}]$	1.87
$E_I[\chi_{\text{signal period},\pi_{t+3}}\tau_{\text{signal period}}]$	1.9	1.245	$E_I[\chi_{\text{signal period},\pi_{t+3}}]$	2.61
$E_I[\chi_{\text{signal period},\pi_{t+4}}\tau_{\text{signal period}}]$	2.235	1.56	$E_I[\chi_{\text{signal period},\pi_{t+4}}]$	3.365
$E_I[\chi_{\text{signal period},\pi_{t+5}}\tau_{\text{signal period}}]$	2.515	1.86	$E_I[\chi_{\text{signal period},\pi_{t+5}}]$	4.135

In principle, we could extrapolate to parameters further into the future, but I will cut off firm planning horizons at 5 years here, for this and the two benchmarks.

For the decomposition in Figure 2, we proceed in a similar fashion. Estimating equation (37) with $E_{it}\pi_{t+3}$ as the dependent variable, and instrumenting $E^{post}_{it}\pi_{t+1}$ with the π_{t+1} treatment only, gives an estimated γ_1 of 0.74. Doing the same with $E_{it}\pi_{t+5}$ as the dependent variable gives an estimated γ_1 of 0.77. Both of these coefficients are significantly different from 0 at the 1% level. Applying the same procedure as described above gives the right panel of Table 4. Note that since we assume $\bar{\tau}_{\pi_t} = 1$, and thus $\tau_{i\pi_t} = 1$ for all i, there is no narrative heterogeneity channel from π_t signals, so we do not need to compute $\bar{\chi}_{\pi_t,\pi_{t+h}}$.

Including cross-learning to marginal costs. Baumann et al. (2024) Table 8 contains results from estimating equation (37), which (as argued in Section 4.1) reflect average cross-learning from the inflation period in the signal to each X_i variable. They report results for $X_i \in \{$1-yr-ahead expected wage, non-wage costs, employment$\}$. To go from those variables to the marginal cost that enters the model, we need to specify the firm's production function, before then applying the temporary-information equilibrium approach. For the exercise here, I will use the simple production function

$$y_{it} = \min\left(\frac{n_{it}}{\alpha_n}, \frac{x_{it}}{\alpha_x}\right) \tag{155}$$

where labor n_{it} and other inputs x_{it} are perfect complements, so the optimal factor ratio is

$$\frac{n_{it}}{x_{it}} = \frac{\alpha_n}{\alpha_x} \tag{156}$$

In this case, real marginal costs are given by

$$mc_t = \alpha_n \frac{w_t}{P_t} + \alpha_x \frac{p_{xt}}{P_t} \tag{157}$$

where w_t is the nominal wage and p_{xt} is the price of x_{it}. Log-linearizing we obtain

$$mc_t = \tilde{\alpha}_n w_t + \tilde{\alpha}_x p_{xt} - (\tilde{\alpha}_n + \tilde{\alpha}_x) P_t \tag{158}$$

where

$$\tilde{\alpha}_n \equiv \frac{\alpha_n w^{ss}}{\alpha_n w^{ss} + \alpha_x p_x^{ss}}, \quad \tilde{\alpha}_x \equiv \frac{\alpha_x p_x^{ss}}{\alpha_n w^{ss} + \alpha_x p_x^{ss}} \tag{159}$$

are the steady-state shares of labor and other costs in marginal costs. Since this model features a constant marginal cost, these are also the steady state shares of labor costs and other costs in total costs.

The average update to expected marginal costs after a π_{t+1} signal is therefore

$$\frac{dE_{it} mc_{t+1}}{d\varepsilon_{\pi_{t+1}}} = \tilde{\alpha}_n \frac{dE_{it} w_{t+1}}{d\varepsilon_{\pi_{t+1}}} + \tilde{\alpha}_x \frac{dE_{it} p_{xt+1}}{d\varepsilon_{\pi_{t+1}}} - \frac{dE_{it} \pi_{t+1}}{d\varepsilon_{\pi_{t+1}}} \tag{160}$$

$$= \underbrace{\left(\tilde{\alpha}_n \chi_{iw_{t+1}\pi_{t+1}} + \tilde{\alpha}_x \chi_{ip_{xt+1}\pi_{t+1}} - 1\right)}_{\equiv \chi_{imc_{t+1}\pi_{t+1}}} \tau_{i\pi_{t+1}} \tag{161}$$

where I have used the observation that $\tilde{\alpha}_n + \tilde{\alpha}_x = 1$, the assumption that $\tau_{i\pi_t} = 1$, the definition $P_{t+1} = \pi_{t+1} + P_{t-1}$,[35] and the fact that P_{t-1} is fixed as it was determined before the shock. I have also assumed that there is no cross-learning from π_t to $E_{it} w_{t+1}$ or $E_{it} p_{xt+1}$, because all estimates for cross-learning from the current-inflation treatment in Baumann et al. (2024) are not significantly different from 0.

The average cross-learning from π_{t+1} to mc_{t+1} can therefore be computed from Baumann et al. (2024) Table 8 as

$$\bar{\chi}_{mc_{t+1}\pi_{t+1}} = \tilde{\alpha}_n \bar{\chi}_{w_{t+1}\pi_{t+1}} + \tilde{\alpha}_x \bar{\chi}_{p_{xt+1}\pi_{t+1}} - 1 \tag{162}$$

$$= \tilde{\alpha}_n \times 0.29 + (1 - \tilde{\alpha}_n) \times 0.90 - 1 \tag{163}$$

$$= -0.4538 \tag{164}$$

where for the final equality I have assumed that the labor share of total costs is $\tilde{\alpha}_n = 0.58$. As this data is not available for the Euro area , this is taken from the 2023 measurement for the US.[36] On average firms therefore decrease their real marginal cost expectations when they learn of higher future inflation. This is because the positive updating found by Baumann et al. (2024) is over the

[35]Recall that π_{t+1} is defined as the cumulative inflation $P_{t+1} - P_{t-1}$ in this model.

[36]Available at https://fred.stlouisfed.org/series/MPU4900141.

components of nominal marginal costs, and those updates are dominated by the increase in inflation, causing real marginal cost expectations to fall.

As in the exercises in Section 4, one could in principle extrapolate from this to cross-learning about marginal costs at longer horizons, but here I assume for simplicity that all cross-learning beyond the horizon we observe (here one year) is zero. Allowing for the cross-learning from π_{t+1} to mc_{t+1}, the partial equilibrium effects of $\bar{E}_t(\pi_{t+1})$ are therefore given by

$$\Omega^{PE}(\pi_{t+1}) = \Omega^{PE}(\pi_{t+1} : \chi_{imc_{t+1}\pi_{t+1}} = 0) + (1-\lambda)(1-\beta\lambda)\beta\lambda \times \left(\bar{\chi}_{mc_{t+1}\pi_{t+1}} + \frac{1}{\bar{\tau}_{\pi_{t+1}}} Cov_I[\chi_{imc_{t+1}\pi_{t+1}}, \tau_{i\pi_{t+1}}] \right) \quad (165)$$

where $\Omega^{PE}(\pi_{t+1} : \chi_{imc_{t+1}\pi_{t+1}} = 0)$ denotes the partial equilibrium effect calculated in Figure 1a under the assumption of no cross-learning to real marginal costs. The general equilibrium effects can then be calculated from this using equation (16). Since I am assuming no cross-learning from π_t to marginal costs, the $\Omega^{PE}(\pi_t)$ term in equation (16) is unaffected.

Note that while we can calibrate $\bar{\chi}_{mc_{t+1}\pi_{t+1}}$ from the data in Baumann et al. (2024), we cannot calibrate the narrative heterogeneity term $Cov_I[\chi_{imc_{t+1}\pi_{t+1}}, \tau_{i\pi_{t+1}}]$. For Figure 6 below I will assume that the narrative heterogeneity term is 0, so the results show what would happen if all firms updated marginal cost expectations in the same way.

The choice of benchmark models is difficult here, because we haven't modeled how wages and costs are determined. Without specifying a full model for those processes, we cannot write down the perfect foresight or rational expectations benchmarks. Instead, I therefore compare the results to the case calibrated to the RCT, but with no cross-learning to future marginal costs. That is, the benchmark I will use is the black line in Figures 1a and 1b. This gives a sense of whether the updates to cost expectations found in the RCT are quantitatively important relative to only looking at the effects on inflation expectations.

Figure 6 therefore plots the partial and general equilibrium effects of $\bar{E}_t(\pi_{t+1})$ with and without cross-learning to marginal costs. Again the quantitative significance of that cross-learning depends on price stickiness. In partial equilibrium (panel a), the effects of inflation expectations are smaller than they are without marginal cost learning at all values of price stickiness λ, because higher expected inflation causes firms to reduce their real marginal cost expectations, which tempers their incentives to raise prices. Proportionally, this effect is stronger when prices are more flexible (λ is small): at the plausible calibration of $\lambda = 0.2$, the PE effect including marginal cost cross-learning is 69% as large as the effect without such cross-learning. At $\lambda = 0.5$, it is 84% as large. This is again because stickier prices imply longer-horizon expectations matter relatively more for pricing decisions. Since the marginal-cost cross-learning is assumed to only affect marginal costs one period in the future, its effect on inflation dynamics declines as λ rises.

In general equilibrium, again the effect is smaller once we allow for marginal cost cross-learning, and the magnitude of the difference is large at low values of λ. However, the differences to the baseline case are smaller than in partial equilibrium: at $\lambda = 0.2$ and 0.5 respectively the GE effect with marginal cost cross-learning is 84% and 92% as large as the baseline case without that cross-learning. This is because of an effect discussed in Section 4: a larger partial equilibrium effect $\Omega^{PE}(\pi_{t+1})$ actually dampens general equilibrium feedback, because it causes a greater scaling of the original shock to ensure that $\bar{E}_t(\pi_{t+1})$ increases by 1 unit. When we introduce cross-learning to marginal costs, that partial equilibrium effect falls, but there is no change on the other component of the general equilibrium effect through $\Omega^{PE}(\pi_t)$. General equilibrium feedback is therefore stronger in the presence of marginal cost cross-learning, but this is purely due to the way we are scaling the shock in computing these causal effects.

This mechanism is also why the general equilibrium effect of expected inflation initially rises with λ once we allow for marginal cost cross-learning. As λ increases the general equilibrium feedback through $\Omega^{PE}(\pi_t)$ becomes larger, while the counter-effect on general equilibrium feedback through $\Omega^{PE}(\pi_{t+1})$ and scaling makes the general equilibrium effect smaller. Marginal cost cross-learning weakens the second of these mechanisms, so that overall the causal effect in general equilibrium is initially increasing in λ.

Figure 6: Partial and general equilibrium effects of a shock to $s_{\pi_{t+1}}$, comparing RCT estimates with and without cross-learning to mc_{t+1}.

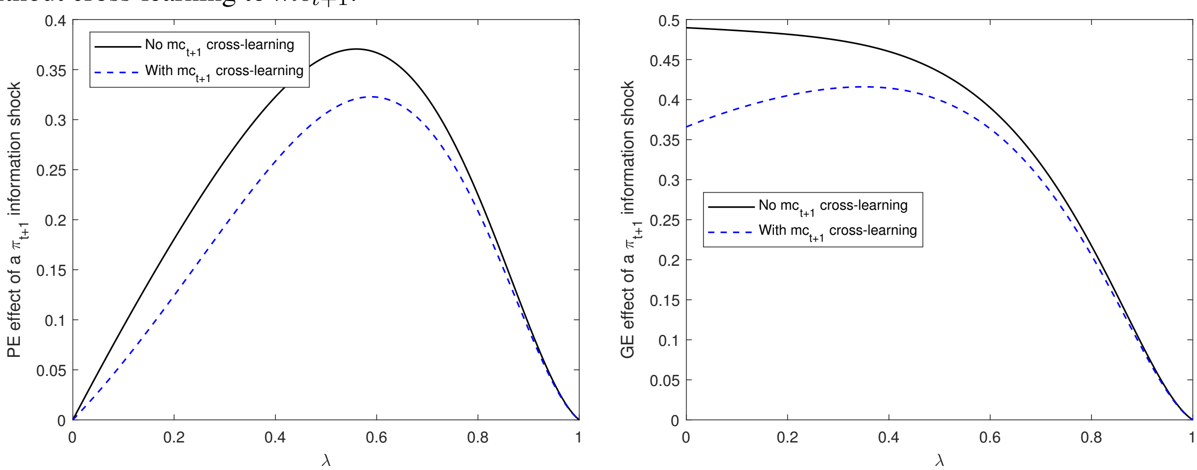

(a) Partial Equilibrium **(b)** General Equilibrium

Note: Figure plots the partial- and general-equilibrium causal effects of an increase in $\bar{E}_t\pi_{t+1}$, following the formulae in Propositions 2 and 3, and a calibration derived from the estimates in Baumann et al. (2024) (detail in Appendix C), with and without allowing cross-learning from π_{t+1} to mc_{t+1}.

The share of fundamental shock transmission due to expectations. Following the strategy in Proposition 4, we now compute the share of the general equilibrium transmission of a fundamental shock that can be attributed to inflation expectations, for the RCT data and both benchmarks (perfect foresight and rational expectations). Since we are assuming $\bar{\tau}_{\pi_t} = 1$, these shares for an i.i.d. mc_t shock would just equal the partial equilibrium effect of a $\bar{E}_t(\pi_t)$ shock (equation (21)). We therefore opt instead to consider a news shock: firms learn in period t of an exogenous increase in mc_{t+1}. This demonstrates the flexibility of the approach developed here.

With this shock the transmission channels are more numerous. Once period $t+1$ arrives, the increased marginal costs will imply inflation above steady state, as described in equation (20). In period t, firms therefore not only learn that mc_{t+1} will be above steady state, their signals about π_{t+1} also endogenously increase. Even without cross-learning from marginal cost expectations to inflation expectations, there is therefore an increase in expected future inflation, which has cross-learning effects on inflation expectations at other horizons. All of these forces contribute to the increase in π_t, which in turn has general equilibrium feedback effects through $s_{it}(\pi_t)$.

Formally, we model the news shock by assuming that $\tau_{imc_{t+1}} = 1$. We also assume that marginal costs are i.i.d., so there is no cross-learning from the news about future marginal cost to expected marginal costs at any other horizon, or to expected inflation. Differentiating equation (11) with respect to mc_{t+1} with these assumptions implies

$$\frac{d\pi_t}{dmc_{t+1}} = (1-\lambda)(1-\beta\lambda)\beta\lambda + \bar{\tau}_{\pi_{t+1}}\Omega^{PE}(\pi_{t+1})\frac{d\pi_{t+1}}{dmc_{t+1}} + \Omega^{PE}(\pi_t)\frac{d\pi_t}{dmc_{t+1}} \tag{166}$$

where we have also used our recurring assumption that $\bar{\tau}_{\pi_t} = 1$.

Next, from equation (20) we have

$$\frac{d\pi_{t+1}}{dmc_{t+1}} = \frac{d\pi_t}{dmc_t} = \frac{(1-\lambda)(1-\beta\lambda)}{1-\Omega^{PE}(\pi_t)} \tag{167}$$

where it is $\Omega^{PE}(\pi_t)$, not $\Omega^{PE}(\pi_{t+1})$ that matters because by the time period $t+1$ arrives, the shock will be contemporaneous, so it is the effect of a contemporaneous shock that determines the contemporaneous inflation rate – i.e. π_{t+1}.

Substituting this in to equation (166) and rearranging, we find

$$\frac{d\pi_t}{dmc_{t+1}} = \frac{(1-\lambda)(1-\beta\lambda)\beta\lambda}{1-\Omega^{PE}(\pi_t)} + \frac{(1-\lambda)(1-\beta\lambda)\bar{\tau}_{\pi_{t+1}}\Omega^{PE}(\pi_{t+1})}{(1-\Omega^{PE}(\pi_t))^2} \tag{168}$$

We define the share of this overall transmission that is due to expectations as in equation (21), where the counterfactual shock transmission is computed assuming firms only observe the mc_{t+1} signal, with no updating of any other expectations. Thus this calculation gives the share of shock transmission due to expectations of inflation specifically, as marginal cost expectations are the same in the overall effect and the counterfactual.

This counterfactual with no response of inflation expectations is found by setting $\Omega^{PE}(\pi_t) = \Omega^{PE}(\pi_{t+1}) = 0$, at which point

$$\left.\frac{d\pi_t}{dmc_{t+1}}\right|_{\sim E} = (1-\lambda)(1-\beta\lambda)\beta\lambda \tag{169}$$

Using these expressions, we calculate the share of the shock transmission to π_t that is due to the response of inflation expectations as

$$\frac{\frac{d\pi_t}{dmc_{t+1}} - \left.\frac{d\pi_t}{dmc_{t+1}}\right|_{\sim E}}{\frac{d\pi_t}{dmc_{t+1}}} = 1 - \frac{\beta\lambda(1-\Omega^{PE}(\pi_t))^2}{\beta\lambda(1-\Omega^{PE}(\pi_t)) + \bar{\tau}_{\pi_{t+1}}\Omega^{PE}(\pi_{t+1})} \tag{170}$$

To construct each line in Figure 7, we replace $\Omega^{PE}(\pi_t)$ and $\Omega^{PE}(\pi_{t+1})$ with values calculated using equations (12) and (13) and the relevant calibrations of $\boldsymbol{\chi}_i$ and $\boldsymbol{\tau}_i$.

In this case, the extent of price stickiness can change the sign of the RCT-implied departure from standard benchmarks, although the quantitative differences remain small. At $\lambda = 0.2$, the share of the shock transmission due to the dynamics of inflation expectations is just under 1% smaller in the RCT calibration than in the rational expectations benchmark. At $\lambda = 0.5$, the RCT-implied share is 3% larger than the rational expectations benchmark. This sign change occurs because the RCT results imply somewhat smaller cross-learning from π_t to π_{t+1} than implied by rational expectations, even if long-horizon cross-learning is stronger (as discussed above). This means that when prices are rather flexible, and short-horizon expectations drive the vast majority of pricing decisions, the mc_{t+1} shock entails a smaller reaction in the most important expectation horizons than is generated with rational expectations. Of course, even the limited short-horizon cross-learning in the RCT calibration is stronger than seen under perfect foresight, so the transmission due to expectations is larger than the perfect foresight case at all levels of price stickiness.

Alternative general equilibrium assumptions. Figure 8 reproduces Figure 1b, except that in all lines the partial equilibrium effect is the one computed using the RCT results, but the general equilibrium feedback is taken from different assumptions on expectations. This incorrect way to calculate the general-equilibrium causal effects of expected inflation highlights how RCT results are

Figure 7: Share of transmission from a mc_{t+1} news shock to π_t due to the adjustment of inflation expectations, comparing RCT estimates to benchmarks

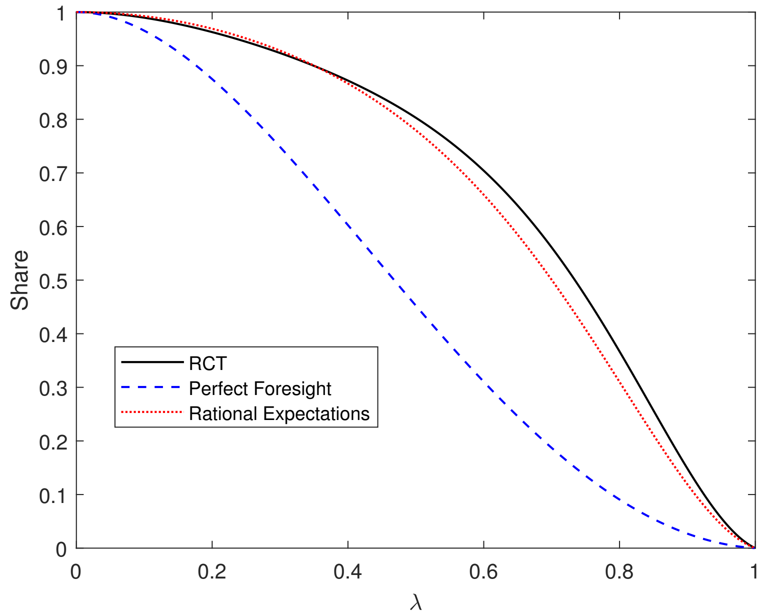

Note: Figure plots the share of the transmission of a mc_{t+1} news shock that is due to the endogenous response of inflation expectations in the Calvo model of Section 2, for calibrations based on perfect foresight, rational expectations, and the RCT results of Baumann et al. (2024) (detail in Appendix C).

important for the intercept (i.e. general equilibrium feedback) as well as the direct effect of interest.

Other models as benchmarks. Recall that the equilibrium New Keynesian Phillips Curve cannot reveal the causal effects of inflation expectations because expectations there simultaneously cause and are caused by dynamics of other variables. The same is true if we attempt to compare different models of expectation formation by studying equilibrium conditions. If, for example, a particular model implies a reduced link between inflation expectations and realized inflation, we cannot tell if that is because expectations have been made less responsive to inflation, or if their causal effect on inflation has been reduced.

The temporary-information equilibrium approach therefore offers a structured way to disentangle why certain models have different implications from one another, and to assess whether they match the equilibrium implications of RCT evidence. I provide an example here, re-computing Figure 1 for diagnostic expectations (Bordalo et al., 2020) and inattentiveness (Gabaix, 2020).

In Gabaix (2020), expectations of a variable x_{t+h} are given by

$$E_{it}x_{t+h} = m^h E_t^* x_{t+h} \tag{171}$$

Figure 8: General equilibrium effects of a shock to $s_{\pi_{t+1}}$, comparing RCT estimates to alternatives with general equilibrium effects from other benchmarks

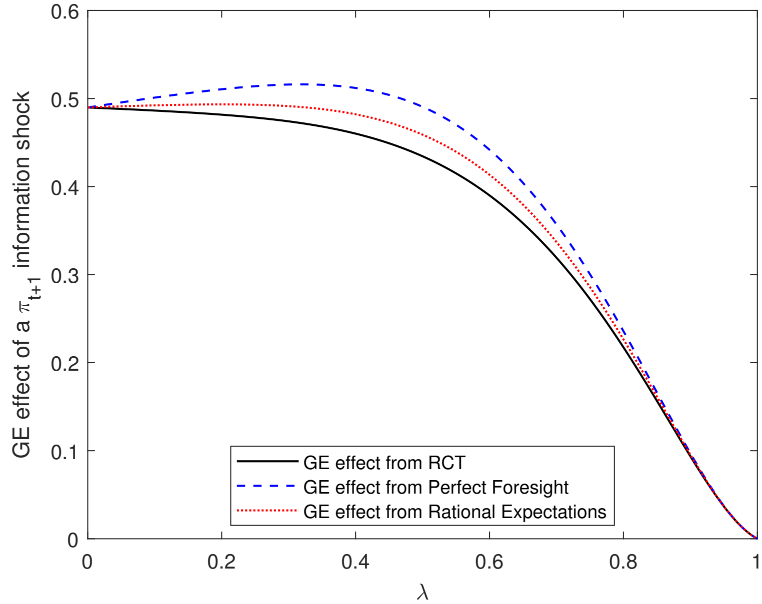

Note: Figure plots (in black) the general-equilibrium causal effects of an increase in $\bar{E}_t \pi_{t+1}$ in the Calvo model, following the formulae in Propositions 2 and 3, and a calibration derived from the estimates in Baumann et al. (2024). The red and blue lines replace the true calibrated general-equilibrium feedback effect with the equivalent calculated using the perfect foresight and rational expectations calibrations (detail in Appendix C).

where $E_t^* x_{t+h}$ is the rational expectation of x_{t+h} formed with information on all variables realized up to period t, and $m \in [0, 1]$ is a parameter governing the degree of inattentiveness. The first implication of this setup is that when $h = 0$, inattentiveness has no effect on expectations. That is, inattentiveness does not affect the mapping from observed information on a variable x_t into beliefs about x_t. It only affects how that information is used to form beliefs about x_{t+h}. Inattentiveness therefore affects cross-learning from x_t to x_{t+h}, but not the information matrix $\boldsymbol{\tau}_i$.

Specifically, we will restrict ourselves here to cases with no cross-learning to or from marginal costs. The inattentiveness case can therefore be summarized by

$$\chi^{\text{IA}}_{i\pi^a_{t+s}\pi^a_t} = m^h \chi^{\text{RE}}_{i\pi^a_{t+s}\pi^a_t} \tag{172}$$

where $\chi^{\text{RE}}_{i\pi^a_{t+s}\pi^a_t}$ is the cross-learning from annual inflation π^a_t to annual inflation π^a_{t+s}. We specify this in terms of per-period inflation to fit with the specification in Gabaix (2020). Translating this to the cumulative inflation used in our model, assuming that $\chi^{\text{RE}}_{i\pi^a_{t+s}\pi^a_t} = \rho^s$, is achieved by setting

$$\chi^{\text{IA}}_{i\pi_{t+s}\pi_t} = \frac{1 - (m\rho)^s}{1 - m\rho} \tag{173}$$

The second model we consider here is diagnostic expectations, as in Bordalo et al. (2020). In this model, if x_t is an AR(1) process with persistence ρ, we have

$$E_{it}x_{t+h} = \rho^h E_{it}x_t = \rho^h \left((1+\theta) \frac{\sigma_x^2}{\sigma_x^2 + \sigma_{\varepsilon x}^2} s_{it}(x_t) \right) \tag{174}$$

where θ is the degree of diagnosticity, $s_{it}(x_t)$ is the (noisy) signal observed about x_t, σ_x^2 is the variance of x_t, and $\sigma_{\varepsilon x}^2$ is the variance of the noise in $s_{it}(x_t)$. In this case, the departure from rational expectations operates solely on the way in which agents map from direct signals on x_t into expectations of x_t. It therefore affects $\boldsymbol{\tau}_i$ and not $\boldsymbol{\chi}_i$. Specifically, diagnostic expectations imply $\boldsymbol{\tau}_i$ is multiplied by $1+\theta$, so for $\theta > 0$ agents overreact to news.

Figure 9 plots the partial and general equilibrium effects of $\bar{E}_t(\pi_{t+1})$ in each of these models, alongside the effects under rational expectations and the RCT benchmark. For the calibration, I set $m = 0.8^4$, as advised by Gabaix (2020) for firm expectations. I set $\theta = 0.37$, the average of the results from the two methods employed by Bordalo et al. (2020) to estimate θ in the context of CPI inflation expectations (taken from their Table 4).

Figure 9: Partial and general equilibrium effects of a shock to $s_{\pi_{t+1}}$, comparing RCT estimates to rational expectations, inattentiveness, and diagnostic expectations.

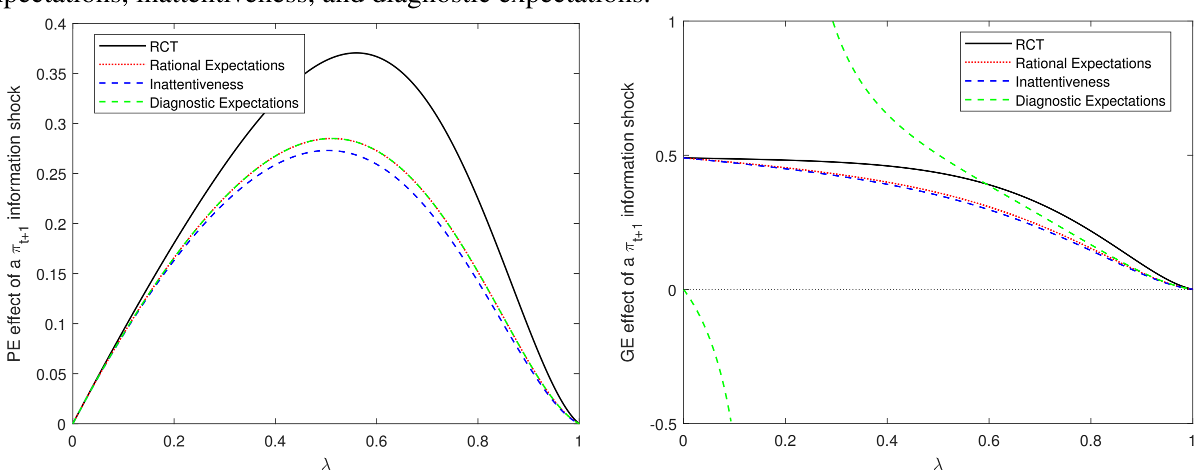

(a) Partial Equilibrium

(b) General Equilibrium

Note: Figure plots the partial- and general-equilibrium causal effects of an increase in $\bar{E}_t\pi_{t+1}$ in the Calvo model, following the formulae in Propositions 2 and 3, and a calibration derived from the estimates in Baumann et al. (2024) (detail in Appendix C). This RCT calibration is compared to calibrations based on perfect foresight, rational expectations, inattentiveness as in (Gabaix, 2020), and diagnostic expectations as in (Bordalo et al., 2020).

Inattentiveness generates similar results to rational expectations, with a small degree of dampening relative to that benchmark in both partial and general equilibrium. The dampening is proportionately larger at higher λ, when longer-horizon expectations are more important for current pricing.

Diagnostic expectations are identical to rational expectations in partial equilibrium, because the only difference to rational expectations is in $\bar{\tau}_{\pi_{t+1}}$. When we scale the shock to ensure a unit response of $\bar{E}_t(\pi_{t+1})$, we therefore cancel out any effects of diagnosticity.

However, in general equilibrium diagnostic expectations deliver radically different results from the other models and the from the RCT results. The over-reaction to news generated by diagnostic expectations makes general equilibrium feedback effects more powerful, as any increase in π_t is met with a further round of over-reaction, entailing even more π_t response, creating powerful strategic complementarities in price-setting. As price stickiness falls from $\lambda = 1$, general equilibrium effects grow in all models. The extra complementarity induced by diagnostic expectations means that as λ falls the causal effects of expectations grow more rapidly with diagnostic expectations than they do with rational expectations. At sufficiently low λ, the strategic complementarities become so strong that an increase in $\bar{E}_t(\pi_t)$ generates a greater than one-for-one response in π_t, implying that in equilibrium greater $\bar{E}_t(\pi_t)$ must in fact be met with a decrease in π_t.

General equilibrium causal effect with and without heterogeneity. Figure 10 repeats the black (RCT-calibration) line from Figure 1b, showing the overall general-equilibrium effect of expected one-year-ahead inflation on realized current inflation. The blue dashed line alongside it shows the same effect if there is no narrative heterogeneity channel in the model, computed using the parameters from the right panel of Table 4.

Figure 10: General equilibrium effects of a shock to $s_{\pi_{t+1}}$, decomposition

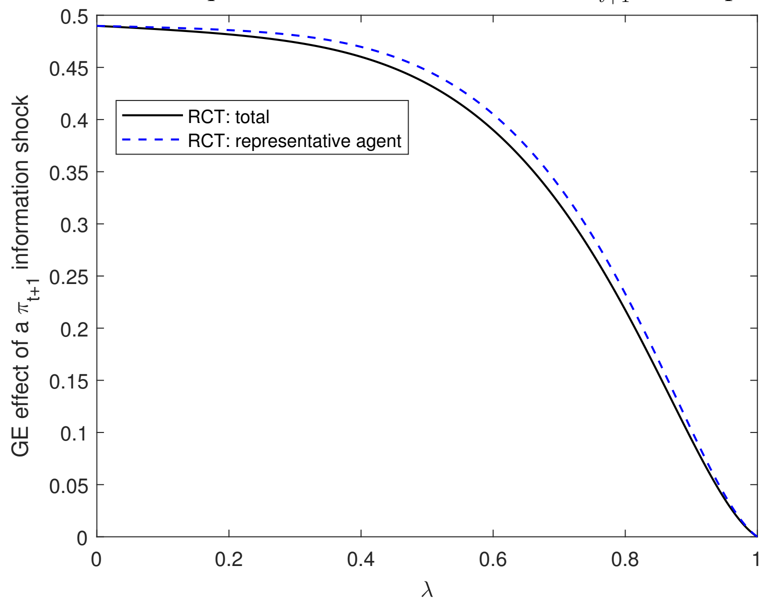

Note: Figure plots (in black) the general-equilibrium causal effects of an increase in $\bar{E}_t\pi_{t+1}$ in the Calvo model, following the formulae in Propositions 2 and 3, and a calibration derived from the estimates in Baumann et al. (2024). The blue line plots the equivalent effect if we assume there is no correlation between $\boldsymbol{\chi}_i$ and $\boldsymbol{\tau}_i$ (detail in Appendix C).

D Empirical Appendix for Section 5.1

D.1 Defining the direct information indicator in the IAS

The full set of questions used to construct the information dummy is set out below, along with the dates at which each was asked and how the answers are mapped into the information indicator used above. Note that my question numbering differs from the labels in the IAS microdata, to aid the logical organization of the paper. All of the questions were only asked in the first quarter of the year(s) indicated. In the main exercises I exclude questions 2e and 2g from the total information variable, to ensure that there are no periods in which two questions are asked. I remove these rather than the short run questions in those periods to keep the majority of questions as short run expectations. The results are robust to including these extra questions. See Appendix D.3 for this, and robustness checks with other variations in the definition of the information indicator.

Question 2a *What were the most important factors in getting to your expectation for how prices in the shops would change over the next 12 months?*

Please select up to 4.

1. *How prices have changed in the shops recently, over the last 12 months*
2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*

3. *Reports of current inflation in the media*
4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *Other factors*
10. *None*

We can divide the possible answers into four categories. First, options 1 and 2 concern past experienced price rises. Options 3 and 4 are direct information about inflation. Options 5-8 concern other macroeconomic variables, either current or expected, and options 9 and 10 are extras. A rational household may well use the information sources in options 1,2 and 5-9 to forecast inflation, but in the decomposition in Proposition 6 this would represent cross-learning from information about other variables. To use the level of interest rates (5) to forecast inflation, for example, a household must employ a model of how interest rates relate to inflation. Similarly, to use past experienced price changes (1-2), households need a model of the persistence of inflation.[37] The only answers that represent the use of direct information about inflation are options 3 and 4.

Question 2b *What were the most important factors that led you [to change (insert their response to how expectation has changed)] your expectation of prices in the shops over the next 12 months?*

Please select up to 4:

1. *How prices have changed in the shops recently, over the last 12 months*
2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*
3. *Reports of current inflation in the media*
4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *The level of the exchange rate (the value of sterling)*
10. *Other factors*

[37] Macaulay and Moberly (2025) find this perceived persistence is very heterogeneous across households. Note that strictly, option 3 also concerns past price changes, so the assumption here is that media reports of inflation tend to discuss both current and future inflation simultaneously. Appendix D.3 shows that the results below are robust to various small changes to this definition of the information indicator.

11. *None*

Asked: 2017
Information indicator: =1 if items 3 or 4 selected, =0 otherwise.

Question 2c *What were the most important factors that led you to change/not change your expectation of prices in the shops in the longer term?*

1. *How prices have changed in the shops recently, over the last 12 months*
2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*
3. *Reports of current inflation in the media*
4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *The level of the exchange rate (the value of sterling)*
10. *Other factors*
11. *None*

Asked: 2018, 2019
Information indicator: =1 if items 3 or 4 selected, =0 otherwise.

Question 2d *When you said prices would go up in the next 12 months, how important were the following things in getting to that answer?*

For each option, possible answers are:

- *Very important*
- *Fairly important*
- *Not very important*
- *Not at all important*
- *Don't know*
- *Refused*

Options:

1. *How prices have changed in the shops in your most recent visits (i.e. the last 1 to 6 months).*
2. *How prices have changed in the shops over the longer term (i.e. the last 12 months or more)*
3. *The current level of interest rates.*
4. *The current strength of the British Economy.*

5. *The inflation target set by the government.*
6. *Reports on inflation outlook in the media.*
7. *Reports of VAT changes in the media.*
8. *Other factor(s).*

Asked: 2009, 2010, 2011, 2013
Information indicator: =1 if 'very important' selected for option 6, =0 otherwise.

Question 2e *And which, if any, of the same factors were important in getting to your expectation of how prices will change over the longer term (say in 5 years time)?*

1. *How prices have changed in the shops in your most recent visits (i.e. the last 1 to 6 months).*
2. *How prices have changed in the shops over the longer term (i.e. the last 12 months or more)*
3. *The current level of interest rates.*
4. *The current strength of the British Economy.*
5. *The inflation target set by the government.*
6. *Reports on inflation outlook in the media.*
7. *Reports of VAT changes in the media.*
8. *Other factor(s).*

Asked: 2011, immediately after Question 2d
Information indicator: =1 if item 6 selected, =0 otherwise.

Question 2f *What were the most important factors in getting to your expectation for how prices in the shops would change over the next 12 months?*

Please select up to 4:

1. *How prices have changed in the shops recently, over the last 12 months*
2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*
3. *Reports of current inflation in the media*
4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *Other factors*
10. *None*

Asked: 2016
Information indicator: =1 if items 3 or 4 selected, =0 otherwise.

Question 2g *And what were the most important factors in getting to your expectation for how prices in the shops would change over the longer term (say in 5 years' time)?*

Please select up to 4:

1. *How prices have changed in the shops recently, over the last 12 months*
2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*
3. *Reports of current inflation in the media*
4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *Other factors*
10. *None*

Asked: 2016
Information indicator: =1 if items 3 or 4 selected, =0 otherwise.

D.2 The relationship of planned consumption with measured information and subjective models

To confirm that the survey measures of information and subjective models uncover meaningful aspects of household beliefs, I consider how they correlate with planned consumption behavior. To this end, I use the following survey question:

Question 3 *Which, if any, of the following actions are you taking, or planning to take, in the light of your expectations of price changes over the next twelve months?*

- *Cut back spending and save more.*

Crucially, this asks about consumption choices which are explicitly driven by expected inflation.[38] A household answering 'yes' to this question, and who reports elsewhere in the survey that they expect prices to rise in the next year, is therefore indicating that $dc_{it}/d\mathbb{E}_{it}p_{t+1} < 0$. A question

[38]Another question in the survey asks if the respondent will "bring forward major purchases such as furniture or electrical goods" as a result of expected inflation. I do not use this for two reasons. First, as Nunes and Park (2020) note, the question refers specifically to durable goods, which may not respond to prices in the same way as aggregate consumption, the object of interest. Second, it is very rarely chosen: just 6% of respondents said they would bring forward major purchases. In contrast, 40% report that they will cut back spending and save more. Any estimation on this variable will therefore be heavily influenced by a small subset of agents.

that only asked about consumption or consumption changes, without reference to the cause of the behavior, would conflate this with reactions to expectations of other variables, which might also be influenced by the same shocks as expected inflation, either directly or through cross-learning. Question 3 is therefore informative about the sign of $\frac{dc_{it}}{d\mathbb{E}_{it}p_{t+1}}$. If current prices are taken as given by the household, then this is the same as the sign of $\tilde{\chi}_i = \frac{dc_{it}}{d\mathbb{E}_{it}\pi_{t+1}}$.

The vast majority of respondents (98%) expect positive inflation over the next 12 months.[39] For these households, yes and no responses to Question 3 respectively indicate that:

$$\frac{dc_{it}}{d\mathbb{E}_{it}p_{t+1}} \begin{cases} < 0 & \text{if answer yes} \\ \geq 0 & \text{if answer no} \end{cases} \tag{175}$$

For the minority who expect deflation, these inequalities are reversed: responding with 'yes' indicates consumption is being cut because of an expected fall in prices. I therefore define the following indicator:

$$\frac{\widetilde{dc_{it}}}{d\mathbb{E}_{it}p_{t+1}} = \begin{cases} 1 & \text{if Q3='no' and } \mathbb{E}_{it}\pi_{t+1} > 0 \\ 0 & \text{if Q3='yes' and } \mathbb{E}_{it}\pi_{t+1} > 0 \\ 1 & \text{if Q3='yes' and } \mathbb{E}_{it}\pi_{t+1} < 0 \\ 0 & \text{if Q3='no' and } \mathbb{E}_{it}\pi_{t+1} < 0 \end{cases} \tag{176}$$

For the large majority who expect inflation, this is equal to 1 if $\frac{dc_{it}}{d\mathbb{E}_{it}p_{t+1}} \geq 0$, and equal to 0 if the reaction to expected price rises is strictly negative. The same is true of the minority who expect deflation, except that any household with $\frac{dc_{it}}{d\mathbb{E}_{it}p_{t+1}} = 0$ would respond 'no' to Question 3, and so is counted as if their response to expected price rises is strictly negative. The mislabeling is not a large issue, as less than 1% of respondents to Question 3 both expect deflation and answer 'no'. The results below are robust to removing the few households who expect deflation (see Table 5 column 2).

Table 5 shows how this is related to the information indicator and the subjective models (responses to Question 1). Column 1 shows the results from estimating a probit regression of $\frac{\widetilde{dc_{it}}}{d\mathbb{E}_{+1}}$ on the information indicator interacted with subjective models (Question 1), plus the standard household controls and time fixed effects used above. The coefficient on information is significantly

[39]The analysis in this section excludes any households who report expecting zero inflation over the next 12 months, or who do not answer the inflation expectation question, as Question 3 is difficult to interpret for these households. I discuss the appropriate counterfactual implicit in the question below. Including these people, 79% of respondents to Question 3 expect positive inflation, 7% expect zero inflation, 2% expect deflation, and 12% do not answer.

negative for those with negative subjective models of inflation, despite the fact that substitution effects imply $\frac{dc_{it}}{d\mathbb{E}_{it}p_{t+1}} \geq 0$ in many standard models. Being informed is therefore associated with a *lower* probability of responding positively to expected inflation for these households.

However, for those who believe inflation makes the economy stronger, being informed is associated with a significantly higher $\Pr(\frac{dc_{it}}{d\mathbb{E}_{it}\pi_{t+1}} \geq 0)$. For those who believe inflation makes no difference, the average value of $\Pr(\frac{dc_{it}}{d\mathbb{E}_{it}\pi_{t+1}} \geq 0)$ with and without information, which is also consistent with the interpretation of these variables as $\frac{\widetilde{dc_{it}}}{d\mathbb{E}_{+1}} = 1$ includes the case where $\frac{dc_{it}}{d\mathbb{E}_{it}\pi_{t+1}} = 0$.

Table 5: Consumption response to inflation correlates with information, by subjective model

	(1)	(2)
	c response to Eπ	c response to Eπ
information indicator=1	-0.213***	-0.224***
	(0.0611)	(0.0613)
end up stronger	0.0108	0.0392
	(0.0891)	(0.0906)
information indicator=1 $\times$ end up stronger	0.348*	0.313*
	(0.185)	(0.186)
make little difference	0.130**	0.157***
	(0.0594)	(0.0600)
information indicator=1 $\times$ make little difference	0.0240	-0.0149
	(0.126)	(0.128)
dont know	0.0958	0.0978
	(0.0833)	(0.0846)
information indicator=1 $\times$ dont know	-0.0158	-0.0342
	(0.186)	(0.187)
Expected Inflation	All	Exclude Deflation
Controls	All	All
Time FE	Yes	Yes
Observations	4940	4871

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the results of probit regressions of the $\widetilde{\frac{dc_{it}}{d\mathbb{E}_{it}\pi_{t+1}}}$ indicator on the information indicator, interacted with responses to Question 1. The omitted category is a household with information indicator=0 who holds the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS.

This is consistent with individuals filtering information through their subjective models of the economy. If a household who believes inflation weakens the economy gets more information about future positive inflation, their subjective model implies that they should cut consumption, because bad times lie ahead. If instead a household believes inflation strengthens the economy, then they will react in the opposite way to the same inflation. The overall correlation of information and consumption response is negative because the majority of households believe inflation makes the economy weaker. This therefore supports the claim that the information indicator and answers to Question 1 reflect the information and subjective models used by households in making their consumption decisions.

The analysis here assumes that when asked whether they will cut back consumption and save more, households are comparing their actions to a counterfactual in which there are no price rises over the next 12 months. An alternative possibility is that they are comparing with a consumption plan made in the past, in which case the relevant counterfactual is where expected inflation is unchanged from the level expected when the plan was made. I consider this in two ways, and find that the qualitative patterns in reported consumption responses to inflation are the same for households expecting inflation to increase or decrease relative to the previous year. It does not therefore appear that past inflation is the relevant counterfactual for most respondents.

First, column 2 of Table 5 re-runs the regression in column 1, excluding any respondent who reports expecting prices to fall over the next year. All results are qualitatively the same as over the full sample, showing that the few respondents expecting deflation are not driving the results.

Second, I split the sample by the sign of the respondent's expected change in inflation, computed as the sign of the difference between 12-month ahead inflation forecast and their perception of inflation over the previous 12 months. The results are in Table 6. The sample sizes in each group are substantially smaller than over the full sample, so some significance is lost, but importantly the signs of the key coefficients remain the same. In each group, households who believe inflation makes the economy weaker are less likely to have $\frac{dc_{it}}{d\mathbb{E}_{it}\pi_{t+1}} \geq 0$ when they get inflation information. For households who believe inflation makes the economy stronger, this effect is reversed. The similarity of these patterns suggests that most respondents use 'no price change' as the counterfactual when answering Question 3, not 'no inflation change'. If the latter was used, we would expect to see changes of sign across the columns in Table 6, as a household expecting a fall in inflation would be reporting $-1 \times \frac{dc_{it}}{d\mathbb{E}_{it}\pi_{t+1}}$, while one expecting a rise in inflation would report $\frac{dc_{it}}{d\mathbb{E}_{it}\pi_{t+1}}$.

Table 6: Consumption response to inflation correlates with information, by subjective model and sign of perceived Eπ change.

	(1) $\text{E}\Delta\pi < 0$	(2) $\text{E}\Delta\pi = 0$	(3) $\text{E}\Delta\pi > 0$
Dc_Dpi			
Information=1	-0.140	-0.305***	-0.257**
	(0.116)	(0.101)	(0.107)
end up stronger	0.0668	-0.178	0.195
	(0.164)	(0.151)	(0.165)
Information=1 × end up stronger	0.586	0.349	0.397
	(0.441)	(0.293)	(0.307)
make little difference	0.165	0.136	0.181
	(0.111)	(0.0957)	(0.112)
Information=1 × make little difference	0.129	-0.300	0.113
	(0.241)	(0.211)	(0.216)
dont know	0.156	0.0293	0.0264
	(0.176)	(0.128)	(0.167)
Information=1 × dont know	-0.141	0.469	0.117
	(0.354)	(0.359)	(0.325)
Controls	All	All	All
Time FE	Yes	Yes	Yes
Observations	1384	1876	1463

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the results of probit regressions of the $\frac{\widetilde{dc_{it}}}{d\text{E}_{it}\pi_{t+1}}$ indicator on the information indicator, interacted with responses to Question 1, split by the sign of the respondent's inflation expectations. The omitted category in all cases is a household with information indicator=0 who holds the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS.

D.3 Cross-sectional patterns in information on inflation

Columns 1-3 of Table 7 show the results of probit regressions of the information indicator on subjective models, controls, and period fixed-effects, for three subsamples. The first only uses questions about the information used to arrive at the respondent's *change* in expected inflation, and the second uses only questions about information used to form point forecasts. The third column excludes questions relating to forecast horizons longer than 12 months. The signs of the marginal effects are the same as in the main exercise in Table 1, though they are not significant in the case of the revisions questions, as the sample size is small.

Table 7: Information correlates with subjective models, split by information question type

	(1) Revision	(2) Point	(3) Short horz.	(4) Extra Qs	(5) Q2d wider	(6) +Other
end up stronger	0.0575	-0.0335	-0.0123	0.00114	-0.00126	-0.00779
	(0.0380)	(0.0218)	(0.0206)	(0.0196)	(0.0196)	(0.0205)
make little difference	-0.0191	-0.0331**	-0.0392***	-0.0310**	-0.0312**	-0.0429***
	(0.0233)	(0.0155)	(0.0141)	(0.0132)	(0.0131)	(0.0139)
dont know	-0.0408	-0.0715***	-0.0622***	-0.0663***	-0.0472***	-0.0663***
	(0.0297)	(0.0206)	(0.0192)	(0.0174)	(0.0180)	(0.0191)
Controls	All	All	All	All	All	All
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2364	5906	6848	8306	8270	8270

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the average marginal effects from estimating probit regressions of the information indicators constructed from subsets of the questions listed in Appendix D.1 on the responses to Question 1. The omitted category is the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS.

The remaining columns of Table 7 repeat the regression for broader definitions of the information dummy than that used in Table 1. In the fourth column, the information indicator includes Questions 2e and 2g. In the fifth column, I extend the criteria for setting the information indicator equal to 1 in Question 2d to account for the fact that some people may be unwilling to select the highest importance box for any information source. I therefore set the information indicator to 1 if in answer to Question 2d, the respondent selects 'very important' for direct inflation information (as before), or if they do not select 'very important' for any option, but do respond that four or fewer options were 'fairly important', and direct inflation information is among them. In the final column, I set the information indicator =1 if the household chooses a direct information source or 'Other', in case this includes direct information sources (e.g. checking the Bank of England published forecasts). In all of these, the results are robust.

To account for possible selection bias from missing observations, I estimate a version of Table 1 amended for selection as in Heckman (1979). As in Michelacci and Paciello (2024),

Table 8: Information correlates with subjective models, with selection correction

	(1) Information	(2) Information
end up stronger	-0.0178	-0.0183
	(0.0172)	(0.0172)
make little difference	-0.0306**	-0.0315***
	(0.0120)	(0.0121)
dont know	-0.0575***	-0.0579***
	(0.0172)	(0.0173)
Inverse Mills ratio	-0.282***	-0.0696**
	(0.0820)	(0.0355)
Selection stage		
Economic Literacy	0.205***	
	(0.0226)	
HH does not know past π		-0.876***
		(0.0365)
r affects π in 1-2 months		0.0334
		(0.0236)
r affects π in 1-2 yrs		0.0882***
		(0.0231)
Pseudo-R^2 (selection)	0.103	0.127
Controls	All	All
Time FE	Yes	Yes
Observations	18026	18026

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the coefficients from estimating a linear regression of the information indicator defined in Section 5.1 on the responses to Question 1, augmented with the inverse Mills ratio from a first-stage probit regression of whether the information indicator is observed on measures of economic literacy defined above. The omitted category is the belief that inflation makes the economy weaker. The selection stage is only run for quarters in which the information questions were asked. The model is run using the 2-step limited information method in Heckman (1979). Time fixed effects and controls as in Table 1 are included in both stages.

I predict observing the relevant survey response using a measure of economic literacy. Here the relevant response is only the information indicator, as there are no missing values for Question 1. Following Michelacci and Paciello (2024), economic literacy is measured with three indicators: the household reports a value for perceived current inflation, and they answer 'agree' or 'strongly agree' to the statements "a rise in interest rates makes prices in the high street rise more slowly in the short term (say a month or two)" and "a rise in interest rates makes prices rise more slowly in the medium term (say a year or two)". I estimate versions of the model with this as an aggregate index (=1 if and only if the household scores on all components), and with the components disaggregated. The results are in Table 8. The predictors used in the first stage are strongly significant. Qualitatively the second-stage results are unchanged from Table 1, and the quantitative differences are small.

D.4 Time series patterns in subjective models of inflation

Bhandari et al. (2024) also study the time series of responses to Question 1, and conclude that households are more pessimistic about inflation when output growth is low. To explore this, I regress the proportion of households responding 'end up weaker' on realized annual CPI inflation and quarterly GDP growth. The results are in column 2 of Table 9. Consistent with Bhandari et al. (2024), the coefficient on GDP growth is significantly negative. However, the R^2 is only slightly higher than that of a regression on inflation only (column 1), so GDP growth does not account for much of the variation in survey answers. Indeed, GDP growth does not have any significant relationship with the proportion of households with a negative view of inflation outside of the four worst months of the Great Recession (column 3).

Table 9: Regressions of the proportion of households answering weaker to Question 1 on aggregate variables.

	(1) Proportion weaker	(2) Proportion weaker	(3) Proportion weaker
Inflation	0.0568***	0.0517***	0.0501***
	(0.00489)	(0.00479)	(0.00469)
GDP growth		-0.0261***	-0.0110
		(0.00869)	(0.0180)
Constant	0.466***	0.487***	0.482***
	(0.0109)	(0.0123)	(0.0152)
Omitted quarters	None	None	2008Q2-2009Q1
R-squared	0.615	0.647	0.554
Observations	70	70	66

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the results of regressing the proportion of households answering Question 1 that inflation makes the economy weaker on annual CPI inflation and quarter-on-quarter real GDP growth. Proportions are computed using survey weights.

To explore which measure of inflation affects subjective models, Table 10 reports the results of regressing an indicator variable for if the respondent reports a negative subjective model on a variety of inflation measures. The first column uses CPI inflation, so is very similar to the time-series regression in Table 9. Columns 2-4 use more granular measures of the inflation rate experienced by different households, split by whether they are above retirement age (65), above median income, and by their housing tenure. Inflation rates split by these characteristics are provided by the ONS.[40] Column 5 uses perceived current inflation. The different realized inflation measures are strongly correlated, so cannot be included jointly. Although the coefficient sizes vary as the different inflation rates have different levels of volatility, in all cases higher inflation is associated with a significantly greater probability of reporting a negative subjective model. The R^2 is highest for perceived inflation, supporting the view that inflation information affects subjective models.

Table 10: Probability of reporting negative subjective model by experienced and perceived inflation

	(1)	(2)	(3)	(4)	(5)
	Weaker	Weaker	Weaker	Weaker	Weaker
Inflation	0.0510***	0.0463***	0.0457***	0.0292***	0.0254***
	(0.00177)	(0.00170)	(0.00165)	(0.00137)	(0.000720)
Inflation measure	CPI	by retirement	by income	by housing	perceived
Controls	All	All	All	All	All
R-squared	0.0303	0.0286	0.0292	0.0237	0.0371
Observations	68269	68269	68269	68269	68269

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the results of estimating a linear probability model of whether a respondent reports that inflation makes the economy weaker in response to Question 1 on various measures of inflation. These are annual CPI inflation, inflation split by whether the respondent is of retirement age, split by whether the respondent has above or below median income, split by the respondent's housing tenure, and finally the respondent's perceived current inflation. Sample begins in 2006 Q1, as this is when the ONS sub-group inflation data is available from. Households not reporting a perceived rate of inflation are dropped in all regressions. All regressions are weighted using the survey weights provided in the IAS.

[40] Finer decompositions of inflation by household characteristics are not reliable, given the data available for the UK (see e.g. Dawber et al., 2022).

Table 11: Regressions of the proportion of households giving each answer to Question 1 on aggregate variables.

	(1) Proportion	(2) Proportion	(3) Proportion
Stronger			
Inflation	-0.0123***	-0.0116***	-0.0108***
	(0.00193)	(0.00215)	(0.00221)
GDP growth		0.00346	-0.00392
		(0.00363)	(0.00646)
Constant	0.104***	0.102***	0.104***
	(0.00431)	(0.00550)	(0.00638)
No difference			
Inflation	-0.0292***	-0.0262***	-0.0257***
	(0.00303)	(0.00313)	(0.00314)
GDP growth		0.0150***	0.0106
		(0.00473)	(0.0107)
Constant	0.277***	0.264***	0.266***
	(0.00772)	(0.00883)	(0.0103)
Don't know			
Inflation	-0.0154***	-0.0139***	-0.0135***
	(0.00249)	(0.00262)	(0.00267)
GDP growth		0.00762*	0.00428
		(0.00423)	(0.00987)
Constant	0.153***	0.147***	0.148***
	(0.00687)	(0.00757)	(0.00884)
Omitted quarters	None	None	2008Q2-2009Q1
Observations	70	70	66

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the results of regressing the proportion of households giving each answer to Question 1 (except 'weaker') on annual CPI inflation and quarter-on-quarter real GDP growth. Proportions are computed using survey weights. The R^2s of the core regressions in column 1 are 0.388 (stronger), 0.534 (no difference), and 0.355 (don't know).

Similar patterns in reverse are observed for the other answers. Table 11 repeats the regressions of Table 9, replacing the dependent variable with the proportion of respondents choosing each of the other answers to Question 1. In all cases, inflation accounts for a large share of the variation in survey answers, and higher inflation is associated with significantly lower proportions giving each answer. Higher GDP growth is associated with higher proportions on these other answers, but that relationship is not significantly different from zero for any answer when excluding the worst of the Great Recession.

To test if the distribution of beliefs about inflation shifts when the economy reaches the ZLB, I estimate an ordered probit regression of subjective models of inflation in the zero lower bound period, and a variety of controls.[41] A response that inflation makes the economy stronger is coded as the highest value, and inflation makes the economy weaker is the lowest value (I exclude the 'don't know' answers). A positive coefficient on the zero lower bound period would therefore imply a shift towards believing inflation makes the economy stronger, as we would expect if households follow a standard New Keynesian model. This is not what the results in Table 12 show: there is no significant shift towards a positive view of inflation in the ZLB period.

Table 12: Ordered probit regressions of subjective models of inflation on whether the economy is at the zero lower bound on nominal interest rates.

	(1) Subjective model	(2) Subjective model	(3) Subjective model
Subjective model			
ZLB	-0.00801	-0.00785	-0.00513
	(0.00937)	(0.00962)	(0.00972)
Controls	None	Household	Household + macro
Observations	83526	83526	83526

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the results of an ordered probit regression of answers to Question 1 on an indicator for whether the UK economy was at the zero lower bound, defined as the period from 2009Q2 to the end of 2019 (end of the sample). The ordering is: "stronger", "no difference", "weaker". Those answering "no idea" are omitted. All regressions are weighted using the survey weights provided in the IAS.

D.5 Perceived and expected inflation across households

To account for possible selection bias from missing observations, I estimate a version of Table 2 amended for selection as in Heckman (1979). As in Appendix D.3, I predict observing perceived and expected inflation using the components of the economic literacy indicator in Michelacci and Paciello (2024), this time removing the component concerning whether perceived inflation is reported as this is closely related to the dependent variables of the regressions. The results are in Table 13. The predictors used in the first stage are strongly significant. Qualitatively the second-stage results are unchanged from Table 2, and the quantitative differences are small.

[41] The first column of Table 12 has no controls, the second includes the set of household-level covariates used throughout the paper, and the third adds inflation and GDP growth.

Table 13: Perceived and expected inflation are higher for those with more negative subjective models, with selection correction

	(1) Perceived inflation	(2) Expected inflation
end up stronger	-0.724***	-0.607***
	(0.0319)	(0.0301)
make little difference	-0.548***	-0.478***
	(0.0213)	(0.0201)
dont know	-0.452***	-0.407***
	(0.0292)	(0.0280)
Inverse Mills ratio	0.714***	0.168
	(0.248)	(0.191)
Selection stage		
r affects π in 1-2 months	0.133***	0.192***
	(0.0209)	(0.0209)
r affects π in 1-2 yrs	0.278***	0.310***
	(0.0205)	(0.0205)
Pseudo-R^2 (selection)	0.0477	0.0536
Controls	All	All
Time FE	Yes	Yes
Observations	95339	95339

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports the coefficients from estimating a linear regression of perceived and expected inflation on the responses to Question 1, augmented with the inverse Mills ratio from a first-stage probit regression of whether the dependent variable is observed on measures of economic literacy. The omitted category is the belief that inflation makes the economy weaker. The model is run using the 2-step limited information method in Heckman (1979). Time fixed effects and controls as in Table 1 are included in both stages.

E New Keynesian model further details

E.1 Proof of Proposition 9

The household maximizes (38) subject to the budget constraint (39). With $Y = B = 1$ normalized in steady state, goods market clearing $Y_t = C_t$ implies $C = 1$ in steady state as well.

The first-order conditions imply

$$C_t^{-\frac{1}{\sigma}} = \beta \bar{E}_t \left[\frac{R_t}{\Pi_{t+1}} C_{t+1}^{-\frac{1}{\sigma}} \right] + \xi B_t^{-\eta}. \tag{177}$$

In steady state ($C = B = 1$, $\Pi = 1$), this yields

$$1 = \beta R + \xi, \qquad \text{so} \quad R = \frac{1-\xi}{\beta}. \tag{178}$$

Log-linearizing (177) around the steady state gives

$$c_t = (1-\xi)\left[\bar{E}_t c_{t+1} - \sigma(r_t - \bar{E}_t \pi_{t+1})\right] + \sigma\eta\xi\, b_t. \tag{179}$$

Log-linearizing the budget constraint (39) around $C = B = Y = 1$ gives

$$c_t + b_t = R(r_{t-1} - \bar{E}_t \pi_t + b_{t-1}) + y_t, \tag{180}$$

where y_t denotes log-deviation of disposable household income (output) from steady state, and $\bar{E}_t \pi_t$ appears because households observe a noisy signal of current inflation.[42]

Recasting the system. Define the exogenous "cash-on-hand" forcing variable as

$$w_t \equiv y_t + R(r_{t-1} - \bar{E}_t \pi_t), \tag{181}$$

so that the budget constraint (180) becomes $b_t = -c_t + R\, b_{t-1} + w_t$. Substituting this into (179) and rearranging, we obtain:

$$\bar{E}_t c_{t+1} = \frac{1+\sigma\eta\xi}{1-\xi}\, c_t - \frac{\sigma\eta\xi\, R}{1-\xi}\, b_{t-1} + \sigma(r_t - \bar{E}_t \pi_{t+1}) - \frac{\sigma\eta\xi}{1-\xi}\, w_t. \tag{182}$$

Together with the budget constraint, this yields the system

$$\bar{E}_t \begin{pmatrix} c_{t+1} \\ b_t \end{pmatrix} = \underbrace{\begin{pmatrix} \frac{1+\sigma\eta\xi}{1-\xi} & -\frac{\sigma\eta\xi\, R}{1-\xi} \\ -1 & R \end{pmatrix}}_{\equiv \mathbf{A}} \begin{pmatrix} c_t \\ b_{t-1} \end{pmatrix} + \begin{pmatrix} f_{1,t} \\ f_{2,t} \end{pmatrix}, \tag{183}$$

where $f_{1,t} = \sigma(r_t - \bar{E}_t \pi_{t+1}) - \frac{\sigma\eta\xi}{1-\xi}\, w_t$ and $f_{2,t} = w_t$. This system has one forward-looking (jump) variable, c_t, and one predetermined (state) variable, b_{t-1}. The determinant and trace of $\mathbf{A}$ are

$$\det(\mathbf{A}) = \frac{R}{1-\xi} = \frac{1}{\beta}, \qquad \operatorname{tr}(\mathbf{A}) = \frac{1+\sigma\eta\xi}{1-\xi} + R. \tag{184}$$

[42] By certainty equivalence in this linear-quadratic environment, the optimal consumption rule has the same form whether or not the household perfectly observes π_t. We simply replace π_t with $\bar{E}_t \pi_t$.

Let λ_1, λ_2 denote the eigenvalues, satisfying

$$\lambda^2 - \mathrm{tr}(\mathbf{A})\,\lambda + \frac{1}{\beta} = 0. \tag{185}$$

For a unique saddle-path solution we require one eigenvalue inside and one outside the unit circle (Blanchard and Kahn, 1980). Evaluating the characteristic polynomial at $\lambda = 1$:

$$p(1) = 1 - \mathrm{tr}(\mathbf{A}) + \frac{1}{\beta} = \xi\left(\frac{1}{\beta} - \frac{1+\sigma\eta}{1-\xi}\right). \tag{186}$$

Saddle-path stability requires $p(1) < 0$, i.e.

$$\sigma\eta > R - 1 = \frac{1-\xi-\beta}{\beta}, \tag{187}$$

which is satisfied for all calibrations in this paper. Under this condition, $0 < \lambda_1 < 1 < \lambda_2$.

Saddle-path solution. Conjecture that the solution takes the form $c_t = \mu\, b_{t-1} + g_t$, where μ is a constant and g_t depends only on current and expected future exogenous variables.

Substituting the conjecture into the budget constraint gives $b_t = (R - \mu)\, b_{t-1} + (w_t - g_t)$, so the stable eigenvalue governs the dynamics of b_t on the saddle path:

$$\mu = R - \lambda_1. \tag{188}$$

Substituting the conjecture into (182) and matching coefficients on b_{t-1} yields the quadratic

$$(1-\xi)\,\mu^2 - \big[(1-\xi)R - 1 - \sigma\eta\xi\big]\,\mu - \sigma\eta\xi\, R = 0, \tag{189}$$

and the saddle-path root is the one satisfying $R - \mu \in (0, 1)$, i.e. $\mu = R - \lambda_1$ where λ_1 is the smaller root of (185).

Matching the remaining (non-b_{t-1}) terms gives

$$\bar{E}_t g_{t+1} = \lambda_2\, g_t + f_{1,t} - \mu\, w_t, \tag{190}$$

where the coefficient on g_t equals λ_2 (verified by $\frac{1+\sigma\eta\xi}{1-\xi} + \mu = \mathrm{tr}(\mathbf{A}) - R + R - \lambda_1 = \lambda_2$). Since

$\lambda_2 > 1$, the unique bounded solution is obtained by solving forward:

$$g_t = \sum_{s=0}^{\infty} \lambda_2^{-(s+1)} \bar{E}_t \left[\mu w_{t+s} - f_{1,t+s}\right]. \tag{191}$$

Expanding notation. Define

$$\zeta \equiv \mu + \frac{\sigma\eta\xi}{1-\xi}. \tag{192}$$

Then

$$\mu w_{t+s} - f_{1,t+s} = \zeta w_{t+s} - \sigma(r_{t+s} - \bar{E}_{t+s}\pi_{t+s+1}). \tag{193}$$

Define $\Lambda \equiv \lambda_2^{-1} < 1$ as the effective discount factor for the consumption function. Substituting the definition of w_{t+s} from (181) into (191) gives:

$$g_t = \sum_{s=0}^{\infty} \Lambda^{s+1} \bar{E}_t \Big[\zeta\big(y_{t+s} + R(r_{t+s-1} - \pi_{t+s})\big) - \sigma(r_{t+s} - \pi_{t+s+1})\Big]. \tag{194}$$

The $s = 0$ term in the summation contains $r_{t-1} - \bar{E}_t\pi_t$, which involves the predetermined nominal rate r_{t-1}. To collect interest-rate terms, note that

$$\sum_{s=0}^{\infty} \Lambda^{s+1} \zeta R\, \bar{E}_t(r_{t+s-1} - \pi_{t+s}) = \zeta R\Lambda\,(r_{t-1} - \bar{E}_t\pi_t) + \zeta R\Lambda^2 \sum_{s=0}^{\infty} \Lambda^s\, \bar{E}_t(r_{t+s} - \pi_{t+s+1}),$$

where the second equality follows from re-indexing the $s \geq 1$ terms. Combining this with the $-\sigma\Lambda \sum \Lambda^s\, \bar{E}_t(r_{t+s} - \pi_{t+s+1})$ term, the coefficient on the discounted sum of expected real interest rates becomes $\zeta R\Lambda^2 - \sigma\Lambda = -\Lambda(\sigma - \zeta R\Lambda)$.

Combining $c_t = \mu\, b_{t-1} + g_t$ and collecting terms yields:

$$c_t = \mu b_{t-1} + \zeta R\Lambda(r_{t-1} - \bar{E}_t\pi_t) + \Lambda \sum_{s=0}^{\infty} \Lambda^s \left[\zeta \bar{E}_t y_{t+s} - (\sigma - \zeta R\Lambda)(\bar{E}_t r_{t+s} - \bar{E}_t\pi_{t+s+1})\right], \tag{195}$$

which is equation (40) in the main text.

E.2 Model description

Households. A representative household chooses consumption C_t, labor supply N_t, and real one-period bond holdings B_t to maximize equation (38) subject to (39). Combining first-order conditions and log-linearizing around the steady state with $Y = C = B = \Pi = 1$ yields:

$$c_t = (1 - \xi)\left[\bar{E}_t c_{t+1} - \sigma(r_t - \bar{E}_t \pi_{t+1})\right] + \sigma\eta\xi\, b_t \tag{196}$$

$$\psi\, n_t = w_t - \frac{1}{\sigma}\, c_t \tag{197}$$

where lower-case letters denote log-deviations from steady state. The derivation of the resulting consumption function (equation (40) in the main text) is presented in Appendix E.1.

Firms. A perfectly competitive final goods producer combines a continuum of intermediate goods varieties $j \in [0, 1]$ using a CES production function:

$$Y_t = \left(\int_0^1 Y_{jt}^{\frac{\varepsilon-1}{\varepsilon}}\, dj\right)^{\frac{\varepsilon}{\varepsilon-1}} \tag{198}$$

where $\varepsilon > 1$ is the elasticity of substitution. Cost minimization yields the demand function for variety j:

$$Y_{jt} = \left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon} Y_t \tag{199}$$

and the aggregate price index $P_t = \left(\int_0^1 P_{jt}^{1-\varepsilon}\, dj\right)^{\frac{1}{1-\varepsilon}}$.

Each intermediate goods firm j produces using a linear production function:

$$Y_{jt} = AN_{jt}, \tag{200}$$

where aggregate productivity A is constant and common across firms. Cost minimization implies that real marginal cost is equal to the real wage for all firms:

$$MC_t = \frac{W_t}{A}. \tag{201}$$

Firms set prices subject to quadratic Rotemberg (1982) adjustment costs. Following Bilbiie (2024), the adjustment cost is paid relative to last period's *aggregate* price level P_{t-1}, rather than the firm's own lagged price. This simplification makes the pricing problem static, allowing us to focus on the core demand mechanisms of interest.

Specifically, firm j chooses P_{jt} to maximize current-period profits net of adjustment costs:

$$\frac{P_{jt}}{P_t} Y_{jt} - W_t N_{jt} - \frac{\theta}{2}\left(\frac{P_{jt}}{P_{t-1}} - 1\right)^2 Y_t \tag{202}$$

subject to the demand function (199) and the production function (200). Here $\theta > 0$ governs the magnitude of the price adjustment cost. The first-order condition for P_{jt} is:

$$(1-\varepsilon)\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon}\frac{Y_t}{P_t}+\varepsilon\frac{W_t}{A}\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon-1}\frac{Y_t}{P_t}-\theta\left(\frac{P_{jt}}{P_{t-1}}-1\right)\frac{Y_t}{P_{t-1}}=0. \tag{203}$$

In a symmetric equilibrium ($P_{jt} = P_t$ for all j), this simplifies to:

$$\Pi_t(\Pi_t - 1) = \frac{\varepsilon}{\theta}\left(MC_t - \mathcal{M}^{-1}\right). \tag{204}$$

where $\mathcal{M} \equiv \varepsilon/(\varepsilon - 1)$ denotes the flexible-proce markup.

Log-linearizing, and adding an exogenous AR(1) markup shock $\nu_{\pi t}$, we obtain:

$$\pi_t = \frac{\varepsilon - 1}{\theta} mc_t + \nu_{\pi t} \tag{205}$$

From (201), (197), (200), and (43), real marginal costs in terms of output are:

$$mc_t = w_t = \psi\, n_t + \frac{1}{\sigma} c_t = \left(\psi + \frac{1}{\sigma}\right) y_t \tag{206}$$

Substituting into (205) yields the Phillips curve in the main text ((41)), with

$$\kappa \equiv \frac{\varepsilon - 1}{\theta}\left(\psi + \frac{1}{\sigma}\right). \tag{207}$$

The markup shock follows $\nu_{\pi t} = \rho_\pi\, \nu_{\pi,t-1} + \epsilon_{\pi t}$, with $\epsilon_{\pi t} \sim \mathcal{N}(0, \sigma_\pi^2)$.

Monetary policy. The central bank sets the nominal interest rate according to the standard Taylor rule (42), where ν_{rt} is an AR(1) monetary policy shock: $\nu_{rt} = \rho_r\, \nu_{r,t-1} + \epsilon_{rt}$, with $\epsilon_{rt} \sim \mathcal{N}(0, \sigma_r^2)$.

Fiscal policy. The fiscal authority issues a constant real supply of bonds $B = 1$ and levies lump-sum taxes to service its debt. The government budget constraint in levels is:

$$T_t = \frac{R_{t-1}}{\Pi_t}B - B = \frac{R_{t-1}}{\Pi_t} - 1 \tag{208}$$

where the second equality uses $B = 1$. That is, the government levies taxes equal to the net real interest payments on its outstanding debt. Since bond supply is constant, in equilibrium $b_t = 0$ for all t, so we do not need to track b_t as a state variable.[43] Log-linearizing (208) yields:

$$t_t = \frac{R}{R-1}(r_{t-1} - \pi_t) \tag{209}$$

Market clearing. The goods market clears when aggregate output equals consumption (43). The labor market clears when aggregate labor demand from firms $\int_0^1 N_{jt}dj$ equals labor supply N_t. In symmetric equilibrium with linear production, this is equivalent to $Y_t = AN_t$, i.e. $y_t = n_t$. The bond market clears by Walras' law given the household and government budget constraints.

E.3 Temporary-information equilibrium

Equilibrium definition. For given distributions of $\tau_{i\pi_t}, \tilde{\chi}_i$, and given exogenous shocks $v_{rt}, v_{\pi t}, \varepsilon_{\pi t}$, the temporary-information equilibrium consists of $\{c_t, r_t, \pi_t, y_t\}$ and signals s_{it} such that:

1. *Households:* households choose c_t to maximize expected lifetime utility (following (40)).
2. *Firms:* firms set prices to maximize expected lifetime profits, implying π_t follows (41).
3. *Monetary Policy:* policymakers choose the nominal interest rate r_t according to (42).
4. *Market Clearing:* the goods and labor markets clear (y_t satisfies equation (43), and $y_t = n_t$).
5. *Signals:* inflation signals are determined as a function of realized inflation and idiosyncratic and aggregate noise according to (44).

Causal effects. To find the causal effects of inflation expectations on y_t in partial equilibrium, we hold π_t, r_t, and signals fixed, except for the public noise in $s_{it}(\pi_{t+1})$. This isolates the pure demand impulse from a change in expected inflation. We then find:

$$\Omega_y^{PE}(\pi_t) = -\frac{\zeta R \Lambda}{1-\zeta\Lambda} + \int \frac{\tilde{\chi}_i \tau_{i\pi_t}}{(1-\zeta\Lambda)\bar{\tau}_{\pi_t}} di \tag{210}$$

where following the notation of Section 3 I use $\Omega_y(\pi_t)$ to denote the causal effect of expectations of π_t on y_t, rescaled such that aggregate expectations move by 1 unit.

In general equilibrium, π_t and r_t respond endogenously, and the change in π_t also feeds back

[43] In the log-linearized consumption function (40), the $\mu\, b_{t-1}$ term therefore drops out in equilibrium, though the coefficients μ, ζ, and Λ derived in Appendix E.1 still enter the other terms.

into signals.[44]

$$\Omega_y^{GE}(\pi_t) = \Omega_y^{PE}(\pi_t) \cdot \frac{1-\zeta\Lambda}{1-\zeta\Lambda+\Lambda(\sigma-\zeta R\Lambda)(\phi_\pi\kappa+\phi_y)} \tag{211}$$

E.4 Calibration

Table 14 gives a full list of the calibrated parameters. Steady state A is omitted, as it is simply adjusted to generate the normalization $Y = 1$.

Table 14: Calibration for Section 6

Parameter	Value	Source	Parameter	Value	Source
β	0.905	Target $R = 1.01$	ϕ_π	1.5	Standard
η	2.5	Kaplan and Violante (2018)	ϕ_y	0.125	Standard
λ	0.8	Standard	χ_+	7.245	Inverse NKPC slope
ξ	0.086	Target $MPC = 0.35$	χ_-	-4.547	Target $\tilde{\chi}_i^2$
ρ	0.441	AR(1) regression	ψ	1	Standard
σ	1	Standard	p_+	0.079	IAS
τ_H	0.539	Target average Kalman gain estimate	p_-	0.588	IAS
τ_L	0	Normalization			

The Euler equation in steady state (178) implies a target for R is not sufficient to pin down β independently of ξ. I therefore also target the MPC out of a change in y_t, which is increasing in ξ. I target an impact MPC of 0.35, the mean of the empirical studies surveyed in Sokolova (2023). These targets imply a substantially lower β than typically used in models without bonds in utility, consistent with the findings in e.g. Auclert et al. (2024). These parameters imply $\mu = 0.353, \zeta = 0.589, \Lambda = 0.594$.

For the slope of the Phillips curve, I make use of the well-known equivalence between Rotemberg- and Calvo-derived New Keynesian Phillips curves (NKPCs) up to first order (Ascari and Rossi, 2012). These two models give the same NKPC if

$$\frac{\varepsilon-1}{\theta} = \frac{(1-\lambda)(1-\beta\lambda)}{\lambda} \tag{212}$$

where λ, as in Section 2, is the probability a firm cannot reset its price each period. I choose $\lambda = 0.8$, which is a standard value in the literature consistent with an average price duration of 5 quarters.

[44]In this case, this is the same as the effect that would be obtained if signals were held constant, because the rescaling to give a unit change in $\bar{E}_t\pi_t$ exactly offsets the signal-feedback channel, as in the causal effects of inflation perceptions in the Calvo model in Section 2.

For ρ, I estimate an AR(1) model on quarterly UK CPI inflation data from 2001 Q2-2019 Q4. To do this, I take two approaches. In both, the raw data are monthly observations of the CPI All Items Index (not seasonally adjusted, base 2015 = 100; ONS series D7BT). I first aggregate to quarterly frequency by averaging the three monthly index values within each quarter. Quarterly log inflation is then $\pi_t^q = \ln(\text{CPI}_t/\text{CPI}_{t-1})$, where CPI_t is the quarterly average of the monthly index.

Since at the time of writing the UK does not publish officially seasonally adjusted CPI data (Dixon and Michail, 2025), I handle seasonality in two ways. Column 1 of Table 15 works with the unadjusted quarterly log inflation series and includes quarter-of-year fixed effects directly in the AR(1) regression. This absorbs the average seasonal pattern within the regression without pre-adjusting the data.

Column 2 instead applies a classical multiplicative decomposition before estimation. A 2×12 centred moving average of the monthly index estimates the trend-cycle component; dividing through gives seasonal–irregular ratios. Monthly seasonal factors are computed as the mean ratio for each calendar month across all available years, normalized so the twelve factors average to 1. The seasonally adjusted monthly CPI is then the raw index divided by the corresponding factor. These adjusted monthly values are averaged within each quarter before computing log inflation and estimating the AR(1) regression without quarter fixed effects.

Both approaches yield an estimated persistence parameter of $\rho = 0.44$, as reported in Table 15. I use $\rho = 0.441$ (Column 1) in the calibration.

Table 15: AR(1) inflation regressions

	(1) CPI inflation	(2) CPI inflation
L.CPI inflation	0.441***	0.440***
	(0.128)	(0.127)
Quarter FE	Yes	No
R-squared	0.549	0.197
Observations	75	75

Robust standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports estimates from regressing quarterly UK CPI log inflation on its own lag. Sample period: 2001 Q2–2019 Q4. In Column 1, the raw (unadjusted) quarterly series is used and quarter-of-year fixed effects absorb seasonal variation within the regression. In Column 2, the monthly CPI is seasonally adjusted prior to quarterly aggregation using a classical multiplicative decomposition (ratio-to-moving-average method with a 2×12 centred moving average), and no quarter fixed effects are included.

The steady state values of p_+ and p_- are taken from the average proportion of households responding “stronger” and “weaker” to Question 1 across all waves of the IAS.

With χ_+ set to the true inverse NKPC slope κ^{-1}, I set χ_- such that $\tilde{\chi}_i^2$ is the same in both cases.

Formally, this implies

$$\chi_- = -\left(\chi_+ + \frac{2(\sigma - \zeta R\Lambda)(\rho - \phi_\pi)}{\zeta - (\sigma - \zeta R\Lambda)\phi_y}\right) \tag{213}$$

Finally, with τ_L set to 0, I set τ_H such that

$$E_I[\tau_i] \equiv (p_+ + p_-)\tau_H + (1 - p_- - p_+)\tau_L = \bar{\tau}^* \tag{214}$$

where $\bar{\tau}^*$ is an estimate of the average Kalman gain in inflation perceptions across all households and survey waves. To compute this, I take Equation (49) and average across households to give:

$$\bar{E}_t\pi_t = \bar{\tau}_{\pi_t}\pi_t + (1 - \bar{\tau}_{\pi_t})\,\rho\bar{E}_{t-1}\pi_{t-1} + \bar{\tau}_{\pi_t}\varepsilon_{\pi t} \tag{215}$$

where I have used the fact that all households are calibrated to have the same ρ. I therefore use average inflation perceptions in the IAS to estimate:

$$\bar{E}_t\pi_t = \gamma_1\pi_t + \gamma_2\bar{E}_{t-1}\pi_{t-1} + \nu_t \tag{216}$$

by OLS, restricting $\gamma_2 = \rho(1 - \gamma_1)$, where ρ is as in Table 14. The estimated γ_1 therefore gives an estimate of the average Kalman gain across the population. This target is $\bar{\tau}^* = 0.359$.