Price Setting Under Inferential Expectations

Timo Henckel* Gordon Menzies†,*

Daniel J. Zizzo‡,*

* Centre for Applied Macroeconomic Analysis, ANU

† University of Technology Sydney

‡ University of East Anglia

17 June 2009

Work in Progress

Please Do Not Quote or Distribute Without Authors' Prior Permission

Abstract

Time-dependent pricing models such as the well-known Calvo (1983) model are deficient in two key respects: they fail to match disaggregated price data and they do not allow for state-dependency of pricing decisions. We propose a new model of pricing based on a new theory of expectations formation called inferential expectations. To overcome the costs associated with information acquisition and processing, firms adopt simple heuristics. In particular, under inferential expectations firms are assumed to perform hypothesis tests over the variance of real marginal costs. Depending on the size of the shocks and amount of available information, firms may remain inactive for an extended period of time and only review their pricing policy when sufficient evidence warrants it. This approach gives rise to an augmented, state-dependent Phillips curve and sheds light onto the dynamics of disaggregated prices.

JEL classification: E30; E50

Keywords: Price stickiness; expectations; monetary policy

1 Introduction

In New Keynesian economies price stickiness is the mechanism commonly purported to be responsible for generating persistent output fluctuations in response to nominal aggregate demand disturbances. Understanding price behavior is crucial to understanding the sources of business cycle fluctations and the effects of monetary and fiscal policy.

The list of suggestions why nominal prices adjust sluggishly ranges from the presence of fixed adjustment costs (menu costs) and fixed term pricing to irrationality and costly information processing. Most models have in common that, in aggregate, only a fraction of firms adjusts prices at any point in time. Combined with some form of monopoly power so that price exceeds marginal cost, such pricing behavior implies that these firms adjust output in response to nominal aggregate disturbances as opposed to prices.[1]

To allow for price rigidity macroeconomists have relied on a couple of workhorse models, notably Calvo (1983) and, to a lesser extent, Rotemberg (1982). The Calvo model in particular is highly tractable and forms the basis for the New Keynesian Phillips curve. But there is growing dissatisfaction with this approach because (a) it only matches aggregate price data for periods of low, stable inflation and (b) recent evidence about prices at the individual product level largely contradicts the estimates of price stickiness obtained for aggregate price data. Moreover, the Calvo model is a time-dependent model, meaning that a firm's decision to change prices is a function of time only, independent of the state of the economy. Many economists consider this property unappealing.

During the past decade numerous authors have offered alternatives to the Calvo model a discussion of which follows in the literature review below. Yet, although considerable progress has been made, a coherent, tractable model that allows for state-dependency while being consistent with the data at the micro and macro level remains elusive.

In this paper we propose a new model of pricing based on a new theory of expectations formation called inferential expectations. Firms operate in a volatile and uncertain environment and information acquisition and processing is burdensome. Reviewing pricing policy every period in response to each new piece of information is therefore inefficient. To overcome the costs associated with information acquisition and processing, firms adopt simple heuristics. In particular, under inferential expectations firms are assumed to perform hypothesis tests over the variance of real marginal costs. When the true state of nature is close to the maintained hypothesis, e.g. when the shocks to real marginal costs are small, and the amount of information is limited, firms remain inactive and do not review their pricing policy. Only when the shocks are sufficiently large or the available information is so big as to reduce much of the uncertainty surrounding real marginal costs, will firms be persuaded to abandon their null

[1] A highly relevant and as of yet unsettled question is why firms may prefer to keep prices unchanged while allowing output to vary as opposed to the other way around. Presumably, the costs associated with changing prices apply to changing output levels too, and possibly more so. We will address this question in future work.

hypothesis and revise their beliefs accordingly.

This way of modelling firms' pricing decisions has a number of advantages over the existing dominant paradigm. The assumed behavior of 'belief conservatism' is considered by many to be intuitive and a good approximation of how agents with limited cognitive abilities navigate through a complex, volatile environment. The model is state-dependent and therefore avoids some of the pitfalls associated with time-dependent models yet it remains tractable and can serve as a basis for deriving a Phillips curve which, of course, can be estimated and used in a wider modelling context. Finally, it offers rich insight into the dynamics of prices at the micro level, allowing a careful comparison of its implications to the recent data on individual prices.

2 Literature Review

With price stickiness being a cornerstone of New Keynesian economics, the literature analyzing and discussing price stickiness goes back several decades. However, in many applications price stickiness is assumed, not explained, with the main purpose of generating a positively sloped short-run Phillips curve (in price-output space). This practice is typically justified by referring to the existence of quadratic price adjustment costs (Rotemberg, 1982) or small menu costs (Mankiw, 1985) or by assuming exogenously given staggering of prices (Fischer, 1977; Taylor, 1980; Calvo, 1983).

Recently, economists have re-evaluated the existing theories and examined in greater detail the empirical evidence, both on the macro as well as on the micro level. Measures of price stickiness obtained from disaggregated data tend to be considerably smaller than those obtained using aggregated data. This has naturally led to a push for new theories to explain the co-existence of price flexibility at the micro level and price rigidity at the macro level.

Theoretically, the literature is divided into time-dependent models (TDMs) and state-dependent models (SDMs). In TDMs firms change all or some of their prices at specific intervals, independent of the state of the underlying economy. In SDMs firms use information about the state of the economy upon which they base their pricing decisions. Historically, TDMs have been the model of choice as SDMs suffer from two technical difficulties: (a) in most state-dependent pricing models, the entire distribution of prices is a state variable, giving rise to an insurmountable fixed-point problem; and (b) state-dependent price adjustment generates discontinuities and non-differentiabilities, making it very difficult, for example, to obtain Taylor expansions.

2.1 Empirical Evidence on Price Stickiness

Empirical studies based on aggregate data, e.g. estimations using vector autoregressions, typically find considerable persistence in prices.[2] Such evidence has

[2] See, for example, Christiano et al. (1999), Fuhrer and Moore (1995) and Galí and Gertler (1999).

been corroborated by some of the early studies focusing on specific wholesale or retail items in which the unconditional duration of goods' prices is recorded. For example, Carlton (1986), Cecchetti (1986), Kashyap (1995) and Levy et al. (1997) find evidence that prices in the United States are fixed for about one year. Blinder et al. (1998), who surveyed firms using questionnaires, come up with a similar figure.[3]

Recent evidence based on disaggregated data, however, revises these estimates downwards. Bils and Klenow (2004), using data on 350 categories of goods and services collected by the U.S. Bureau of Labor Statistics, report that the median time between price changes is 4.3 months, claiming that sectoral inflation rates are far more volatile than implied by standard sticky-price models.[4]

Boivin et al. (2009), using a factor-augmented VAR that discriminates between macroeconomic and sector-specific shocks, present similar evidence but argue that, "the apparent persistence of aggregate inflation may reflect heterogeneity across sectors or a structural break in the mean inflation during the sample." (p. 351) The reason is that, according to Boivin et al.'s estimates, only 15 percent of the variation in monthly individual prices is attributable to macroeconomic factors.

This contradictory pattern of behavior—relative price rigidity at the aggregate level combined with relative flexibility at the disaggregated level—is a recurring theme in the recent pricing literature and responsible for the confusion surrounding the debate about whether prices are flexible or sticky. However, there is broad agreement that looking at individual price behavior is essential to fully understand an economy's inflation process. The main findings from recent micro studies on prices we address in this paper can be summarized as follows:[5]

1. Prices change frequently (every 4–7 months, depending on whether or not sale prices are included) and the majority of adjustments are large in absolute value (on the order of 10-12%) although for a large subset price adjustments are much smaller (5% or less). (Klenow and Kryvtsov, 2008, Midrigan, 2008, Nakamura and Steinsson, 2008, Wolman, 2007)[6][7]

[3] In the Euro area, according to Rumler and Vilmunen (2005), the average time between price changes is 13 months.

[4] There is considerable variation in the estimates. For goods alone (which comprise 30.4% of the CPI) the median is 3.2 months, for services (40.8% of the CPI) the median is 7.8 months. For individual items the variation is even more pronounced.

[5] For a comprehensive, non-technical review of the empirical literature on price adjustment see Wolman (2007).

[6] Nakamura and Steinsson (2008) emphasize the importance of distinguishing between sale and nonsale price changes by noting, for example, that some types of sales (price reductions) are orthogonal to macroeconomic conditions, and they argue that price changes due to product substitutions are fundamentally different from the regular price changes macroeconomists usually have in mind. For example, the spring and fall clothing seasons in apparel and the new model year for cars are associated with a large number of price changes due to the introduction of new products.

[7] Midrigan's explanation for the price dispersion among individual products comes from extending a simple menu cost model to a multi-product setting in which the size and frequency

2. Approximately, one-third of nonsale price changes are price decreases. (Nakamura and Steinsson, 2008)

3. The fraction of price increases and decreases typically offset each other, so that movements in aggregate inflation largely reflect movements in the size of price changes (the intensive margin) rather than the fraction of items changing price (the extensive margin). (Klenow and Kryvtsov, 2008)

The first two stylized facts roughly give rise to a tri-modal distribution of price changes. One of the key contributions of this paper is to replicate this feature of the price data.

Additional findings from recent studies are listed here for the sake of completeness but in this paper we have little to say about them. They are:

4. The size and timing of price changes vary considerably for a given item, but the size and probability of a price change are unrelated to the time since the last price change, e.g. prices do not appear to exhibit time-dependence. (Klenow and Kryvtsov, 2008, Nakamura and Steinsson, 2008)

5. The frequency of price increases covaries strongly with inflation, whereas the frequency of price decreases and the size of price increases and price decreases do not. (Nakamura and Steinsson, 2008)

6. The frequency of price change is highly seasonal: it is highest in the first quarter and then declines. (Nakamura and Steinsson, 2008)

7. The observed dispersion of price changes is heavily determined by the degree of market power and price changes among tradeables are more frequent than among non-tradeables. (Boivin et al., 2009)

8. Prices in sectors with volatile and persistent idiosyncratic shocks react more rapidly to (aggregate) monetary policy shocks. (Boivin et al., 2009)

9. Correlations between idiosyncratic components of prices and quantities tend to be negative, suggesting that sector-specific shocks are supply shocks or measured price data are subject to considerable sampling error. (Boivin et al., 2009)

10. Price behavior changes with the business cycle; in particular, markups tend to be counter-cyclical.

Time-dependent models and first-generation state-dependent models are unable to account for these empirical observations. For example,

of price changes for the products depend on each other. In his model a two-product firm faces a fixed cost of changing its entire menu of prices but, conditional on paying this cost, can change the prices of each individual good costlessly. In such an economy the aggregate price level is much less flexible in response to monetary shocks, a consequence of the smaller selection effect.

- the Taylor model cannot match the variable length in price spells for individual items;
- the Taylor and Calvo models predict bigger absolute price changes for older prices, when no such pattern exists in the data;
- they also fail to generate enough movement in the fraction of price increases vs. decreases with inflation;
- Dotsey, King and Wolman (1999), an early state-dependent pricing model, produces only small price changes and predicts far too big a role for the extensive margin in inflation movements;
- Golosov and Lucas (2007), another state-dependent model, does not generate enough small price changes;
- of the stylized facts above, the first, second, and fifth facts are consistent with a benchmark menu-cost model (a SDM), whereas the fourth and sixth facts are not.

Klenow and Kryvtsov conclude that only the newer, second generation state-dependent pricing models, such as Dotsey, King and Wolman (2006), Midrigan (2008) and Gertler and Leahy (2008) are able to broadly replicate the observed price facts.

A good theory of pricing should be broadly consistent with the above stylized facts. It will therefore require a minimum degree of disaggregation to capture the behavior of idiosyncratic price changes while allowing for simple aggregation that delivers a recognizable Phillips curve and sizeable rigidity at the aggregate level.

2.2 Time-Dependent Models

The most commonly used time-dependent model is Calvo (1983). Firms' pricing decisions are temporaly fixed, irrespective of the economic environment. As Dotsey et al. (1999) point out, TDMs such as Fischer (1977), Taylor (1980), and Calvo (1983) typically imply short-run monetary non-neutrality, followed by long-run neutrality. Caplin and Leahy (1991) note that TDMs assume that, "between price adjustments firms are not allowed to respond even to extreme changes of circumstances. This makes it difficult to know whether the qualitative effects of money in these models are the result of nominal rigidities per se or of the exogenously imposed pattern of price changes."

Several other authors provide further criticism of TDMs, focusing on their inability to match the facts. Ball (1994) shows that in a world with time-dependent pricing credible disinflations surprisingly lead to booms rather than recessions. Mankiw (2001) claims that shocks to monetary policy tend to have a delayed and gradual effect on inflation, contrary to the prediction of TDMs, while Fuhrer and Moore (1995) and Mankiw and Reis (2002) argue that TDMs cannot explain inflation persistence (or *inflation* stickiness), only price stickiness.

Mankiw and Reis (2002) propose a new model that is supposed to fit the facts better. The essence of their model is, "that information about macroeconomic conditions diffuses slowly through the population. This slow diffusion could arise because of either costs of acquiring information or costs of reoptimization. In either case, although prices are always changing, pricing decisions are not always based on current information." (p. 1296). They refer to their model as one of sticky information instead of sticky prices. Because some price setters are setting prices based on old decisions and old information, the model's dynamic response to monetary shocks closely resembles that of an old Keynesian Phillips curve with backward-looking expectations, even though expectations are rational and credibility matters. The key difference is the timing of expectations—in the standard sticky-price model inflation is co-determined by current expectations of future inflation while in the Mankiw-Reis model of sticky information inflation is co-determined by past expectations of current inflation.

Sticky information models and other proposed alternatives to standard sticky price models have one thing in common: they embrace recent behavioral evidence which suggests that people do not have an infinite capacity to costlessly collect and process information. Agents therefore rationally choose to disregard some information or they base their decisions on computationally inexpensive heuristics or both. Hence, agents are still assumed to be rational but they face informational constraints.

Reis (2006) expands his earlier work by introducing a new constraint, namely a cost to acquire, absorb, and process information in forming expectations. He claims that this yields additional benefits, most notably that the framework provides a micro-foundation for time-contingent adjustment as price adjustment is invariant to the state of the economy. "Time-contingent adjustment is appealing relative to its state-contingent alternative because it typically implies larger and longer real effects of monetary policy and it can reproduce the delayed and hump-shaped response of inflation to shocks." (p. 794). We maintain that a complete absence of state-contingency is not only theoretically unappealing but also counterfactual.

The Mankiw-Reis model of sticky information and its offshoots (e.g. Ball, Mankiw, and Reis, 2005, Reis, 2006, and others) have made some headway into finding a coherent theoretical explanation for puzzling dynamic effects of aggregate demand on output and inflation. Yet the sticky-information model is still a TDM, considered a virtue by the authors as it, "more naturally [justifies] the widely assumed time-contingent adjustment process." (Mankiw and Reis, 2002, p. 1319) And while it may indeed be more natural to assume that firms set prices at fixed intervals when they face costs of collecting information and choosing optimal plans, the above authors admit that,

> "information processing is more complex than the time-contingent adjustment assumed here. [...] It seems clear that when circumstances change in large and obvious ways, people alter the mental

resources they devote to learning and thinking about the new aspects of the world. Developing better models of how quickly people incorporate information about monetary policy into their plans, and why their response is faster at some times than at others, may prove a fruitful avenue for future research on inflation-output dynamics." (Mankiw and Reis, 2002, p. 1319)

Thus, a key challenge remains—to find a coherent explanation for price setting behavior that is able to encompass both periods of sluggish adjustment, observationally equivalent to time-dependent pricing, with periods of discrete, large price changes that respond to significant changes in the state of the economy. At the same time the model should be able to replicate considerable degree of price flexibility at the micro-level. We argue in this paper that a model of price setting under inferential expectations goes a long way towards meeting this challenge successfully.

2.3 State-Dependent Models

Recognizing the limitations of TDMs, a number of authors have offered models in which firms' pricing decisions depend on the state of the economy. Early contributions include include Blanchard and Kiyotaki (1987) and Ball and Romer (1990), building on the basic model of menu costs. With menu costs, a firm will only change its price if the benefit of changing prices, e.g. the additional profit resulting from a realignment of price to its optimum, exceeds the cost, a fixed amount incurred such as reprinting the menu in a restaurant. Mankiw (1985) showed how even relatively small menu costs can have large welfare consequences in aggregate but the empirical literature has been less kind to this theory and many economists find the menu cost theory unappealing for intuitive reasons.

Dotsey et al. (1999) find a simple way of modelling state-dependent pricing by assuming that firms face fixed costs of changing prices that are random across firms and drawn from a continuous distribution, implying that in equilibrium a fraction of firms adjusts prices. They find that, "the time-dependent approach captures the main mechanisms that lead to monetary nonneutrality under state-dependent pricing, although the magnitude of nonneutralities is often overstated when adjustment timing is assumed invariant to shocks." (Dotsey et al. 1999, p. 657) Furthermore, money is neutral in the long-run and, "the evolving distribution of price setters plays a central role in dictating how monetary shocks affect the dynamics of prices and real activity."

Maćkowiak and Wiederholt (2007) present a model in which monopolistically competitive firms have limited scope to process price information. In the spirit of Lucas (1972), firms must optimally decide what information to pay attention to: idiosyncratic, firm-specific information or aggregate information. The upshot is that when idiosyncratic conditions are more volatile or of greater significance than aggregate conditions, firms will pay more attention to idiosyncratic conditions and less to aggregate conditions. While shedding considerable light onto the complex problem of price setting, the model cannot explain why prices remain fixed for some time; they change every period. The authors propose adding a menu cost. However, this runs into the same problems that most other menu cost models do.

A different approach to modeling state-dependency is adopted by McAdam and Willman (2007). These authors question the intuitive validity of time-dependent pricing and point out that, "three salient features of [US business cycle] data—non-zero and, occasionally, highly volatile inflation rates as well as the likelihood of time-varying price stickiness—does not appear to tally with the Calvo-NKPC framework." (p. 7) McAdam's and Willman's innovation is to make the Calvo price-adjustment signal a function of inflation and market structure, viz. the Calvo signal itself is state-dependent. As a consequence the conventional Calvo probability parameter is not the average price contract length but its upper limit. The authors argue that this may help explain why, "empirical findings appear to over-estimate contract length." (p. 7)[8]

Golosov and Lucas (2007) re-examine the role menu costs play in the persistence of inflation and the real effects of monetary shocks. They remind us that in the Calvo (1983) model firms choose the size of price changes but not their timing and they argue that this is inconsistent with the facts. In particular: repricing is more frequent in high-inflation regimes than in low-inflation regimes (Calvo predicts a constant repricing rate) and those firms that do change prices are the ones whose prices are most out of line (in the Calvo model the firms that change prices are selected at random each period).[9]

Golosov and Lucas present a dynamic *Ss* model that respects these facts. Crucial to this is the assumption that firms are not only subject to aggregate inflation shocks but also mean-reverting idiosyncratic productivity shocks. As a consequence, Golosov and Lucas find that, "even though monetary shocks have almost no impact on the rate at which firms change prices, the shocks' real effects are dramatically less persistent than in an otherwise comparable economy with time-dependent price adjustment." (p. 173) Furthermore,

> "a positive aggregate shock induces the lowest-priced firms to increase prices. At the same time, it offsets negative idiosyncratic shocks, and some firms that would otherwise have decreased prices choose to wait. As a result, the lowest priced firms do most of the adjusting, their adjustments are large and positive, and the economywide price level increases quickly to reflect the aggregate shock. In the Calvo setting, in contrast, firms get the opportunity to reprice

[8] Another big advantage to their approach is that it is possible to derive a Phillips curve expressed in terms of the log-levels of variables, viz. without linearizing the system around an arbitrary steady-state, thus eliminating any kind of regime dependency. Simulations suggest that the economy's response to an inflation shock is highly regime dependent, a feature which leads traditional New Keynesian Phillips curve models to under-estimate inflation volatility and over-estimate inflation persistence.

[9] An extreme version of this was already modelled by Caplin and Spulber (1987), in which there exists a stationary distribution of firms' relative prices, making a monetary expansion neutral as firms at the low end of the distribution reprice to the high end.

randomly, many firms reprice even though they were already close to their desired price, and the average response of prices to the shock is much smaller. It takes longer for the monetary shock to be reflected in prices, and impulse responses become more persistent." (p. 173)

Gertler and Leahy (2008) develop a state-dependent pricing model that, with the help of some important simplifying assumptions, is able to generate a New Keynesian Phillips curve. Unlike Golosov and Lucas, the authors allow for strategic complementarities in price setting, often referred to in the literature as "real rigidities". Such real rigidities reenforce any kind of nominal inertia and are critical to generating an empirically reasonable degree of nominal rigidity. They derive a New Keynesian Phillips curve whose slope is steeper than under time-dependence and therefore exhibits greater price flexibility. The reason is the same as in Golosov and Lucas and other *Ss* models: the selection effect. Through the choice of suitable parameters the model can deliver the (relatively high) degree of price rigidity assumed in the time-dependent pricing literature and yet remain consistent with the recent microeconomic evidence on price adjustment.

3 Inferential Expectations

3.1 The Concept of Inferential Expectations

The theory of Inferential Expectations (IE) states that people adhere to beliefs for as long as the available information supports these beliefs. Only when there is sufficient evidence to the contrary will they abandon their beliefs and replace them with new ones that fit the data better. Such behaviour captures the mental (and physical) cost of acquiring, processing and interpreting information in a noisy, uncertain world. Rational expectations (RE), the dominant paradigm of expectations formation, demands a very high degree of sophistication that probably far exceeds the faculty of the average citizen. To cope, people adopt "fast and frugal heuristics" (see Gigerenzer et al., 1999) which allow them to navigate through life in a way that is mentally less taxing while still obeying over-arching maxims of rationality and optimality.

Following Menzies and Zizzo (2009), we consider a belief formation algorithm based on a Neyman-Pearson hypothesis test. We assume that when a belief is overturned agents switch to the RE solution. The RE solution can either be the 'all-knowing' RE of the macro modeller who understands the full system or the best possible use of information, as in a signal extraction model.

The adoption of RE when agents reject the null ensures that IE converges to RE in the limiting case when the test size $\alpha = 1$. If agents are unconcerned about mistakenly changing their beliefs they reject the null—and converge to RE—without fail. The nesting of RE within IE grounds IE theory in the structure of the model, retaining the much celebrated discipline of RE. Furthermore, α becomes a metric for rationality. If $\alpha = 0$, agents are completely unresponsive to evidence, while if $\alpha = 1$ they make the best possible use of evidence, converging to the RE solution.

An early application of IE is provided in Menzies and Zizzo (2009) who embed IE into the Dornbusch (1976) overshooting model, by assuming that agents do not believe a monetary expansion is permanent until it has endured for an extended period of time. The model generates 'double overshooting', with a second jump depreciation occurring at the moment when the only parcel of news is that the money shock has remained in place for an additional period. Agents realize monetary laxity is permanent and depreciate the currency further.

As a matter of definition, all IE models contain a cognitive target (the variable about which a null and alternative hypothesis is believed), a sample, a test statistic defined on the sample, a rejection region and a test size.[10]

3.2 Inferential Expectations and Firm Pricing Behavior

We propose a new theory of price setting based on inferential expectations. At the core of this theory are three basic assumptions:

- Information is noisy.
- Agents (firms) are characterized by belief conservatism and require sufficient contrarian evidence to switch beliefs because information gathering and processing is costly.
- If they don't have sufficient evidence, they adopt shortcuts in forming their H_0 beliefs.

Firms are free to set prices in every period, e.g. there is no physical adjustment (menu) cost. However, the capacity to collect, process, and correctly interpret information, which is necessary to find the optimal product price, is limited. It is too costly to collect all relevant information and it is computationally too taxing to calculate a new price in each period. Moreover, the information that is collected may be of dubious quality if it is random. As a consequence firms adopt a heuristic, a rule-of-thumb approach to processing a continuous stream of new information. This heuristic is captured by a process of hypothesis testing, which enables firms to hold onto existing beliefs about the economic environment until evidence to the contrary is (statistically) so compelling that the beliefs are revised.

To motivate the use of a hypothesis test, consider a signal extraction problem in which firms receive a noisy signal $\widehat{X}$ about some variable that feeds directly into their profit maximization problem. If there is too much noise, agents default to some mean level μ_X but if the noise is trivial relative to the signal, they use $\widehat{X}$. As is well known the 'best' (i.e. MMSE) thing to do is to update the prior

[10] In Menzies and Zizzo (2009), for example, the cognitive target is the long run level of money, the sample (and the test statistic) is the enumeration of time since the expansion began and the rejection region is the amount of time that convinces agents the expansion is permanent given their test size.

μ_X using $\widehat{X}$ with weights determined by the relative variances of the signal and the noise:

$$X_{\text{Signal Extract}} = \Omega\mu_X + (1-\Omega)\,\widehat{X},$$

where Ω is the optimal weight based on the relative variances of the signal and the noise.

However, there are two major problems with this approach. The first is that the variance of the shocks may be unknown, rendering the above formula impractical.

Second and more philosophically, Mayo (1996) argues forcefully that human beings are not natural Bayesian information processors. Even strong familiarity with probabilistic models is insufficient to overcome natural judgment mechanisms. More generally, she notes that scientists do not compare theories by assigning their posterior probability assignments to one hypothesis compared with others (op. cit. pg. 90).

We therefore propose the following recasting of the solution in terms of IE. We consider a situation in which the variance of the signal is typically small, i.e. Ω is close to one, but occasionally it is big. Agents use the mean level X as a null hypothesis unless they obtain overwhelming evidence that the shocks to the economy are large, in which case they reject H_0 and embrace H_1. The cognitive target in this model is a parameter measuring the variance of aggregate shocks to X. Under H_1 agents make the best use of information. If Ω is unknown, $\widehat{X}$ is the best guess, and we refer to this as RE.

What classifies as real world 'noise' in the model of prices in this paper? It can be any kind of information that is of variable or poor quality or that is open to multiple interpretations. Such information could be anything which is ultimately relevant to the firm's operations—macroeconomic data such as official inflation figures, interest rates, tax rates and wages; other firms' prices, product characteristics and demand; consumer sentiment reports and population surveys; even informal information such as rumours and hearsay.

Firms must then decide whether this new information is consistent with the null hypothesis, in which case they will leave prices unchanged, or whether the new information suggests that economic conditions have changed. In the latter case the null hypothesis needs to be replaced by a new belief about the state of the economy which in turn calls for a new optimal price to be computed and set.

The readiness with which firms revise their beliefs is determined by the parameters of the hypothesis test, the cognitive target, the signal, the test statistic, the rejection region and the test size. We will provide these details in the following section, together with some short-cuts agents take in forming their beliefs under H_0.

In what follows, the economy-wide price level is determined by the parameters associated with the hypothesis tests of groups of firms. Although we do not explore this in the current paper, one considerable benefit of the IE approach is that one can adjust the parameters of the hypothesis tests to reflect differing firm and industry characteristics. Firms operating in a highly competitive environment might be forced to collect many data points to allow them to adjust their beliefs frequently. Firms operating in protected industries with high rents can afford to collect less data, or they might have to contend with little variation in new data, making it less likely that they will hold onto the status quo. By making reasonable assumptions about industry and firm structure it is then possible to model aggregate price behavior from a "bottom-up" approach.

Alternatively, it may prove possible in future work to empirically test the microeconomic plausibility of aggregate price behavior. IE pricing shares with competing theories such as Mankiw and Reis (2002) the assumption that information is costly to collect and process. This reflects a growing consensus that price stickiness is not the result of direct adjustment costs (see, for example, Bils and Klenow, 2004, and Klenow and Kryvtsov, 2008). However, unlike Mankiw and Reis, IE pricing is state-, not time-dependent.

We believe that state-dependency is desirable for two reasons. First, it is intuitively appealing, for time-dependency implies that firms may not even react to the largest of shocks. Second, it matches observed pricing behaviour better. With regard to the latter, the IE framework below generates (a) a positively sloped Phillips curve (in inflation-output space); (b) the existence of *mistakenly* set prices; (c) the positive extrinsic margin correlation between the proportion of positive price changes and negative price changes; and (d) a 'tri-modal' distribution of large (positive and negative) price changes and small price changes.

Finally, we emphasize that while IE is state-dependent, it can take on a Calvo-appearance in a steady-state in which H_0 is true. In stable times, when the null of a small underlying shock to the economy is *true*, a fixed proportion of agents (equal to the test size) will optimally choose prices as they fall into the rejection region. Thus, for a 5% test size, the 'Calvo fairy' makes a guest appearance for 5% of firms each period. In more turbulent times, however, the fairy is at the beck and call of the state of the economy and will tap many more firms with its wand.

4 The Model

4.1 Economic Environment

We assume a single producer of a final good in a centralized economy. This 'aggregator' demands a continuum of intermediate goods y_{it} on $i \in [0,1]$ and combines them using a constant-elasticitiy-of-substitution (CES) production function:

$$y_t = \left[\int_0^1 y_{it}^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}, \qquad \theta > 1. \tag{1}$$

The aggregator solves the following profit maximization problem:

$$\max_{y_{it}} \left\{ p_t y_t - \int_0^1 p_{it} y_{it} di \right\},$$

the solution of which entails the following demand function for each input y_{it},

$$y_{it} = y_t \left(\frac{p_{it}}{p_t} \right)^{-\theta} \equiv y_t \left(\rho_{it} \right)^{-\theta}, \tag{2}$$

where ρ_{it} denotes the relative price of input i.

Substituting (2) into (1) and eliminating y_t gives a standard expression for the aggregate price level:

$$p_t^{1-\theta} = \int_0^1 p_{it}^{1-\theta} di. \tag{3}$$

4.2 Stochastic Environment

Intermediate firms' true real marginal costs are assumed constant across all firms but may vary over time. Thus, real marginal costs are given by

$$rmc_t = rmc_{ss} \left(1 + e_t \right),$$

where the subscript ss refers to the steady state and e_t is an i.i.d., mean-zero innovation term. Firms obtain n i.i.d. noisy signals of rmc_t, which they average to obtain a test statistic on rmc_t. The idiosyncratic perception of rmc_t for firm i is $\widehat{rmc}_{it}$:

$$\widehat{rmc}_{it} = \frac{1}{n} \sum_{j=1}^{n} rmc_{ss}(1 + e_t + \nu_{ijt}) = rmc_{ss}(1 + e_t + \bar{\nu}_{it}),$$

with

$$\begin{bmatrix} e \\ \bar{v} \end{bmatrix} \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \psi_t \sigma^2 & 0 \\ 0 & \frac{\sigma^2}{n} \end{bmatrix} \right),$$

and where v_{ijt} denotes a single noise innovation for firm i at time t.

Given these variances, the best guess of $rmc_{ss}(1 + e_t)$ is Lucas's optimal signal extraction, namely:

$$rmc_{Lucas,it} = \frac{1}{1 + n\psi_t} rmc_{ss} + \frac{n\psi_t}{1 + n\psi_t} \widehat{rmc}_{it}.$$

As shown below, firms want to set prices with reference to real marginal costs and aggregate prices according to $p_{it} = rmc_t p_t$. Thus, the full RE solution to this model would be for all firms to use $rmc_{Lucas,it}$ and solve the full model to obtain the correct value of p_t. However, IE may be a better framework because of agents' documented reluctance to use probability-weighted solutions, their ignorance of ψ_t and their proclivity to use shortcuts for p_t to avoid calculating a full-model solution.

Under IE firms conduct the following hypothesis test:

$$\begin{aligned} H_0 &: \psi_t = \psi \qquad (\psi \text{ small}) \\ H_1 &: \psi_t > \psi \end{aligned}$$

An auxiliary assumption is that agents minimize their calculation efforts under H_0 and set their guess of rmc_t equal to rmc_{ss}. Thus,

$$\begin{aligned} H_0 \text{ accepted:} &\quad rmc_{ss} \text{ used because } e_t \text{ small} \\ H_0 \text{ rejected:} &\quad \widehat{rmc}_{it} \text{ used to gauge } e_t \end{aligned}$$

Below we impose that when H_0 is accepted it not only implies that rmc_{ss} is used for rmc_t but also a state of mind in which the agent uses short-cuts to avoid a full calculation of p_t (which requires a full RE solution to the model).

Returning to the hypothesis test on rmc_t, agents use a standard normal test statistic:

$$\frac{\widehat{rmc}_{it} - rmc_{ss}}{rmc_{ss}\sigma\sqrt{\psi + 1/n}}$$

To see that the null implies that test statistic, note that when $\psi_t = \psi$, the variance of $\widehat{rmc}_{it}$ becomes

$$\psi\sigma^2 + \frac{\sigma^2}{n}.$$

We assume that the variance σ^2 is known to agents but ψ_t is not. That is, they know the margin of error associated with their own casual calculation, but they do not know the period-by-period volatility of the economy-wide shock.

The relevant components for this test are therefore: the null and alternative *hypotheses* about ψ_t (the *cognitive target*), the observations of $rmc_{ss}(1+e_t+v_{ijt})$ (the *sample*) which form the *test statistic* $\widehat{rmc}_{it}$, and the *rejection region* which is given by critical values in a normal test:

$$\text{Reject } H_0 \text{ if } \left| \frac{\widehat{rmc}_{it} - rmc_{ss}}{rmc_{ss}\sigma\sqrt{\psi + 1/n}} \right| > Z_{\alpha/2},$$

where Z refers to the standard normal distribution. The actual proportion of agents who reject H_0 (assuming independent hypothesis testing) depends upon the size of $n\psi_t$. Since the true standard deviation of $\widehat{rmc}_{it}$ is $rmc_{ss}\sigma\sqrt{\psi_t + 1/n}$, this can be calculated easily:

$$\Pr\left(\left|\frac{\widehat{rmc}_{it}-rmc_{ss}}{rmc_{ss}\sigma\sqrt{\psi+1/n}}\right|>Z_{\alpha/2}\right)=\Pr\left(\left|\frac{\widehat{rmc}_{it}-rmc_{ss}}{rmc_{ss}\sigma\sqrt{\psi_t+1/n}}\right|>\frac{\sqrt{\psi+\frac{1}{n}}}{\sqrt{\psi_t+\frac{1}{n}}}Z_{\alpha/2}\right) \tag{4}$$

$$=\Pr\left(|Z|>\frac{\sqrt{n\psi+1}}{\sqrt{n\psi_t+1}}Z_{\alpha/2}\right)=pvalue(\frac{\sqrt{n\psi+1}}{\sqrt{n\psi_t+1}}Z_{\alpha/2})\geq\alpha$$

This proportion is by definition equal to α under H_0 but if H_0 is false (that is, $\psi_t > \psi$), then

$$\frac{\sqrt{n\psi+1}}{\sqrt{n\psi_t+1}}<1$$

and the p-value will be greater than α. We call this proportion $\mu_t\alpha$ where $\alpha \leq \mu_t\alpha \leq 1$. Note that μ_t is indexed by t as it is a function of ψ_t.

Table 1 shows the proportion of agents who reject H_0 for various test sizes with various true $n\psi_t$'s. Numerical calculations require a value of ψ under the null, and we use $\psi = 0$ for simplicity. From the table, if $\alpha = 0.5$, and $n\psi_t = 0.4$, $\mu_t = 0.10/0.05 = 2$. That is, twice as many agents reject the null than they would if $H_0 : \psi = 0$ were true. The last column shows how long a particular firm is likely to believe H_0, given the rejection proportions, when the test size is 0.05.

Table 1

More Agents Reject the H_0 than the Nominal Test Size

$\psi = 0$ under the null

		Proportion Rejecting H_0				Duration
	α	**0.01**	**0.05**	**0.10**	**0.20**	of H_0
	$Z_{\alpha/2}$	**2.576**	**1.960**	**1.645**	**0.674**	$\alpha = 0.05$
$n\psi_t$	**0.0**	0.01	0.05	0.10	0.50	20
	0.2	0.02	0.07	0.13	0.54	14
	0.4	0.03	0.10	0.16	0.57	10
	1	0.07	0.17	0.24	0.63	6
	2	0.14	0.26	0.34	0.70	4
	5	0.29	0.42	0.50	0.78	2
	100	0.80	0.85	0.87	0.95	1

If $n\psi_t >> n\psi$ (= 0 above) the p-value will approach unity, as in the bottom row of the table. In that case, everyone adopts H_1 and uses $\widehat{rmc}_{it}$. The definition of $\widehat{rmc}_{it}$ prohibits ψ_t from being too large (rmc_t cannot be negative), so $n\psi_t >> n\psi$ must occur through lots of information on a shock being available. That is, if $n \to \infty$ *and* H_0 is false, everyone adopts H_1 and $\widehat{rmc}_{it}$ is used.[11]

[11] Clearly if H_0 is true, then by definition the proportion rejecting H_0 will be unaffected by

In what follows, we repeatedly use the result that $1-\mu_t\alpha$ is the proportion of agents who believe H_0 after the hypothesis test and $\mu_t\alpha$ is the fraction who believe H_1.

We now return to (3) which is pivotal to defining the Phillips curve. We index firms who believe H_0 over $[0, 1-\mu_t\alpha]$ and those who believe H_1 over $[1-\mu_t\alpha, 1]$. This allows us to write a probability weighted average of H_0 and H_1 prices.

$$\begin{aligned} p_t^{1-\theta} &= \int_0^1 p_{it}^{1-\theta} di = \int_0^{1-\mu_t\alpha} p_{it}^{1-\theta} di + \int_{1-\mu_t\alpha}^1 p_{it}^{1-\theta} di \\ &= \int_0^{1-\mu_t\alpha} p_{it}^{H_0\,1-\theta} di + \int_{1-\mu_t\alpha}^1 p_{it}^{H_1\,1-\theta} di \\ &= (1-\mu_t\alpha)\int_0^1 p_{it}^{H_0\,1-\theta} di + \mu_t\alpha \int_0^1 p_{it}^{H_1\,1-\theta} di \end{aligned} \tag{5-7}$$

4.3 Optimization

We define p^{H_0} and p^{H_1} with reference to optimal prices. Firms choose p_{it} to maximize current profits.[12] Nominal marginal costs are denoted by mc_{it}. Firm i's profit function is given by

$$\Pi_{it}^{real} = \frac{(p_{it} - mc_{it})\, y_{it}}{p_t}.$$

Substituting (2) for y_{it},

$$\begin{aligned} \Pi_{it}^{real} &= (\rho_{it} - rmc_{it})\, y_t\, (\rho_{it})^{-\theta} \\ &= \left(\rho_{it}^{1-\theta} - \rho_{it}^{-\theta} rmc_{it}\right) y_t \end{aligned} \tag{9}$$

The corresponding necessary first-order condition is:[13]

$$\frac{d\Pi_{it}^{real}}{d\rho_{it}} = \left\{(1-\theta)\rho_{it}^{-\theta} + \theta\rho_{it}^{-\theta-1} rmc_{it}\right\} y_t = 0 \tag{10}$$

increasing the sample size; the critical values of the Z test statistic head towards zero at the same rate as the standard deviation, preserving the proportion who reject H_0 at α. From (4), under the null the probability of being in the rejection region stays at p-value($Z_{\alpha/2}$) which is clearly $\alpha/2$.

[12] This assumption is made for simplicity. We could model IE firms who realize that there is a chance that they will not change future prices because they will not fall in a rejection region. The infinite horizon profit maximization problem shares certain features of the Calvo setup (see Appendix 3).

[13] The second derivative of (9) is negative, fulfilling the second-order-condition for a local maximum.

Multiply by ${\rho_{it}}^{1+\theta}/\left(1-\theta\right)y_t$ to obtain

$$\rho_{it} = \frac{\theta}{\theta-1} rmc_{it},$$

or, equivalently,

$$p_{it} = \frac{\theta}{\theta-1} rmc_{it} p_t. \tag{11}$$

4.4 Beliefs Under H_0 and H_1

We can now be specific about beliefs and shortcuts used to calculate (11) under IE. Agents need to pin down rmc_{it} and p_t.

Under H_0 Agents set rmc_{it} equal to its unconditional expected value rmc_{ss}. Furthermore, they adopt one of two rules-of-thumb for p_t on the RHS of (10):[14]

1. *Indexers*: A share β set $p_t = \left(1+\pi_{t-1}\right)p_{t-1}$ in (11).[15]
2. *Nonindexers*: A share $1-\beta$ set $p_t = p_{t-1}$ in (11).

Under H_1 Agents under H_1 work out the full model RE solution for p_t in (11) and use $\widehat{rmc}_{it}$ for rmc_{it}.

Returning to (7), we can expand the price index as a convex combination of three price integrals.

[14] IE is itself a rule of thumb, but we obtain modeling discipline by the use of the HMZ rule which states that "the number of rules-of-thumb for a given agent cannot exceed the number of thumbs".

[15] It is possible to index off p_{it} and we do so in the infinite horizon problem of Appendix 3. However, it becomes more complex when analyzing distributions, so we adopt the simplification in the text.

$$
\begin{aligned}
p_t^{1-\theta} &= (1-\mu_t\alpha)\int_0^1 p_{it}^{H01-\theta} di + \mu_t\alpha \int_0^1 p_{it}^{H11-\theta} di \\
&= \underbrace{\beta(1-\mu_t\alpha)\int_0^1 [\frac{\theta}{\theta-1} rmc_{ss}\{1+\pi_{t-1}\}p_{t-1}]^{1-\theta} di}_{H_0 \text{ indexers}} \\
&\quad \underbrace{+ (1-\beta)(1-\mu_t\alpha)\int_0^1 [\frac{\theta}{\theta-1} rmc_{ss} p_{t-1}]^{1-\theta} di}_{H_0 \text{ non-indexers}} \\
&\quad \underbrace{+ \mu_t\alpha \int_0^1 [\frac{\theta}{\theta-1} \widehat{rmc}_{it} p_t]^{1-\theta} di}_{H_1}
\end{aligned}
$$

Appendix 2 proves that

$$
\int_0^1 \left[\frac{\theta}{\theta-1} \widehat{rmc}_{it} p_t^{1-\theta}\right] di = \left[\frac{\theta}{\theta-1} rmc_t p_t\right]^{1-\theta},
$$

so:

$$
\begin{aligned}
p_t^{1-\theta} &= \beta(1-\mu_t\alpha)\left[\frac{\theta}{\theta-1} rmc_{ss}\{1+\pi_{t-1}\}p_{t-1}\right]^{1-\theta} \\
&\quad + (1-\beta)(1-\mu_t\alpha)\left[\frac{\theta}{\theta-1} rmc_{ss} p_{t-1}\right]^{1-\theta} + \mu_t\alpha\left[\frac{\theta}{\theta-1} rmc_t p_t\right]^{1-\theta}
\end{aligned} \tag{12}
$$

Divide both sides by $(p_{t-1})^{1-\theta}$ to simplify:

$$
\begin{aligned}
(1+\pi_t)^{1-\theta} &= \beta(1-\mu_t\alpha)\left[\frac{\theta}{\theta-1} rmc_{ss}\{1+\pi_{t-1}\}\right]^{1-\theta} \\
&\quad + (1-\beta)(1-\mu_t\alpha)\left[\frac{\theta}{\theta-1} rmc_{ss}\right]^{1-\theta} + \mu_t\alpha\left[\frac{\theta}{\theta-1} rmc_t (1+\pi_t)\right]^{1-\theta}
\end{aligned}
$$

In the SS inflation is zero, so:

$$
\begin{aligned}
1 &= (1-\mu_t\alpha)\left[\frac{\theta}{\theta-1} rmc_{ss}\right]^{1-\theta} + \mu_t\alpha\left[\frac{\theta}{\theta-1} rmc_{ss}\right]^{1-\theta} \\
&\Rightarrow \quad \frac{\theta}{\theta-1} rmc_{ss} = 1.
\end{aligned}
$$

4.5 The Inferential Expectations Phillips Curve

Using the above expressions one can obtain the IE Phillips Curve. To do this it is necessary to log-linearize around the steady-state. Denote log deviations of variables from their steady state by a tilde such that, for real marginal costs,

$$rmc_t = rmc_{ss} e^{\widetilde{rmc}_t}.$$

Now we can rewrite (12) as

$$\begin{aligned}(1+\pi_t)^{1-\theta} &= \beta(1-\mu_t\alpha)\{1+\pi_{t-1}\}^{1-\theta} \\ &+ (1-\beta)(1-\mu_t\alpha) + \mu_t\alpha\left[\frac{\theta}{\theta-1} rmc_{ss}.e^{\widetilde{rmc}_t}(1+\pi_t)\right]^{1-\theta} \\ &= \beta(1-\mu_t\alpha)\{1+\pi_{t-1}\}^{1-\theta} \\ &+ (1-\beta)(1-\mu_t\alpha) + \mu_t\alpha\left[(1+\widetilde{rmc}_t)(1+\pi_t)\right]^{1-\theta}\end{aligned}$$

Simplifying further yields,

$$\begin{aligned}(1+(1-\theta)\pi_t) &= \beta(1-\mu_t\alpha)\{1+(1-\theta)\pi_{t-1}\} + (1-\beta)(1-\mu_t\alpha) \\ &+ \mu_t\alpha(1+(1-\theta)\{\pi_t+\widetilde{rmc}_t\}) \\ &= (1-\mu_t\alpha) + \beta(1-\mu_t\alpha)(1-\theta)\pi_{t-1} \\ &+ \mu_t\alpha(1+(1-\theta)\{\pi_t+\widetilde{rmc}_t\}) \\ &= 1+\beta(1-\mu_t\alpha)(1-\theta)\pi_{t-1} + \mu_t\alpha(1-\theta)\{\pi_t+\widetilde{rmc}_t\}\end{aligned}$$

Solving for π_t gives,

$$\pi_t = \beta(1-\mu_t\alpha)\pi_{t-1} + \mu_t\alpha\{\pi_t+\widetilde{rmc}_t\},$$

or

$$\pi_t = \beta\pi_{t-1} + \frac{\mu_t\alpha}{1-\mu_t\alpha}\widetilde{rmc}_t. \tag{13}$$

This is the IE Phillips curve for the case where IE agents merely perform static profit maximization. The first, and most important observation is that the 't' subscript on μ indicates that it this is a state-dependent model. Indeed, if $\mu_t\alpha$ approaches unity, (13) collapses (by multiplying both sides by $(1-\mu_t\alpha)$) to $\widetilde{rmc}_t = 0$, which implies that the aggregate price level moves one-for-one with nominal marginal costs so their proportional deviations from the SS are always the same. This is a full RE solution.

This is not surprising from (11) calculated under H_1. The observation noise is averaged out over the whole economy (cf. Appendix 2) leaving an 'average' firm with $rmc_t p_t$ on the RHS of (11). The p_t's cancel out, giving nominal marginal costs times the constant markup.

In what state of nature does $\mu_t\alpha$ approach unity? The key requirement is that the signal is relatively clear to firms ($n\psi_t >> n\psi$). When firms know 'what is going on' the introduction of an 'obvious' cost shock (one where the parcels of information, n, are numerous and $\psi_t > \psi$) will be incorporated immediately into prices, unlike the Calvo model. In fact, as $n \to \infty$ the 'information deluge' will align prices with nominal marginal costs for arbitrarily small shocks to nominal marginal costs. This could be thought of as a metaphor for a small shock, about which there is perfect information, such as a change in a consumption tax rate.

But there are other situations, or states, where RE does not hold. When H_0 is true, $\mu_t = 1$, so agents adopt something close to backwardlooking indexation, with the degree of indexation determined in a straightforward manner by the proportion of firms who index when they fail to reject H_0. Thus, in the IE framework inertia in aggregate inflation can coexist with a large dispersion of individual price changes (see next section). Clearly if $\beta = 0$, no agents index under H_0 and there is no impact of lagged inflation.

In the backward indexation solution, recall that under H_0 a fraction of agents α *do* reject the null, thereby doing a proper calculation each period. Hence, there is always a role for fundamentals in the IE Phillips Curve. Of course, if α goes to zero, this effect disappears and pure backward-looking inflation occurs—to an extent determined by indexation.

4.6 The Distribution of Prices

In this section we derive the distribution of prices and of price changes. We find, in keeping with the recent micro evidence, that:

- Price changes for some firms are large but small for others.
- Prices can decrease, as well as increase.
- The extensive margin exhibits a positive correlation between the number of firms making positive changes and negative price changes.

One implication of the first two findings is that with IE firms occasionally make mistakes. Not only do firms adopt a rule-of-thumb to economize on information costs but occasionally they draw a wrong inference from the data presented to them, viz. they commit type I errors. These errors will typically be corrected relatively quickly but they highlight that movements in prices are not always desirable. This stands in stark contrast to traditional RE models.

The model delivers a distribution of prices, based on firms' different draws of $\widehat{rmc}_{it}$. The steady-state analytic distribution of prices, and price changes, can be examined graphically.

A steady-state in this context is defined as as $p_t = p_{t-1}$ (i.e. $\pi_t = 0$). More subtly, we assume a constant $\mu_t \alpha$ (the proportion of agents rejecting the null) for all time periods whence we calculate the distribution of price changes. Finally, in keeping with the monopolistic competition literature, we assume that θ is large, so $\theta / (\theta - 1)$ is close to one.

We can use the relative price form of (11) since p_t, $(1 + \pi_{t-1})p_{t-1}$, p_{t-1}—which are all used for p_t by agents under different hypotheses—are identical in the steady state. Divide both sides of (11) by p_t to obtain

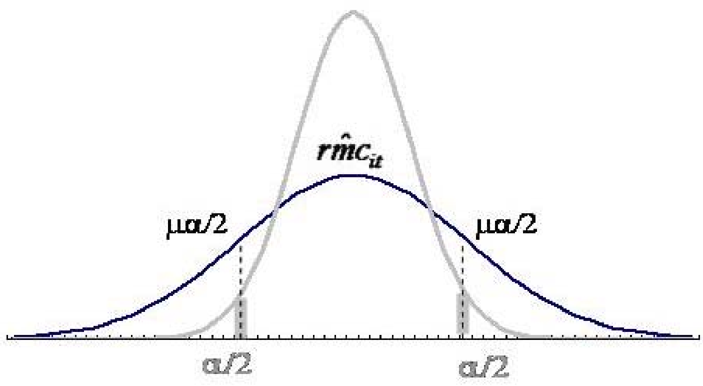

Figure 1: Distribution of rmc Under H_0 and True rmc-Distribution

$$\rho_{it} = \frac{\theta}{\theta - 1} rmc_{it}.$$

Since $\theta / (\theta - 1)$ is close to one, the steady-state relative prices under H_0 and H_1 are equal to the value used for rmc and differ only insofar as a different guess for rmc is adopted under H_0 and H_1. That is,

$$\begin{array}{ll} \rho_{it} \approx rmc_{ss} & \text{under } H_0 \\ \rho_{it} \approx \widehat{rmc}_{it} & \text{under } H_1 \end{array}$$

This permits a direct graphical mapping of the distribution of $\widehat{rmc}_{it}$ onto the distribution of relative prices and relative price changes.

In figure 1, consider firm i in receipt of a signal $\widehat{rmc}_{it}$. Its belief about the sampling distribution of rmc is the light grey pdf, and it rejects H_0 in the tails. However, the actual probability of rejection is the p-value of the shaded critical values, namely $\mu\alpha$. (We drop the t-subscript on μ in the steady state.)

From now on, all that is relevant is the actual proportion rejecting, and this is determined by the solid sampling distribution, as shown in figure 2.

The distribution of (relative) prices can be derived easily from figure 2. A proportion $\mu\alpha$ fall in the two tails and their estimate is $\widehat{rmc}_{it}$. So, the distribution of prices has exactly the same tails. However, under H_0 agents believe $\rho = rmc_{ss}$ so probability mass $1 - \mu\alpha$ is a discrete random variable at the mean rmc_{ss} of the sampling distribution. This is summarized in figure 3.

We can now derive the distribution of price *changes* assuming that each probability mass in the above distribution is randomly allocated across prices in the next period, in the proportions above. That is, the proportion of agents $\mu\alpha/2$ in, say, the left tail is divided up next period across the left, middle and

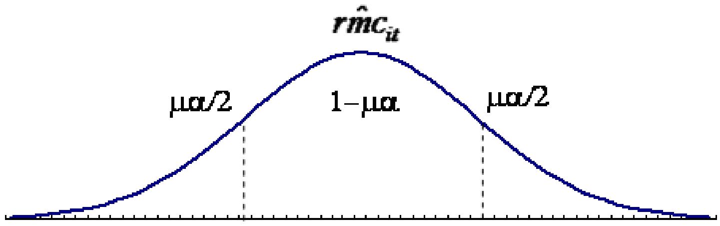

Figure 2: True Distribution of *rmc* with Rejection Regions

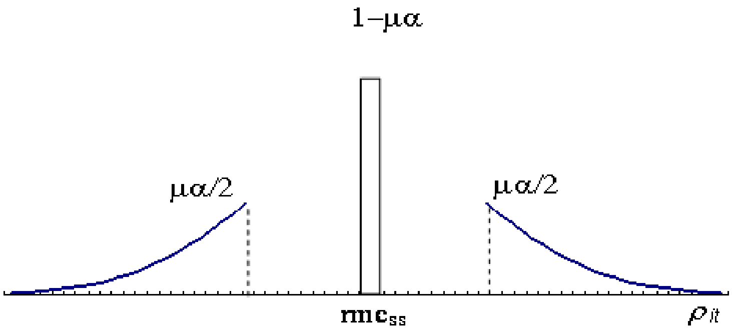

Figure 3: The Steady-State Distribution of Prices

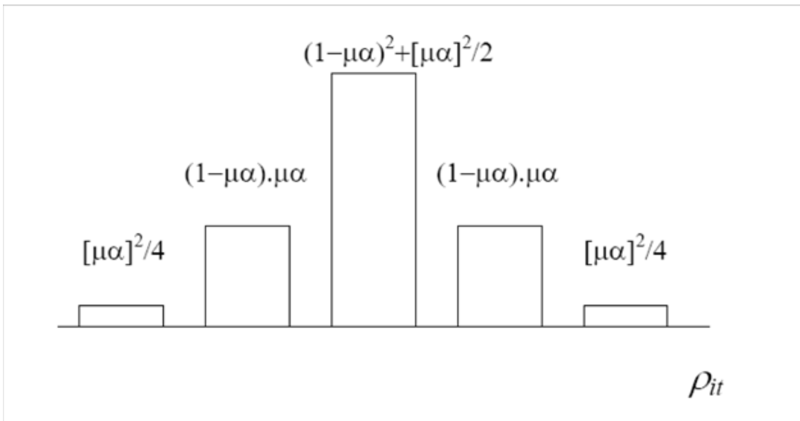

Figure 4: Graphical Representation of Steady-State Distribution Price Changes

right with proportions $(\mu\alpha/2)(\mu\alpha/2)$, $(\mu\alpha/2)(1-\mu\alpha)$, and $(\mu\alpha/2)(1-\mu\alpha)$. Table 2 shows the entire transition matrix, with rows giving the position of the price in period t and columns the position of the price in period $t+1$.

Table 2

Transition Matrix for Price Changes

	Left Tail$_{t+1}$	Middle$_{t+1}$	Right Tail$_{t+1}$
Left Tail$_t$	$(\mu\alpha)^2/4$	$(\mu\alpha/2)(1-\mu\alpha)$	$(\mu\alpha)^2/4$
Middle$_t$	$(1-\mu\alpha)(\mu\alpha/2)$	$(1-\mu\alpha)^2$	$(1-\mu\alpha)(\mu\alpha/2)$
Right Tail$_t$	$(\mu\alpha)^2/4$	$(\mu\alpha/2)(1-\mu\alpha)$	$(\mu\alpha)^2/4$

We denote an unchanged position between t and $t+1$ as '0', a shift up one region (from the left tail to the middle or from the middle to the right tail) as '+' and a shift from the left tail to the right tail as '++'. Using the same convention for '−' and '−−' we obtain a matrix of qualitative changes, summarized in table 3.

Table 3

Qualitative Transition Matrix for Price Changes

	Left Tail$_{t+1}$	Middle$_{t+1}$	Right Tail$_{t+1}$
Left Tail$_t$	0	+	++
Middle$_t$	−	0	+
Right Tail$_t$	−−	−	0

A graphical representation of this is given in figure 4.

If α is small, '++' and '−−' are second-order and crossing from one tail to the opposite becomes nigh impossible. (See figure 5.) The qualitative transition matrix thus collapses to a tri-modal distribution of price changes as depicted in figure 6. At the same time, increases in μ (higher proportion of rejecters)

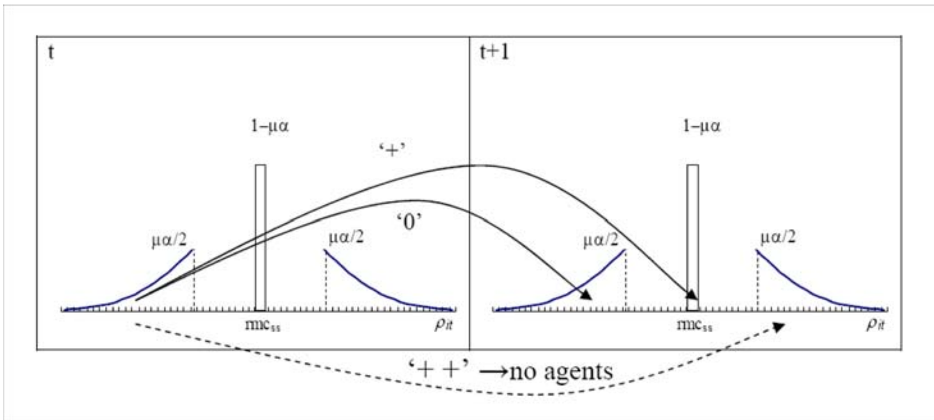

Figure 5: Graphical Representation of Price Changes

will fatten the tails, supporting the empirical evidence that the fraction of price increases is roughly offset by the fraction of price decreases.

4.7 Filling in the Hollows of the Price Distribution

One of the stylized facts about price changes is that they are roughly 'tri-modal'. That is, some prices change by a relatively large amount, both up and down, and some change by only a small amount. The distribution we have shown so far has two 'hollows' and a spike at rmc_{ss} implying that the fraction of firms experiencing zero price changes could be very large (the mapping from the middle in t to the middle in $t+1$ occurs with probability $(1-\mu\alpha)^2$). In this section, we explore whether the central spike in the distribution of prices could become a wider probability mass. If so, it would allow 'tri-modal' price changes.

The transition matrices in the previous section described probabilities of moving between particular ranges. The actual 'point-to-point' changes (as opposed to region-to-region changes) would be continuous, except in the case of rmc_{ss} to rmc_{ss}, and there is no presumption that they would have zero probability mass over particular regions.

However, even if we confine our attention to levels, it turns out that IE does not necessarily imply the stark distribution in the diagram above. First, we must consider what happens out of steady state. Then, we must consider what happens if there is a more nuanced null hypothesis.

We used the steady state as a simplification earlier, removing the need to consider the share β of indexers and the share $1-\beta$ of non-indexers. But that led us to the distribution above, and so it is natural to ask if the relaxation of this assumption will create a different distribution.

Out of the steady-state the central column will be divided into two subgroups: the indexers will maintain the mean rmc_{ss}, but the non-indexers will

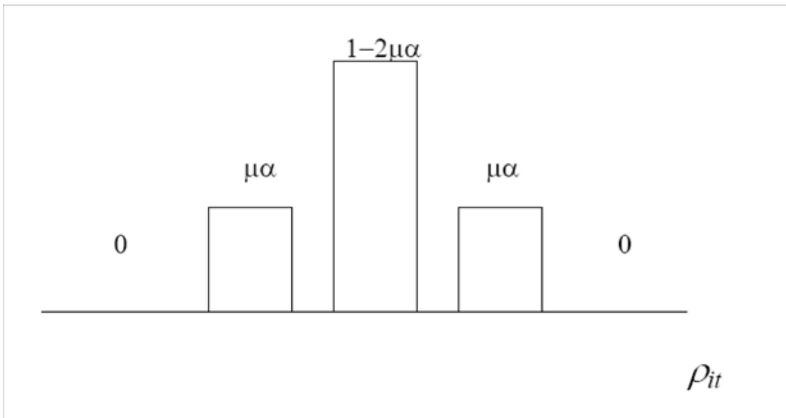

Figure 6: Graphical Representation of Steady-State Distribution of Price Changes for $\alpha \to 0$.

have mean $rmc_{ss}p_{t-1}/p_t$.[16]

One could go further with numerical simulations. A richer model could be obtained from (11) if one assumed that each firm under the null indexed to its *own* past price, rather than to aggregate prices. Namely,

1. *Indexers*: A share β set $p_t = (1+\pi_{t-1})\, p_{it-1}$ in (11).
2. *Nonindexers*: A share $1-\beta$ set $p_t = p_{it-1}$ in (11).

That way, the past distribution of prices would feed into the present one, and the hollows would be filled out, but in a way that is hard to predict without running simulations.

One can generate a more nuanced distribution even in the steady state. Let us suppose that under the null agents do a Lucas-style signal extraction problem, so that under the H_0-value of ψ we have

$$rmc_{Lucas} = \frac{1}{1+n\psi_t} rmc_{ss} + \frac{n\psi}{1+n\psi}\widehat{rmc}_{it}$$

Since we are in the steady state we describe the two sorts of relative prices, under H_0 and H_1:

$$\begin{array}{ll} \rho_{it} \approx \frac{1}{1+n\psi_t} rmc_{ss} + \frac{n\psi}{1+n\psi}\widehat{rmc}_{it} & \text{under } H_0 \\ \rho_{it} \approx \widehat{rmc}_{it} & \text{under } H_1 \end{array}$$

The rejection regions of $\widehat{rmc}_{it}$ still map one-to-one onto the distribution of prices but now the H_0 probability mass of $1-\mu\alpha$ has to be mapped onto

[16] From (11), for small volatility in p_t and p_{t-1} the former mean is $rmc_{ss}p_{t-1}(\pi_{t-1}+1/p_t) = rmc_{ss}$ while the latter mean $rmc_{ss}p_{t-1}/p_t$.

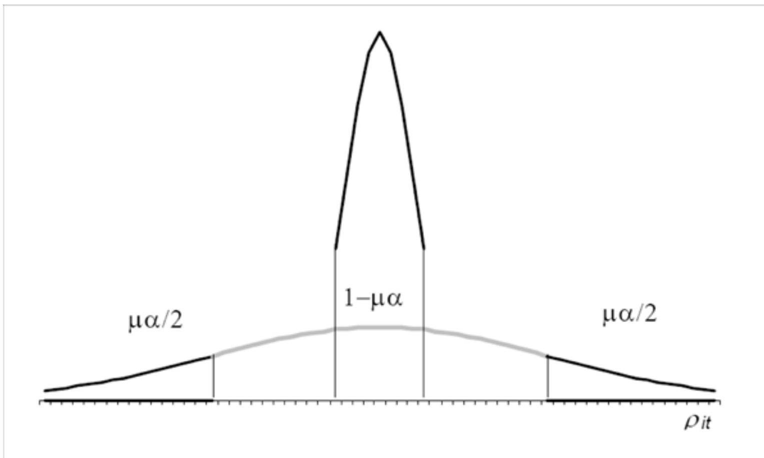

Figure 7: Steady-State Distribution of Prices with Lucas Agents Under H_0

$\frac{1}{1+n\psi_t} rmc_{ss} + \frac{n\psi}{1+n\psi}\widehat{rmc}_{it}$. What does this random variable look like? Let σ^2_{rmc} be $\sigma\sqrt{\psi + 1/n}$. Then it follows,

$$
\begin{aligned}
&\widehat{rmc}_{it} \sim N\left(rmc_{ss}, \sigma^2_{rmc}\right) \\
&\qquad \Rightarrow \frac{n\psi}{1+n\psi}\widehat{rmc}_{it} \sim N\left(\frac{n\psi}{1+n\psi} rmc_{ss}, \left(\frac{n\psi}{1+n\psi}\right)^2 \sigma^2_{rmc}\right) \\
&\qquad \Rightarrow \frac{1}{1+n\psi_t} rmc_{ss} + \frac{n\psi}{1+n\psi}\widehat{rmc}_{it} \sim N\left(rmc_{ss}, \left(\frac{n\psi}{1+n\psi}\right)^2 \sigma^2_{rmc}\right)
\end{aligned}
$$

The agents who fail to reject H_0 will find themselves spread out over a $1 - \mu\alpha$ confidence interval with standard deviation $\frac{n\psi}{1+n\psi}\sigma_{rmc}$. The final distribution of prices will have a cluster of ‘small’ values, a probability mass of extreme values, and an absence of middle values. (See figure 7.) Clearly, as ψ approaches zero the Lucas H_0 solution puts all the weight on rmc_{ss} and we return to our earlier ‘hollow’ solution with a spike at rmc_{ss}.

As before, the distribution of changes will not necessarily have a gap in the probability mass.

5 Conclusion

Time-dependent pricing models such as the well-known Calvo (1983) model are deficient in two key respects: they fail to match disaggregated price data and they do not allow for state-dependency of pricing decisions. We propose a new model of pricing based on a new theory of expectations formation called inferential expectations. To overcome the costs associated with information acquisition and processing, firms adopt simple heuristics. In particular, under inferential expectations firms are assumed to perform hypothesis tests over the variance of real marginal costs. Depending on the size of the shocks and amount of available information, firms may remain inactive for an extended period of time and only review their pricing policy when sufficient evidence warrants it. This approach gives rise to an augmented, state-dependent Phillips curve and sheds light onto the dynamics of disaggregated prices.

The model presented here is only a first pass at incorporating inferential expectations into a pricing model. It is still lacking in certain dimensions and not able to redress all the deficiencies identified with the status quo but it brings across the basic idea and shows the way for future research.

References

[1] Ball, Laurence (1994). "Credible Disinflation with Staggered Price Setting," *American Economic Review*, Vol. 84, 282-289.

[2] Ball, Laurence, N. Gregory Mankiw, and Ricardo Reis (2005). "Monetary policy for inattentive economies," *Journal of Monetary Economics*, Vol. 52, No. ??, 703-725.

[3] Ball, Laurence, and David Romer (1990). "Real Rigidities and the Non-Neutrality of Money," *Review of Economic Studies*, Vol. 57, 183-203.

[4] Ball, Laurence, and David Romer (1991)."Sticky prices as coordination failure," *American Economic Review*, Vol. 81, No. 3, 539-552.

[5] Bils, Mark, and Peter J. Klenow (2004). "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, Vol. 112, No. 5, 947-985.

[6] Blanchard, Olivier J., and Nobuhiro Kiyotaki (1987). "Monopolistic Competition and the Effects of Aggregate Demand," *American Economic Review*, Vol. 77, 647-666.

[7] Blinder, Alan S., Elie R. D. Canetti, David E. Lebow, and Jeremy B. Rudd (1998). *Asking About Prices: A New Approach to Understanding Price Stickiness*. New York: Russell Sage Foundation.

[8] Boivin, Jean, Marc P. Giannoni, and Ilian Mihov (2009). "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data," *American Economic Review*, Vol. 99, No. 1, 350-384.

[9] Bonomo, Marco, and Carlos Carvalho (2004). "Endogenous Time-Dependent Rules and Inflation Inertia," *Journal of Money, Credit and Banking*, Vol. 36, No. 6, 1015-41.

[10] Caballero, Ricardo J., and Eduardo Engel (1993a). "Microeconomic Rigidities and Aggregate Price Dynamics," *European Economic Review*, Vol. 37, 697-717.

[11] Caballero, Ricardo J., and Eduardo Engel (1993b). "Microeconomic Adjustment Hazards and Aggregate Dynamics," *Quarterly Journal of Economics*, Vol. 108, 359-383.

[12] Caballero, Ricardo J., and Eduardo Engel (2006). "Price Stickiness in Ss Models: Basic Properties," *mimeo*, MIT.

[13] Caballero, Ricardo J., and Eduardo Engel (2007). "Price Stickiness in Ss Models: New Interpretations of Old Results," *Journal of Monetary Economics*, Vol. 54 (supp. 1), 100-121.

[14] Calvo, Guillermo A. (1983). "Staggered Prices In a Utility-Maximizing Framework," *Journal of Monetary Economics*, Vol. 12, No. 3, 383-398.

[15] Caplin, Andrew, and John Leahy (1991). "State-Dependent Pricing and the Dynamics of Money and Output," *Quarterly Journal of Economics*, Vol. 106, 683-708.

[16] Caplin, Andrew, and John Leahy (1997). "Aggregation and Optimization with State-Dependent Pricing," *Econometrica*, Vol. 65, 601-625.

[17] Caplin, Andrew, and Daniel Spulber (1987). "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, Vol. 102, 703-725.

[18] Carlton, Dennis W. (1985). "The Rigidity of Prices," *American Economic Review*, Vol. 76(4), 637-658.

[19] Cecchetti, Stephen G. (1986). "The Frequency of Price Adjustment: A Studyof the Newsstand Prices of Magazines," *Journal of Econometrics*, Vol. 31(3), 255-274.

[20] Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans (1999). "Monetary Policy Shocks: What Have We Learned and to What End?" in: *Handbook of Macroeconomics*, Volume 1A, edited by John B. Taylor and Michael Woodford, New York: Elsevier Science, 65-148.

[21] Danziger, Leif (1999). "A Dynamic Economy with Costly Price Adjustments," *American Economic Review*, Vol. 89, 878-901.

[22] Dotsey, Michael, and Robert G. King (2005). "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models," *Journal of Monetary Economics*, Vol. 52, 213-242.

[23] Dotsey, Michael, Robert G. King, and Alexander L. Wolman (1999). "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics*, Vol. 114, No. 2, 655-90.

[24] Dotsey, Michael, Robert G. King, and Alexander L. Wolman (2006). "Inflation and Real Activity with Firm-Level Productivity Shocks: A Quantitative Framework," *mimeo*, Federal Reserve Bank of Richmond.

[25] Fischer, Stanley (1977). "Long-Term Contracts, Rational Expectations and the Optimal Money Supply Rule," *Journal of Political Economy*, Vol. 85, No. 1, 163-190.

[26] Frydman, Roman, and Michael D. Goldberg (2003). "Imperfect Knowledge Expectations, Uncertainty-Adjusted Uncovered Interest Parity, and Exchange Rate Dynamics," in: Aghion, Philippe, Roman Frydman, Joseph Stiglitz, and Michael Woodford (eds.). *Knowledge, Information, and Expectations in Modern Macroeconomics.* Princeton: Princeton University Press.

[27] Frydman, Roman, and Michael D. Goldberg (2007). *Imperfect Knowledge Economics: Exchange Rates and Risk.* Princeton: Princeton University Press.

[28] Fuhrer, Jeffrey, and George Moore (1995). "Inflation Persistence," *Quarterly Journal of Economics*, Vol. 110, 127-160.

[29] Gertler, Mark, and John Leahy (2008). "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, Vol. 116, No. 3, 533-572.

[30] Gigerenzer, Gerd (2007). *Gut Feelings: The Intelligence of the Unconscious.* New York: Viking Press.

[31] Golosov, Mikhail, and Robert E. Lucas Jr. (2007). "Menu Costs and Phillips Curves," *Journal of Political Economy*, Vol. 115, No. 2, 171-199.

[32] John, A. Andrew, and Alexander L. Wolman (2008). "Steady-state equilibrium with state-dependent pricing," *Journal of Monetary Economics*, Vol. 55, 383-405.

[33] Kashyap, Anil K. (1995). "Sticky Prices: New Evidence from Retail Catalogs," *Quarterly Journal of Economics*, Vol. 110(1), 245-274.

[34] Kiley, Michael T. (2000). "Endogenous Price Stickiness and Business Cycle Persistence," *Journal of Money, Credit, and Banking*, Vol. 32, No. 1, 28-53.

[35] Klenow, Peter J., and Oleksiy Kryvtsov (2008). "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics*, Vol. 123, No. 3, 863-904.

[36] Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable (1997). "The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains," *Quarterly Journal of Economics*, Vol. 112(3), 791-825.

[37] Lucas, Robert E. Jr. (1972). "Expectations and the Neutrality of Money," *Journal of Economic Theory*, Vol. 4, 103-124.

[38] Maćkowiak, Bartosz, and Mirko Wiederholt (2007). "Optimal Sticky Prices under Rational Inattendance," *mimeo*, Northwestern University, March.

[39] Mankiw, N. Gregory (1985). "Small menu costs and large business cycles: A macroeconomic model of monopoly," *Quarterly Journal of Economics*, Vol. 100, 529-539.

[40] Mankiw, N. Gregory (2001). "The Inexorable and Mysterious Tradeoff between Inflation and Unemployment" *Economic Journal*, Vol. 111, C45-C61.

[41] Mankiw, N. Gregory, and Ricardo Reis (2002). "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, Vol. 117, No. 471, 1295-1328.

[42] Mankiw, N. Gregory, and Ricardo Reis (2003). "Sticky Information: A Model of Monetary Nonneutrality and Structural Slumps," in: Aghion, Philippe, Roman Frydman, Joseph Stiglitz, and Michael Woodford (eds.). *Knowledge, Information, and Expectations in Modern Macroeconomics.* Princeton: Princeton University Press.

[43] Mankiw, N. Gregory, and Ricardo Reis (2006). "Pervasive Stickiness," *American Economic Review*, Vol. 96, No. 2, 164-169.

[44] Mankiw, N. Gregory, and Ricardo Reis (2007). "Sticky Information in General Equilibrium," *Journal of the European Economic Assocation*, Vol. 5, Nos. 2-3, 603-613.

[45] Mayo, Deborah G., and Aris Spanos (2006). "Severe Testing as a Basic Concept in a Neyman-Pearson Philosophy of Induction," *The British Journal for the Philosophy of Science*, 1-35.

[46] McAdam, Peter, and Alpo Willman (2007). "State-Dependency and Firm-Level Optimization - A Contribution to Calvo Price Setting," ECB Working Paper No. 806, Frankfurt.

[47] Menzies, Gordon D., and Daniel J. Zizzo (2009). "Inferential Expectations," *The B.E. Journal of Macroeconomics*, forthcoming.

[48] Menzies, Gordon D., and Daniel J. Zizzo (2007). "Exchange Rate Markets and Conservative Inferential Expectations," *mimeo*, University of Technology, Sydney, February.

[49] Midrigan, Virgiliu (2008). "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations," *mimeo*, New York University.

[50] Nakamura, Eki, and Jón Steinsson (2008). "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, Vol. 123, No. 4, 1415-1464.

[51] Reis, Ricardo (2006). "Inattentive Producers," *Review of Economic Studies*, Vol. 73, No. ??, 793-821.

[52] Rotemberg, Julio J. (1982). "Monopolistic price adjustment and aggregate output," *Review of Economic Studies*, Vol. 49, 517-531.

[53] Rumler, F., and J. Vilmunen (2005). "Price setting in the euro area: some stylized facts from micro consumer data," ECB Working Paper #524.

[54] Sims, Christopher A. (1998). "Stickiness," Carnegie-Rochester Conference Series on Public Policy, Vol. 49, 317-356.

[55] Sims, Christopher A. (2003). "Implications of Rational Inattention," *Journal of Monetary Economics*, Vol. 50, 665-690.

[56] Sims, Christopher A. (2006). "Rational Inattention: A Research Agenda," *mimeo*, Princeton University, March.

[57] Taylor, John B. (1980). "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, Vol. 88(1), 1-24.

[58] Wolman, Alexander L. (2007). "The Frequency and Costs of Individual Price Adjustment," *Managerial and Decision Economics*, Vol. 28, 531-552.

[59] Woodford, Michael (2003). "Imperfect Common Knowledge and the Effects of Monetary Policy," in: Aghion, Philippe, Roman Frydman, Joseph Stiglitz, and Michael Woodford (eds.). *Knowledge, Information, and Expectations in Modern Macroeconomics.* Princeton: Princeton University Press.

[60] Woodford, Michael (2008a). "Inattention as a Source of Randomized Discrete Adjustment," *mimeo*, Columbia University, April.

[61] Woodford, Michael (2008b). "Information-Constrained State-Dependent Pricing," NBER Working Paper No. 14620, December.

[62] Zbaracki, Mark, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen (2004). "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," *Review of Economics and Statistics*, Vol. 86, 514-533.

6 Appendices

Appendix 1: Using Signal Extraction under H_0 In the text agents use rmc_{ss} under H_0 as an estimate for rmc_t in (11). Alternatively, they could use the minimum MSE estimator:

$$\begin{aligned} MSE \equiv E(\Gamma^2) &= E(rmc_{Lucas} - rmc_t)^2 \\ &= E(\Omega rmc_{ss} + (1-\Omega)\widehat{rmc}_{it} - rmc_t)^2 \\ &= E\left[(1-\Omega) rmc_{ss}\bar{v}_{it} - \Omega rmc_{ss}(e_t)\right]^2 \end{aligned}$$

Now $E(\Gamma) = 0$, so

$$E(\Gamma^2) = Var(\Gamma) = (1-\Omega)^2 rmc_{ss}^2 \frac{\sigma^2}{n} + \Omega^2 rmc_{ss}^2 \psi\sigma^2.$$

Setting $dE(\Gamma^2)/d\Omega = 0$, we obtain

$$\Omega = \frac{1}{1+n\psi_t}$$

and

$$1-\Omega = \frac{n\psi}{1+n\psi}.$$

When rmc_{Lucas} is used under the null, we obtain the following expression for the estimate of rmc under the null:

$$\begin{aligned} rmc_{Lucas} &= \Omega rmc_{ss} + (1-\Omega)\widehat{rmc}_{it} \\ &= rmc_{ss} + (1-\Omega) rmc_{ss}(1 + e_t + \bar{v}_{it}) \\ &= rmc_{ss}(1 + (1-\Omega)[e_t + \bar{v}_{it}]). \end{aligned}$$

This expression could be used in the null hypothesis price based on equation (11). That is, instead of using

$$\frac{\theta}{\theta-1} rmc_{ss}\left(1+\pi_{t-1}\right) p_{t-1}$$

we use

$$\frac{\theta}{\theta-1} rmc_{ss}(1 + (1-\Omega)\left(e_t + \bar{v}_{it}\right))\left(1+\pi_{t-1}\right) p_{t-1}.$$

The remainder of the derivation is unchanged, with the $\bar{v}_{it}$ terms integrated out over the economy. The resultant IE Phillips curve is given by

$$\pi_t = \beta\pi_{t-1} + \frac{\mu_t\alpha}{1-\mu_t\alpha}\widehat{rmc}_{it} + \frac{n\psi}{1+n\psi}\widehat{rmc}_{it},$$

which contains an additional $\widehat{rmc}_{it}$-term. The first rmc-term is a direct consequence of $\widehat{rmc}_{it}$ being used under H_1 while the second, new term appears because $\widehat{rmc}_{it}$ features in rmc_{Lucas}-solutionuses under H_0.

Appendix 2: Averaging Shocks Across an Economy We seek to prove that

$$\int_0^1 \left[\frac{\theta}{\theta-1}\widehat{rmc}_{it}p_t\right]^{1-\theta} di = \left[\frac{\theta}{\theta-1}rmc_t p_t\right]^{1-\theta}$$

Preliminary:

$$\int_0^1 v_i di = 0 \qquad \text{if} \qquad v_i \sim i.i.d.(0,\sigma^2)$$

Proof: As a matter of definition,

$$\int_{i=0}^1 v_i di \equiv \lim_{r\to\infty} \sum_{i=1}^r v_{\left(\frac{i}{r}\right)}\frac{1}{r},$$

but changing the index increment of a countable sum does not change the sum, so that this sum clearly must equal zero:

$$\lim_{r\to\infty} \sum_{i=1}^r v_{\left(\frac{i}{r}\right)}\cdot\frac{1}{r} = \lim_{r\to\infty} \sum_{i=1}^r v_i\frac{1}{r} = \lim_{r\to\infty} \frac{\sum_{i=1}^r v_i}{r} = \lim_{r\to\infty} \overline{v} = 0 \qquad \text{Q.E.D.}$$

We now use this preliminary result to establish the main result:

$$\begin{aligned}\int_0^1 \left[\frac{\theta}{\theta-1}\widehat{rmc}_{it}p_t\right]^{1-\theta} di &= \int_0^1 \left[\frac{\theta}{\theta-1}rmc_{ss}(1+e_t+\overline{v}_{it})p_t\right]^{1-\theta} di \\ &= \int_0^1 \left(\frac{\theta}{\theta-1}rmc_{ss}p_t\right)^{1-\theta}(1+e_t+\overline{v}_{it})^{1-\theta} di \\ &= \left(\frac{\theta}{\theta-1}rmc_{ss}p_t\right)^{1-\theta}\int_0^1 (1+e_t+\overline{v}_{it})^{1-\theta} di\end{aligned}$$

Assume that the shocks to real marginal costs are small in the sense that

$$(1+e_t+\overline{v}_{it})^{1-\theta} \approx 1 + (1-\theta)(e_t+\overline{v}_{it})$$

$$
\begin{aligned}
\int_0^1 \left[\frac{\theta}{\theta-1}\widehat{rmc}_{it}p_t\right]^{1-\theta} di &\approx \left[\frac{\theta}{\theta-1}rmc_{ss}p_t\right]^{1-\theta} \int_0^1 (1+(1-\theta)(e_t+\bar{v}_{it}))di \\
&= \left[\frac{\theta}{\theta-1}rmc_{ss}p_t\right]^{1-\theta} \left(1+(1-\theta)\left(e_t+\int_0^1 \bar{v}_{it}di\right)\right) \\
&= \left[\frac{\theta}{\theta-1}rmc_{ss}p_t\right]^{1-\theta} \left(1+(1-\theta)\left(e_t+\int_{i=0}^1 \frac{1}{n}\sum_{j=1}^n v_{ijt}di\right)\right) \\
&= \left[\frac{\theta}{\theta-1}rmc_{ss}p_t\right]^{1-\theta} \left(1+(1-\theta)\left(e_t+\frac{1}{n}\sum_{j=1}^n \int_{i=0}^1 v_{ijt}di\right)\right) \\
&= \left[\frac{\theta}{\theta-1}rmc_{ss}p_t\right]^{1-\theta} \left(1+(1-\theta)\left(e_t+\frac{1}{n}\sum_{j=1}^n 0\right)\right) \\
&= \left[\frac{\theta}{\theta-1}rmc_{ss}p_t\right]^{1-\theta} (1+(1-\theta)e_t)
\end{aligned}
$$

The second last line comes from the preliminary result. We now 'go back' using the small e_t assumption to reintroduce the exponential function. Note, however that this does not 'undo' the earlier approximation; rather it is another approximation, possibly taking us further away from the strictly correct expression:

$$
\begin{aligned}
\int_0^1 \left[\frac{\theta}{\theta-1}\widehat{rmc}_{it}p_t\right]^{1-\theta} di &\approx \left[\frac{\theta}{\theta-1}rmc_{ss}p_t\right]^{1-\theta} (1+(1-\theta)e_t) \\
&\approx \left[\frac{\theta}{\theta-1}rmc_{ss}p_t\right]^{1-\theta} (1+e_t)^{1-\theta} \\
&= \left[\frac{\theta}{\theta-1}rmc_{ss}(1+e_t)p_t\right]^{1-\theta} \\
&= [\frac{\theta}{\theta-1}rmc_t p_t]^{1-\theta} \qquad \text{Q.E.D.}
\end{aligned}
$$

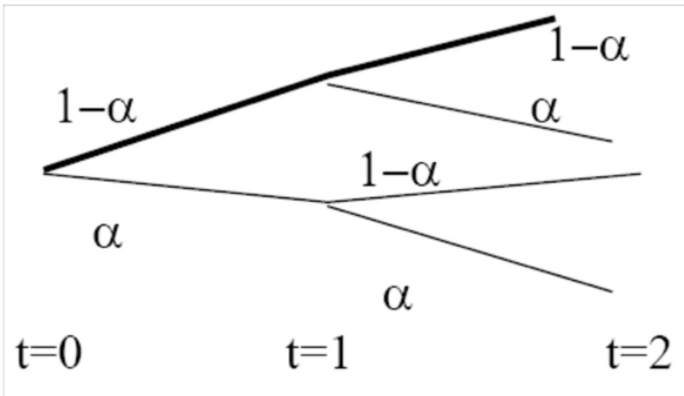

Figure 8: Possible States of Nature

Appendix 3: Maximizing an Intertemporal Stream of Profits In this appendix we solve the price setting problem for a firm that maximizes an intertemporal stream of profits. Under H_0 we assume agents adopt the solution to the single period problem in (11); namely, replace rmc_t with rmc_{ss}, replace p_t with $(1+\pi_{t-1})p_{t-1}$, and, unlike the solution in the text, replace p_{t-1} with $p_{i,t-1}$.[17] We also assume the steady state relationship

$$\frac{\theta}{\theta-1}rmc_{ss}=1$$

holds, giving $p_{it}=(1+\pi_{t-1})p_{i,t-1}$ under H_0.

At the moment of rejecting H_0 firms believe they have a probability $1-\alpha$ that they will believe H_0 in each future period. This is $1-\alpha$ and not $1-\mu_t\alpha$ because they are doing future hypothesis tests under the H_0 value of ψ.[18]

At each point in time, the calculation of p^{H_1} need only take into account the state of nature along an arm of possibilities where the firm chooses not to reject H_0, just like in the Calvo problem. (See figure 8.)

A firm maximizes expected future profits across all the possibilities. However, as far as choosing prices now are concerned, only the bold arm is relevant. Choices made off the bold arm take the agent to a place where a new optimal price is calculated *de novo*, without any impact from the period zero price. Thus, in the maximization problem, all the optimally chosen probability-weighted future profits off the bold arm are constants, and can be ignored.

Equation (7) is still relevant and repeated here for convenience,

[17] These 'ultra-lazy' agents do not observe last period's aggregate price and instead use their own. From a modeling point of view, this is necessary to cast the problem in the Calvo mould.

[18] Of course, one could argue that a RE agent would know that a large future shock would enable them to optimize.

$$p_t^{1-\theta} = (1-\mu_t\alpha)\int_0^1 p_{it}^{H_0\,1-\theta} di + \mu_t\alpha \int_0^1 p_{it}^{H_1\,1-\theta} di. \qquad (7)$$

We now choose p^{H_1} to maximize the current and future stream of profits. (The index 'j' in the following expressions is not the same as the index over n pieces of information in the main text.)

$$\Theta = \sum_{j=0}^{\infty} (1-\alpha)^j \frac{(p_{i,t+j} - mc_{i,t+j})\, y_{i,t+j}}{p_{t+j}}. \qquad \text{(A3.1)}$$

The variable $p_{i,t+j}$ is complex. Its value in period t is $p_t^{H_1}$ and that is what must be chosen to maximize Θ. However, after that, it evolves according to lagged inflation indexation under H_0 along the bold arm of Figure 8. Algebraically, this means,[19]

$$p_{i,t+j} = p_{it}^{H_1}(1+\pi_t)(1+\pi_{t+1})\cdots(1+\pi_{t+j-1}), \qquad j > 0.$$

Now let the relative price between the H_1 and the aggregate price in period t be $\rho_{it} = p_{it}^{H_1}/p_t$. We can simplify $p_{i,t+j}/p_{t+j}$ as follows:

$$\frac{p_{i,t+j}}{p_{t+j}} = \rho_{it} \qquad \text{if } j = 0;$$

$$\begin{aligned} \frac{p_{i,t+j}}{p_{t+j}} &= \frac{p_{i,t}^{H_1}(1+\pi_t)(1+\pi_{t+1})\cdots(1+\pi_{t+j-1})}{p_t(1+\pi_{t+1})\cdots(1+\pi_{t+j-1})(1+\pi_{t+j})} \qquad \text{if } j > 0 \qquad \text{(A3.2)} \\ &= \frac{p_{i,t}^{H_1}(1+\pi_t)}{p_t(1+\pi_{t+j})} \equiv \rho_{it}\delta_{t+j} \end{aligned}$$

As before, we re-write the objective function in real terms.[20]

$$\Theta = \sum_{j=0}^{\infty} (1-\alpha)^j \{\rho_{it}\delta_{t+j} - rmc_{i,t+j}\}\, y_{i,t+j} \qquad \text{(A3.3)}$$

The aggregator's demand for intermediate good i in period $t+j$ produced by a firm choosing prices optimally will be given by (2) with $\rho_{it}\delta_{t+j}$ replacing the relative price. Substituting (2) for y_{it+j} in (A3.3) gives

$$\Theta = \sum_{j=0}^{\infty} (1-\alpha)^j \left\{\delta_{t+j}^{1-\theta}\rho_{it}^{1-\theta} - \delta_{t+j}^{-\theta}\rho_{it}^{-\theta} rmc_{i,t+j}\right\} y_{t+j}. \qquad \text{(A3.4)}$$

[19] Note that the following expression would not be true off the bold arm in Figure 1, and, that if we had an additional category of H_0 agents as in the main model who do not index, then the choice of price now would stand forever, and we would have to calculate the maximization problem for those agents too.

[20] We ignore discounting because $(1-\alpha)$ acts on the problem in the same way. Adding a discount factor is straightforward.

Differentiation w.r.t. ρ_{it} yields[21]

$$\frac{d\Theta}{d\rho_{it}} = \sum_{j=0}^{\infty} (1-\alpha)^j \left\{ (1-\theta)\delta_{t+j}^{1-\theta}\rho_{it}^{-\theta} + \theta\delta_{t+j}^{-\theta}\rho_{it}^{-\theta-1} rmc_{i,t+j} \right\} y_{t+j} \qquad \text{(A3.5)}$$

Multiply by ${\rho_{it}}^{1+\theta}/(1-\theta)$ and setting the expression equal to zero gives the necessary FOC for a maximum:

$$\sum_{j=0}^{\infty} (1-\alpha)^j \left\{ \delta_{t+j}\rho_{it} + \frac{\theta}{1-\theta} rmc_{i,t+j} \right\} \delta_{t+j}^{-\theta} y_{t+j} = 0 \qquad \text{(A3.6)}$$

The current relative price can be taken out of the summation, as it does not depend on j. This gives us the optimal relative price for good i in period t:

$$\rho_{it} = \frac{\theta}{\theta-1} \frac{\sum_{j=0}^{\infty} (1-\alpha)^j\, rmc_{i,t+j}\delta_{t+j}^{-\theta} y_{t+j}}{\sum_{j=0}^{\infty} (1-\alpha)^j\, \delta_{t+j}\delta_{t+j}^{-\theta} y_{t+j}} \qquad \text{(A3.7)}$$

(Note that $\theta > 1$, so $\theta - 1$ is positive.)

If model solutions are provided for the RHS variables, (A3.7) is the model--consistent RE solution for ρ_{it}. We now linearize it around the steady state. Let $Y_{t+j} = \delta_{t+j}^{-\theta} y_{t+j}$. (We let the log deviation of y from its steady state be $\tilde{y}$ since capitals in exponents are cumbersome. The terms will disappear anyway.) Thus,

$$\rho_{it} = \rho_{ss} e^{\tilde{\rho}_{it}}, \qquad rmc_{it} = rmc_{ss} e^{\widetilde{rmc}_{it}}, \qquad \delta_t = \delta_{ss} e^{\tilde{\delta}_t}, \qquad Y_t = Y_{ss} e^{\tilde{y}_t},$$

and we can write

$$\rho_{ss} = \frac{\theta}{\theta-1} \frac{\sum_{j=0}^{\infty} (1-\alpha)^j\, rmc_{ss} Y_{ss}}{\sum_{j=0}^{\infty} (1-\alpha)^j\, \delta_{ss} Y_{ss}}$$

$$\Rightarrow \rho_{ss} = \frac{\theta}{\theta-1} \frac{rmc_{ss} \sum_{j=0}^{\infty} (1-\alpha)^j}{\delta_{ss} \sum_{j=0}^{\infty} (1-\alpha)^j} = \frac{\theta}{\theta-1} \frac{rmc_{ss}}{\delta_{ss}}.$$

In the SS both the relative price ρ and δ will be unity, so

[21] The second derivative of (A3.4) is negative, fulfilling the second-order-condition for a local maximum.

$$\frac{\theta}{\theta-1}rmc_{ss}=1,$$

as before. Making use of the above definitions of the variables, (A3.7) may be rewritten as

$$\rho_{ss}e^{\tilde{\rho}_{it}}=\frac{\theta}{\theta-1}\frac{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,rmc_{ss}e^{\widetilde{rmc}_{i,t+j}}Y_{ss}e^{\tilde{y}_{t+j}}}{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,\delta_{ss}e^{\tilde{\delta}_{t+j}}Y_{ss}e^{\tilde{y}_{t+j}}}$$

which simplifies to

$$e^{\tilde{\rho}_{it}}=\frac{\theta}{(\theta-1)}rmc_{ss}\frac{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,e^{\widetilde{rmc}_{i,t+j}}e^{\tilde{y}_{t+j}}}{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,e^{\tilde{\delta}_{t+j}}e^{\tilde{y}_{t+j}}}=\frac{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,e^{\widetilde{rmc}_{i,t+j}}e^{\tilde{y}_{t+j}}}{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,e^{\tilde{\delta}_{t+j}}e^{\tilde{y}_{t+j}}}$$

$$\Rightarrow 1+\tilde{\rho}_{it}=\frac{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,(1+\widetilde{rmc}_{i,t+j}+\tilde{y}_{t+j})}{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,(1+\tilde{\delta}_{t+j}+\tilde{y}_{t+j})}$$

We now need to expand $\tilde{\rho}_{it}$. Recall that it is the ratio of chosen prices to general prices in period t, log differences w.r.t. to the SS give us the following:

$$\rho_{it}=\frac{p_{it}^{H_1}}{p_t}\qquad\Rightarrow\tilde{\rho}_{it}=\tilde{p}_{it}^{H_1}-\tilde{p}_t$$

We then thus write,

$$1+\tilde{p}_{it}^{H_1}-\tilde{p}_t=\frac{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,(1+\widetilde{rmc}_{i,t+j}+\tilde{y}_{t+j})}{\sum_{j=0}^{\infty}(1-\alpha)^{j}\,(1+\tilde{\delta}_{t+j}+\tilde{y}_{t+j})}.\tag{A3.8}$$

Noting that δ_t is equal to unity by definition, the deviation from SS of δ_{t+j} must be zero when $j=0$. Thus,

$$\sum_{j=0}^{\infty}(1-\alpha)^{j}\,(1+\tilde{\delta}_{t+j}+\tilde{y}_{t+j})=\frac{1}{\alpha}+\sum_{j=1}^{\infty}(1-\alpha)^{j}\,\tilde{\delta}_{t+j}+\sum_{j=0}^{\infty}(1-\alpha)^{j}\,\tilde{y}_{t+j}.$$

Using this in (A3.8) yields

$$1+\tilde{p}_{it}^{H_1}-\tilde{p}_t=\frac{1+\alpha\sum\limits_{j=0}^{\infty}(1-\alpha)^j_{i,t+j}\,\widetilde{rmc}_{i,t+j}+\alpha\sum\limits_{j=0}^{\infty}(1-\alpha)^j\,\tilde{y}_{t+j}}{1+\alpha\sum\limits_{j=1}^{\infty}(1-\alpha)^j\,\tilde{\delta}_{t+j}+\alpha\sum\limits_{j=0}^{\infty}(1-\alpha)^j\,\tilde{y}_{t+j}}$$

Since all the deviations from SS are small, in the mathematical sense, the Y-deviations cancel (after all that!) and we are left with

$$\tilde{p}_{it}^{H_1}-\tilde{p}_t=\alpha\sum_{j=0}^{\infty}(1-\alpha)^j_{i,t+j}\,\widetilde{rmc}_{i,t+j}-\alpha\sum_{j=1}^{\infty}(1-\alpha)^j\,\tilde{\delta}_{t+j} \tag{A3.8}$$

Now define the (across firms) average H_1 price deviation, $\tilde{p}_t^{H_1}$. (Note the absence of the i-subscript). Since $\widehat{rmc}_{it}=rmc_{ss}(1+e_t+\bar{v}_{it})$, the log-deviation of this from its SS is $\widetilde{rmc}_{it}=e_t+\bar{v}_{it}$.[22] Integrate both sides of (A3.8) over $i\in[0,1]$ and note that $\bar{v}_{it}=0$ (Appendix 2) to obtain

$$\tilde{p}_t^{H_1}=\tilde{p}_t+\alpha\sum_{j=0}^{\infty}(1-\alpha)^j\,\widetilde{rmc}_{t+j}-\alpha\sum_{j=1}^{\infty}(1-\alpha)^j\,\tilde{\delta}_{t+j},$$

where $\widetilde{rmc}_{t+j}=e_{t+j}$. The above equation can be rewritten as a first-order (forward) difference equation:

$$\tilde{p}_t^{H_1}-\tilde{p}_t=(1-\alpha)\left(\tilde{p}_{t+1}^{H_1}-\tilde{p}_{t+1}\right)+\alpha\left(\widetilde{rmc}_t-(1-\alpha)\tilde{\delta}_{t+1}\right) \tag{A3.10}$$

To obtain our Phillips curve, first linearize (7):

$$\begin{aligned}p_t^{1-\theta}&=(1-\mu_t\alpha)\int_0^1 p_{it}^{H_0\,1-\theta}di+\mu_t\alpha\int_0^1 p_{it}^{H_1\,1-\theta}di\\&=(1-\mu_t\alpha)(1+\pi_{t-1})^{1-\theta}\int_0^1 p_{i,t-1}^{1-\theta}di+\mu_t\alpha\int_0^1 p_{it}^{H_1\,1-\theta}di\\&=(1-\mu_t\alpha)(1+\pi_{t-1})^{1-\theta}p_{t-1}^{1-\theta}+\mu_t\alpha\int_0^1 p_{it}^{H_1\,1-\theta}di\end{aligned}$$

We now need an expression for the average price under H_1 which is not a function of i and can be taken out of the integral and directly into a steady-state form:

$$\left(p_{ss}e^{\tilde{p}_t}\right)^{1-\theta}=(1-\mu_t\alpha)(1+\pi_{t-1})^{1-\theta}\left(p_{ss}e^{\tilde{p}_{t-1}}\right)^{1-\theta}+\mu_t\alpha\left(\bar{p}^{H_1}e^{\hat{p}_t^{H_1}}\right)^{1-\theta}$$

[22]Recall that hat represents an estimator and tilde a log deviation from the steady-state.

Divide by steady state $p_{ss}^{1-\theta}$ and note that relative prices are unity in the SS and inflation is close to zero:

$$\left(e^{\tilde{p}_t}\right)^{1-\theta} = (1-\mu_t\alpha)(1+\pi_{t-1})^{1-\theta}\left(e^{\tilde{p}_{t-1}}\right)^{1-\theta} + \mu_t\alpha\left(e^{\tilde{p}_t^{H_1}}\right)^{1-\theta}$$

After tedious algebra we obtain an expression for deviations in aggregate prices:

$$\tilde{p}_t = (1-\mu_t\alpha)\{\pi_{t-1}+\tilde{p}_{t-1}\} + \mu_t\alpha\tilde{p}_t^{H_1}.$$

Solve this for the (log difference of the) relative price of the H_1-price in preparation for substitution into (A3.9). In solving, we use $(1+\pi_t) = p_t/p_{t-1}$. Note that $\tilde{p}_t - \tilde{p}_{t-1} = \pi_t$ and let $(1-\mu_t\alpha) = \varepsilon$ which is the share of H_0 agents. (We suppress the t-subscript in ε_t until the end.) This gives rise to the following relationships:

$$\tilde{p}_t = \varepsilon\left(\pi_{t-1}+\tilde{p}_{t-1}\right) + (1-\varepsilon)\,\tilde{p}_t^{H_1}$$
$$\tilde{p}_t^{H_1} = \frac{\tilde{p}_t - \varepsilon\left(\pi_{t-1}+\tilde{p}_{t-1}\right)}{1-\varepsilon}$$

and

$$\tilde{p}_t^{H_1} - \tilde{p}_t = \frac{\varepsilon\left(\pi_t - \pi_{t-1}\right)}{1-\varepsilon} \tag{A3.11}$$

Substitute this into (A3.10) to give

$$\tilde{p}_t^{H_1} - \tilde{p}_t = (1-\alpha)\left(\tilde{p}_{t+1}^{H_1} - \tilde{p}_{t+1}\right) + \alpha\left(\widetilde{rmc}_t - (1-\alpha)\tilde{\delta}_{t+1}\right)$$

or

$$\frac{\varepsilon\left(\pi_t-\pi_{t-1}\right)}{1-\varepsilon} = (1-\alpha)\frac{\varepsilon\left(\pi_{t+1}-\pi_t\right)}{1-\varepsilon} + \alpha\left(\widetilde{rmc}_t - (1-\alpha)\tilde{\delta}_{t+1}\right).$$

Note that $\tilde{\delta}_{t+1} = \ln(\delta_{t+1}) - \ln(\delta_{ss}) = \ln(\delta_{t+1})$ and $\ln(\delta_{t+1}) = \ln(1+\pi_t) - \ln(1+\pi_{t+1}) = \pi_t - \pi_{t+1}$. On substitution we obtain a difference equation in π_t:

$$\frac{\varepsilon\left(\pi_t-\pi_{t-1}\right)}{1-\varepsilon} = (1-\alpha)\frac{\varepsilon\left(\pi_{t+1}-\pi_t\right)}{1-\varepsilon} + \alpha\left\{\widetilde{rmc}_t - (1-\alpha)\left(\pi_{t+1}-\pi_t\right)\right\}$$

Collecting terms,

$$\begin{aligned}&\varepsilon\pi_t + (1-\alpha)\,\varepsilon\pi_t + \alpha\,(1-\alpha)\,(1-\varepsilon)\,\pi_t \\ &\quad = \varepsilon\pi_{t-1} + (1-\alpha)\,\varepsilon\pi_{t+1} + \alpha\,(1-\alpha)\,(1-\varepsilon)\,\pi_{t+1} + \alpha(1-\varepsilon)\widetilde{rmc}_t\end{aligned}$$

Letting $\Phi \equiv (1-\alpha)\,\varepsilon + \alpha\,(1-\alpha)\,(1-\varepsilon)$, rewrite the above equation as

$$\pi_t = \frac{\varepsilon_t}{\varepsilon_t+\Phi}\pi_{t-1} + \frac{\Phi}{\varepsilon_t+\Phi}E_t\pi_{t+1} + \frac{\alpha\,(1-\varepsilon_t)}{\varepsilon_t+\Phi}\widetilde{rmc}_t. \tag{A3.12}$$

This is the forward-looking IE Phillips curve (with the t-subscript returned to remind us it is a state-dependent model) which bears some resemblance to the standard (hybrid) New Keynesian Phillips curve but also differs in some important aspects. When all agents reject the null, $\varepsilon = 0$ and $\Phi = \alpha(1-\alpha)$ and we get

$$\widetilde{rmc}_t = (1-\alpha)(\pi_{t-1} - \pi_{t+1}).$$

Interpreting this, substitute forward to obtain

$$\pi_t = \frac{1}{1-\alpha}\left(\widetilde{rmc}_t + \widetilde{rmc}_{t+1} + \widetilde{rmc}_{t+2} + \ldots\right)$$

The ratio $1/(1-\alpha)$ is the expected duration of being 'stuck' in H_0. Forward-looking agents (firms) in principle take account of the entire expected path of future marginal costs and the more inertia there is expected to be (because, say, α is low), the more they respond in aggregate.[23]

When $\alpha = 0$, $\Phi = 0$ and $\varepsilon = 1$, so

$$\pi_t = \pi_{t-1}.$$

That is, if H_0 is never rejected, inflation is fully indexed and becomes independent of current-period news.

[23] In principle, because our assumptions about the generation of rmc imply that expected future values of rmc are just equal to rmc_{ss}, the expected future deviations from the steady-state are zero.