UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Essays on Macroeconomic Expectations

Permalink
https://escholarship.org/uc/item/4mz3q3x8

Author
Caruso Bloeck, Martin

Publication Date
2024

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Essays on Macroeconomic Expectations

by

Martín Caruso Bloeck

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Yuriy Gorodnichenko, Chair

Professor David Romer

Professor Jón Steinsson

Spring 2024

Essays on Macroeconomic Expectations

Copyright 2024

by

Martín Caruso Bloeck

Abstract

Essays on Macroeconomic Expectations

by

Martín Caruso Bloeck

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Yuriy Gorodnichenko, Chair

This dissertation studies the economic effects of the formation of inflation expectations. It uses both theoretical and empirical tools to bridge gaps in our understanding of the process of expectation formation and its macroeconomic implications. Chapter 1 is motivated by increasing evidence that consumers have particularly limited knowledge about inflationary trends and that they rely heavily on their observation of everyday prices to form their inflation expectations. This chapter asks how the fact that consumers learn from prices affects the price-setting process and the implications for the allocative efficiency of the price system. In environments of high inflation uncertainty, prices provide mostly information about aggregate inflationary shocks and little information about relative prices. This feature makes consumers less price-sensitive in equilibrium, allowing firms to set higher markups. Additionally, these incentives are stronger for high-cost relative to low-cost firms, increasing cross-sectional price dispersion. These distortions to price-setting have first-order welfare implications and explain essentially all of the welfare losses in the model. Chapter 2 causally estimates the effect of information about a monetary policy reform on firm inflation expectations in Uruguay. A causal quantification of this effect has been previously infeasible given data considerations, despite its importance for the cost of disinflation in a broad class of New-Keynesian models. We overcome these problems by conducting a randomized control trial (RCT) providing information to a subset of firms shortly after the reform was announced. The results indicate a strong and long-lived effect of treatment on firm inflation expectations and that firms expect temporarily lower growth as a result of treatment. Taken together, these causal estimates imply a relatively low cost of disinflation. Finally, Chapter 3 inquires whether conventional measures of inflation expectations are a reasonable guide to determine the appropriate monetary policy stance in a model of endogenous expectation formation. The model generates forecast errors that are consistent with the empirical evidence, and I use the model to evaluate inflation and expectations dynamics under different monetary policy rules. Inflation-targeting rules can successfully keep expectations anchored by offsetting second-round effects of inflationary shocks. However, expectations-targeting rule cannot generally keep expectations anchored given that monetary policy can only lean against inflationary shocks if expectations become unanchored. In a dynamic extension, this results in

gradual inflation in inflation-targeting regimes, while inflation is volatile and front-loaded in expectations-targeting regimes. In all, these papers contribute to a growing and exciting field within macroeconomics.

To my father, mother, sister, wife, and daughter.

Acknowledgements

I am grateful to my advisors Yuriy Gorodnichenko, David Romer, and Jón Steinsson for continuous guidance, support, and dedication throughout my time at Berkeley. I also wish to thank Emi Nakamura, Ben Schoefer, Maury Obstfeld, and Pierre-Olivier Gourinchas whose engagement and enthusiasm in discussions, seminars, and classes have contributed enormously to my intellectual growth. I am also thankful to many friends I have made at UC Berkeley, especially Felipe Arteaga, Tomás Breach, Bernardo Candia, Andrea Cerrato, Abhi Gupta, Collin Jones, and Maximiliano Lauletta for support and camaraderie during the graduate program. Finally, above all I want to thank my wife for her unconditional support and continuous encouragement.

Contents

Introduction **1**

1 Informational Costs and Allocative Inefficiency of Inflation Uncertainty **5**
1.1 Introduction . . . 5
1.2 Partial Equilibrium Model . . . 8
1.2.1 Setup . . . 9
1.2.2 Equilibrium . . . 13
1.2.3 The Effects of Inflation Uncertainty . . . 18
1.3 Inflation Uncertainty in General Equilibrium . . . 24
1.3.1 Setup . . . 24
1.3.2 Equilibrium . . . 28
1.4 Welfare . . . 37
1.4.1 Welfare Criterion . . . 37
1.4.2 The Welfare Effects of Inflation Uncertainty . . . 38
1.4.3 Welfare under Counterfactual Behaviour of Households and Firms . . . 41
1.5 Conclusions . . . 50

2 News of Disinflation and Firms' Expectations: New Causal Evidence **51**
2.1 Introduction . . . 51
2.2 Related Literature . . . 53
2.3 The Randomized Control Trial Experiment . . . 55
2.3.1 Data Source: The Inflation Expectations Survey . . . 55
2.3.2 Monetary Policy Framework and Reform . . . 56
2.3.3 Treatment Description . . . 58
2.3.4 Pretreatment Characteristics . . . 59
2.4 Results . . . 60
2.4.1 Effect on Inflation Expectations . . . 60
2.4.2 Placebo Test . . . 64
2.4.3 Heterogeneous Effects . . . 65
2.4.4 Other Variables of Interest . . . 66
2.4.5 Cost Inflation Expectations . . . 67
2.5 Final remarks . . . 68

3 Unreliable Expectations **70**
3.1 Introduction . . . 70

3.2 Predictive Power of Professional Inflation Forecasts 73
3.2.1 Predictive Accuracy of Professional Forecasts 73
3.2.2 Inflation and Forecast Errors 74
3.2.3 Summary of Empirical Evidence 77
3.3 Monetary Policy Rules with Endogenous Expectations 78
3.3.1 Model 78
3.3.2 Equilibrium Analysis 80
3.3.3 Inflation Dynamics 83
3.4 Concluding Remarks 86

Bibliography **88**

Appendix **94**
A.1 Household Problem in General Equilibrium 94
A.2 Firm Problem in General Equilibrium 96
A.3 General Equilibrium Steady State 97
A.4 Solution to the General Equilibrium 99
A.5 Derivation of Welfare Function 100
A.6 Proof of Proposition 1.4.2 and 1.4.3 102

Introduction

There is hardly any area in macroeconomics in which expectations do not play a crucial role. In environments that are characterized by uncertainty, economic agents' decisions must rely on beliefs about the state of the world, events far into the future, the actions of their competitors, and many other factors that may influence their outcomes. Furthermore, the predictions of our macroeconomic models can vary substantially depending on how expectations are introduced. This highlights the importance of both understanding and accurately modeling the process of belief formation.

Inflation expectations are of utmost relevance for central banks, and likely for other policymakers as well. Inflation expectations are central to inflation dynamics in essentially all forward-looking macroeconomic models, and therefore the formation of expectations can either facilitate or hinder the accomplishment of central bank objectives. Moreover, the formation of inflation expectations is important for consumer behavior, both in the cross-section and over time.

This dissertation studies different aspects of the formation of beliefs about inflation. Two of the chapters contained within it are theoretical but motivated by empirical evidence, and a third chapter provides novel empirical evidence to address a theoretical question. As such, its content should be understood as an effort to fill in necessary gaps to integrate theory and empirics.

My first chapter is called *Informational Costs and Allocative Inefficiency of Inflation Uncertainty*. It is motivated by growing empirical evidence that consumers have rather limited knowledge about inflation and that they rely heavily on their observation of everyday prices such as gasoline or groceries to learn about price dynamics (Coibion & Gorodnichenko 2015*b*, Binder 2018, D'Acunto et al. 2021, Dietrich 2023). This chapter first asks: if consumers learn about inflation from prices, how do those prices arise in the first place? Second, what are the implications of these prices and this learning process for the allocation of resources and, ultimately, welfare?

To answer these questions, I develop a search model in which consumers have incomplete information about the distribution of prices. This means that when a consumer observes a price, she does not know with certainty whether it is low or high relative to unobserved prices because inflation confounds relative prices. Rather, consumers use observed prices to form beliefs about the distribution of prices, and this inference process is central to the behavior of consumers. Moreover, firms internalize how these frictions affect the behavior of consumers in their price setting. This introduces an interaction between consumers' information frictions and market structure.

The main mechanism in the model is that inflation uncertainty makes consumers' search responses less price-sensitive in equilibrium. Intuitively, consumers' search decisions depend on their capacity to infer the relative component from observed prices. Inflation uncertainty changes the informational content of prices. When inflation is predictable, prices contain mostly information about relative prices and consumers can make their search choices without making large errors. However, as inflation uncertainty rises, prices contain more information about the aggregate price level, making it more difficult for consumers to differentiate high- and low-price sellers.

This has two effects on pricing. First, markups rise because consumer search responses are a component of the demand elasticity. The rise in markups is an optimal response to the fall in the demand elasticity given the lower price-sensitivity of search. Second, cross-sectional price dispersion also rises. Intuitively, firms' demand elasticity depends on a weighted average of the elasticities of two types of consumers: flexible (searchers) and captive (non-searchers). High-cost firms lose flexible consumers easily, their customer base is skewed towards captive consumers, and they find it optimal to exploit their captive consumers by setting higher markups. Increasing inflation uncertainty makes it harder for consumers to infer that firms have high prices, increasing the weight of captive consumers in the customer base of high-cost firms even more. Similarly, consumers also have difficulty identifying low-cost firms who are therefore unable to retain captive consumers, and whose customer base is skewed towards flexible consumers. These differential effects across the price distribution raise cross-sectional price dispersion.

The second part of the chapter conducts a welfare analysis in a general equilibrium version of the model. In the welfare analysis, I show that equilibrium search effort is generally inefficient. By this, I mean that a perturbation to search behavior would generally have a first-order welfare effect, which may be positive or negative. The intuition is as follows. A perturbation to search behavior causes a second-order loss on consumers' optimization problem given that consumer utility is flat around the optimum. However, consumer search is guided by private incentives ruled by the endogenous distribution of relative prices. Because relative prices are distorted from relative costs, private incentives to search differ from social costs and search is inefficient.

Moreover, I show that the distortions that inflation uncertainty has on price-setting (markups and the cross-sectional price dispersion) are also first-order. Markups have a first-order welfare effect because they distort the relative price of consumption and labor, and create a wedge in households' consumption and labor optimization conditions. This results in inefficiently low labor and production in equilibrium.

Cross-sectional price dispersion has a first-order welfare effect because it distorts households' incentive to substitute in the cross-section of firms. A rise in cross-sectional price dispersion reallocates consumers to low-cost firms, but requires an increase in consumption. If the increase in consumption is larger than the efficiency gain, both consumption and labor rise with price dispersion. The overall welfare effect of rising price dispersion depends on the magnitude of the distortions that are already in place. If markups are high and consumption and labor are low, an increase in price dispersion that increases consumption and labor can be welfare improving. However, if price dispersion is already high, the increase in consumption and labor induced by higher price dispersion is welfare-reducing.

I end the paper illustrating the magnitude of these effects in a parametrized version of the model. Inflation uncertainty causes a reduction in search, a rise in markups, and a rise cross-sectional price dispersion. Because price dispersion is large, the reduction in search alone is welfare-improving. However, the distortions to price-setting generate sizable welfare losses due to resource misallocation and explain essentially all of the welfare effects.

My second chapter is called *News of Disinflation and Firms' Expectations: New Causal Evidence*, it is a coauthored paper with Miguel Mello and Jorge Ponce of the Central Bank of Uruguay (BCU), and is published in the *Journal of International Money and Finance*[1]. In this paper, we carry out a causal evaluation of an informational treatment about a change in the inflation target.

The motivation behind the paper is that in a broad class of macroeconomic models, the cost of disinflation depends crucially on how firms form expectations about future inflation. Although this is a fundamental question in macroeconomics, we lacked a causal evaluation of this expectational response for two main reasons. First, events in which central bankers exogenously carry out disinflationary efforts are extremely rare. Second, a causal analysis requires that surveys of firm inflation expectations be carried out in real-time, and these surveys have only recently become more widespread (see Candia et al. 2023).

We overcome both of these obstacles by conducting a randomized control trial (RCT) in a unique setting caused by a monetary policy reform in Uruguay. The reform consists of a significant overhaul of the monetary policy framework and was announced at the end of 2020. The most important element of the reform for our study is that the reform is intended to lower inflation significantly from the level of approximately 10% around which it had hovered in the previous decade to around 5%. This is a sizable disinflationary effort that is arguably exogenous. Second, Uruguay is one of the few countries that had a running firm inflation expectation survey at the time, in which we introduced the RCT.

The RCT consists of providing an informational treatment with the BCU forecast for inflation several years into the future to a subset of the survey respondents. The remaining respondents receive a placebo treatment. The random assignment of the informational treatment allows us to estimate its causal effect. We can estimate the effect on expectations up to six months after the treatment.

The results indicate that the informational treatment lowers inflation expectations by approximately half a percentage point at both the 12-month and 24-month horizons. Moreover, this effect persists until the end of our sample. Treated firms also anticipate a temporarily slower growth rate. These results resonate with the Keynesian intuition that disinflation is costly, although the implied cost of disinflation is relatively low given the change in inflation expectations (Hazell et al. 2022).

[1] See Caruso Bloeck et al. (2023)

My last chapter is called *Unreliable Expectations* and it asks whether conventional measures of inflation expectations are a suitable guide to determine the appropriate stance of monetary policy. The paper is motivated by previous work in Adler et al. (2022) in which we document a slow response of inflation expectations to inflation shocks in a panel of 36 emerging markets and advanced economies. In these circumstances, anchored expectations may prompt central bankers to adopt an accommodative monetary policy stance. However, the paper argues that this line of reasoning overlooks the fact that inflation expectations are themselves endogenous to the way monetary policy is carried out.

To answer this question, I develop a model of endogenous expectation formation that is consistent with the empirical evidence and use it to analyze how the formation of expectations changes depending on the type of monetary policy rule. In particular, I focus on comparing standard inflation targeting (IT) rules with rules that target inflation expectations (ET). The formation of expectations depends on the incentives that agents have to acquire accurate information about inflation dynamics. When inflation is highly volatile, agents have stronger incentives to form expectations accurately. However, under suitable monetary policy responses, the cost of acquiring information may be hiring than the cost of holding biased beliefs about inflation.

The main result is that IT regimes can always keep inflation expectations anchored, while ET regimes generally cannot. The intuition for this is keeping inflation expectations anchored requires monetary policy to lean against inflationary pressures. In the model, IT regimes can do this by compensating second-round effects of inflationary shocks. However, ET regimes cannot stabilize inflation unless expectations become unanchored, and hence require some level of attention to trigger a stabilizing monetary policy response.

The results question the usefulness of inflation expectations for monetary policy decisions. While inflation expectations may contain useful information, the results suggest that inflation expectations should not be taken as given and that careful consideration of their endogeneity is important.

These chapters provide new perspectives about the relevance of inflation expectation formation for macroeconomic dynamics. This is an exciting area of research that challenges our understanding of macroeconomics, provides novel insights into classical questions, and promises to be active in the future.

Chapter 1

Informational Costs and Allocative Inefficiency of Inflation Uncertainty

1.1 Introduction

Milton Friedman once used a pencil to illustrate the power of prices to direct allocations. While this beautiful mechanism is a marvel of the market economy, the efficiency with which prices allocate resources can break down when movements in relative and aggregate prices are confounded and, more generally, when economic agents are uncertain about future, current, or even past inflation.

Economists have long appreciated that these informational channels can degrade the allocative efficiency of the price system (e.g., Friedman 1977, Fischer 1981). However, a treatment of the interaction between firms' pricing incentives and consumers' formation of beliefs in inflationary environments that is tractable, portable, and sufficiently rich has proven elusive. In light of this, we lack a framework that allows us to comprehensively study the implications of these informational effects for pricing, and their broader macroeconomic implications.

To this end, this paper builds a tractable search model that combines (i) consumers with incomplete information who use observed prices to infer the distribution of prices and decide whether to search for substitutes, and (ii) firms that set prices strategically understanding their effect on consumers' beliefs and likelihood of searching. Because consumers base their search decisions on beliefs about the price distribution, the informational effect of prices enters directly into the elasticity of demand. In equilibrium, the informational content of prices and market structure are determined jointly by the optimal actions of both types of agents. The model is portable to any environment in which belief formation interacts with market structure and scales easily to the addition of more shocks, dynamics, or sequential learning. In my main application, I nest the search model into a general equilibrium framework to analyze how inflation uncertainty affects welfare through its effect on pricing and market structure.

The key channel through which inflation uncertainty affects market structure is by making consumers' search responses less price-sensitive and their demand more inelastic. Intuitively, the intensity with which consumers want to search depends on their belief about relative prices. As inflation uncertainty rises, prices contain more information about the price level and less information about relative prices. Beliefs about relative prices therefore become relatively unresponsive to observed prices and consumers differ less in their search response across producers. The diminished search response to prices allows firms to raise their markups because the possibility of losing consumers to search is a factor that enters the demand elasticity.

In addition, cross-sectional price dispersion typically rises with inflation uncertainty given the effect of observed prices on beliefs about the distribution of prices. This has opposite incentives in the price setting of high- and low-cost firms. High-cost firms find it easier to retain customers as observing high prices makes them more pessimistic about the prices they would find by searching. On the other hand, low-cost firms find it harder to dissuade their customers from searching and find it optimal to lower their prices to attract customers from other stores. These differential effects across the price distribution guide the allocative effects of prices, as high-cost firms extract a greater surplus from their customers while low-cost firms make a greater effort to attract customers.

After developing this intuition in a partial equilibrium setting, I extend the model to a general equilibrium framework with the objective of studying its welfare implications. By construction, the model is set up in such a way that inflation uncertainty has an insignificant effect on the volatility of real aggregate variables. Welfare losses attributable to business cycle fluctuations are therefore trivial as in Lucas (2003). Rather, welfare effects in the model arise mostly from the way in which households and firms change their behavior in response to inflation volatility rather than because monetary shocks have a large effect on real activity. To isolate the welfare effect of the agents' behavior, I compute the welfare effect of perturbations to the behavior of households and firms holding everything else constant.

The central result of the welfare analysis is that distortions in the behavior of households and firms caused by inflation uncertainty have first-order welfare effects. First, consumer search is inefficient. Intuitively, consumers search until the private marginal benefit equals the private marginal cost. Given that inflation uncertainty distorts relative prices, private marginal benefits of search are not aligned with the social marginal cost of reallocating sales across producers. Perturbations to search behavior therefore cause first-order welfare effect due to changes in allocative efficiency. These allocative effects trump losses due to consumer optimization errors, which are second-order small given that expected utility is flat around the optimum.

Second, the rise in markups causes a first-order welfare loss due to a distortion of households' consumption and labor optimization conditions. Rising markups reduce incentives to work, lowering equilibrium production and consumption. The welfare effect of these changes is negative because the marginal utility of consumption is higher than the marginal disutility of labor given depressed real wages.

Finally, the effect of price dispersion on welfare depends on the relative magnitude of the distortions that are in place. An increase in price dispersion causes an increase in consumption due to a higher likelihood of searchers finding low prices. Generally, this also requires an increase in labor because the relative prices typically overstate relative costs. The increase in labor and consumption may therefore offset the distortion generated by markups. However, a high level of price dispersion causes cross-sectional misallocation by generating excessive private incentives to search. If the latter distortion is sufficiently severe, the effect of rising price dispersion is a first-order welfare loss.

I perform several numerical exercises quantifying the strength of the different channels. This analysis reveals that changes in markups and increasing price dispersion are the main channels through which inflation uncertainty affects welfare. Together, they account for essentially all of the welfare loss generated by inflation uncertainty. The remaining channels have minor effects on welfare. Particularly, monetary shocks have a negligible contribution to aggregate volatility. Taken together, these results indicate that inflation uncertainty can have meaningful welfare effects even if money is close to neutral due to a deterioration of resource allocation.

The main contribution of my paper is to develop a theory in which the costs of inflation volatility arise from degrading the informational role of prices. This resembles the seminal work of Lucas (1972) and more recent work analyzing the macroeconomic effects of dispersed information (Woodford 2001, Angeletos & La'O 2010, Angeletos & Lian 2016, 2018) and learning from prices or other aggregate variables (Lucas 1972, Angeletos & La'O 2008, Amador & Weill 2010, Lorenzoni 2009, 2010). The key distinction with this body of work is that prices arise endogenously as an optimal response to the process of consumer belief formation, rather than aggregating dispersed information. As a result of this process, the endogenous price distribution plays a central role in the allocation of resources, both affecting the consumption-labor decisions and the cross-sectional allocation between firms.

The relation between inflation uncertainty and market competitiveness is the subject of an older literature. Benabou & Gertner (1993) develop a non-linear search model in which consumers infer common and idiosyncratic cost shocks from prices as in my framework. That model can only be solved numerically, making it difficult to draw sharp analytical results. By contrast, the linear structure of my model allows me to conduct a transparent analysis of the key parameters behind each equilibrium element. Moreover, the linear structure of the model makes the distribution of beliefs tractable, making it easier to scale to richer environments, such as including more shocks, dynamics, or more general forms of sequential learning.

Other papers that analyze the implications of informational effects of prices are Tommasi (1994) and Ball & Romer (2003). Relative to these, my model features greater internal consistency by explicitly solving for optimal pricing and belief formation that are best responses to each other. In addition, these papers are set up in customer market environments and a key assumption is that the information embedded in prices has dynamic value. My model requires arguably weaker assumptions given that the possibility of cross-sectional substitution through search suffices to introduce an informational role of prices.

Finally, my paper is related to a large body of work studying the relationship between inflation and price dispersion (Van Hoomissen 1988, Nakamura et al. 2018, Drenik &

Perez 2020, Sheremirov 2020, Sara-Zaror 2021). These papers differ in the source of price dispersion, which can be nominal rigidities (e.g., Nakamura et al. 2018, Sara-Zaror 2021), dispersed information (Drenik & Perez 2020), or search frictions (Sara-Zaror 2021). In this paper, price dispersion arises from a combination of informational effects of prices and search frictions. The introduction of search frictions makes my model resemble recent papers in which price setting depends on the demand composition of different types of consumers who differ in their elasticity (e.g., Pytka 2018, Kaplan et al. 2019, Nord 2022). The key difference with these papers is that the compositional effects are endogenous to the information embedded in prices in my model. To introduce this feature, I develop a novel protocol of sequential search in which the problem of belief formation is tractable. This makes the model scalable to additional instances of learning, such as dynamics or more flexible shopping patterns. Furthermore, my framework includes cross-sectional heterogeneity across producers, which allows me to examine implications for cross-sectional allocative efficiency.

The remainder of the paper is structured as follows. Section 1.2 develops the partial equilibrium model and analyzes the equilibrium effects of inflation uncertainty. Next, Section 1.3 extends the model to a general equilibrium framework by adding production, labor supply, and exogenous money supply shocks that are the source of inflation. I analyze the welfare implications of the model in Section 1.4. Finally, Section 1.5 offers some concluding remarks. Proofs and derivations are featured in the appendix, while complimentary analysis is available in Caruso Bloeck (2024).

1.2 Partial Equilibrium Model

This section develops the model in a partial equilibrium setting. I will provide a concise overview of the model before describing the assumptions in detail.

The model features two types of agents: consumers and stores. Store costs have two components: one that is common across stores and one that is idiosyncratic. Because of the latter, stores will differ in the prices they set. However, store prices will be correlated because of the common component of costs.

Consumers begin the period by visiting a store and observing the price. If they think the price of that store is high, they can choose to search at a different store. However, consumers have incomplete information about the distribution of prices, i.e., they know the likelihood of observing prices but they do not know whether prices are high or low across the board. Given that common cost components imply prices are correlated, consumers update their beliefs about the distribution of prices in the direction of observed prices.

Stores maximize profits by setting optimal markups given their demand elasticity. The elasticity of demand depends on the likelihood that consumers search, which itself depends on the information consumers infer from prices. The equilibrium therefore requires solving for an optimal pricing scheme and consumer inference that are best responses to each other.

This section develops the model and solves for the equilibrium. The model is written in linear form, allowing me to develop a clear intuition about the economic mechanisms behind the results. The linear model is also easily scalable to a general equilibrium setting, which is the subject of Section 1.3.

Section 1.2.3 examines the effect of rising inflation uncertainty on the equilibrium which is the primary focus of this paper.

1.2.1 Setup

The model features firms (which I will call stores) and consumers. Stores face shocks to their costs, and these shocks have a common component that affects all stores and idiosyncratic components. Consumers have incomplete information about the distribution of costs, which they infer before deciding between accepting the price quotes they received or searching for an alternative quote. Stores understand the nature of this inference problem and the effect of quote prices on the likelihood that consumers search and incorporate it into their pricing decisions.

Consumers. There is a mass one of consumers indexed $i \in (0, 1)$. Consumers have constant relative risk aversion utility over consumption:

$$U(\mathcal{C}_i) = \frac{\mathcal{C}_i^{1-\gamma} - 1}{1 - \gamma}$$

where γ is the coefficient of relative risk aversion. The assumption regarding CRRA preferences is not crucial. The main reason for including it is that it makes the analysis as comparable as possible to the general equilibrium framework developed in the following section. However, consumers would face the same kind of inference problem regardless of their preferences, and the main results about market structure would hold, for example, if consumers were risk-neutral.

Consumers have a common disposable income of W and their main decision is where to shop. They are initially uniformly and randomly assigned to stores $f(i)$, observe the level of prices set by that store $P_{f(i)}$ and decide whether to make all their purchases at that store or to search. Therefore, consumption for a consumer who shops at her first store is:

$$\mathcal{C}_i^C = \frac{W}{P_{f(i)}}$$

where the superscript C stands for "captive." The demand from captive consumers has unitary elasticity. This means that stores only facing captive consumers would find it optimal to raise their prices infinitely since they would minimize their production costs given the constant revenue of W.

Consumers also have the option to search for substitutes at other stores, in which case they draw a second price $P_{s(i)}$ where the subscript s stands for substitute. I will make the following distributional assumption regarding the probability density function (p.d.f.) of prices drawn from search:

$$g(P_s) \propto P_s^{1-\varsigma}$$

where $\varsigma > 1$. This assumption indicates that searchers are more likely to obtain draws from stores setting low prices. A possible interpretation is that consumers acquire some information about prices and are more likely to select stores with lower prices. Integrating over the distribution of prices, the p.d.f. can be written as:

$$g(P_s) = \left(\frac{P_s}{P}\right)^{1-\varsigma}$$

where $P = \left(\int P_f^{1-\varsigma} df\right)^{1/(1-\varsigma)}$. The setup resembles a constant elasticity of substitution (CES) demand system where the elasticity of substitution is given by ς and P is the CES price index. I will call searching consumers "flexible" given that $\varsigma > 1$ implies their demand is more elastic than that of captive consumers.

The main benefit of introducing search this way is to generate convenient analytical expressions for the demand faced by firms. Store demand will be composed of captive and flexible consumers and it is assumed that stores cannot discriminate between the two. However, stores anticipate how consumers responds to their prices, which will depend on a single behavioral parameter for each type. In the case of flexible consumers, the parameter is the elasticity of substitution parameter ς. For captive consumers, whose demand elasticity was already assumed to be unitary, the central parameter will rule consumers' marginal propensity to search and it will be introduced further below as μ.

By comparison, alternative assumptions would generally not lead to equally simple expressions, despite their popularity. For example, the assumption that stores sell identical goods and could return to their first store would make demand depend on the hazard rate of the cost distribution. Alternatively, if consumers could combine purchases at their first store with purchases at other stores, demand would depend on consumption shares that would arise from combinations of stores. In my setting, the demand elasticity of individual consumers is either one or ς, and the consumer search problem depends solely on expectations about the mean of the distribution.

Searching and becoming a flexible consumer is costly. I assume search costs have iceberg form such that $\tau > 0$ units are purchased per unit consumed. The choice of iceberg costs is a convenient way to introduce search costs in such a way that these neither reveal information about prices nor affect the relative prices between observed and unobserved stores. Moreover, I will introduce an idiosyncratic non-pecuniary factor of consumption $A_i \geq 0$ to generate smooth search responses to prices. I will assume A_i has an exponential distribution with shape parameter $\mu > 0$. This distributional assumption generates convenient analytical expressions in discrete choice settings, with the shape parameter ruling the sensitivity of choices to differences in the attractiveness of different choices. In my setting, the parameter μ determines the degree to which consumers respond to higher observed prices by searching, the reason for which I call this parameter the "search elasticity." I do not introduce a scale parameter because it would be unidentified given the iceberg cost τ. This means that an alternative interpretation of the model is that it is search costs that are idiosyncratic.

Hence, consumption for flexible consumers is:

$$\mathcal{C}_i^F = \frac{W}{\tau P_{s(i)}}$$

Consumers will search if:

$$U(\mathcal{C}_i^C) \leq \mathbb{E}_i \left[U(A_i \mathcal{C}_i^F) \right]$$

where $\mathbb{E}_i \left[U(\mathcal{C}_i^F) \right]$ is the expected utility from search given the information held by consumer i. Because consumers obtain information from the price they observe at the store they are originally assigned to, it will be useful to define $V(P_f)$ as:

$$V(P_f) \equiv \mathbb{E} \left[P_s^{-(1-\gamma)} | P_f \right]^{-1/(1-\gamma)}$$

The expression $\mathbb{E} \left[P_s^{-(1-\gamma)} | P_f \right]$ is the expected utility from one unit of income given $A_i/\tau = 1$. Raising this value to the power $-1/(1-\gamma)$ gives a price that would make consumers indifferent between searching and not searching. In other words, $V(P_f)$ is the certainty-equivalent price of searching given expectations elicited by the price P_f and absent costs or preferences for search.

Hence, a consumer will search if:

$$A_i \geq \frac{V(P_f)\tau}{P_f} \tag{1.1}$$

Given the assumption that A_i has an exponential distribution, the probability that a consumer searches is:

$$Pr(S_i = 1 | P_f) = \exp \left\{ -\mu\tau \frac{V(P_f)}{P_f} \right\} \tag{1.2}$$

where S_i equals one if consumer i searches and zero otherwise.

The right-hand side of Equation (1.1) is a cutoff that determines which consumers search and which consumers do not. The cutoff is simply the ratio between the certainty-equivalent price of searching $V(P_f)$ adjusted for search costs τ and the price set by the firm P_f. Given that $V(P_f)$ depends on expectations elicited by the price P_f, the cutoff will depend on beliefs about the observed price relative to the average price in the distribution.

Stores. There is a mass one of stores indexed by $f \in (0, 1)$, and stores are the suppliers in the economy. Stores have constant marginal costs $C_f = CE_f$, where C is a common component of costs and E_f is an idiosyncratic component.

Stores are profit maximizers and therefore their optimal prices are:

$$P_f = \frac{\varepsilon_f}{\varepsilon_f + 1} C_f \tag{1.3}$$

where $\varepsilon_f = \frac{\partial q_f}{\partial P_f}\frac{P_f}{q_f}$ is the elasticity of demand, where q_f are units sold by store f. Note that under this sign convention, $\varepsilon_f < 0$.

Each store receives one unit of consumers who observe the price and decide whether to purchase or search. Let S_f be the fraction of consumers who search. Using the expression above for the probability of search, the fraction of consumers who do not search and therefore become captive is

$$1 - S_f = 1 - \exp\left\{-\mu\tau\frac{V(P_f)}{P_f}\right\} \tag{1.4}$$

Captive consumers spend all of their income at store f. Quantities sold to captive consumers q_f^C equal:

$$q_f^C = (1 - S_f)\frac{W}{P_f} \tag{1.5}$$

Firms also sell to searchers who are matched to stores. Let $S \equiv \int S_f df$ be the aggregate share of flexible consumers. Given the assumption that determines how flexible consumers are matched to firms, the fraction of flexible consumers attracted by a firm setting price P_f is $(P_f/P)^{1-\varsigma}$ and the expenditure of flexible consumers at store f equals $W(P_f/P)^{1-\varsigma}$. Hence, let q_f^F be the units sold by store f to flexible consumers, then:

$$q_f^F = S\frac{W}{P_f}\left(\frac{P_f}{P}\right)^{1-\varsigma} \tag{1.6}$$

Total sales by store f therefore equal

$$q_f = (1 - S_f)\frac{W}{P_f} + S\frac{W}{P_f}\left(\frac{P_f}{P}\right)^{1-\varsigma} \tag{1.7}$$

Equation (1.7) shows that demand is the sum of the two types of consumers. These consumers differ in their demand elasticity, implying that the composition of consumer types in demand will be central to determining the optimal price for the firm.

Deriving the demand elasticity from Equation (1.7) is straightforward:

$$\varepsilon_f = -\frac{q_f^C}{q_f} - \varsigma\frac{q_f^F}{q_f} - \mu\left(1 - \rho\right)\frac{S_f}{1 - S_f}\tau\frac{V(P_f)}{P_f}\frac{q_f^C}{q_f} \tag{1.8}$$

where $\rho \equiv \frac{V'(P_f)}{V(P_f)} P_f$ is the elasticity of the function $V(\cdot)$. Because the function $V(P_f)$ is the expected indirect utility of search given price P_f, its elasticity ρ captures the marginal effect that the price P_f has on beliefs about the price P that is available by searching. For convenience, I will refer to ρ as the "informational effect" of prices. This informational effect arises because prices are correlated given that they share a common component of costs C.

Equation (1.8) shows how learning about inflation, captured by the informational effect ρ, affects the demand elasticity. Intuitively, the demand elasticity depends on a weighted average of the elasticities of captive and flexible consumers, captured by the first and second term, and the middle term which captures the marginal loss of consumers to search. This term shows that the search elasticity μ is attenuated by a factor ρ. Intuitively, this attenuation occurs because consumers update their beliefs about P in the direction of P_f. For example, consumers observing a high price will be more pessimistic about the price they are likely to observe after searching and are less likely to search than they would be absent the informational effect.

Equation (1.8) is also useful to understand how cross-sectional price dispersion arises in the model. Intuitively, price dispersion depends on how the demand elasticity varies across the cost distribution. The distribution of prices could be more or less dispersed than costs depending on whether the demand elasticity rises or falls across the price distribution. To figure out which is the relevant case, consider a high-cost firm that sets a higher price than a low-cost firm. The firm with higher prices will sell less to both captive and flexible consumers. However, because flexible consumers are more elastic, demand for this type of consumer falls more than proportionally. Raising prices therefore skews sales towards captive consumers which lowers the demand elasticity. Firms with high costs therefore have an incentive to charge higher markups.

Shocks and Information Structure. I assume common and idiosyncratic costs have a log-normal distribution with variances σ_c^2 and σ_f^2. Without loss of generality, I assume the mean of both cost shocks is zero.

Stores have full information. Consumers, on the other hand, have incomplete information, in the sense that they know the know the distributions of cost shocks but not their realizations. They therefore use prices to update their beliefs about common and idiosyncratic costs, and the average price P.

I will assume consumers are fully Bayesian, although alternative behavioral frameworks of expectation formation could be incorporated in a straightforward manner.

1.2.2 Equilibrium

This section formally defines the equilibrium in the market. Schematically, an equilibrium requires that firms set their prices optimally given demand and the distribution of prices set by other firms, and consumers form beliefs about prices in a way that is consistent with firm price setting and search optimally given these beliefs.

Below is the formal definition of the equilibrium. Because the space of actions is price setting for firms and searching for consumers, I define the equilibrium in terms of these actions.

Definition (Equilibrium). *Let $\omega_f \equiv (E_f, C)$ be the realized components of firm costs, $g_\omega(\omega)$ be the distribution that captures consumers' common prior beliefs about the distribution of the cost components ω, and let P_i be the price observed by consumer i. Then, the functions $\mathcal{P}(\omega)$ and $\mathcal{S}(P_i)$ constitute a symmetric Bayesian Nash equilibrium in pure strategies if:*

1. *$\mathcal{P}(\omega)$ maximizes profits following (1.3) when other stores set their prices with the function $\mathcal{P}(\omega)$ and consumers search is ruled by $\mathcal{S}(P_i)$,*
2. *$\mathcal{S}(P_i)$ satisfies the optimal search condition given by equation (1.1) under posterior beliefs characterized by*

$$g_{P|P_i}(P) = \frac{g_P(P_i, P)}{g_P(P_i)}$$

where $g_P(P_i)$ is the unconditional distribution of the price index P and $g_P(P, P_i) \equiv g_\omega\left(\omega : \mathcal{P}(\omega) = P_i \wedge \left[\int \mathcal{P}(\omega)^{1-\varsigma} df\right]^{1/(1-\varsigma)} = P\right)$ is the joint distribution of the price index and individual prices.

The task of finding functions $\mathcal{P}(\omega)$ and $\mathcal{S}(P_i)$ that satisfy the equilibrium conditions is non-trivial. In particular, it requires finding a fixed point between optimal price setting and beliefs generated by said price setting.

One approach to finding a solution that satisfies the equilibrium conditions is numerical. It can be shown that the distribution of prices satisfies a differential equation that resembles optimal bidding in auctions as in Milgrom & Weber (1982); this is the strategy pursued by Benabou & Gertner (1993). However, beliefs about the possible distributions of prices will, in their most general form, be high- or potentially infinite-dimensional objects. While feasible to solve, the numerical approach comes with the cost that it is harder to draw crisp and robust predictions or conclusions from the model. Moreover, the non-linear model is hard to scale, particularly when considering dynamic extensions and with many shocks. This is because the distribution of prices would not generally be one that leads to tractable distributions of beliefs.

The approach in this paper is to use perturbation methods to study the approximate equilibrium response around an initial equilibrium. I will log-linearize the model around the equilibrium in which common costs equal their unconditional mean and consumers' prior beliefs equal the unconditional distribution of costs. Because cost shocks were assumed to be log-normal, beliefs in the log-linear model will always be within the log-normal family, making the model tractable.

I will focus on the pure strategy symmetric Bayesian Nash equilibrium, although alternative equilibria may exist. For example, it is known that search markets may have an equilibrium where homogeneous firms with constant marginal costs set their price through a mixed strategy (e.g., Burdett & Judd 1983).

Despite the difficulties in implementing this definition of the general model, the definition provides valuable insights into the scope of the model. Namely, the definition requires that prices set by firms be optimal given the prices set by their competitors, consumers' demand, and their total costs. Specifically, the breakdown between common and idiosyncratic costs does not affect optimal prices given the distribution of prices set by other firms and consumer demand. This means that firms need not know the breakdown between common and idiosyncratic costs by assumption. Rather, it suffices that firms learn about these components in equilibrium and that they can condition their prices on the information that is revealed.

Log-Linear Equilibrium

For the remainder of the paper, I will be solving for the linear equilibrium of the model. Following the equilibrium definition (1.2.2), the linear equilibrium consist on a pair of equations governing price setting and search that satisfy optimality and consistency conditions. Let these functions be:

$$\hat{p}_f = \beta_c \hat{c} + \beta_e e_t$$
$$\hat{S}_f = \psi \hat{p}_f$$

where the hats indicate log-deviations from the steady state.

The use of perturbation methods requires first solving for a steady state. I will focus on the equilibrium in which firms' common costs equal the unconditional mean, which I will set to one without loss of generality, and consumers' prior beliefs are common and equal to the unconditional distribution. The steady state is such that households form expectations under the belief that what they observe could have been generated by shocks. In this sense, the steady state resembles the stochastic (Juillard & Kamenik 2005) or risky steady-state (Coeurdacier et al. 2011). However, in this setting the steady state also features cross-sectional heterogeneity, given that firms differ in their marginal costs even absent aggregate shocks.

The log-linear equilibrium is obtained by log-linearizing equations (1.3) to (1.8) around the steady state. The steady state is characterized by the following expressions:

$$\bar{P}_f = \overline{markup}_f \tag{1.9a}$$
$$\bar{q}_f^C = \frac{1-\bar{S}_f}{\bar{P}_f} \tag{1.9b}$$
$$\bar{q}_f^F = \frac{\bar{S}}{\bar{P}_f} \exp\left\{-\frac{1}{2}(1-\varsigma)^2 \sigma_{p_f}^2\right\} \tag{1.9c}$$
$$\bar{\varepsilon}_f = -\frac{\bar{q}_f^C}{\bar{q}_f^C + \bar{q}_f^F}\left[1 + \mu(1-\rho)\bar{v}\frac{\bar{S}_f}{1-\bar{S}_f}\right] - \frac{\bar{q}_f^F}{\bar{q}_f^C + \bar{q}_f^F}\varsigma \tag{1.9d}$$
$$\bar{S}_f = \exp\left\{-\mu \bar{v}\right\} \tag{1.9e}$$
$$\bar{S} = \exp\left\{-\mu\bar{v} + \frac{1}{2}\left[\mu\bar{v}(1-\rho)\right]^2 \sigma_{p_f}^2\right\} \tag{1.9f}$$
$$\bar{v} = \tau \exp\left\{-\frac{\varsigma-\gamma}{2}\sigma_{p_f}^2 + \frac{\gamma-1}{2}\frac{\sigma_{p_f}^2}{1+\sigma_{p_f}^2/\sigma_p^2}\right\} \tag{1.9g}$$
$$\overline{markup}_f = \frac{\bar{\varepsilon}_f}{\bar{\varepsilon}_f + 1} \tag{1.9h}$$

where bars denote the value of these variables in the steady state, $\sigma_{p_f}^2$ is the cross-sectional variance of prices, σ_p^2 is the variance of the price level, $\overline{markup}_f$ is the steady-state markup which depends on the demand elasticity $\bar{\varepsilon}_f$, and $\bar{v}$ is the certainty-equivalent price of searching. The values σ_p^2, $\sigma_{p_f}^2$ and the informational effect ρ depend on stores' and consumers' optimal actions which are equilibrium objects. For this reason, I will solve for equilibrium parameters jointly with the steady state.

The steady-state demand elasticity $\bar{\varepsilon}_f$ depends on multiple equilibrium parameters. As argued previously, the effect of observed prices on beliefs about unobserved prices given by ρ tends to reduce the elasticity of demand. Intuitively, this is because the expected difference between observed and unobserved prices is smaller the higher ρ is and the lower expected difference in prices implies a smaller search response. The elasticity also depends on the certainty-equivalent price of searching $\bar{v}$, both independently and through its effect on $\bar{S}$. It is easy to show that the effect of an increase in $\bar{v}$ on the elasticity depends on the sign of the expression $1 + (1-\rho) - (1-\rho)\mu\bar{v} - \varsigma$. Typically, we will assume that ς is a relatively high value, in which case an increase in $\bar{v}$ makes demand more inelastic.

As explained in Section 1.2.2, an equilibrium requires solving simultaneously for a pricing function and the beliefs elicited by said pricing function such that prices are set optimally given the effect of beliefs on consumer search. These components are expressed formally in the following lemma.

Lemma 1.2.1 (Log-Linear equilibrium). *A log-linear equilibrium is characterized by a parameters β_c, β_e, ρ and ψ such that firms set their prices following $\hat{p}_f = \beta_e\hat{e}_f + \beta_c\hat{c}$, consumers form expectations following $\hat{\mathbb{E}}_f[p] \equiv \mathbb{E}[\hat{p}|\hat{p}_f] = \rho\hat{p}_f$, where $\rho \equiv \frac{\beta_c^2\sigma_c^2}{\beta_c^2\sigma_c^2+\beta_e^2\sigma_f^2}$ is the correlation of prices, and search following $\hat{S}_f = \psi\hat{p}_f$, and the following optimization conditions are satisfied:*

$$\hat{q}_f = \frac{\bar{q}_f^C}{\bar{q}_f^C+\bar{q}_f^F}\hat{q}_f^C + \frac{\bar{q}_f^F}{\bar{q}_f^C+\bar{q}_f^F}\hat{q}_f^F \tag{1.10a}$$

$$\hat{q}_f^C = -\hat{p}_f - \mu\bar{v}\frac{\bar{S}_f}{1-\bar{S}_f}\left[\hat{p}_f - \hat{\mathbb{E}}_f[p]\right] \tag{1.10b}$$

$$\hat{q}_f^F = \hat{S} - \varsigma\hat{p}_f + (\varsigma - 1)\hat{p} \tag{1.10c}$$

$$\hat{S} = \mu\bar{v}\left[\hat{p} - \hat{\mathbb{E}}[p]\right] \tag{1.10d}$$

$$\hat{\varepsilon}_f = \omega_C^\varepsilon(\hat{q}_f^C - \hat{q}_f) + \bar{S}_f\frac{\mu\bar{v}(1-\rho)^2(1-\mu\bar{v})}{-\bar{\varepsilon}_f}\hat{p}_f + (1-\omega_C^\varepsilon)(\hat{q}_f^F - \hat{q}_f) \tag{1.10e}$$

$$\hat{markup}_f = \frac{1}{\bar{\varepsilon}+1}\hat{\varepsilon}_f \tag{1.10f}$$

$$\hat{p}_f = \hat{markup}_f + \hat{c} + \hat{e}_f \tag{1.10g}$$

where $\omega_C^\varepsilon = \dfrac{1+\mu\bar{v}(1-\rho)\frac{\bar{S}_f}{1-\bar{S}_f}}{1+\mu\bar{v}(1-\rho)\frac{\bar{S}_f}{1-\bar{S}_f}+\varsigma\frac{\bar{q}_f^F}{\bar{q}_f^C}}$.

According to Lemma 1.2.1, finding an equilibrium requires solving for three values, β_e, β_c and ρ. The values β_e and β_c are the pass-throughs of idiosyncratic and common components of costs to prices and determine the informational content of prices.

The value $\rho = \frac{\beta_c^2 \sigma_c^2}{\beta_c^2 \sigma_c^2 + \beta_e^2 \sigma_f^2}$ is the expression for the informational effect discussed in Equation (1.8). It has multiple possible interpretations. The value ρ equals the correlation of prices or the ordinary least squares (OLS) coefficient of a regression between two log-prices. It also equals the Kalman gain and the reduction in the variance of prior beliefs about aggregate prices. These properties stem from the assumption that expectations are formed rationally given the pricing function $\hat{p}_f = \beta_e \hat{e}_f + \beta_c \hat{c}$. When prices are highly correlated, they provide little information about relative prices and therefore induce a smaller search response. This is key to understanding the link between the informational effect ρ and equilibrium markups.

Furthermore, firms' optimal pass-through of idiosyncratic costs to prices β_e also rules the cross-sectional dispersion of prices. Unlike common models in which cross-sectional price dispersion arises from nominal rigidities, in this setting cross-sectional dispersion depends on firms' incentives to either exploit captive consumers or serve a broader consumer base. To the extent that price dispersion reflects the exploitation of captive consumers, consumers will not benefit from price dispersion even if indirect utility is quasiconvex in prices. Moreover, price dispersion is a force that may encourage search and increase market competitiveness, for example as in Benabou (1988), this effect needs to be very strong to compensate the lower mobility of consumers.

Following Lemma 1.2.1, solving for the pricing parameters β_c and β_e for a given set of steady-state values and an expectational effect ρ is straightforward. This is because the optimization conditions are simply a series of linear equations. However, the difficulty lies in that both the steady state and ρ are equilibrium objects that are determined jointly with β_c and β_e.

Unfortunately, the equilibrium pass-throughs of β_e, β_c, and the informational effect ρ cannot be solved for in closed form. The reason is that the system of equations that define the equilibrium is highly non-linear in the equilibrium values. This is a typical feature of models with endogenous in which the solution requires solving simultaneously for the response to signals and the information embedded in those signals (e.g., Lorenzoni 2009, Amador & Weill 2010).

The following proposition describes an algorithm that finds the fixed point between firms' actions, households' expectations and the steady state.

Proposition 1.2.2 (Numerical Solution to the Log-Linear Equilibrium)**.** *The following algorithm can be used to find the solution to the log-linear equilibrium:*

- *conjecture an initial guess* $(\beta_c)_0, (\beta_e)_0$,
- *find the value of* $(\rho)_0$ *consistent with that conjecture given by* $(\rho)_0 = \frac{((\beta_c)_0)^2 \sigma_c^2}{((\beta_c)_0)^2 \sigma_c^2 + ((\beta_e)_0)^2 \sigma_f^2}$,
- *solve for the steady-state values* $(SS)_0$ *consistent with* $(\beta_c)_0, (\beta_e)_0$ *and* $(\rho)_0$,
- *use the system of equations (1.10a)-(1.10g) to generate a new guess* $(\beta_c)_1, (\beta_e)_1$,
- *iterate until convergence.*

I will defer the analysis of the equilibrium to the following section.

1.2.3 The Effects of Inflation Uncertainty

This section performs comparative statics to assess the effect of uncertainty about the shock to common costs σ_c^2 on equilibrium objects. The parametrization used is shown in Table 1.1.

Paremeter	Value
γ	2
ς	5
σ_f	0.2
τ	0.5
μ	1.5
ς (alternative)	2

Table 1.1: Baseline Parametrization - Partial Equilibrium

I will begin by discussing the main objects of interest ρ, β_c, and β_e, which are the effect of observed prices on beliefs about unobserved prices $\mathbb{E}[\hat{p}|\hat{p}_f] = \rho\hat{p}_f$, and the elements of the pricing function $\hat{p}_f = \beta_c\hat{c} + \beta_e e_f$. I will then discuss the effect of cost inflation uncertainty on the other equilibrium variables.

Strategic pricing and expectations

Figure 1.1 shows how the equilibrium values for ρ, β_c and β_e change as the variance of common shocks σ_c^2 increases. I will begin analyzing the how the informational effect ρ varies with inflation volatility. Given the assumption that shocks are log-normal and that consumers are rational, the value of ρ is:

$$\rho = \frac{(\beta_c)^2\sigma_c^2}{(\beta_e)^2\sigma_f^2 + (\beta_c)^2\sigma_c^2} \tag{1.11}$$

For expositional purposes, assume for now that β_c and β_e are fixed. From this expression, it is clear that ρ is zero when σ_c^2 is zero, that is when there is no uncertainty about common cost shocks, and ρ converges to one as σ_c^2 grows arbitrarily large. Furthermore, it is easy to prove that the first derivative of ρ with respect to σ_c^2 is positive and that the second derivative is negative. Intuitively, common components become more likely to be source of observed prices as the variance of common shocks σ_c^2 becomes larger.

The left panel of Figure 1.1 shows how the informational effect ρ changes with uncertainty in common cost shocks. While the coefficients β_e and β_c that weigh the variance terms in Equation (1.11) are changing with the equilibrium, β_e^2/β_c^2 would have to rise by a large magnitude to overturn the effect of raising σ_c^2. The graph shows that the effect of σ_c^2 dominates, and the informational effect is increasing, concave and converges to one with as inflation uncertainty rises.

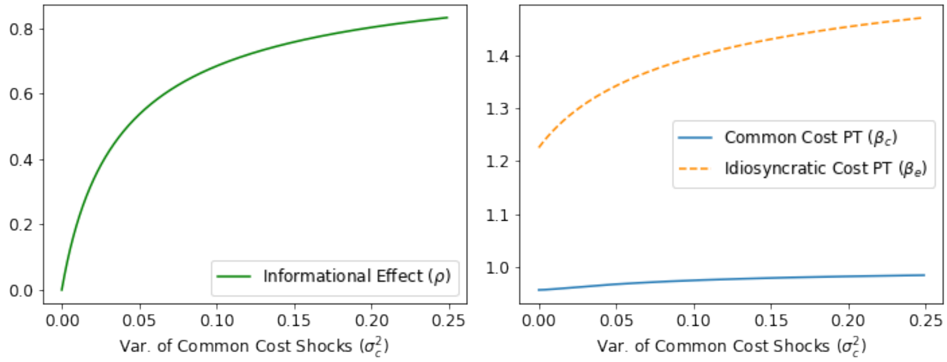

Figure 1.1: Effect of Inflation Uncertainty on Pricing and Beliefs

Next, let us study how the pass-through of common cost shocks β_c varies with inflation uncertainty. It will be useful to use the following expression which contains the demand elasticity of the average firm changes when there is a common shock.

$$\hat{\varepsilon}_f = \frac{\bar{S}_f \mu \bar{v}(1-\rho)}{-\bar{\varepsilon}_f} \left[\kappa_\varsigma (\varsigma - 1) - (\rho - 1) - \mu \bar{v}(1-\rho) \left(1 + \kappa_\varsigma \frac{\bar{S}_f}{1 - \bar{S}_f} \right) \right] \hat{p}_f$$

for a constant $\kappa_\varsigma = \frac{\bar{q}_f^C}{\bar{q}_f^F + \bar{q}_f^C} \frac{\bar{q}_f^F}{\bar{q}_f^F + \bar{q}_f^C} \frac{1}{(1-\bar{S}_f)\bar{S}_f}$.

This expression is obtained by combining equations (1.10a) to (1.10e) of firms' optimization conditions and setting $e_f = 0$. We can see that $1 - \rho$ is a factor on the left-hand-side. As ρ approaches one, the left-hand-side expression approaches zero, meaning that the elasticity of demand is unresponsive to the change in price of the average firm. Because the elasticity is approximately constant, so is the markup and the pass-through of common costs is approximately one for high values of ρ.

Now, let us consider the case where ρ is not one. For expositional purposes, it will be useful to think of κ_ς as approximately one, in which case the parenthesis reduces to $\varsigma - \rho - \mu \bar{v} \frac{1-\rho}{1-\bar{S}_f}$.[1] In my baseline parametrization, ς is a large value and the parenthesis is positive. In this case, the elasticity rises in response to higher prices driven by common cost shocks. Because the elasticity rises with prices, the pass-through of common shocks is less than one as shown in Figure 1.1. Moreover, the responsiveness of the elasticity to prices will typically fall with ρ as consumers infer that the likelihood that observed prices reflect common components is higher.

Under an alternative parametrization in which ς is low relative to the other terms in the parenthesis, it is possible for the pass-through of common shocks to be higher than one for low values of ρ, as shown in Figure 1.2. Under $\kappa_\varsigma \approx 1$, this happens if ς is lower than $\frac{\mu \bar{v}}{1-\bar{S}}$. In this case, as the average firm raises prices given a common shock,

[1]The steady-state equations (1.9b) and (1.9c) show that $\bar{q}_f^F / (\bar{q}_f^F + \bar{q}_f^C)$ resembles the value $\bar{S}_f$ with the addition of variance adjustments. In my baseline parametrization, those adjustments make κ_ς somewhat smaller than one although values larger than one are feasible under alternative reasonable parametrizations. In either case, the same intuition is similar for the purpose of understanding the effect of ρ on β_c.

it loses its captive consumers but receives searchers from other firms. For $\varsigma < \frac{\mu\bar{v}}{1-\bar{S}_f}$ the loss of captive consumers happens at a rate that is high enough for the demand elasticity to fall even as firms receive more flexible consumers. Intuitively, this is because the demand for flexible consumers is actually more inelastic than the demand from consumers assigned to the store for a low value of ς. In this case, aggregate shocks increase search and the demand elasticity of the average firm falls as it receives these consumers who are relatively inelastic. Markups rise as the elasticity falls, generating a pass-through larger than one. Still, this can only happen for small values of ρ; as ρ rises the pass-through of common costs converges to one.

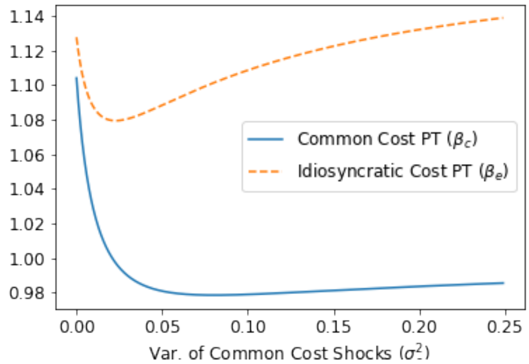

Figure 1.2: Effect of Inflation Uncertainty on Pricing - Alternative Parametrization

Now I turn to analyzing the pass-through of idiosyncratic shocks β_e. Intuitively, β_e depends on how differences in the composition of demand across low- and high-cost firms affect the demand elasticity. Take a high-cost firm that is setting higher prices than a low-cost counterpart. As the high-cost firms raise prices, they lose flexible consumers who have a high elasticity at a faster rate than captive consumers. The demand composition of high-cost firms is therefore skewed towards captive consumers, demand is more inelastic, and high-cost firms therefore set higher markups.

To see how the choice of β_e is affected by inflation uncertainty, I proceed as before by combining equation (1.10a) to (1.10e) but taking common cost shocks to be zero. The expression for the elasticity is as follows:

$$\hat{\varepsilon}_f = -\frac{\bar{q}_f^C}{\bar{q}_f^C + \bar{q}_f^F}\frac{\bar{q}_f^F}{\bar{q}_f^C + \bar{q}_f^F}\frac{\left[\varsigma - 1 - \mu\bar{v}(1-\rho)\frac{\bar{S}_f}{1-\bar{S}_f}\right]^2}{-\bar{\varepsilon}_f}\hat{p}_f + \bar{S}_f\frac{\mu\bar{v}(1-\rho)^2(1-\mu\bar{v})}{-\bar{\varepsilon}_f}\hat{p}_f$$

The first term is clearly negative and depends on the squared difference of elasticities between flexible consumers ς and consumers assigned to the store $1 + \mu\bar{v}(1-\rho)\frac{\bar{S}_f}{1-\bar{S}_f}$. This term captures how the composition of demand varies across the price distribution. Assuming ς is high as in our baseline parametrization, a high value of $\varsigma - 1 + \bar{S} - \mu\bar{v}(1-\rho)\frac{\bar{S}_f}{1-\bar{S}_f}$ implies that a firm with high prices will lose sales to flexible consumers at a higher rate than it loses captive consumers. In this case, the composition of sales of high-cost firms becomes skewed towards captive consumers whose demand is inelastic. This compositional effect lowers the demand elasticity and raises markups for high-cost firms, resulting in a pass-through of idiosyncratic costs that is larger than one.

The sign of the second term depends on the sign of $1 - \mu\bar{v}$. This value captures whether the loss of captive consumers as prices rise happens at a rate that is fast enough for the expenditure of non-searchers to fall. Under my baseline parametrization which requires a relatively high value of search to generate reasonable markups, $\mu\bar{v}$ is low and the second term is also negative.

Figure 1.2 shows that for a low value of ς idiosyncratic cost pass-through may fall for low values of the expectational effect ρ. This happens if $\varsigma < 1 + \mu\bar{v}\frac{\bar{S}_f}{1-\bar{S}_f}$. In this case, as ρ rises the difference in the bracket initially falls but eventually rises as the search becomes less responsive to prices.

The effect of inflation uncertainty on other equilibrium variables

Figure 1.3 shows the effect of cost inflation uncertainty on multiple equilibrium values.

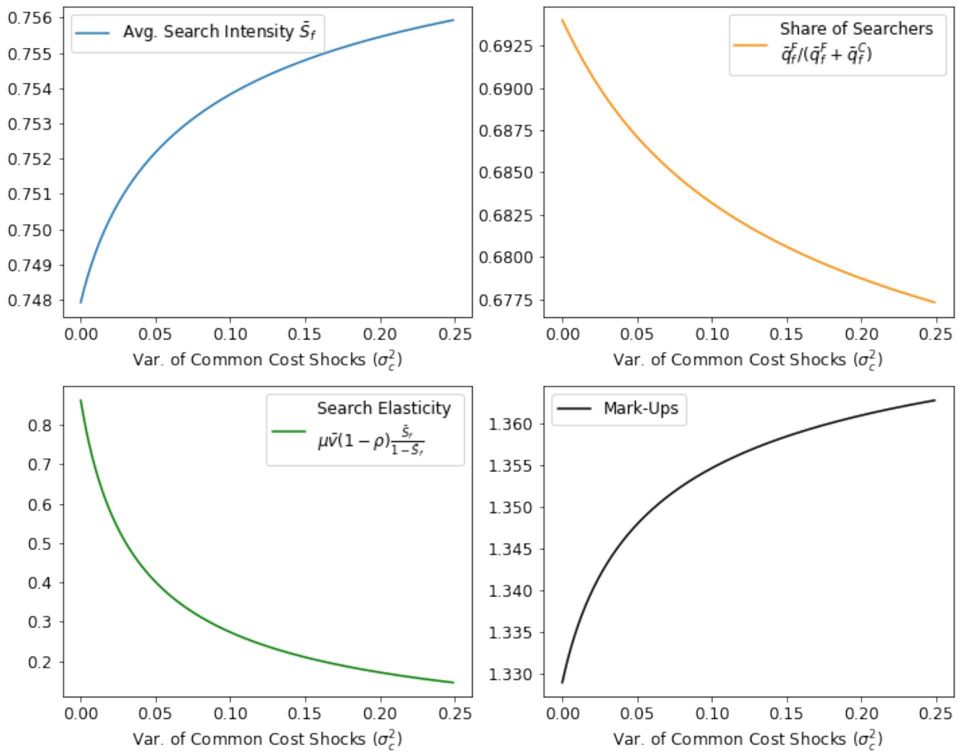

Figure 1.3: Effect of Inflation Uncertainty on Other Equilibrium Values

The top left panel shows the effect of inflation uncertainty on $\bar{S}_f$ which is the fraction of consumers visiting the average firm that search. I will refer to this value as the search intensity of the average consumer or simply average search intensity. Perhaps surprisingly, $\bar{S}_f$ rises with the variance of common costs. To make sense of this, it is useful to consider the expression for the average search intensity $\bar{S}_f$ and the certainty-equivalent price of search relative to the average price $\bar{\upsilon} \equiv \tau \frac{\bar{V}(\bar{P}_f)}{\bar{P}_f}$ which are featured in Equations (1.9e) and (1.9g):

$$\bar{S}_f = \exp\{-\mu\bar{\upsilon}\}$$
$$\bar{\upsilon} = \tau \exp\left\{-\frac{\varsigma-\gamma}{2}\sigma^2_{p_f} + \frac{\gamma-1}{2}\frac{\sigma^2_{p_f}}{1+\sigma^2_{p_f}/\sigma^2_p}\right\}$$

It is clear that average search intensity rises when the value $\bar{\upsilon}$ falls. The certainty equivalent price depends on two equilibrium variances. The first is the cross-sectional variance of prices $\sigma^2_{p_f} = (\beta_e)^2\sigma^2_f$. As the cross-sectional variance of prices rises, the certainty equivalent price falls because the CES price index falls when prices are more dispersed. Intuitively, higher price dispersion increase the probability of finding very low prices, to which an agent with a high elasticity of substitution will respond strongly. The second term $(1-\rho)(\beta_c)^2\sigma^2_c$ is the posterior variance of beliefs about the common component of prices. If $\gamma > 1$, this term raises the certainty equivalent price because search is subject to greater uncertainty. It is easy to show that $(1-\rho)(\beta_c)^2\sigma^2_c$ equals $(\beta_e)^2\sigma^2_f\rho$ which is strictly smaller than the cross-sectional distribution of prices. As ρ approaches one, average search intensity rises with the cross-sectional dispersion of prices if $\varsigma > \gamma$, as is the case in the baseline parametrization.

The top right panel of Figure 1.3 shows the effect of inflation uncertainty on the fraction of sales to flexible consumers for the average firm. To understand the intuition behind this expression, we can combine the expressions for $\bar{q}^F_f$ and $\bar{q}^C_f$ from the steady-state equilibrium Equations (1.9b) and (1.9c).

$$\frac{\bar{q}^F_f}{\bar{q}^F_f + \bar{q}^C_f} = \frac{\bar{S}_f \exp\left\{\frac{1}{2}\left[\mu\bar{\upsilon}(1-\rho)\right]^2\sigma^2_{p_f} - \frac{1}{2}(1-\varsigma)^2\sigma^2_{p_f}\right\}}{\bar{S}_f \exp\left\{\frac{1}{2}\left[\mu\bar{\upsilon}(1-\rho)\right]^2\sigma^2_{p_f} - \frac{1}{2}(1-\varsigma)^2\sigma^2_{p_f}\right\} + 1 - \bar{S}_f}$$

This value depends on the value $\bar{S}_f$ and a variance adjustment. The variance adjustment decreases the weight on $\bar{S}_f$ given that a large value of ς ensures the expression in the exponential is negative. Inflation uncertainty increases the cross-sectional dispersion of prices $\sigma^2_{p_f}$ given that β_e rises, making the weight on $\bar{S}_f$ smaller. Intuitively, because searchers have a CES demand, increasing price dispersion concentrates a larger fraction of sales to low-price stores and reducing the sales to flexible consumers for the average firm.

The bottom left panel of Figure 1.3 shows the search elasticity. In contrast to the average search intensity $\bar{S}_f$, which is the likelihood of search for a consumer who visits the average store, the search elasticity captures the differential search response of consumers who observe a price different from the average price. The dominant factor driving the search elasticity is the expectational effect ρ which rises with inflation uncertainty. As ρ approaches one, the search elasticity falls to zero. Intuitively, this is because for arbitrarily high values of inflation uncertainty, prices no longer reveal information about the relative components of prices which matter for the search decision. Because consumers' beliefs about relative components barely move with prices, they essentially behave as if they had no information about relative prices and search at about the average rate.

Finally, I will discuss markups. The expressions for markups and the steady-state elasticity featured in the steady-state Equations (1.9d) and (1.9h) will be useful.

$$\bar{\varepsilon}_f = -\frac{\bar{q}_f^C}{\bar{q}_f^C + \bar{q}_f^F}\left(1 + \mu\bar{v}(1-\rho)\frac{\bar{S}_f}{1-\bar{S}_f}\right) - \frac{\bar{q}_f^F}{\bar{q}_f^C + \bar{q}_f^F}\varsigma$$
$$\overline{markup}_f = \frac{\bar{\varepsilon}_f}{\bar{\varepsilon}_f + 1}$$

It is clear that the search elasticity $\mu\bar{v}(1-\rho)\frac{\bar{S}_f}{1-\bar{S}_f}$ analyzed previously is a component of the elasticity. As inflation uncertainty rises, this component vanishes and demand becomes more inelastic. Furthermore, the fraction of sales to searchers $\frac{\bar{q}_f^F}{\bar{q}_f^C + \bar{q}_f^F}$ is also falling, as shown in the top left panel of Figure 1.3. Both of these factors increase markups, as shown in the bottom right panel of Figure 1.3.

Summary

I end the section by summarizing the equilibrium effects of rising inflation uncertainty. First, inflation uncertainty increases the effect of observed prices on beliefs about the aggregate price level. As prices reveal more information about the price level and less about relative prices, consumer search becomes less price-sensitive. Moreover, firms generally face a higher share of captive consumers as inflation uncertainty rises. Both factors prompt firms to raise their markups.

The pass-through of common cost shocks to prices is typically less than one given that higher nominal prices increase the aggregate share of flexible consumers. However, as inflation uncertainty rises, consumers infer that prices reflect changes in the price level, moderating the aggregate search response until it vanishes when inflation uncertainty is very high. At this point, the pass-through of common cost shocks converges to one. A pass-through higher than one is possible, but it likely requires a parametrization that is unrealistic.

The pass-through of idiosyncratic shocks to prices is generally larger than one because demand becomes more inelastic as firms raise prices. The fall in the price sensitivity of search caused by inflation uncertainty has opposing effects on low- and high-cost firms. The marginal consumers for high- and low-cost firms become more captive and flexible, respectively, as inflation uncertainty rises. This raises cross-sectional price dispersion.

In all, the partial equilibrium model allows for a transparent and comprehensive analysis of the effects of inflation uncertainty on the market equilibrium. The next section extends the model to a general equilibrium framework with the objective of analyzing the welfare effect of these channels in Section 1.4.

1.3 Inflation Uncertainty in General Equilibrium

The previous section showed the distortions generated by inflation uncertainty. This section builds a general equilibrium model by adding production and labor supply to the partial equilibrium model developed in Section 1.2. The goal of developing a general equilibrium model is to carry out welfare evaluations that are comparable to standard macroeconomic models (e.g., Woodford 2003) which is the object of Section 1.4.

In the general equilibrium model, aggregate inflation is caused by two shocks: aggregate TFP and money supply shocks. These shocks have similar implications for prices (money shocks raise prices while TFP shocks lower them) but differ in that TFP shocks also affect the aggregate supply of goods. That being said, the inclusion of aggregate TFP shocks is inessential and the model works similarly absent them. Money supply shocks are used to finance an exogenous fiscal deficit that is transferred to households, and the main exercise in the general equilibrium model is to determine how increasing the variance of money supply shocks affects the equilibrium.

The setup is built to resemble the partial equilibrium model as much as possible. Periods begin with aggregate shocks which are known to firms but not to households. Stores differ in their idiosyncratic TFP and set prices anticipating the likelihood of consumers searching. One difference between the general equilibrium framework and the partial equilibrium framework is that consumers will observe a sample of prices rather than a single price from an individual store. Consumers use samples to form beliefs about aggregate and relative components of prices and decide how intensely to search for substitutes, i.e. they decide to search for a subsample of the prices they observed.

After consumers have made their search decisions and found prices at which they can buy goods, they observe the nominal wage rate, receive a transfer from the government and firm profits, and decide how much to work and consume. Labor and consumption choices are made frictionlessly given that all uncertainty has been resolved by this time, and goods and labor markets clear.

I will assume that households have multiplicatively separable preferences in consumption, leisure and real money balances. The inclusion of money in the utility function isnecessary to generate an incentive to hold the money supply. Nevertheless, desired real money balances are proportional to consumption in equilibrium, so the utility function admits a representation without real money balances. Moreover, expected utility in the log-linearized equilibrium has a mean-variance representation, making the setup analytically tractable. The mean-variance representation of expected utility will also be very useful in Section 1.4 when I use the model to evaluate welfare from inflation uncertainty.

The following section describes the model in detail.

1.3.1 Setup

The model features households, stores, and a central bank. There is a continuum of households indexed by $i \in (0, 1)$ who supply labor and consume, a continuum of stores indexed by $f \in (0,1)$ who hire labor, produce and set prices, and the central bank sets the money supply.

I will begin by discussing the sequence of events before discussing the actions of each agent in greater detail below. There is a single period that begins with aggregate TFP and money supply shocks which are known to firms. Given that they have perfect information, firms set prices and correctly anticipate the equilibrium.

Consumers have incomplete information, meaning that they know the distribution that generates shocks but do not know the actual realizations. Once store prices are set, consumers visit stores and use observed prices to form expectations about the likelihood of realized shocks. Given their beliefs about relative prices, consumers decide how intensely to search for substitutes. Search choices determine the prices at which households can purchase goods.

Once households have searched, the period ends with a series of actions that happen simultaneously. These are production, consumption and labor decisions, and money transfers from the central bank to households.

Production. Firms use labor services to produce final goods with constant returns to scale production technology. Firms' total factor productivity has a component that is common across firms and one that is idiosyncratic, as shown below.

$$\begin{aligned} Q_f &= Z_f L_f \\ z_f &= z + \eta_f^z \\ z &= \mu_z + \eta^z \end{aligned}$$

where Z_f if firm f total factor productivity, L_f is the amount labor services hired by firm f, z_f is the log of firm TFP which has a common component z and an idiosyncratic component η_f^z. As before, shocks η_f^z and η^z are i.i.d. mean-zero normally distributed with variances σ_f^2 and σ_z^2, respectively.

Stores are profit maximizers. It is assumed that consumers have no information on the productivity of stores. In that case, profit maximization is reduced to the usual static profit maximization rule, as follows:

$$P_f = \frac{\epsilon_f}{\epsilon_f + 1} MC_f \tag{1.12}$$

where ϵ_f is the demand elasticity which will be derived from consumer search, and $MC_f \equiv W/Z_f$ is the marginal or unit cost of production, with W being the wage rate. Firms set prices with full information about aggregate shocks and perfectly anticipate consumers' responses to observed variables.

Households provide homogenous labor services such that

$$L_f = \int_{i \in (0,1)} N_{if} di$$

where N_{if} is the amount of labor services provided by household i to firm f.

Because households provide homogenous labor services, the wage rate W is common across households and will be such that the labor market clears.

Households. Households maximize their expected utility. I assume flow utility has a multiplicatively separable form. The advantage of using this function is that expected utility has mean-variance representation in the log-linearized model, thus making it analytically tractable.

$$\mathbb{E}_i [U_i] = \mathbb{E}_i \left[-\frac{1}{\gamma - 1} \frac{C_i^{1-\gamma}}{1 - \frac{N_i^{1+\eta}}{1+\eta}} \left(\frac{M_i}{P_i} \right)^{-\lambda} \right] \tag{1.13}$$

where $\gamma > 1$, $\lambda > 0$ and $\eta > 0$. The value P_i is the household's price index which arises from prices found through search and about which I will elaborate below. The introduction of money in the utility will imply that the money supply is the nominal anchor in the economy. Households maximize Equation (1.13) subject to the following budget constraint:

$$WN_i + M_i^s + \Pi = C_i P_i + M_i$$

where M_i are money holdings which do not yield any interests, Π are profits from firms and M_i^s is a transfer from the central bank.

Consumption C_i is Cobb-Douglas aggregate over a subset $J_i^C \subset (0, 1)$ of products. A possible interpretation of J_i^c is that households differ in their consumption baskets.

$$\ln C_i = \int_{f \in J_i^c} \ln C_i(f) df$$

The Cobb-Douglas assumption implies that demand for products is proportional to aggregate expenditure in a way that resembles the setup in the partial equilibrium model. Alternative assumptions such as CES aggregation over products do not have a meaningful effect on how the model works.

Given the assumptions about the sequence of actions, households observe prices in the set J_i^c before observing either aggregate wages W or money transfers M^s. Assumptions regarding the sequence in which households observe different objects are meaningful to determine the size of the informational effect of prices. Schematically, prices have a stronger informational effect the less information households have about underlying shocks. This assumption also captures the possibility that prices may be more flexible than other sources of information such as wages, reason for which they may contain more information about the latest shocks affecting the economy.

Following Lorenzoni (2010), subsets J_i^c are such that the idiosyncratic components of firm productivity are correlated within the sample, that is $\mathbb{E}[\eta_f^z | f \in J_i^c] = \eta_i^c$ which has a variance σ_c^2. The value η_i^c is a sampling error, and it is the reason why samples do not reveal the aggregate price level to households. The sampling error must be such that its variance is strictly smaller than the variance of firm idiosyncratic TFP, that is $\sigma_c^2 < \sigma_f^2$. It will be useful to define the ratio of these variances as $\chi = \sigma_c^2 / \sigma_f^2$, which must satisfy $\chi \in (0,1)$. A higher value of χ implies that idiosyncratic TFP is more correlated within samples.

Households can search for substitutes for the products f in their consumption basket. Substitutes are drawn from a set $J_i^s \subset (0,1)$ and the likelihood of drawing a price P_s from that sample is proportional to $P_s^{1-\varsigma}$, as assumed in the partial equilibrium setup. This implies that demand from searchers has a constant elasticity of substitution given by ς. As was the case with J_i^c, subsets J_i^s have sampling errors which will be called η_i^s with variance σ_s^2 and are independent of η_i^c. It is assumed that household i draws substitutes for products f in J_i^c from the same sample of substitutes J_i^s.

Households observe all of the prices in the set J_i^c and use that information to form beliefs before searching. This is somewhat different from the partial equilibrium framework in which consumers form beliefs based on one single price. Despite this difference, I will show that consumers' demand in the general equilibrium model is very similar to that of the partial equilibrium model. The only difference is that the information effect is caused by the sample rather than an individual price.

Consumers form their household price index P_i after searching. Given the Cobb-Douglas assumption, the household price index is:

$$\ln P_i = \int_{f \in J_i^c} \ln P_i(f) df$$

Each price $P_i(f)$ is either the price P_f of the initial sample if the consumer does not search or a price P_i^s if the consumer searches. As in the partial equilibrium framework, search is subject to an iceberg cost τ and also a non-pecuniary factor $A_{if} > 0$. A_{if} has an exponential distribution with parameter μ and is i.i.d. in both i and f. The effective price consumers pays for each good in their consumption bundle can be written as:

$$P_i(f) = (1 - S_{if})P_f + S_{if}\tau \frac{P_i^s}{A_{if}}$$

where $S_{if} \in \{0,1\}$ is a decision variable that equals one if the consumer searches or zero otherwise.

Because the non-pecuniary benefits of search A_{if} range from zero to infinity, all consumers search for substitutes for at least a small fraction of products. This perfectly reveals the distribution of prices in the set J_i^s and all consumers end the period with the same amount of information. This rules out the possibility that consumers search solely for the purpose of gaining more information.

I assume households can only condition search on the non-pecuniary factor A_{if} and the expectations elicited by the observed prices. This assumption will simplify the solution to the search problem considerably while keeping the essence that search depends on beliefs about relative prices. An alternative assumption is that households can condition on both the non-pecuniary factor A_{if} and the price of each good in their J_i^c basket, P_f. This would complicate the solution to the search problem which would depend on three factors, the non-pecuaniary element A_{if}, expectations elicited by the sample J_i^c and individual prices within the sample.

Having searched and having formed their price indices P_i, households receive the transfer M_i^s, observe the wage rate W and deciding how much to work and consume. The labor decision that arises from the household first-order conditions is:

$$\frac{1}{1-\frac{N_i^{1+\eta}}{1+\eta}} = (\gamma - 1)\frac{W}{P_i}\frac{N_i^{-\eta}}{C_i}$$

Finally, the following expression combines the first-order conditions for consumption and money holdings:

$$\frac{\lambda}{\gamma - 1}\frac{M_i}{P_i} = C_i$$

This property of the household optimization conditions will turn out to greatly simplify results as household expenditure will be equal to the money supply. It will also imply that real money balances can be substituted from the utility function which allows me to write it in such a way that depends only on consumption and labor.

Central Bank. The central bank sets an exogenous money supply. The distribution for the log of the money supply is the following:

$$m^s = \bar{m}^s + d$$

where d is the log of the deficit to be financed by money issuance. This value is mean-zero normally distributed with a variance σ_m^2.

Shocks and Information Structure. All shocks have log-normal distribution and are independent. Firms have full information while consumers know the distribution from which shocks arise but not their realization.

1.3.2 Equilibrium

This section develops the solution to the general equilibrium model and analyzes the effect of rising inflation uncertainty on the equilibrium.

Finding a solution to the model requires solving for the optimal actions of households and firms given their information sets. Firms have full information and their information set is the realized shocks. The household optimization problem is set up in two stages. At the beginning of the period, households have incomplete information about aggregate shocks, observe a first sample of prices, and decide how intensely to search. This search choice depends on the beliefs elicited by observed prices. At the end of the period, households work, consume, and decide how much money to hold. These decisions are made frictionlessly given that households observe all relevant prices when making these choices.

Like in the partial equilibrium setup, consumers form beliefs optimally given firms' optimal pricing strategies and firms set prices optimally given the beliefs formed by households and their ensuing search response. The solution to the model is developed in Section 1.3.2. The main text delivers a more intuitive treatment, while technical details are provided in appendices.

Section 1.3.2 analyzes the effect of rising volatility in the money supply on the equilibrium. Money supply volatility is the main source of inflation uncertainty in the model. The effects on pricing are qualitatively similar to those of the partial equilibrium model. The main addition of the general equilibrium setting is that production and labor are endogenous. The effect of shocks on aggregate variables are not particularly interesting by themselves, although these will be important when studying welfare effects in Section 1.4 which is the main goal of the paper.

Solution to the General Equilibrium Model

I begin this section by defining the equilibrium. As before, I will solve for a log-linearized equilibrium around a steady state and use the notation $\hat{x}$ to be the log-deviation of x from its steady state value. The equilibrium is defined below.

Definition (Log-Linear General Equilibrium). *Let $x_i = (\hat{c}_i\ \hat{n}_i\ \hat{m}_i)'$ be a vector of household consumption, labor and money holding choices. A series of coefficients $\{\psi\}$ and $\{\beta\}$ constitute an equilibrium if the actions*

$$\hat{S}_i = \psi^S_{p^c}\mathbb{E}_i[\hat{p}^c_i - \hat{p}^s_i] \tag{1.14a}$$

$$\hat{x}_i = \psi_p\hat{p}_i + \psi_d\hat{m} + \psi_w\hat{w} + \psi_\Pi\hat{\Pi} \tag{1.14b}$$

$$\hat{p}_f = \beta_z\hat{z} + \beta_m\hat{m}^s + \beta_e\eta^z_f \tag{1.14c}$$

maximize household expected utility and firm profits, and households form beliefs $\mathbb{E}_i[\hat{p}^s_i]$ that are consistent with Equation (1.14c).

The description of how to find a solution is below. The problem of solving for optimal prices given how consumers search is similar to that of the partial equilibrium framework. For households, the biggest technical challenge is finding the optimal search choice. Fortunately, the assumptions made simplify the search problem significantly, which can be written as a mean-variance optimization problem. Finally, the model is closed by market-clearing conditions.

Households. The problem for households consists of two steps. At the beginning of the period, households observe prices in the set J^c_i and decide how intensely to search. Having obtained a price index P_i for their consumption basket, they decide how much to consume, work, and how much money to hold. These choices are made frictionlessly, and all households work at the same wage W and receive the same transfers from firms Π and the central bank M^s.

The solution to the search problem depends on how consumers believe the period will end. Therefore, I begin by solving for consumption, labor, and money holdings in the following lemma.

Lemma 1.3.1. *The consumption, labor and money-holding decisions at the end of the period have the following form:*

$$\hat{c}_i = -\hat{p}_i + \psi_w^c \hat{w} + \psi_\Pi^c \hat{\Pi} + \psi_m^c \hat{m} \quad (1.15a)$$

$$\hat{n}_i = \psi_w^n \hat{w} + \psi_\Pi^n \hat{\Pi} + \psi_m^n \hat{m} \quad (1.15b)$$

$$\hat{m}_i = \psi_w^m \hat{w} + \psi_\Pi^m \hat{\Pi} + \psi_m^m \hat{m} \quad (1.15c)$$

The result of Lemma 1.3.1 is very useful for the purpose of finding the solution to the search problem. That is because consumption is the only choice that depends on the price index which is formed by searching. Intuitively, a fall in the household price index increases consumption proportionally. Because the first order conditions indicate that consumption equal real money balances, nominal balances are not affected by the household price index. The result for labor supply is a property of the utility function used. A fall in the price index increases real wages and increases labor supply due to a substitution effect. However, this effect on labor supply is canceled out by the income effect (see King et al. 1988).

Households therefore only need to consider the effect that prices have on consumption and on deflating nominal balances. The following lemma shows that these facts allow us to write the search problem in a simple form.

Lemma 1.3.2. *The search problem can be written in the following form. Let $S_i \in (0,1)$ be the fraction of goods in their consumption basket for which they search for substitutes. Then, the optimal value of S_i satisfies the following condition:*

$$S_i^* = \arg\min_{S_i} \left[\mathbb{E}_i[\ln P_i(S_i)] + \frac{\gamma + \lambda - 1}{2} \mathbb{V}_i[\ln P_i(S_i)] \right]$$

where:

$$\ln P_i(S_i) = (1 - S_i) \ln P_i^c + \Gamma(S_i) + S_i \ln(\tau P_i^s)$$

for a function $\Gamma(\bullet)$ that ensures an interior solution.

Lemma 1.3.2 states that the optimal level of search is one that satisfies a mean-variance optimization problem. Intuitively, the mean-variance representation of the problem is a consequence of the multiplicatively separable utility function and log-linearity. Under these assumptions, expected utility has a mean-variance representation. The variance penalty is $\gamma + \lambda - 1$ which measures the curvature of the expected utility function, with $\gamma - 1$ being the curvature of utility for changes in consumption and λ being the curvature for changes in the deflator of money holdings.

The following lemma states the solution of the search problem.

Lemma 1.3.3. *The log-linearized solution to the search problem has the following solution. The* *steady-state search effort is characterized implicitly by:*

$$\bar{S}_i = \frac{\ln \bar{v} - \mu \ln \tau}{\mu(\gamma + \lambda - 1)\left[(1-\rho)\sigma_{p_c}^2 + \sigma_{p_s}^2\right]} \tag{1.16}$$

where $\bar{S}_i = \exp\{-\mu\bar{v}\}$*. Further, the log-deviation of the search decision from the steady state is:*

$$\hat{S}_i = \psi_{p^c}^S (1-\rho)\hat{p}_i^c \tag{1.17}$$

where $\psi_{p^c}^S \equiv \frac{\mu^2\bar{v}}{1+\mu^2\bar{v}\bar{S}(\rho+\lambda-1)\left[(1-\gamma)\sigma_{p_c}^2+\sigma_{p_s}^2\right]}$ *and* $\rho \equiv \frac{\beta_z^2\sigma_z^2+\beta_m^2\sigma_m^2}{\beta_z^2\sigma_z^2+\beta_m^2\sigma_m^2+\beta_e^2\sigma_c^2}$ *is the optimal belief given the pricing function.*

Finally, the log-linearized price index for households is:

$$\hat{p}_i = (1-\bar{S})\hat{p}_i^c + \bar{S}\hat{p}_i^s - \bar{S}_i \frac{1}{\mu}(\ln \bar{v} - \mu \ln \tau)\hat{S}_i \tag{1.18}$$

Equations (1.16) and (1.17) shows that uncertainty about the price $\ln P_i^s$, given by the term $(1-\rho)\sigma_c^2 + \sigma_s^2$, has two effects on search. The first is to discourage average search $\bar{S}_i$, as shown by Equation (1.16). If $(1-\rho)\sigma_c^2 + \sigma_s^2 = 0$, consumers would search until $\ln \bar{v} = \mu \ln \tau$, which is the point at which the marginal non-pecuniary benefit of search equals the search cost. For this reason, search decreases the price level on the margin in Equation (1.18).

The other effect is to reduce the responsiveness of search to expected price differences, as shown by the denominator of the Equation (1.17). Intuitively, this is because search is risky, and the household faces risk about uncertainty in the aggregate price level, given by $(1-\rho)\sigma_{p_c}^2$ as well as sampling risk σ_s^2.

The linear expression for log prices (1.18) can be combined with the consumption equation in (1.15a) to obtain the solution to household decisions.

Firms. The log-linear expressions on firm price-setting are similar to those of the partial equilibrium model of Section 1.2 with two main differences.

The first difference is that the expenditure of flexible and captive consumers is no longer fixed. Rather, expenditure varies with aggregate shocks. Fortunately, because households' nominal expenditure equals money balances, aggregate nominal expenditure will be random but essentially exogenous.

The second is that the informational effect is slightly different than in the partial equilibrium setup. This is because firms are now grouped into samples, and households form expectations based on the average price in the entire sample. Sampling errors bias consumers' expectations toward the price of individual firms.

As a result, a firm f will face consumers whose average expectations for the common component of prices in equilibrium are:

$$\mathbb{E}_{i:f\in J_i^c}(\hat{p}_t) = \rho(\hat{p} + \chi(\hat{p}_f - \hat{p}))$$

This expression shows that households' expectations are shifted towards the relative price of the firm f by the factor χ. Intuitively, a higher value of χ implies that idiosyncratic TFP is more correlated within samples and so price $\hat{p}_f$ is generally in baskets with prices closer to $\hat{p}_f$.

Despite these distinctions, the conditions determining firm pricing are similar to the partial equilibrium setting. The following lemma contains these conditions.

Lemma 1.3.4. *The log-linear equations that determine firms' optimal prices are given by:*

$$\hat{q}_f = \frac{\bar{q}_f^C}{\bar{q}_f}\hat{q}_f^C + \frac{\bar{q}_f^F}{\bar{q}_f}q_f^F \tag{1.19a}$$

$$\hat{q}_f^C = -\hat{p}_f + \hat{d}_f^C \tag{1.19b}$$

$$\hat{q}_f^F = -\varsigma\hat{p}_f + \hat{d}_f^F \tag{1.19c}$$

$$\hat{\varepsilon}_f = -\frac{\bar{q}_f^C/\bar{q}_f}{\bar{\varepsilon}_f}(\hat{q}_f^C - \hat{q}_f) - \varsigma\frac{\bar{q}_f^F/\bar{q}_f}{\bar{\varepsilon}_f}(\hat{q}_f^F - \hat{q}_f) \tag{1.19d}$$

$$\hat{p}_f = \frac{1}{\bar{\varepsilon}_f+1}\hat{\varepsilon}_f - \hat{z} - \eta_f^z + \hat{w} \tag{1.19e}$$

where $\hat{d}_f^C$ and $\hat{d}_f^F$ are terms that capture the average consumption of captive and flexible consumers of firm f, respectively. These values are:

$$\hat{d}_f^C = -\frac{\bar{S}_f}{1-\bar{S}_f}\psi_{p^c}^S(1-\rho)(\hat{p} + \chi\beta_e\eta_f^z) + \hat{m}$$

$$\hat{d}_f^F = \hat{S} + (\varsigma - 1)(\hat{p} + \chi\beta_e\eta_f^z) + \hat{m}$$

Note that the expression for elasticity in the general equilibrium case is written differently from that of the partial equilibrium framework. Equation (1.19d) expresses the demand elasticity as a weighted average between captive and flexible consumers. The informational effect of prices is included in the demand factors $\hat{d}_f^C$ and $\hat{d}_f^C$ that affect $\hat{q}_f^C$ and $\hat{q}_f^F$.

The assumption that firm prices are included in samples generates a subtle difference between the partial and general equilibrium models. In general equilibrium, households form beliefs based on samples instead of individual prices. From the perspective of the firm, the beliefs elicited by the sample are a shock to its mass of captive consumers which the firm takes as given. Despite this difference, the informational effect ρ enters demand in both setups in the same way. Namely, a higher value of ρ allows high-cost firms to retain more captive consumers and low-cost firms to retain fewer captive consumers than they would absent the informational effect.

Note that the system of Equations (1.19a)-(1.19e) solves for firm quantities, $\hat{q}_f$ as a function of consumer nominal demand. This implies that the goods market clears. The remaining market-clearing equations that close the model are below.

Market clearing. There are three markets that must clear: the money market, the goods market, and the labor market.

The money market clearing condition is that the total supply of money must equal money holdings: $\hat{m} = \hat{m}^s$. The market clearing condition for goods is that goods demanded must equal goods supplied. This is already taken care of in the system of equations that determine firm prices, which use demand for goods as an input.

Finally, the labor market clearing condition is that labor demand equals labor supply. Aggregate labor demand comes from production:

$$\hat{q} = \hat{z} + \hat{l}$$

And labor supply comes from households' optimization problem:

$$\hat{c} + \tilde{\eta}\hat{n} = \hat{w} - \hat{p}_{\bullet}$$

where $\hat{p}_{\bullet}$ is the log-deviation of household price indices which differs from the log-deviation of prices set by firms $\hat{p}$ because of non-pecuniary benefits of search.

In equilibrium, the wage rate $\hat{w}$ must be such that the labor market clears, or

$$\hat{l} = \hat{n}$$

Steady state and solution. As was the case in the partial equilibrium model, the steady state is determined jointly with the equilibrium parameters.

To find the solution to the equilibrium conditions, we can proceed by starting with an initial guess for prices, solving for beliefs consistent with that guess, and iterating.

Having laid out how to solve the model, the next section analyzes the effect of rising inflation uncertainty on the equilibrium.

The Effect of Inflation Uncertainty in General Equilibrium

This section shows how inflation uncertainty affects firms' optimal prices and the informational effect. While there are two sources of aggregate inflation in the model, aggregate TFP shocks and money supply shocks, I will focus on inflation uncertainty stemming from the latter given that TFP shocks are arguably unrelated to policy. Hence, the main exercise in this section is to evaluate how changes to money supply volatility affect price setting and the informational effect. The parametrization used is presented in Table 1.2.

Paremeter	Value
γ	2
η	1
λ	1
ς	5
σ_f	0.2
σ_z	0.05
τ	0.5
μ	1.5
χ	0.8

Table 1.2: Baseline Parametrization - General Equilibrium

For ease of exposition, recall that optimally set prices and the informational effect have the following form:

$$\hat{p}_f = \beta_z \hat{z} + \beta_m \hat{m}^s + \beta_e \eta_f$$
$$\rho = \frac{\beta_z^2 \sigma_z^2 + \beta_m^2 \sigma_m^2}{\beta_z^2 \sigma_z^2 + \beta_m^2 \sigma_m^2 + \beta_e^2 \sigma_c^2}$$

where β_z, β_m, and β_e indicate how aggregate TFP shocks, monetary shocks, and idiosyncratic TFP shocks affect optimal prices, respectively. The informational effect ρ is the Kalman gain that is consistent the price setting equation.

Figure 1.4 shows how increases in the volatility of money supply shocks affect firms' optimal prices and the informational effect. The top left panel shows the effect of aggregate TFP shocks on prices. This value is usually negative reflecting that prices fall given the reduction in costs associated with a positive productivity shock. The effect is less than one in absolute value but converges to one with the variance of money supply shocks. As was the case in partial equilibrium, the pass-through of aggregate TFP shocks is less than one-for-one in the baseline parametrization because firms retain more captive consumers as they lower prices and firms find it optimal to raise their markups given a more captive customer base.

It is somewhat paradoxical that the pass-through of aggregate TFP shocks converges to its full information value of -1 although households would not believe observed prices are driven by TFP shocks. This is because when aggregate price level volatility is very high consumers infer that prices are driven mostly by common components. Given a TFP shock, that belief is actually correct although households believe it is more likely to be caused by money supply shocks than a TFP shock.

The top right panel shows the effect of an idiosyncratic TFP shock on prices. This magnitude is negative, reflecting the fall in costs associated with a positive productivity shock. The magnitude is larger than one. As was the case in the partial equilibrium setup, this reflects the compositional change in firm demand as they raise or lower their prices with their costs. Higher prices will tend to disproportionately reduce sales to flexible consumers, increase the share of sales to captive consumers which motivates firms to raise markups.

The bottom left panel shows the effect of money supply shocks which mirrors the effect of TFP shocks. There are two reasons for this. The first is that, given that this symmetry is true, a positive money shock elicits the same expectations and search response as a negative TFP shock. Second, the labor supply response is also the same given the assumption that utility is multiplicatively separable in consumption and labor. This is a well-known property of this utility function, in which substitution and wealth effects on labor supply cancel out, making it a popular choice to generate balanced growth paths in models of growth.

Finally, the bottom right panel shows the informational effect ρ. As was the case in partial equilibrium, a higher variance of the common components of prices increases the effect of observed prices on expectations.

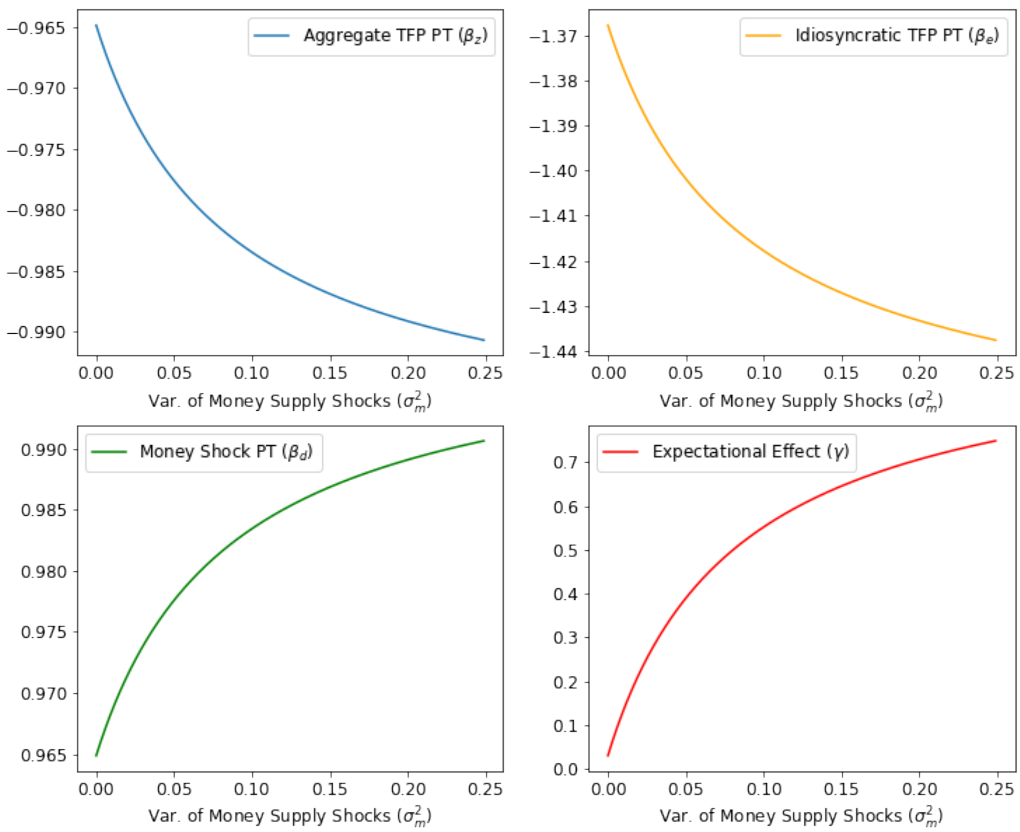

Figure 1.4: Effect of Money Supply Volatility

The top two panels of Figure 1.5 show the effect of TFP shocks on consumption and labor. Consumption rises given a positive TFP shock because of the greater abundance of goods. However, it rises less than one-for-one with TFP shocks because part of the fall in prices is offset by lower search activity. Because consumption rises less than TFP, less labor services are required to carry out production as shown in the right panel.

The bottom panels of Figure 1.5 show the effect of money supply shocks. While a positive money supply shock has the same effect on prices and aggregate search as a negative TFP shock, the effect on consumption differs in that TFP shocks increase the supply of goods in the economy even absent labor responses, while money supply shocks do not. In this case, the expansion in consumption and labor generated by a money supply shock is due to the offsetting effect that search has on prices. Namely, an increase in search induces firms to pass through costs less than one-for-one, increasing real consumption. Even so, the effect of money on consumption and labor is two orders of magnitude smaller than the expansion in money supply.

This will turn out to be important when studying the welfare effects of inflation uncertainty. Despite increasing uncertainty in money supply, money supply shocks themselves do not generate a lot of volatility in consumption and labor, which is what matters for household welfare. The next section studies the welfare implications in detail.

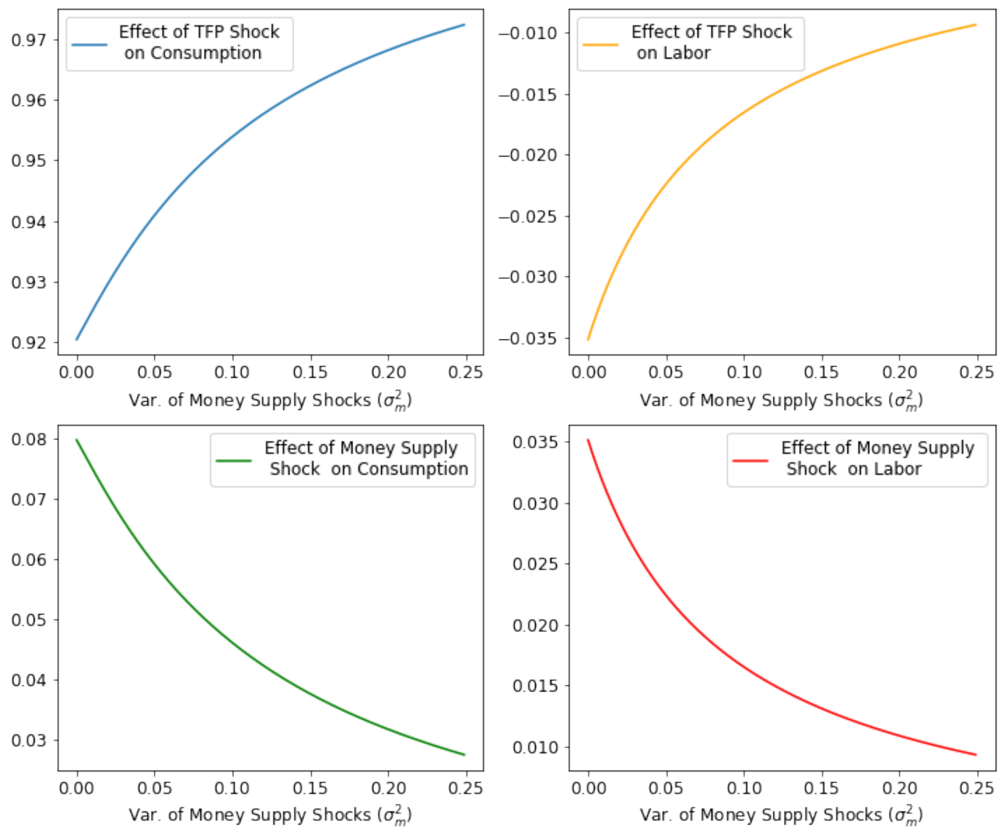

Figure 1.5: Effect Money Supply Volatility on Consumption and Labor

1.4 Welfare

This section derives a welfare function and analyzes its components. I will focus on conducting the welfare analysis in the general equilibrium model developed in Section 1.3. The purpose of carrying out welfare evaluations in general equilibrium is to make the results comparable to those of other macroeconomic models, although welfare analysis in partial equilibrium setups is certainly feasible.

I will further assume that the only source of cross-sectional household heterogeneity is given by price samples and their effect on search. This narrows the welfare effects to those that result only from uncertainty about relative prices.

1.4.1 Welfare Criterion

I will use the ex-ante expected utility of households as the appropriate welfare criterion. This value is as follows:

$$\mathcal{W} = \mathbb{E}\left[\frac{C_i^{1-\gamma}}{1-\gamma}\frac{1}{1-\frac{N_i^{1+\eta}}{1+\eta}}\left(\frac{M_i}{P_i}\right)^{-\lambda}\right]$$

where the i subscript is dropped given that households are ex-ante identical in terms of endowments and beliefs.

Given normal shocks, the linear solutions developed throughout the paper, and the first-order conditions of households' utility maximization problem, it is possible to show that the welfare function takes a convenient mean-variance representation.

Proposition 1.4.1 (Welfare function)**.** *The welfare function can be written as*

$$\mathcal{W} = \bar{\mathcal{W}}_i \exp\left\{\frac{1}{2}\left[\Theta_p \mathbb{V}(\hat{p}) + \Theta_{p_f}\mathbb{V}(\hat{p}_f - \hat{p})\right]\right\}\mathcal{W}^{IP} \tag{1.20}$$

where $\bar{\mathcal{W}}_i = -\bar{C}_i^{1-(\gamma+\lambda)}\left(1-\frac{\bar{N}_i^{1+\eta}}{1+\eta}\right)^{-1}$, $\mathbb{V}(\hat{p})$ *is the variance of aggregate prices,* $\mathbb{V}(\hat{p}_f - \hat{p})$ *is the cross-sectional price dispersion, and* $\mathcal{W}^{IP}$ *are factors independent of policy.*

Proposition 1.4.1 shows the expression that will be used to make welfare evaluations. Expected utility has a mean-variance representation, where the mean is given by $\ln \bar{\mathcal{W}}_i$, and the variance is given by the expression inside the exponential. The value Θ_p can be thought of as a welfare weight on aggregate inflation uncertainty. Similarly, Θ_{p_f} can be thought of as a welfare weight on cross-sectional price dispersion.

A feature of the model is that the welfare weights Θ_p and Θ_{p_f} depend on the inflationary regime. In particular, the variance of money supply shocks rises and the volatility of aggregate inflation $\mathbb{V}(\hat{p})$ rises, the term Θ_p will fall in such a way that the product $\Theta_p\mathbb{V}(\hat{p})$ is stationary. Intuitively, as shown in Figure 1.5, money supply shocks have a very small effect on consumption and labor, and this effect shrinks to zero as the volatility of money supply shocks rises. In consequence, welfare losses associated to aggregate volatility in this model will be limited.

Equation (1.20) indicates that inflation uncertainty affects welfare through its effect on either the average level of consumption and labor, or variance terms associated with aggregate volatility and price dispersion. While business cycle models typically focus on the latter, allocative distortions affecting average labor and consumption will turn out to be significant sources of welfare losses. The next section performs an accounting decomposition of the contribution of each of these factors.

1.4.2 The Welfare Effects of Inflation Uncertainty

This section studies the welfare costs of inflation uncertainty and its main channels in the model described in Section 1.3.

Figure 1.6 shows welfare in consumption equivalent terms. As expected, welfare falls with inflation uncertainty. The welfare loss is approximately equivalent to a 1.4 percentage point fall in consumption. The welfare loss is small because agents make decisions frictionlessly, with the exception of search which is subject to informational frictions. Even so, it is orders of magnitude larger than welfare costs associated with aggregate volatility in a similar environment (e.g., Lucas 2003).

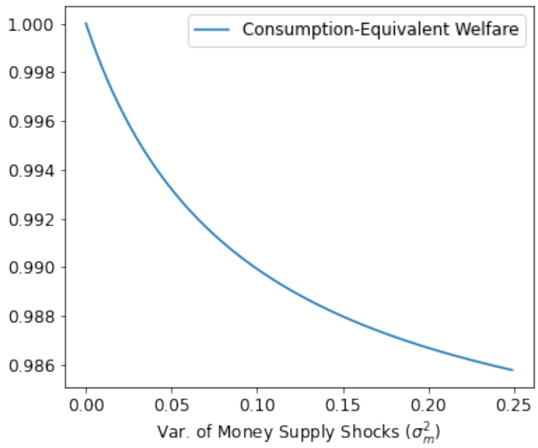

Figure 1.6: Effect Money Supply Volatility on Welfare

A useful feature of the welfare function in Equation (1.20) is that welfare effects can be decomposed into four distinct factors: average consumption $\bar{C}_i$, average labor $\bar{N}_i$, inflation volatility $\Theta_p \mathbb{V}(\hat{p})$, and cross-sectional price dispersion $\Theta_{p_f} \mathbb{V}(\hat{p}_f - \hat{p})$. An intuitive way to perform a decomposition is to evaluate the welfare function changing one factor and keeping the rest fixed. While these factors are interrelated in equilibrium, such a decomposition is insightful as an initial diagnostic of the drivers of welfare effects.

Figure 1.7 performs this decomposition. The top left panel shows that the fall in consumption $\bar{C}_i$ explains about 1.2 percentage points of the welfare loss. This accounts for about 85% of the welfare loss. The top right panel shows that labor supply $\bar{N}_i$ generates a modest welfare gain. This implies that labor effort falls in equilibrium with inflation uncertainty, potentially explained by falling incentives to work due to higher markups.

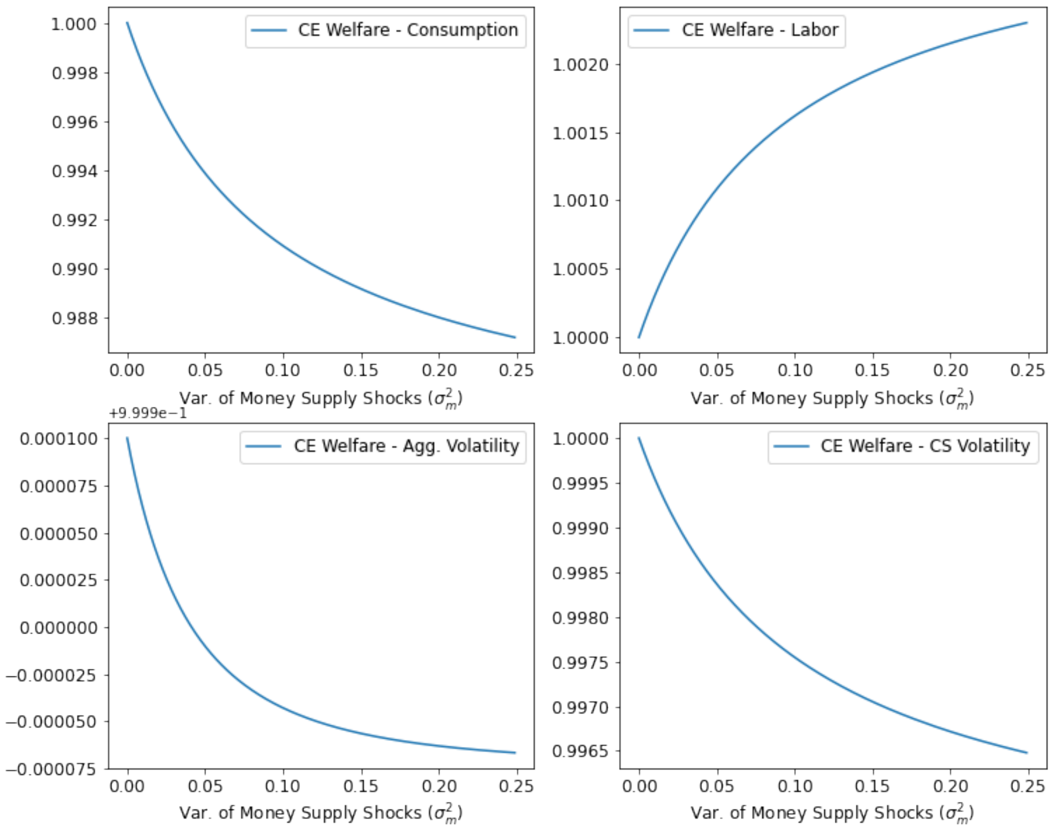

Figure 1.7: Welfare Decomposition

The bottom left panel shows the welfare loss attributable to aggregate inflation volatility $\Theta_p \mathbb{V}(\hat{p})$. In this case, I consider the joint effect of the weight Θ_p, as well as the variance $\mathbb{V}(\hat{p})$, given that the former is endogenous and changes with the equilibrium. The graph shows that this factor accounts for a negligible welfare loss which is in the order of 10^{-4}. Furthermore, the welfare loss levels out despite inflation volatility rising almost linearly with the variance of money supply shocks σ_m^2 as shown in the top left panel of Figure 1.8. The top right panel of Figure 1.8 shows that this is because the weight on the variance of aggregate prices Θ_p shrinks to zero. Intuitively, money supply shocks have smaller real effects and become neutral as the volatility in money supply rises, as shown in Figure 1.5. This makes the product $\Theta_p \mathbb{V}(\hat{p})$ level off despite the rise in aggregate price volatility.

An interesting implication of the asymptotic neutrality of money is that the welfare effects of inflation uncertainty are bounded. Intuitively, this is because money supply shocks only affect allocations if consumers are uncertain about their likelihood. As the volatility of money supply shocks grows arbitrarily large, prices reveal money supply shocks to consumers and the model works as if money supply shocks had been anticipated. In that extreme case, money supply shocks have no effect on real allocations, raising nominal prices one-for-one. This result is akin to the Lucas-supply curve, by which countries with volatile inflation tend to have steep supply curves.

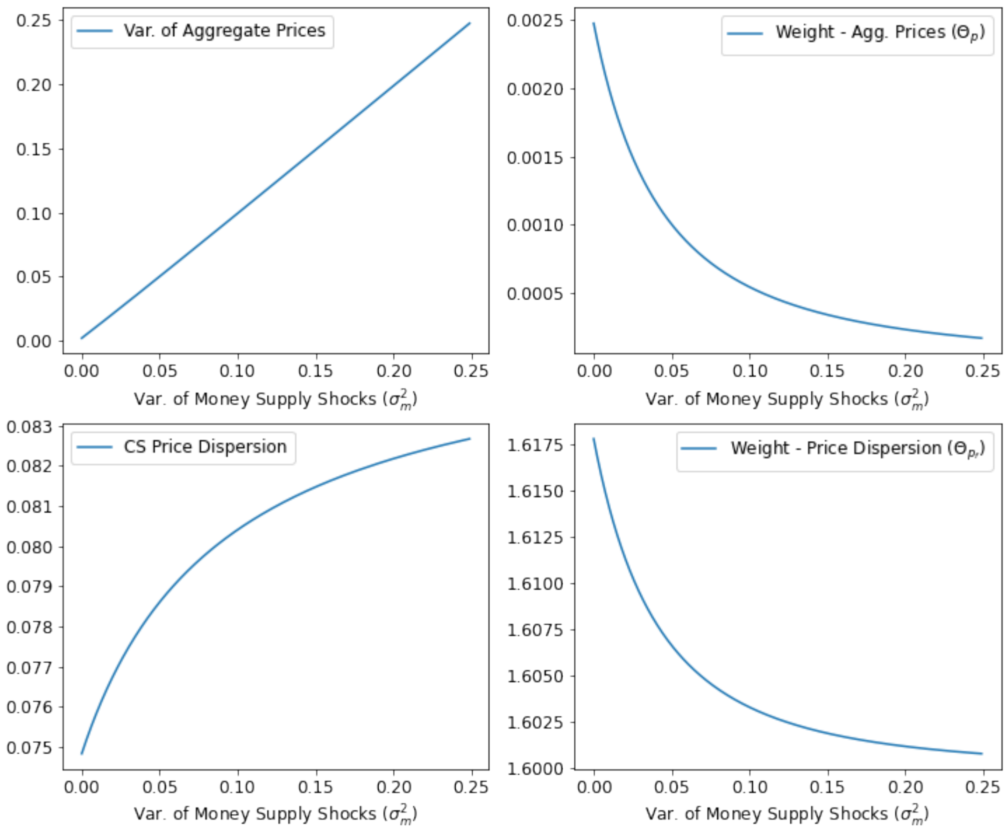

Figure 1.8: Welfare Decomposition: Variance Terms

The bottom right panel of Figure 1.7 shows the welfare loss attributable to cross-sectional price dispersion term $\Theta_{p_f}\mathbb{V}(\hat{p}_f - \hat{p})$. This factor accounts for about 25% of the total welfare loss. The bottom right panel of Figure 1.8 shows that the weight Θ_{p_f} changes very slightly with inflation uncertainty and the rise in price dispersion accounts for most of the welfare loss.

It is worth pointing out that this accounting exercise only attributes welfare losses that originate in cross-sectional consumption dispersion to price dispersion. However, cross-sectional price dispersion also directs the steady-state allocation of purchases across firms. Therefore, the factors $\bar{C}_i$ and $\bar{N}_i$ also depend on cross-sectional price dispersion. This point is discussed at greater length in Section 1.4.3.

This decomposition exercise highlights that the fall in steady-state consumption is the most important factor behind the welfare effects of inflation uncertainty. This is an atypical property of the model. In mainstream models, inflation is costly because of its effect on aggregate volatility rather than distorting the steady-state allocation. The next section delves more deeply into how these distortions arise.

1.4.3 Welfare under Counterfactual Behaviour of Households and Firms

The previous section performed an accounting decomposition of welfare effects associated with each component of Equation (1.20). While that exercise is useful, the terms in this equation are interrelated by the optimal actions of households and firms. This section performs counterfactuals based on the effect that changes to the environment have on the optimal actions of households and firms.

To fix ideas, let Λ^F be a vector that collects the optimal actions of firms and Λ^H be a vector that collects the optimal action of households. These actions are optimal responses to each other and to the variance matrix of aggregate shocks, which I will call Σ. The welfare function can then be written as:

$$\mathcal{W}(\Lambda^H, \Lambda^F, \Sigma)$$

This section studies the counterfactual welfare effects that arise from changes in the behavior of households Λ^H and firms Λ^F. The intention behind computing these counterfactual calculations is to identify what behavioral distortions generated by inflation uncertainty are most responsible for the welfare effect.

I will consider two kinds of counterfactuals. The first series of counterfactuals is based on the perturbation of the behavior of either households or firms. To fix ideas, consider a perturbation of the behavior of households from Λ^H to $\Lambda^{H'}$. Then, the welfare effect of the perturbation is:

$$\mathcal{W}(\Lambda^{H'}, \Lambda^F, \Sigma) - \mathcal{W}(\Lambda^H, \Lambda^F, \Sigma) \tag{1.21}$$

The purpose of computing this kind of counterfactual is to determine factors that make the equilibrium inefficient. Furthermore, perturbing the behavior of some agents and keeping the behavior of others fixed, allows me to derive insightful theoretical results which will be discussed in Section 1.4.3.

However, note that the welfare calculations in Equation (1.21) involve allocations that would not typically arise in equilibrium. This is because changes in household behavior bring about a strategic response from firms. Section 1.4.3 shows how to attribute welfare losses to these strategic responses.

Counterfactual Behavior and Welfare Effects

This section computes counterfactual welfare based on Equation (1.21). In all cases, markets clear given the counterfactual matches between consumers and firms and counterfactual prices.

While the counterfactual $\mathcal{W}(\Lambda^{H'}, \Lambda^F, \Sigma)$ in Equation (1.21) is not an object that would arise in equilibrium, computing these counterfactuals has the advantage that useful theoretical results can be derived. The first of these results concerns the effect of a perturbation to consumer search.

Proposition 1.4.2 (Welfare Effects - Search)**.** *Distortions to households' search decisions have a second-order effect on households' optimization problem. However, if $\beta_e \neq -1$ the reallocation of purchases across producers due to search has a first-order welfare effect. Furthermore, if $\beta_e < -1$, equilibrium search is inefficiently high, while equilibrium search is inefficiently low if $\beta_e > -1$.*

The intuition for the second-order welfare effects for households is as follows. Households search maximizing the expected value of utility conditional on the information that is available to them. Because the welfare function is concave, the function is (locally) flat at the maximum. A small movement around the maximum therefore typically implies a small optimization error.

However, a perturbation to search causes a reallocation of purchases. Intuitively, search increases the likelihood that consumers are matched with efficient firms, generating a welfare improvement due to productive efficiency gains. However, because search is costly, a marginal change in search is accompanied by a change in production. If the rise in production equals the gain in productive efficiency, customers can be served without increasing aggregate labor effort. This happens if $\beta_e = -1$.

Consider what happens as search reallocates sales across firms if $\beta_e < -1$. To fix ideas, consider a high-cost firm with price P^H and marginal cost MC^H that loses sales to a low-cost firm with P^L and MC^L. Under $\beta_e < -1$, the high-cost firm has a higher markup than the low-cost firm, implying that $P^H/P^L > MC^H/MC^L$. Given that consumers' nominal expenditure is fixed, a reallocated consumer switches M/P^H units of the high-cost firm for M/P^L units at the low-cost firm with no gain in utility. However, the social cost SC of this change is $\Delta SC = \frac{1}{P^H} MC^H \left[\frac{P^H}{P^L}\frac{MC^L}{MC^H} - 1\right] > 0$. The reallocation therefore requires a net increase in labor effort for markets to clear at P^H and P^L. Conversely, under $\beta_e > -1$ an increase in search would generate a net saving in labor.

More generally, $\beta \neq -1$ implies that there is cross-sectional misallocation. Cross-sectional misallocation arises because consumers search until the point at which they exhaust profitable arbitrages provided by relative market prices. However, because relative prices are distorted from marginal costs, counterfactual reallocation of consumers would have first-order welfare effects. Under $\beta_e < -1$, the social marginal cost of labor exceeds the marginal utility of consumption, and a reduction in both labor and consumption would be socially desirable.

Similarly, distortions to firm pricing decisions also have a first-order welfare effect.

The following proposition states this result.

Proposition 1.4.3 (Welfare Effects - Firm Pricing)**.** *Distortions to firms' markups and the pass-through of idiosyncratic TFP shocks have a first-order welfare effect unless optimally set taxes are in place.*

The intuition behind the welfare effect of markups in Proposition 1.4.3 is relatively straightforward. Markups distort the consumption and labor optimization conditions by introducing a wedge between the price at which consumers can acquire goods and the cost of production. Therefore, households would be willing to work more in exchange for more goods. Moreover, rising markups discourage labor given that the less consumption goods can be purchased with the marginal labor effort. If the marginal utility of consumption is higher than the marginal disutility of labor, a rise in markups implies a welfare loss due to inefficiently low labor effort and consumption.

The pass-through of idiosyncratic TFP shocks β_e affects welfare through multiple channels, and its welfare effects depend on the distortions that are in place. To understand why, consider first the effect of an increase in price dispersion (equivalently, a fall in β_e or a rise in its absolute value) on allocative efficiency. The rise in price dispersion lowers the price index $\bar{P}_i^s$ for searchers. Intuitively, because $\bar{P}_i^s$ is a CES price index with $\varsigma > 1$, it falls with price dispersion. This reflects the greater likelihood of finding a low price as price dispersion increases. The fall in $\bar{P}_i^s$ increases consumption.

However, the rise in consumption requires more labor if $\beta_e < -1$. The intuition is the same as the case of changes in search: if $\beta_e < -1$ relative prices overstate relative costs and the efficiency difference between firms is not enough to offset the rise in consumption as consumers are attracted to low-price firms. The rise in price dispersion therefore causes a rise in consumption and labor for the case of $\beta_e < -1$.

The welfare effect of changes to β_e therefore depends on the extent to which it offsets the distortions generated by markups and inefficient search. Recall that markups reduce labor and consumption, which may be offset by a rise in price dispersion. On the other hand, consumption and labor are inefficiently high under $\beta_e < -1$, and this may be offset by a fall in price dispersion.

The effect of β_e on allocative efficiency can be shown mathematically in the following expression derived in Caruso Bloeck (2024):

$$\frac{d\mathcal{W}^{Alloc}}{d\beta_e} \approx U_C \bar{C}_i \left[\bar{S}_i(\varsigma - 1)\left(1 - \frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i}\right)\beta_e + \frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i}\bar{\varepsilon}_f(\beta_e + 1)\right]$$

The first term in the bracket $\bar{S}_i(\varsigma - 1)\left(1 - \frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i}\right)\beta_e$ is negative unless $\frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i} = 1$, so an increase in price dispersion (a fall in β_e) causes a welfare gain through this term. This term captures the social marginal benefit of increasing labor and consumption when there are markups (and $\frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i} < 1$). The second term $\frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i}\bar{\varepsilon}_f(\beta_e + 1)$ captures the excess in labor that is necessary to raise consumption through price dispersion. This term is positive given $\bar{\varepsilon}_f < 0$ and $\beta_e < -1$, and a rise in price dispersion causes a negative welfare effect given by this term. The allocative effect of changes in β_e depends on how the two terms measure up. Naturally, in the case of no distortions ($\frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i} = 1$, $\beta_e = -1$) a marginal changes in β_e have a welfare effect that is second order.

Finally, costs associated with cross-sectional consumption dispersion always favor an allocation with less price dispersion than allocations that only consider allocative efficiency. This is simply due to risk-aversion, and is a result featured in other frameworks in which cross-sectional heterogeneity exists (e.g., see Lorenzoni 2010).

It is worth noting that the welfare effects discussed in Propositions 1.4.2 and 1.4.3 are properties of the environment in which the problem of learning about prices is set. That is to say, these welfare effects arise because the protocol that matches consumers to firms generates some degree of monopolistic competition. However, the results in this paper indicate that the introduction of learning about prices generally makes these distortions worse. The numerical exploration in the remaining section shows that this interaction between learning about prices and distortions associated with monopolistic competition is key to magnifying the cost of inflation uncertainty.

Counterfactual Welfare Effects: Numerical Exploration

This section computes counterfactual welfare effects following Equation (1.21).In all cases, I compute counterfactuals assuming $\sigma_m^2 = 0$ and I calculate welfare for the range of values of equilibrium values between $\sigma_m^2 = 0$ and $\sigma_m^2 = 0.25$. To save space, this section computes the welfare effects of firm idiosyncratic pass-through, markups, and the search elasticity.

As will become clear, this analysis gives us a sense of which distortions are most responsible for the observed welfare effects.

Pass-through of Idiosyncratic TFP Shocks

Figure 1.9 shows the effect of changing the pass-through of idiosyncratic TFP shocks β_e on welfare. The left panel shows that the pass-through rises in absolute with inflation volatility. The right panel shows that the effect of rising idiosyncratic TFP pass-through is negative for welfare. The effect is quantitatively large, in fact the welfare effect of the idiosyncratic pass-through alone is larger than the total welfare cost of inflation uncertainty.

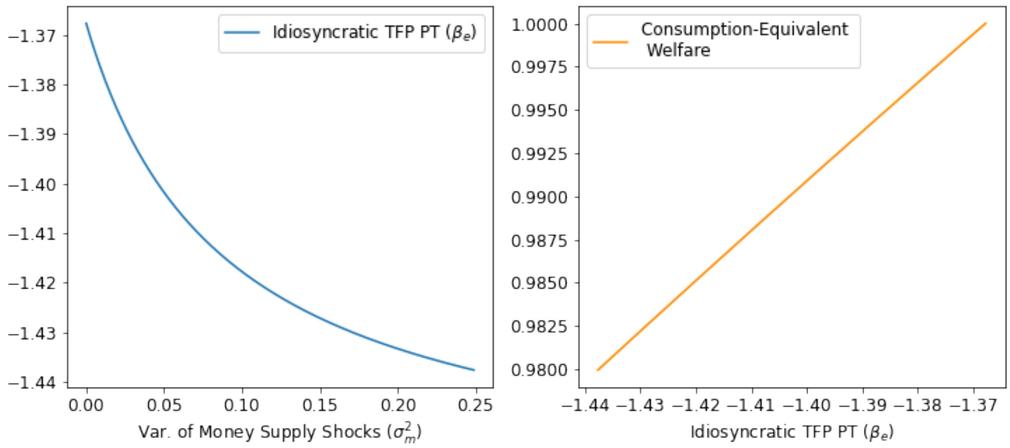

Figure 1.9: Welfare Effect of Idiosyncratic TFP Pass-through

The rise in pass-through has two effects on welfare. The first is that it has reallocative effects. The rise in price dispersion reallocates purchases of searchers to low-cost producers, which has a positive effect on welfare. However, the rise in price dispersion also affects wedges between prices and marginal costs across the firm distribution as in Restuccia & Rogerson (2008) and Hsieh & Klenow (2009). In consequence, the increase in units sold as purchases are reallocated from high- to low-cost producers can increase labor inputs required for production while the increase in consumption is diluted by search costs. The left panel of Figure 1.10 shows that the second effect dominates. Intuitively, starting from $\beta_e < -1$, a rise in the absolute value of β_e raises the wedge between incentives for cross-sectional substitution provided by relative prices and relative marginal costs.

In addition, a higher pass-through of idiosyncratic TFP shocks increases cross-sectional consumption. This is because consumers differ in their price indices, and this variation is larger when there is a larger cross-sectional dispersion in prices. The right panel of Figure 1.10 shows that the welfare effect of rising β_e is coming from higher cross-sectional consumption is somewhat larger than half a percentage point.

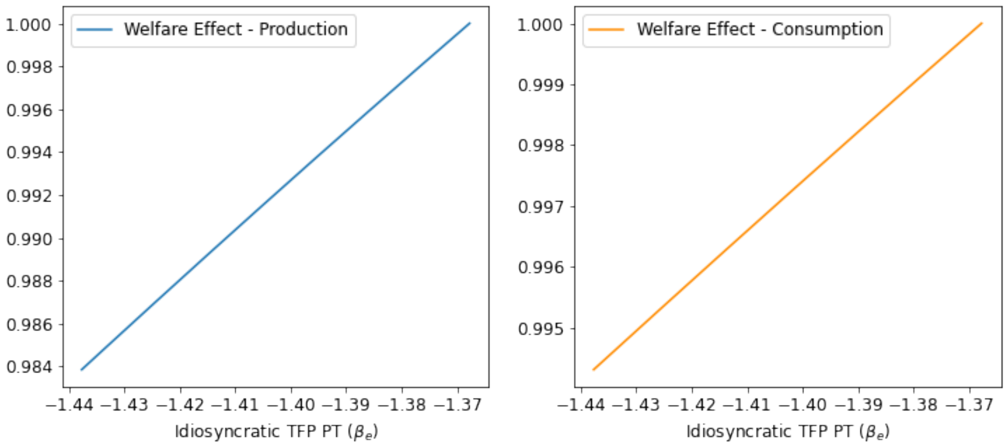

Figure 1.10: Welfare Effect of Pass-through: Channels

Firm markups

The effect of rising markups on welfare is shown in Figure 1.11. The left panel shows that markups rise and the right panel shows that the ensuing welfare loss. The reason is that markups generate a wedge in households' first-oder conditions. The effect of this distortion is equivalent to a fall of about 0.7 percentage points in consumption.

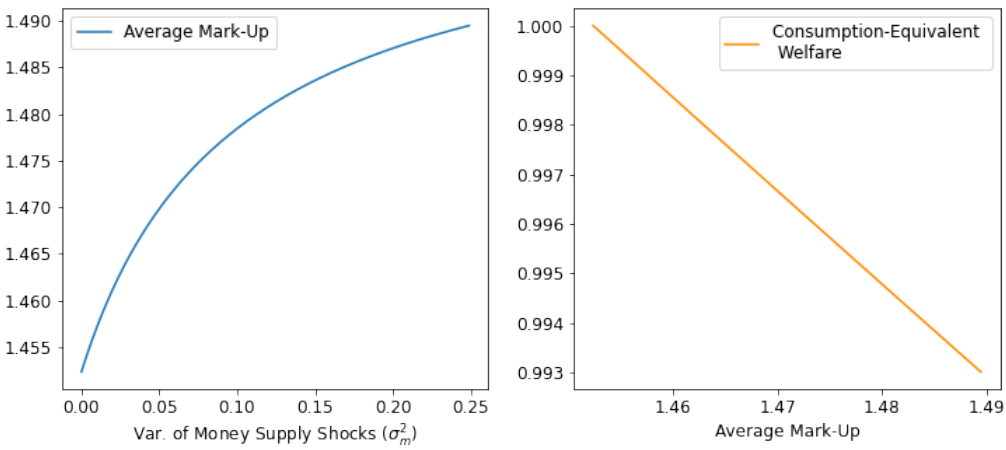

Figure 1.11: Welfare Effect of Markups

Household Search Elasticity

Figure 1.12 shows the change in search elasticity defined as $\psi_{p^c}^S(1-\rho)$ and its welfare effect. The left panel that the search elasticity falls with inflation uncertainty. Intuitively, as inflation uncertainty rises, prices provide less information about relative components and consumers therefore become less responsive. The panel on the right shows that this increases welfare by more than 0.8 percentage points. The effect of falling average search intensity $\bar{S}_i$ is shown in Caruso Bloeck (2024), and equals an increase of 0.35 percentage points in consumption. These results are consistent the allocative effect discussed in Proposition 1.4.2: given that $\beta_e < -1$, search is socially inefficient on the margin.

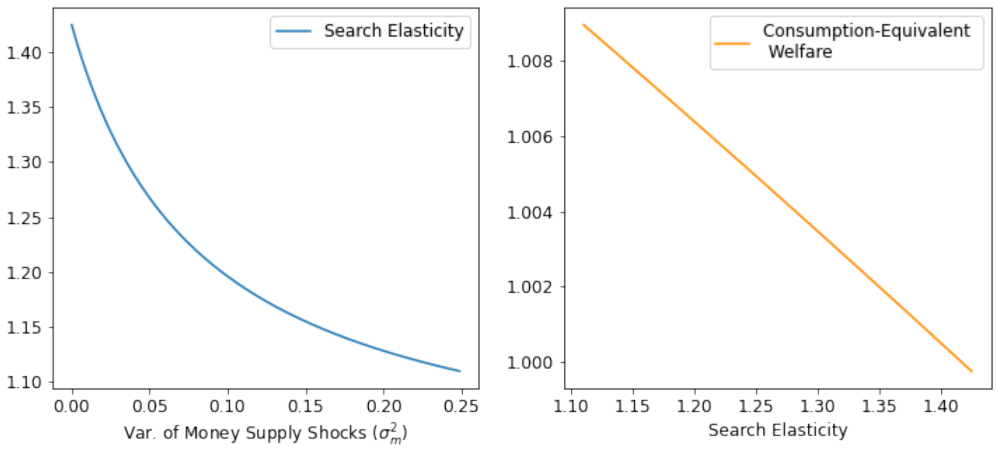

Figure 1.12: Welfare Effect of Search Elasticity

The Effect of Firms' Strategic Response

The previous sections considered the welfare effect from perturbations of the behavior of households and firms. These exercises are useful to understand how each component is affecting welfare. That being said, those exercises are still artificial, in the sense that they do not represent equilibria that would arise from optimal actions of households or firms.

By contrast, this section considers welfare effects that arise as agents reoptimize. More concretely, recall that the welfare function could be written as:

$$\mathcal{W}(\Lambda^H, \Lambda^F, \Sigma)$$

where Λ^H and Λ^F are parameters ruling the optimal actions of households and firms. Because these actions are best responses to each other the environment given by the covariance matrix Σ, we can write:

$$\Lambda^H = \Lambda^H(\Lambda^F, \Sigma)$$
$$\Lambda^F = \Lambda^F(\Lambda^H, \Sigma)$$

The exercise in this section consists on manipulating the actions which firms can reoptimize as the environment and households actions change. To do this, I will compute the following object:

$$\mathcal{W}(\Lambda^H(\bar{\Lambda}^F, \Sigma), \bar{\Lambda}^F, \Sigma) \tag{1.22}$$

where $\bar{\Lambda}^F$ indicates that I am constraining the actions over which firms can reoptimize. Note that while firms are not reoptimizing, I assuming that households are acting optimally given price-setting and aggregate volatility. I will begin by assuming firms set prices as if there were no volatility in the money supply and progressively add margins of reoptimization.

No Firm Reoptimization

Consider first the case in which firms are not able to change their pricing decisions. In this case, it would be expected that households would be subject to two sources of welfare effects. The first one is that as inflation uncertainty rises, consumers still have greater difficulty discerning between the price level and relative prices and leads to a welfare loss. Second, the rise in money supply volatility generally leads to less searching activity. This can have a positive or negative welfare effect depending on whether search is inefficiently high or low.

Figure 1.13 shows the welfare effect of rising inflation uncertainty absent firm reoptimization. Part of these welfare losses can be attributed informational effects that arise because inflation uncertainty makes it harder for consumers to discern between high- and low-cost prices. However, the cost of these errors is typically small. Most of the welfare loss arises because of the allocative effects. Intuitively, given that $\beta_e < -1$ reducing search is socially desirable. Absent firm reoptimization, the effect of inflation uncertainty is an increase in welfare that is equivalent to about one percentage point.

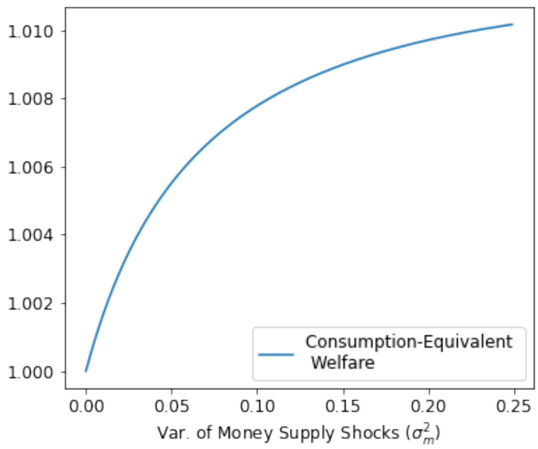

Figure 1.13: Welfare Effect of Rising Money Supply Volatility

Firm Reoptimization

I now consider what happens when firms are able to change their pricing function. I will consider two cases. First, I will allow firms to optimally reset their idiosyncratic cost pass-through β_e. Then, I will allow firms to reset their markups. The reason for allowing these changes is that these were the factors that had a greater impact on welfare based on the analysis in section 1.4.3.

The effects are shown in Figure 1.14. The left panel shows the effect of only changing markups in the solid green line and only changing the idiosyncratic pass-through in the dash-dotted blue line. Markups lower welfare gain by about 0.6 percentage points relative to the case of no firm reoptimization. In contrast, allowing firms to reset their idiosyncratic pass-through β_e more than compensates the gain from lower search. This result implies that the distortion of relative prices leads to an overall worse allocation of consumers to firms.

The right panel of Figure 1.14 shows that allowing firms to reset their idiosyncratic pass-through and markups accounts for essentially all of the welfare losses. In fact, these actually slightly overshoot the original welfare loss because the fall in money shock pass-through increases welfare. Still, aggregate pass-throughs have a minimal welfare effect. Practically all of the welfare effects can be attributed to the loss in market competitiveness, lower allocative efficiency and the rise in cross-sectional consumption generated by inflation uncertainty.

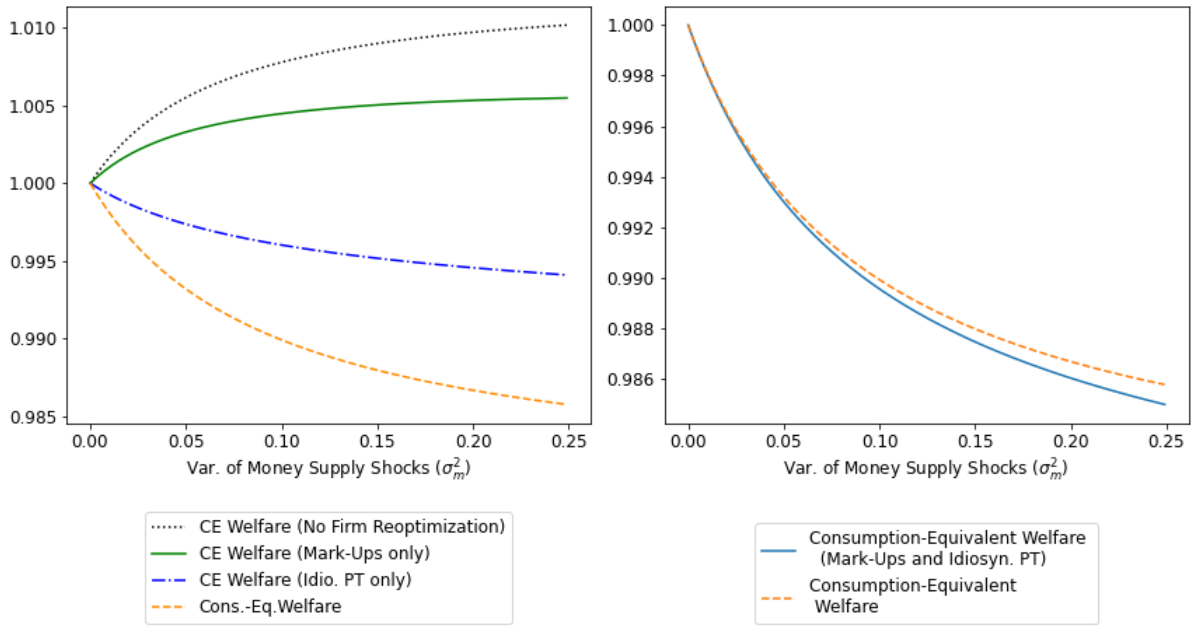

Figure 1.14: Welfare Effect due to Firm Response

Summary

The numerical exploration in this section indicates that the main contributors to welfare losses are the rise in markups, and the rise in cross-sectional price dispersion. These factor distort households' first order conditions and deteriorate the matching process consumers to firms. Taken together, these factors explain practically all of the welfare losses given that additional losses attributable to aggregate volatility or inference errors are small.

These welfare effects are considerably higher than those that could be attributed to aggregate volatility alone. The key force that is magnifying the welfare effects is the interaction between learning about prices and allocative distortions that are a feature of the environment being considered. Namely, while these distortions exist absent inflation uncertainty, the addition of inflation uncertainty makes these distortions worse.

These mechanisms differ markedly from other microfoundations for the costs of inflation, such as nominal rigidities or the redistributive effects of inflation. the following section offers a more detailed discussion of this comparison and offers closing remarks.

1.5 Conclusions

This paper proposes a theory in which inflation uncertainty affects the efficiency of the price system. These effects arise because inflation conceals information about relative prices and firms adapt their price setting strategies as consumers change their behavior. In equilibrium, the informational content of prices and market structure arise from optimal actions of consumers and firms given the environment.

The study of the implications of inflation uncertainty for market competitiveness is an area that has not received much attention in the last 30 years. It differs markedly from the way in which the profession currently thinks about the costs of inflation which emphasize costs associated with nominal rigidities or redistributional effects of inflation.

The results indicate that the way households and consumers behave in volatile inflation environments have rich implications for the allocation of resources and welfare. Generally, inflation uncertainty makes it harder for consumers to infer relative prices, making their search response less sensitive to price signals. This allows firms to raise markups without fear of losing customers. The rise in markups is an important channel through which inflation uncertainty reduces welfare in the model.

Inflation uncertainty also increases cross-sectional price dispersion. In my model, this affects welfare negatively by generating cross-sectional misallocation and increasing cross-sectional consumption dispersion. Together with the effect of markups, these channels explain almost all of the welfare effects of inflation uncertainty.

These channels are the main drivers of welfare because the model is set up in such a way that money supply volatility actually has a very small effect on real allocations. This implies the model generates non-trivial welfare effects despite the fact that money supply shocks themselves do not contribute markedly to aggregate volatility.

The recent inflationary experience around the world is likely to bring new data that could be used to study how inflation affects consumers' cross-sectional substitution patterns. This paper indicates that rigidities in cross-sectional substitution patterns may be responsible for non-trivial welfare effects, particularly for consumers. While this paper provides a theoretical underpinning highlighting why these channels are likely to be important, empirical exploration is likely to be an interesting and fruitful direction for future research.

Chapter 2

News of Disinflation and Firms' Expectations: New Causal Evidence

This chapter is based on published work with Miguel Mello and Jorge Ponce of the Banco Central del Uruguay (BCU) in the *Journal of International Money and Finance*. The views and opinions therein are those of the authors and do not necessarily represent those of the institutions to which they are affiliated. See Caruso Bloeck et al. (2023).

2.1 Introduction

Expectations about future inflation are a first-order driver of inflation dynamics in broad classes of forward-looking macroeconomic models. Nevertheless, empirical evidence on the factors behind the formation of inflationary expectations is limited to relatively recent literature. Due to data availability, most studies are restricted to countries with fixed monetary policy frameworks and consolidated monetary regimes with low and stable inflation rates. Examples are Coibion et al. (2018) for New Zealand, Coibion, Gorodnichenko & Ropele (2020) for Italy, Coibion et al. (2022*a*) for households in the U.S, and Hunziker et al. (2022) for Switzerland. As a consequence, the empirical evidence is silent about how inflation expectations respond to news of policies aimed at permanently lowering inflation. Importantly, this is a fundamental question for central bankers trying to lower inflation without an unnecessarily large recession. The cost in terms of activity depends crucially on the response of expectations to news of permanent disinflation, which varies markedly across macroeconomic models (see, for instance, Taylor 1982, Buiter & Miller 1983, Fischer 1986, Ball 1994*a*). Existing empirical evidence studying episodes of disinflation can help distinguish between these models (Sargent 1982, Easterly 1996, Ball 1994*b*, Hazell et al. 2022). Even so, these studies rely on observational data which makes it impossible to establish a causal relationship between the news of disinflation and expectations. This task would require external variation in the news of disinflation to credibly identify the effect of the news. A randomized control trial (RCT) experiment would generate such random variation.

In this paper, we carry out an RCT experiment to causally quantify the effect of the news of disinflation on the economic expectations of firms. More precisely, we randomly split into two groups the sample of a monthly economic expectations survey to firms in Uruguay. During the wave of June 2021, the treatment group receives a short paragraph consisting of a summary of a published article informing about the new monetary policy framework and the policymakers' forecast of inflation for the coming years. A quote from the president of the central bank stating that lower inflation will make it easier for locals to trust the local currency and a link to additional information was also included. Overall, the treatment consists of information about a comprehensive reform in the monetary policy framework, which is a major difference from previous studies. Then, both the treatment and the control group, which does not receive the information, answer the same questions about inflation and economic expectations as in the other waves of the survey.

We find that treated firms lower their inflation expectations by about half a percentage point at both the 12- and the 24-month horizons. This represents a fall of about 6 percent on an average inflation expectation of around 8%. While the magnitude of the effect varies with time, the effect seems to persist until the end of our sample. Consistently, similar causal impacts occur in the firms' own cost expectations. In addition to that, treated firms also anticipate lower growth. Unlike inflation and cost expectations which fell for both the 12 and 24-month horizons, treated firms anticipate temporarily slower growth only for the current year.

We also study how the results vary depending on the attentiveness of firms to monetary policy. To proxy for attentiveness, we ask firms whether they had read the newsletter on monetary policy which had been distributed three months before the treatment, i.e. in March 2021. The results suggest that attentive firms are generally more responsive to the treatment. Unfortunately, given that the sample of attentive firms is small, we are often unable to reject differences between attentive firms and the rest.

The empirical evidence provided in this paper highlights the importance of central bank communication to move firms' inflation expectations towards the central bank's objectives. Moreover, the causal effects of the news of disinflation persist over time with only a short-term impact on expected activity. Although extrapolation of these results to other circumstances is not direct, one could reasonably argue that similar results will hold in other economies, especially in those with long records of lower inflation and higher credibility of the central bank than in our case study.

Our experiment circumvents two major obstacles that have not permitted a study like this one before. First, events of comprehensive reforms in monetary policy frameworks intended to lower inflation are very rare in the recent period. However, in our case study, the reform introduced in 2020 consisted of a change in monetary policy instrument and a change in the inflation target band which became lower and narrower. In addition, the reform includes several improvements in terms of transparency, communication, and the information provided by the central bank to the general public. The intention behind the reform is to bring down inflation significantly in the coming years and reduce the reliance of the local economy on foreign currency.

Second, an RCT requires collecting information on expectations in real-time. Very few countries systematically collect information on firms' expectations, Uruguay being one the few exceptions. We use the Uruguayan firm inflation expectations survey

(IES) to carry out the RCT. The survey is conducted monthly since 2009 to a sample of firms, which is representative of the universe of private firms in Uruguay with at least 50 employees. It is one of the few systematic and quantitative surveys of firms' macroeconomic expectations (Candia et al. 2021). The survey has fixed and rotating modules, with an open module in June that we use for research purposes. Importantly, the IES is monthly conducted by the Instituto Nacional de Estadísticas, the national statistics office, on behalf of the Banco Central del Uruguay, the central bank.[1] Hence, the RCT is embedded in a routine survey, which reduces to the minimum the likelihood that firms actually understand that they are being treated as part of an academic project and respond strategically.

The RCT is based on the assumption that at least some of the information provided is not known to firms. If this assumption did not hold and firms knew the whole content of the information provided, we would expect results to be null. As a consequence, we would be unable to determine whether the reason for null results is that the treatment is not credible or that expectations already include that information. An extensive literature documents that firms are not aware of all information, even if it is part of the public domain (Andrade & Le Bihan 2013). In our setting, Borraz & Mello (2020), using the same survey we use in this paper, document that many firms are unaware of the inflation target and also the inflation rate. In spite of concerns about whether firms already know the information provided to them, our results show that expectations do respond to the treatment in ways that cannot be explained by being previously informed. This is comforting, as it clears doubts about the appropriateness of our assumption.

The remainder of this paper is structured as follows. Section 2.2 discusses the relevant literature and our contribution. Section 2.3 presents the source of data, explains the reform to the monetary policy framework, and describes the RCT exercise. Section 2.4 shows the effects of treatment. Finally, Section 2.5 discusses our interpretation of the results and offers concluding comments.

2.2 Related Literature

This paper is motivated by rich literature studying how the acquisition of information and the formation of expectations affect the economy. The seminal paper by Mankiw & Reis (2002) introduces the notion of information stickiness in the form of friction in the acquisition of information. Information stickiness has important implications for macroeconomic phenomena, for example, generating costly disinflations and hump-shaped impulse responses to monetary policy shock. Reis (2006*a*,*b*) studies consumers' and producers' incentive to acquire costly information, which endogenously generates inattention and sluggish dissemination of information.

Empirical work is also rich and active on this topic. Carroll (2003) presents a model in which households probabilistically acquire information from rational agents and tests whether the model can match household inflation expectations. Coibion & Gorod-

[1]We are extremely grateful to colleagues at the Instituto Nacional de Estadísticas for their effort and support in making the RCT exercise possible.

nichenko (2015*a*) propose a test for whether expectations are formed under full-information rational expectations. The paper finds expectations underreact to information, consistent with models of sticky information. Similarly, Andrade & Le Bihan (2013) show that professional forecasters in the Eurozone do not systematically update their forecasts following new information and disagree even when they do.

Furthermore, the use of randomized control trials (RCTs) has significantly contributed to our understanding of how information is incorporated into expectations. For example, Coibion et al. (2022*a*), Haldane & McMahon (2018) and Binder & Rodrigue (2018) use RCTs to analyze the effect of communication about inflation on households expectations.

While there is a widespread view that firm expectations should be particularly important for inflation dynamics, information on firm expectations is particularly scarce (Bernanke 2007, Candia et al. 2021, Coibion, Gorodnichenko, Kumar & Pedemonte 2020). For firms, Coibion et al. (2018), Coibion, Gorodnichenko & Ropele (2020) and Hunziker et al. (2022) carry out experiments in New Zealand, Italy and Switzerland, respectively. The former paper studies the extent to which firms update their beliefs when presented with new information, while the latter one tracks effects on firm decisions.

This literature focuses on countries with stable inflation and consolidated inflation targeting regimes. By contrast, in this paper, we study the effect of information about long-run inflation and a comprehensive review of the monetary policy framework. Episodes such as these are very rare and a rigorous study based on an RCT requires having a running survey collecting information on firm inflation expectations. The number of countries with such surveys is extremely limited (see Candia et al. 2021, Coibion, Gorodnichenko, Kumar & Pedemonte 2020). To the best of our knowledge, Coibion, Gorodnichenko, Knotek II & Schoenle (2020) is the only paper studying a change in the monetary policy framework.

Furthermore, the overall inflation environment in our case of study, Uruguay, is markedly different from those of previous papers. Inflation had been consistently above 20% in the 1990s and has been around 8% in the last 10 years. As a consequence, firms are more informed about inflation relative to countries such as Italy and New Zealand (see Borraz & Zacheo 2018). As a consequence, the results of our experiment are more likely to be informative about the possible effect of communication in a typical developing country than in the previous literature.

Moreover, inflation expectations in Uruguay are anchored consistently above the inflation target. A recent paper by Shapiro & Wilson (2021) suggests that this may also be the case in the US, although the inflation target is not explicit as in the case of Uruguay. Similarly, Coibion et al. (2022*b*) show that US household inflation expectations are also consistently above target, and conduct an RCT to quantify the effect of different communication strategies. Like us, they find that communication is effective, although the effect depends on the precise form of communication.

The question of how expectations react to information about disinflation is key to understanding the cost of said effects of disinflation. In theory, immediate disinflations can be costless if agents are forward-looking (Buiter & Miller 1983), or can even be expansionary if disinflation is credible (Ball 1994*a*). While empirical evidence can help us distinguish between these models, it is impossible to determine causal drivers of inflation expectations without an exogenous source of information. This paper contributes to filling that gap.

2.3 The Randomized Control Trial Experiment

We conduct a randomized control trial (RCT) experiment which assigns an informational treatment to a random sub-sample of firms in the Uruguayan Inflation Expectations survey (IES) conducted to firms. In this Section we describe the IES, the overall macroeconomic setting and monetary policy reform, and the treatment exercise.

2.3.1 Data Source: The Inflation Expectations Survey

Uruguay is one of the few countries that routinely survey firms' inflation and cost expectations. The survey is commanded by the central bank and carried out by the Instituto Nacional de Estadísticas (INE), the national statistics office. The survey is sent to 500 firms every month with an average response ratio of 64%.

The survey is representative of the universe of Uruguayan private companies with at least 50 employees, being one of the few systematic and quantitative surveys of firms' macroeconomic expectations (Candia et al. 2021). The survey, however, does not cover the agricultural and financial sectors. Table 2.1 from Borraz & Mello (2020) shows that the survey is broadly representative of the Uruguayan economy.

Sector	Sample	Population
Manufacturing	41.48	46.60
Trade and commerce	29.99	23.06
Services	18.40	14.96
Health	4.47	11.61
Primary activities	2.36	1.06
Education	1.99	1.73
Utilities	0.58	0.74

Source: Borraz & Mello (2020)

Table 2.1: Firms' distribution by sectors - sample and population in percentage

The IES began in October 2009 and continues towards the present. In this project, we use the data from October 2020, when the sample was refreshed, to December 2021. The treatment takes place in June 2021. The resulting data set is an unbalanced long panel with a total of 15 months and 3,973 observations. During the sample period, 333 firms completed the survey at least once, while 256 of the firms answered the questionnaire more than 66% of the time (10 months).

The IES surveys firms' inflation and cost expectations every month. In addition, there are rotating modules that survey the frequency of price and cost changes, and access to financing among other variables. Finally, there is an open module in June, in which we introduced questions about expectations on GDP growth, firms' investment, and exchange rate forecasts.

The IES is a vital input to the BCU assessment of inflation expectations, which in turn is used to inform monetary policy in the country. That makes this data source extremely unique since a different survey prepared specifically for an experiment such as this one would not have this feature.

2.3.2 Monetary Policy Framework and Reform

The starting point of this paper is a significant reform in a monetary policy framework that includes a change in the policy instrument, from monetary aggregates to the short interest rate, and a change in the inflation target band to one that is narrower and lower. In addition, the central bank revised its strategy for communication and information sharing with the public. Further details on each of these elements are presented below.

In our case of study, the policy instrument was changed from monetary aggregates to the short-term interest rate starting in September 2020. This was the second time the central bank changed the monetary policy instrument in the last decade, which had been the short-term rate since the adoption of the inflation targeting regime until mid-2013. The main motivation behind the change in the instrument is the greater transparency of the interest rate. As a policy tool, the interest rate is more efficient and easier to observe, although the transmission channels are essentially the same.

The inflation target range was changed from [3-7] to [3-6] percent.[2] At face value, this looks like a relatively minor change. However, as Figure 2.1 shows, inflation expectations have been consistently above the upper bound of the target. Inflation has also been above the target for most of the last 15 years despite being relatively stable. Bringing inflation within the target range, therefore, requires lowering inflation and inflation expectations significantly. Like the change in the instrument, the target range had been changed in 2013, broadening from [4-6] to [3-7].

An important aspect of the reform is the improvement of the communication of the monetary policy. The frequency of meetings of the Monetary Policy Committee was doubled from four annual meetings to eight. Minutes of the Committee meetings started to be published in April 2020. These minutes highlight the assessment supporting decisions of the Committee and individual members and information about the policy recommendations of technical staff. These actions are intended to increase transparency and build credibility. Technical quarterly monetary policy reports started to include detailed descriptions of the projection models used to inform policymakers in addition to their projection results. Technical boxes were introduced to explain the methodology and assumptions supporting the analysis.

[2] See https://www.gub.uy/presidencia/comunicacion/noticias/uruguay-modificara-rango-meta-inflacion-segundo-semestre-2022

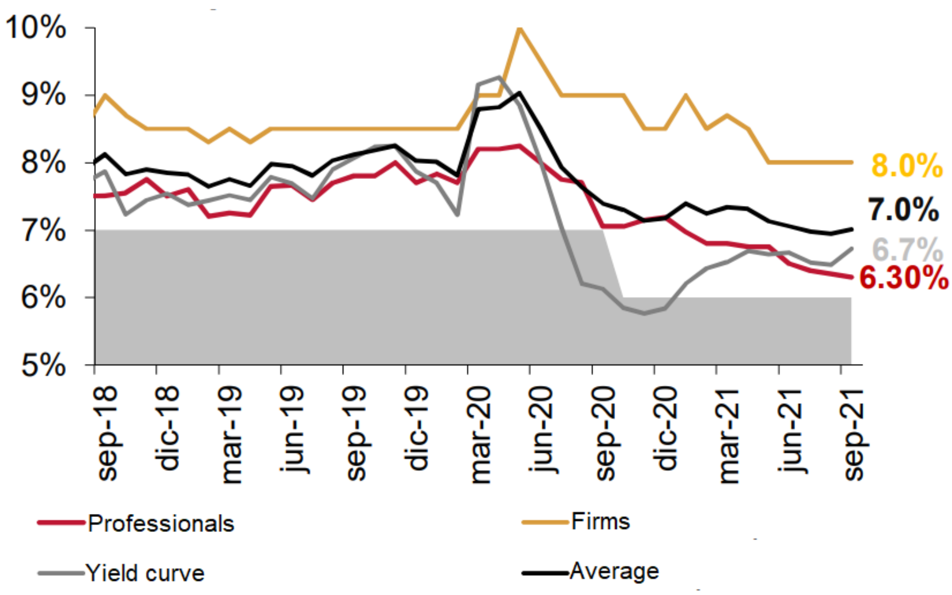

Figure 2.1: 24-Months Inflation Expectations

While the transparency and communication efforts mentioned in the previous paragraph are mainly directed to professionals and markets, a substantial effort was made to open communication channels with the general public. A newsletter using simple language to communicate macroeconomic trends, monetary policy objectives, and decisions started to be published in March 2021.[3] The first newsletter was also distributed to commercial chambers and unions in an attempt to directly reach firms.

The main objective of the reform is to bring down inflation significantly in the country. Uruguay has a history of high inflation: inflation has hovered around 10% in the last 10-15 years. Figure 2.1 shows that inflation and inflation expectations have been consistently above the upper bound of the inflation target band in this period. However, inflation expectations seem anchored below 10%.

The reform in the monetary policy framework is taken in parallel with additional measures to encourage the use of local currency and decrease the reliance on the local economy in US dollars. Dollarization is a phenomenon that is extremely widespread. About 75% of deposits in the banking sector are denominated in US dollars, 90% of exports and 65% of imports are invoiced in US dollars (Boz et al. 2022), and even dollar invoicing of domestic transactions is frequent (Licandro & Mello 2019). Licandro & Mello (2016) document that a common practice in Uruguay is to express large sums of money in US dollars, a phenomenon they call cultural dollarization.

De-dollarization is an important aspect of the reform agenda given that reliance on the dollar weakens the transmission of monetary policy to the local economy. As a consequence, policymakers in Uruguay will often argue that monetary policy will be more robust once a regime with lower inflation has been achieved.

[3] See https://www.bcu.gub.uy/Politica-Economica-y-Mercados/Paginas/La-inflacion-y-sus-expectativas.aspx.

2.3.3 Treatment Description

The informational treatment consists of a short paragraph at the beginning of the survey which tells firms that policymakers forecast inflation will fall significantly in the coming years and that this forecast is based on improvements in monetary policy. We also include a quote from the governor of the central bank saying low inflation is important for citizens to trust their currency.

The experiment was carried out in the June 2021 wave of the inflation expectations survey (IES). Firms were randomly assigned into a treatment and a control group. For the treatment group, the survey began with the following paragraph:

Inflation in the 12 months to April was 6.76%. In 2020 inflation had been 9.41%. Authorities forecast a gradual decline toward 3.7% in 2024 as a result of the monetary policy measures adopted. The president of the Central Bank of Uruguay stated that "with inflation at 3% or 4%, it will be easier for the Uruguayan to trust their currency." For more information, you can consult this link:
https://www.gub.uy/ministerio-economia-finanzas/comunicacion/noticias/rango-meta-inflacion-sera-3-6-ciento-2022-se-preve-37-ciento-para-2024

For the control group, the survey began with the following text:

Inflation in the 12 months to April was 6.76%. In 2020 inflation had been 9.41%.

We can see that the treatment consists of several parts. The first is information about inflation, which is shared by both treatment and control groups. As a consequence, differences in responses capture the effect of the rest of the paragraph. The next sentence claims authorities forecast inflation to decline significantly. The paragraph also tells firms that the forecast is based on monetary policy measures. We chose this over a technical explanation of the link between monetary policy and inflation to avoid making the treatment overly long and complicated. The next sentence gives the respondents a quote by the president of the central bank saying that lower inflation will increase confidence in the local currency. This sentence is easy to understand for Uruguayans, as they are familiar with the pervasiveness of dollarization and probably find the benefits of a more reliable currency intuitive. Last, there is a link to an article which is an extended version of the paragraph included in the treatment. We are able to check whether respondents click on the link, but unfortunately, too few respondents click to draw any meaningful conclusions.

It is evident that the treatment is a complicated object with many parts. Ideally, we would have multiple treatment arms and we would add each element sequentially to isolate the effect of each part. However, because of the small sample size, around 280 respondents each month, this was infeasible.

The text for the control group also contains some information that may not be known to firms and their expectations may reflect this information. In previous experiments, similar treatments have been shown to affect firm expectations (Coibion, Gorodnichenko & Ropele 2020). Even so, we designed the experiment this way so that we could isolate the effect of the information on disinflation, which required comparing inflation at the time with the forecast of future inflation. In addition, we find that the distribution of expectations becomes less dispersed in the control group in June. The reduction in dispersion is a common result in this literature and may reflect information incompleteness, although we lack a proper baseline that identifies this causal effect.

The experimental design requires that the information provided is not known to firms. If it were, we would not expect to see any differences between groups. More importantly, we would be unable to determine the cause of null results, which may arise either because the information provided is not credible to firms or because the information is already incorporated into firms' expectations.

There is growing empirical evidence that firms make their forecasts based on incomplete information, in particular about monetary policy. Coibion et al. (2018) find that firms with larger incentives to track and process information are more attentive to inflation information. In a context of moderate inflation like Uruguay, we would expect firms to be more attentive to inflation. Candia et al. (2021) find that firms' expectations in Uruguay respond to even transitory changes in inflation, which they argue is an indication that firms pay attention to inflation information. Even so, Borraz & Mello (2020) document that while about 60% of firms in Uruguay are informed about the inflation rate, the share of firms informed about the inflation target is roughly one-third.

Our results do show responses to expectations. Moreover, the results are consistent with the previous literature studying inflation expectations with RCTs. This clears the doubt that the treatment does not contain any new information.

2.3.4 Pretreatment Characteristics

This section shows statistical tests to evaluate the process of randomization. The INE, which carries out the survey, assigned the treatment randomly. While this minimizes the risk that the treatment was unbalanced, we still test for systematic differences between treatment and control groups. More precisely, we estimate the following regression equation by OLS:

$$Y_i = \alpha + \beta T_i + \varepsilon_i, \tag{2.1}$$

where Y_i is a variable of interest and T_i is a dummy variable that takes the value one if firm i is in the treatment group and zero otherwise. We estimate this equation for a broad array of firm characteristics previous to the treatment, i.e. for the month of May 2021. Under successful randomization, we would expect β to be zero.

The results of the regression are presented in Table 2.2. The table shows that the null hypothesis is not rejected for any of the variables at conventional levels. Overall, the results are consistent with successful randomization.

Variable	Estimated β
Lag 12-Mo. Inflation Expectations	-0.022
	(0.213)
Lag 24-Mo. Inflation Expectations	-0.008
	(0.196)
Read Newsletter	0.042
	(0.054)
Small Firm	-0.031
	(0.055)
Medium Firm	0.015
	(0.062)
Services	0.023
	(0.059)
Manufacturing	-0.074
	(0.050)
Obs	259

Statistical significance: * 10%, ** 5%, *** 1%.

Table 2.2: Pretreatment Characteristics

2.4 Results

In this section, we present the main results.

2.4.1 Effect on Inflation Expectations

This section shows the results of the treatment effect on firms' inflation expectations for the next 12 and 24 months. Visual evidence of the treatment effect is in Figures 2.2 and 2.3. The graphs show that the distribution of 12-month inflation expectations moves to the left for the treatment group in June and moves in the opposite direction for the control group. At the 24-month horizon, the distribution of non-treated firms becomes more concentrated around the same center, while the mode of the distribution of the treated firms moves to the left.

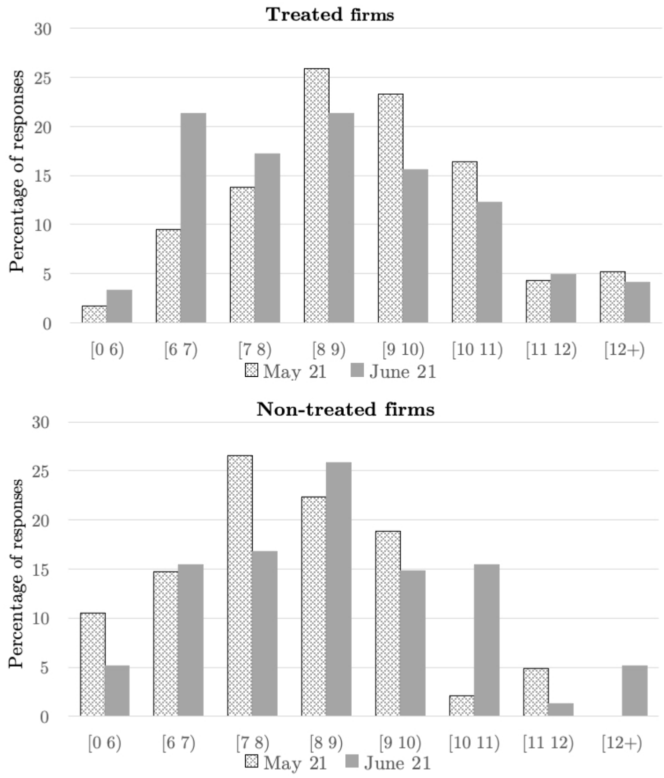

Figure 2.2: Histograms of 12-Month Inflation Expectations

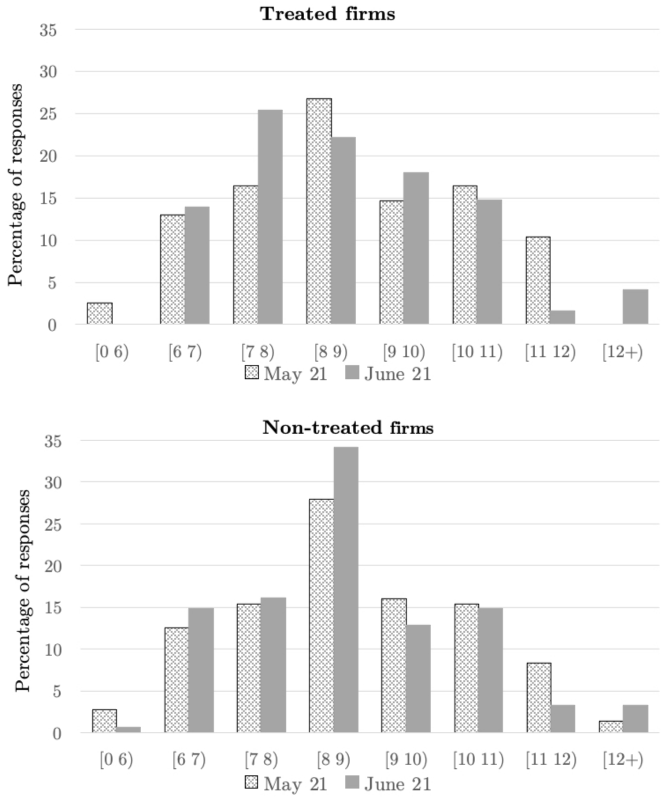

Figure 2.3: Histograms of 24-Month Inflation Expectations

Given that the preliminary evidence looking at the distributions of inflation expectations indicates some reduction in inflation expectations, we proceed to statistically test whether the treatment effect is different from zero. In order to do that, we estimate the following regression model:

$$E_{it}[\pi_{i,t+h}] = \alpha + \beta\, T_i\, D_t + \gamma_t + \lambda_i + \varepsilon_{it}, \tag{2.2}$$

where $E_{it}[\pi_{i,t+h}]$ is firm i's inflation expectation between periods t and $t+h$, α is a constant, γ_t and λ_i are time and firm fixed effects, T_i is a dummy indicating the firm was treated and D_t is a dummy variable taking the value one for the month of the treatment, i.e. June 2021, and after. The parameter of interest is β in Equation 2.2, which is a difference-in-difference estimator of the treatment effect. Given that our experimental framework assigns the treatment exogenously, the treatment effect can be estimated by ordinary least squares. In addition, we cluster standard errors by firm given concerns that errors may be serially correlated.

The results are shown in Table 2.3. The treatment is successful in reducing inflation expectations, which fall by about half a percentage point at both horizons on average after the treatment.

	12-Month Inflation Expectations	24-Month Inflation Expectations
β	-0.566***	-0.544***
	(0.094)	(0.103)
Obs	3754	3754
R^2	0.621	0.588
Num. Firms	293	293

Statistical significance: * 10%, ** 5%, *** 1%.
Standard errors clustered by firm.

Table 2.3: Treatment Effect on Inflation Expectations

Next, we are interested in seeing how the effect of treatment evolves with time, in particular, whether it is persistent or short-lived. In order to do that, we estimate the following regression equation:

$$E_{it}[\pi_{i,t+h}] = \alpha + \beta\ T_i\ D_t + \beta^{Jul}\ T_i\ D_t^{Jul} + \beta^{Aug}\ T_i\ D_t^{Aug} + \gamma_t + \lambda_i + \varepsilon_{it}, \qquad (2.3)$$

where D_t^{Jul} and D_t^{Aug} are dummies indicating the observation on and after July and August, respectively. The parameters β^{Jul} and β^{Aug}, therefore, measure the incremental impact of the treatment after one and two months respectively.

The results for Equation 2.3 are displayed in Table 2.4. The table shows the effect on impact is smaller for measures of inflation expectations than in the pooled estimate. For the 12-month inflation expectations, about three-fifths of the effect is reversed in July, but differences between the treatment and control group broaden again in August. For the 24-month inflation expectations, we see that differences in inflation expectations are widening as time goes by.

	12-Month Inflation Expectations	24-Month Inflation Expectations
β	-0.490***	-0.376***
	(0.128)	(0.132)
β^{Jul}	0.290**	-0.154
	(0.138)	(0.125)
β^{Aug}	-0.517***	-0.310**
	(0.125)	(0.155)
Obs	3754	3754
R^2	0.623	0.589
Num. Firms	293	293

Statistical significance: * 10%, ** 5%, *** 1%.
Standard errors clustered by firm.

Table 2.4: Dynamics Treatment Effect on Inflation Expectations: Incremental Impact

To have another look at the persistence of the effect of the treatment on inflation expectations we estimate the following regression equations:

$$E_{it}[\pi_{i,t+h}] = \alpha + \beta^M\ T_i\ D_t^M + \gamma_t + \lambda_i + \varepsilon_{it}, \qquad (2.4)$$

where D_t^M is a dummy variable equal to one for observation on and after month M. For instance, for $M = June$, then $D_t^M = D_t$ in Equation 2.2. Table 2.5 shows the results of estimating Equation 2.4 for each month between June 2021, the month in which the treatment takes place, and December 2021, the end of our sample. These results add evidence to the persistence of the effects of the treatment. Inflation expectations of treated firms are significantly lower than those of the control group after the treatment and up to the end of the sample in December 2021.

Estimated β^M for $M =$	12-Month Inflation Expectations	24-Month Inflation Expectations
June	-0.566***	-0.544***
	(0.094)	(0.103)
July	-0.506***	-0.520***
	(0.093)	(0.112)
August	-0.590***	-0.565***
	(0.096)	(0.116)
September	-0.689***	-0.597***
	(0.089)	(0.101)
October	-0.843***	-0.703***
	(0.091)	(0.100)
November	-0.709***	-0.590***
	(0.091)	(0.102)
December	-0.644***	-0.496***
	(0.101)	(0.110)
Obs	3754	3754
Num. Firms	293	293

Statistical significance: * 10%, ** 5%, *** 1%.
Standard errors clustered by firm.

Table 2.5: Dynamics Treatment Effect on Inflation Expectations: Persistence

2.4.2 Placebo Test

The previous estimates of treatment effects are based on the assumption of parallel trends. Differences in inflation expectations between the treated and control group that predate the month of treatment would imply a failure of randomization and invalidate our empirical approach while still generating statistically significant estimates. While this possibility is unlikely in an experimental setting, we can carry out placebo tests to verify that differences in inflation expectations arise after June, the month of treatment. The placebo tests consist of estimating Equation 2.4 for months before the treatment. If the parameter β^M was different from zero, that would be an indication that differences between the treatment and control group predate the experiment, casting doubt on our results.

Table 2.6 shows the results of estimating Equation 2.4 for the three months prior to treatment: March, April, and May. For comparison, we also include the months of the treatment, i.e. June 2021. While the table shows some differences between the treated and control group before the date of the experiment, the differences are not statistically significant. In addition, the differences are small in magnitude compared to the treatment effect which is approximately a half of percentage point.

Estimated β^M for $M =$	12-Month Inflation Expectations	24-Month Inflation Expectations
March	0.152	0.079
	(0.112)	(0.145)
April	-0.062	-0.032
	(0.130)	(0.126)
May	-0.088	-0.122
	(0.118)	(0.136)
June	-0.566***	-0.544***
	(0.094)	(0.103)
Obs	3754	3754
Num. Firms	293	293

Statistical significance: * 10%, ** 5%, *** 1%.
Standard errors clustered by firm.

Table 2.6: Dynamics Treatment Effect on Inflation Expectations: Placebo Test

2.4.3 Heterogeneous Effects

One key dimension of heterogeneity between firms is whether they pay attention to monetary policy events. The fact that some firms are more attentive to monetary policy is likely to reflect the usefulness of this information for firms, so the fact that some firms pay more attention to monetary policy is very likely to be endogenous.

We proxy for firm attentiveness to monetary policy by asking firms whether they had read the BCU newsletter which had been distributed in April. The newsletter is a piece of communication about a monetary policy designed specifically for firms. For that reason, firms reading the newsletter are likely to be better informed about monetary policy, in addition to being more attentive to it.

A priori, it is not clear whether firms that read the newsletter should react more or less strongly than inattentive firms. For example, attentive firms may already be aware of the information provided to them, in which case we would expect no effect. On the other hand, firms may be attentive to information on monetary policy because it is easier for them to process the information, in which case attentive firms may respond more strongly to the treatment.

To estimate whether the treatment affects attentive and inattentive firms differently, we estimate the following equation:

$$E_{it}[\pi_{i,t+h}] = \alpha + \beta^{NL}\, T_i\, NL_i\, D_t + \gamma_t + \lambda_i + \varepsilon_{it}, \tag{2.5}$$

where NL_i is a dummy indicating the firm read the newsletter. The parameter β^{NL} estimates whether attentive firms are more sensitive to the treatment than non-attentive firms. We also estimate the following equation to assess the incremental impact of reading the newsletter:

$$E_{it}[\pi_{i,t+h}] = \alpha + \beta\, T_i\, D_t + \beta^{NL}\, T_i\, NL_i\, D_t + \gamma_t + \lambda_i + \varepsilon_{it}. \tag{2.6}$$

The results are in Table 2.7. The table shows that firms reading the newsletter are more sensitive to the treatment, in particular for the 12-Month horizon: -0.645 versus -0.566 in the baseline model. However, the difference is not large enough to be statistically significant. Moreover, the incremental effect of having read the newsletter is not statistically significant in the results from estimating Equation 2.6 (see columns (2) and (4) in Table 2.7).

	12-Month Inflation Expectations		24-Month Inflation Expectations	
	(1)	(2)	(3)	(4)
β		-0.508***		-0.512***
		(0.105)		(0.115)
β^{NL}	-0.645***	-0.248	-0.539***	-0.138
	(0.164)	(0.183)	(0.199)	(0.219)
Obs	3754	3754	3754	3754
R^2	0.616	0.622	0.584	0.589
Num. Firms	293	293	293	293

Statistical significance: * 10%, ** 5%, *** 1%.
Standard errors clustered by firm.

Table 2.7: Heterogeneous Treatment Effect on Inflation Expectations

2.4.4 Other Variables of Interest

We have single cross-sections for the remaining variables of interest other than inflation and cost expectations. Therefore, we estimate the effect of treatment by estimating the following equation by OLS:

$$X_i = \alpha + \beta T_i + \varepsilon_i, \tag{2.7}$$

where X_i is each variable of interest.

The results are shown in Table 2.8. We can see that firms that receive the treatment adjust their growth forecast both for the current and the following year. The point estimate is sizable and is consistent with models of costly disinflation. The response to the question on plans to expand delivers a similar conclusion: the point estimate is surprisingly large although we cannot reject the null of zero with 95% confidence.

	2021 Growth Forecast	2022 Growth Forecast	Plans to Expand	ER 2021	ER 2022	Last Price Change (Months)	No Price Change
β	-0.621**	-0.261	-0.107*	0.551	0.899	0.357	0.036
	(0.272)	(0.251)	(0.056)	(0.606)	(0.610)	(0.364)	(0.035)
Obs	279	279	279	279	279	279	279
R^2	0.018	0.018	0.014	0.002	0.007	0.004	0.001

Statistical significance: * 10%, ** 5%, *** 1%.
Huber-Eicker-White standard errors.

Table 2.8: Effect of Treatment on Other Variables

Table 2.8 indicates that the treatment has no effect on expectations about the exchange rate although a considerable part of the treatment message involves foreign currency and the effect on other variables which should, at least in theory, affect the exchange rate. Not only is the point estimate small, but the confidence intervals indicate modest fluctuations in the exchange rate, at best. Last, the confidence intervals on the questions on price changes are too large to find statistically significant effects.

2.4.5 Cost Inflation Expectations

In addition to inflation expectations, the IES asks firms about their expectations about their own costs in the next 12 and 24 months. While this is not a measure of price inflation per se, firms have a strong incentive to be informed about cost changes that are more directly linked to their profits. Moreover, the correlation between cost and price inflation expectations is about 50% in our data, indicating both measures have common factors but are not quite the same. This section, therefore, explores whether the treatment has any effect on expected costs for the firms.

Table 2.9 shows the results when we estimate Equation 2.2 with the expected change in costs as the left-hand-side variable. The table shows that the effect on 12-month cost inflation expectation is almost -0.40 percentage points. The effect is larger in magnitude for the 24-month horizon, at -0.52 percentage points.

	12-Month Cost Inflation Exp.	24-Month Cost Inflation Exp.
β	-0.384**	-0.521***
	(0.184)	(0.194)
Obs	3265	3754
R^2	0.697	0.661
Num. Firms	293	293

Statistical significance: * 10%, ** 5%, *** 1%.
Standard errors clustered by firm.

Table 2.9: Treatment Effect on Cost Inflation Expectations

We also estimate the dynamic effect on cost expectations. Table 2.10 shows the estimates from Equation 2.3 with costs inflation expectations as the left-hand-side variable. Unfortunately, the estimates are no longer significant at the 5% level when we look at the incremental contribution of each month. However, the table does show the same qualitative pattern as Table 2.4. At the 12-month horizon, the treatment generates an initial negative effect, part of which is reversed in July, but that broadens in August. At the 24-month horizon, differences in cost inflation expectations grow across groups with time.

	12-Month Cost Inflation Exp.	24-Month Cost Inflation Exp.
β	-0.302	-0.368*
	(0.203)	(0.198)
β^{Jul}	0.114	-0.094
	(0.201)	(0.197)
β^{Aug}	-0.323	-0.290
	(0.204)	(0.218)
Obs	3265	3754
R^2	0.697	0.661
Num. Firms	293	293

Statistical significance: * 10%, ** 5%, *** 1%.
Standard errors clustered by firm.

Table 2.10: Dynamics Treatment Effect on Cost Inflation Expectations: Incremental Impact

The persistence of the cost expectations also follows a similar pattern to inflation expectations. Table 2.11 shows the results of estimating Equation 2.4 with the firms' expected change in costs as the left-hand-side variable. Although the level of statistical significance decreases through time, the effect of the treatment on firms' cost expectations seems to persist in the month following the treatment.

Estimated β^M for $M =$	12-Month Cost Inflation Exp.	2*	24-Month Cost Inflation Exp.
June	-0.384**		-0.521***
	(0.184)		(0.194)
July	-0.367**		-0.494**
	(0.184)		(0.202)
August	-0.426**		-0.543**
	(0.196)		(0.215)
Obs	3265		3754
Num. Firms	293		293

Statistical significance: * 10%, ** 5%, *** 1%.
Standard errors clustered by firm.

Table 2.11: Dynamics Treatment Effect on Cost Inflation Expectations: Persistence

2.5 Final remarks

This paper presents a novel experiment in which a randomly selected sub-sample of firms is informed about changes in monetary policy and the intention of the central bank to reduce trend inflation significantly in the following years. Our results show that in spite of being small, the treatment is sufficiently strong to affect inflation and own costs expectations. While the magnitude of the effect varies with time, the fall in expectations as a consequence of the treatment seems to persist until the end of the sample. Moreover, our results show that firms that are more attentive to inflation are more responsive to the treatment. These firms have lower inflation expectations and are more likely to postpone price changes as a result of treatment. To the best of our knowledge, this is the first evaluation of communication of this type of policy change.

Macroeconomic theory predicts ambiguous results as to whether there is a cost in terms of output to disinflation changes in monetary policy. Our results indicate that news of disinflation also makes firms adjust their growth expectations in addition to those of inflation. Unlike inflation and cost expectations which fell for both the 12 and 24-month horizons, treated firms anticipate temporarily slower growth only for the current year.

Although the message with which firms were treated emphasizes the beneficial effect of disinflation on demand for domestic currency, we find that there is no effect on exchange rate expectations. Even the confidence intervals suggest modest exchange rate variation at best. This is in spite of rich literature suggesting that monetary policy should be relevant to exchange rates.

There are a number of aspects that make our experiment novel and a series of questions that remain open. First, there is limited research on central bank communication in emerging markets, particularly in one characterized by relatively high inflation. Our results indicate that communication is effective even though firms have a stronger incentive to have knowledge of inflation than in advanced economies with low and stable inflation. Second, we show that the treatment lowers growth expectations in tandem with inflation expectations. This result contrasts with evidence in developed economies where firms seem to associate inflation with stagnation. Future research will need to assess whether these lower growth expectations impact firm decisions such as hiring and investment.

Overall, the empirical evidence provided in this paper highlights the importance of central bank communication to move firms' inflation expectations towards the central bank's objectives. Moreover, the causal effects of the news of disinflation persist over time with only a short-term impact on expected activity. Although to extrapolate these results to other circumstances is not direct, one could reasonably argue that similar results will hold in other economies, especially in those with long records of lower inflation and higher credibility of the central bank than in our case study.

Chapter 3

Unreliable Expectations

3.1 Introduction

Despite the rise in inflation following the reopening of economies from the COVID-19 pandemic, inflation expectations around the world were considerably well anchored around central bank targets around the world. Figure 3.1 shows how inflation in advanced economies and emerging markets started rising around mid-2020 to about 3 percentage points above central bank targets by the end of 2021. Figure 3.2 shows that this rise in inflation was accompanied by a modest rise in inflation expectations. In particular, a proxy for 12-month ahead inflation expectations is remarkably close to central bank targets even by the end of 2021.

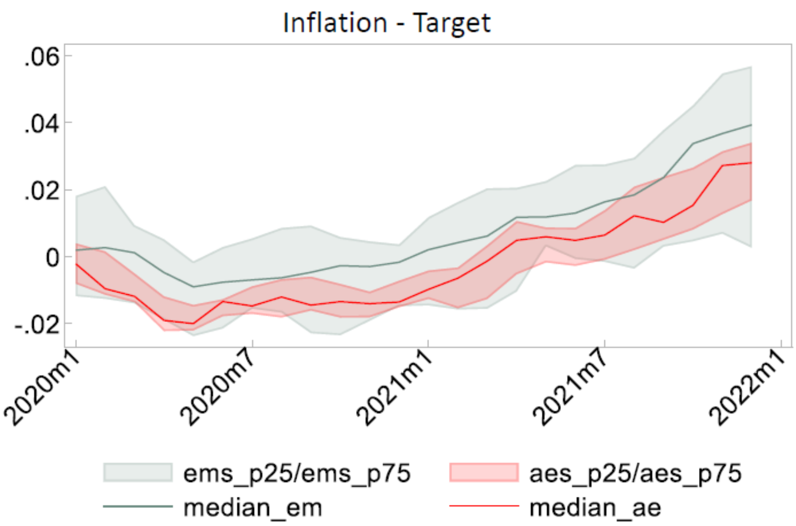

Figure 3.1: Inflation in Advanced Economies and Emerging Markets
Source: Adler, Caruso Bloeck, and Fratto (2022).

The fact that forward-looking measures of inflation expectations were close to targets motivated many central banks around the world to pursue an accommodative monetary policy stance. However, this outlook changed rather quickly in 2022 as inflation started rising more rapidly following the Russian invasion of Ukraine and also spread to the service sector. The swift change in scenarios around the world motivated the question of whether central banks should have acted earlier and been more proactive in leaning against inflationary pressures.

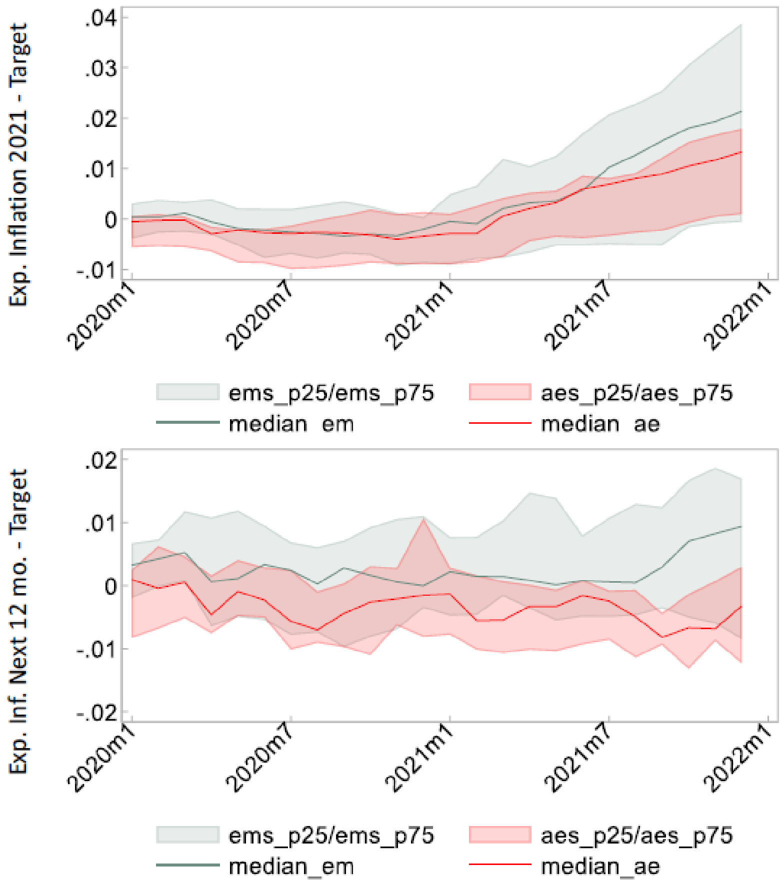

Figure 3.2: Inflation Expectations in Advanced Economies and Emerging Markets
Source: Adler, Caruso Bloeck, and Fratto (2022).

This paper argues that using private sector expectations to determine the monetary policy stance makes inflation dynamics inherently more unstable. To make this case, I first document a series of facts to assess the informational content of inflation expectations of professional forecasters in a panel of 36 advanced economies and emerging markets. These forecasts barely beat standard forecasting models such as random walks or the unconditional mean. Moreover, I document a generalized pattern of expectations lagging inflation.

To assess the implications of expectations-based monetary policy rules, I develop a simple model of expectation formation that can replicate these qualitative facts. I use this model to study the effect of inflationary shocks under counterfactual monetary policy rules. The key feature of the model is that differences in inflation dynamics and their implied monetary policy response determine agents' incentives to form accurate expectations.

The main difference between inflation-targeting (IT) rules and expectation-targeting (ET) rules is that the former can always guarantee anchoring of inflation expectations, while the latter generally cannot. The reason for this result is that these monetary policy rules differ in the implied monetary policy response given a shock. IT rules lean against inflationary shocks, reducing agents' incentives to acquire information about aggregate shocks. In contrast, ET rules cannot neutralize inflationary shocks unless agents learn about them, implying that some unanchoring of expectations is necessary to stabilize inflation.

These results are based on two main assumptions. The first is that agents form expectations more accurately when they have a greater incentive to do so. I model this friction in expectation formation in a way that resembles models with endogenous attention (e.g., Sims 2003, Maćkowiak & Wiederholt 2009, Gabaix 2014). In this framework, agents do not pay attention if the payoff from doing so is sufficiently low, resulting in anchoring in inflation expectations.

The second assumption is that agents understand the partial equilibrium effect of shocks and form expectations about their general equilibrium implications. This assumption resembles models in which agents' capacity to iterate over other agents' responses is limited (e.g., Farhi & Werning 2019), with the added feature that the depth of iterations is endogenous in my framework. A byproduct of this assumption is that it yields an easily interpretable condition under which IT frameworks anchor expectations, which is that the monetary policy rule has to offset the second-order effects in price-setting. Other frictions such as imperfect common knowledge may require monetary policy responses that are arbitrarily sensible to deviations from the target to guarantee anchoring.

This paper contributes to a rich literature analyzing the interplay between monetary policy and the formation of expectations (Maćkowiak & Wiederholt 2015, Gabaix 2020). In particular, a series of papers in this literature document biases in expectation formation and incorporate these biases into macroeconomic models to conduct counterfactuals (Ryngaert 2018, Sung 2022, Dietrich 2023). To the best of my knowledge, counterfactuals in which monetary policy is based on inflation expectations, rather than inflation *per se*, have not been an object of analysis. The results in this paper indicate that such counterfactual monetary regimes have qualitatively different implications for inflation and expectations dynamics.

My results resonate with those of Bernanke & Woodford (1997) who show that ET regimes can lead to indeterminacy in New-Keynesian models. By comparison, my model resolves the indeterminacy problem given that a sufficient level of volatility causes agents to learn about aggregate shocks and this process of expectation formation induces a stabilizing monetary policy response. However, IT and ET frameworks in my model differ in that the former does not require this expectational response to stabilize (optimal) prices.

The remainder of this paper is structured as follows. Section 3.2 shows the empirical evidence on the predictive power of inflation expectations. Section 3.3 develops the model and extensions. Finally, section 3.4 offers the conclusions and final remarks.

3.2 Predictive Power of Professional Inflation Forecasts

This section reviews the evidence on the predictive power of professional inflation forecasts. The goal of this section is to evaluate the usefulness of professional forecasts for predicting inflation. This section is a review of the evidence presented in Adler et al. (2022).

3.2.1 Predictive Accuracy of Professional Forecasts

Figure 3.3 shows the distribution of current-year forecast errors of professionals and that of simple forecasting models, namely a random walk and the sample unconditional mean. The top panels show that professional forecasts outperform unconditional means in both advanced economies and emerging markets. The lower panels show that they also outperform random walks, although the difference is smaller, particularly in advanced economies.

The fact that professional forecasts outperform these simple benchmarks should not be surprising. Current-year forecasts include several months of observed inflation data. This implies forecasters often know with certainty a sizable component of the object being forecasted. With this in mind, the degree of outperformance relative to random walks seems small, particularly in advanced economies.

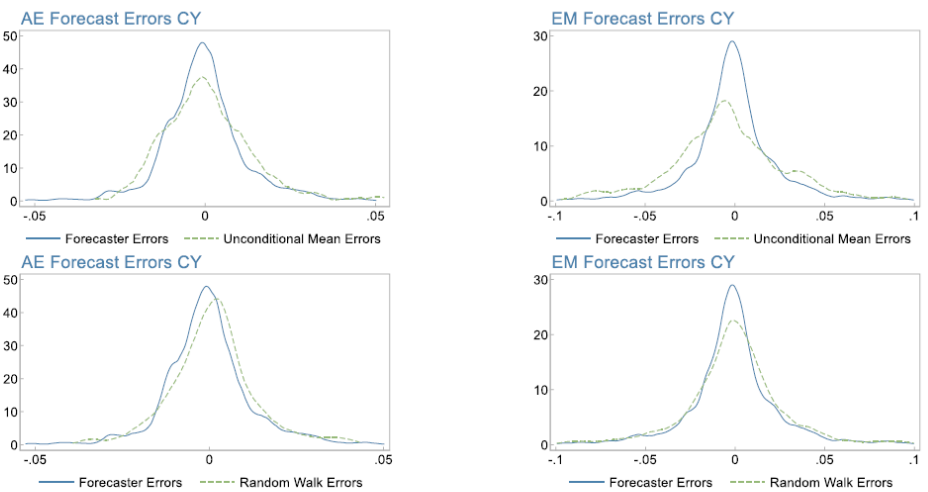

Figure 3.3: Forecast Errors, Current Year
Source: Adler, Caruso Bloeck, and Fratto (2022).

Figure 3.4 compares professional next-year forecast errors with those of the same simple forecasting models. It is worth noting that next year forecast errors should be significantly more relevant for the purpose of determining the stance of monetary policy, given the lags with which the latter typically operates. The figure shows that professional forecast errors are comparable to those of the unconditional mean in both advanced economies and emerging markets. This implies that the professional forecasting ability is mostly limited to very short time horizons.

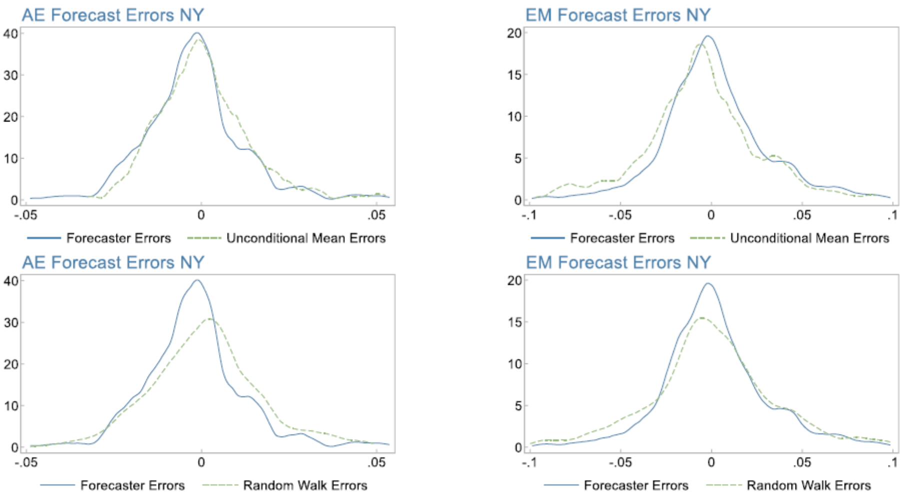

Figure 3.4: Forecast Errors, Current Year
Source: Adler, Caruso Bloeck, and Fratto (2022).

The fact that professional forecasts do not markedly outperform these simple models beyond the current year is consistent with the evidence on forecasting inflation. Namely, inflation is commonly viewed as a comparatively hard macroeconomic aggregate to forecast.

Figures 3.3 and 3.4 indicate that professional forecasts are unlikely to reflect significant amounts of additional information that central banks could not infer from other data sources.

3.2.2 Inflation and Forecast Errors

While the previous subsection evaluated the average predictive accuracy of professional forecasts, this section is dedicated to whether professional forecasts can identify changes in inflation trends. This requires studying the time-series properties of professional forecasts given changes in inflation.

I begin by estimating the association between changes in inflation and forecast errors, which is a variant of the specification proposed in Kohlhas & Walther (2021). Specifically, the panel specification that I consider is as follows:

$$\pi_{c,\tau} - \mathbb{E}_{c,t+\gamma}[\pi_{c,\tau}] = \alpha_c + \beta_\gamma(\pi_{c,t} - \pi_{c,t-\phi}) + \epsilon_{c,t} \tag{3.1}$$

where c indexes countries, τ is an index for a forecast period, $t+\gamma$ is the period in which forecasts are made, and $\pi_{c,t} - \pi_{c,t-\phi}$ is the change in inflation between $t-\phi$ and t. The results of this equation show the association between recent changes in inflation in the last ϕ months and forecast errors made γ months into the future. The coefficients of interest are the values β_γ. A positive value of $\beta_\gamma = 0$ indicates that changes in inflation are not associated with forecast errors γ months into the future, implying that forecasters have incorporated this information into their forecasts. However, a value of $\beta_\gamma > 0$ implies that forecasters under-react to the acceleration in inflation while a value of $\beta_\gamma < 0$ indicates over-reaction.

The results are displayed in Figure 3.5. Both graphs show an initial pattern of under-reaction. It is noteworthy that the graphs show under-reaction even after several months following a change in inflation, by which time these changes in inflation are public information. Moreover, Figure 3.6 shows these patterns are generally true in individual countries, although using smaller samples reduces the precision in the estimated β_γ coefficients.

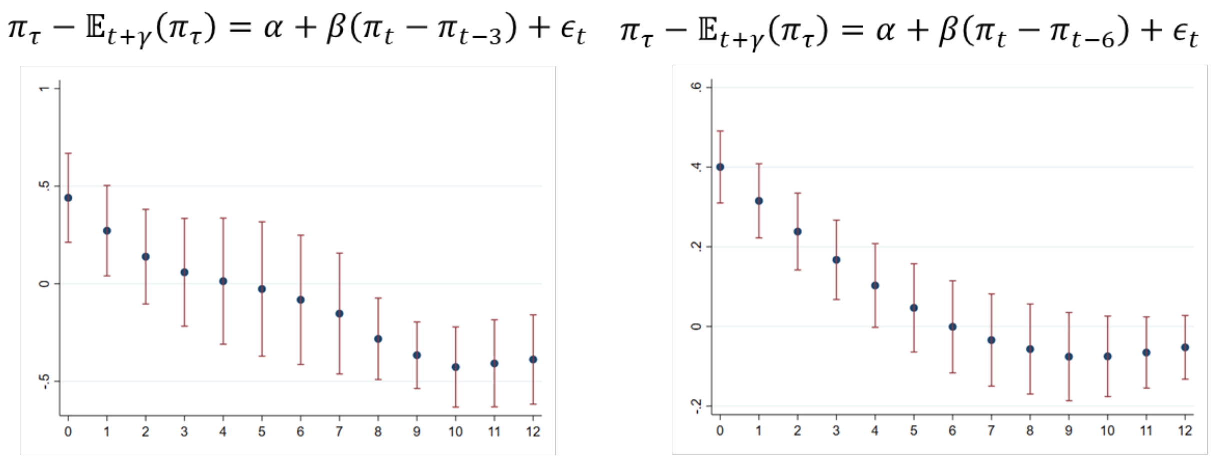

Figure 3.5: Forecast Errors and Inflation Changes
Source: Adler, Caruso Bloeck, and Fratto (2022).

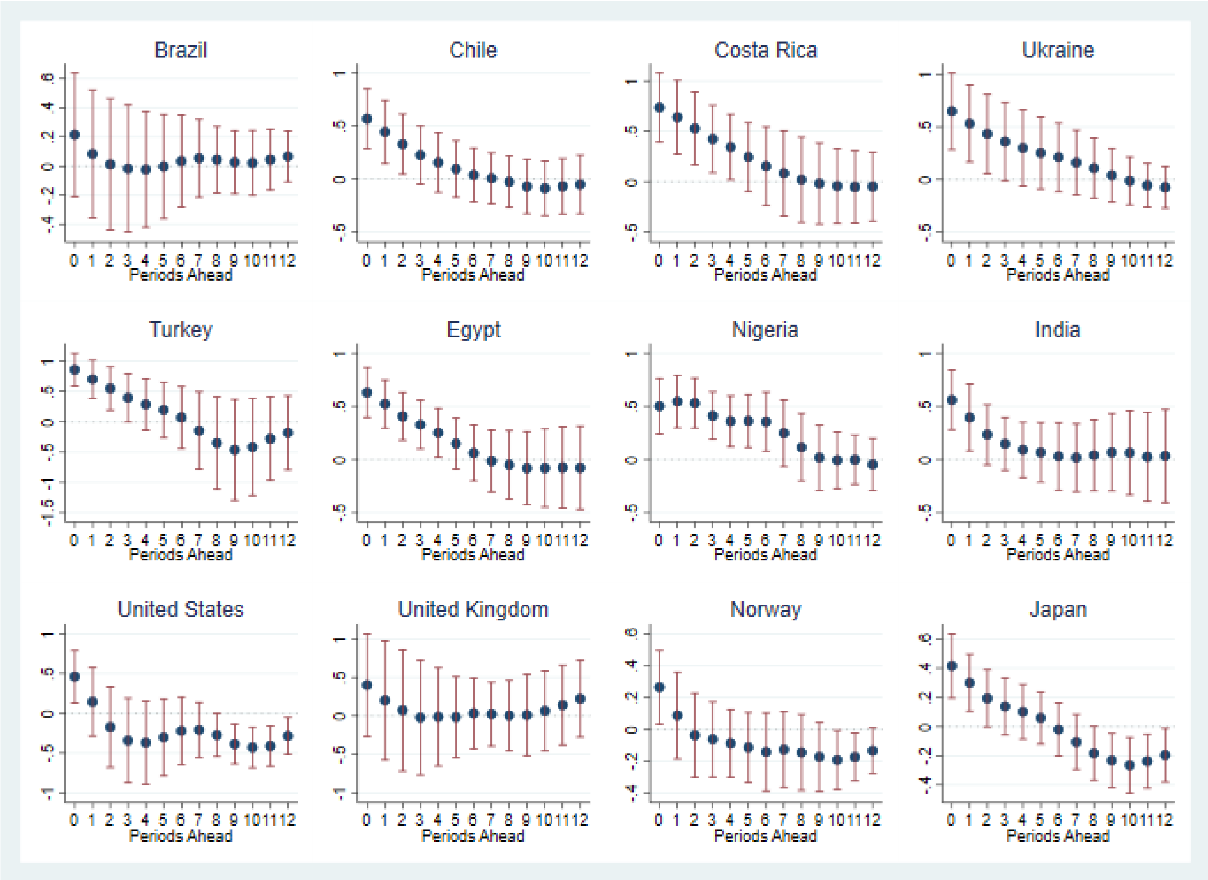

Figure 3.6: Forecast Errors and Inflation Changes, Individual Countries
Source: Adler, Caruso Bloeck, and Fratto (2022).

Furthermore, I present evidence in the form of the Coibion & Gorodnichenko (2015*a*) regression which estimates the relation between forecast revisions and forecast errors. The results are featured in Figure 3.7. The graph shows the regression coefficients given 3- and 6-month revisions. Both graphs show an initial under-reaction and a subsequent over-reaction of expectations. The result on under-reaction is consistent with that displayed in Figure 3.5 while the evidence on the subsequent over-reaction is stronger.

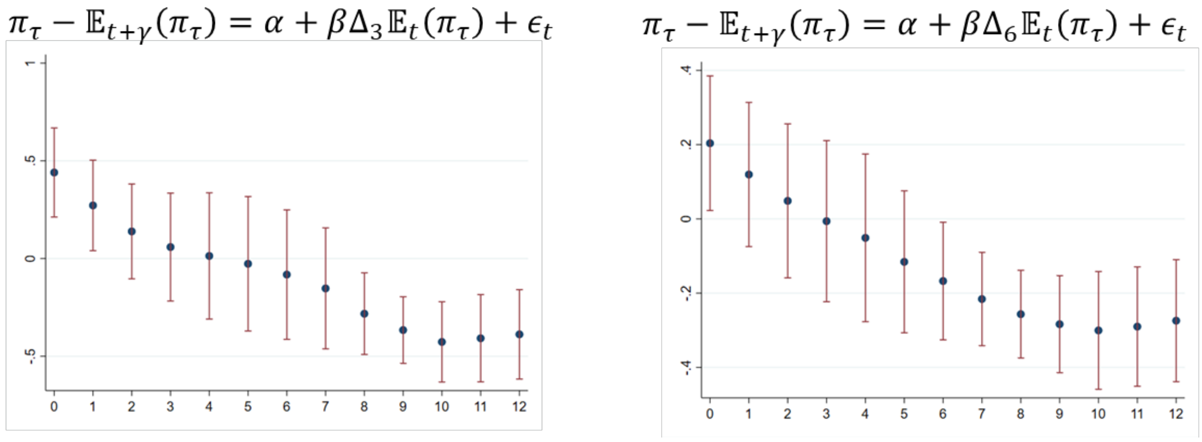

Figure 3.7: Forecast Revisions and Forecast Errors
Source: Adler, Caruso Bloeck, and Fratto (2022).

3.2.3 Summary of Empirical Evidence

The empirical results in this section indicate that inflation expectations, even when formulated by professional forecasters, do not contain particularly useful information about inflation trends. Private forecasts do not markedly outperform simple models, and lag changes in inflation for several months.

What are the implications of this empirical evidence for monetary policy? If the process of learning about inflation dynamics is exogenous, monetary policy guided by private forecasts may exhibit the same delayed response as private inflation expectations. By itself, this process need not generate deanchoring of inflation expectations. As long as private agents understand that monetary policy will counter inflationary pressures, they should build said monetary policy response into their forecasts as they learn about inflationary shocks.

Alternative expectation formation processes may have different implications for inflation expectations and dynamics. In expectation formation processes that have extrapolative features, such as adaptive expectations or diagnostic expectations, inflation surprises may lead agents to overestimate the persistence of inflation. In these frameworks, monetary policy affects the expectation formation process by affecting the size of inflation surprises. However, this is broadly inconsistent with the evidence on aggregate under-reaction shown earlier. Moreover, it is not obvious that expectational monetary policy cannot replicate standard Taylor rules.

The next section considers a model of endogenous information acquisition in the form of rational inattention. In contrast with the previous frameworks, in a rational inattention model monetary policy affects the expectation formation process by ruling the benefits of acquiring information. The model shows that different Taylor rules have different implications for the incentives to acquire information, and that expectational monetary policy rules cannot generally replicate a Taylor rule.

3.3 Monetary Policy Rules with Endogenous Expectations

This section develops a simple rational inattention framework which is used to study how expectation formation depends on the monetary policy rule. The main exercise is to compare a standard Taylor rule against one in which monetary policy is set based on private agents' inflation expectations.

3.3.1 Model

Setup. The setup of the model follows the monetary model of Angeletos & Lian (2016). The model features consumer-producer households and a central bank. Households are subject to shocks that raise their demand for goods and the central bank sets the money supply in such a way that it can stabilize nominal aggregate demand. This section describes the main equations that determine agents' decision problems.

Households are indexed by $i \in (0, 1)$ and decide the nominal price of their varieties p_i. Following Maćkowiak & Wiederholt (2009), it is straightforward to express the loss function as:

$$U_i \approx \bar{U} - b\mathbb{E}_i(p_i - p_i^*)^2 \tag{3.2}$$

where p_i^* is the optimal price that the consumer-producer would set if she had full information, given the information of other consumer-producers. Given expression 3.2, firms will set prices that are equal to their expected optimal price:

$$p_i = \mathbb{E}_i[p_i^*]$$

The optimal price for the firm is:

$$p_i^* = \alpha g_i + (1 - \alpha)\bar{p} - r$$

where g_i is a measure of demand for product i, $\bar{p}$ is the average price set across firms, and r is the monetary policy rate. The first term αg_i can be thought of as capturing partial equilibrium effects. The second and third terms capture general equilibrium effects due to strategic interaction amongst firms and the monetary policy response, respectively.

I assume the level of demand g_i is known to households and is an exogenous random shock. Moreover, I assume g_i is the sum of a common shock g, and household-specific noise η_i, both of which are drawn from zero-mean normal distributions with variances σ_g^2 and σ_i^2, respectively. Hence, we have:

$$g_i = g + \eta_i$$

I will consider two types of monetary policy rules. The first is an inflation target rule r^{IT} and the second one is an inflation-expectation target rule r^{ET}. Hence, I will write:

$$\begin{aligned} r^{IT} &= \beta^{IT} p \\ r^{ET} &= \beta^{ET} \bar{\mathbb{E}}(p) \end{aligned}$$

where $\bar{\mathbb{E}}(p)$ is the average expected inflation across individuals.

Expectation Formation. The last step is to determine how households form expectations. Households observe the partial equilibrium effect g_i and must form expectations about the general equilibrium components $p^*_{GE} \equiv (1-\alpha)\bar{p} - r$. In order to adopt a flexible framework, I assume that expectation formation has the following property:

$$\mathbb{E}_i[p^*_{GE}] = \mu p^*_{GE} + (1-\mu)\eta^s_i \tag{3.3}$$

where $\mu \in [0,1]$ and η^s_i is a zero-mean idiosyncratic signal. The vital property of this expectation formation process is that aggregate expectations are biased toward and *can be* strictly zero, as can be seen in the following expression:

$$\bar{\mathbb{E}}[p^*_{GE}] = \mu p^*_{GE} \tag{3.4}$$

We could certainly entertain more general processes for expectation formation. For example, we could assume that the first term of Equation (3.3) has additional noise, that the weights in the two factors need not add up to one, or that the idiosyncratic signal η^s_i is correlated with the idiosyncratic demand shock η_i. As long as property 3.4 is satisfied, the main result regarding the feasibility that different monetary policy rules anchor expectations will follow.

Property 3.4 is consistent with the empirical evidence documented previously on initial under-reaction of expectations. However, several frameworks of information rigidity do not satisfy the property that μ can be strictly zero, even if they generate under-reaction. For example, frameworks such as noisy information where expectations about p^*_{GE} are formed costlessly with the observed shock g_i have $\mu > 0$ in Equation 3.4. The possibility that $\mu = 0$ can be associated with rigidities in understanding the general equilibrium interaction across players (e.g., as in Farhi & Werning 2019). The possibility that $\mu = 0$ is important because it implies that expectations can be strictly anchored while inflation differs from target, in which case the two monetary policy rules respond differently.

The value μ is an endogenous object. Specifically, a higher value of μ decreases households' tracking errors and raises expected utility (3.2). However, obtaining a higher value of μ is costly. I use a flexible cost function $\mathcal{C}(\mu)$ that resembles Gabaix (2014) that is differentiable with the following properties:

$$\mathcal{C}'(0) = \kappa > 0$$
$$\mathcal{C}''(\bullet) > 0$$
$$\lim_{\mu \to 1} \mathcal{C}'(\mu) = +\infty$$

The first property is the one that generates anchoring of expectations. Specifically, because the marginal cost of paying attention to the aggregate shock is positive at zero, households will only pay attention if the marginal benefit is above the initial cost κ. The remaining conditions ensure the solution to the problem is well-behaved.

While I will use a specific function for the cost of attention, any function that satisfies these conditions will generate qualitative similar results. Namely, I will use the following cost function which resembles the rational inattention literature (e.g., Maćkowiak et al. 2023):

$$\mathcal{C}(\mu) = \frac{\Gamma}{2}\log_2\left(\frac{1}{1-\mu}\right)$$

Finally, households pick the value μ that minimizes the following expression:

$$b\mathbb{E}_i(p_i - p_i^*)^2 + \mathcal{C}(\mu) \tag{3.5}$$

Solution. A solution to the rational inattention problem can be obtained following the method of undetermined coefficients. Namely, I conjecture that the equilibrium price level has the following form:

$$\bar{p} = \gamma g$$

The equilibrium is characterized by a fixed point between the conjecture and the formation of expectations given that the conjecture is true. The result is expressed in the following lemma.

Lemma 3.3.1 (Equilibrium conditions). *The equilibrium is characterized by values $\{\gamma, \lambda^R\}$, ruling the effect exogenous shocks on aggregate prices and the formation of expectations, respectively, that satisfy the following conditions:*

$$\gamma = \alpha + \rho^R\lambda^R \tag{3.6}$$

$$\lambda^R = \max\left\{0, 1 - \frac{\Gamma}{\ln 2}\frac{1}{(\rho^R)^2\sigma_g^2}\right\} \tag{3.7}$$

where $R \in \{IT, ET\}$, and $\rho^{IT} = (1-\alpha)\gamma - \beta\gamma$, $\rho^{ET} = (1-\alpha)\gamma - \lambda^{ET}\beta\gamma$.

The following section analyzes the equilibrium of the model in the two regimes and carries out comparative static exercises.

3.3.2 Equilibrium Analysis

Equilibrium Equations

I begin the analysis by plotting Equations (3.6) and (3.7). Figure 3.8 shows these equations for an inflation-targeting monetary policy rule. The graph shows that the value λ in Equation (3.7) is zero until some cutoff value of γ, and is increasing and concave for higher values of γ. The graph of Equation (3.6) shows a negative relation between λ and γ, which ensures that the equilibrium is unique. This requires the parameter β^{IT} of the reaction function to be larger than $1-\alpha$. If β^{IT} is lower than $1-\alpha$, the graph of Equation (3.6) would be upward-sloping, and multiple equilibria may be possible depending on the relative curvature of the curves.

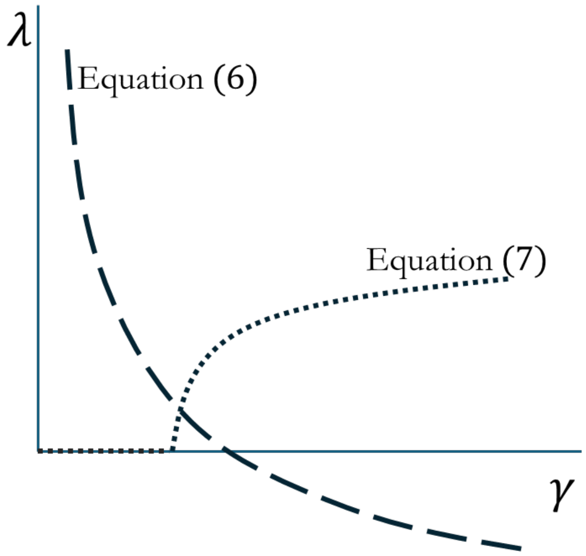

Figure 3.8: Equilibrium Equations with Inflation Targeting Rule

Intuitively, the relation between γ and λ in Equation (3.6) depends on the strength of monetary policy response. A strong monetary policy response to inflation decreases the desired passthrough of aggregate shocks, and the latter is lower when firms are more attentive to aggregate shocks. This explains the negative slope in Figure 3.8. Conversely, a weaker monetary policy response raises the desired passthrough to an aggregate shock, which would make the relation between γ and λ positive.

Figure 3.9 contains the graph for an expectation-targeting regime. The graph of Equation (3.7) looks qualitatively similar to the inflation targeting case. The main difference, however, is the graph of Equation (3.6). The graph shows that lowering λ from a high value initially increases γ. As argued in the previous paragraph, this reflects the fact that a strong monetary policy response generates a negative relation between γ and λ. However, the graph shows that for a sufficiently low level of λ, the relation between the two variables becomes positive. As explained previously, this reflects a weak monetary policy response. Intuitively, this happens because low attention weakens the monetary policy response, which requires agents to pay greater attention.

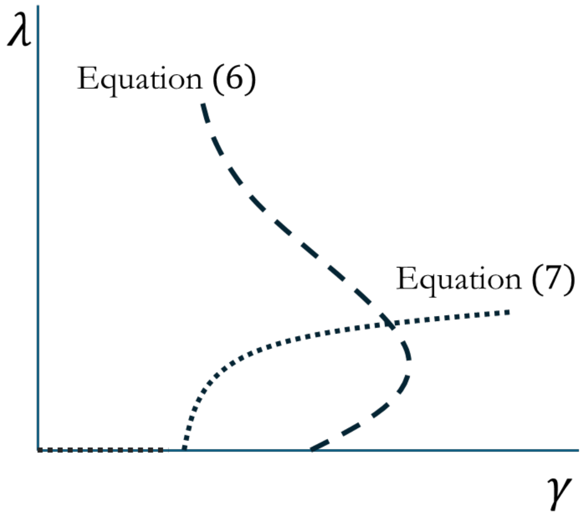

Figure 3.9: Equilibrium Equations with Expectation Targeting Rule

The Effect of Monetary Policy Rules

This section analyzes the effect of changing the monetary rule parameters β^{IT}, β^{ET}. This is the main comparative static of interest and the goal is to study qualitative differences in the inflation and the expectation formation process.

Figure 3.10 shows the effect of raising β^{IT} in the inflation targeting regime. The rise in the reaction function coefficient causes the γ schedule to pivot around the horizontal intercept. In addition, the effect on the λ schedule is to raise the second element of the maximum operation. Intuitively, starting from a value of $\beta^{IT} > 1 - \alpha$, raising β^{IT} increases the volatility of the aggregate component of the optimal price and hence the payoff of accurate expectations.

In the graph, the effect of the rise in β^{IT} is to move the equilibrium from the intersection of the solid lines to the intersection of the dashed lines. In the graph, the change in equilibrium is such that inflation expectations become unanchored given that the equilibrium value of λ increases from zero to a positive value. However, anchoring expectations is always possible given that Equation (3.7) converges to the horizontal axis and Equation (3.6) becomes vertical as β^{IT} converges to $1 - \alpha$. Hence, this setting requires the monetary policy rule to offset the general equilibrium effects alone to ensure anchoring of expectations, which contrasts with other settings in which β^{IT} may need to be arbitrarily high.

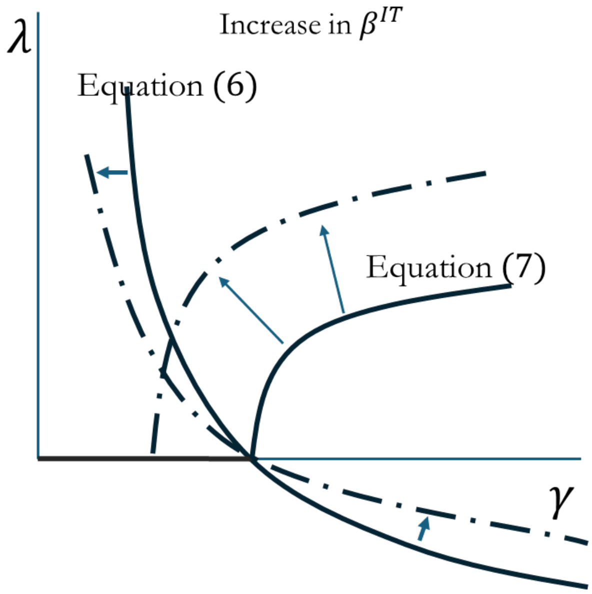

Figure 3.10: The Effect of Changes in β^{IT}

Figure 3.11 shows how raising the coefficient β^{ET} affects both curves in the expectations-targeting regime. The most important feature of this graph is that the horizontal intercepts of both curves are unchanged by the reaction function coefficient. Intuitively, this is because monetary policy does not react to inflationary shocks around $\lambda = 0$. Furthermore, this implies that the equilibrium is either one in which expectations are anchored or one where expectations are not anchored, regardless of the monetary policy rule.

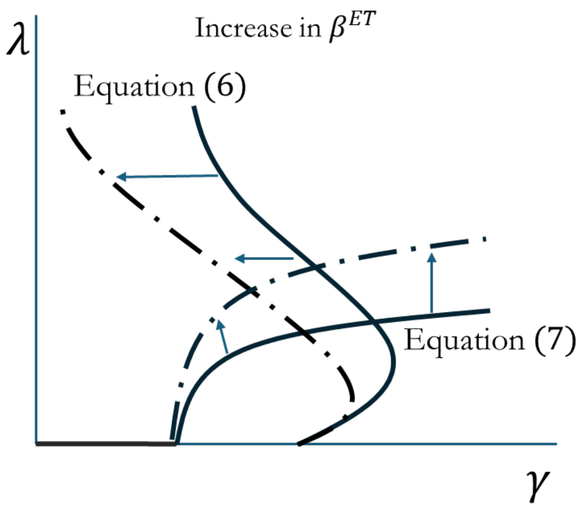

Figure 3.11: The Effect of Changes in β^{ET}

3.3.3 Inflation Dynamics

I now extend the model to a dynamic environment to analyze the implications of the model for inflation and expectations dynamics. I will develop a partial equilibrium setup in the sense that I will take the reduced-form expressions of the firm pricing problem as given without fully specifying the demand side of the economy.

Model Setup

Firms. I assume time is infinite and indexed by t and firms minimize the present discounted value of the expected quadratic deviations from an optimal price.

$$U_{it} = \mathbb{E}_{it}\left[\sum_{s=0}^{\infty}\beta^s(p_{it+s} - p^*_{it+s})^2\right] \tag{3.8}$$

As before, optimal prices will be given by:

$$p^*_{it} = \alpha g_{it} + (1-\alpha)\bar{p}_t - r_t$$

where g_{it} is a nominal expenditure shock observed by the firm and it is the sum of an aggregate shock and an idiosyncratic shock $g_{it} = g_t + \eta_{it}$. I assume the common component of g_{it} is a random walk, so that nominal demand shocks have a level effect on prices in the long run:

$$g_t = g_{t-1} + \eta_t$$

In this dynamic environment, policy rules will be a function of inflation and inflation expectations (and not price levels). Because of that, it will be more convenient to define inflation as $\pi_t \equiv \bar{p}_t - \bar{p}_{t-1}$ and to write optimal prices as:

$$p^*_{it} = \alpha g_{it} + (1-\alpha)\bar{p}_{t-1} + (1-\alpha)\pi_t - r_t$$

Firms set prices flexibly, which are optimally set equal to their expected optimal price, $p_{it} = \mathbb{E}_{it}[p^*_{it}]$.

Furthermore, I will assume that firms observe past price levels and therefore know past aggregate shocks. This means firms only face uncertainty about shocks in the current period and they are subject to the same frictions as before. From a modeling perspective, this assumption implies that the variance of prediction errors is finite even if firms do not formulate general equilibrium expectations. If firms did not observe past price levels, firms would eventually have to pay attention given that the price level is non-stationary and the variance in prediction errors is infinite without attention.

As in Section 3.3.1, I define $\bar{p}^*_{GE,t} \equiv (1-\alpha)\pi_t - r_t$ which is the general equilibrium part of the optimal price that is unknown to the firm. Firm form expectations equal to:

$$\mathbb{E}_{it}[p^*_{GE,t}] = \mu p^*_{GE,t} + (1-\mu)\eta_{it} \tag{3.9}$$

Finally, firms minimize Equation (3.8) minus the cost of forming expectations which is the same as that of Section 3.3.1:

$$\mathcal{C}(\mu) = \frac{\Gamma}{2}\log_2\left(\frac{1}{1-\mu}\right)$$

Monetary Policy. I will consider two monetary policy rules which correspond to inflation targeting and expectation targeting. The rules are the following:

$$\begin{aligned} r_t^{IT} &= \beta^{IT}\pi_t \\ r_t^{ET} &= \beta^{ET}\bar{\mathbb{E}}_t[\pi_t] \end{aligned}$$

For simplicity, I will also assume that the monetary policy coefficients β would fully offset inflation under full information. This happens if $\beta^{IT} = \beta^{ET} = 1-\alpha$. This assumption is not essential, but it simplifies the algebra given that agents face uncertainty only about current shocks.

Equilibrium. To find the equilibrium, let us conjecture that inflation is proportional to the aggregate nominal expenditure shock:

$$\pi_t = \sum_{l=0}^{\infty}\gamma_l g_{t-l}$$

The solution is found and analyzed in the following subsection.

Equilibrium in the Dynamic Model

The following lemma expresses the solution to the coefficients in recursive form:

Lemma 3.3.2 (Dynamic solution)**.** *The coefficient γ_0 satisfies the following conditions:*

$$\gamma_0 = \alpha + \rho^R\lambda^R \tag{3.10}$$

$$\lambda^R = \max\left\{0, 1 - \frac{\Gamma}{\ln 2}\frac{1}{(\rho^R)^2\sigma_g^2}\right\} \tag{3.11}$$

where $R \in \{IT, ET\}$, and $\rho^{IT} = (1-\alpha)\gamma_0 - \beta\gamma_0$, $\rho^{ET} = (1-\alpha)\gamma_0 - \lambda^{ET}\beta\gamma_0$.

Furthermore, for $l \geq 1$, we have that:

$$\gamma_l = 1 - \alpha^l + \alpha^l\gamma_0 \tag{3.12}$$

The first part of Lemma 3.3.2 states that the solution for the passthrough on impact is analogous to that of the partial equilibrium model. Equation (3.12) shows the long-run passthrough converges to unity. Intuitively, the reason for that is that it was assumed that shocks are permanent, and monetary policy only offsets current inflation.

Figure 3.12 shows the impulse response of a nominal demand shock. The graph shows that inflation arises and falls gradually in an inflation targeting regime in which expectations are initially anchored. In the expectations-targeting regime, firms learn about the shock on impact which generates a stronger initial passthrough.

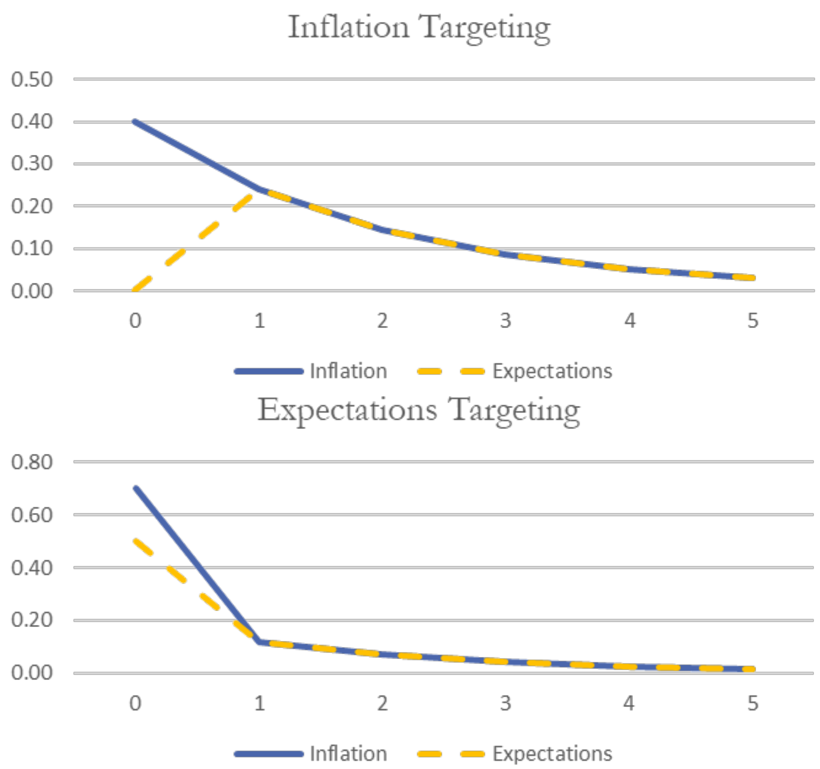

Figure 3.12: Impulse Response to Nominal Demand Shock

Schematically, the dynamic version of the model shows how inflation dynamics differ in the two regimes. Despite inflation shocks quickly becoming common knowledge after a single period, the initial difference in incentives to acquire information affects short-run passthrough significantly. Given that long-run passthrough is the same in both regimes, the main difference is that a given inflationary shock is more front-loaded in the expectations targeting regime.

3.4 Concluding Remarks

This short paper highlights the risks associated with the use of conventional measures of inflation expectations. There is a large and growing body of evidence highlighting that these measures are systematically biased and incorporate inflationary shocks slowly.

This paper presents a model in which this process of expectation formation is endogenous to monetary policy. Gradual learning about inflation reflects the effectiveness of monetary policy in offsetting these shocks. However, counterfactual monetary policy rules which do not sufficiently counterbalance inflationary shocks can cause unachoring of inflation expectations.

The main result of the paper is that monetary policy rules based on agents' expectations cannot generally guarantee anchoring of expectations. This is because these rules require unachoring in the first place to generate a stabilizing monetary policy response. Furthermore, whether expectations are anchored or not is a feature of the monetary regime and therefore independent of the reaction function of the central bank.

These results imply that the pass-through of shocks to prices can depend sensibly on the process of expectation formation. When expectations are unanchored, passthrough is usually quick and inflation is more volatile, while anchored expectations generate gradual inflation.

These results indicate the promise of pursuing greater empirical research on the formation of expectations under different inflationary settings and monetary regimes. These efforts will be fruitful to understand how expectations are formed out of sample.

Bibliography

Adler, G., Caruso Bloeck, M. & Fratto, C. (2022), The cost of missing the inflation target, IMF mimeo.

Amador, M. & Weill, P.-O. (2010), 'Learning from prices: Public communication and welfare', *Journal of Political Economy* **118**(5), 866–907.

Andrade, P. & Le Bihan, H. (2013), 'Inattentive professional forecasters', *Journal of Monetary Economics* **60**(8), 967–982.

Angeletos, G.-M. & La'O, J. (2008), 'Dispersed information over the business cycle: optimal fiscal and monetary policy'.

Angeletos, G.-M. & La'O, J. (2010), 'Noisy business cycles', *NBER Macroeconomics Annual* **24**(1), 319–378.

Angeletos, G.-M. & Lian, C. (2016), Incomplete information in macroeconomics: Accommodating frictions in coordination, *in* 'Handbook of macroeconomics', Vol. 2, Elsevier, pp. 1065–1240.

Angeletos, G.-M. & Lian, C. (2018), 'Forward guidance without common knowledge', *American Economic Review* **108**(9), 2477–2512.

Ball, L. (1994*a*), 'Credible disinflation with staggered price-setting', *The American Economic Review* **84**(1), 282–289.

Ball, L. (1994*b*), What determines the sacrifice ratio?, *in* 'Monetary policy', The University of Chicago Press, pp. 155–193.

Ball, L. M. & Romer, D. (2003), 'Inflation and the informativeness of prices', *Journal of Money, Credit, and Banking* **35**(2), 177–196.

Benabou, R. (1988), 'Search, price setting and inflation', *The Review of Economic Studies* **55**(3), 353–376.

Benabou, R. & Gertner, R. (1993), 'Search with learning from prices: does increased inflationary uncertainty lead to higher markups?', *The Review of Economic Studies* **60**(1), 69–93.

Bernanke, B. (2007), Inflation expectations and inflation forecasting, Technical report, Board of Governors of the Federal Reserve System (US).

Bernanke, B. S. & Woodford, M. (1997), 'Inflation forecasts and monetary policy', *Journal of Money, Credit, and Banking* pp. 653–684.

Binder, C. C. (2018), 'Inflation expectations and the price at the pump', *Journal of Macroeconomics* **58**, 1–18.

Binder, C. & Rodrigue, A. (2018), 'Household informedness and long-run inflation expectations: Experimental evidence', *Southern Economic Journal* **85**(2), 580–598.

Borraz, F. & Mello, M. (2020), Communication, information and inflation expectations, Technical report, Banco Central del Uruguay.

Borraz, F. & Zacheo, L. (2018), 'Inattention, disagreement and internal (in) consistency of inflation forecasts', *Banco Central del Uruguay Working Papers* **7-2018**.

Boz, E., Casas, C., Georgiadis, G., Gopinath, G., Le Mezo, H., Mehl, A. & Nguyen, T. (2022), 'Patterns of invoicing currency in global trade: New evidence', *Journal of International Economics* **136**, 103604.

Buiter, W. H. & Miller, M. H. (1983), Costs and benefits of an anti-inflationary policy: questions and issues, Technical report, National Bureau of Economic Research.

Burdett, K. & Judd, K. L. (1983), 'Equilibrium price dispersion', *Econometrica: Journal of the Econometric Society* pp. 955–969.

Candia, B., Coibion, O. & Gorodnichenko, Y. (2021), 'The inflation expectations of U.S. firms: Evidence from a new survey', *NBER working paper 28836* .

Candia, B., Coibion, O. & Gorodnichenko, Y. (2023), The macroeconomic expectations of firms, *in* 'Handbook of Economic Expectations', Elsevier, pp. 321–353.

Carroll, C. D. (2003), 'Macroeconomic expectations of households and professional forecasters', *The Quarterly Journal of Economics* **1**(118), 269–298.

Caruso Bloeck, M. (2024), 'Informational cost and allocative inefficiency of inflation uncertainty', *Available at SSRN* .

Caruso Bloeck, M., Mello, M. & Ponce, J. (2023), 'News of disinflation and firms' expectations: New causal evidence', *Journal of International Money and Finance* **137**, 102914.

Coeurdacier, N., Rey, H. & Winant, P. (2011), 'The risky steady state', *American Economic Review* **101**(3), 398–401.

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

Coibion, O. & Gorodnichenko, Y. (2015*b*), 'Is the phillips curve alive and well after all? inflation expectations and the missing disinflation', *American Economic Journal: Macroeconomics* **7**(1), 197–232.

Coibion, O., Gorodnichenko, Y., Knotek II, E. S. & Schoenle, R. (2020), Average inflation targeting and household expectations, Technical report, National Bureau of Economic Research.

Coibion, O., Gorodnichenko, Y. & Kumar, S. (2018), 'How do firms form their expectations? new survey evidence', *American Economic Review* **108**(9), 2671–2713.

Coibion, O., Gorodnichenko, Y., Kumar, S. & Pedemonte, M. (2020), 'Inflation expectations as a policy tool?', *Journal of International Economics* **124**, 103297.

Coibion, O., Gorodnichenko, Y. & Ropele, T. (2020), 'Inflation expectations and firm decisions: New causal evidence', *The Quarterly Journal of Economics* **135**(1), 165–219.

Coibion, O., Gorodnichenko, Y. & Weber, M. (2022*a*), 'Monetary policy communications and their effects on household inflation expectations', *Journal of Political Economy* **6**(130), 1537–1584.

Coibion, O., Gorodnichenko, Y. & Weber, M. (2022*b*), 'Monetary policy communications and their effects on household inflation expectations', *Journal of Political Economy* **130**(6), 000–000.

Dietrich, A. M. (2023), 'Consumption categories, household attention, and inflation expectations: Implications for optimal monetary policy'.

Drenik, A. & Perez, D. J. (2020), 'Price setting under uncertainty about inflation', *Journal of Monetary Economics* **116**, 23–38.

D'Acunto, F., Malmendier, U., Ospina, J. & Weber, M. (2021), 'Exposure to grocery prices and inflation expectations', *Journal of Political Economy* **129**(5), 1615–1639.

Easterly, W. (1996), 'When is stabilization expansionary? evidence from high inflation', *Economic policy* **11**(22), 65–107.

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

Fischer, S. (1981), 'Relative shocks, relative price variability, and inflation', *Brookings Papers on Economic Activity* **1981**(2), 381–441.

Fischer, S. (1986), 'Indexing, inflation, and economic policy', *MIT Press Books* **1**.

Friedman, M. (1977), 'Nobel lecture: inflation and unemployment', *Journal of political economy* **85**(3), 451–472.

Gabaix, X. (2014), 'A sparsity-based model of bounded rationality', *The Quarterly Journal of Economics* **129**(4), 1661–1710.

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

Haldane, A. & McMahon, M. (2018), Central bank communications and the general public, *in* 'AEA papers and proceedings', Vol. 108, pp. 578–83.

Hazell, J., Herreno, J., Nakamura, E. & Steinsson, J. (2022), 'The slope of the phillips curve: evidence from us states', *The Quarterly Journal of Economics* **137**(3), 1299–1344.

Hsieh, C.-T. & Klenow, P. J. (2009), 'Misallocation and manufacturing tfp in china and india', *The Quarterly journal of economics* **124**(4), 1403–1448.

Hunziker, H.-U., Raggi, C., Rosenblatt-Wisch, R. & Zanetti, A. (2022), 'The impact of guidance, short-term dynamics and individual characteristics on firms' long-term inflation expectations', *Journal of Macroeconomics* **71**, 103380.

Juillard, M. & Kamenik, O. (2005), Solving sdge models: approximation about the stochastic steady state, Technical report, Society for Computational Economics.

Kaplan, G., Menzio, G., Rudanko, L. & Trachter, N. (2019), 'Relative price dispersion: Evidence and theory', *American Economic Journal: Microeconomics* **11**(3), 68–124.

King, R. G., Plosser, C. I. & Rebelo, S. T. (1988), 'Production, growth and business cycles: I. the basic neoclassical model', *Journal of monetary Economics* **21**(2-3), 195–232.

Kohlhas, A. N. & Walther, A. (2021), 'Asymmetric attention', *American Economic Review* **111**(9), 2879–2925.

Licandro, G. & Mello, M. (2016), 'Cultural and financial dollarization of households in Uruguay', *Financial decisions of households and financial inclusion: Evidence for Latin America and the Caribbean* .

Licandro, G. & Mello, M. (2019), 'Foreign currency invoicing of domestic transactions as a hedging strategy: evidence for uruguay', *Journal of Applied Economics* **22**(1), 622–634.

Lorenzoni, G. (2009), 'A theory of demand shocks', *American economic review* **99**(5), 2050–2084.

Lorenzoni, G. (2010), 'Optimal monetary policy with uncertain fundamentals and dispersed information', *The Review of Economic Studies* **77**(1), 305–338.

Lucas, R. E. J. (1972), 'Expectations and the neutrality of money', *Journal of economic theory* **4**(2), 103–124.

Lucas, R. E. J. (2003), 'Macroeconomic priorities', *American economic review* **93**(1), 1–14.

Maćkowiak, B., Matějka, F. & Wiederholt, M. (2023), 'Rational inattention: A review', *Journal of Economic Literature* **61**(1), 226–273.

Maćkowiak, B. & Wiederholt, M. (2009), 'Optimal sticky prices under rational inattention', *American Economic Review* **99**(3), 769–803.

Maćkowiak, B. & Wiederholt, M. (2015), 'Business cycle dynamics under rational inattention', *The Review of Economic Studies* **82**(4), 1502–1532.

Mankiw, G. & Reis, R. (2002), 'Sticky information versus sticky prices: a proposal to replace the new Keynesian Phillips curve', *The Quarterly Journal of Economics* (117), 1295–1328.

Milgrom, P. R. & Weber, R. J. (1982), 'A theory of auctions and competitive bidding', *Econometrica: Journal of the Econometric Society* pp. 1089–1122.

Nakamura, E., Steinsson, J., Sun, P. & Villar, D. (2018), 'The elusive costs of inflation: Price dispersion during the us great inflation', *The Quarterly Journal of Economics* **133**(4), 1933–1980.

Nord, L. (2022), 'Shopping, demand composition, and equilibrium prices', *Available at SSRN 4178271* .

Pytka, K. (2018), 'Shopping effort in self-insurance economies', *Manuscript, University of Mannheim* .

Reis, R. (2006*a*), 'Inattentive consumers', *Journal of Monetary Economics* (53), 1761–1800.

Reis, R. (2006*b*), 'Inattentive producers', *Review of Economic Studies* (73), 793–821.

Restuccia, D. & Rogerson, R. (2008), 'Policy distortions and aggregate productivity with heterogeneous establishments', *Review of Economic dynamics* **11**(4), 707–720.

Ryngaert, J. (2018), 'What do (and don't) forecasters know about us inflation?', *Journal of Money, Credit and Banking* .

Sara-Zaror, F. (2021), 'Expected inflation and welfare: The role of consumer search', *Available at SSRN 4127502* .

Sargent, T. J. (1982), *2. The Ends of Four Big Inflations*, University of Chicago Press.

Shapiro, A. H. & Wilson, D. (2021), Taking the fed at its word: A new approach to estimating central bank objectives using text analysis, Federal Reserve Bank of San Francisco.

Sheremirov, V. (2020), 'Price dispersion and inflation: New facts and theoretical implications', *Journal of Monetary Economics* **114**, 59–70.

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

Sung, Y. (2022), 'Macroeconomic expectations and cognitive noise'.

Taylor, J. B. (1982), 'Union wage settlements during a disinflation'.

Tommasi, M. (1994), 'The consequences of price instability on search markets: Toward understanding the effects of inflation', *The American Economic Review* pp. 1385–1396.

Van Hoomissen, T. (1988), 'Price dispersion and inflation: Evidence from israel', *Journal of Political Economy* **96**(6), 1303–1314.

Woodford, M. (2001), 'Imperfect common knowledge and the effects of monetary policy'.

Woodford, M. (2003), *Interest and prices: Foundations of a theory of monetary policy*, Princeton University Press.

Appendix

A.1 Household Problem in General Equilibrium

I start by solving the optimization problem for households at the end of the period. The log-linearized household decisions are:

$$\begin{aligned}
&\hat{c}_i + \tilde{\eta}\hat{n}_i = \hat{w} - \hat{p}_i \\
&\hat{c}_i = \hat{m}_i - \hat{p}_i \\
&\bar{P}\bar{C}(\hat{p}_i + \hat{c}_i) + \bar{M}\hat{m}_i = \bar{W}\bar{N}(\hat{w} + \hat{n}_i) + \bar{M}^s\hat{m}^s + \bar{\Pi}\hat{\Pi}
\end{aligned} \tag{A.13}$$

where $\tilde{\eta} = \eta - \frac{\bar{N}^{1+\eta}}{1-\frac{\bar{N}^{1+\eta}}{1+\eta}}$.

Replacing $\frac{1}{\tilde{\eta}}(-\hat{c}_i - \hat{p}_i + \hat{w}) = \hat{n}_i$ and $\hat{p}_i + \hat{c}_i = \hat{m}_i$ into the budget constraint, we have:

$$\bar{P}\bar{C}(\hat{m}_i) + \bar{M}\hat{m}_i = \bar{W}\bar{N}(\hat{w} + \frac{1}{\tilde{\eta}}(-\hat{m}_i + \hat{w})) + \bar{M}^s\hat{m}^s + \bar{\Pi}\hat{\Pi}$$

Collecting terms, we get the nominal money balance equation in (1.15a). The first and second equations follow from $\hat{p}_i + \hat{c}_i = \hat{m}_i$ and $\hat{m}_i + \tilde{\eta}\hat{n}_i = \hat{w}$. This complete Lemma 1.3.1.

To prove Lemma 1.3.2, I begin by defining the search problem. The search problem requires finding the subset $\mathcal{S}_i \subset J_i^c$ such that the consumer searches for all goods in the set $\mathcal{S}_i$.

By assumption, consumers can only condition search on the non-pecuniary factor A_{if}. Because utility is increasing in A_{if}, the consumer will always search for high A_{if} values and will not search for A_{if} sufficiently low. We must then find the cutoff that determines where consumers are indifferent between searching and not searching. We can therefore write $\mathcal{S}_i = \{f : A_{if} \geq A_i\}$ for a value A_i that is cutoff between searching and not searching. I will define $S_i \equiv \exp\{-\mu A_i\}$ which the optimal fraction of products for which consumers search.

Then, the expression for the household price index is:

$$\ln P_i(S_i) = (1 - S_i)\ln P_i^C + \Gamma(S_i) + S_i \ln(\tau P_i^S)$$

where the function $\Gamma(S_i)$ is given by:

$$\Gamma(S_i) = -\int_{-\ln(S_i)/\mu}^{\infty} \ln A \exp\{-\mu A\}\, dA$$

The value $S_i \in (0,1)$ is the fraction of products for which household i searches. The function $\Gamma(S_i)$ captures the non-pecuniary benefits of searching. It has the properties that $\Gamma(0)$ is a positive constant, $\lim_{S\to 0}\Gamma'(S) = -\infty$, $\lim_{S\to 1}\Gamma'(S) = +\infty$ and is convex. These conditions ensure that households always search for at least a fraction of the goods. Intuitively, these properties arise because there is always some value for the non-pecuniary benefit of search A_{if} that is sufficiently high so that consumers search for substitutes to product f and also some value A_{if} so low that households will not search.

The search problem is to select the cutoff S_i that maximizes:

$$\mathbb{E}_i\left[\exp\{(1-\gamma)\hat{c}_i + \tilde{\eta}\hat{n}_i - \lambda(\hat{m}_i - \hat{p}_i)\}\right]$$

By Lemma 1.3.1, $\hat{m}_i$ and $\hat{n}_i$ are independent of $\hat{p}_i$. Replacing the conditional expectation given $\hat{p}_i^c$ we have:

$$\mathbb{E}_i\left[\hat{U}_i\right] = \exp\left\{(\gamma+\lambda-1)\mathbb{E}_i(\hat{p}_i) - \frac{1}{2}(\gamma+\lambda-1)^2\mathbb{V}_i(\hat{p}_i)\right\}\tilde{U}_i$$

where $\tilde{U}_i$ collects the expected values for $\hat{m}_i$ and $\hat{n}_i$ that are independent of search.

The optimal level of search must then solve:

$$S_i^* = \arg\min_{S_i}\left[\mathbb{E}_i[\ln P_i(S_i)] + \frac{\gamma+\lambda-1}{2}\mathbb{V}_i[\ln P_i(S_i)]\right]$$

This gives Lemma 1.3.2.

I now solve for the expectation and the variance:

$$\begin{aligned}
\mathbb{E}_i[\ln P_i(S_i)] &= \ln P_i^C + S_i\ln(\tau) + \Gamma(S_i) + S_i\mathbb{E}_i\left[\ln P_i^s - \ln P_i^c\right] \\
\mathbb{V}_i[\ln P_i(S_i)] &= (S_i)^2\left[(1-\rho)\sigma_{p_c}^2 + \sigma_{p_s}^2\right]
\end{aligned}$$

Taking the first order condition for the optimization problem, I arrive at the following implicit solution:

$$\ln(A_i) - \mu(\gamma+\lambda-1)S_i\left[(1-\rho)\sigma_{p_c}^2 + \sigma_{p_s}^2\right] = \mu\mathbb{E}_i\left[\ln P_i^s - \ln P_i^c\right] + \mu\ln(\tau) \qquad \text{(A.14)}$$

where $S_i = \exp\{-\mu A_i\}$.

At the steady state, we have that $S_i = \bar{S}$, $A_i = \bar{A}$ and $\mathbb{E}_i\left[\ln P_i^s - \ln P_i^c\right]$ is a constant.

Then, Equation (A.14) becomes:

$$\bar{S} = \frac{\ln \bar{v} - \mu \ln \tau}{\mu(\gamma + \lambda - 1)\left[(1-\rho)\sigma_{p_c}^2 + \sigma_{p_s}^2\right]}$$

Last, log-linearizing Equation (A.14) gives:

$$\hat{S}_i = \mu^2 \bar{v} \frac{(1-\rho)\hat{p}_i^c}{1 + \mu^2 \bar{v} \bar{S}(\rho + \lambda - 1)\left[(1-\gamma)\sigma_{p_c}^2 + \sigma_{p_s}^2\right]}$$

These are the expressions featured in Lemma 1.3.3.

A.2 Firm Problem in General Equilibrium

As in Section 1.2, firms sell to captive and flexible consumers:

$$q_f = q_f^C + q_f^F$$

Log-linearizing this expression, we get $\hat{q}_f = \frac{\bar{q}_f^C}{\bar{q}_f}\hat{q}_f^C + \frac{\bar{q}_f^F}{\bar{q}_f}q_f^F$.

Now, demand from captive consumers is:

$$q_f^C = \frac{1}{P_f} D_f^C$$

where D_f^C is the nominal demand from captive consumers. Its value equals:

$$D_f^C = \int_{i:f \in J_i^c} (1 - S_i) P_i C_i di$$

Log-linearizing, we have that:

$$\hat{d}_f^C = -\frac{\bar{S}_f}{1 - \bar{S}_f}\psi_{p^c}^S (1-\rho)\mathbb{E}(\hat{p}_i | i : f \in J_i^c) + \mathbb{E}(\hat{p}_i + \hat{c}_i | i : f \in J_i^C)$$

Given the sampling assumption and equilibrium price setting, we have that $\mathbb{E}(\hat{p}_i | i : f \in J_i^c) = \hat{p} + \chi\beta_e\eta_f^z$. Further, from the household first order conditions we know that $\hat{p}_i + \hat{c}_i = \hat{m}_i$. Moreover, it was shown in Equation (1.15a) that $\hat{m}_i$ is common across households, and the last term is simply $\hat{m}$.

Demand from flexible consumers is:

$$q_f^F = \int_{i:f \in J_i^s} \frac{1}{P_f}\left(\frac{P_f}{P_i^s}\right)^{1-\varsigma} S_i di$$

Log-linearizing, we have:

$$\hat{q}_f^F = -\varsigma \hat{p}_f + \mathbb{E}((\varsigma - 1)\hat{p}_i^s + \hat{S}_i | i : f \in J_i^s)$$

Note that the sampling assumption implies $\hat{S}_i$ is independent of $f \notin J_i^c$. Hence, $\mathbb{E}\hat{S}_i = \hat{S} = \psi_{p^c}^S (1-\rho)\hat{p}$. Moreover, $\mathbb{E}(\hat{p}_i^s | i : f \in J_i^s) = \hat{p} + \chi \beta_e \eta_f^z$.

The demand elasticity for the firm is:

$$\varepsilon_f = \frac{q_f^C}{q_f}\varepsilon_f^C + \frac{q_f^F}{q_f}\varepsilon_f^F$$

where ε_f^C and ε_f^F are the demand elasticities of captive and flexible consumers, respectively. These values are constant and equal 1 and ς. Log-linearizing the elasticity, we have:

$$\hat{\varepsilon}_f = -\frac{\bar{q}_f^C / \bar{q}_f}{\bar{\varepsilon}_f}(\hat{q}_f^C - \hat{q}_f) - \varsigma \frac{\bar{q}_f^F / \bar{q}_f}{\bar{\varepsilon}_f}(\hat{q}_f^F - \hat{q}_f)$$

Finally, log-linearizing $P_f = \frac{\varepsilon_f}{\varepsilon_f + 1}\frac{W}{Z_f}$ gives:

$$\hat{p}_f = \frac{1}{\bar{\varepsilon}_f + 1}\hat{\varepsilon}_f - \hat{z} - \eta_f^z + \hat{w}$$

These are the expression in Lemma 1.3.4.

A.3 General Equilibrium Steady State

The following expressions show the steady state value for the general equilibrium model. For expositional simplicity, I will present the equilibrium equations by groups.

First, consumer search and price indices are given by:

$$\bar{S}_i = \frac{\ln \bar{v} - \mu \ln \tau}{\mu(\gamma + \lambda - 1)\left[(1-\rho)\sigma_{p_c}^2 + \sigma_{p_s}^2\right]} \tag{A.15}$$

$$\bar{S}_i = \exp\{-\mu \bar{v}\} \tag{A.16}$$

$$\ln \bar{P}_i = (1 - \bar{S}_i)\ln \bar{P}_f + \Gamma(\bar{S}_i) + \bar{S}_i \ln(\tau \bar{P}_i^s) \tag{A.17}$$

$$\bar{P}_i^s = \bar{P}_f \exp\left\{\frac{1-\varsigma}{2}\beta_e^2 \sigma_f^2\right\} \tag{A.18}$$

Next, the following equations show the components of demand, the elasticity and the optimal price:

$$\bar{q}_f^C = \frac{\bar{M}}{\bar{P}_f}(1-\bar{S}_f) \tag{A.19}$$

$$1-\bar{S}_f = (1-\bar{S}_i)\exp\left\{\frac{1}{2}\left(\frac{\bar{S}_i}{1-\bar{S}_i}\psi_{p^c}^S(1-\rho)\chi\beta_e\right)^2(1-\chi)\sigma_f^2\right\} \tag{A.20}$$

$$\bar{q}_f^F = \frac{\bar{M}}{\bar{P}_f}\bar{S}_f\exp\left\{\frac{1}{2}\left[\left(\psi_{p^c}^s(1-\rho)\chi\beta_e\right)^2-((\varsigma-1)\beta_e)^2\right](1-\chi)\sigma_f^2\right\} \tag{A.21}$$

$$\bar{\varepsilon}_f = -\frac{\bar{q}_f^C}{\bar{q}_f^C+\bar{q}_f^F}-\varsigma\frac{\bar{q}_f^F}{\bar{q}_f^C+\bar{q}_f^F} \tag{A.22}$$

$$\bar{P}_f = \frac{\bar{\varepsilon}_f}{\bar{\varepsilon}_f+1}\frac{\bar{W}}{\bar{Z}} \tag{A.23}$$

Next, I show labor demand, supply and market clearing:

$$\bar{L} = \frac{\bar{q}_f}{\bar{Z}}\exp\left\{\frac{1}{2}(\beta_q-1)^2\sigma_f^2\right\} \tag{A.24}$$

$$\frac{\bar{N}_i^\eta}{1-\frac{\bar{N}_i^{1+\eta}}{1+\eta}} = (\gamma-1)\frac{\bar{W}}{\bar{M}} \tag{A.25}$$

$$\bar{L} = \bar{N}_i \tag{A.26}$$

where $\beta_q = -\beta_e\left[\frac{\bar{q}_f^C}{\bar{q}_f}\left(1+\frac{\bar{S}_f}{1-\bar{S}_f}\psi_{p^c}^S(1-\rho)\chi\right)+\frac{\bar{q}_f^F}{\bar{q}_f}\left(\varsigma-(\varsigma-1)\bar{S}_i\chi\right)\right]$.

And finally, we have goods market clearing and the household budget constraint:

$$\bar{M} = \bar{P}_f\bar{q}_f \tag{A.27}$$

$$\frac{\lambda}{\gamma-1}\bar{M} = \bar{P}_i\bar{C}_i \tag{A.28}$$

$$\bar{P}_i\bar{C}_i = \bar{P}_f\bar{q}_f\exp\left\{\frac{1}{2}(\beta_e+\beta_q)^2\sigma_f^2\right\} \tag{A.29}$$

A.4 Solution to the General Equilibrium

The following system collects the equations that characterize the general equilibrium model.

$$\hat{c}_i + \tilde{\eta}\hat{n}_i - \hat{w} + \hat{p}_i = 0 \tag{A.30}$$

$$\hat{p}_i + \hat{c}_i - \frac{\bar{W}\bar{N}}{\bar{P}\bar{C}}(\hat{w} + \hat{n}_i) - \frac{\bar{\Pi}}{\bar{P}\bar{C}}\hat{\Pi} + \frac{\bar{M}}{\bar{P}\bar{C}}\hat{m}_i = \frac{\bar{D}}{\bar{P}\bar{C}}\hat{d} \tag{A.31}$$

$$\hat{c}_i - (\hat{m}_i - \hat{p}_i) = 0 \tag{A.32}$$

$$\hat{p}_i + \psi^S_{p^C}(1-\rho)\Xi\hat{p}^C_i - (1-\bar{S})\hat{p}^C_i - \bar{S}\hat{p}^S_i = 0 \tag{A.33}$$

$$\hat{q}_f - \frac{\bar{q}^C_f}{\bar{q}_f}\hat{q}^C_f - \frac{\bar{q}^F_f}{\bar{q}_f}\hat{q}^F_f = 0 \tag{A.34}$$

$$\hat{q}^C_f + \hat{p}_f - \hat{d}^C_f = 0 \tag{A.35}$$

$$\hat{q}^F_f + \varsigma\hat{p}_f - \hat{d}^F_f = 0 \tag{A.36}$$

$$\hat{d}^C_f + \frac{\bar{S}_f}{1-\bar{S}_f}\psi^S_{p^c}(1-\rho)(\hat{p} + \chi\beta_e\eta^z_f) - \hat{m} = 0 \tag{A.37}$$

$$\hat{d}^F_f - \psi^S_{p^c}(1-\rho)\hat{p} - (\varsigma - 1)(\hat{p} + \bar{S}_i\chi\beta_e\eta^z_f) - \hat{m} = 0 \tag{A.38}$$

$$\hat{\varepsilon}_f + \frac{\bar{q}^C_f/\bar{q}_f}{\bar{\varepsilon}_f}(\hat{q}^C_f - \hat{q}_f) + \varsigma\frac{\bar{q}^F_f/\bar{q}_f}{\bar{\varepsilon}_f}(\hat{q}^F_f - \hat{q}_f) = 0 \tag{A.39}$$

$$\hat{p}_f - \frac{1}{\bar{\varepsilon}_f + 1}\hat{\varepsilon}_f - \hat{w} = -\hat{z} - \eta^z_f \tag{A.40}$$

$$\hat{m} = \hat{d} \tag{A.41}$$

$$\hat{q}_f - \hat{l}_f = \hat{z} + \eta^z_f \tag{A.42}$$

$$\hat{l} - \hat{n} = 0 \tag{A.43}$$

To find the solution, I proceed by iteration as before. It will be useful to postulate a guess for firm quantities as well as prices. First, I will postulate an initial guess:

$$\hat{p}_f = (\beta_m)^0\hat{m} + (\beta_z)^0\hat{z} + (\beta_e)^0\eta^z_f$$
$$\hat{q}_f = (\beta^q_m)^0\hat{m} + (\beta^q_z)^0\hat{z} + (\beta^q_e)^0\eta^z_f$$

These give $(\hat{p})^0 = (\beta_m)^0\hat{m} + (\beta_z)^0\hat{z}$ and $(\hat{q})^0 = (\beta^q_m)^0\hat{m} + (\beta^q_z)^0\hat{z}$. The effect on households' beliefs is $(\rho)^0 = \frac{((\beta_z)^0)^2\sigma_z^2 + ((\beta_m)^0)^2\sigma_m^2}{((\beta_z)^0)^2\sigma_z^2 + ((\beta_m)^0)^2\sigma_m^2 + ((\beta_e)^0)^2\sigma_c^2}$.

These allow me to solve for the remainder equilibrium variables $(\hat{c})^0$, $(\hat{n})^0$, $(\hat{w})^0$, $(\hat{p}_\bullet)^0$, $(\hat{l})^0$, $(\hat{\Pi})^0$ as a function of $(\hat{p})^0$, $(\hat{q})^0$. I will also use the equilibrium condition that $\hat{m}_t =$

$\hat{d}_t$. The system of equations for aggregate variables is:

$$\hat{c} + \tilde{\eta}\hat{n} - \hat{w} + \hat{p}_{\bullet} = 0$$
$$\hat{c} - (\hat{m} - \hat{p}_{\bullet}) = 0$$
$$\hat{p}_{\bullet} = (1 - \psi^S_{p^c}(1 - (\rho)^0)\bar{S}\Xi)(\hat{p})^0$$
$$\hat{p}_{\bullet} + \hat{c} - \tfrac{\bar{W}\bar{N}}{\bar{P}\bar{C}}(\hat{w} + \hat{n}) - \tfrac{\bar{\Pi}}{\bar{P}\bar{C}}\hat{\Pi} + \tfrac{\bar{M}}{\bar{P}\bar{C}}\hat{m} = \tfrac{\bar{D}}{\bar{P}\bar{C}}\hat{d}$$
$$\hat{m} = \hat{d}$$
$$\hat{l} = (\hat{q})^0 - \hat{z}$$
$$\hat{l} - \hat{n} = 0$$

Next, I solve for the values $(\hat{d}^C_{ft})^0$ and $(\hat{d}^F_{ft})^0$ that together with $(\hat{w}_t)^0$ are the only factors that determine how general equilibrium features affect firms' profit maximization conditions. The values are:

$$(\hat{d}^C_f)^0 = \hat{m} - \tfrac{\bar{S}}{1-\bar{S}}\psi^S_{p^c}(1 - (\rho)^0)((\hat{p}_t)^0 + \chi(\beta_e)^0\eta^z_{ft})$$
$$(\hat{d}^F_f)^0 = \psi^S_{p^c}(1 - (\rho)^0)(\hat{p})^0 + (\varsigma - 1)((\hat{p})^0 + \bar{S}\chi(\beta_e)^0\eta^z_f) + \hat{m}$$

Finally, we can generate a new guess by combining firms' profit maximization conditions:

$$\hat{q}_f - \frac{\bar{q}^C}{\bar{q}}\hat{q}^C_f - \frac{\bar{q}^F}{\bar{q}}\hat{q}^F_f = 0$$
$$\hat{q}^C_f + \hat{p}_f = (\hat{d}^C_f)^0$$
$$\hat{q}^F_f + \varsigma\hat{p}_f = (\hat{d}^F_f)^0$$
$$\hat{\varepsilon}_f + \frac{\bar{q}^C/\bar{q}}{\bar{\varepsilon}}(\hat{q}^C_f - \hat{q}_{ft}) + \varsigma\frac{\bar{q}^F/\bar{q}}{\bar{\varepsilon}}(\hat{q}^F_f - \hat{q}_f) = 0$$
$$\hat{p}_f - \frac{1}{\bar{\varepsilon}+1}\hat{\varepsilon}_f = -\hat{z} - \eta^z_f + (\hat{w})^0$$

With these equation, we can find a new guess $(\beta_m)^1, (\beta_z)^1, (\beta_e)^1$ and $(\beta^q_m)^1, (\beta^q_z)^1, (\beta^q_e)^1$. The process is iterated until the desired convergence is achieved.

A.5 Derivation of Welfare Function

Equation 1.20 is derived by plugging in the solution to the log-linearized model into the expected utility.

$$\mathcal{W} = \mathbb{E}\left[\frac{\bar{C}_i^{1-\gamma}}{1-\gamma}\frac{1}{1 - \frac{\bar{N}_i^{1+\eta}}{1+\eta}}\left(\frac{\bar{M}_i}{\bar{P}_i}\right)^{-\lambda}\exp\left\{(1-\gamma)\hat{c}_i + \tilde{\eta}\hat{n}_i - \lambda(\hat{m}_i - \hat{p}_i)\right\}\right] \tag{A.44}$$

We can replace $\bar{C}_i = \frac{\bar{M}_i}{\bar{P}_i}$ and $\hat{c}_i = \hat{m}_i - \hat{p}_i$. Given the solution to household's consumption, labor and money holdings given by (1.15a), we have that:

$$\begin{aligned} \hat{n}_i &= \hat{n} \\ \hat{c}_i &= \hat{c} + (\hat{c}_i - \hat{c}) \\ (\hat{c}_i - \hat{c}) &= -(\hat{p}_i - \hat{p}_\bullet) \end{aligned}$$

The last expression implies that cross-sectional consumption dispersion is orthogonal to aggregate consumption because $\hat{p}_i - \hat{p}_\bullet$ depends only on sampling shocks which are i.i.d.. Replacing these expressions into equation (A.44) gives:

$$\mathcal{W} = \bar{\mathcal{W}}_i \mathbb{E}\left[\exp\left\{\frac{1}{2}\left[(1-\gamma-\lambda)\hat{c} + \tilde{\eta}\hat{n} + (1-\gamma-\lambda)(\hat{c}_i - \hat{c})\right]\right\}\right] \tag{A.45}$$

I now proceed to explain how to express the welfare function in terms of prices. From the equilibrium conditions, we have that:

$$\hat{p}_\bullet = (1 - \psi^S_{p^c}(1-\rho)\bar{S}_i\Xi)\hat{p} \equiv \phi_{p_\bullet}\hat{p}$$

Further, given that $\beta_m = \beta_z$:

$$\hat{p} = \beta_m(\hat{m} - \hat{z}) \equiv \beta_m \hat{p}^{FI}$$

where $\hat{p}^{FI} = \hat{m} - \hat{z}$ is the full-information price level.

Now, let us solve for $\hat{q}$ and insert this in labor $\hat{n} = \hat{q} - \hat{z}$:

$$\begin{aligned} \hat{q} &= -\hat{p} + \left[-\frac{\bar{q}^C_f}{\hat{q}_f}\frac{\hat{S}_f}{1-\hat{S}_f} + \frac{\bar{q}^F_f}{\bar{q}_f}\right]\psi^S_{p^c}(1-\rho)\hat{p} + \hat{m} \\ \hat{n} &= \tilde{\phi}_n\hat{p} + \hat{m} - \hat{z} = \phi_n\hat{p} \end{aligned}$$

where $\phi_n = -1 + \left[-\frac{\bar{q}^C_f}{\hat{q}_f}\frac{\hat{S}_f}{1-\hat{S}_f} + \frac{\bar{q}^F_f}{\bar{q}_f}\right]\psi^S_{p^c}(1-\rho) + \frac{1}{\beta_m}$.

Then, we can solve for consumption:

$$\begin{aligned} \hat{c} &= \hat{m} - \hat{p}_\bullet = \hat{p}^{FI} - \phi_{p_\bullet}\hat{p} + \hat{z} \\ &= \phi_c\hat{p} + \hat{z} \end{aligned}$$

where $\phi_c = 1/\beta_m - \phi_{p_\bullet}$.

This gives us the expression for aggregate variables:

$$(1-\gamma-\lambda)\hat{c} + \tilde{\eta}\hat{n} = \left[(1-\gamma-\lambda)\phi_c + \tilde{\eta}\phi_n\right]\hat{p} + (1-\gamma-\lambda)\hat{z}$$

Now, let me work on the terms related to cross-sectional price dispersion. Recall that

$(\hat{c}_i - \hat{c}) = -(\hat{p}_i - \hat{p}_\bullet)$. Then:

$$\begin{aligned} \hat{p}_i &= (1 - \bar{S}_i - \tilde{\mu})\hat{p}_i^c + \bar{S}_i\hat{p}_i^s \\ &= (1 - \bar{S}_i - \tilde{\mu} + \bar{S}_i)\hat{p} + (1 - \bar{S}_i - \tilde{\mu})(\hat{p}_i^c - \hat{p}) + \bar{S}_i(\hat{p}_i^s - \hat{p}) \\ \hat{p}_i - \hat{p}_\bullet &= (1 - \bar{S}_i - \tilde{\mu})(\hat{p}_i^c - \hat{p}) + \bar{S}_i(\hat{p}_i^s - \hat{p}) \end{aligned}$$

Then, equation A.45 can be written as:

$$\begin{aligned} \mathcal{W} &= \bar{\mathcal{W}}_i\mathbb{E}\Bigg[\exp\Bigg\{ \left[(1-\gamma-\lambda)\phi_c + \tilde{\eta}\phi_n\right]\hat{p} + (1-\gamma-\lambda)\hat{z} + \\ &\qquad (1-\gamma-\lambda)((1-\bar{S}_i-\tilde{\mu})(\hat{p}_i^c - \hat{p}) + \bar{S}_i(\hat{p}_i^s - \hat{p}))\Bigg\}\Bigg] \end{aligned}$$

Now, we take the expectation which equals. Given the assumption that shocks are normal, we have that:

$$\begin{aligned} \mathcal{W} &= \bar{\mathcal{W}}_i\mathbb{E}\Bigg[\exp\Bigg\{ \frac{1}{2}\left[(1-\gamma-\lambda)\phi_c + \tilde{\eta}\phi_n\right]^2\sigma_p^2 \\ &\qquad + \left[(1-\gamma-\lambda)\phi_c + \tilde{\eta}\phi_n\right](1-\gamma-\lambda)\mathbb{C}(\hat{p},\hat{z}) \\ &\qquad + \frac{1}{2}(1-\gamma-\lambda)^2\sigma_z^2 \\ &\qquad + \frac{1}{2}(1-\gamma-\lambda)^2((1-\bar{S}_i-\tilde{\mu})^2 + (\bar{S}_i)^2)\chi\sigma_{p_f}^2\Bigg\}\Bigg] \end{aligned}$$

Finally, note that we can write $\mathbb{C}(\hat{p},\hat{z})$ as:

$$\mathbb{C}(\hat{p},\hat{z}) = \frac{1}{\beta_m}(1 - (R_{m,p})^2)\sigma_p^2$$

where $R_{m,p}^2 \equiv \frac{\beta_m^2\sigma_m^2}{\beta_m^2\sigma_m^2 + \beta_z^2\sigma_z^2}$ is the fraction of the inflation variance explained by money shocks.

The last step involves collecting terms and noting that $(1-\gamma-\lambda)^2\sigma_z^2$ is independent of policy.

A.6 Proof of Proposition 1.4.2 and 1.4.3

I will begin by showing that perturbations to consumer search have a second-order welfare effect on consumer welfare. To fix ideas, in the log-linearized equilibrium household search decisions can be written as:

$$S_i = a + b\hat{p}_i^c$$

The log-linearized price index for households is:

$$p_i = \bar{p}_i^c + \hat{p}_i^c + (a + b\hat{p}_i^c)(\bar{p}_i^s - \bar{p}_i^c) + (a + b\hat{p}_i^c)(\hat{p}_i^s - \hat{p}_i^c) + \Gamma(a) + b\Gamma(a)'\hat{p}_i^c$$

Households wish to minimize:

$$\mathbb{E}\left[p_i - \frac{\gamma + \lambda - 1}{2}(S_i)^2\left((1-\rho)\sigma_p^2 + \sigma_{p^s}^2\right)\Big|\hat{p}_i^c\right]$$

This is the log-linear version of the solution developed in Appendix A.1. Hence, the solutions are $a = \bar{S}_i$ and $b = \psi_{p^c}^S(1-\rho)$, and the derivatives with respect to a and b must be zero given that the solution is interior.

To show the first-order effects, I will show that the proposed changes in parameters distort some of the first-order conditions for households. In equilibrium we have that:

$$\bar{P}_i\bar{C}_i = \bar{P}_f\bar{q}_f \exp\left\{\tfrac{1}{2}(\beta_e + \beta_q)^2\sigma_f^2\right\} \tag{A.46}$$

$$\bar{N}_i = \tfrac{\bar{q}_f}{\bar{Z}}\exp\left\{\tfrac{1}{2}(\beta_q - 1)^2\sigma_f^2\right\} \tag{A.47}$$

$$\frac{\bar{N}_i^\eta}{1 - \frac{\bar{N}_i^{1+\eta}}{1+\eta}} = (\gamma - 1)\frac{\bar{W}}{\bar{P}_i\bar{C}_i} = \lambda\frac{\bar{W}}{\bar{M}} \tag{A.48}$$

where $\beta_q = -\beta_e\left[\frac{\bar{q}_f^C}{\bar{q}_f}\left(1 + \frac{\bar{S}_f}{1-\bar{S}_f}\psi_{p^c}^S(1-\rho)\chi\right) + \frac{\bar{q}_f^F}{\bar{q}_f}\left(\varsigma - (\varsigma - 1)\bar{S}_i\chi\right)\right]$.

To show this change in search activity has a first-order effect on welfare through the reallocation of resources, first note that changes to either $\bar{S}_i$ or $\psi_{p^c}^S(1-\rho)$ affect the elasticity of output with respect to idiosyncratic TFP β_q. The search elasticity $\psi_{p^c}^S(1-\rho)$ generally raises the value of β_q given that $\beta_e < 0$, and $\bar{S}_i$ also raises value of β_q through its effect on the weights $\bar{q}_f^F/\bar{q}_f$ and $\bar{q}_f^C/\bar{q}_f$. For notational simplicity, I will write $\beta_q \equiv \beta_e\delta$ where $\delta < 0$, and consider the case of a fall in δ.

Combining the market clearing condition for goods with the first-order conditions for households, we get:

$$\bar{N}_i\lambda\frac{1}{\bar{M}} = \frac{1}{\bar{P}_f}\frac{1}{\bar{Z}}\exp\left\{\frac{1}{2}\left[(\beta_e\delta - 1)^2 - (\beta_e)^2(1+\delta)^2\right]\sigma_f^2\right\} \tag{A.49}$$

The derivative of the expression in the exponential with respect to δ is $-(\beta_e^2 + \beta_e) < 0$ for $\beta_e < -1$. Therefore, $\bar{N}_i$ rises following the fall in δ.

Combining the previous expression with the labor demand, we have that:

$$\kappa^0 = \frac{\bar{q}_f}{\bar{Z}}\exp\left\{\frac{1}{2}\left[(\beta_e)^2(1+\delta)^2\right]\sigma_f^2\right\}$$

where κ^0 is a constant. By (A.46), this implies consumption is not affected by changes in production and the welfare effect arises from the change in equilibrium labor effort.

To see why this effect is first-order, we have that:

$$\begin{aligned} d\mathcal{W} &= \mathbb{E}\left[U_S d\bar{S}_i\right] + U_N d\bar{N}_i \\ &= U_N d\bar{N}_i \end{aligned}$$

where the first effect is zero to a first order, and so the effect on equilibrium labor remains. The effect of a change in the search elasticity is analogous. This concludes the proof of Lemma 1.4.2.

Now, consider the welfare effect of β_e. The parameter β_e generates cross-sectional consumption volatility and also affects average consumption $\bar{C}_i$. Further, the parameter β_e also affects the steady-state equilibrium, which can be reasoned analogously as the effect on the elasticity. To see the effects on average consumption and cross-sectional consumption volatility, recall that consumption is inversely proportional to marginal changes in the household price index P_i. The change in average consumption will then be inversely proportional to the change in $\bar{P}_i$. This change is given by:

$$\frac{d\bar{P}_i}{d\beta_e} = \bar{P}_i \bar{S}_i (1-\varsigma)\sigma_f^2 \beta_e$$

This derivative is negative given that $\varsigma > 1$. It is intuitive that the price index should fall, since search component of the price index is a CES aggregate which falls with price dispersion.

The effect on cross-sectional consumption volatility is given by:

$$\frac{d\mathbb{V}(\hat{p}_i)}{d\beta_e} = 2\left[(1-\bar{S}_i-\tilde{\mu})^2 + (\bar{S}_i)^2\right]\chi\sigma_f^2\beta_e$$

Third, β_e affects cross-sectional production. To simplify the analysis, I will assume away any differential search responses caused by the change in β_e. That is, I will consider that the factor $\beta_q = \beta_e \bar{\varepsilon}_f - (\beta_e)^0 \left[\frac{\bar{q}_f^C}{\bar{q}_f}\left(\frac{\bar{S}_f}{1-\bar{S}_f}\psi_{p^c}^S(1-\rho)\chi\right) - \frac{\bar{q}_f^F}{\bar{q}_f}(\varsigma-1)\bar{S}_i\chi\right]$ and take the derivative with respect to β_e around $\beta_e = (\beta_e)^0$. This implies the fraction of searchers and non-searchers do not change, only the distribution of searchers among firms. This makes the analysis simpler by not having to consider welfare costs associated to differential search responses. Accounting for these would not change the overall conclusions, given that search responses also carry distortions.

Combining the market clearing conditions for labor, we have that:

$$\frac{d\bar{N}_i}{d\beta_e} = \bar{N}_i\left(-2\beta_e\bar{\varepsilon}_f - \beta_e - \bar{\varepsilon}_f\right)\sigma_f^2$$

Putting this together, the effect on expected utility is:

$$\frac{d\mathcal{W}}{d\beta_e} = \mathcal{W}\Bigg[(\gamma+\lambda-1)\bar{S}_i(\varsigma-1)\sigma_f^2\beta_e - (\gamma+\lambda-1)^2((1-\bar{S}_i-\tilde{\mu})^2+(\bar{S}_i)^2)\chi\sigma_f^2\beta_e + \frac{\bar{N}_i^{1+\eta}}{1-\frac{\bar{N}_i^{1+\eta}}{1+\eta}}(2\beta_e\bar{\varepsilon}_f+\bar{\varepsilon}_f+\beta_e)\sigma_f^2\Bigg] \tag{A.50}$$

While it is difficult to compare the magnitudes in this expression, it is generally not zero. I will elaborate on the different terms and finally show that the derivative is not zero numerically.

To build intuition as to why these terms are not zero, let us abstract from risk adjustments and consider only the allocative effect of β_e by analyzing its effect on labor and consumption. Then, the welfare effect can be written as:

$$\begin{aligned} d\mathcal{W}^{Alloc} &= U_C\bar{C}_i\left[\frac{1}{\bar{C}_i}d\bar{C}_i + \frac{U_N\bar{N}_i}{U_C\bar{C}_i}d\bar{N}_i\right] \\ &= U_C\bar{C}_i\left[\bar{S}_i(\varsigma-1)\beta_e + \frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i}(2\beta_e\bar{\varepsilon}_f+\beta_e+\bar{\varepsilon}_f)\right]\sigma_f^2 \\ &\approx U_C\bar{C}_i\left[\bar{S}_i(\varsigma-1)\left(1-\frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i}\right)\beta_e + \frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i}\bar{\varepsilon}_f(\beta_e+1)\right]\sigma_f^2 \end{aligned}$$

where the last term uses the fact that $\bar{S}_i \approx \bar{q}_f^F/\bar{q}_f$ and so $\bar{S}_i(\varsigma-1) \approx -\bar{\varepsilon}_f + 1$. The terms in the bracket capture how changes in price dispersion affect the allocative distortions that are in place (see discussion in the main text). There is generally no condition that links the two terms, and the sign of this expression is undetermined.

Last, risk adjustments always increase the marginal cost of price dispersion. This is given by the term $(\gamma+\lambda-1)^2((1-\bar{S}_i-\tilde{\mu})^2+(\bar{S}_i)^2)\chi\sigma_f^2\beta_e$ in Equation (A.50).

Last, the effect on markups can be shown by expressing $\bar{P}_f = \overline{markup}_f\frac{\bar{W}}{\bar{Z}}$.

$$\frac{\bar{N}_i^{1+\eta}}{1-\frac{\bar{N}_i^{1+\eta}}{1+\eta}} = (\gamma-1)\frac{1}{\overline{markup}_f}\exp\left\{\frac{1}{2}(\beta_e\bar{\varepsilon}_f-1)^2\sigma_f^2 - \frac{1}{2}(\beta_e)^2(1+\bar{\varepsilon}_f)^2\sigma_f^2\right\} \tag{A.51}$$

In this case, a rise in markups lowers equilibrium labor, all else equal. Because labor falls, consumption must fall by the same proportion. To see the effect on welfare, I can write:

$$
\begin{aligned}
d\mathcal{W} &= U_c d\bar{C}_i \left[\frac{U_N}{U_C} \frac{d\bar{N}_i}{d\bar{C}_i} + 1 \right] \\
&= U_c d\bar{C}_i \left[\frac{U_N \bar{N}_i}{U_C \bar{C}_i} + 1 \right] \\
&= U_c d\bar{C}_i \left[-\frac{\bar{W} \bar{N}_i}{\bar{P}_i \bar{C}_i} + 1 \right]
\end{aligned}
$$

The second line holds because given constant returns to scale, consumption moves proportionately with labor, and the third line comes from households' first-order condition. Note that $\frac{\bar{W}\bar{N}_i}{\bar{P}_i\bar{C}_i} \neq 1$, so the welfare effect is first-order.