Essays on Sticky Prices and High Inflation Environments

Daniel Villar

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2016

ABSTRACT

Essays on Sticky Prices and High Inflation Environments

Daniel Villar

It has been well established for a long time that sticky prices are fundamental to our understanding of monetary policy. Indeed, sticky prices are a common micro-foundation in models of monetary policy and nominal aggregate fluctuations, as monetary variables typically do not have real economic effects if prices are fuly flexible. This is why price stickiness has been the focus of much research, both theoretical and empirical. A particularly exciting development in this literature has been the recent availability of large, detailed, micro data sets of individual prices, which allow us to observe when and how often the prices of individual goods and sevices change. This type of data has greatly improved our ability to discipline the theoretical models that are used to analyze monetary policy, and advances in sticky price modelling have also provided important questions to ask of the data. The most common data set used in this literature has been the micro data underlying the U.S. Consumer Price Index. While work with this data has produced important results, an important limitation is that it has, until recently, only been available going back to 1988. This is a limitation because it means that the data set only cover periods of low and stable inflation, which limits the types of questions that the price data can help answer.

In this dissertation, I present an extension to this data set: in work carried out with Emi Nakamura, Jón Steinsson and Patrick Sun, we re-constructed an older portion of the data to extend it back to 1977. With this new sample, we can study the high inflation periods of the late 1970's and early 1980's, and in this dissertation I explore various questions related

to monetary policy, and show that several important insights can be gained from this new data set.

Chapter 1, "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation", presents the extended CPI data set and addresses a key policy question: How high an inflation rate should central banks target? This depends crucially on the costs of inflation. An important concern is that high inflation will lead to inefficient price dispersion. Workhorse New Keynesian models imply that this cost of inflation is very large. An increase in steady state inflation from 0% to 10% yields a welfare loss that is an order of magnitude greater than the welfare loss from business cycle fluctuations in output in these models. We assess this prediction empirically using a new dataset on price behavior during the Great Inflation of the late 1970's and early 1980's in the United States. If price dispersion increases rapidly with inflation, we should see the absolute size of price changes increasing with inflation: price changes should become larger as prices drift further from their optimal level at higher inflation rates. We find no evidence that the absolute size of price changes rose during the Great Inflation. This suggests that the standard New Keynesian analysis of the welfare costs of inflation is wrong and its implications for the optimal inflation rate need to be reassessed. We also find that (non-sale) prices have not become more flexible over the past 40 years.

Chapter 2, "The Skewness of the Price Change Distribution: A New Touchstone for Sticky Price Models", documents the predictions of a broad class of existing price setting models on how various statistics of the price change distribution change with the rate of aggregate inflation. Notably, menu cost models uniformly feature the price change distribution becoming less dispersed and less skewed as inflation rises, while in the Calvo model both relations are positive. Using a novel data set, the micro data underlying the U.S. CPI from the late 1970's onwards, we evaluate these predictions using the large variation in inflation over this period. Price change dispersion does indeed fall with inflation, but skewness does not, meaning that menu cost models are at odds with these empirical patterns. The Calvo model's prediction on price change skewness are consistent with the data, but it fails to match the positive relationship between inflation and the frequency of price change, and the negative relationship between inflation and price change dispersion. Since the negative correlations for dispersion and skewness are driven by the selection effect in menu cost models, the evidence presented suggests that selection is less substantial than in menu cost models.

Chapter 3, "The Selection Effect and Monetary Non-Neutrality in a Random Menu Cost Model", presents a random menu cost model that nests the Golosov and Lucas (2007) and Calvo (1983) models as extreme cases, as well as intermediate cases, depending on the distribution of menu costs. This model includes idiosyncratic technology shocks and aggregate demand shocks, so it can be applied to price micro data, and to evaluate the degree of monetary non-neutrality implied by different kinds of menu cost distributions. This model can match the empirical patterns presented in Chapter 2. I find that a random menu cost model with a much weaker selection effect (than in existing menu cost models) no longer predicts such a negative relationship between inflation and price change skewness, but still predicts that the frequency of price change rises with inflation, as in the data, and contrary to the Calvo model. This model also predicts a very high degree of monetary non-neutrality, and the results overall provide evidence in favor of high non-neutrality.

Chapter 4, "The State-Dependent Price Adjustment Hazard Function: Evidence from High Inflation Periods", considers a model-free approach to understanding sticky prices and non-neutrality. The price adjustment hazard function has been used to establish the relationship between individual firms' price setting behavior (micro-level price stickiness) and the response of the aggregate price level to monetary shocks (aggregate stickiness, or monetary non-neutrality), but scant work has been done to estimate the function empirically. We show first that various types of hazard functions (with widely different levels of implied aggregate stickiness) can match the unconditional moments that have been the focus of empirical work on sticky prices (such as the average frequency and size of price changes). However, the relationship between inflation and the shape of the price change distribution over time provides considerable information on the shape of the hazard function. In particular, we find that in order to match the positive inflation-frequency correlation, and the non-negative inflation-price change skewness correlations, the hazard function has to be asymmetric around zero (price increases are overall more likely than decreases) and relatively flat for small to intermediate values of the desired price gap. The latter feature means that our estimated hazard function implies a large degree of aggregate flexibility.

Contents

List of Tables

List of Figures

Acknowledgements

This dissertation has been the product of several years of work, much of it collaborative. It is therefore natural that many people deserve my gratitude for its completion. I am especially thankful to my advisors, Emi Nakamura and Jón Steinsson. Not only have they provided me with encouragement, guidance, and support, but none the work that I have done would have been possible if they had not introduced me to the world of empirical sticky price research, and brought me onto the project (and adventure) of re-constructing the historical CPI data on which this dissertation is based. Patrick Sun's help in getting myself familiarized with the data was also extremely valuable. Two of the chapters are jointly written with Shaowen Luo, who has been an outstanding co-author. Working together has given me a great introduction to comparative advantage put in practice. I am also extremely grateful to Ricardo Reis and Mike Woodford for their advice and support.

My work has also been greatly enriched by discussions with and comments from professors and fellow students at Columbia, such as Serena Ng, Jennifer La'O, Gauti Eggertsson, Steve Zeldes, Savi Sundaresan, Sergey Kolbin, Stéphane Dupraz, Oskar Zorrilla, Keshav Dogra, Jorge Mejía, and many others. Since most of the empirical work was carried out by the BLS, I also owe a very special debt of gratitude to everyone there who not only helped me gain access to the data and with my questions and requests, but who also made my time there welcoming and enjoyable. I want to thank my co-ordinators Ted To and John Molino, as well as Jong Grenlees, Tim Erickson, Fred Rowland, Anna Sanders, Rob Cage, Greg Kurtzon, Erick Sager, Steve Grandits, and Mina Kim.

I am also extremely grateful for having been part of such a friendly and supportive group of classmates, without whom these past six years would not have been the same. Similarly, I cannot thank my friends and family enough for giving me so much love and for sharing in all the different experiences of the last few years. My parents, in particular, have provided me with unconditional love and support, with encouragement, and helped me stay grounded and optimistic. I cannot fully express how grateful I am for Michelle's love and presence in my life. She is always there for me, and makes every day a joyful one, and every challenge that must be faced, worthwhile. I dedicate this dissertation to her, and to my grandparents. whose love and example continue to inspire me.

Finally, I thank and praise God for the gift of my life, for all the opportunities and people He has placed in it, and for his grace. I pray that this dissertation, and all the work that I will do, will bring glory to his name, and may, in some way, contribute to the good of all. "For from Him and through Him and to Him are all things. To Him be glory for ever. Amen" (Rom 11:36)

Chapter 1

The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation

Emi Nakamura, Jón Steinsson, Patrick Sun and Daniel Villar

1.1 Introduction

Recent years have seen a resurgence of interest in the question of the optimal level of inflation. In the years before the Great Recession, there was a growing consensus emerging among policymakers that good policy consisted of targeting an inflation rate close to zero. One manifestation of this was that many countries adopted explicit inflation targets concentrated around 2% per year. Within academia, prominent studies argued for still lower rates of inflation even after having explicitly taken account of the zero lower bound (ZLB) on nominal interest rates (Coibion et al., 2012; Schmitt-Grohe and Uribe, 2011). The Great Recession has led to reconsideration of this consensus view with an increasing number of economists arguing for targeting a higher inflation rate of say 4% per year (see, e.g., Ball, 2014; Blanchard et al., 2010; Blanco, 2015b; Krugman, 2014).

An important concern with targeting higher inflation is that this will increase price dispersion and thereby distort the allocative role of the price system. Intuitively, in a high inflation environment, relative prices will fluctuate inefficiently as prices drift away from their optimal value during intervals between price adjustment. As a consequence relative prices will no longer give correct signals regarding relative costs of production, leading production efficiency to be compromised.

In standard New Keynesian models—the types of models used in most formal analysis of the optimal level of inflation—these costs are very large even for moderate levels of inflation. Calibrating such a model in a relatively standard way, we show that the consumption equivalent welfare loss of moving from 0% inflation to 10% inflation is roughly 10%. For comparison, the welfare costs of business cycle fluctuations in output—even including large recessions like the Great Depression and Great Recession—are an order of magnitude smaller in these same models.[1] No wonder these models strongly favor virtual price stability.

Measuring the sensitivity of inefficient price dispersion to changes in inflation is challenging for several reasons. One challenge is the small amount of variation we have seen in the inflation rate in the U.S. over the last few decades. Existing BLS micro-data on U.S. consumer prices have been influential in establishing basic facts about the frequency and size of price changes (Bils and Klenow, 2004; Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008). These data, however, have the substantial disadvantage that they span only the post-1987 Greenspan-Bernanke period of U.S. monetary history, when inflation was low and stable. This seriously limits their usefulness in studying how variation in inflation affects the economy.

To overcome this challenge, we have extended the BLS micro-dataset on U.S. consumer prices back to 1977. This allows us to analyze a period when inflation in the U.S. rose sharply—peaking at roughly 14% per year in 1980—and was then brought down to a lower level in dramatic fashion by the Federal Reserve under the leadership of Paul Volcker (see Figure 1.1). We constructed these new data from original microfilm cartridges found at the BLS by first scanning them and then converting them to a machine-readable dataset using custom optical character recognition software. This effort took several years to complete partly because the data were never allowed to leave the BLS building in Washington DC.

[1] The models used to analyze the costs of welfare in the New Keynesian literature typically assume a representative agent with constant relative risk aversion preferences and output fluctuations that are trend stationary. In this case, Lucas (2003) shows that the consumption equivalent welfare loss of business cycle fluctuations in consumption over the period 1947-2001 are 0.05% if consumers are assumed to have log-utility. Redoing Lucas' calculation with a coefficient of relative risk aversion of 2 and considering fluctuations in annual per capita consumption around a linear trend over the period 1920-2009 implies a welfare loss of 0.4%. A substantial literature has since argued that Lucas' calculation substantially understates the true costs of business cycle fluctuations (see, e.g., Barro, 2009, Krusell et al., 2009).

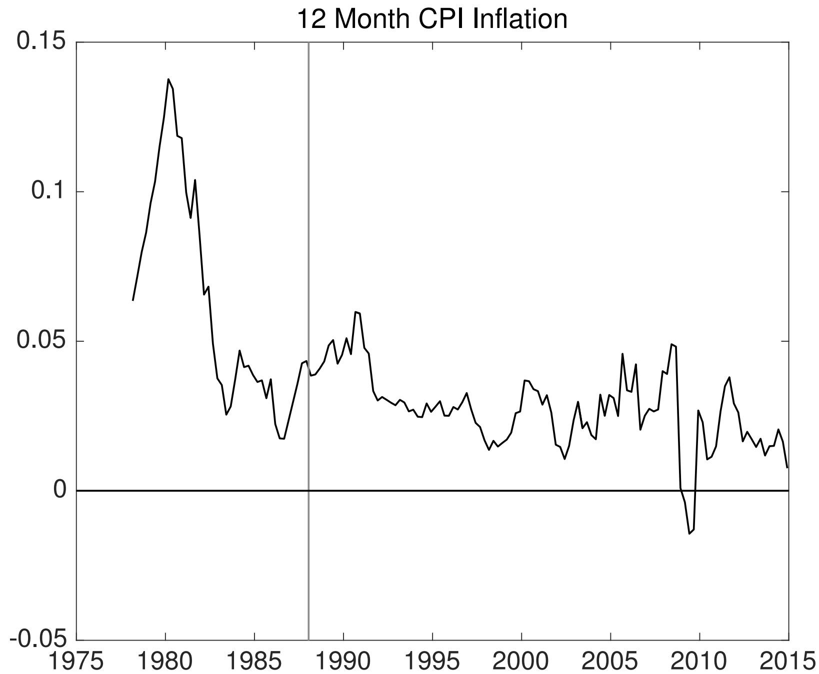

Figure 1.1: CPI Inflation in the U.S.

A second challenge to measuring the sensitivity of inefficient price dispersion to changes in inflation is that much of the cross-sectional dispersion in prices—even within narrowly defined product categories—likely results from heterogeneity in product size and quality (e.g., a can of soda versus a 2 liter bottle, organic versus non-organic milk, Apple's iPhone 6S versus LG's G4 smartphone). The simplest way to empirically assess price dispersion is to calculate the standard deviation of prices within a narrow category. But this approach will lump together desired price dispersion resulting from heterogeneous product size and quality and inefficient price dispersion resulting from price rigidity. In fact, the amount of desired price dispersion within even narrow product categories is likely to dwarf inefficient price dispersion at moderate levels of inflation.

To overcome this challenge, we assess the sensitivity of inefficient price dispersion to changes in inflation by looking at how the absolute size of price changes varies with inflation.

Intuitively, if inflation leads prices to drift further away from their optimal level, we should see prices adjusting by larger amounts when they adjust. The absolute size of price adjustments should reveal how far away from optimal the adjusting prices had become before they were adjusted. The absolute size of price adjustment should therefore be highly informative about inefficient price dispersion.

We show that the mean absolute size of price changes in the U.S. is essentially flat over our entire sample period. There is, thus, no evidence that prices deviated more from their optimal level during the Great Inflation period when inflation was running at higher than 10% per year than during the more recent period when inflation has been close to 2% per year.

We conclude from this that the main costs of inflation in the New Keynesian model are completely elusive in the data. This implies that the strong conclusions about optimality of low inflation rates reached by researchers using models of this kind need to be reassessed. It may well be that inflation rates above 2% have other important costs. A strong consensus for low inflation being optimal must rely on these other costs outweighing the benefits of higher inflation.

Rather than seeing an increase in the absolute size of price changes during the Great Inflation, we see a substantial increase in the frequency of price change. The behavior of both the absolute size and frequency of price change as inflation varies in our sample line up much better with the predictions of menu cost models than they do with the predictions of the workhorse Calvo (1983) model. Intuitively, in the menu cost model, prices never drift too far from their optimal level since firms find it optimal to pay the (relatively small) menu cost before this happens. This greatly limits the extent to which price dispersion rises with inflation in the menu cost model and, as a result, the welfare loss from increasing inflation is small in this model, a point emphasized by Burstein and Hellwig (2008).

A second dramatic result of our analysis is that, despite all of the technological change that has occurred over the past four decades, regular prices (excluding temporary sales) do not seem to have become more flexible over this period, controlling for inflation. We show that a simple menu cost model with a fixed menu cost over the entire sample period can match the empirical relationship between the frequency of price change and inflation. Menu costs are, of course, a veil for a variety of deeper frictions in the price adjustment process arising from technological, managerial, or customer-related factors. Whatever these costs are, they do not appear to be going away over time.

In sharp contrast, we show that the frequency of temporary sales has increased substantially over the past four decades. Temporary sales occur only in a subset of sectors. But their frequency has increased substantially in all of these sectors. Whether this has important implications for aggregate price flexibility is a question of active research over the past decade. The empirical literature has emphasized that temporary sales have quite different empirical properties from those of regular prices. Sales are much more transitory than other price changes, and less responsive to macroeconomic conditions. These characteristics substantially limit the contribution of temporary sales to aggregate price flexibility.[2] Moreover, this growth in temporary sales leaves the large and growing "gorilla in the room" sector—the service sector—untouched.

Relatively little work has been done on the sensitivity of price dispersion to changes in

[2]These arguments are made in Nakamura and Steinsson (2008), Guimaraes and Sheedy (2011), Kehoe and Midrigan (2015), Anderson et al. (2015). See also Nakamura and Steinsson (2013) for a discussion of these ideas.

inflation in the United States. Reinsdorf (1994) uses BLS micro data for the period 1980-1982 (a subset of our data) and finds that price dispersion rose when inflation fell. Sheremirov (2015) uses scanner price data for the relatively low inflation period of 2002-2012. He finds that price dispersion rises with inflation. Alvarez et al. (2011) study the relationship between price dispersion and inflation during the Argentinian hyperinflation in 1989-1990. They find that the elasticity of price dispersion with inflation is roughly 1/3 at high inflation rates, in line with a simple menu cost model.

An earlier literature studied the relationship between inflation and the dispersion of sectoral inflation rates (Debelle and Lamont, 1997; Fischer, 1981; Glejser, 1965; Mills, 1927; Parks, 1978; Vining and Elwertowski, 1976). However, such measures are very sensitive to sectoral shocks such as oil price shocks (Bomberger and Makinen, 1993). Van Hoomissen (1988) and Lach and Tsiddon (1992) study the the relationship between inflation and the dispersion of price changes within sector during high inflation periods in Israel in the 1970's and early 1980's. These papers argue that menu cost models yields similar implications for "relative price variability" (the dispersion in product-level inflation rates) as for price dispersion itself. This is not, however, the case in the models we study. On a related note, Vavra (2014) studies the cyclical properties of the dispersion of price changes.

The paper proceeds as follows. Section 1.2 discusses the welfare loss resulting from inflation in different models with price rigidity. Section 1.3 describes the construction of our new micro-dataset on consumer prices. Section 1.4 presents evidence based on this data that inefficient price dispersion was no higher when inflation was high in the late 1970's and early 1980's than it has been since then. Section 1.5 discussed the evolution of the frequency of price change over our sample period. Section 1.6 discusses our results on the evolution of price flexibility. Section 1.7 concludes.

1.2 Costs of Inflation in Sticky Price Models

To understand what drives the large costs of inflation in standard sticky price models, it is useful to lay out a simple model of the type used in the literature. The model economy is populated by households, firms, and a government. Consider first the households. There are a continuum of identical households that seek to maximize discounted expected utility given by

$$E_t \sum_{j=0}^{\infty} \beta^j \left[\log C_{t+j} - L_{t+j}\right], \tag{1.1}$$

where E_t denotes the expectations operator conditional on information known at time t, C_t denotes household consumption of a composite consumption good, and L_t denotes household supply of labor. Households discount future utility by a factor β per period. The composite consumption good C_t is an index of household consumption of individual goods produced in the economy given by

$$C_t = \left[\int_0^1 c_{it}^{\frac{\theta-1}{\theta}} di\right]^{\frac{\theta}{\theta-1}}, \tag{1.2}$$

where c_{it} denotes consumption of individual product i. The parameter $\theta > 1$ denotes the elasticity of substitution between different individual products.

Households earn income from two sources: their labor and ownership of the firms in the economy. The household's budget constraint is therefore

$$P_t C_t + Q_{it} B_{it} \leq W_t L_t + (D_{it} + Q_{it}) B_{it-1}, \tag{1.3}$$

where P_t denotes the price of the final good, W_t denotes the wage rate, D_{it}, Q_{it}, and B_{it} denote the dividend, price, and quantity purchased and sold of asset i. The assets in the economy include ownership claims to the firms in the economy and may include other assets such as a risk-free nominal bond and Arrow securities although these will not play any role in our analysis. To rule out "Ponzi schemes", we assume that household financial wealth must always be large enough that future income suffices to avert default.

Households take the prices of the individual goods p_{it} as given and optimally choose to minimize the cost of attaining the level of consumption C_t. This implies that their demand for individual product i is given by

$$c_{it} = \left(\frac{p_{it}}{P_t}\right)^{-\theta} C_t, \tag{1.4}$$

where

$$P_t = \left[\int_0^1 p_{it}^{1-\theta} di\right]^{\frac{1}{1-\theta}} \tag{1.5}$$

is the cost-of-living price index.

Optimal choice of labor by the household taking the wage W_t as given yields a labor supply equation

$$\frac{W_t}{P_t} = C_t. \tag{1.6}$$

Household optimization also yields expressions for the household's valuation of all assets that exist in the economy. For the purpose of calculating the equilibrium in our model, it will be useful to have an expression for the household's valuation at time t of an uncertain dividend payment from firm i at time $t + j$, i.e., a j-period "dividend strip" for firm i. Lets denote the value of this dividend strip as V_{it}^j. Its value is

$$V_{it}^j = E_t \left[\beta^j \left(\frac{C_{t+j}}{C_t} \right)^{-1} D_{t+j} \right]. \tag{1.7}$$

Other conditions for household optimization do not play a role in determining the equilibrium.

There exist a continuum of firms in the economy that each produce a distinct individual product using the production function

$$y_{it} = A_{it} L_{it}. \tag{1.8}$$

Here A_{it} denotes the productivity level of firm i and L_{it} is the amount of labor demanded by firm i. The logarithm of firm productivity varies over time according to the following AR(1) process

$$\log A_{it} = \rho \log A_{it-1} + \epsilon_t, \tag{1.9}$$

where $\epsilon_t \sim \mathrm{N}(0, \sigma_\epsilon^2)$ are independent over time and across firms.

Firms commit to meet demand for their products at the price they post. They hire labor on the economy-wide labor market at wage rate W_t in order to satisfy demand. The marginal cost of firm i is $MC_{it} = W_t / A_{it}$. The firms are monopoly suppliers of the goods they produce. Their main decision is how to price these products. We assume that changing prices is costly and consider several different assumptions about these costs below.

Finally, to keep our model as simple as possible so that we can focus on the driving forces underlying costs of inflation in a sticky price setting, we assume that the monetary authority is able to control nominal output $S_t = P_t C_t$. Specifically, the monetary authority acts so as to make nominal output follow a random walk with drift in logs:

$$\log S_t = \mu + \log S_{t-1} + \eta_t \tag{1.10}$$

where $\eta_t \sim \mathrm{N}(0, \sigma_\eta^2)$ are independent over time. We will refer to S_t either as nominal output or as nominal aggregate demand.

1.2.1 The Flexible Price Benchmark

Let's begin by considering the equilibrium of this economy when prices are completely flexible. In this case, the firms will set the price of the good they produce equal to a markup over marginal cost

$$p_{it} = \frac{\theta}{\theta - 1} \frac{W_t}{A_{it}}, \tag{1.11}$$

This price setting equation can then be used to show that at the aggregate level

$$P_t = \frac{\theta}{\theta - 1} \frac{W_t}{A_f}, \tag{1.12}$$

where

$$A_f = \left[\int_0^1 A_{it}^{\theta - 1} di \right]^{\frac{1}{\theta - 1}}. \tag{1.13}$$

We refer to A_f as flexible price aggregate labor productivity.[3] Using firm i's production function—equation (1.8)—the demand curve for firm i's output—equation (1.4)—the price

[3] Since we abstract from aggregate productivity shocks, if the cross-sectional distribution of idiosyncratic firm productivity A_{it} starts off at its ergodic distribution, the integral on the right-hand-side of equation (1.13) will remain constant.

setting equation for firm i—equation (1.11)—and integrating over i yields the following aggregate production function

$$Y_t = A_f L_t. \tag{1.14}$$

where Y_t denotes aggregate output (which is equal to aggregate consumption $Y_t = C_t$). See Appendix A.1 for derivations of equations (1.11)-(1.14).

We can now see that in this flexible price version of our model, the equilibrium value of output, labor, and real wages is determined by the following three simple equations:

$$\text{Labor Supply:} \qquad \frac{W_t}{P_t} = Y_t \tag{1.15a}$$

$$\text{Production Function:} \qquad Y_t = A_f L_t \tag{1.15b}$$

$$\text{Markup:} \qquad P_t = \Omega_f \frac{W_t}{A_f}. \tag{1.15c}$$

where $\Omega_f = \theta/(\theta - 1)$ denotes the amount by which firms choose to markup their products' prices over marginal costs when prices are fully flexible. Notice that these three equations only determine W_t/P_t, not the level of P_t and W_t individually. To pin down the level of nominal prices and wages, one must add $S_t = P_t Y_t$ to the system.

Using equations (1.15a)-(1.15c) to solve for output and labor supply yields

$$Y_t = \Omega_f^{-1} A_f \tag{1.16a}$$

$$L_t = \Omega_f^{-1}. \tag{1.16b}$$

Notice that this solution is independent of the rate of inflation and also independent of the history of shocks to nominal aggregate demand. The only distortion that moves this economy away from a first-best outcome is the monopoly power of the firms. This distortion leads the firms to set prices above marginal costs. As a consequence, output is inefficiently low.[4]

1.2.2 Equilibrium with Sticky Prices

When prices are sticky, the determination of equilibrium is more complicated and will depend on the exact nature of the price adjustment costs. We will consider several different assumptions about the nature of price adjustment costs including the case of a constant fixed cost (the menu cost model) and a case where the cost is zero with some probability in each period and infinite with a complementary probability (the Calvo (1983) model). We assume that firms maximize the value of their stochastic stream of dividends. The methods we use to solve for the equilibrium in these models are those described in detail in Nakamura and Steinsson (2010).

To help build understanding about the costs of inflation that result from sticky prices, it is useful to compare the equilibrium in the sticky price case with the flexible price equilibrium.[5] To this end, we consider an analogous set of equations to equations (1.15) for the sticky price

[4]In a fully competitive version of the economy described above—i.e., one in which the markets for all the goods are competitive—all prices would be set equal to marginal cost and the markup Ω_f would therefore be one. Output would therefore be higher. We know from the first welfare theorem that this is the efficient level of output.

[5]This exposition builds on related analysis in Blanco (2015a) as well as earlier work in Burstein and Hellwig (2009).

case:

$$\text{Labor Supply:} \qquad \frac{W_t}{P_t} = Y_t \tag{1.17a}$$

$$\text{Production Function:} \qquad Y_t = A_t(\bar{\pi})(L_t - L_t^{pc}) \tag{1.17b}$$

$$\text{Price Setting:} \qquad P_t = \Omega_t(\bar{\pi})\frac{W_t}{A_t(\bar{\pi})}. \tag{1.17c}$$

The same labor supply equation continues to hold in the sticky price case. The aggregate production function—equation (1.17b)—however, differs from its flexible price analog—equation (1.15b)—in two ways. First, some labor is needed to change prices and does not produce output. We use L_t^{pc} to denote this extra labor. This is one source of costs of price rigidity. Second, aggregate labor productivity is lower in the sticky price economy than under flexible prices because in the sticky price model the relative prices of different goods do not accurately reflect the goods' relative marginal cost of production. In Appendix A.1 we show that the value of aggregate labor productivity when prices are sticky is

$$A_t(\bar{\pi}) = \left[\int_0^1 \left(\frac{p_{it}}{P_t}\right)^{-\theta} A_{it}^{-1} di\right]^{-1}. \tag{1.18}$$

Note that aggregate labor productivity is not only a function of the physical productivity of the individual firms, but also a function of the relative prices of the goods they sell. This occurs because we use a utility based definition of aggregate output (see equation (1.2) and note that $Y_t = C_t$) and the marginal utility households derive from each individual product falls as they consume more of that product relative to other products. Consider for simplicity a case where all products have the same physical productivity. In this case, products that have low relative prices—and are therefore consumed in greater quantities—will contribute less to aggregate output on the margin than products with high relative prices unless. This will lower aggregate productivity of labor.[6] More generally, the more intensive consumption of the low priced goods will lower aggregate productivity of labor, unless their lower relative prices are offset by higher physical productivity.

When prices are sticky, the relative price of a particular product drifts downward as time passes. This is one source of divergence between relative prices and relative productivity that results in lower aggregate labor productivity. This drift is more pronounced the higher is the level of inflation. For this reason, aggregate labor productivity is a decreasing function of the average level of inflation. To emphasize this, we make explicit the dependence of A_t on $\bar{\pi}$ by writing $A_t(\bar{\pi})$.

Equation (1.17c) is most usefully thought of as defining $\Omega_t(\bar{\pi})$. We will refer to $\Omega_t(\bar{\pi})$ as the aggregate markup in the sticky price case. Variation in $\Omega_t(\bar{\pi})$ reflects the degree to which the price level rises more or less rapidly than $A_t(\bar{\pi})$ falls as inflation changes.[7]

Manipulating equations (1.17) yields

$$Y_t = \Omega_t(\bar{\pi})^{-1} A_t(\bar{\pi}) \tag{1.19a}$$

$$L_t = \Omega_t(\bar{\pi})^{-1} + L_t^{pc} \tag{1.19b}$$

This shows that output under sticky prices will differ from its level under flexible prices for

[6] In the special case of equal idiosyncratic productivity, the variance of prices is a second order approximation for labor productivity.

[7] It should be noted that $\Omega_t(\bar{\pi})$ does not measure how the average markup of firms over *physical* marginal costs change as inflation changes since $A_t(\bar{\pi})$ is not a measure of physical productivity but rather is also affected by the distribution of relative prices.

two reasons: 1) aggregate labor productivity will be lower, 2) the aggregate markup may be different; and aggregate labor supply will also differ from its level under flexible prices for two reasons: 1) Some labor is needed to change prices, 2) the aggregate markup may be different.

Welfare in the economy, in turn, depends on output and labor through equation (1.1). As is common in the literature, we will report welfare differences across models and levels of inflation in terms of consumption equivalent welfare changes. I.e., when comparing welfare in model economy A with welfare in model economy B we will solve for the value of Λ that yields

$$E\left[\log\left((1+\Lambda)C_t^A\right) - L^A\right] = E\left[\log\left(C_t^B\right) - L^B\right]. \tag{1.20}$$

The value Λ then measures the percentage change in consumption needed to make households in economy A equally well off as households in economy B.

1.2.3 Model Calibration

Below we calculate equlibrium outcomes for a menu cost model and a Calvo model and compare them to the flexible price benchmark. These models are calibrated as follows. A unit of time is meant to correspond to a month. We set the subjective discount factor to $\beta = 0.96^{1/12}$. The baseline value that we use for the elasticity of substitution between intermediate goods is $\theta = 4$. This value is roughly in line with estimates of the elasticity of demand for individual products in the industrial organization and international trade literatures (Berry et al., 1995; Broda and Weinstein, 2006; Nevo, 2001). This is, however, at the low end of values for θ that have been used in the macroeconomics literature on the welfare costs of inflation. We will also present results for $\theta = 7$, which is the values used by Coibion et al. (2012).

In the menu cost model, we calibrate the level of the menu cost and the standard deviation of the idiosyncratic shocks to match the median frequency of price change of 10.1% per month and the median absolute size of price changes of 7.5% over the (relatively low inflation) period 1988-2014.[8] The resulting parameter values are $K = 0.019$ for the menu cost and $\sigma_\epsilon = 0.037$ for the standard deviation of the idiosyncratic shocks.[9] In the Calvo model, we set the frequency of price change equal to the median frequency of price change in the data and the standard deviation of the idiosyncratic shocks to the same value as in the menu cost model. In both models, we assume that the first-order autoregressive parameter in the process for idiosyncratic productivity is $\rho = 0.7$, the same value as in Nakamura and Steinsson (2010).

We calibrate the standard deviation of shocks to nominal aggregate demand to be $\sigma_\eta = 0.0039$ based on the standard deviation of changes in U.S. nominal GDP over the period 1988-2014. We present results for a range of values of average change in nominal aggregate demand (i.e., a range of values for average inflation).

1.2.4 Numerical Results on the Costs of Inflation

Figure 1.2 plots the consumption equivalent welfare loss experienced by households when prices are sticky as a function of the inflation rate. The welfare loss is calculated relative to welfare in an economy with flexible prices. The difference in results between the menu

[8] More specifically, we first calculate the average frequency and absolute size of price changes within ELIs (see section 1.3 for a discussion of what an ELI is) in each year. We then take a median across ELIs in each year. We then take an average of these medians over years.

[9] This menu cost implies that 0.019 units of labor are needed to change a price. For comparison, the non-stochastic steady state level of labor each month is $\Omega_f^{-1} = 0.75$.

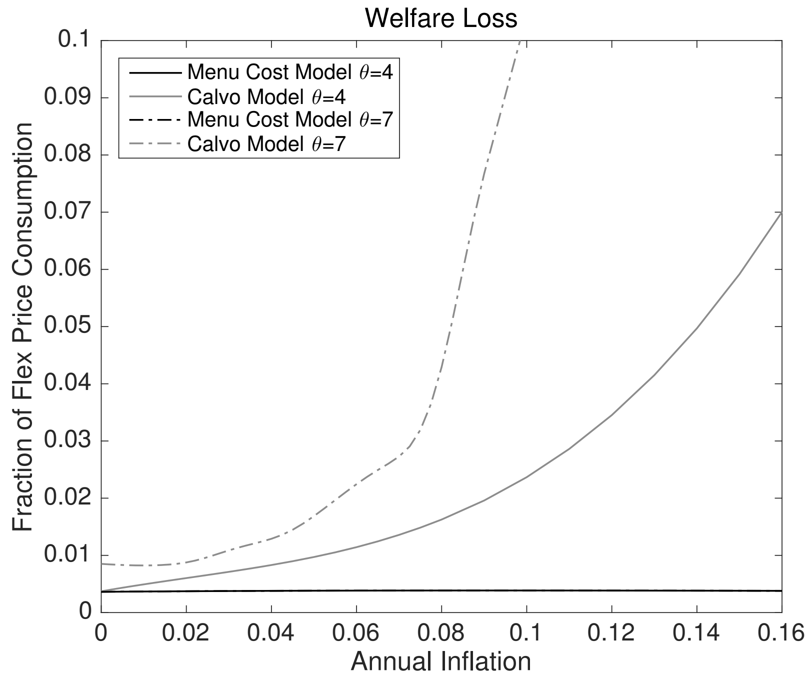

Figure 1.2: Welfare Loss

Note: The figure plots the consumption equivalent loss of welfare in each model as a function of the inflation rate relative to welfare when prices are completely flexible.

cost model and the Calvo model is striking. For the menu cost model, the welfare loss is small and virtually completely constant at about 0.5% as a function of inflation. This is true both when $\theta = 4$ and $\theta = 7$. For the Calvo model, however, the costs of price rigidity rise sharply with inflation. Consider first our baseline case of $\theta = 4$. When inflation is zero, these costs are similar in magnitude to those in the menu cost model. When inflation is 10% per year, these costs have risen to 2.4%; and when inflation is 16% per year, these costs are a staggering 7.7% per year. When $\theta = 7$, these welfare losses rise even faster. In this case, the welfare loss hits 10% when inflation is roughly 10%. Clearly the exact nature of price rigidities matter a great deal when assessing the costs of inflation.

To gain further insight, Figure 1.3 plots the level of output and labor supply in the menu cost model and Calvo model as a function of inflation (again relative to the level of these

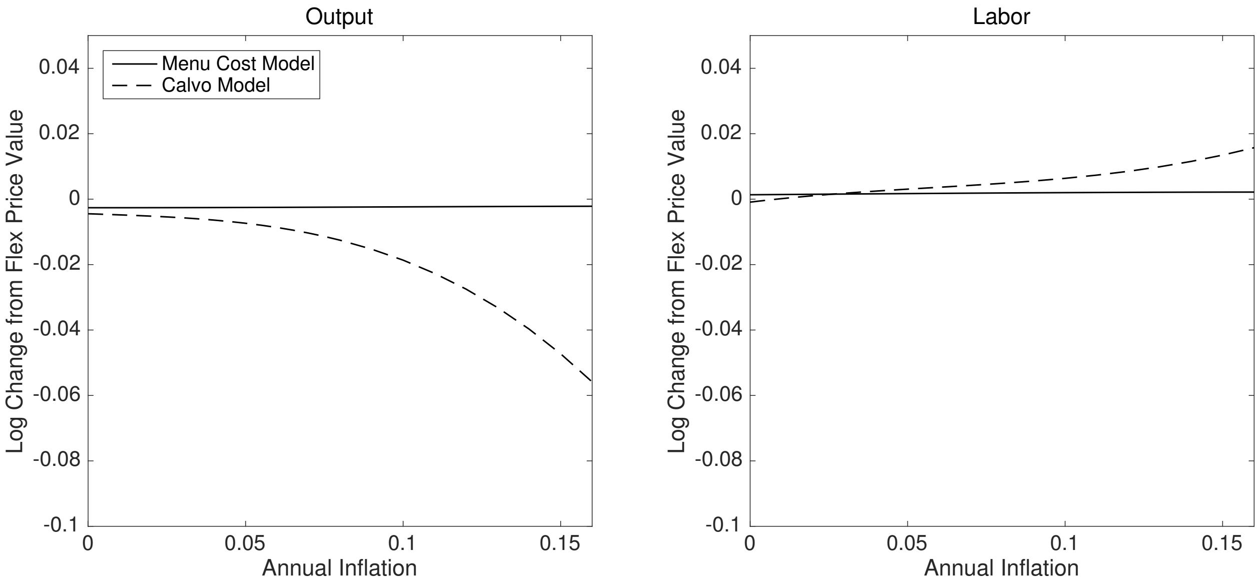

Figure 1.3: Output and Labor Supply
Note: The left panel plots average output and the right panel plots average labor supply. In each case, these variables are plotted as a function of inflation relative to their level when prices are flexible.

variables when prices are flexible). This figure shows that in the Calvo model, an increase in inflation from 0% to 10% leads to a fall in output of about 1.5%. But the amount of labor needed to produce this lower output is actually greater by 0.7% due to a fall in labor productivity. As with welfare, these changes in output and labor supply grow increasingly rapidly as inflation rises above 10%.

Figure 1.4 plots labor productivity directly as well as the aggregate markup (again relative to the level of these variables when prices are flexible). From this figure we see that the welfare loss that results from a higher rate of inflation in the Calvo model come entirely from a sharp fall in labor productivity. Labor productivity falls by 2.1% when inflation rises from 0% to 10%. The other two potential sources of welfare losses discussed in section 1.2.2 are non-existent or actually increase welfare in the Calvo model. First, in the Calvo model, firms face no direct costs when they change their prices. Second, the aggregate markup actually falls

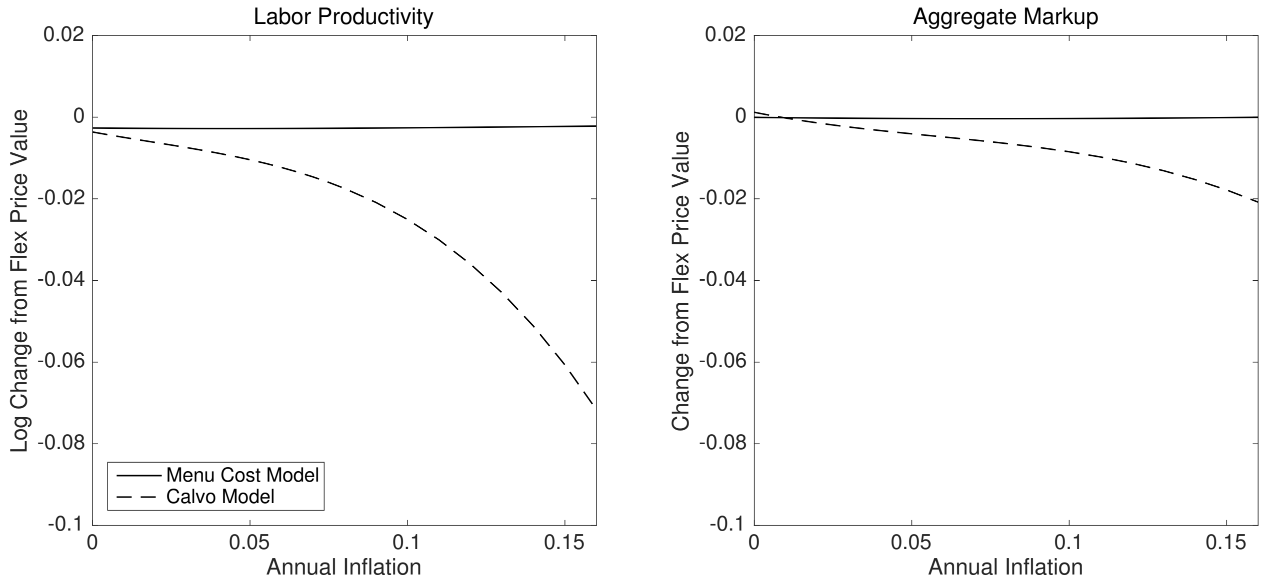

Figure 1.4: Labor Productivity and Markup
Note: The left panel plots average labor productivity and the right panel plots the aggregate markup. In each case, these variables are plotted as a function of inflation relative to their level when prices are flexible.

when inflation rises (by roughly 0.9 percentage points, when inflation rises from 0% to 10%). In other words, the price level relative to the wage rate does not rise quite as rapidly as labor productivity falls implying the output does not fall as rapidly as labor productivity.[10]

At an intuitive level, the loss of labor productivity in the Calvo model when inflation rises is due to an increase in inefficient price dispersion. In fact, these two concepts are equal up to a second order approximation when we abstract from idiosyncratic productivity shocks. To drive home this point, Figure 1.5 plots inefficient price dispersion in the Calvo model and the menu cost model as a function of inflation. We see that the pattern for inefficient price dispersion is very similar to the pattern for labor productivity (and the overall welfare loss). This is useful in terms of providing us with guidance regarding what statistics we should

[10]In the menu cost model we analyze, the aggregate markup does not change much with inflation over the range we consider. Benabou (1992) studies a menu cost model with consumer search in which the higher price dispersion resulting from inflation leads to more search, which in turn makes markets more competitive an lowers markups.

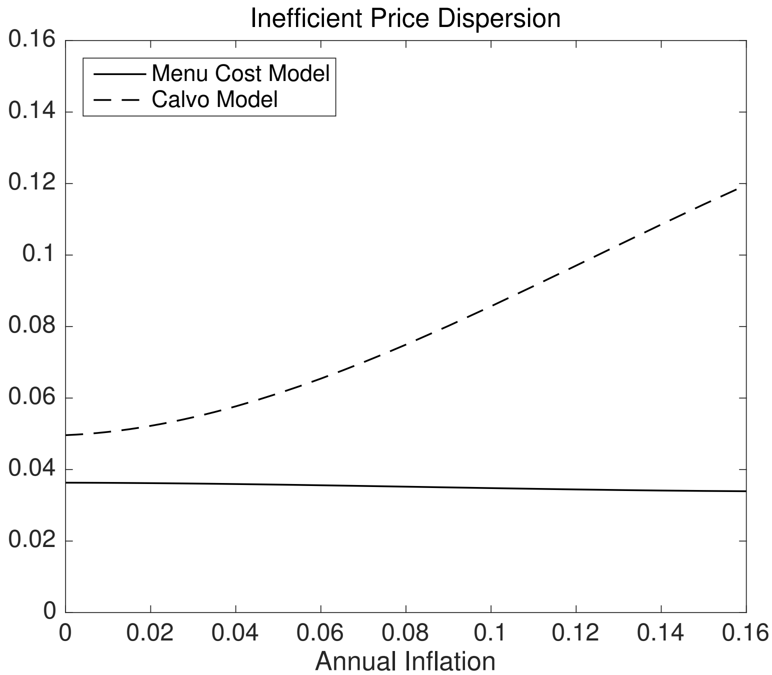

Figure 1.5: Inefficient Price Dispersion

Note: The measure of inefficient price dispersion that we plot here is the standard deviation across firms of the firm's price relative to the price it would charge if it had flexible prices.

calculate to shed empirical light on the costs of inflation.

1.3 New Micro-Data on Consumer Prices during the Great Inflation

Our analysis is based on a new dataset that we developed with the help and support of staff at the Bureau of Labor Statistics.[11] The dataset contains the individual price quotes underlying the U.S. Consumer Price Index for the period from May 1977 to October 1986 and May 1987 to December 1987. Prior to our project, the BLS CPI Research Database contained data starting only in 1988. It therefore had the important disadvantage that it did not cover the most eventful period in post-WWII U.S. monetary history: the "Great

[11] We owe a huge debt to Daniel Ginsberg, John Greenlees, Michael Horrigan, Robert McClelland, John Molino, Ted To and numerous others at the Bureau of Labor Statistics whose efforts were crucial in making this project possible.

Inflation" of the late 1970's and early 1980's.

The construction of the dataset involved two main phases. First, we worked with the BLS staff to scan the physical microfilm cartridges to convert them to digital images. This step was difficult because the microfilm cartridges were sufficiently old that modern scanners could not read them. Fortunately, we were able to find a company that was willing and able to retrofit a modern microfilm reader to read these outdated cartridges.

This original process of scanning the price cartridges left us with roughly 1 million images of "Price Trend Listings" that needed to be converted to machine-readable form. The BLS' high standards of confidentiality imply that all processing of the data must be done on-site at the BLS in Washington DC. This made it infeasible to outsource this step to a professional data-entry firm for manual data entry. The alternative available to use was to use optical character recognition (OCR) software for this conversion process. This was challenging because leading commercial software solutions turned out to be both prohibitively expensive and too slow. After considerable search, we, however, found a firm that was able to create custom software that ensured high quality and high enough speed to convert the large number of images we had.[12]

The raw microfilm cartridges we found at the BLS contain images of Price Trend Listings, starting in May 1977 and ending in October 1986. Each Price Trend Listing contains prices for the previous 12 months for a given product—a feature of the data that we make considerable use of in checking for errors, as we describe below. All cartridges from the period 1977-1981 we scanned, while cartridges from every other month for the period 1982-1986

[12]In overcoming these practical obstacles, we benefited greatly from Patrick's tireless work and ingenuity. The rest of us are very grateful for his efforts in this regard.

were scanned. This choice was motivated by the higher quality of the images on the more recent cartridges. It is possible that even older CPI micro-data exists at the BLS. However, the CPI underwent a major revision in 1978. We conjecture that data collection was revised as a part of this revision and the May 1977 start date of our data is the start date of the new data collection system put in place at this time. We obtained separate, already digitized data on prices for the months of May to December in 1987. This left us with a short gap in coverage for the period November 1986 to April 1987.

The information contained on the Price Trend Listing images includes 1) an internal BLS category label called an Entry Level Item or ELI, 2) a location (city) identifier, 3) an outlet identifier, 4) a product identifier, 5) the product's price, 6) the percentage change in the product's price between collection periods, 7) a "sales flag" indicating whether the product's price was temporarily marked down at the time of collection, 8) an "imputation flag" indicating whether the price listed is truly a collected price or was imputed by the BLS, and 9) several additional flags that we do not use. From this we see that each product in the dataset is identified at a very detailed level—for example, a 2-Liter Diet Coke at a particular Safeway store in Chicago.

BLS employees collect the data by visiting outlets on a monthly or in some cases bi-monthly basis. Somewhere between 80,000 and 100,000 observations are collected per month. Prices of all items are collected monthly in the three most populous locations (New York, Los Angeles, and Chicago). Prices of food and energy are collected monthly in all other locations as well. Prices of other items are collected bimonthly. We focus on the monthly data in our analysis.

Fortunately, there are numerous redundancies in the raw data that allow us to check for errors in our OCR procedure. The first form of redundancy arises from the fact that prices for a particular product in a particular month appear on multiple Price Trend Listings because the Price Trend Listings include 12 months of previous prices, as we note above. We can use this redundancy to verify that the price observations obtained from different Price Trend Listings are, in fact, the same. The second form of redundancy arises because each Price Trend Listing includes both the price and a percentage change variable. We can therefore verify that the percentage change in prices obtained when we calculate this directly based on the converted prices is the same as the one reported in the percentage change variable. We describe both of these procedures in detail in Appendix A.2. We err on the side of caution: all of the price observations included in our final dataset have been "accepted" by either of the procedures described above.

The order of the Price Trend Listings allows us to check for errors in the OCR conversion of the product label and category variables. Each microfilm cartridge corresponds to a particular "collection period" when the prices were collected. On each cartridge, images are sorted first by product category (ELI), then by outlet, then by quote (a specific product, such as a 2-Liter Diet Coke), and finally by the version (used when a product is replaced by other very similar product). The order of the Price Trend Listings means that if our OCR procedure fails to convert a particular ELI value, we can easily fill it in using the surrounding values of ELI's. We use a similar procedure to fill in missing values of the outlet, quote and version variables. The algorithm we use for this is described in Appendix A.2. Errors in converting the product identifiers will lead to a spuriously large number of products. Appendix A.2 discusses this in more detail and describes a procedure we use to verify that this does not bias our results.

Finally, our process for converting the sales and imputation flag variables makes use of the limited set of values taken by these values (e.g., "I" stands for imputation). We also make use of the fact that, like in the case of prices, the flags for a given product-month appear on multiple Price Trend Listings. We discuss these procedures in greater detail in Appendix A.2.

The BLS has changed the organization of the consumer price micro-data two times since 1978. The first change occurred in 1987 and the second, more substantial change, occurred in 1998. We have created a concordance to harmonize the ELI categories across these different time periods. To do this, we first used the category descriptions to match the ELI categories used in the 1977-1986 and the ones used in 1987-1997. We then used the descriptions available in the BLS' documentation for the CPI Research Database to match the ELI's for 1987-1997 and 1998 onwards. A detailed set of concordances is available on our websites.

We will hand over the new dataset we have constructed to the BLS so that they can make it available to researchers in the same way the existing BLS CPI Research Database is available. The data we will make available will include the original scanned images (in PDF format), the raw dataset that resulted from our OCR conversion (with all the redundancies discussed above), and the final dataset we constructed using the procedures discussed above and in Appendix A.2. The availability of all three of these versions of the data will therefore allow future researchers to improve on our data construction effort (both the OCR conversion and the accuracy verification of the OCR output).

In our analysis of this data, we drop all imputed prices. Whenever we observe a price change that is larger than one log point ($\log(p_t/p_{t-1}) > 1$), we set the price change variable and price change indicator to missing (i.e., we drop these large price changes). Only 0.04% of observations are dropped because they are large. We frequently calculate weighted means and medians of various statistics across ELIs by year. In all cases, we hold fixed the expenditure weights at their value in 2000. While the accuracy of our data conversion methods seems high for most of our sample, we are not fully confident in the quality of the resulting data at the very beginning of our sample. For this reason, we drop the data from 1977. Our sample period is therefore 1978 to 2014.

1.4 Price Dispersion and the Size of Price Changes

We have seen in section 1.2 that the costs of inflation in sticky price models are largely due to increases in inefficient price dispersion. Figure 1.6 plots the evolution of a simple measure of price dispersion for U.S. consumer prices over the period 1978-2014. We first calculate the interquartile range of prices within Entry Level Items (ELI's)—narrow product categories defined by the BLS such as "salad dressing"—for each year. We then and then take the expenditure weighted median across ELI's. We calculate this measure of price dispersion both including and excluding temporary sales.

Figure 1.6 shows that this simple measure of price dispersion has increased steadily over the past 40 years. This is driven by dramatic increases in price dispersion within ELI for unprocessed food, processed food, and travel services. This increase in dispersion is, of course, the opposite result from what one might have expected given that inflation has fallen sharply over this period (see Figure 1.1). As we discuss in the introduction, a key empirical challenge is that much of the cross-sectional dispersion in prices—even within narrowly defined product categories—likely results from heterogeneity in product size and

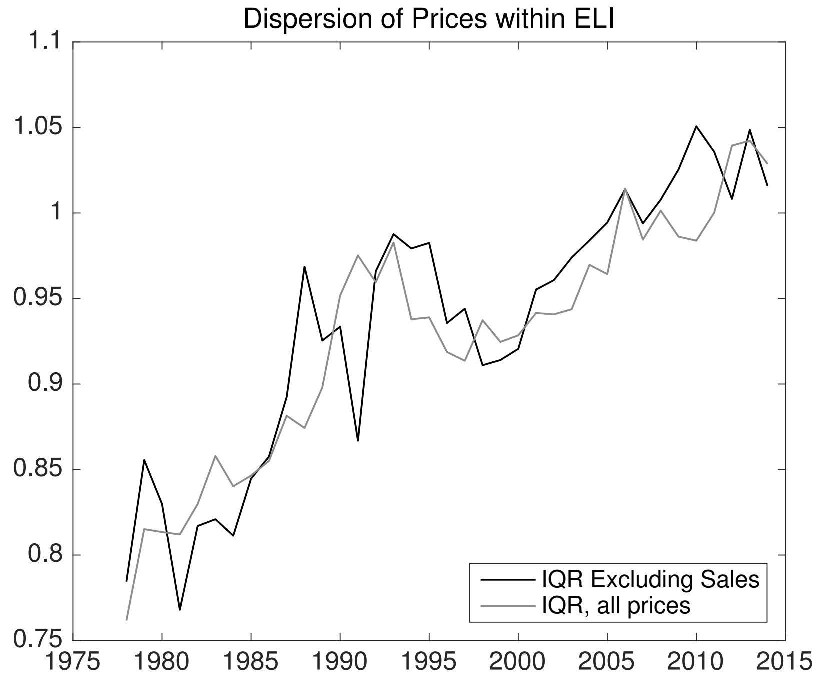

Figure 1.6: Dispersion of Log Prices within ELI

Note: To construct the series plotted in this figure, we first calculate the interquartile range of log prices in each ELI for each year. We then take the weighted median across ELI's. We do this both for prices including and excluding temporary sales.

quality. Time variation in the cross-sectional dispersion in prices may therefore come from time variation in product heterogeneity as opposed to time variation in inefficient price dispersion. The pattern revealed in Figure 1.6 suggests that a large increase in product variety over the past 40 year for example among food products has lead to an increase in price dispersion within ELI that was large enough to dwarf any (relatively small) changes in inefficient price dispersion associated with price rigidity.

We therefore consider next a gauge of the extent of inefficient price dispersion that "differences out" fixed product characteristics. We focus on the absolute size of price changes. The basic intuition for this measure is that, if higher inflation truly leads prices to drift further from their efficient levels due to price rigidity, we should observe larger price changes when firms finally have an opportunity to adjust.

Figure 1.7 illustrates this by comparing the relationship between inflation and the average absolute size of price changes in the Calvo model and in the menu cost model. We do this in two ways. First, we plot the steady state average absolute size of price changes for different steady state values of inflation. These are the two lines in the figure. For the Calvo model, the average absolute size of price changes rises sharply with inflation, while this is not the case in the menu cost model. Recall that this is exactly the pattern that holds for inefficient price dispersion. Studying the absolute size of price changes thus provides an indirect, yet powerful, way to measure the extent to which inefficient price dispersion rises with inflation.

A concern with the steady state calculation that underlies the two lines in Figure 1.7 is that the Great Inflation was a somewhat transitory event. Perhaps inflation was not high for long enough during the Great Inflation to create the degree of price dispersion needed to substantially raise the average absolute size of price changes. We can address this concern by simulating the response of the average absolute size of price change to the actual evolution of inflation in the U.S. in the two models over our sample period. These are the two sets of points in the figure. Each point gives the average absolute size of price changes and the average inflation in a particular year between 1978 and 2014.[13] This exercise, while somewhat noisier, gives results very similar to the steady state calculation discussed above.

Figure 1.8 plots the evolution of the absolute size of price changes over the period 1978-2014. We first calculate the mean absolute size of price changes within ELI by year and then take an expenditure weighted median across ELIs for each year. We again report results

[13]The simulations are done at a monthly frequency and the results then time-averaged to annual observations. We start the simulation in January 1960 to make sure that the distribution of relative prices in 1978 reflects the actual U.S. inflation history leading up to that point. In these simulations, we assume for simplicity that the real wage is constant and feed in the observed inflation rate. The perceived law of motion for the price level is a random walk with drift. We calibrate the perceived average drift and perceived standard deviation of monthly changes in the price level to their sample analogies over the period 1960 to 2014.

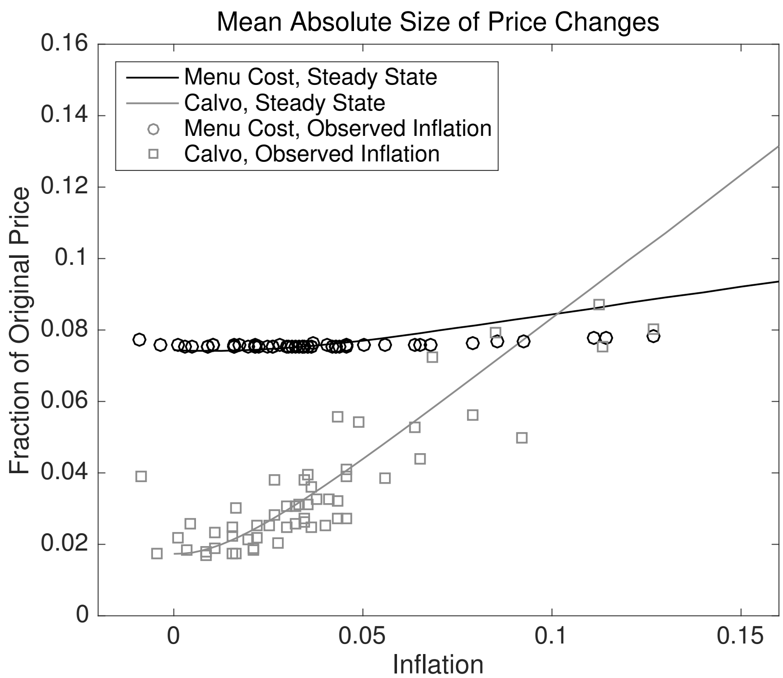

Figure 1.7: Absolute Size of Price Changes in Sticky Price Models

Note: The lines plot the mean absolute size of price changes as we vary the steady state level of inflation in the menu cost and Calvo models. The circles and squares plot the mean absolute size of price changes in the menu cost model and Calvo model, respectively, from a monthly simulation using the actual inflation rate from 1978 to 2014. Each point is the average from a particular year in the simulation.

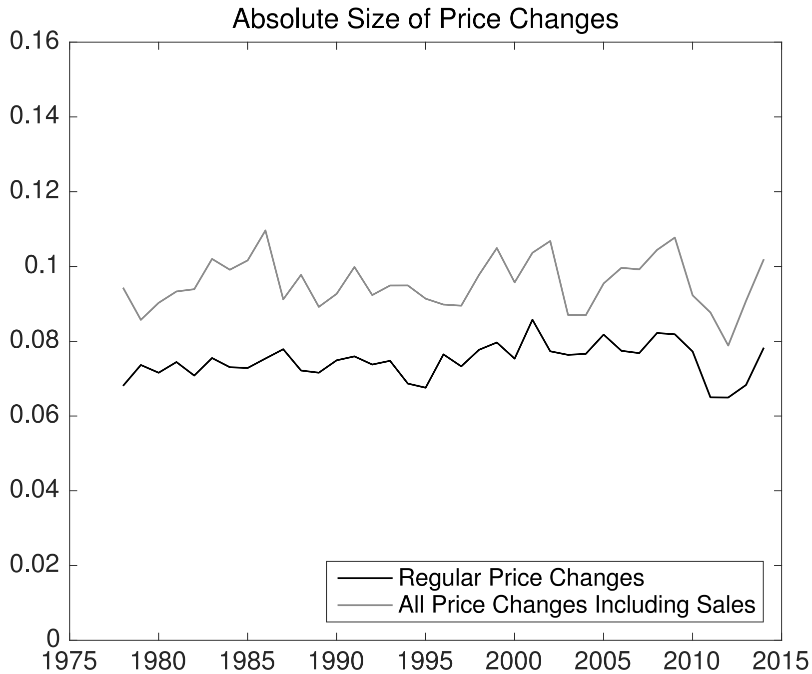

Figure 1.8: Absolute Size of Price Changes in U.S. Data

Note: To construct the series plotted in this figure, we first calculate the mean absolute size of price changes in each ELI for each year. We then take the weighted median across ELI's. We do this both for prices including and excluding temporary sales.

both including and excluding temporary sales. Even though the inflation rate has fallen sharply over our sample period, the absolute size of price changes has remained essentially unchanged over this time period at roughly 8%. If anything, there is actually a slight upward trend over the sample period. The evolution of the absolute size of price changes, therefore, provides no evidence that prices strayed farther from their efficient levels during the Great Inflation than in the low-inflation Greenspan-Bernanke period.[14]

The welfare costs of price rigidity depends non-linearly on the extent to which individual prices differ from their efficient level. Prices that are very far from optimal contribute disproportionately to welfare losses. The random nature of the timing of price adjustment in the Calvo model implies that as inflation rises the distribution of relative prices becomes

[14]Figure A.1 in the appendix shows that the average size is flat (or slightly upward sloping) within sector for six of the most important sectors in our data. Figure A.2 shows that there is nothing special about the median across ELIs. The 10th, 25th, 75th and 90th quantiles tell essentially the same story.

highly dispersed. The distribution of relative prices has a long left tail with some firms having wildly inappropriate prices because they have not been able to change their price for a long period. When these prices finally change, they change by large amounts. In contrast, in the menu cost model, the absolute size of firms' price changes is small because all price changes are clustered around the sS bounds.

Conditional on the average absolute size of price changes, the standard deviation of the absolute size of price changes provides information about the dispersion of the distribution of relative price changes and, in particular, information about how many firms have wildly inappropriate prices. Figure 1.9 plots the standard deviation of the absolute size of price changes as a function of inflation in the Calvo model and in the menu cost model (using the same two methods as we do for the mean absolute size of price changes in Figure 1.7). We see that the standard deviation of the absolute size of price changes rises by more than a factor of six in the Calvo model, as the inflation rate rises from 0 to 16% per year. In contrast, this measure rises only slightly in the menu cost model.

Figure 1.10 reports the standard deviation of the absolute size of price changes in the data over our sample period of 1978-2014. Like with the average absolute size statistic, we first calculate the standard deviation of the absolute size of price changes within ELI for each year and then take a weighted median across ELI's for each year. This statistic turns out to be quite stable over time in the data. It varies between roughly 4.5% and 6% for the majority of the sample period. If anything, it displays a slightly upward trend. This statistic, again, provides us with no evidence that the distribution of relative prices became more dispersed during the Great Inflation.

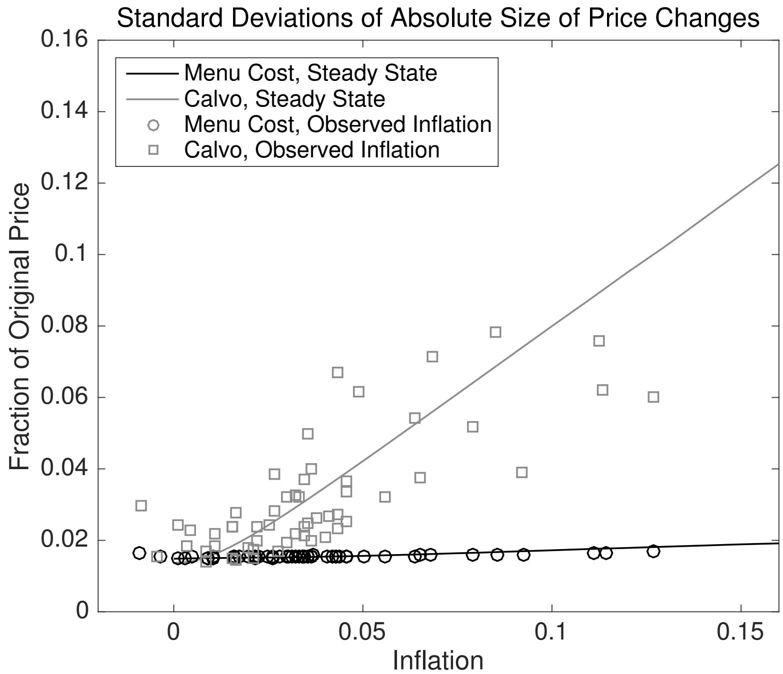

Figure 1.9: Standard Deviation of Absolute Size of Price Changes

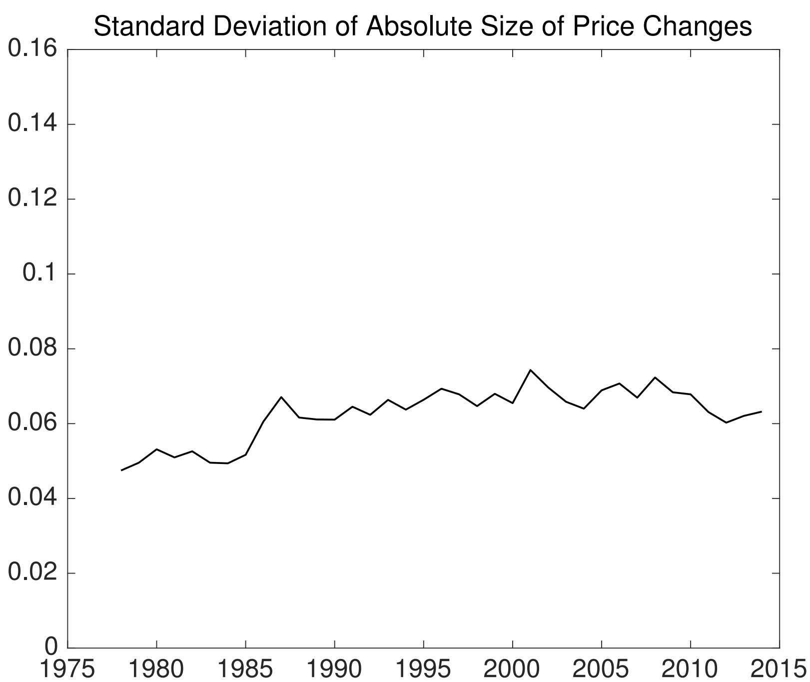

Figure 1.10: Standard Deviation of Absolute Size of Price Changes: Data

Note: To construct the series plotted in this figure, we first calculate the standard deviation of the absolute size of price changes in each ELI for each year. We then take the weighted median across ELI's.

1.5 Frequency of Price Change

Thus far in the paper, we have focused on the size of price changes because of its relation to price dispersion and welfare. The frequency of price change is in some sense the flip side of the coin. If inflation rises but the size of price changes don't, the frequency of price change must be changing.[15] Earlier research by Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008) has studied the time series behavior of the frequency of price change and its relationship to inflation in the U.S. using data on prices since 1988. The inference one can draw from these papers is, however, limited by the fact that inflation in the U.S. has been low and stable over the post-1988 period. The data from the Great Inflation that we analyze in this paper have much more power to distinguish among different pricing theories.[16]

Figure 1.11 plots the relationship between the frequency of price change and inflation in the menu cost model and the Calvo model. As with the size statistics discussed in section 1.4, we calculate this relationship for different steady state levels of inflation (the lines in the figure) and for a simulation based on the actual history of inflation in the U.S. over the period 1978 to 2014 (the points in the figure). Not surprisingly, the menu cost model implies that the frequency of price change rises with inflation. When inflation rises from 0% to 16%, the frequency of price change rises by more than half, going from 10% to 16%. In contrast, the frequency of price change is constant in the Calvo model by assumption.

Figure 1.12 plots the frequency of price change for consumer prices in the U.S. over our

[15] A subtlety here is that at low levels of inflation the frequency and absolute size of price changes can be relatively constant as inflation rises if the fraction of price changes that are increases is rising with inflation. In this case, the behavior of the average size and the average absolute size can be quite different. At higher levels of inflation, most price changes are increases and this distinction is less important (Gagnon, 2009).

[16] Important evidence on this topic is also available from a number of other (mostly middle income) countries with more volatile inflation rates. See, in particular, Gagnon (2009) and Alvarez et al. (2011) for evidence from Mexico and Argentina, respectively, and Wulfsberg (2015) for evidence from Norway. Nakamura and Steinsson (2013) discuss this literature in more detail.

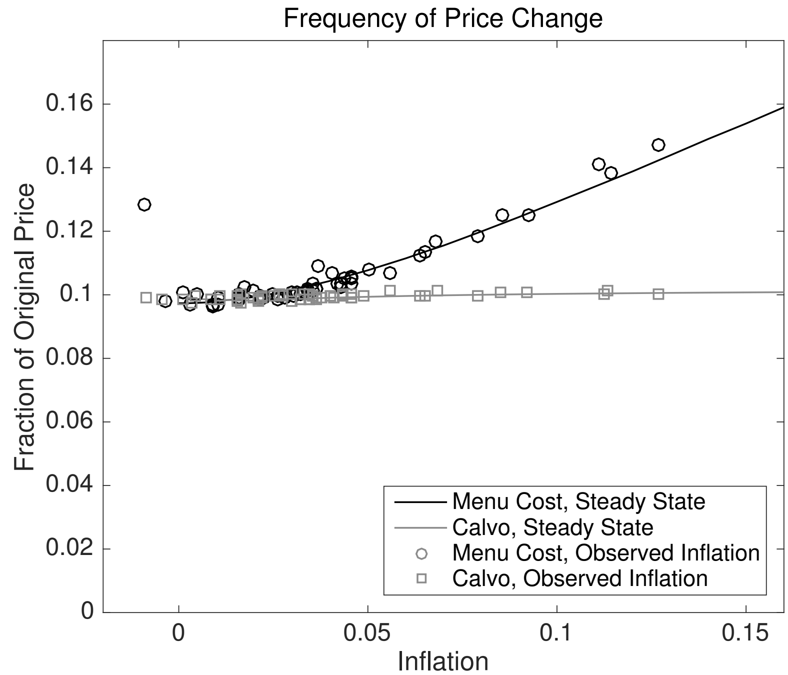

Figure 1.11: Frequency of Price Changes in Sticky Price Models
Note: The lines plot the frequency of price change as we vary the steady state level of inflation in the menu cost and Calvo models. The circles and squares plot the frequency of price change in the menu cost model and Calvo model, respectively, from a monthly simulation using the actual inflation rate from 1978 to 2014. Each point is the average from a particular year in the simulation.

sample period of 1978-2014 along with the CPI inflation rate. To construct this series, we first calculate the mean frequency of price change by ELI for each year. We then take an expenditure weighted median across ELIs for each year. The figure clearly shows that the frequency of price change comoves strongly with inflation. This data therefore strongly favors the menu cost model over the Calvo model.

Figure 1.13 separates the frequency of price increases and the frequency of price decreases. Here we plot the 12 month average frequency of price change at a quarterly frequency to see a bit more detail. The most striking feature of this figure is that it is the frequency of price increases that varies with the inflation rate, while the frequency of price decreases is unresponsive. Nakamura and Steinsson (2008) show that this asymmetry arises naturally in the menu cost model when prices are drifting upward due to a positive average inflation rate.

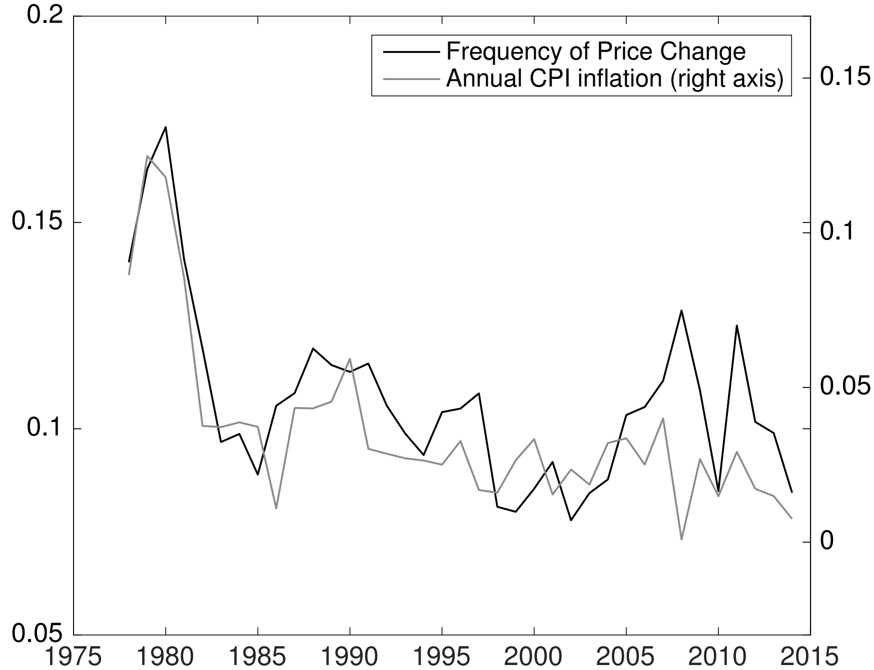

Figure 1.12: Frequency of Price Changes in U.S. Data
Note: To construct the frequency series plotted in this figure, we first calculate the mean frequency of price changes in each ELI for each year. We then take the weighted median across ELI's.

In this case, prices tend to "bunch" toward the bottom of their inaction region. Because of this bunching, when there is an aggregate shock that changes desired prices, there is a large response of the frequency of price increases (reflecting the relatively large mass at the bottom of the band), but a much smaller response of the frequency of price decreases. This is the same argument as the one described by Foote (1998) for why job destruction will be more volatile than job creation in declining industries.[17]

One curious feature of Figure 1.12 is the spike in the frequency of price changes that occurs in 2008. Looking at Figure 1.13 and especially the analogous plot for food in Figure A.3 in the appendix, we see however, that inflation was highly volatile in 2008. It first spiked up due to the commodity price boom early in that year, and then fell dramatically with the

[17]Figure A.3 presents figures analogous to Figure 1.13 for two important sectors in our data: food and services. In this figure, the inflation rate that we plot on each panel is the sectoral inflation rate in that sector. In both sectors, the frequency of price increases covaries strongly with inflation, while the frequency of price decreases is largely flat.

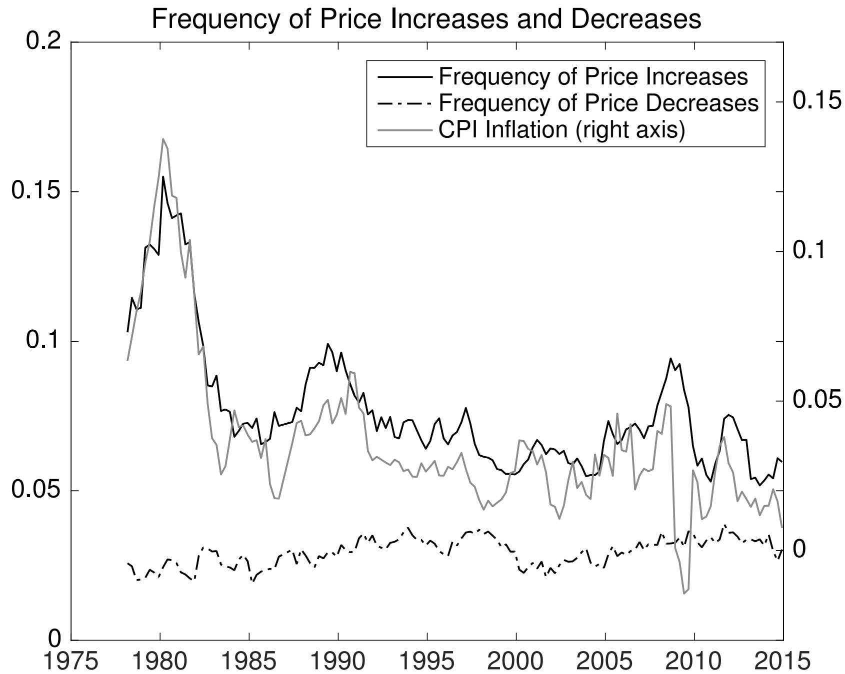

Figure 1.13: Monthly Frequency of Price Change

Note: To construct the frequency series plotted in this figure, we first calculate the mean frequency of price increases and decreases in each ELI for each month. We then take the weighted median across ELI's.

onset of the recession and the collapse of commodity prices. In light of this unusual volatility of inflation, the spike in the frequency of price changes in 2008 seems less puzzling.

1.6 Have Prices Become More Flexible over Four Decades?

The "menu cost" in the menu cost model is best thought of as a stand-in for a variety of costs associated with price adjustment. Though economists have failed to settle on what exactly this menu cost represents, many theories have been considered, including adverse customer reactions to price changes, limited managerial attention, and the actual costs of changing price tags or reprinting menus. Given all of the technological advancement that has occurred over the past half-century, it seems natural to conjecture that some of the costs of changing prices may have fallen, allowing prices to become more flexible.

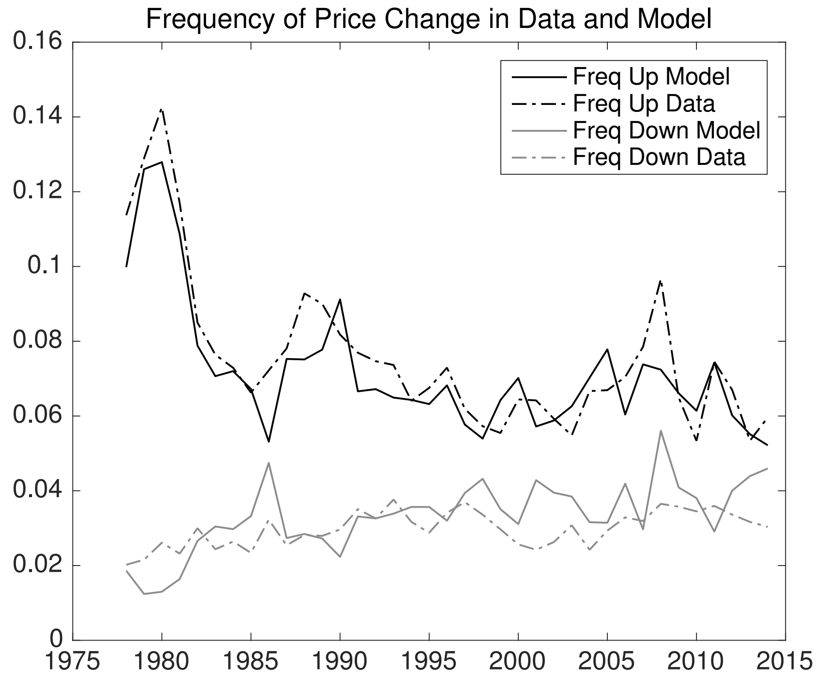

Figure 1.14: Predicted and Actual Frequency of Price Changes

Yet, there is no evidence that prices (excluding sales) have become more flexible over time. Figure 1.12 shows that the frequency of price change (excluding sales) has actually fallen over the past 40 years. Of course, the benefits of changing prices frequently have also fallen over this period since inflation has fallen. For this reason, the evolution of the frequency of price change is not an ideal measure of the evolution of price flexibility.

An alternative (arguably better) measure of price flexibility is the menu cost needed to match the frequency of price change at a particular point in time given the level of inflation at that time. If the menu cost model is able to match the frequency of price change over time with a constant menu cost, this would indicate that prices (excluding sales) have not become more flexible over time.

Figure 1.14 presents the results of this type of exercise. The broken lines in the figure are the frequency of price increases and decreases in the data. The solid lines are the frequency of price increases and price decreases from a simple menu cost model with a constant menu cost.[18] Evidently, the frequency of price change in the data tracts the model implies frequency of price change quite well over time as inflation rises and falls. If the costs of price adjustment had trended down over the past four decades, one would expect that our model would systematically underpredict the frequency of price change toward the end of the sample period. This is not the case.

Since our simple menu cost model with a fixed cost of price adjustment can explain the overall trend in the frequency of price change over the sample period, we conclude that there is no evidence that prices (excluding sales) have become more flexible over time. One might worry that these facts about price changes excluding temporary sale might be somehow contaminated by the increasing frequency of sales (discussed below); but the same downward trend in the frequency of price change is visible even in sectors with essentially no sales, such as the service sector.

One way in which prices *have* become more flexible over time is that the frequency of temporary sales has increased. Temporary sales are distributed very unequally across sectors, occurring frequently in processed and unprocessed food, apparel, household furnishings, and recreation goods, but quite infrequently in other sectors of the economy (see Nakamura and Steinsson (2008)). Figure 1.15 plots the evolution of the frequency of sales in the sectors in which sales are prevalent. In all five sectors, there has been a dramatic increase in the frequency of sales over our sample. In some categories, the increase seems to continue unabated, while in others (especially apparel and household furnishing) the frequency of

[18]For simplicity, in this exercise, we feed the inflation rate into the model directly (as opposed to feeding in a process for nominal aggregate demand and having inflation be an endogenous outcome). The model we use in this exercise is therefore a partial equilibrium model.

sales seems to have plateaued. This trend increase in the prevalence of sales may, in fact, go back considerably before the start of our sample period. Pashigian (1988) documents a trend increase in the frequency of sales going back to the 1960's.

1.7 Conclusion

In this paper, we develop a new comprehensive price dataset going back four decades for the U.S. to study the costs of inflation. We find little evidence that the Great Inflation of the late 1970's and early 1980's led to a substantial increase in price dispersion—the costs of inflation emphasized in standard New Keynesian models of the economy. The frequency of price change varies substantially with the inflation rate, in line with the predictions of standard menu cost models. We find no evidence that the costs of price adjustment have fallen over these four decades, despite the many technological improvements that have occurred over this period–suggesting that the barriers to price adjustment are not purely technological in nature. The frequency of temporary sales has increased dramatically over this period in many sectors of the economy, rising at an almost linear pace. However, the service sector–which has grown over time–exhibits both sticky prices and almost no sales.

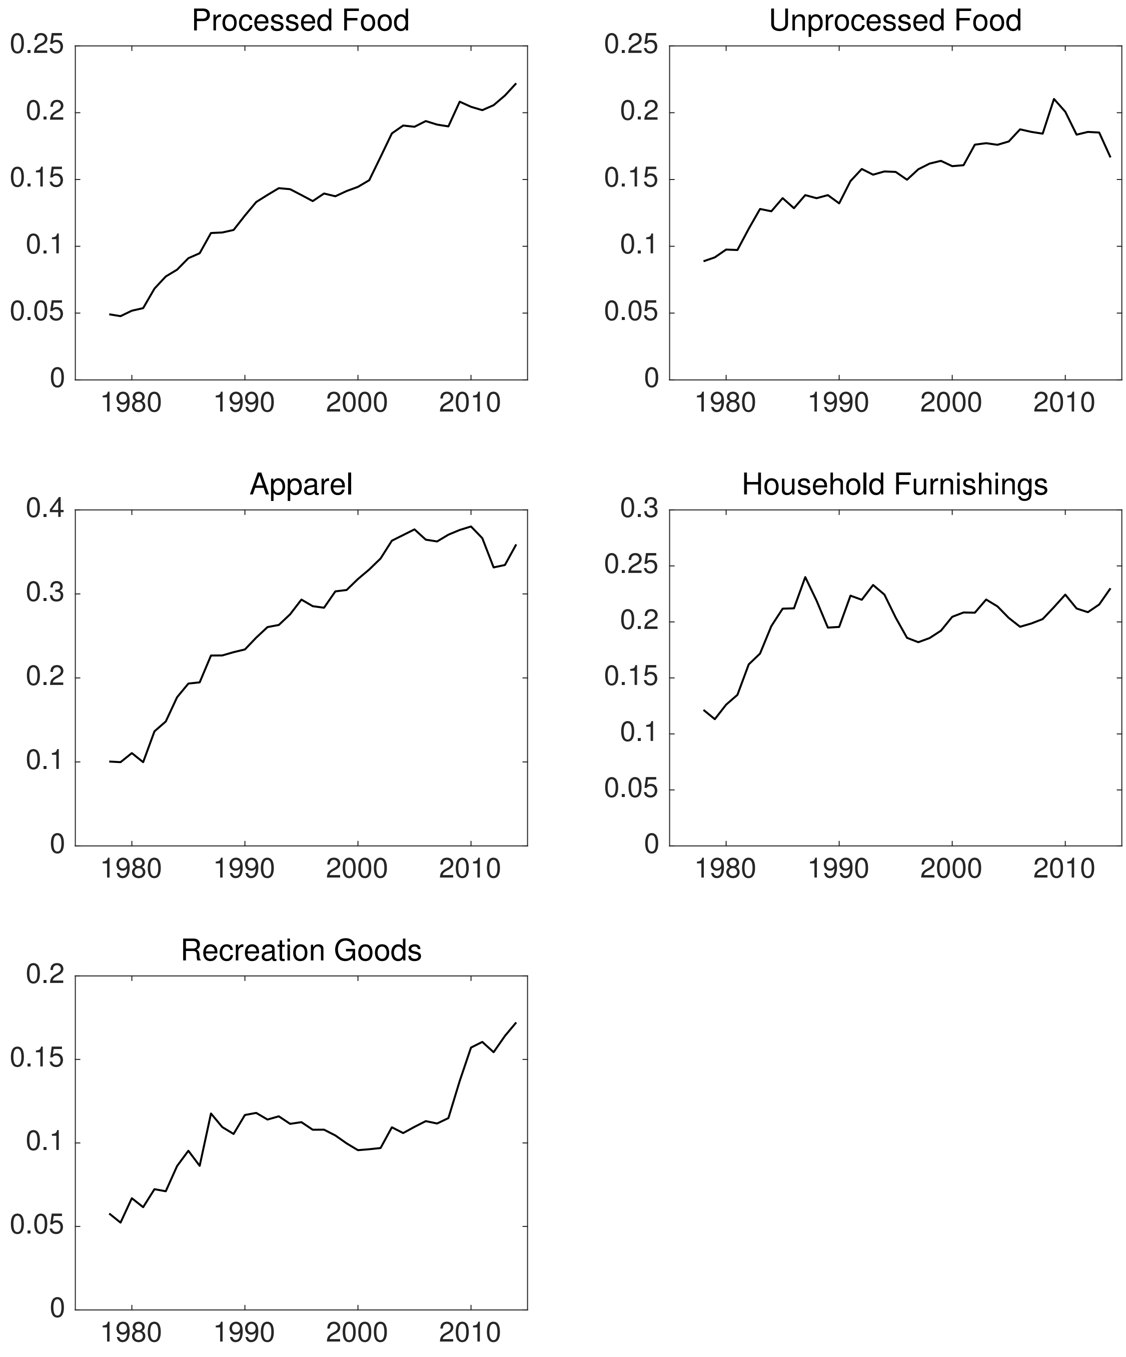

Figure 1.15: Frequency of Temporary Sales

Note: To construct the series plotted in this figure, we first calculate the mean frequency of temporary sales in each ELI for each year. We then take the weighted mean across ELI's.

Chapter 2

The Skewness of the Price Change Distribution: A New Touchstone for Sticky Price Models

Shaowen Luo and Daniel Villar

2.1 Introduction

The dynamics of price changes (when, how, and why firms change the prices of the goods and services that they sell) have been a major focus of the study of monetary economics for the past several decades. It is indeed well known that monetary variables have no influence on real economic activity (monetary neutrality) if all prices can be freely re-set at any point in time. This has drawn attention to the study of frictions in the price-setting process for a long time: Barro (1972) and Sheshinski and Weiss (1977) characterized the pricing behaviour of a firm that faces a fixed price adjustment cost, while Calvo (1983) did so for a firm facing the random opportunity to change its price. What has also become well established is that the distinction between these two approaches in modelling price change dynamics matters greatly for monetary non-neutrality. While central banks have widely adopted Calvo-style staggered price setting into the models that they use to evaluate the effects of their policies, much of the literature has highlighted how this considerably over-states the effectiveness of monetary policy, compared to what it would be if prices are set based on adjustment (or menu) costs.

The literature has emphasized that monetary non-neutrality depends not only on how often prices change, but also crucially on which prices change. Caplin and Spulber (1987) and especially Golosov and Lucas (2007) demonstrated this by showing that if prices are sticky because of menu costs, money is close to neutral. These seminal studies showed that in the presence of menu costs, only relatively large price changes will justify the payment of the cost and occur at all, which makes the aggregate price level considerably more responsive to nominal shocks than in the Calvo model. This mechanism came to be known as the selection effect, and much research has been devoted to re-evaluating the results of Golosov and Lucas (2007), and the strength of the selection effect, in light of new empirical findings established with price micro data sets (most notably, Nakamura and Steinsson (2010) and Midrigan (2011)).

Understanding the selection effect, and to what extent it plays an important role, is necessary to determine the true degree of monetary non-neutrality, but this mechanism cannot be observed directly. It would be very difficult to observe whether the prices that change are those predicted by the selection effect, so its presence and strength must be inferred indirectly from observable price change statistics. The existing work in the field has done this primarily by bringing quantitative price setting models together with the price data that has become available in the past decade. However, an important limitation with these studies is that they have, for the most part, only used unconditional moments of the price change distribution (such as the frequency or size of price changes, averaged over time) to discipline the models used. In this paper, we show that conditional moments, which have been seldom used, are extremely informative and yield new insights on the selection effect. In particular, we find that the selection effect makes very strong predictions about how the shape of the price change distribution should change with aggregate inflation. Using a new data set, the price data underlying the U.S. CPI from 1977 onwards, we show that these predictions are not supported empirically.

In menu cost models, the presence of a fixed adjustment cost induces a selection effect: only price changes that are large enough to justify the cost occur, leaving an inaction region of changes (centered at zero) that are too small to be justified. A positive monetary shock (raising nominal demand) will induce prices that were otherwise already strongly mis-aligned to change, meaning that average price changes would respond relatively strongly to such a shock. This implies, in turn, that the aggregate price level will be very responsive to monetary shocks, eliminating much of the effect of the monetary shock on real activity (money is close to neutral). We exploit the fact that this logic also has strong implications for how the distribution of price changes responds to such shocks: an inflationary shock will push more price changes out of the inaction region to the positive side, and into the inaction region from the negative side. There will therefore be more price changes concentrated on the positive side of the inaction region, leaving a price change distribution that is less dispersed and more asymmetric (negatively skewed). Indeed, all existing menu cost models, because of the selection effect created by the presence of an adjustment cost, imply a very strong negative correlation between inflation and both dispersion and skewness of price changes, and these are implications that can be empirically tested.

The literature on sticky prices has faced thus far been unable to test these types of predictions because the kind of price data that is necessary has only been available for periods of low and stable inflation. Although some studies (such as Alvarez et al. (2011a); and Gagnon (2009)) have used price data from countries that experienced high inflation, they used this data to determine how the frequency of price change behaves at high inflation, without considering the higher moments of the price change distribution. For the U.S., the main source of price data in this line of work, the micro data underlying the Consumer Price Index, was, until recently, only available going back to 1988 (while other commonly used data sets go back even less far). However, we use the data set recently presented in Chapter 1, which extends the C.P.I micro data back to 1977, to evaluate whether the dispersion and skewness of price changes do indeed fall with inflation. Since the newly recovered period includes the highest inflation episodes in the post-war U.S., as well as the disinflation period initiated by the Federal Reserve under Paul Volcker, our data set is particularly well suited for the tests that we propose.

We find that while the dispersion of price changes does go down considerably in high inflation periods, the skewness does not, contrary to the strong predicitons of menu cost models. Since the counter-factual predictions are driven by the mechansim behind the selection effect, these results cast doubt on whether this effect, which has been emphasized as a source of aggregate price flexibility, is really very significant. However, the fact that the frequency of price change rises with inflation (which we find), and that the dispersion falls with inflation, contrary to the assumptions and predictions of the Calvo model (in which there is no selection), seems to indicate that the selection effect is acting to some extent. The quantitative implications of these results for the strength of the selection effect and monetary non-neutrality will be the subject of the following chapter.

The rest of the paper is organized as follows. In what remains of the introduction we provide a more detailed overview of the work done in this branch of the literature. In section 2.2, we present the predictions of a large class of sticky price models, and explain why time- and state-dependent models give such different predictions. Section 2.3 describes the data set that we use and evaluates the predictions of the different models based on the data. Finally, Section 2.4 provides some concluding remarks.

Literature Review

While a few empirical studies of price stickiness in certain industries have been around for some time (e.g. Cecchetti (1986); Carlton (1986); Kashyap (1995)), it is only starting with Bils and Klenow (2004) that monetary economists have been able to start measuring statistics related to price stickiness for the economy as a whole. The facts established by Bils and Klenow and the subsequent empirical studies on price stickiness (most notably, Klenow and Kryvtsov (2008); and Nakamura and Steinsson (2008)) have enriched the discussion on monetary non-neutrality by providing the models that evaluate monetary non-neutrality with a standard by which to be measured.

Caplin and Spulber (1987)had used a very stylized model to show that if prices are sticky, state-dependent pricing implies that monetary shocks can still have little or no effect on economic activity. Golosov and Lucas (2007) then incorporated this mechanism into a quantitative menu cost model that was calibrated to match the new empirical facts of the sticky price literature, and they confirmed that under state-dependent pricing, monetary policy is close to neutral. The model matched the fraction of prices that change (frequency of price change) estimated by the empirical papers, but also the observation that when prices do change, the changes tend to be large. Since, under menu costs, firms will only change their prices when they really need to, and so will not bother incurring a menu cost for a small price change, this latter fact in particular lent credibility to the adoption of a menu cost as the foundation of price stickiness.

Since then, the literature has continued to combine quantitative, micro-founded, price setting models with empirical facts from micro price datasets, and in this way the non-neutrality debate has advanced. While the Golosov and Lucas model matched the frequency and average size of price changes, much subsequent work has modified the model to match other aspects of the distribution of price changes, generally finding that the degree of monetary non-neutrality predicted ends up being much larger than in the original model (for example, Nakamura and Steinsson (2010); Midrigan (2011); Alvarez et al. (2014)). In a slightly different style, Vavra (2014) showed that the frequency and dispersion of price changes are counter-cyclical in the U.S., and introduced counter-cyclical dispersion shocks to match this.

With the exception of Vavra (2014), however, the papers mentioned thus far match moments that are price change statistics averaged across time. Yet all the statistics that they consider can be computed period by period, as they pertain to a distribution of price changes, which is observed period by period. Obviously, focusing on averages across time abstracts from the time series variation in these statistics, which is observed to be quite significant in the data, and this misses out on potentially informative patterns. Our paper departs from most of the existing literature by focusing on the variation of price change statistics over time to evaluate sticky price models. These models are aimed at understanding how the dynamic pricing behaviour of firms aggregates up to the response of aggregate inflation to monetary shocks. A natural way to use the time series variation of price stickiness statistics is therefore to see how they co-move with inflation, both in models and empirically. However, as mentioned earlier in this section, most existing studies have faced the limitation of working with price data sets that only cover periods of low and stable inflation. It is in this way that our data set is novel, as it makes it possible to measure price stickiness statistics at high and low inflation.

Nevertheless, evaluating sticky price models with this kind of time series variation is not unprecedented. For example, Gagnon (2009) and Alvarez et al. (2011a) use price data from high inflation episodes in Mexico and Argentina, respectively, to show that the frequency of price change rises with inflation. This fact is consistent with menu cost models, but it goes against the core assumption of the Calvo pricing model, that firms face a constant probability of changing their prices over time. Our paper confirms this result, but documents more patterns based on other statistics that paint a more nuanced picture. While the relation between the frequency of price change and inflation provides strong evidence against the strict assumptions of the Calvo model, changes in the shape of the price change distribution (measured by its dispersion and skewness) are also informative to distinguish between the models.

Ultimately, we find that neither menu cost nor Calvo models are able to match all the patterns in the data that we present. In particular, the menu cost model makes very strong predictions about the shape of the price change distribution: the dispersion and the skewness fall sharply with inflation. In the data the dispersion of price changes does fall with inflation, but the skewness does not. We are not the first to find empirical failures of this model: Nakamura and Steinsson (2010) and Midrigan (2011) had already pointed out problems with some of the predictions of the Golosov and Lucas model, and shown that changes to the model that corrected these problems overturned the result of low monetary non-neutrality. However, we show that even these modifications to the Golosov and Lucas model, though they reconcile the menu cost framework with the data in some ways, are also inconsistent with the facts that we present. Finally, we also consider models of imperfect information in which firms adjust their prices infrequently (Alvarez et al. (2011b), Woodford (2009)), and find that these also fail to match the data, although each in different ways.

2.2 The Skewness of Price Change in Sticky Price Models

We begin by presenting the models that we will be evaluating, and describing the predictions that we will focus on testing. Our analysis will consider the models that have been used in the sticky price literature, including the Calvo model, the Golosov and Lucas menu cost model and the variants of it that have appeared since. First, we describe the set-up of the various models, both the common framework and the differences that set them apart, before explaining how we derive the predictions, and we finally summarize the predictions.

2.2.1 General Set-Up

All the sticky price models that we consider have certain features in common, that are also used in the sticky price literature in general. First, households maximize expected discounted utility of the following form:

$$E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left[logC_{\tau+t} - \omega L_{\tau+t} \right].$$

All our analysis will focus on the firm's dynamic price setting, so the set up of the household problem matters for our purposes insofar as it determines the relationship between aggregate consumption and the real wage, which will be the firm's main cost. There is then a continuum of monopolistically competitive firms, indexed by z, producing a differentiated product, and aggregate consumption is given by a constant elasticity of substitution aggregator, meaning that each firm faces the standard demand function for its good:

$$c_t(z) = \left(\frac{p_t(z)}{P_t}\right)^{-\theta} C_t.$$

, where θ is the elasticity of demand, and P_t is the CES price aggregator. Firms produce output based on a linear production function, with labor as the only input:

$$y_t(z) = A_t(z)L_t(z).$$

Productivity is subject to idiosyncratic shocks, which have been an important feature of sticky price models since Golosov and Lucas (2007). Large idiosyncratic shocks make it possible for such models to match the large heterogeneity and high average size of price changes observed in the data, which was documented notably by Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008). These shocks are typically modelled as first-order autoregressive processes with normal innovations, but Midrigan (2011) argues that such a process yields a distribution of price changes with tails that are too thin, relative to what is observed in the data. He therefore introduces Poisson shocks in the productivity process in the following way:

$$logA_t(z) = \begin{cases} \rho logA_{t-1}(z) + \epsilon_t, & Probability = p_\epsilon \\ logA_{t-1}(z), & Probability = 1 - p_\epsilon \end{cases} , \ \epsilon_t \overset{iid}{\sim} N(0, \sigma_\epsilon^2).$$

This set-up nests the standard AR(1) productivity, which can be obtained by simply setting the probability of a shock occurring (p_ϵ) to 1. Since we will consider various models with

AR(1) productivity, as well as Midrigan's model with Poisson shocks, we maintain this set-up, and cover the different models by adjusting the relevant parameters.

In order to generate aggregate fluctuations, the sticky price models that we look at incorporate a stochastic process for nominal aggregate demand. Again, we stick to what is most often used in the literature by modelling nominal output as a log random walk with drift:

$$logP_tC_t = logS_t = \mu + logS_{t-1} + \eta_t,\ \eta_t \overset{iid}{\sim} N(0, \sigma_\eta^2).$$

This process stands in for monetary policy in these models: nominal output is determined exogenously, and firms' price responses to these shocks determine how inflation, and how real output respond. We will use the same parameter values for this process (to match the behaviour of US aggregate activity) across the different models, and we define monetary non-neutrality as the variation in aggregate real consumption induced by the nominal shocks. This has become the main way of introducing monetary variables in the menu cost literature because it lends itself much more easily to the global solution methods that are used for such models than explicitly incorporating systematic monetary policy. Although Blanco (2015b) developed a menu cost model with a Taylor-type policy rule, we do not attempt this for the models in this section. Our goal is to show how the price change distribution changes with inflation under different sticky price models, and the aggregate demand process that we use enables us to do this. Next, we describe the price setting problem faced by firms, which is the main dimension along which the different models vary.

2.2.2 Price-Setting

In the standard Golosov and Lucas (2007) menu cost model, firms must pay a fixed cost (in units of labor) whenever they change their price. The period profit function therefore takes the following form:

$$\Pi_t(z) = p_t(z)y_t(z) - W_tL_t(z) - \chi W_t I\{p_t(z) \neq p_{t-1}(z)\}.$$

The menu cost (χ) can then be calibrated to match the frequency of price changes, while the standard deviation of the idiosyncratic shocks can be set to match the average size of price changes (we also set the probability of an idiosyncratic shock occurring, p_ϵ, to 1 to make the process an AR(1), as in the original model). This is, in a way, the most "state-dependent" model, as under the fixed menu cost firms are fully in control of the decision of when to change the price for each good (subject to the constant menu cost). It is this feature that makes prices very responsive to aggregate demand shocks, and that famously yields very low monetary non-neutrality.

The first extension to the menu cost model that we consider is the Nakamura and Steinsson (2010) multi-sector menu cost model. The only change here is that firms are separated into sectors, with firms in different sectors facing different menu costs, and a different variance of idiosyncratic shocks. This reflects the fact, documented in the paper and in Nakamura and Steinsson (2008), that the frequency of price change varies considerably across sectors, as does the average size of price changes. Golosov and Lucas (2007) calibrated their model to match the average frequency of price changes across sectors, and Nakamura and Steinsson show that calibrating sector by sector makes a major difference for the degree of monetary non-neutrality in the models, as the multi-sector model predicts much higher non-neutrality than the standard model.

Midrigan (2011) modified the standard menu cost model in two ways: first by changing the idiosyncratic shock process so that it would feature fat tails (which we described above), and giving firms a motive to make small price changes. In the standard model, since a firm always has to pay a fixed cost to change its price, there will be a threshold for the size of the price change, such that changes below a certain size are not profitable and do not occur. Midrigan (2011) models multi-product firms that can change the prices of all their products for the payment of the menu cost. Because of this, a firm might choose to pay the menu cost to change the product of a particularly mis-aligned product price, and then also take the opportunity to change the price of another product by a small amount. This enables the model to match the considerable fraction of small price changes that are observed in the data, but it also makes the model much more difficult to solve. We therefore follow Vavra (2014) in simplifying the Midrigan model by assuming that, instead of producing multiple products, firms each period are randomly given the possibility of changing their price for free (with a low probability), or by paying a menu cost. This adds, as an additional parameter to calibrate, the probability of drawing a zero menu cost (free price change): p_z. With the additional parameters in this model, we target the fraction of price changes that are small, as in Midrigan (2011).[1]

[1] Midrigan (2011) defines a small price change as a price change that is less than half, in absolute value, of the average size of price change. Due to the variation in the average size of price changes over time and across sectors, we prefer to use an absolute measure, and focus instead on the fraction of price changes that are smaller than 1% in absolute value. Finally, Midrigan (2011) also emphasized the failure of the Golosov and Lucas model to match the kurtosis of the price change distribution, and the introduction of Poisson idiosyncratic shocks helps to get the kurtosis in the model closer to what it is in the data. However, it turns out to be very difficult to match (it seems to be very high in the data), and Midrigan (2011)1 does not achieve it completely. We therefore do not match the kurtosis either.

We also consider a Calvo model, which has the set-up described above, except that firms, instead of facing a menu cost, have a fixed probability every period of receiving the opportunity to freely change their price (otherwise, they do not get to change price). This is equivalent to the simplified Midrigan model that we describe, but with the high menu cost set to infinity, and the probability of a free price change set to equal the average frequency of price change in the data. This model includes idiosyncratic shocks to obtain a distribution of price changes, and we also set the variance of these shocks to match the average size of price changes. The variance needs to be higher than in menu cost models, because menu costs induce the selection effect that naturally leads to large price changes to be more likely.

Finally, we also include two models involving imperfect information: the Alvarez et al. (2011b) model of observation and menu costs, and the rational inattention model of Woodford (2009). In the former, firms must pay a fixed cost to observe the relevant state (or conduct a "price review"), and a menu cost to change their price. Facing such costs, firms conducting a price review choose the date of the next review, and a price plan until that date. Woodford (2009) considers the same type of price-setting problem, but within the rational inattention framework proposed by Sims (2003): firms face a cost based on how much information they process, and therefore choose to receive limited information based on which they choose when to review prices. In this model, the cost of processing information is a crucial parameter, and both the Calvo model and standard menu cost model are nested as extreme cases of the information cost in this set-up (infinite and zero, respectively). Furthermore, intermediate values of the information cost result in what is described as a "generalized Ss model": while a simple Ss model involves a threshold rule for price adjustment, a generalized Ss model features a probability of price adjustment as a function of the degree of price mis-alignment. This is the kind of model that we work with in Chapter 3, and we view the rational inattention framework as a potential micro-foundation for this.

2.2.3 Solution and Simulation

We solve each of the models mentioned above by value function iteration, mostly with the parameter values used by the original authors, which were set for the models to match various features of the micro price data. One difficulty in solving these models is that in all of them the price level (P_t) is an aggregate endogenous variable whose evolution depends on the behaviour of all firms. This means that, in principle, every firm's relevant state should include the state of every other firm, which makes for an infinitely large state space. As done elsewhere in the literature, we use an approach analogous to Krusell and Smith (1998) to solve the model assuming a relationship between the price level and a small number of variables, and to then verify that the resulting solution is consistent with the assumed relationship. In the appendix, we provide more details about the procedure, as well as the calibration of the different models. The parameters of the process for nominal aggregate demand, described above, are calibrated to match the average growth and volatility of U.S. nominal GDP, and the same values are used for all the models.

The first aim of our paper is to document what these different models imply for the price change distribution at different inflation rates. Our approach is to simulate each model, for 1,000 periods (months) and 40,000 firms. From these, we obtain a simulated series for aggregate inflation (determined by the endogenous response of prices to the nominal aggregate demand shocks) and a distribution of price changes for each period. Since the models are calibrated to match the frequency of price change that is observed empirically, the vast majority of prices do not change every period. Our analysis is therefore based on the distribution of price changes, conditional on a non-zero price change, and this applies for the rest of the paper, including in our empirical work. We compute various moments of each period's price change distribution, giving us a time series for each moment, and compute correlations between inflation and each of the moments, and this is how we determine how the price change distribution changes with inflation.

As mentioned in the introduction, the studies that have examined price change statistics in high inflation environments have mostly focused on whether the frequency of price change rises with inflation, as the menu cost model predicts. We present the correlation between frequency and inflation in the models, but also consider other correlations with other moments: the standard deviation of price changes, and the skewness of price changes. As we will show, the menu cost models have very strong and clear implications for these correlations that are markedly different from those of the Calvo model. Furthermore, as seen in Midrigan (2011), the shape of the price distribution can be very informative about the importance or presence of the mechanisms that weaken the role of monetary shocks, and it is therefore to be expected that the way in which the shape of this distribution changes (as described by the dispersion and skewness) with inflation would also be informative about these mechanisms.

We present a summary of these theoretical results in Table 2.1, indicating whether each correlation is positive (+), close to zero (0), or negative (-) in the different models:

In order to further illustrate these results, we present scatter plots between inflation and the different moments from the simulations (in which one point represents one period in the model simulations). Figure 2.1 shows the correlations for the frequency of price change,

Table 2.1: Moments correlation with inflation

Model	corr(π, frequency)	corr(π, std. deviation)	corr(π, skewness)
Calvo (1983)	0	+	+
Golosov and Lucas (2007)	+	-	-
Nakamura and Steinsson (2010)	+	-	-
Midrigan (2011)	+	-	-
Alvarez et al. (2011b)	+	0	-
Woodford (2009)	+	0	+

while Figures 2.2 and 2.3 do so for the dispersion and skewness of price changes, respectively. These bring out the fact that in the menu cost models, the relationships between inflation and dispersion and skewness are very clear and strong (especially in the Golosov and Lucas model for the dispersion). In contrast, the same relations in the Calvo and imperfect information models are not so strong.

Although the relationships come out very clearly in these simulations, it could be a concern that the higher moments that we are estimating might not be well defined in the distributions that we are working with. In addition, estimates of higher moments are very sensitive to outliers, which would be of concern particularly when we estimate from the data. That is why we also consider alternative measures for the dispersion and skewness of price change: the inter-quartile range (for dispersion) and Kelly's coefficient of skewness[2](as opposed to "moment skewness", which is what we have been estimating so far). Since these statistics are quantile-based, they are well-defined for any distribution, and they are also less sensitive to outliers. The correlations are similar for all the models (inter-quartile range compared with standard deviation, and moment skewness with Kelly Skewness). Figure 2.4 shows scatter

[2]These statistics are defined as follows, with Q_i representing the i^{th} percentile. Inter-quartile range $= Q_{75} - Q_{25}$. Kelly Skewness $= \frac{(Q_{90}-Q_{50})-(Q_{50}-Q_{10})}{Q_{90}-Q_{10}}$. Kelly skewness essentially measures the degree of asymmetry in a distribution, comparing the size of the right and left tails.

plots of Kelly Skewness in the different models.

Another concern could be that these simulations all assume that the value of steady-state inflation is held constant throughout the simulated time period. This could be problematic in terms of testing the predictions on data, as the U.S. clearly went from a moderate to a low inflation regime over our sample period. To address this, we also conduct the following exercise: we solve each model for different values of the trend inflation parameter (μ), and for each solution compute the average dispersion and skewness of price change (either from the stationary distribution of price changes, or averaging over simulated time periods; they are almost the same). In Figure 2.5, we plot the results.

What the scatter plots show is that, as in the "short-run" analysis, the dispersion and skewness of price changes fall with trend inflation in the menu cost model (we are only plotting results for the Golosov and Lucas model, but the same pattern holds for the other menu cost models). Here too, the Calvo model predicts weak positive relations for both moments. This will be important when comparing the skewness of price change between the low and high inflation periods in the data.

To conclude our theoretical analysis, we emphasize that the correlations that we consider all have the same sign in the four menu cost models (Golosov and Lucas, Nakamura and Steinsson, Midrigan, and observation costs). The scatter plots show that the values taken by moments we report do vary across the models (for example, in the Golosov and Lucas model the skewness of price changes takes a wider range of values than in the other models), but the fact that the sign and strength of the correlations across the models are similar is notable. Indeed, the Nakamura and Steinsson and Midrigan menu cost models were developed as extensions of the Golosov and Lucas model to make it match new empirical facts, and the

Figure 2.1: Simulated frequency and inflation from different models

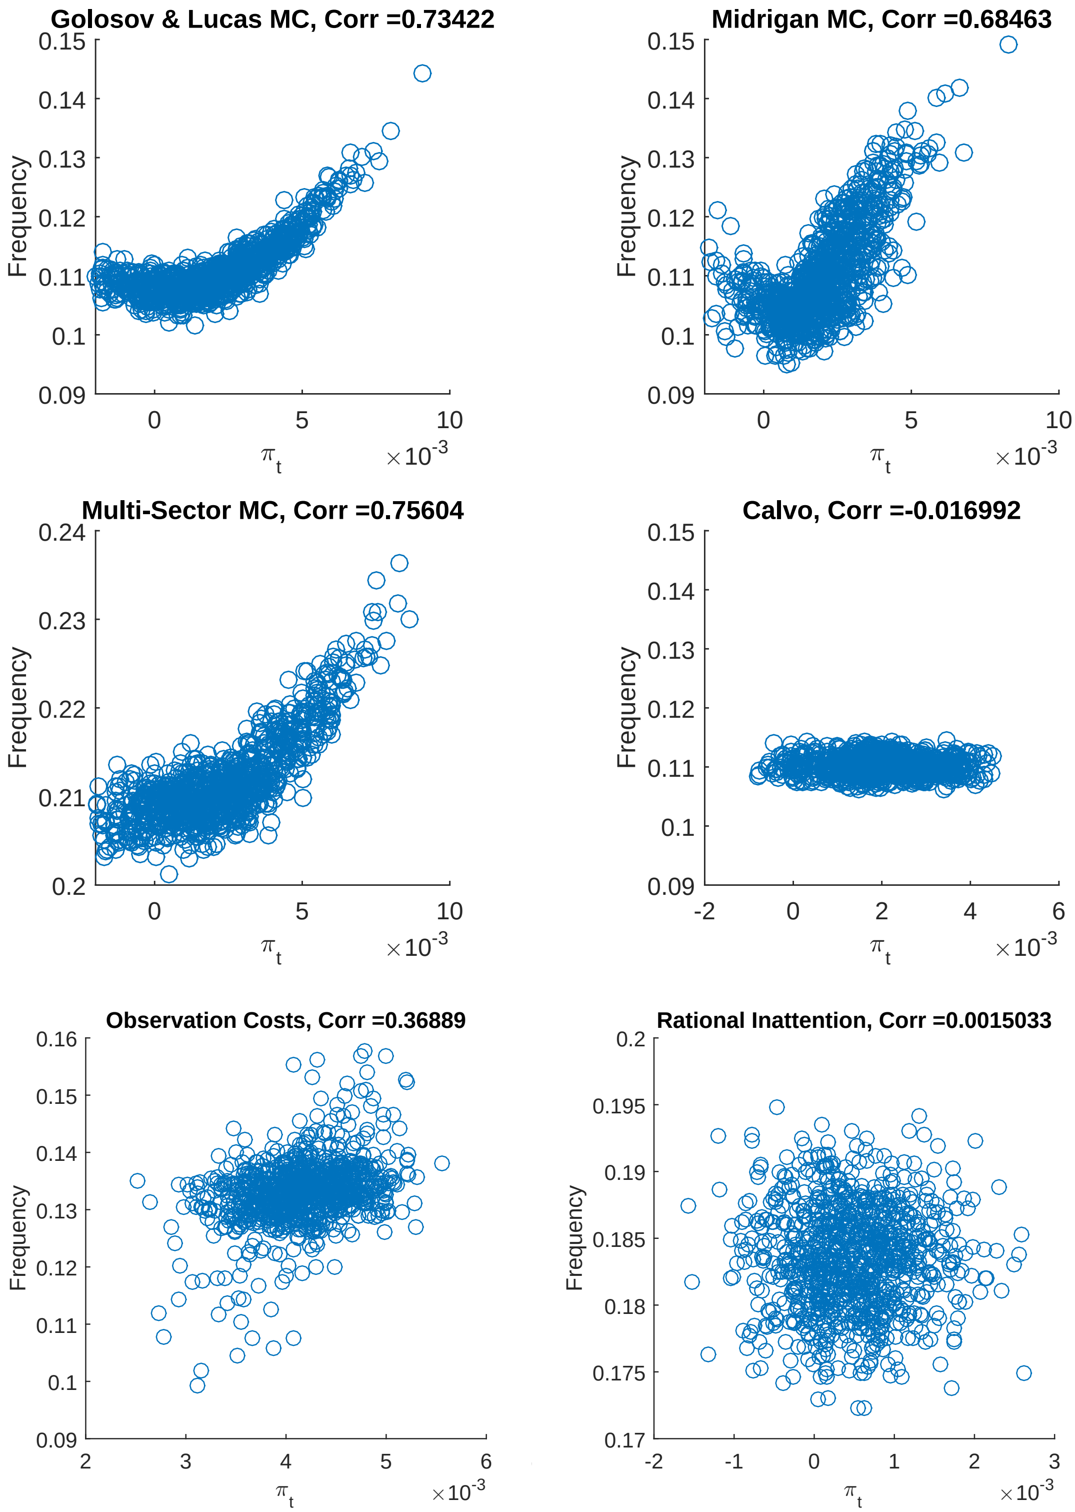

Figure 2.2: Simulated dispersion and inflation from different models

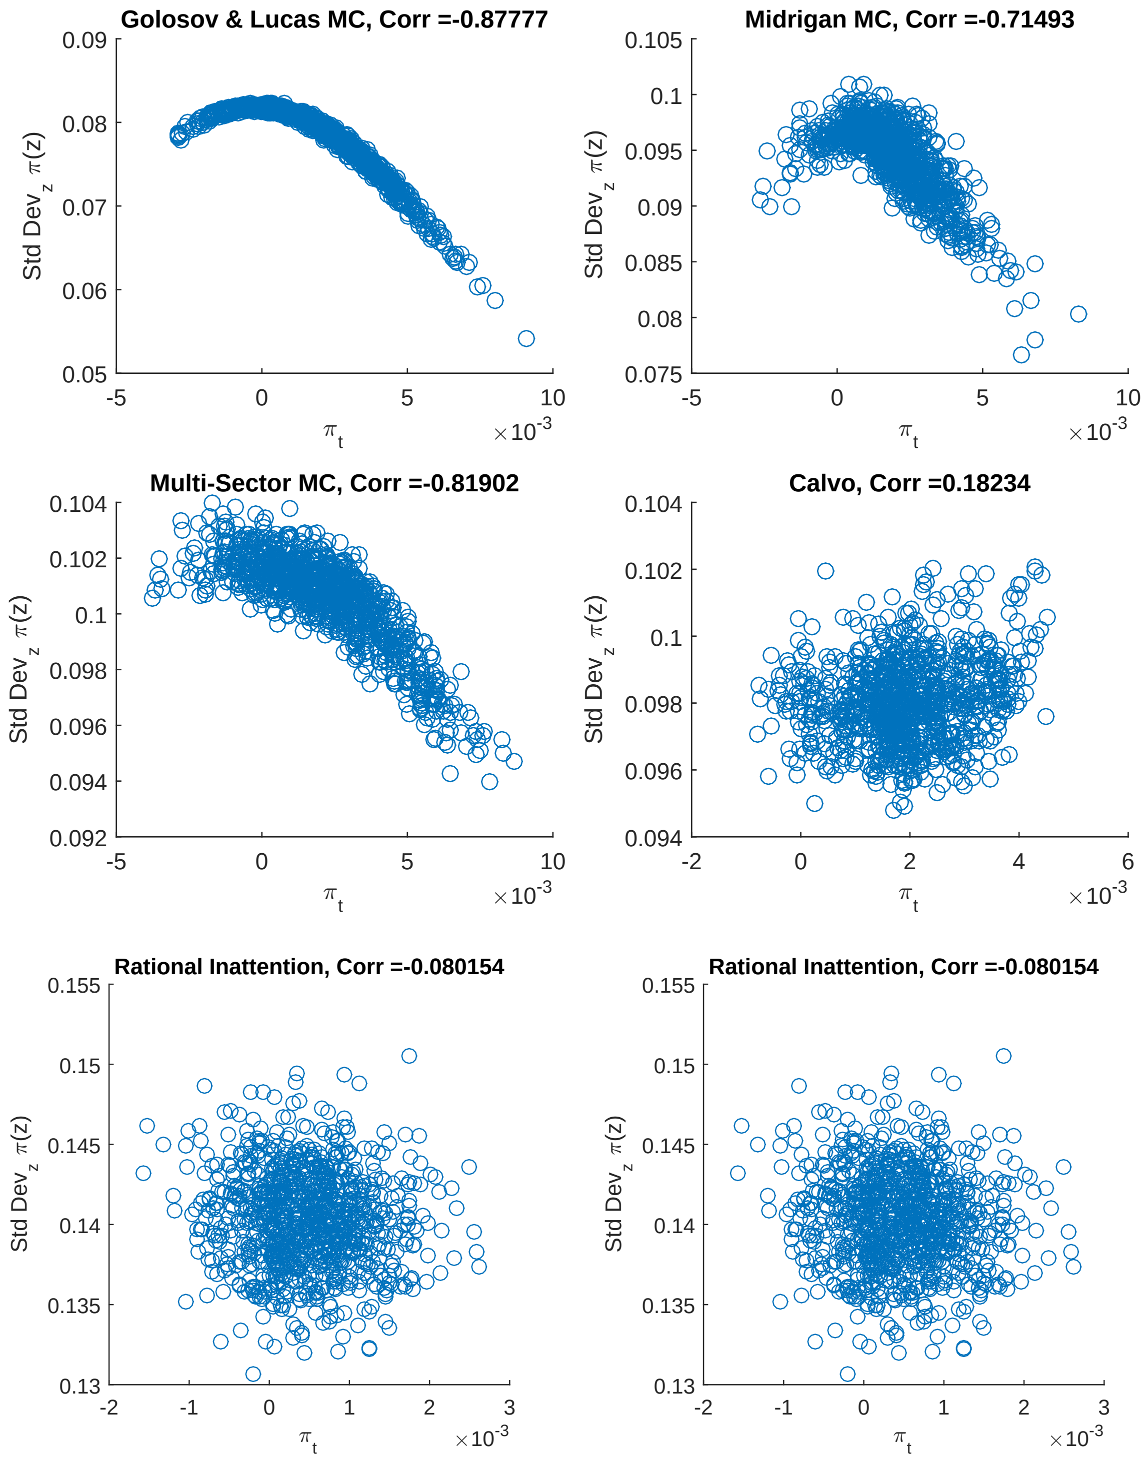

Figure 2.3: Simulated skewness and inflation from different models

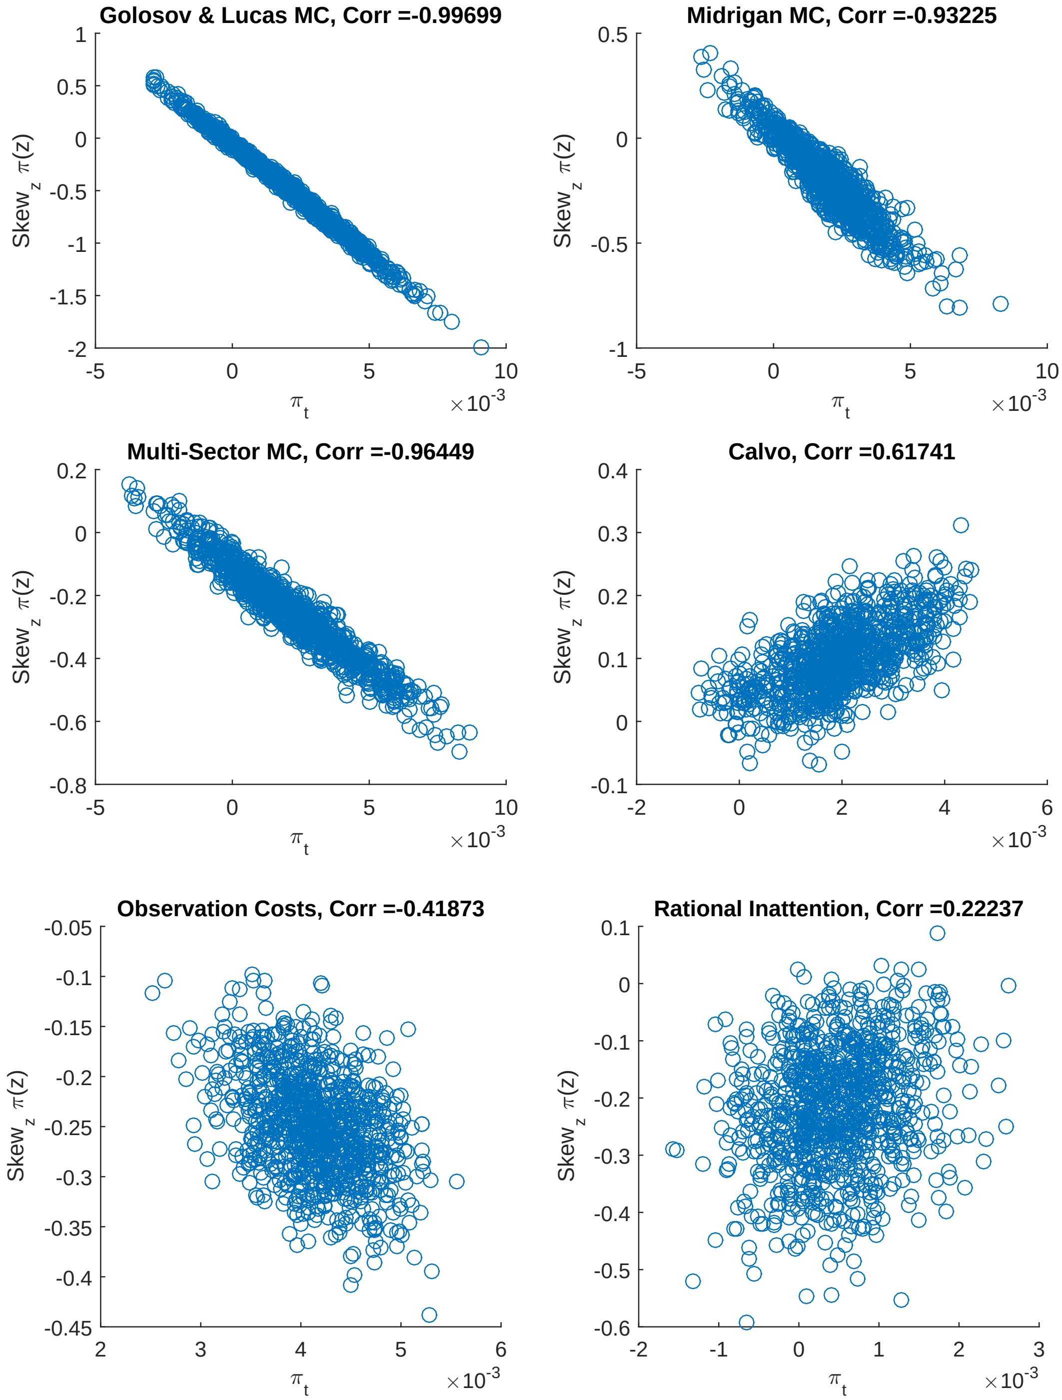

Figure 2.4: Simulated Kelly skewness and inflation from different models

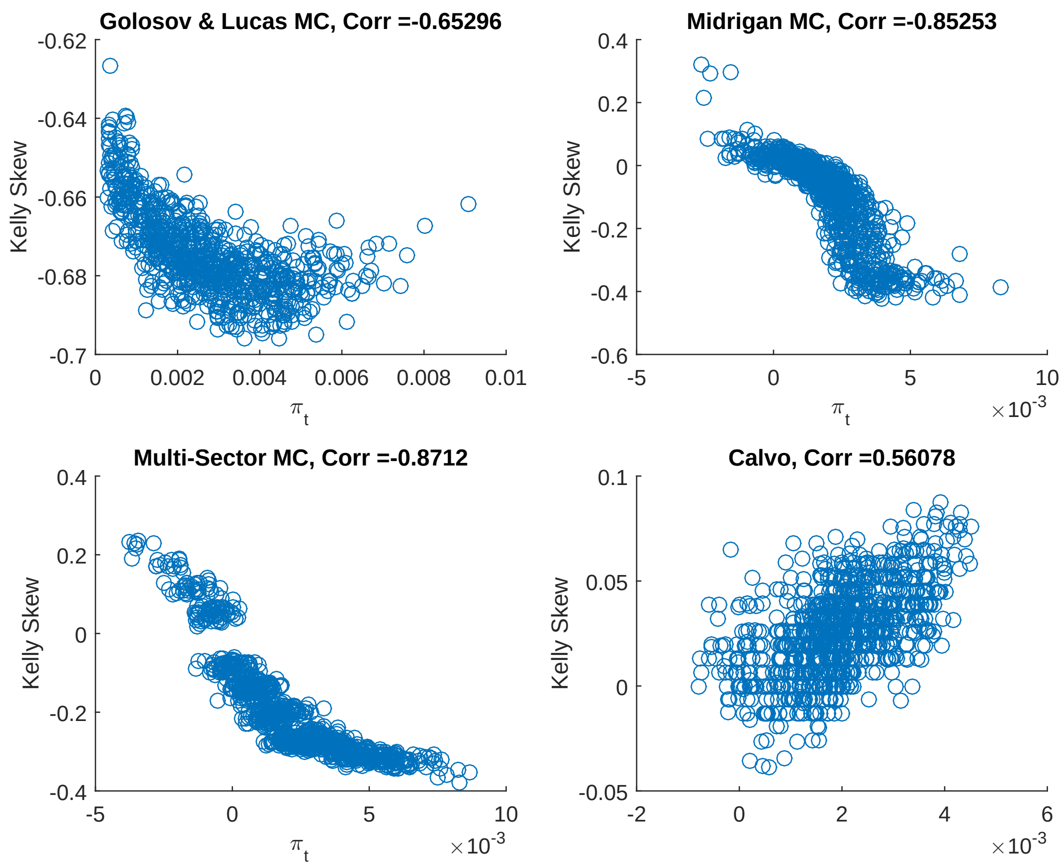

Figure 2.5: Simulated long run statistics from different models

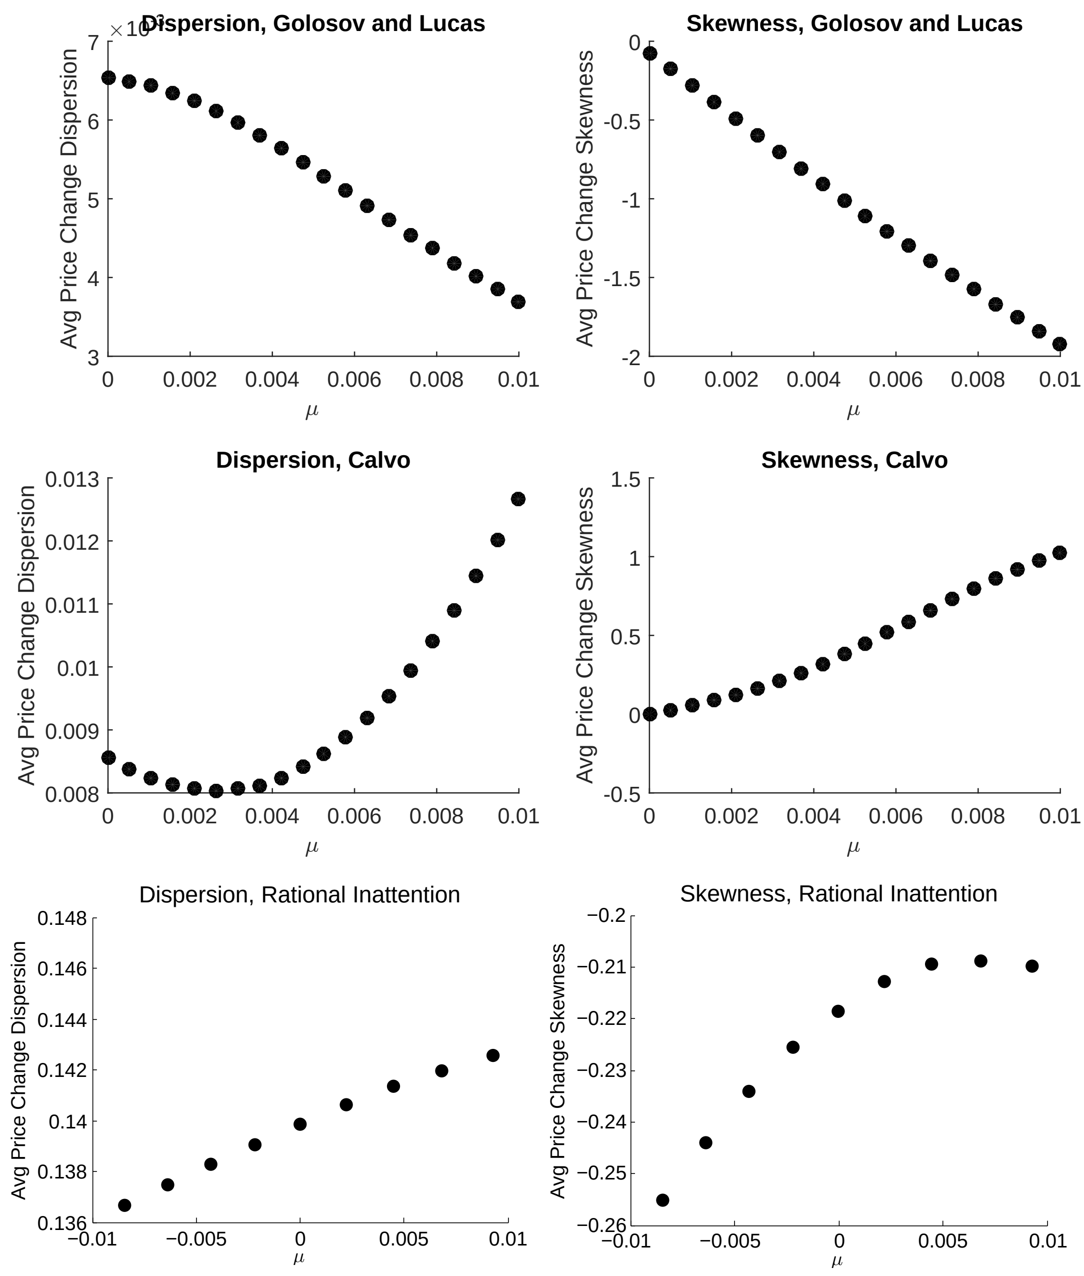

changes made considerably weakened the selection effect that reduces the importance of monetary shocks. However, what we find here is that, despite the important changes made to the baseline menu cost model, they all have the same implications along the dimensions that we are considering. Next, we discuss the intuition behind these theoretical results.

2.2.4 Intuition for the Menu Cost Model

Menu cost models are often also known as "Ss" models, due to the fact that they tend to feature an inaction region for price changes (the edges of which can be labelled with "S" and "s"), and this makes it easier to understand the theoretical correlations between inflation and the moments of the price change distribution that we find in this section. Price change dynamics in the menu cost model can be thought of in the following way: both idiosyncratic and aggregate nominal shocks give a distribution of desired price changes (the price change a firm would choose if it changed its price, or in the absence of price change frictions). The presence of a menu cost means that only desired prices above a certain size (positive and negative) will actually occur, as only those will yield a benefit to the firm big enough to compensate for the menu cost. The realized price change distribution in this model is therefore the underlying distribution with a band containing 0 removed, as illustrated in Figure 2.6 below.

The presence of idiosyncratic shocks yields variation in firms' desired prices, and nominal aggregate shocks move the position (average) of the underlying distribution. For example, a positive aggregate shocks moves the distribution to the right, which also leads to realized prices being higher on average, resulting in higher inflation (the reverse is true for negative

Figure 2.6: Intuition for the Menu Cost Model

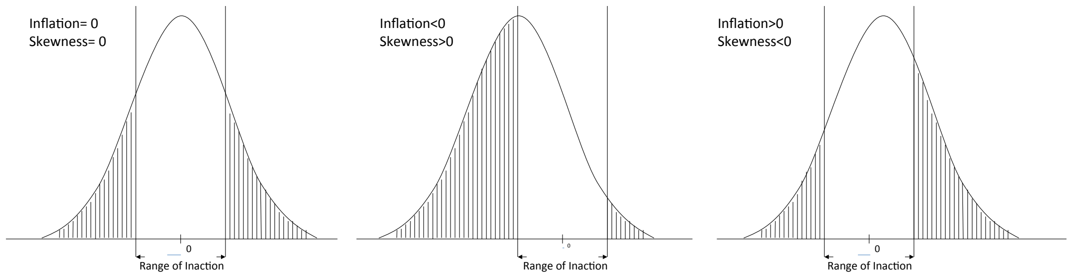

aggregate shocks). Such shocks also result in a higher fraction of price changes being positive, which are separated from the negative ones by the inaction region. This reduces the dispersion of price changes because a bigger fraction of them are on one side of the inaction region, and therefore relatively close to each other. It is when the share of price changes on either side of the inaction region is equal that the dispersion is highest, and by the same logic, higher than when inflation is negative (when more price changes are decreases), which is what we see in the dispersion plots (Figure 2.2) for the menu cost model: dispersion decreasing with inflation in the positive region, and increasing in the negative region, with the maximum attained at zero inflation.

The logic for why the skewness falls with inflation is related. The skewness, as a statistic, measures the asymmetry of a distribution, or the relative sizes of the right and left tails. As a positive aggregate shock raises the average desired price change, and the average realized price change, some negative price changes (to the left of the inaction region) remain and form the left tail. This makes the skewness negative: the resulting distribution has a left tail (price decreases relatively distant from the average price change, which is positive), without a corresponding right tail (as price increases are to the right of the inaction region and relatively close to each other). Furthermore, for the range of values that inflation takes in our simulations (which corresponds roughly to the historical range for inflation since the late 1970's), there is always a significant proportion of negative price changes. This means that as inflation rises (due to larger positive aggregate shocks), these negative price changes form a left tail in the price change distribution that is further and further (to the left) of the average of the price change distribution, leading to a skewness that is more negative.[3] What this also implies is that the relationship between skewness and inflation is monotonic, decreasing for positive and negative values of inflation.

It is important to emphasize that these correlations have to do with the central mechanism of the menu cost model: the selection effect. When firms face a fixed cost to changing their price, only relatively large price changes will occur, leading to the presence of the inaction region. As the average of the underlying distribution rises (moved by aggregate shocks), there is a large response of inflation because there is a large share of price increases that are marginal: without the shock they would not occur, but are pushed outside the inaction region (and many marginal price decreases do not occur with the shock), leads to a relatively large rise in inflation, muting the real effect of the aggregate shock. This is the logic for why state-dependent models are known to imply low levels of monetary non-neutrality.

However, what we show is that this same mechanism leads to predictions that are in principle observable: the presence of an inaction region means that positive aggregate shocks should lead to not only more price increases, but to a distribution with price changes more

[3] This also means that if the aggregate shock were so high that all price changes were positive (to the right of the inaction region), the relationship would break down, as price decreases would no longer be separated from price increases. However, this would also mean that all prices would change, and that inflation would be extremely high. This kind of situation, or anything resembling it, never occurs in the period we are considering.

concentrated on the right, leading to a declining dispersion and skewness. This does not occur in a Calvo model: in such a model every desired price change has a fixed probability of being realized, so as the desired price changes rise, the shape of the realized price change distribution does not change in a meaningful way.

The intuition for this theoretical result is easiest to explain in the case of the "standard" Golosov and Lucas model, or in general any menu cost model with a single fixed menu cost. The other menu cost models that we consider feature a richer structure of menu costs that led to some very different empirical predictions. However, we have shown that these models also imply negative correlations for the dispersion and skewness of price changes, and the intuition for this is the same as for the standard model. In the multi-sector menu cost model, different sectors face different menu costs, and this can be thought of as sectors facing different inaction regions, with each sector behaving as described for the standard menu cost model. Therefore, the aggregate price change distribution behaves similarly to how each sector's distribution does.

The Midrigan model involves firms randomly facing either a positive or zero menu cost. This weakens the selection effect, because there is now a positive probability that a firm will change its price even if it will be a small change, so that price changes are no entirely "selected" based on how out of line the original price is. However, the selection effect is still present to a certain extent, because it is only relatively large price changes that will happen with certainty (as those will be the only ones for which a firm will be willing to pay the positive menu cost, when it faces the positive menu cost). It is this difference between small and large price changes that makes the same mechanism present in this model and drives the correlations, even though small price changes do occur (as they do in the data, but do not in the Golosov and Lucas model).

We have shown that menu cost models, under the assumptions commonly made in the literature, make clear, consistent predictions about how the shape of the price change distribution changes with inflation, and that these do not change much based on the type of menu cost model in question, and that the predictions are strikingly different from those of the Calvo model. Furthermore, these are predictions that can be tested with the price data available to us, which enables us to evaluate this broad class of sticky price models. In the following section, we do this by presenting the empirical counterpart to the correlations presented in Table 2.1, and we discuss how each of the models falls short of matching the data.

2.3 Empirical Evidence from High Inflation Periods

In the previous section, we documented the predictions made by various sticky price models on the behaviour of price changes at different inflation rates. In this section, we describe the data set that we will use to test these predictions, and report that while the inflation-dispersion correlation is consistent with the empirical evidence, the inflation-skewness correlation is not.

2.3.1 Previous Empirical Work

The micro data that underlies the U.S. Consumer Price Index (CPI), gathered by the Bureau of Labor Statistics, is one of the most widely used data sets in the literature on monetary price-setting models. Bils and Klenow (2004) were the first to use this data set to provide estimates for the frequency of price change. Since then, other studies have documented additional features of the price change distribution using this data set (e.g. Nakamura and Steinsson (2008); Klenow and Kryvtsov (2008)). The availability of a large, representative data set that makes it possible to observe the price changes of very specific products has lead monetary economists to develop models that match the behaviour of price changes as closely as possible.

The data set that has been used in this line of work covers the period 1988 to the present, as 1988 marked a major revision of the structure of the CPI. However, a limitation of the data used thus far is that throughout this period, aggregate inflation has been relatively low and stable, especially compared to the years before. Since 1988, the maximum twelve month change in the headline CPI has been 6.2% (4.6% for the Core PCE), and the average has been 2.8% (2.2% for the Core PCE). Partly because of this, most research on sticky price models up until now has focused on matching moments of the price change distribution that are averaged over time (the main exception being Vavra (2014), who uses the CPI micro data to investigate the cyclicality of price change moments). But as we showed in the previous section, the models imply that these moments would change over time, and in a way that is closely related to aggregate inflation, with implications that differ strongly across models.

Motivated by this, a few studies have used data from other countries that experienced episodes of high inflation, such as Argentina (Alvarez et al. (2011a)) and Mexico (Gagnon (2009)). These studies also used the micro data that underlies the CPI's of these countries, and reported how various price change statistics change as inflation goes from low, to moderate, to high. They find that the frequency of price change is fairly constant, and not very responsive to inflation, at low levels of inflation (below 10% annual). Once inflation rises even higher, however, the frequency of price change begins to rise sharply with inflation. In addition, they show that a standard menu cost model matches this relationship very well. What this shows is that, first, the Calvo (1983) assumption of a constant frequency, while possibly approximately valid for low inflation, becomes problematic when inflation rises beyond a certain level. Second, the evidence presented in these papers is shown to be consistent with standard menu cost models, suggesting that they better explain the behaviour of price changes when inflation is high. However, Alvarez et al. (2011a) and Gagnon (2009) do not look at the higher moments of the price change distribution that we emphasized in the previous section, which is what we do in this paper.

2.3.2 Data Set and Construction of Statistics

The CPI Research Database collected and maintained by the U.S. BLS contains about 80,000 monthly prices collected from around the U.S, classified into about 300 categories called Entry Level Items (ELI's). As mentioned before, the data going back to 1988 has been available for a little over a decade. The data going back to 1977 has recently become available, and this is the novel part of the data set that we use extensively. This new data set has thus far only been used by the work included in this dissertation, and Chapter 1 also describes in detail just how the data set was made available. As explained in the BLS handbook of Methods, there were several changes made to how the BLS samples prices and computes the CPI. While there are many variables present in the post-1988 data set that are not available for the older period, we are able to study the price change distribution in a way that is consistent throughout the whole period, and with the theoretical framework that we are testing. First, we have access to the variables that identify specific products, and that reveal when a substitution has occurred (when a new version of a product has replaced the old one). Second, the data set contains information on when any given price is a temporary sale, or imputed (or not properly collected). Because of this, we are confident that we are observing the price changes of identical products and services, with the price being actually observed; and all of this with the same standards throughout the sample period.

The empirical literature on price setting has emphasized the importance of identifying "pure" or regular price changes, as opposed to price changes coming from temporary sales or substitutions. The reason is that sales and substitutions have features that make them different in terms of their relevance for the study of the role of monetary policy and aggregate shocks. Indeed, when a product goes on sale, its price will change, but it is not clear that this happens in response to any changes in aggregate conditions. What's more, products on sale tend to revert back quickly to their pre-sale price. This distinction was pointed out notably by Nakamura and Steinsson (2008), and Anderson et al. (2015) document the ways in which sale prices behave differently from regular prices. In a similar way, the distinction between regular price changes and substitutions is made because a price change coming from a product substitution could reflect the changes in product characteristics or in quality that could be behind the substitution. Although it is possible in some cases to estimate the contribution of quality or characteristic changes to a substitution price change (and the BLS does for certain products), we prefer to use the product identifiers to focus on price changes involving identical products.

In order to test the predictions that we presented in the previous section, we use the data set to construct distributions of price changes for each month, and a few observations on how these are constructed are in order. First, since the vast majority of prices do not change in any given month, these distributions only include non-zero price changes (which corresponds to what we look at in the theoretical results). Second, because estimates of higher moments are very sensitive to outliers, we follow other empirical work in excluding price changes whose absolute value is above a certain value (e.g. Klenow and Kryvtsov (2008); Alvarez et al. (2014)), (our threshold is one log point). Third, Eichenbaum et al. (2013) have shed light on problems with the methods of reporting and collecting prices in some of the product categories of data sets such as the CPI. They show that this leads to erroneous small price changes appearing in the data, price changes that come from the price collection methods, and that do not reflect actual price changes. This is particularly important for us, as estimates of dispersion and skewness will be sensitive to the relative amounts of small and large price changes. We deal with this by constructing statistics that exclude very small price changes ($< 1\%$ in absolute value) in the ELI's that Eichenbaum et al. flagged as problematic as a robustness check.

Finally, it has been pointed out by Nakamura and Steinsson (2008 and 2010) that there is significant heterogeneity of price change statistics across sectors. To report the average overall frequency of price change, they estimate the frequency for each ELI, and then take a weighted average of each frequency (with the expenditure weights that go into the CPI). The same method is used by many of the other cited empirical studies. For the frequency of price change, we use the same method, considering both the weighted median and mean frequency[4]. For the dispersion and skewness, we follow a similar approach: we first estimate each moment by sector-month. However, as ELI's are fairly narrow categories, most of them have a handful of price change observations in any given month, fewer than would be necessary to estimate higher moments with any precision. We therefore do not use ELI's as our definition of sectors, but instead separate products into 13 "major groups", which are listed in the appendix. While this sectoral classification is fairly broad, it allows us to separate goods and services into similar categories, while leaving enough observations in each sector-month to obtain good estimates of the dispersion and skewness., and then for each month take weighted averages of the statistics.

This approach then leaves us with monthly series of the different moments of the price change distribution. We believe that our approach, following the empirical price setting literature, gives us the most valid estimates to compare with those from model simulations. Indeed, the models that we are testing involve "pure" price changes, and abstract from temporary sales and product substitutions, which is why we try as much as possible to include only regular price changes in our empirical estimates. Perhaps more importantly, the models do not allow for differences across sectors. Such differences, such as sector-specific shocks, have the potential to strongly affect the shape of the overall price change distribution (when all price changes across sectors are pooled together), in turn affecting the higher moments of the distribution. Because of this, we might see the moments of the

[4]Nakamura and Steinsson (2008) highlight the difference between the mean and the median, arising from the fact that the distribution of frequencies by ELI is very skewed to the right, with a few ELI's having very high frequencies. They argue that the median is a better measure of the average frequency in the sense that a single-sector menu cost model calibrated to match the median frequency is a much better approximation of a multi-sector model, of the kind described in Section 2.2. In this way, the median frequency is a statistic that better describes the degree of price stickiness (as it relates to monetary non-neutrality). This is also why we calibrate all the single sector models to match the median frequency.

"pooled" distribution of price changes vary over time due to such sector-specific shocks, which would be unrelated to the mechanisms that are behind the predictions of the models that we described in the previous section. For this reason, we attempt to "control for" these kinds of effects by computing statistics sector by sector.

2.3.3 Results

The goal of our empirical work is to determine whether the theoretical patterns documented above are borne out by the data. As in the theoretical section, we focus on the correlations between aggregate inflation and price change dispersion, and between inflation and price change skewness. The price change moments are calculated as described above, and our preferred measure for aggregate inflation is monthly core PCE inflation. We prefer to use core inflation because the sharp changes in headline inflation tend to be driven by changes in the global market prices of food and commodities, which would not be well described by the price-setting models that we are working with. However, we will also compute correlations with headline inflation as a robustness check (as well as using estimates of the moments excluding price changes from food and energy categories). Finally, to control for seasonality in the inflation and moment series, we calculate the correlations after removing month dummies from the series, and after applying a moving average smoother to them.

The price data is monthly, and inflation series are monthly, so we can compute the correlations at a monthly frequency. However, the drawback of using monthly series is that each period's moment estimates are based on relatively few observations, making them less precise (this is especially important for higher moments such as the dispersion or skewness). The alternative is to group price change observations by quarters or years (but still separating them by sector) and to estimate the moments from these samples, which gives us more precise estimates (as they are based on distributions with more observations), but only quarterly or annual moment series. Since quarterly and annual inflation averages also have the advantage of containing less noise than monthly inflation series, we consider monthly, quarterly, and annual correlations. We present the results in two ways: first, with raw correlations and scatter plots, as with the models, to give a simple illustration of the signs and strength of the relationships in the data, and a qualitative comparison with the models. which we correlate with inflation series of the corresponding frequency. Secondly, we estimate these relationships with regressions. This allows us to more formally test for significance, and to control for other variables that might conceivably affect the price change distribution.

Correlations

Our sample period for the price data is 1977-2014, and the early, high inflation, part of the period is particularly important. We want to answer whether the dispersion and skewness of price changes move inversely with aggregate inflation, as predicted by most sticky price models, and in order to do this it is very helpful to see how the statistics change when inflation was high. However, we first verify that the frequency of price change rises with inflation, as found by Gagnon (2009) and Alvarez et al. (2011a). Table 2.2 reports correlations between the frequency of price change and inflation, and Figure 2.7 is the empirical counterpart to Figure 2.1 from the simulations (scatter plots of the average frequency and inflation for the months in the sample, for both the weighted mean and median frequency).

The table and figure confirm that there is a positive association between the frequency and inflation, although this is considerably clearer for the median than the mean frequency. As argued in the previous studies that had looked into this relation, this provides strong evidence against the Calvo assumption of time-dependent price setting. Figures 2.8 and 2.9 illustrate the other correlations that are presented in Table 2.2: those involving quarterly and annual averages of inflation and the frequency, and here the same pattern holds.

Next, we look at the results for the moments that our discussion has focused on: the dispersion and skewness of price changes. Tables 2.3 and 2.4 report the correlations for the dispersion and skewness respectively. Our main results is that while there does seem to be a clear negative relationship between inflation and dispersion, there is no such relation between inflation and skewness. Indeed, for both measures of skewness (moment skewness and Kelly skewness; "Skewness" in the tables and graphs refers to moment skewness), the correlation is either strongly positive (over the whole sample period) or close to zero (post-1984). This can also be seen in Figures 2.13 and Figures 2.16, [5], which are scatter plots illustrating the correlations (with each period corresponding to a month).

Figures 2.11-2.18 show these correlations with the quarterly and annual measures (including the Kelly Skewness correlation using annual data), illustrating how the same patterns hold.

[5] In the following figures, the left panel uses the statistics estimated using all available observations, while the right panel uses the estimates that exclude price changes below 1% in absolute value in ELI's deemed problematic by Eichenbaum et al. (2013).

Table 2.2: Corr(Frequency, Inflation)

	Weighted Median					
	Monthly		Quarterly		Annual	
	1977-2014	1985-2014	1977-2014	1985-2014	1977-2014	1985-2014
Raw	0.575	0.399	0.671	0.536	0.764	0.618
Smoothed	0.769	0.552	0.785	0.628	-	
	Weighted Mean					
Raw	0.311	-0.019	0.314	-0.216	0.374	-0.243
Smoothed	0.371	-0.337	0.36	-0.295	-	

Table 2.3: Corr(IQR, Inflation)

	All Observations					
	Monthly		Quarterly		Annual	
	1977-2014	1985-2014	1977-2014	1985-2014	1977-2014	1985-2014
Raw	-0.602	-0.446	-0.716	-0.665	-0.776	-0.751
Smoothed	-0.675	-0.706	-0.719	-0.742	-	
	EJRS					
Raw	-0.666	-0.434	-0.711	-0.689	-0.775	-0.779
Smoothed	-0.792	-0.701	-0.709	-0.769	-	

Table 2.4: Corr(Skewness, Inflation)

	All Observations					
	Monthly		Quarterly		Annual	
	1977-2014	1985-2014	1977-2014	1985-2014	1977-2014	1985-2014
Raw	0.265	0.084	0.345	0.067	0.473	0.122
Smoothed	0.506	0.136	0.474	0.133	-	
	EJRS					
Raw	0.272	0.068	0.327	0.053	0.447	0.102
Smoothed	0.462	0.144	0.452	0.105	-	

Table 2.5: Corr(Kelly Skewness, Inflation)

	All Observations					
	Monthly		Quarterly		Annual	
	1977-2014	1985-2014	1977-2014	1985-2014	1977-2014	1985-2014
Raw	0.584	0.069	0.674	-0.106	0.744	-0.165
Smoothed	0.696	-0.067	0.697	-0.199	-	

Figure 2.7: Frequency of Price Change and Inflation, Monthly

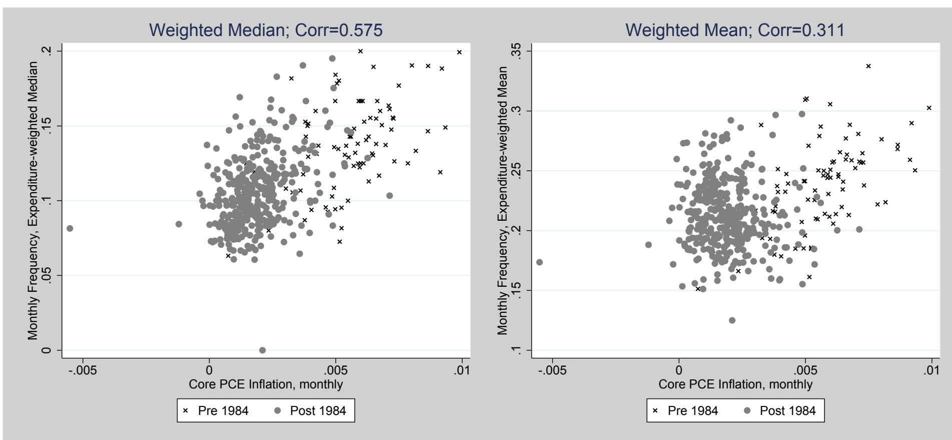

Figure 2.8: Frequency of Price Change and Inflation, Quarterly

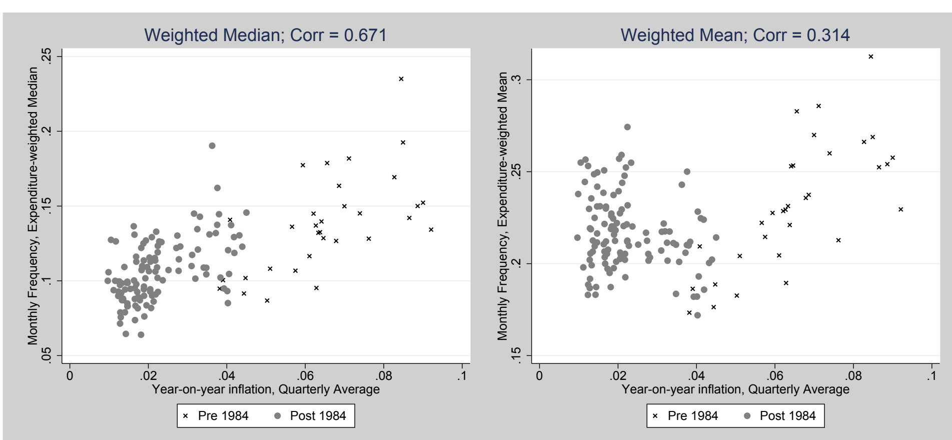

Figure 2.9: Frequency of Price Change and Inflation, Annual

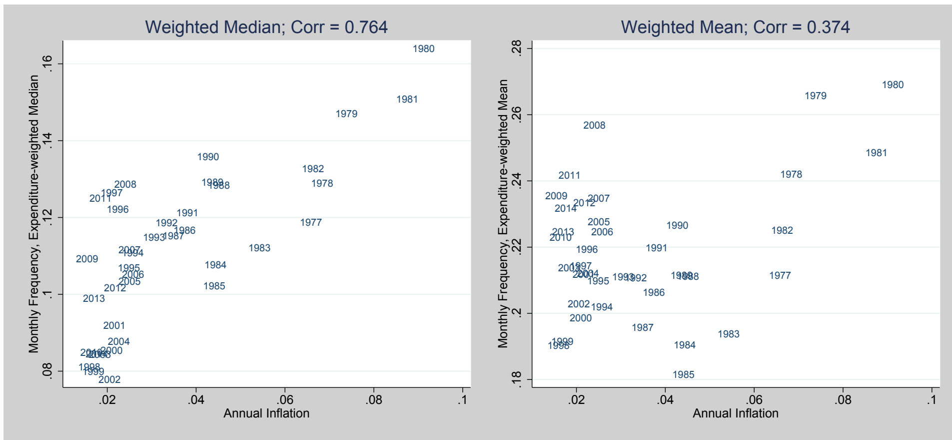

Figure 2.10: IQR & Inflation, Monthly

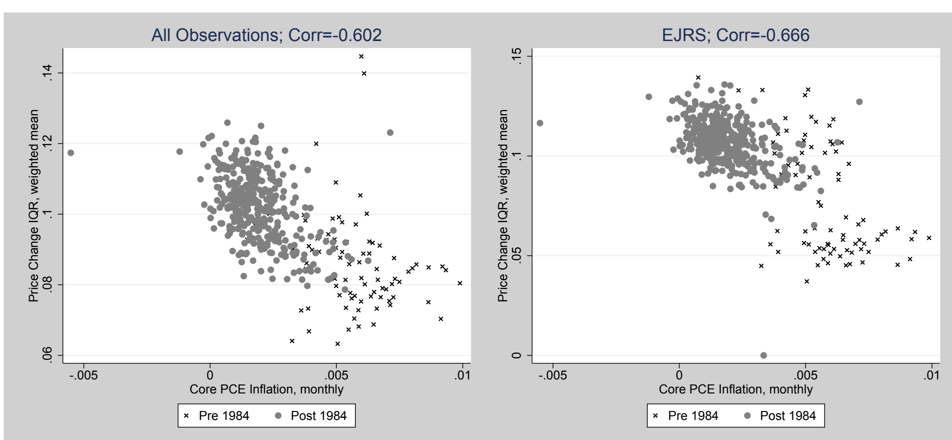

Figure 2.11: IQR & Inflation, Quarterly

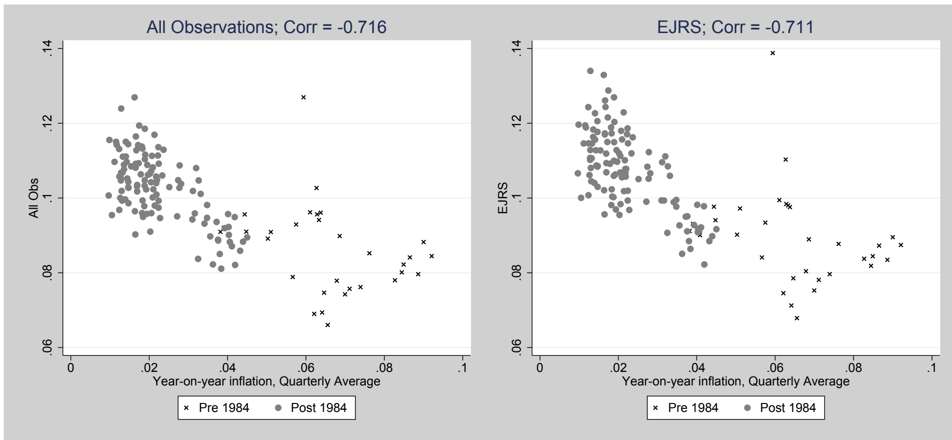

Figure 2.12: IQR & Inflation, Annual

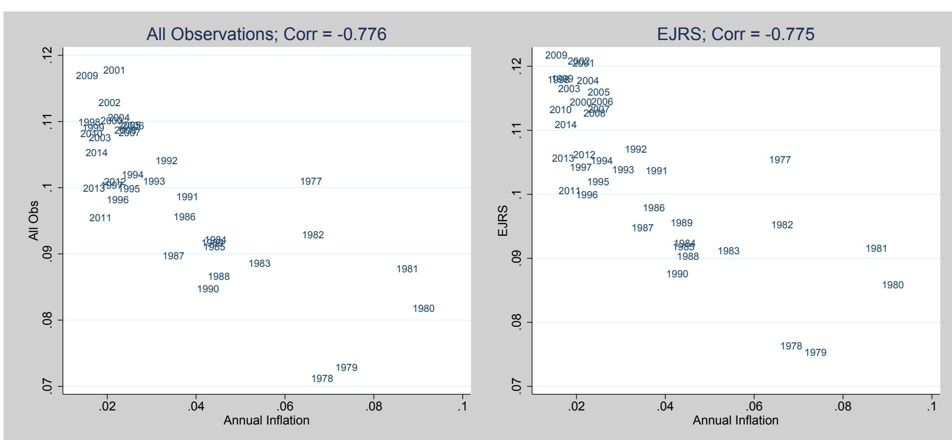

Figure 2.13: Skewness & Inflation, Monthly

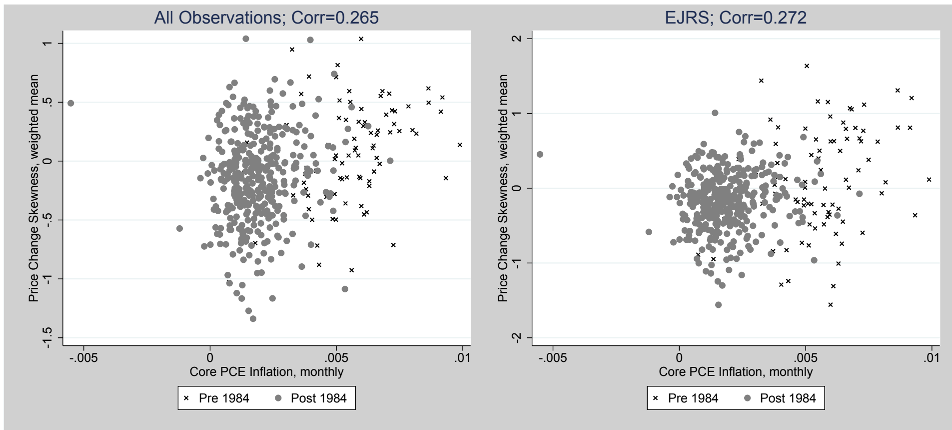

Figure 2.14: Skewness & Inflation, Quarterly

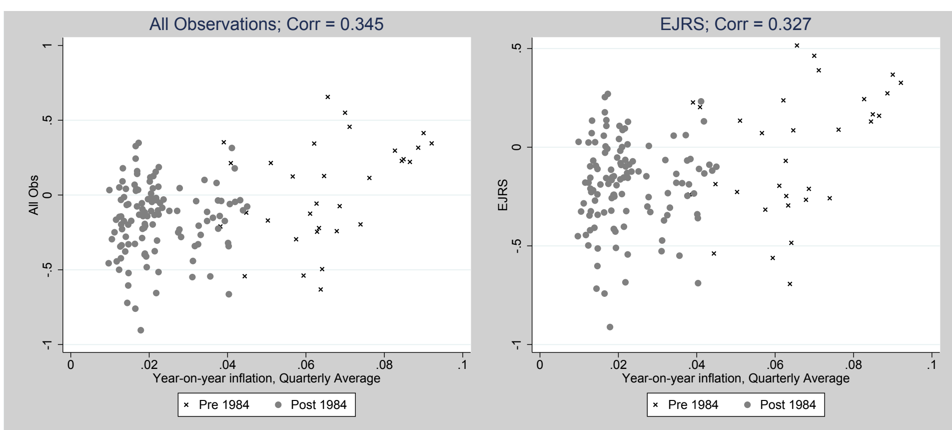

Figure 2.15: Skewness & Inflation, Annual

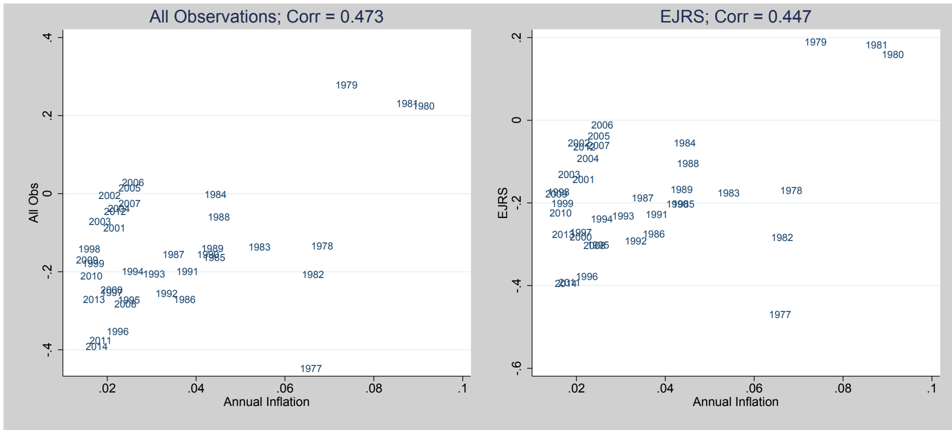

Figure 2.16: Kelly Skewness & Inflation, Monthly, correlation= 0.58

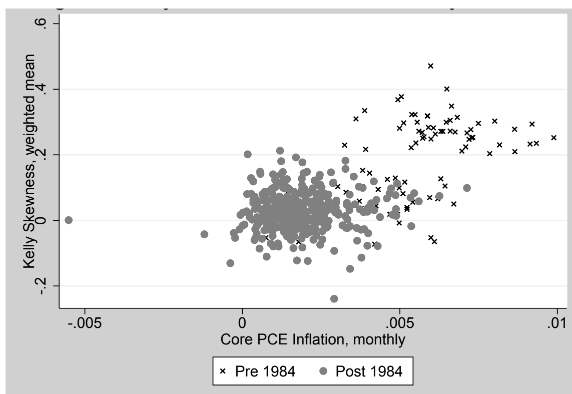

To summarize, we find first that the dispersion of price changes falls sharply with inflation throughout the sample period. Second, the skewness, while varying over time, does change with inflation in a systematic way for low levels of inflation. However, there does seem to be a positive relationship when inflation is high. We see this from the different correlations for the different sample periods (which roughly correspond to the high and low inflation periods). Finally, all these patterns hold true regardless of whether we exclude potentially spurious small price changes or apply seasonal adjustment and smoothing to the data series.

Figure 2.17: Kelly Skewness & Inflation, Quarterly, correlation= 0.67

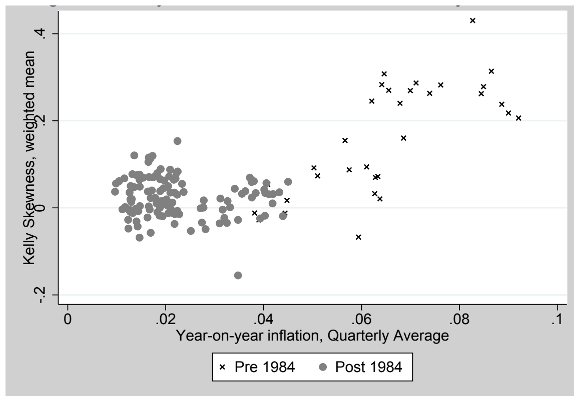

Figure 2.18: Kelly Skewness & Inflation, Annual, correlation= 0.7

Next, we formalize this analysis with linear regressions.

Regressions

Although the correlations and scatter plots provide a general picture of what the data shows on the relationships in question, we turn to regressions to determine whether these correlations are statistically significant, and to consider different control variables. However, the question of interest for us is not merely whether they are statistically significant from zero, but also whether they are significantly different from what the models predict. To do this, we estimate regressions of the frequency, dispersion (inter-quartile range) and skewness (both moment and Kelly skewness) of the price change distribution on inflation, with different specifications allowing for different sets of controls and sample periods. As before, we run the regressions both on the whole sample period and on only after 1984. This allows us to see if the relationship looks different between the low and high inflation periods. The regressions all take the following form:

$$y_t = \alpha + \beta\pi_t + \gamma Controls_t + e_t,$$

where y_t denotes the different price change moments (frequency, dispersion, and skewness). Controls are included to address the fact that many important changes occurred in the U.S. monetary environment over our sample period, which could conceivably have a direct effect on the price change distribution. Since expected inflation could affect firms' price setting decisions separately from present realized nominal shocks, we include expected inflation (measured by the University of Michigan Survey of Consumers) as a control. We also inlcude dummy variables for the different Federal Reserve chair's times in office, to control for differences in the conduct of monetary policy. The different specifications cover different combinations of controls (no controls, Fed dummies only, or Fed dummies with expected inflation) and the different periods. Tables 2.6 to 2.9 show the estimates for β from these different specifications, with the standard errors below them. All standard errors are calculated according to Newey and West (1987), and allow for serial correlation in the residuals.

Table 2.6: Coefficients for Frequency Regressions

	Weighted Median		Weighted Mean	
Specification	1977-2014	1985-2014	1977-2014	1985-2014
All	0.708***	0.777***	0.164	0.018
	(0.071)	(0.224)	(0.203)	(0.196)
Fed Dummies	0.728***	0.810***	0.686***	0.339**
	(0.095)	(0.208)	(0.104)	(0.167)
Inflation Only	0.771***	0.587**	0.438***	-0.087
	(0.237)	(0.252)	(0.108)	(0.236)

Note: Significant at 1% level (** at 5% level; * at 10% level). This table reports the correlation coefficients from regressions of the weighted average (median and mean) frequency of price changes on aggregate CPI inflation. The regressions are run using quarterly series, where quarterly inflation is defined the mean of the 12-month log changes in the CPI for the three months in every quarter. The different cells indicate different specifications, which change with respect to the sample period used and inclusion exclusion of small price changes (columns), and what controls are used. Standard errors that are consistent for heteroskedasticity and auto-correlation of the residuals (Newey-West) are reported. The same observations apply to the other regression tables, which report coefficients on inflation in regressions with other dependent variables.

These results confirm what the correlations showed: the frequency of price change rises with inflation (although for the mean frequency this is not so clear), the relationship between dispersion and inflation is negative and statistically significant in all specifications and sample

Table 2.7: Coefficients for IQR Regressions

	All Observations		EJRS	
Specification	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	-0.296***	-0.428***	-0.327***	-0.491***
	(0.042)	(0.070)	(0.046)	(0.082)
Fed Dummies	-0.186***	-0.414***	-0.204***	-0.476***
	(0.038)	(0.077)	(0.044)	(0.089)
Fed & Expected Infl	-0.257***	-0.222**	-0.261***	-0.224***
	(0.089)	(0.086)	(0.095)	(0.092)

Table 2.8: Coefficients for Skewness Regressions

	All Observations		EJRS	
Specification	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	3.936***	1.732	3.501***	1.108
	(0.827)	(1.641)	(0.828)	(1.534)
Fed Dummies	4.309***	1.541	3.928***	1.130
	(1.012)	(1.857)	(0.966)	(1.705)
Fed & Expected Infl	2.665	3.634	1.947	2.963
	(2.788)	(3.279)	(2.538)	(2.985)

Table 2.9: Coefficients for Kelly Skewness Regressions

	All Observations	
Specification	1977-2014	1985-2014
Inflation Only	2.499***	0.320
	(0.354)	(0.454)
Fed Dummies	2.439***	0.710*
	(0.363)	(0.423)
Fed & Expected Infl	1.658	0.942
	(0.948)	(0.595)

periods. The skewness correlation, however, is significantly positive for the whole sample, but not significantly different from zero when the early, high-inflation period is excluded (and this applies for both measures of skewness). This indicates (as we can also see from the scatter plots), that this relation is close to flat for low inflation periods, but clearly positive for high inflation periods. The fact that the skewness of price change is higher on average in high inflation periods is important, because it also goes against the menu cost models' predictions at high values of steady-state inflation, as we showed in Figure 2.5. It is also notable that the skewness coefficients change considerably when expected inflation is included as a regressor. Since expected inflation is very highly correlated with realized inflation, the estimates are much less precise (as shown by the high standard errors), so this is not surprising. However, this makes little difference to the comparisons with the coefficients predicted by the models, which we turn to with Table 2.10.

Table 2.10: Coefficients on Inflation for Price Change Moment

Model	Frequency	IQR	Skewness	Kelly Skewness
Golosov & Lucas	0.139	-0.937	-17.7	-0.40
Multisector Menu Cost	0.143	-0.218	-5.39	-4.33
Midrigan	0.348	-0.896	-9.84	-6.53
Observation and Menu Costs	.268	-0.071	-4.32	
Calvo	-0.003	0.040	2.93	1.00
Rational Inattention	0.001	0.007	3.03	

The table presents the coefficients on inflation from regressions of the same type, but run on simulated data from the different models. The first four models (menu cost models) have negative coefficients for the inter-quartile range, although for all but the multi-sector model, they are outside the 95% confidence intervals of the coefficients that we estimate. However, the disagreement with the data is much starker with the skewness coefficients. These are all very far outside the confidence intervals that we estimate for the skewness coefficients under all specifications, and the same is true for Kelly skewness[6]. We summarize our findings in Table 2.11 below, which "updates" Table 2.1 by adding the signs of the correlations in the data to those predicted by the models.

Table 2.11: Correlation of inflation and moments

Model	corr(π, frequency)	corr(π, Std. Deviation)	corr(π, Skewness)
Calvo	0	0	+
Golosov and Lucas	+	-	-
Multi Sector Menu Cost	+	-	-
Midrigan	+	-	-
Observation and Menu Costs	+	0	-
Rational Inattention	+	0	+
Data	+	-	+

For each model, the signs that match the data are colored in blue, while those that do not are red. We do this to highlight the fact that in the broad class of state-dependent price setting models that we consider, none match the data in all the dimensions that we have presented. In particular, this highlights the usefulness of the inflation-skewness correlation as a statistic to test the existing menu cost models. As we have already argued, these models make a counterfactual prediction with this statistic because of the state-dependence that underlies them. It is also worth noting that the Calvo and rational inattention models have the same predictions, and therefore disagree with the data in the same ways.

[6]The one exception is the coefficient for the Golosov and Lucas model, which is much smaller in magnitude than in the other menu cost models, and is marginally accepted in the specification that restricts the sample to the post-1984 period and uses only Fed chair controls. It appears that the value of the Kelly Skewness is extremely sensitive to the unusual shape of the price change distribution (bi-modal) in this model, leading to this weak relationship. The model's Kelly Skewness coefficient is still rejected in all the other specifications, however.

2.4 Conclusion

The literature on sticky prices has made extensive use of price micro data to discipline models of price setting, and in this way the models have conformed more and more to important aspects of the dynamics of price changes. This line of work has notably enabled the study of monetary non-neutrality to be more grounded in data. However, an important limitation of the work done so far is that it has mostly used data for low inflation environments. Since the models in question are designed to study how prices respond to aggregate shocks, it is helpful to be able to observe the behaviour of price changes under large aggregate shocks and high inflation.

Our paper contributes to this by using price data from the U.S. going back to the late 1970's to compare how the price change distribution changes with inflation, to the predictions of a wide range of sticky price models. Our main finding is that the menu cost models that have been most used in the literature fail to match the positive relationship between inflation and the skewness of price changes in the data, because they uniformly predict a sharp negative relationship. In addition, we argue that this relationship, although not obvious at first site, follows very intuitively from the selection effect that is central to menu cost models. We also show how a model with random menu costs can overcome this problem when the distribution of menu costs features a significant probability of very high and very low menu costs, making it resemble a Calvo model and weakening the selection effect. Finally, this model predicts a degree of monetary non-neutrality that is considerably higher than what is predicted by the Golosov and Lucas model, and higher still than the Midrigan model.

The distinction between menu cost and Calvo models, or between state- and time-dependent pricing models has taken an important place in this literature. Much work has been done to show how these two ways of modelling pricing stickiness yield such different implications on monetary non-neutrality, and to determine which models are best at matching empirical facts. Our paper contributes to this line of work by introducing statistics not previously considered that are very useful to discriminate between the different models. In addition, we follow Woodford (2009) in presenting the distinction between time- and state-dependent models as a continuum, or spectrum. Woodford (2009) shows how different values for the firm's cost of processing information leads to a different point on this spectrum. In contrast, our approach is agnostic as to what ultimately underlies the randomness of menu costs that allows our model to span the time versus state dependent spectrum. Instead, we determine what point on the spectrum is most consistent with the data. Future research could combine these two approaches to gain a better understanding into the nature and importance of the informational constraints that underly price rigidity. For now, along with Nakamura and Steinsson (2010) and Midrigan (2011), we show that the assumption made by Golosov and Lucas (2007) of firms facing a single, constant menu cost is starkly at odds with many aspects of the price data, and that monetary policy can be expected to have substantial and persistent effects on real economic activity.

Chapter 3

Monetary Non-Neutrality and the Selection Effect in a Random Menu Cost Model

3.1 Introduction

Much of the literature on sticky prices (including the other chapters in this work) have studied the distinction between state-dependent and time-dependent price stickiness. This distinction is extremely important, because the two different classes of models make extremely different predictions on how large the real effect of monetary policy (or monetary non-neutrality) is. One approach to this line of work, notably taken by Woodford (2009), has been to consider state-dependent (or menu cost models) and time-dependent (for example, the Calvo model) as specific cases of a more general approach to price stickiness. More specifically, Woodford (2009) applies the rational inattention framework to the standard monopolistically competitive price setting problem, and shows that the resulting model can nest the Calvo and menu cost models as extreme cases, spanned by different values for the cost of processing information. Indeed, with an infinite cost of information, firms behave as in a Calvo model, choosing to receive random signals about when to change their prices; while in the case of free information, the model is equivalent to a standard menu cost model. Beyond rational inattention, Caballero and Engel (1993) proposed a hazard function approach to studying the link between discrete micro-level decisions by agents and aggregate variable dynamics, applying it to sticky prices in particular. The hazard function, in this context, is the probability of a price changing, as a function of the gap between the current price and the optimal price. Different types of sticky price models will imply different hazard functions, making this approach very general. Caballero and Engel (2007) showed how the hazard function can shed insight on the relationship between price stickiness and monetary non-neutrality, and the following chapter of this work will show how this function can be estimated from price micro data.

In this chapter, I consider a different type of generalized model that nests the menu cost and Calvo cases, with a very particular goal in mind. Chapter 2 had shown how the selection effect in sticky price models induces very strong patterns between aggregate inflation and the shape of the price change distribution. It was shown, notably, that in menu cost models the skewness and dispersion of price changes fall very sharply with inflation, while in the data only the dispersion seems to fall with inflation. This strongly suggests that the degree of selection implied by menu cost models is too high. However, the fact that the frequency of price change appears to rise sharply with inflation is also strong evidence that the complete absence of selection assumed by the Calvo model is inconsistent with the data. I therefore present an extension to the standard Golosov and Lucas (2007) model in which the menu cost that firms must pay if they choose to change their price is random over time and across firms.

Inspired by the Dotsey et al. (1999) random menu cost model, this model can also nest the Calvo and standard menu cost models, depending on the distribution of menu costs. Indeed, a degenerate menu cost distribution of menu costs reduces the model to a standard, fixed, menu cost model. On the other hand, if the menu cost can only take the values of zero and infinity (with a certain probability), the model is identical to a Calvo model. Such a model would feature no selection because the only determinant as to whether a price changes would be the luck of the draw of the menu cost (as in Calvo). I use this model to find a distribution of menu costs that allows the model to match both the average frequency and size of price change (which have been used to calibrate many of the existing sticky price models), and the different inflation correlations presented in the previous chapter. The key factor that will distinguish distributions for these purposes is the amount of selection implied, and the distribution that allows the model to match these moments (and in particular to predict a non-negative inflation-skewness correlation) implies a very weak selection effect. This, in turn, means that the model predicts a very high degree of monetary non-neutrality: around six times greater than in Golosov and Lucas (2007), two times greater than in Midrigan (2011), but around 70% as high as in Calvo (which, featuring no selection, still represents the upper bound of non-neutrality).

The remainder of this chapter is organized as follows. Section 3.2 presents the random menu cost model, and further discusses the relationship with existing generalized approaches to sticky prices. Section 3.3 presents the solution method to the model, the moments used for calibration, and the results on monetary non-neutrality. Finally, section 3.4 presents some concluding remarks.

3.2 Random Menu Cost Model

The model in this paper fits into the general set-up of the models examined in the previous chapter. Households choose consumption and labor supply to maximize expected discounted utility:

$$E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left[logC_{\tau+t} - \omega L_{\tau+t} \right]$$

Consumption consists of a composite of individual varieties produced by monopolistically competitive firms. Aggregate consumption is a CES aggregator, which yields the usual demand function for individual varieties, indexed by z (of which there is a continuum with unit mass):

$$c_t(z) = \left(\frac{p_t(z)}{P_t}\right)^{-\theta} C_t$$

Firms produce their variety using only labor, with a linear production technology:

$$y_t(z) = A_t(z)L_t(z)$$

The demand system and technology faced by the firm is the same as in the existing models, but we generalize the price setting problem in the following way: the menu cost faced by each firm every period is random. Formally, the period profit function of the firm takes on this form:

$$\Pi_t(z) = p_t(z)y_t(z) - W_tL_t(z) - \chi_t(z)W_tI\{p_t(z) \neq p_{t-1}(z)\},\ \chi_t(z) \overset{iid}{\sim} G(\chi)$$

The difference with the Golosov and Lucas (2007) model is that now the menu cost can vary over time and across firms, the difference with the Midrigan (2011) model is that the distribution of menu costs is generalized, and as opposed to the Nakamura and Steinsson model, the menu cost for any given firm here varies over time. The first thing that can be shown is how this model nests some of the models presented in the previous chapter:

$$Golosov\,\&\,Lucas:\ \chi_t(z) = \bar{\chi}\forall t, z$$

$$Calvo: \ \chi_t(z) = \begin{cases} 0, & Prob = f \\ \\ \infty, & Prob = 1 - f \end{cases}$$

$$Midrigan: \ \chi_t(z) = \begin{cases} 0, & Prob = p_z \\ \\ \chi, & Prob = 1 - p_z \end{cases}$$

One can already see how these different cases imply different degrees of selection in price setting. In the degenerate distribution case, the menu cost never varies, such that the firm's decision as to whether to change its price will depend entirely on how mis-aligned its price is. In the Calvo case, however, firms either get to change their price for free, or not at all. In this case, the price change decision depends entirely on what value was drawn for the menu cost in that period. In Midrigan (2011), we have an intermediate case, in which firms with very mis-aligned prices will be more likely to change their prices, but some price changes will also occur simply because the firm randomly received a free price change opportunity. The assumption of random menu costs is similar to that made by Dotsey et al. (1999)[1], but this model still fits into the general framework of Golosov and Lucas (2007), which will make it possible to calibrate the model using micro price data and evaluate the degree of monetary non-neutrality.

Finally, the processes for shocks are the same as the ones used previously. Firms' idiosyn-

[1]The key differences with Dotsey et al. (1999) are that their model does not include idiosyncratic shocks, that it does include capital as an input to production, and that they did not have a way of using information from price micro data to place restrictions on the menu cost distribution, which is what the present exercise is about.

cratic productivity shocks follow a Poisson process:

$$logA_t(z) = \begin{cases} \rho logA_{t-1}(z) + \epsilon_t, & Probability = p_\epsilon \\ \\ logA_{t-1}(z), & Probability = 1 - p_\epsilon \end{cases}, \ \epsilon_t \overset{iid}{\sim} N(0, \sigma_\epsilon^2)$$

Nominal aggregate demand follows a random walk with drift:

$$logP_tC_t = logS_t = \mu + logS_{t-1} + \eta_t, \ \eta_t \overset{iid}{\sim} N(0, \sigma_\eta^2)$$

As before, it will be the aggregate shocks that drive variations in inflation.

3.2.1 Background on Random Menu Costs

I choose to modify the model in this way for several reasons. The first is that it yields a very general model that nests the menu cost models we have considered, as described above. In addition, this approach has a close relation to another, even more general approach already pursued by Caballero and Engel in a series of papers (1993, 2006, 2007). They propose thinking about price adjustment through the price adjustment hazard function[2], which is the probability of a price change occurring as a function of the deviation of the current price from its optimal value (p^*):

$$H(x) = P(\Delta p | \, p^* - p = x)$$

[2] It must be noted that this is distinct from the hazard function of price change estimated by Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008), which is the function $\lambda(t)$ that gives the probability that a price will change after t periods, given that it has already stayed constant for t periods. This function gets at timing of price changes, and at whether prices become more likely to change the longer they have stayed constant.

Any of the models considered so far will imply a price adjustment hazard function, so the hazard function can be a helpful object to summarize the important features of each model. This function also describes the degree of the selection effect, as it indicates to what extent prices have a higher chance of adjusting the more mis-aligned they are. Caballero and Engel (2007) show that, within this framework, aggregate price flexbility (or the inverse of monetary non-neutrality) can be expressed as the sum of two components:

$$\int H(x)f(x)dx + \int xH'(x)f(x)dx$$

where $f(x)$ is the probability density of the desired price gap, x. The first term in this sum is the frequency of price change, and the second is what Caballero and Engel refer to the extensive margin, which incorporates the selection effect. To illustrate how this works, the hazard function corresponding to the Calvo model is a constant (the average frequency of price change), so that the first term is equal to that constant, and the second is zero. This shows that in general, and as long as the hazard function is increasing in the absolute value of the price gap (which is true for all models considered), the Calvo model gives a lower bound on the degree of flexibility. In the random menu cost model, a particular menu cost distribution will imply a particular hazard function, and will therefore determine aggregate flexibility (and monetary non-neutrality) as shown by the expression above. In this way, there is a very tight relation between these approaches[3].

[3]It would naturally also be interesting to directly estimate this hazard function. In this paper we continue to work in the menu cost framework to maintain the structure of a General Equilibrium model and obtain a quantitative response to the question of monetary non-neutrality. However, Caballero and Engel (2006) attempt to do this, although the price change moments that they had access to were limited in the informational value they provided for this. By the same logic we have put forth in this paper, the higher moment price change correlations will be very informative to estimate the hazard function, and we pursue this in chapter 4.

A more structural approach to price stickiness that is also related to mine is Woodford (2009)'s model of rational inattention. He shows that by varying the cost of processing information, price setting under rational inattention in the style of Sims (2003) can also nest, as extreme cases, the single menu cost model (free information) and the Calvo model (infinitely costly information), as well as the spectrum in between, which he also describes with the adjustment hazard function implied by different information costs. In addition, there will also be a mapping between the value for the cost of information in that framework, and a specific distribution of menu costs in ours. While I do not derive this mapping, it is possible that the rational inattention framework (or another type of informational constraint) could provide a micro-foundation for the distribution of menu costs that I assume to be general at first, and then adjust to allow the model to fit the empirical facts[4].

3.2.2 The Distribution of Menu Costs

Introducing random menu costs allows one to determine the extent of state-dependence present in the model, or to what extent firms choose when to change their prices. One extreme case of this is of course perfect price flexibility, or firms being free to change their prices every period without facing any kind of cost for doing so (although this is inconsistent with the fact that most prices do not change on any given month). But right after this comes a menu cost environment such as the one in Golosov and Lucas (2007): firms are

[4] As Woodford (2009) also points out, the direct empirical evidence on the actual costs of price adjustment put forth by Zbaracki et al. (2004) indicates that the most important part of those costs are related to the process of gathering the necessary information for a price review. In addition, Anderson and Simester (2010) give evidence on how price changes can antagonize consumers, which introduces costs to changing prices. To the extent that the menu costs in the menu cost framework represent these costs, I believe that it is plausible that the menu costs are random to some extent, and vary across firms and time. This lends plausibility to the random menu costs assumption, although I leave the explicitly modeling of the informational constraints or consumer considerations that underly it to future research.

still able to choose when to change their prices, but are subject to a fixed cost (that is small in typical calibrations, to match the frequency of price change in the data). Adding randomness to the menu cost makes the price change decision more exogenous to the firm, as an additional dimension of the problem (how much changing the price will cost) is now outside the firm's control (with the extreme being the Calvo model, where the opportunity to change price is completely exogenous). The Midrigan model (both in Midrigan (2011), and the simplification of it that we present) goes in this direction, and as a result the degree of monetary non-neutrality in that model is much higher. The results from chapter 2 suggest that a model would need even more exogeneity (but less than the Calvo model) to match the empirical facts that were presented. Therefore, I parametrize the distribution of menu costs in a way that makes it possible to do this.

There are two important features that the menu cost distribution will need in order to achieve this: a positive probability of the menu cost being zero (of a free price change), which eliminates the "Ss" band or inaction region in the price setting problem, as some firms, facing a free price change, will choose to change their prices even if it is by a small amount. However, the Midrigran model already includes this, and as we have shown it also predicts a counterfactual inflation-skewness correlation. The other feature is that there must also be a positive and considerable probability that the menu cost will be very high, so high that firms will not choose to change their price when faced with these menu costs. This is important, because in the existing models, the skewness of price changes falls with inflation because a positive aggregate shock induces more firms that face a positive menu cost to pay it, effectively pushing them over a threshold, leading to an important shift in the shape of the distribution. Having a positive probability of very high menu costs means that fewer firms will be pushed over this threshold, weakening this effect. It is also helpful to note that the Calvo model contains both of these features in the extreme, as it gives a positive probability of a free price change, and in all other cases the menu cost is infinite. Because of this, one can consider that the menu cost distribution in the generalized model will incorporate a strong "Calvo feature", without going all the way to the Calvo extreme.

In order to achieve this, I present a relatively flexible distribution for menu costs. I assume that menu costs are iid across time and firms, so that every period each firm draws a menu cost χ from a mixed distribution. First, with a certain probability, the menu cost is zero, and otherwise it is drawn from a continuous distribution:

$$\chi = \begin{cases} 0, & Prob = p_z \\ \tilde{\chi}, & Prob = 1 - p_z \end{cases}, \; where\ F(k) = P(\tilde{\chi} \leq k) = 1 - e^{-\lambda k^{\alpha}}$$

In the simplified version of the Midrigan model, the menu cost was either zero or a fixed positive value. The difference here is that instead of the positive value being fixed, it is drawn from a non-degenerate distribution. This distribution is a transformation of the exponential distribution (it is the same when $\alpha = 1$), and shares the important feature that the random variable is always positive. The difference is that α governs the curvature of the distribution function, which roughly corresponds to the fatness of the tails. Figure 3.1 shows how the shape of the cumulative distribution function changes with α:

For our purposes, what is important is that for low values of α, the probability of very low values is relatively high, but the probability of very high values is also quite high. When α is high, these extreme probabilities are low, and as α rises, the density concentrates on

Figure 3.1: Shape of Menu Cost CDF for Different α

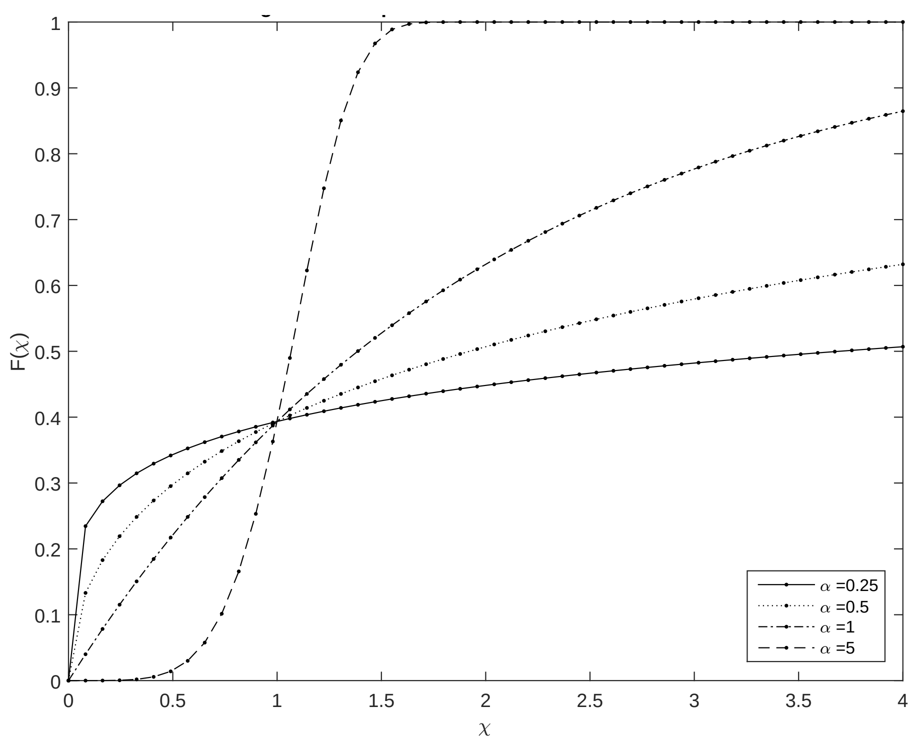

one value, approximating the case of a unique menu cost.

3.3 Solution and Results

This set-up has introduced new parameters, relative to the models that we considered in chapter 2: the inverse of the average menu cost (λ), and the curvature of the menu cost distribution (α). The other parameters important for the firm's price setting problem are the variance of the idiosyncratic shocks (σ_ϵ^2), the arrival probability of the shocks (p_ϵ), and the probability of a free price change (p_z) which was used in the Midrigan model. I set these parameters so that the model can match the empirical facts that were discussed previously, which can be divided into two categories:

1. From existing models: although these have not been the focus of the discussion, all the existing models match the average monthly frequency of price change and the average size of price change. My model therefore matches the median of these statistics measured in our data. In addition, the empirical work has confirmed that, as previous studies had shown, the correlation between inflation and the frequency of price change is positive, so this model also matches this fact.

2. New moments: like the existing menu cost models, and consistent with the data, this model will imply a strongly negative correlation between inflation and the dispersion of price changes. The novelty will be that the implied correlation between inflation and the skewness of price changes will be non-negative, as in the data.

Table 3.1 presents the parameter values that we choose to match these moments, and Table 3.2 shows the moments attained by the model, compared to their empirical values.

Table 3.1: Parameter values

Parameter	Description	Value
λ	Inv. average menu cost	0.1925
α	Fatness of tails of MC	0.27
p_z	P(zero MC)	0.056
p_ϵ	P(idio. shock)	0.345
σ_ϵ	Size of idio. shocks	0.101

Table 3.2: Simulation results

Moment	Model	Data
Avg. Frequency	11.3%	11.3%
Avg. Size	8.0%	8.0%
Corr(IQR,π)	-0.59	-0.70
Corr(Skew,π)	0.05	0.39
Corr(Freq,π)	0.58	0.63

The first two moments are matched almost exactly. For the empirical value of the correlations, we present the results for the quarterly correlations involving the raw data, including all time periods, and excluding suspicious small price changes (for dispersion and skewness), and the weighted median for the frequency. The model matches the dispersion and frequency correlations quite closely. However, the skewness correlation in the model is close to zero, while it is clearly positive in the data for the whole sample. Before explaining this in more detail, I illustrate these correlations with scatter plots for the generalized model under the calibration above in Figures 3.2-3.4.

While the skewness correlation in this model is lower than in the data, for the range of inflation that occurs in the simulations (0-6%), the correlation also appears to be close to zero in the data. I also consider a "long-run" correlation: solving the model for different values of

Figure 3.2: Frequency & Inflation in Generalized MC model, corr=0.58613

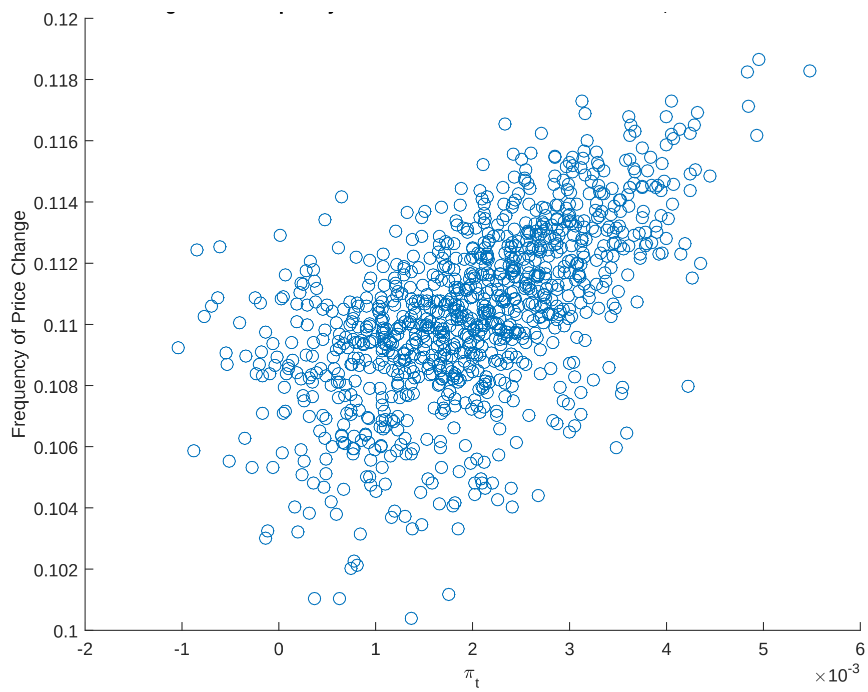

Figure 3.3: Dispersion & Inflation in Generalized MC model, corr=-0.45496

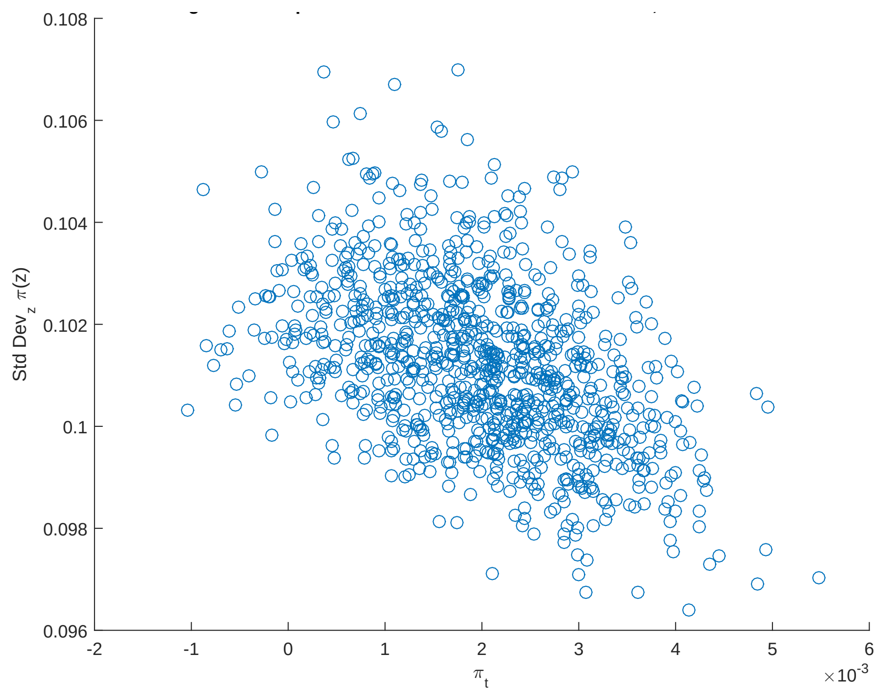

Figure 3.4: Skewness & Inflation in Generalized MC model, corr=0.048446

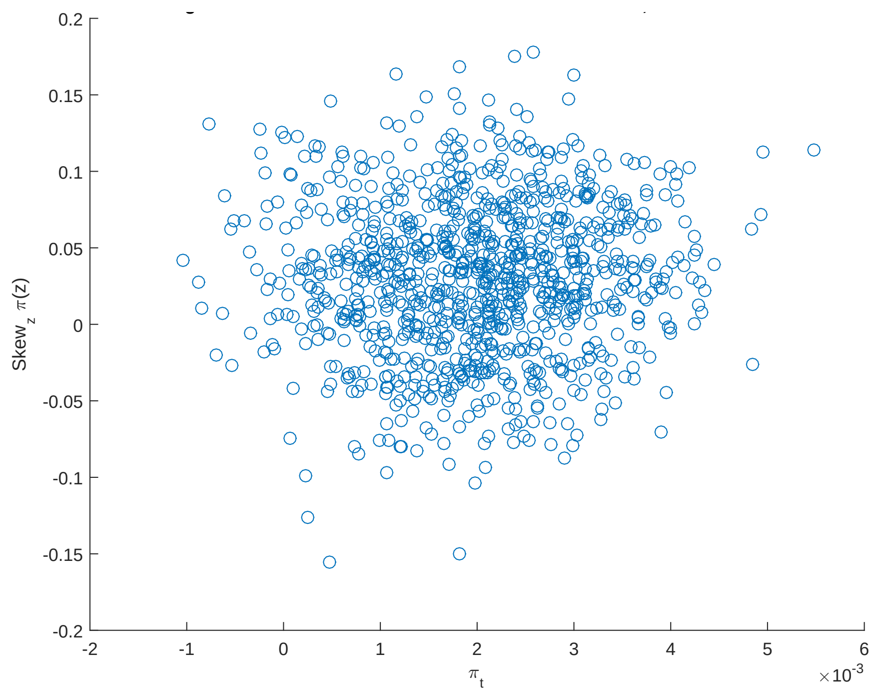

trend inflation, and computing the average skewness of price change for each of these values. I find that for higher steady-state inflation, the average level of skewness in the price change distribution rises, and the correlation between period-by-period price change skewness and inflation (the same correlations we have been focusing on so far) also rises. This result makes the model even more consistent with the data, as it shows that when steady-state inflation is higher (as it surely was in the early, high-inflation part of our sample), we should expect to see the skewness rising with inflation. This is indeed what we saw in the empirical analysis in chapter 2. In addition, this also makes this model stand out even more from the existing ones, as the other menu cost models feature a declining average price change skewness as steady-state inflation rises (and a period-by-period skewness correlation that is always negative). Figure 3.5 below shows this clearly by plotting the steady-state skewness correlations for the Midrigan model and the heteroegenous menu cost model separately.

Figure 3.5: Steady-State Skewness Correlation

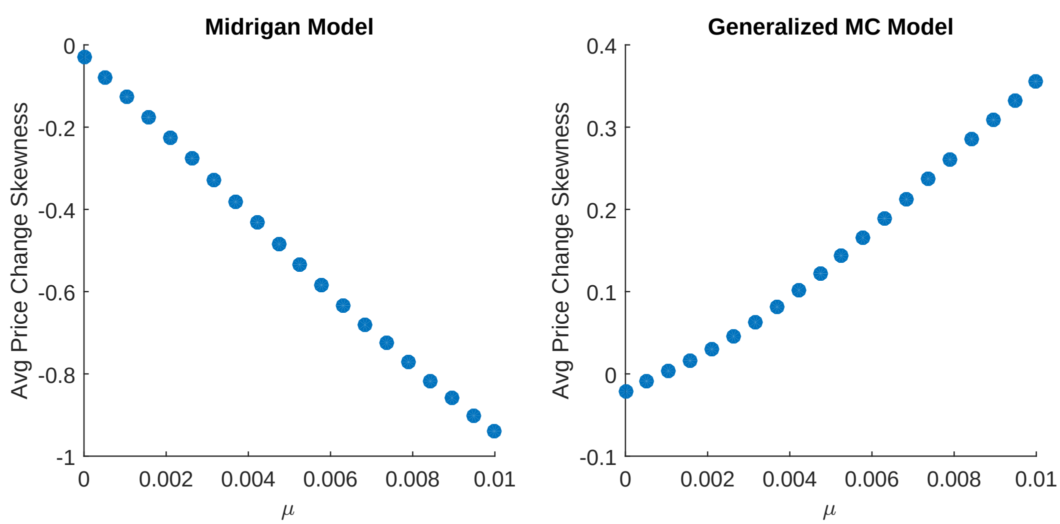

This pattern highlights how the steady-state (or trend) inflation plays an important role behind the model's non-negative skewness correlation. Indeed, positive trend inflation leads firms to expect positive future inflation when considering whether to re-set their prices. This will lead them to be less likely to cut their prices, even when facing an idiosyncratic (or aggregate) shock that would reduce their current desired price. This asymmetry in firms' willingness to cut prices also means that the left tail of the price change distribution will be less responsive to aggregate shocks, weakening the mechanism that led to the negative skewness correlation in the existing models.

What these results and figures make clear is that the generalized menu cost model that we presented, in making menu costs random in a way that weakens the selection effect, matches the important empirical facts that have been the focus of previous work on sticky prices as well as the existing models, and overturns the counterfactual prediction of these models that were emphasized previously. I now show what this means for the degree of monetary non-neutrality.

3.3.1 Monetary Non-Neutrality

Monetary non-neutrality in these models is defined as the variation in real consumption induced by the nominal aggregate demand shocks, which are the only aggregate shocks, and I compare this statistic for the Calvo model, the Golosov and Lucas and Midrigan menu cost models, and the generalized menu cost model. As mentioned above, making the menu cost distribution random in the way that I have proposed weakens the selection effect that is at work in menu cost models, so it is to be expected that this model would imply a greater degree of monetary non-neutrality. Table 3.3 below provides a quantitative illustration of this.

Table 3.3: Monetary Non-Neutrality

Model	$\text{Var}(C_t) * 10^4$
Golosov and Lucas	0.05778
Midrigan	0.17588
Generalized Menu Cost	0.33617
Calvo	0.47696

As Golosov and Lucas (2007) had famously shown, their model features a trivial amount of monetary non-neutrality compared to the Calvo model. Between the menu cost models, the major difference is between the baseline (Golosov and Lucas (2007)) and the others. Allowing for small price changes, as the Midrigan model does, leads to a very large increase in monetary non-neutrality, and this was emphasized by Midrigan (2011). However, the generalized model goes further by giving firms a large probability of effectively not being able to change their price, and yields an even higher level of non-neutrality. The Calvo model still has a higher degree of monetary non-neutrality. To further illustrate the differences between the models, Figure 3.6 shows the impulse response of real aggregate consumption to a one percentage point increase in nominal aggregate demand in the same four models:

Not only is the effect on real activity greater on impact in the random menu cost model, but the effect is also considerably more persistent than in the existing menu cost models. It is clear that while in the Golosov and Lucas model the real effect of a nominal shock is small and transient, it is not so in my model, which has used the inflation-skewness correlation to evaluate the strength of the selection effect. Finally, the Calvo model gives a closer approximation to monetary non-neutrality than the other menu cost models.

Figure 3.6: Impulse Responses in Models

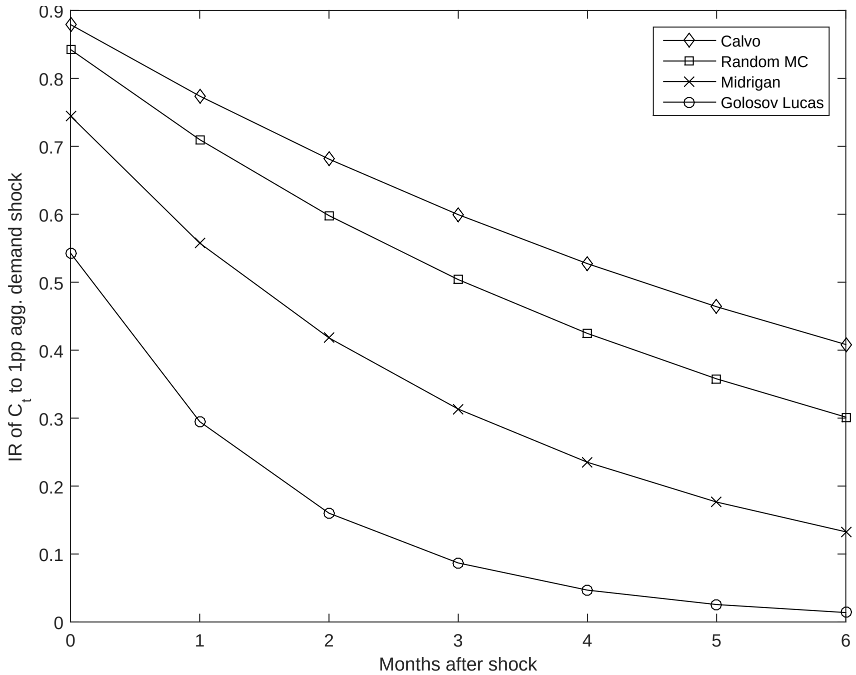

3.4 Conclusion

This chapter has presented a generalized random menu cost model that nests the standard menu cost model and the Calvo model. The model's flexibility allows me to determine the strength of the selection effect by adjusting the distribution of menu costs. Chapter 2 had shown that standard menu cost models predict a counter-factual negative relationship between aggregate inflation and price change skewness, and I find here that the random menu cost model can match the signs of all the different inflation correlations, as long as the distribution of menu costs implies a very weak selection effect. This exercise makes it possible to pin down the required strength of the selection effect quite precisely, and this result means that the model predicts a very high degree of monetary non-neutrality: considerably higher than in the Midrigan or Golosov and Lucas menu cost models, and slightly lower than in the

Calvo model. Beyond providing a precise answer to the question of monetary non-neutrality, an important contribution of of this work is to develop a quantitative random menu cost model, that generalizes the distinction between time and state-dependent models, and that can be readily applied to price micro data.

In addition, I follow Woodford (2009) in presenting the distinction between time- and state-dependent models as a continuum, or spectrum. Woodford (2009) shows how different values for the firm's cost of processing information leads to a different point on this spectrum. In contrast, the approach taken here is agnostic as to what ultimately underlies the randomness of menu costs that allows the model to span the time versus state dependent spectrum. Instead, I determine what point on the spectrum is most consistent with the data. Future research could combine these two approaches to gain a better understanding into the nature and importance of the informational constraints that underly price rigidity. For now, along with Nakamura and Steinsson (2010) and Midrigan (2011), I show that the assumption made by Golosov and Lucas (2007) of firms facing a single, constant menu cost is starkly at odds with many aspects of the price data, and that monetary policy can be expected to have substantial and persistent effects on real economic activity.

Chapter 4

The State-Dependent Price Adjustment Hazard Function: Evidence From High Inflation Periods

Shaowen Luo and Daniel Villar

4.1 Introduction

The size of monetary policy's effect on real economic activity is a question that has received much attention in the monetary economics literature. And with the new availability of large and detailed price data sets over the past decade, what was once a mostly theoretical debate has become more grounded in empirical facts. Indeed, different models have been proposed for sticky prices, and they have yielded very different implications for monetary policy. But price micro data sets (and notably the data underlying the U.S. Consumer Price Index, first used by Bils and Klenow (2004)) have made it possible to measure how sticky prices really are, and to discipline the key parameters of structural models. In chapter 2, we addressed the distinction between time-dependent and state-dependent models, which has been much of the focus of the monetary literature. In doing so, we worked with the two main structural models of sticky prices: the menu cost model and the Calvo model (and their variants), which we evaluated using an extension of the CPI micro data. In this paper, we analyze price stickiness taking a different approach: the hazard function approach originally proposed by Caballero and Engel (1993), which is, in many ways, model-free.

The debate about time-dependent models (such as the Calvo model) and state-dependent models (such as the menu cost model) has centered on the fact that in state-dependent models, firms that choose to change their price are those whose prices are most mis-aligned, while such selection is absent by assumption from time-dependent models.Caplin and Spulber (1987) and Golosov and Lucas (2007) showed that this has enormous implications for monetary policy: monetary non-neutrality (the real effect of monetary shocks) is substantially lower in menu cost models than in the Calvo model (calibrating the models to match the same frequency of price change). Since then, various different models have been developed to match different empirical facts on sticky prices, and to provide richer micro-foundations for sticky prices. This has been done, for example, by Nakamura and Steinsson (2010), Midrigan (2011), and Alvarez et al. (2011b), among others.

However, Caballero and Engel (1993) suggested an approach to sticky prices that does not require explicitly modelling the firm's price setting problem. They consider an adjustment hazard function, which is the probability of a price adjustment occurring for any given product, as a function of the deviation between the current price and optimal price absent any nominal rigidities. The concept of the adjustment hazard function can be applied to many other cases in which discrete microeconomic decisions drive fluctuations in aggregate variables (such as labor market adjustment, or lumpy investment in capital). But it is particularly well suited to analyze the mapping between sticky prices and monetary policy. Indeed, the existing "model-based" literature has shown that the selection effect, or whether more mis-aligned prices are more likely to change, is a crucial determinant of monetary non-neutrality. This is exactly what the hazard function approach adresses, as the shape of the hazard function makes explicity how much selection there is. The advantage of this approach is that it is extremely tractable, does not require solving a model of optimizing agents, and is flexible and general. Any given hazard function implies a certain amount of monetary non-neutrality (or aggregate rigidity), and price change statistics that can be compared to the values observed in price micro data.

In a more recent paper, Caballero and Engel (2007) showed how various properties of the price adjustment hazard function are related to monetary non-neutrality. And while they included an estimate of the hazard function in an earlier version of that paper (Caballero and Engel (2006)), the estimate was based on only a few price change statistics that are not able to adequately discipline the important features of the hazard function. In this paper, we re-visit the exercise of estimating the price adjustment hazard function. In particular, we propose to use the CPI micro data from high inflation periods, and especially the relationship between the higher moments of the price change distribution and inflation (the same relations we had presented in chapter 2) to estimate the hazard function much more precisely.

We find that without the conditional moments that we focus on, a wide range of hazard functions can match the moments that have been the focus of the literature until now, and that these hazard functions imply very different levels of monetary non-neutrality. However, once we include the conditional moments in the estimation, the hazard functions pertaining to the existing models can be rejected. We non-parametrically find a hazard function that can match all the moment, and that implies a much greater degree of non-neutrality than the ones that could have been previously considered. This is closely related to the insight in chapters 2 and 3 of this work, in which we showed that only a model with a very small role for the selection effect (and therefore a large degree of predicted non-neutrality) could match the facts we presented. However, our estimated hazard function provides us with additional insights. For example, the hazard function makes it clear that the probability of price adjustment stays far below 1 even for large price mis-alignments (contrary to what would be obtained from menu cost, or hybrid menu cost models, such as Midrigan (2011)). In addition, the hazard function is asymmetric: for an equivalent magintude of price mis-alignment, a price cut is less likely than a price increase. Both of these features are very important for the hazard function to be able to match all the empirical facts, and are consistent with a rational inattention model with an intermediate cost of processing information (Woodford

(2009)).

The rest of this paper is organized as follows. Section 2 formalizes the hazard function approach, and illustrates it with various examples. Section 3 describest the estimation procedure and presents the results. In section 4, we derive results on monetary non-neutrality from the estimated hazard function. Finally, we provide some concluding remarks in section 5.

4.2 The Price Adjustment Hazard Function

4.2.1 Basic Set-Up

We formally present here the price adjustment hazard function proposed by Caballero and Engel (1993). There are a large number of firms in the economy, and the hazard function gives the probability of a firm's price adjusting given the difference between the current price and the optimal price:

$$H(x) = P(\Delta p_t | x \equiv p_t^* - p_{t-1})$$

Furthermore, when a price adjusts, it is set to its current optimal level (p_t^*), such that the desired price gap (or price imbalance) is closed. This set-up is extremely general, and nests existing sticky price models, as we will show below. It is important to note, however, that we are here not modelling firms' (or households') optimization problems. Rather, we are just laying out a general way of describing the behavior that firms would choose, and that will aggregate up to determine inflation. In a model, firms would optimally choose how to determine the likelihood of adjustment for different price imbalances, given certain constraints on changing prices. Furthermore, it would be optimal for firms to set their price equal to the target at all times (and thus for the hazard function to be always equal to 1) absent any pricing constraints.

Before continuing our discussion of the price adjustment hazard function, it is important to draw a distinction with another hazard function that has received attention in the sticky price literature: the time-dependent hazard function. This function, evaluated at time t, gives the probability of a price change occurring at time t, given that t periods have passed since the last price adjustment. Different models will also imply different types of time-dependent hazard functions, and some work has been done to estimate these using price micro data (such as

citeklenowandkryvtsov2008 and Nakamura and Steinsson (2008)). However, in this paper we only investigate the state-dependent price adjustment hazard function described in this section. The fact that this latter function involves an unobserved variable (the desired price gap) makes it particularly difficult to estimate, although estimating the time-dependent hazard function involves other significant empirical challenges.

Next, we must specify a process for the desired price. We do not explicitly model the production technology or demand system that would lead to a firm's desired price. Instead, we simply assume that the log of the desired price is the sum of an idiosyncratic (z) and aggregate (m) component (each of which are expressed in logs):

$$p_t^*(i) = z_t(i) + m_t$$

An idiosyncratic component, with shocks, is important in order to match the fact that, even within a given period, there is a wide variation of price changes (as shown notably by

Bils and Klenow (2004), Nakamura and Steinsson (2008), and citeklenowandkryvtsov2008). We assume that the idiosyncratic component follows a log-AR(1) process, as the idiosyncratic productivity shocks in menu cost models like Golosov and Lucas (2007):

$$z_t(i) = \rho z_{t-1}(i) + \epsilon_t(i),\ \epsilon_t \sim N(0, \sigma_\epsilon)$$

The aggregate shock follows a random walk with drift:

$$m_t = \mu + m_{t-1} + \eta_t,\ \eta_t \sim N(0, \sigma_\eta)$$

The drift parameter (μ) accounts for the fact that inflation is positive, on average. The aggregate shocks lead to variations in inflation, and represent monetary, or aggregate demand, shocks. While we are not modelling firms' optimal response to these shocks, the hazard function (along with the distribution of idiosyncratic shocks) will determine how the aggregate price level responds to them. It is in this way that the hazard function allows us to assess the degree of aggregate price flexibility, or the inverse of monetary non-neutrality. In the remainder of this section, we will make more explicit just how the hazard function determines the degree of non-neutrality, and illustrate how different models map into this approach with examples.

4.2.2 Relation to Aggregate Flexibility

Here we present the theoretical results of Caballero and Engel (2007), who derived the relationship between various aspects of the hazard function and aggregate price flexibility. We do this to illustrate the importance of having a precise estimate of the hazard function, and of knowing what shape it takes, in particular.

In what follows, we denote the cross-sectional density of price imbalances (x) at time t as $f_t(x)$. Although we have made assumptions about the distribution of the imbalance's components, we present the theoretical results for a general distribution. The change in the aggregate price level (or the average price change, or inflation) can therefore be expressed as:

$$\Delta p_t = \int xH(x)f_t(x)dx$$

We are interested in how this change will depend on the aggregate shock, and so we consider the change in the price level as a function of the change in m_t (since m is the aggregate component of the desired price, we can assume that $f_t(x)$ is the density of price imbalances absent any change in m):

$$\Delta p_t(\Delta m_t) = \int xH(x)f_t(x + \Delta m_t)dx = \int (x - \Delta m_t)H(x - \Delta m_t)f_t(x)dx$$

Aggregate flexibility is then defined as the derivative of inflation with respect to Δm, evaluated at $(\Delta m = 0)$. Based on the previous equation, this is equal to:

$$\Delta p_t'(\Delta m = 0) = \int H(x)f_t(x)dx + \int xH'(x)f_t(x)dx$$

[1]

The first term in this expression is simply the frequency of price change in period t. Although the hazard function clearly plays a role, the frequency of price change can be estimated simply and directly from price micro data, and can therefore be assessed independetly of the hazard function. However, the second term does not have a simple relation to anything measurable in the data. In particular, since it is the derivative of the hazard function that enters the second term, its shape has an important influence on monetary non-neutrality. Caballero and Engel (2007) refer to this term as the role of the extensive margin, and show that it is typically twice as large as the frequency, in a wide variety of models associated with increasing hazard functions. Next, we illustrate the importance of these concepts with a few examples of hazard functions implied by different sticky price models.

4.2.3 Examples from Existing Models

The simplest, and perhaps most used, sticky price model is the one due to Calvo (1983), in which every period a fixed fraction of firms is randomly assigned the opportunity to change its price. Because it is randomly determined whether a price changes or not, the probability of a price change is constant, regardless of the degree of price mis-alignment. This means that the hazard function associated with this model is simply a constant equal to the frequency of price change, which we denote f:

$$H(x)^{Calvo} = f$$

[1]Note that the average sensitivity of the price level to the change in the aggregate component takes the same expression, with the ergodic density of price imbalances, $f_E(x)$ replacing $f_t(x)$.

Applying the expression for aggregate flexibility to this function makes clear why the Calvo model has a much lower degree of monetary non-neutrality than other models it is compared to. Indeed, since the hazard function is constant, the derivative is zero everywhere, meaning that the contribution of the extensive margin to aggregate flexibility is zero. Almost any micro-founded price setting model will imply that the probability of price adjustment is at least weakly increasing in the size of the price imbalance, meaning that the hazard function will also be increasing in the absolute value of x. For any such model, the extensive margin term will be positive.

At the other extreme in terms of flexibility is the class of menu cost, or Ss, models (such as Golosov and Lucas (2007)). When firms have to pay a fixed cost to change their nominal price, they optimally choose a threshold rule, under which they re-set their optimal price if and only if their price mis-alignment is outside of some interval (S,s):

$$H(x)^{MC} = \begin{cases} 0, & x \in (S, s) \\ 1, & x \notin (S, s) \end{cases}$$

While the derivative of this hazard function is not well defined, Caballero and Engel (2007) show that in this case the extensive margin term is clearly positive, and equal to:

$$|S| f_t(S) + s f_t(s)$$

A menu cost has an extreme extensive margin, because only the most mis-aligned prices adjust. Since they adjust by a large amount, the average price response to monetary shocks is very large, making the aggregate price level very flexible.

An example of a hybrid between the two extremes is Midrigan (2011) 's menu cost model. In this model, each firms produces multiple products, but once it pays the menu cost, it can change the price of all its products. Once a particular product's price is mis-aligned beyond a certain amount, the firm pays the menu cost to re-set the price (so beyond a threshold, the hazard function is 1). When this occurs, however, the firm's other prices will also be re-set, essentially for free, and these will not necessarily be very mis-aligned. This means that all prices will have a positive probabily of adjusting, no matter how small their imbalance, because of the possibility that their firm's other prices will lead it to pay the menu cost and adjust all its prices. This leads to a hazard function of the following form:

$$H(x)^{Mid} = \begin{cases} p_z, & x \in (L, U) \\ 1, & x \notin (L, U) \end{cases}$$

In this case, the extensive margin term is equal to:

$$(1 - p_z)[|L| f_t(L) + U f_t(U)]$$

While it is not a priori clear how this compares to the standard Ss extensive margin effect, we will show that once the parameters are calibrated to match some key price change statistics, this term is much smaller for the Midrigan model (due to the fact that the thresholds will be bigger, and coincide with points of lower density). This is consistent with the fact that Midrigan's model exhibits a much higher degree of monetary non-neutrality than the standard menu cost model.

The models that we have presented so far yield hazard functions that are extremely simple. However, richer and more interesting hazard functions can be obtained from models of imperfect information. We also present the example of the rational inattention model due to Woodford (2009). In this model, firms face cost to process relevant information (measured in terms of entropy reduction) to make their pricing decisions. A key insight provided by this framework is that different values for the tightness of the information constraint make the model change drastically. Indeed, the model can nest the Calvo model (in the no information case) and the menu cost model (free information, or no constraint), as well as intermediate cases. This can be seen in the different hazard functions implied by different values for the cost of information (denoted by θ), shown in Figure 4.1

An infinite cost of information corresponds to a Calvo-like hazard function, while free information leads to a trough-shaped one, as in the standard menu cost model. Intermediate values for the cost of information yield smooth funtions increasing in the absolute value of the price imbalance (denoted q in the figure). It is also noteworthy that these functions are asymmetric around zero, so that for a given size of the price imbalance, a price increase is more likely than a decrease. In this model, this is due to the asymmetry of the profit function, which makes it more costly to the firm to have its price be too low. This is consistent with the hazard function estimated by Caballero and Engel (2006), and we will also show that an asymmetric hazard function fits the data best.

As these examples show, various sticky price models can be analyzed under the hazard function approach, and this allows us to clearly see and understand the differences in the degree of aggregate flexibility that they imply. In the following section, we will present the estimation of the hazard function, which we will then use to revisit the aggregate flexibility

Figure 4.1: Rational Inattention Hazard Functions

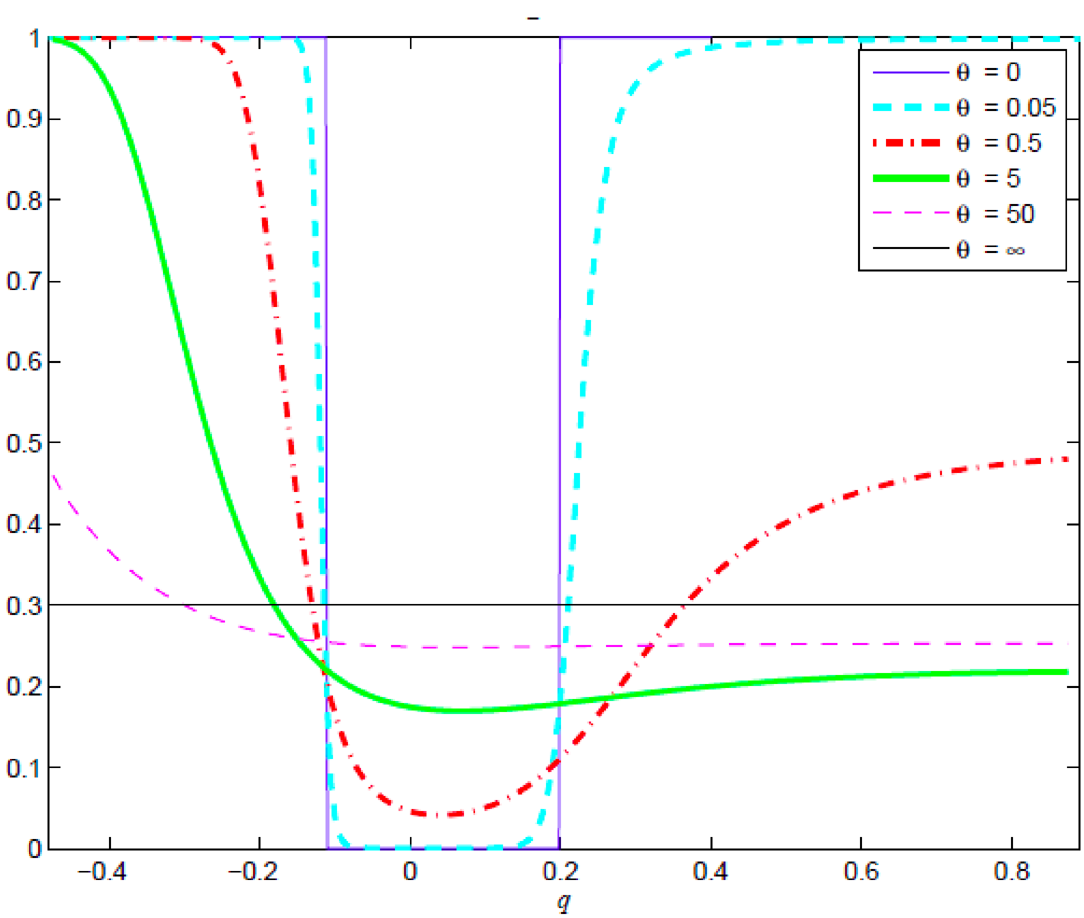

results of the different models.

4.3 Hazard Function Estimation

In order to estimate the hazard funciton, we will simply compare the price change moments implied by different guesses of the hazard function (obtained from simulations) with the moments from the data, and choose the hazard function that matches these moments best. For this to be feasible, we need to choose a finite number of parameters to characterize the hazard function. One way to do this is to impose a functional form on the hazard function, and to estimate the parameters that characterize it. This can be applied to the models in section 2, for example, as most of them imply hazard functions that are summarized by a small number of parameters. In this way, we can actually approximately estimate the underlying models without having to solve them. The resulting estimates do not recover the structural parameters of the models, but they do allow us approximately asess the model's implications for various empirical patterns, and the degree of aggregate flexibilty. We will implement this method for the standard menu cost model and for the Midrigan model (we can also do this trivially for the Calvo model).

A second approach to estimating the hazard function is non-parametric, and this is the one that will constitute our main results. The approach consists of discretizing the space for the price imbalance (x), and estimating the value of the hazard function (the price adjustment probability), essentially making each point on the hazard function a parameter. We prefer this approach, because it has the obvious and important advantage of not imposing parametric restrictions on the hazard function, so that we are not a priori choosing between a menu cost or Calvo model, for example. The disadvantage is that it of course requires estimating considerably more parameters, but we find that we can implement this for a relatively fine grid in a reasonable amount of time. We now describe the data and the moments that we use to carry out the estimation.

4.3.1 Data and Moments

As in chapters 1 and 2, we use the micro data underlying the U.S. Consumer Price Index for the period 1977-2014. Since it contains a very large number of individual prices tracked over time, it enables us to construct statistics related to individual price changes. We can then compare these statistics with the results of simulating firms adjusting prices according to various hazard functions. To estimate the statistics, we use the same method as in chapter 2, and the same restrictions on which observations to include in the estimates.

The first key moment that we use in the estimation is the average frequency of price change. As shown above, monetary non-neutrality depends on the frequency of price change, and on the extensive margin effect. The latter has to do with the extent to which prices that are more mis-aligned are more likely to change (and is closely related to the concept of the selection effect that has received much discussion in the literature). However, the size of this effect cannot be observed. That is why a model, or a hazard function, is needed, and this is the focus of our exercise. However, the frequency of price change can be directly observed and estimated with the data that we have, and we ensure that our estimated hazard function matches its correct value. The second moment that we use is the average absolute value of price changes (measured in percentage changes, conditional on a non-zero change occurring).

While this statistic does not enter directly into the expressions for aggregate flexbility, it is key to discipline the variance of idiosyncratic component of desired price changes. It has also been consistently found that price changes are large on average (Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008)). We construct these estimates as described in chapter 1, and we estimate the frequency and average size of price increases separately, and use those as target moments.

These can be considered the basic moments that any sticky price model or hazard function should match. And while they are necessary to place restrictions on the key parameters of these models (such as the price adjustment probability, or the width of the inaction bands), it is also important that almost any model can match these. Another moment that can provide additional information is the fraction of price changes that are smaller than a certain threshold. Midrigan (2011) had shown that the standard menu cost model cannot match this moment when the threshold is small, as under a fixed menu cost firms will never be willing to pay the cost to carry out small changes. Midrigan's model provides additional flexibility that allows it to match this, and this leads it to predict a much higher degree of monetary non-neutrality than the standard menu cost model. We will also require our hazard function to match this fraction, with the threshold for a price change being small is 1% (in absolute value).

Caballero and Engel (2006) only used the frequency and size moments to estimate the hazard function. They were able to do this by restricting the function to a quadratic function that is potentially asymmetric around 0. While they are able to match these moments, we will show that other hazard functions can match those moments, as well as others that would not be matched by their estimated hazard function. This is even the case when the distribution of idiosyncratic shocks is restricted to be normal, as we do. We will estimate hazard functions pertaining to some of the existing models using these moments to illustrate this point.

As we argued in chapter 2, these unconditional moments do not provide enough information to conclusively discriminate between models that predict very different degrees of monetary non-neutrality. However, conditional moments that describe the relationship between inflation and the shape of the price change distribution can be much more informative. Indeed, we have found that menu cost models predict that the distribution of price changes should become less dispersed and more negatively skewed as inflation increases, while other models (such as the Calvo model, or the rational inattention model) do not make these predictions. We therefore include these correlations among the moments to be matched, and show that it is precisely these moments that allow us to reject hazard functions that would not have been rejected based on the unconditional moments. We will also include the correlation between inflation and the frequency of price change, as it creates a very clear and simple distinction between state-dependent models (in which more firms choose to change their prices when inflation is high) and the Calvo model (in which the same fraction of firms change their prices in every period, by assumption).

We will use all of the moments mentioned to estimate the non-parametric hazard function, and this will constitute our central result. In addition, we will also include the average over time of the dispersion and skewness of price changes, to add further discipline to the shape of the hazard function. We will show that while this function matches the unconditional moments that have been considered by the literature before us, it is also able to match the correlations that we have emphasized here and in chapter 2, in a way that no existing model

Table 4.1: Target Moments

Moment	Value
Avg. Frequency of Increases	0.0805
Avg. Frequency of Decreases	0.031
Avg. Size of Increases	0.072
Avg. Size of Decreases	0.079
Fraction of Small Changes	13.2%
Avg. Dispersion (IQR)	0.099
Dispersion of Price Increases	0.077
Dispersion of Price Decreases	0.087
Avg. Skewness	-0.14
Corr(Frequency,π)	0.523
Corr(IQR,π)	-0.419
Corr(Skewness,π)	0.201

has. In Table 4.1, we list the moments that we will target, as well as their values in the data:

Using these moments, we will estimate various estimates of the price adjustment hazard function. We do this by running simulations for 60,000 firms and 1,000 periods. We simulate the shock processes, and use a candidate hazard function to simulate price changes. We then calculate moments based on the price change distribution, and compare them to their empirical coutnerparts. We then settle on the hazard function that matches the moments best. First, we will use a subset of the moments to estimate parametric functions related to some of the existing models, and then present the non-parametric estimate using all the moments.

4.3.2 Estimates from Existing Models

The simplest sticky price model is the Calvo model, which yields a simple constant hazard function. The only parameter of the hazard function that needs to be set is the adjustment

Table 4.2: Calvo Hazard Function

Parameter	Value
p	0.114
σ_ϵ	0.055
ρ	0.6
Moments	
Avg. frequency	0.114
Avg. absolute price change	0.074

probability. This can simply be set equal to the overall price adjustment frequency, which in our data is 0.114. While this fully describes the hazard function, we also need to set values for the shock parameters. We set the aggregate shock process to have a drift parameter (μ) of 0.002 and a standard deviation of 0.0037 (to match the time series properties of U.S. nominal GDP). There is no clear reference to calibrate the parameters of the idiosyncratic shock process (the persistence ρ, and the standard deviation σ_ϵ), so we set them to match the average size of price changes and the ratio between price increases and decreases. The size of price changes, in particular, is largely determined by σ_ϵ. In Table 4.2, we show the parameters that we select for the Calvo hazard function, and the moments that they imply.

This hazard function can easily match the overall frequency of price change, and the average absolute value of all price changes (inreases and decreases). However, when the frequency and average size are decomposed into increases and decreases, the match is no longer as good.

We next consider the hazard function corresponding to the Golosov and Lucas menu cost model, featuring an inaction region. The parameters to estimate here are the bounds of the

Table 4.3: Ss Hazard Function

Parameter	Value
S (left bound)	-0.074
s (right bound)	0.054
σ_ϵ	0.028
ρ	0.8
Moments	
Avg. frequency of increases	0.079
Avg. frequency of decreases	0.038
Avg. size of increases	0.07
Avg. size of decreases	0.09

inaction region (S and s), and again the idiosyncratic shock parameters. Table 4.3 shows the parameters we set, and the moments obtained

This hazard function matches the overall frequency and average size of all increases quite closely. The frequency and size of price decreases are slightly too high, but the fact that price increases are considerably more frequent and slightly smaller on average is captured. Note that in order to achieve this, the inaction region is very asymmetric around zero, which is also a common feature in the solution to menu cost models.

We also present the hazard function corresponding to the Midrigan model. This function has one additional parameter relative to the previous one, as the function features a positive probability of price adjustment for small price imbalances, which will lead the model to predict the occurrence of small price changes. Another extension introduced in Midrigan's model is that the idiosyncratic variable (firm productivity, in the model) follows a Poisson process. We introduce this modification by using the following process for the idiosyncratic

Table 4.4: Midrigan Hazard Function

Parameter	Value
L (left bound)	-0.095
U (right bound)	0.065
p_z	0.05
p_ϵ	0.25
σ_ϵ	0.058
ρ	0.7

Moments	
Avg. frequency of increases	0.077
Avg. frequency of decreases	0.040
Avg. size of increases	0.071
Avg. size of decreases	0.083
Fraction of small changes	0.133

component of the desired price:

$$z_t(i) = \begin{cases} \rho z_{t-1}(i) + \epsilon_t(i), & Prob = p_\epsilon \\ \\ z_{t-1}(i), & Prob = 1 - p_\epsilon \end{cases}$$

This extension allows the model to match the fact that the distribution of price changes in the data has fat tails, and extending the hazard function set-up in this way has the same effect. This also adds one parameter to set: the probability of a shock occurring (p_ϵ). Table 4.4 shows the parameter values resulting from our calibration, and the implied moments.

This hazard function matches the frequency of price increases and decreases quite closely, and the fraction of price changes that are small (less than 1% in absolute value). In order to achieve this, the bounds beyond which price changes are certain to occur are asymmetric around zero. The difference between the average size of price increases and decreaes is again larger than in the data, but the average size of all price changes is as in the data. In addition, the kurtosis of price changes (around 4 on average) is closer to the empirical value than the Golosov and Lucas simulations.

Finally, we consider the one existing direct estimate of a price adjustment hazard function that we are aware of: the one in Caballero and Engel (2006). Their approach was to use the frequency and average size of price increases and decreases to estimate a simple asymmetric, quadratic hazard function:

$$H(x) = \begin{cases} \lambda_n x^2 & , x \leq 0 \\ \\ \lambda_p x^2 & , x > 0 \end{cases}$$

We estimate this form of the hazard function using the relevant moments from Table reftable: target moments (the values are slightly different than in their paper because their values were based on a shorter time period), and present the results in the table below. As the previous hazard functions, this one is reasonably successful in matching the frequency and size statistics. It is clear that the function has to be very strongly asymmetric around zero (with price increases being more likely) in order to match the considerably higher fraction of price increases. Caballero and Engel had obtained the same result with their original estimate.

In Figure 4.2 below, we plot the hazard functions estimated so far, to illustrate the differences between them.

The hazard functions that we have presented so far are quite successful at matching the targeted unconditional moments, which have been the focus of most of the literature on sticky prices until now. Indeed, they can all exactly match the overall frequency and

Table 4.5: Quadratic Hazard Function

Parameter	Value
λ_n	15
λ_p	100
σ_ϵ	0.035
ρ	0.7

Moments	
Avg. frequency of increases	0.074
Avg. frequency of decreases	0.034
Avg. size of increases	0.07
Avg. size of decreases	0.095

Figure 4.2: Existing Model Hazard Functions

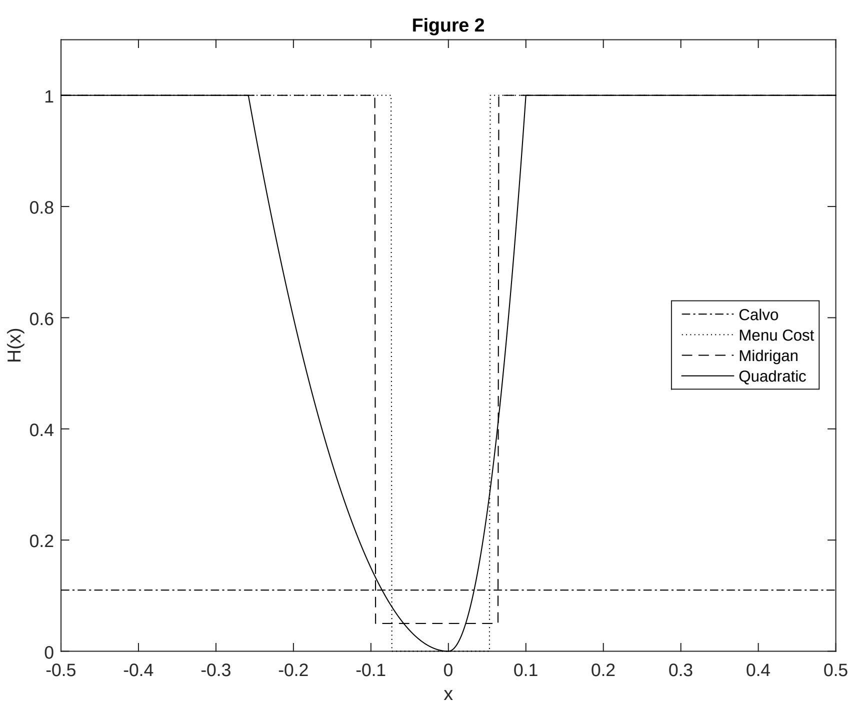

absolute value of price changes, and it is mostly in matching the averages decomposed into increases and decreases that there are marginal differences between the hazard functions (as well as the fraction of small price changes). However, as we will show now, following our findings in chapter 2, the correlations between inflation and various price change statistics show striking differences between some of the hazard functions.

In a sense, this has already been known before our results. Indeed, a very common criticism of the Calvo model is the assumption that firms are randomly assigned the opportunity to change their price, with a constant probability of price adjustment. While the Calvo model can easily match the frequency of price change, and even the average size of price changes (once the model is augmented with idiosyncratic shocks, as we have done), it was always understood that a simple way to reject this assumption would be to show that the frequency of price change rises with inflation. This has indeed been shown by Gagnon (2009) and Alvarez et al. (2011a), among others. In this way, the performance of models at different rates of inflation can provide important information about how they work and how plausible the assumptions underlying them are. The result on the frequency-inflation correlation has, on its own, vindicated the general class of menu cost models by highlighting the necessity for state-dependence. However, we show that by looking at the higher moments of the price change distribution, in particular the skewness, we can find important information on what kind of (or how much) state-dependence is realistic.

In Table 4.6, we report the values of the correlations from the simulations under the different hazard functions.

None of the hazard functions are able to match all three correlations. In particular, while the menu cost hazard functions match the positive frequency correlation, and the negative

Table 4.6: Moment Correlations for Existing Model Hazard Functions

	Table 5		
	Corr(Frequency,π)	Corr(IQR,π)	Corr(Skewness,π)
Calvo	-0.05	0.42	0.56
Midrigan	0.86	-0.97	-0.99
Golosov & Lucas	0.85	-0.68	-0.99
Caballero & Engel	0.91	-0.89	-0.99
Data	0.52	-0.42	0.20

dispersion correlation, they imply a counter-factual skewness correlation. As we explain in chapter 2, this has to do with the fact that, in these models, prices adjust with certainty once they reach a threshold for the mis-alignment. This is a natural consequence of a menu cost, as there will always be a point beyond which it is worth paying the fixed menu cost to adjust a price. This means that as the average of the desired price change distribution rises, a big share of the mass of realized price changes concentrate right beyond the edges of the positive adjustment threshold, inducing more negative skewness. The negative inflation-skewness correlation that this induces is not supported by the data. The Calvo model does not feature this kind of effect, as the hazard function implied by it is flat. While this allows it to match the right sign of the skewness correlation, it fails (by assumption) to match the fact that the frequency rises with inflation. We will therefore look for a hazard function that will match each of these correlations.

4.3.3 Non-Paramteric Hazard Function

Here we present the main result of our paper: the non-parametric price adjustment hazard function estimated using both unconditional and conditional price change moments. The way we implement this is by selecting 51 "points" on the grid of price imbalances, equally spaced between -0.5 and 0.5. We then search for the value of the the candidate function (which is the probability of price adjustment) at the 51 grid points, and assign the values between the grid points by linearly interpolating between them. In other words, the hazard functions we are considering are a subset of the space of real-valued functions spanned by a basis consisting of the 51 points on the space of price imbalances. The hazard function is then based on the candidate function, h(x) in the following way:

$$H(x) = \begin{cases} h(-0.5), & if\ x < -0.5 \\ h(x), & if\ x \in [-0.5, 0.5] \\ h(x), & if\ x > 0.5 \end{cases}$$

We further assume that $h(x) \in [0,1],\ h'(x) \leq 0\ if\ x < 0\ and\ h'(x) \geq 0\ if\ x \geq 0$. However, we notably do not impose that the hazard function be symmetric around zero. In addition to the hazard function values, we also have to set a value for the variance and persistence of the idiosyncratic shocks, as well as the arrival probability of idiosyncratic shocks (p_ϵ), since we use the more flexible Poisson process (this particularly helps to match the frequency of small price changes).

The process we use to find the hazard function values that best fit the data is as follows. First, we fix the values for the parameters p_ϵ and σ_ϵ(as an initial guess), and then aim to solve the following optimization problem:

$$\underset{\{H(x_i)\}}{min} \sum_{t \in \{Moments\}} (1 - \frac{t}{t^*})^2$$

, where t is the value of a given moment by a particular combination of the 51 hazard function values ($\{H(x_i)\}$), and t^* is the empirical value. We are therefore simply looking for the hazard function values that minimize the sum of squared deviations across the moments that we traget, with equal weighting for all the moments. In order to do this, we use the pattern search optimization procedure, which is intended for optimization problems in which the gradient of the objective is not defined. This is exactly the case for our problem, as the values of the model-implied moments can only be computed by simulation. Davidon (1991) describes this specific procedure in more detail, and we stop the process once the value of the objective function changes by less than 10^{-6} across successive iterations. With the values of the hazard function fixed, we then adjust the values of p_ϵ and σ_ϵ manually to match the average size and fraction of small price changes. Since there is a very close relationship between those two parameters and moments, this step is relatively simple to implement manually.

In implementing this non-parametric estimation, it is true that that number of probability points ("parameters") that we are choosing is greater than the number of moments being targeted (14), so in this sense the estimation cannot be considered to be identified, even though the assumptions we make about the hazard function also provide some additional restrictions. Our procedure is in practice a calibration, as we cannot be certain that the paramater values that we settle on are the only ones that match the data. Nevertheless, in searching for the values that best match the moments that we target, we can rule out several types of hazard functions. We find that hazard functions with the "Ss" feature, or with a rapidly rising hazard near x = 0, consistently predict a negative inflation-skewness correlation, and can thus be rejected. Similarly, a flat hazard can be rejected based on the restriction that the frequency of price change must rise with inflation. It is in this sense that our approach, by including sign restrictions on the moment correlations, allows us to recover important features of hazard function.

In light of the concern with the fact that our approach is under-identified, as an estimation procedure, we also implement a calibration of the hazard function searching for only 9 values on the grid of price imbalances. In this way, the number of parameters being chosen is smaller than the number of moments targetted. As we show below, the results for this hazard function are very similar, and the fit with the data is almost as good. Alternatively, we could choose a flexible parametric form for the hazard function, with a small number of parameters, and estimate those parameters using these (or additional moments).

In Figure 4.3, we present the hazard function that best matches the moments we had set as targets. This goes along with the following parameter values for the idiosyncratic shock process: $\rho = 0.7,\ \sigma_\epsilon = 0.058,\ p_\epsilon = 0.4$.

Figure 4.4 the estimated hazard function with only 9 grid points (the accompanying parameter values are: $\rho = 0.7,\ \sigma_\epsilon = 0.056,\ p_\epsilon = 0.4$). We only estimated points on the price imbalance space ranging from -0.3 to 0.3, as imbalances with larger values occur very rarely). The values of the hazard function are overall quite similar to those estimated on a finer grid.

Our estimated hazard function (under both specifications) is clearly increasing in the

Figure 4.3: Non-Parametric Price Adjustment Hazard Function

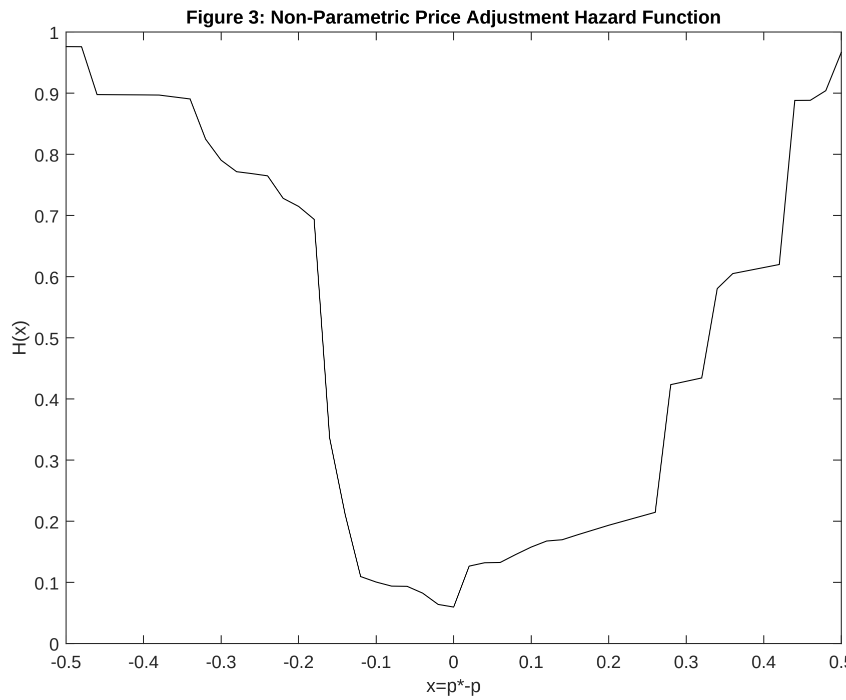

Figure 4.4: Non-Parametric Hazard Function (9 points)

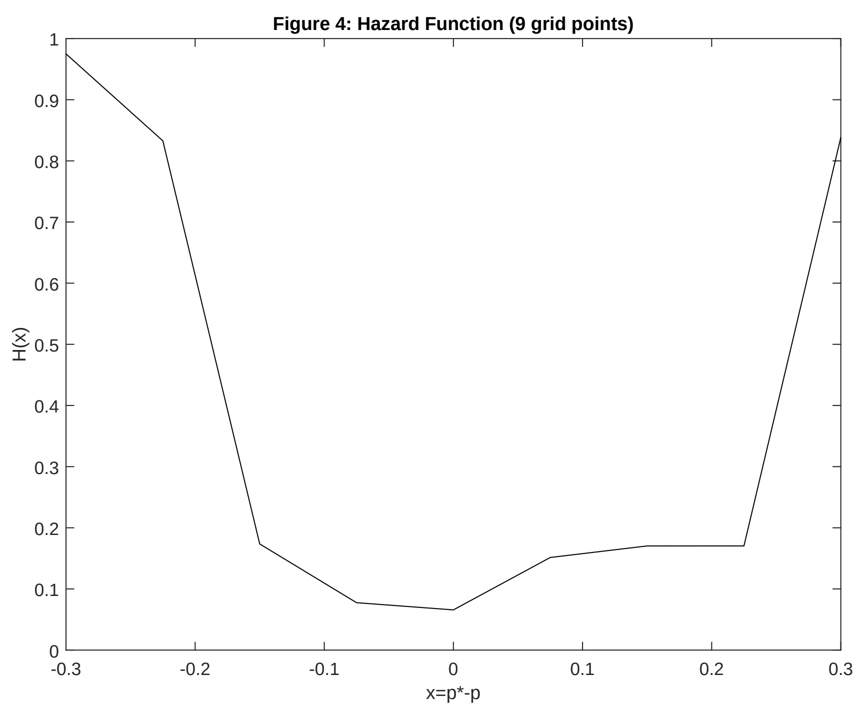

absolute value of the price imbalance. While it does not feature any inaction region, like the the menu cost models do, there is clearly a sharp, large rise in the probability of price adjustment after a certain threshold. However, there is a significant asymmetry in this rise between negative and positive imbalances: the point beyond which the probability rises in the negative range is much smaller than in the positive range. However, there is another type of asymmetry: for imbalances smaller than this "threshold", the probability of adjustment is smaller for desired price decreases than for increases (keeping the size of the desired change constant).

Clearly, the shape of this hazard function is quite different from any of the model-based functions we have presented so far (with the possible exception of some cases of the rational inatention model). Here, we highlight some of the key diffences. First, as mentioned above, the price adjustment probability is asymmetric around zero. This asymmetry can be summarized in the following way: small price and large price increases are more likely than price decreases of equivalent size, but intermediate sized price decreases are more likely (although this is a narrow range).

Second, even for price imbalances of zero, there is a significant probability of price adjustment (around 10%). It turns out that this feature (which we refer to as a Calvo feature, because the Calvo model gives the same probability of adjustment to imbalances of zero) is crucial to achieving a non-negative correlation between price change skewness and inflation. The Midrigan hazard function has a similar feature, but we find that the probability of adjustment at an imbalance of zero has to be higher. While this feature might appear at first sight to make the price adjustment process more flexible, it is actually more likely to lead to less aggregate flexibility. That is because, if the probability of adjustment has to be high for small imbalances, then it must be lower (than it otherwise would be) for larger imbalances, because the overall frequency of price change is, so to speak, fixed by the value that we estimate in the data. This means that a larger share of price changes are smaller price changes that will be less responsive to aggregate shocks.

Finally, we find that the probability of price adjustment does not go all the way to one unless the price imbalances are very large. In other words, even relatively large price imbalances (of about 20%) can still have a high probability of not inducing an adjustment (and for desired price increases, this probability is significantly greater than 50%). This also comes out of matching the non-negative inflation-skewness correlation. Indeed, if large or intermediate desired price changes (that compose the tails of the price change distribution) occur with certainty, the tails of the distribution will be very sensitive to the aggregate shock, which drives inflation. This is what makes the skewness of price changes vary so sharply with inflation in the menu cost models. But this effect is muted if the probability of adjustment is significantly below one for these price changes, which is the case in our estimated hazard function. From our attempts to match the correlations, we have learned that this feature, as well as the previous ones, are necessary in order for the hazard function to simulatenously match the signs of the three correlations.

In Table 4.7 below, we present the moments implied by both estimated hazard functions. They matche the unconditional price change statistics as well as the menu cost hazard functions, but succeed in no longer predicting a negative inflation-skewness correlation. Naturally, the hazard function with 51 estimated points matches the moments slightly better.

By using a combination of unconditional and conditional moments in the estimation, we have been able to place significant restrictions on important features of the price adjustment

Table 4.7: Hazard Function Calibration

Moment	Simulations (51 pts)	Simulations (9 pts)	Data
Avg. Frequency of Increases	0.072	0.069	0.0805
Avg. Frequency of Decreases	0.042	0.039	0.031
Avg. Size of Increases	0.073	0.075	0.072
Avg. Size of Decreases	0.082	0.077	0.079
Fraction of Small Changes	13.3%	13.6%	13.2%
Avg. Dispersion (IQR)	0.12	0.114	0.099
Avg. Skewness	-0.12	-0.13	-0.13
Corr(Frequency,π)	0.76	0.76	0.523
Corr(IQR,π)	-0.90	-0.94	-0.419
Corr(Skewness,π)	0.14	0.38	0.201

hazard function. In the following section, we will show what this means for the degree of aggregate flexibility implied by the hazard function.

4.4 Monetary Non-Neutrality

The degree of monetary non-neutrality, or aggregate flexibility, can be computed given a hazard function and parameters for the shock processes. In Section 2, we showed the analytical expressions for aggregate flexibility based on the hazard function and the distribution of price imbalances. However, since our estimated hazard function only takes values at discrete points, and is thus non-differentiable, we compute aggregate flexibility by simulation. We do this by using the original simulations that yielded the implied moments, and compute the variance (across time) of log real consumption, where log real consumption is defined as $c_t = m_t - p_t$. The aggregate component of the desired price, m_t, simply follows the random walk process described above, and the aggregate price level is solved for using the hazard function. This is the measure for monetary non-neutrality (the inverse of aggregate flexi-

Table 4.8: Monetary Non-Neutrality

Hazard Function	$Var(c_t) \times 10^4$
Calvo	0.484
Non-Parametric	0.291
Midrigan	0.167
Caballero & Engel	0.137
Golosov & Lucas	0.055

bility) most commonly used in the sticky price literature (e.g. Golosov and Lucas (2007), Nakamura and Steinsson (2010), and Midrigan (2011)), as it measures the variation in real activity induced aggregate nominal shocks. With full price flexibility, real activity should not vary as prices would respond one-for-one to aggregate shocks. In the hazard function framework, that would be the case if the probability of price adjustment was always 1. In Table reftable: monetary non neutral, we present the results for the Calvo and menu cost hazard functions, the asymmetric quadratic hazard function based on Caballero and Engel (2006) as well as our non-parametric estimate.

The degree of monetary non-neutrality implied by our estimated hazard function is relatively high: it is considerably higher than those based on menu cost models, and it is about half as high as the Calvo hazard function. These results are generally in line with our findings in chapter 3, which showed that the non-neutrality predicted by the random menu cost model was also between that in the Calvo and Midrigan models. Our results here reiterate the fact that taking into account how the shape of the price change distribution varies with inflation provides evidence in favor of greater non-neutrality than would be expected by simply looking at unconditional moments.

While we do not analytically evaluate the extensive margin component of aggregate flexibility from our estimated hazard function (and the distribution of underlying shocks), the expression for this term helps make sense of the results in Table 7:

$$Extensive\ Margin: \int xH'(x)f_t(x)dx$$

This term is relatively small under our estimated function, because the hazard function is relatively flat (with a small $H'(x)$) at the imbalances that have the most density (which are mostly those such that $|x| \leq 0.1$). In contrast, the "Ss" type hazard functions feature a very large increase in the price adjustment probability at smaller values of x, for which there is a high density, giving them a very strong extensive margin effect. This is why monetary non-neutrality is high under the estimated hazard function, and this result comes from the features of the hazard function that are captured by our estimation.

4.5 Conclusion

As has been shown by Caballero and Engel (2007), the shape of the price adjustment hazard function is closely related to, and provides important information on, the degree of aggregate flexibility (or monetary non-neutrality) implied by micro-level price stickiness. While the question of the significance of monetary non-neutrality has been extensively studied using sticky price models, less attention has been paid to the hazard function approach to this question. This may be in part due to the face that, since it is not grounded in optimizing firm-level behavior, there are very few restrictions that can be placed on the shape of the hazard function a priori. Furthermore, while Caballero and Engel (2007) have derived the exact relationship between the hazard function and aggregate flexibility, they did not consider what empirical patterns could be used to discipline the key features of the hazard function. In this paper, we have attempted to fill this gap, by showing which moments can be used to estimate this function. In particular, we have emphasized that the relationship between inflation and the shape of the price change distribution provides a great amount of information on what shape the hazard function can take, and how much aggregate flexiblity it can realistically imply.

We have found that while "Ss" type hazard functions (featuring an inaction region, and a threshold beyond which desired price changes occur with certainty) can successfully match statistics related to the average frequency and size of price changes, they imply a very strong, and counter-factual, negative relationship between inflation and price change skewness. Starting from a very general form for the hazard function, we find one that is able to match both the average size and frequency moments, and the correlations with inflation. In order to match the correlations, and the non-negative inflation-skewness correlation in particular, the hazard function has to include three important properties. First, the probability of a price adjustment at a price imbalance of zero must be positive. Second, even for relatively large price imbalances (of up to 20%), the probability of price adjustment must be considerably lower than 50%. Put differently, this means that the threshold beyond which price changes are very likely is high. Finally, price increases are overall somewhat more likely than price decreases, for an equal size of the price imbalance. The first two properties, in particular, imply that aggregate flexibility is relatively low, and much lower than what would be predicted by "Ss" type hazard functions. The degree of aggregate flexibility is, moreover, broadly consistent with what we find in a random menu cost model in chapter 3.

The main contribution of this paper has been to provide a new estimate for the price adjustment hazard function using a richer set of data and empirical moments than in Caballero and Engel (2006), yielding different results on aggregate flexibility. While the hazard function framework that we have been working under is very flexible, there are several variations to our estimation procedure that we could attempt. Indeed, a richer set of processes for the idiosyncratic shocks could be considered, as it would be helpful to know how sensitive the hazard function estimates are to changes in the shock process. Specifically, we have shown that the skewness of the price change distribution can provide important information on what shape the hazard function should take, so it would make sense to work with asymmetric distributions of the desired price change distribution to see what that could mean for the results. In addition, this same procedure could be used to estimate a hazard function for different sectors (such as food, services, household goods), by using empirical moments pertaining to particular sectors. Finally, we could also derive the hazard functions implied by other sticky price models that have been proposed, and use these to empirically evaluate the models.

Bibliography

Alvarez, Fernando, Herve Le Bihan, and Francesco Lippi (2014), "Small and large price changes and the propagation of monetary shocks." Working Paper.

Alvarez, Fernando, Martin Gonzalez-Rozada, Andy Neumeyer, and Martin Beraja (2011a), "From hyperinflation to stable prices: Argentina's evidence on menu cost models." Working Paper, University of Chicago.

Alvarez, Fernando, Francesco Lippi, and Luigi Paciello (2011b), "Optimal price setting with observation and menu costs." *Quarterly Journal of Economics*, 126, 1909–1960.

Anderson, Eric, Emi Nakamura, Duncan Simester, and Jon Steinsson (2015), "Information rigidities and the stickiness of temporary sales." Working Paper, Columbia University.

Anderson, Eric T. and Duncan I. Simester (2010), "Price stickiness and customer antagonism." *Quarterly Journal of Economics*, 125, 729–765.

Ball, Laurence (2014), "The case for a long-run inflation target of four percent." IMF Working Paper.

Barro, Robert J. (1972), "A theory of monopolistic price adjustment." *The Review of Economic Studies*, 39.

Barro, Robert J. (2009), "Rare disasters, asset prices, and welfare costs." *American Economic Review*, 99, 243–264.

Benabou, Roland (1992), "Inflation and efficiency in search markets." *Review of Economic Studies*, 59, 299–329.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995), "Automobile prices in market equilibrium." *Econometrica*, 63, 841–890.

Bils, Mark and Peter J. Klenow (2004), "Some evidence on the importance of sticky prices." *Journal of Political Economy*, 112, 947–985.

Blanchard, Olivier, Dell'Ariccia, and Paolo Mauro (2010), "Rethinking macroeconomic policy." *IMF Staff Position Note*, 10, 1–18.

Blanco, Julio A. (2015a), "Costs and benefits of inflation in models with price rigidities." Working Paper, New York University.

Blanco, Julio A. (2015b), "Inflation targeting in economies with menu costs." Working Paper, New York University.

Bomberger, William A. and Gail E. Makinen (1993), "Inflation and relative price variability: Parks' study reexamined." *Journal of Money, Credit and Banking*, 25, 854–861.

Broda, Christian and David E. Weinstein (2006), "Globalization and the gains from variety." *Quarterly Journal of Economics*, 121, 541–585.

Burstein, Ariel and Christian Hellwig (2008), "Welfare costs of inflation in a menu cost model." *American Economic Review*, 98, 438–443.

Burstein, Ariel and Christian Hellwig (2009), "Inflation and welfare with nominal rigidities." Slides from talk given at Yale University on March 24th 2009.

Caballero, Ricardo J. and Eduardo Engel (1993), "Microeconomic adjustment hazards and aggregate dynamics." *Quarterly Journal of Economics*, 108.

Caballero, Ricardo J. and Eduardo Engel (2006), "Price stickiness in ss models: Basic properties." Working Paper.

Caballero, Ricardo J. and Eduardo Engel (2007), "Price stickiness in ss models: New interpretations of old results." *Journal of Monetary Economics*, 54.

Calvo, Guillermo A. (1983), "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics*, 12, 383–398.

Caplin, Andrew and Daniel Spulber (1987), "Menu costs and the neutrality of money." *Quarterly Journal of Economics*, 102, 703–725.

Carlton, Dennis W. (1986), "The rigidity of prices." *American Economic Review*, 76, 637–658.

Cecchetti, Stephen G. (1986), "The frequency of price adjustment: A study of the newsstand prices of magazines." *Journal of Econometrics*, 31, 255–274.

Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland (2012), "Optimal inflation rate in new keynesian models: Should central banks raise their inflation targets in light of the zero lower bound?" *Review of Economic Studies*, 79, 1371–1406.

Davidon, William C. (1991), "Variable metric method for minimization." *SIAM Journal on Optimization*, 1, 1–17.

Debelle, Guy and Owen Lamont (1997), "Relative price variability and inflation: Evidence from u.s. cities." *Journal of Political Economy*, 105, 132–152.

Dotsey, Michael, Robert King, and Alexander Wolman (1999), "State-dependent pricing and the general equilibrium dynamics of money and output." *Quarterly Journal of Economics*, 114, 655–690.

Eichenbaum, Martin, Nir Jaimovich, Sergio Rebelo, and Josephine Smith (2013), "How frequent are small price changes." Working Paper.

Fischer, Stanley (1981), "Relative shocks, relative price variability and inflation." *Brookings papers on Economic Activity*, 2, 381–431.

Foote, Christopher L. (1998), "Trend employment growth and the bunching of job creation and destruction." *Quarterly Journal of Economics*, 113, 809–834.

Gagnon, Etienne (2009), "Price setting under low and high inflation: Evidence from mexico." *Quarterly Journal of Economics*, 124.

Glejser, H. (1965), "Inflation, productivity, and relative prices – a statistical study." *Review of Economics and Statistics*, 47, 76–80.

Golosov, Mikhail and Robert E. Lucas (2007), "Menu costs and phillips curves." *Journal of Political Economy*, 115, 171–199.

Guimaraes, Bernardo and Kevin Sheedy (2011), "Sales and monetary policy." *American Economic Review*, 101, 844–876.

Kashyap, Anil K. (1995), "Sticky prices: New evidence from retail catalogs." *Quarterly Journal of Economics*, 110, 245–274.

Kehoe, Patrick and Virgiliu Midrigan (2015), "Prices are sticky after all." *Journal of Monetary Economics*, 75, 35–53.

Klenow, Peter J. and Oleksiy Kryvtsov (2008), "State-dependent or time-dependent pricing: Does it matter for recent u.s. inflation." *Quarterly Journal of Economics*, 123, 863–904.

Krugman, Paul (2014), "Inflation targets reconsidered." Paper presented at ECB Sintra Conference in May 2014.

Krusell, Per, Toshihiko Mukoyama, Ayşegül Şahin, and Anthony A. Smith, Jr. (2009), "Revisiting the welfare effects of eliminating business cycles." *Review of Economic Dynamics*, 12, 393–404.

Krusell, Per and Anthony Smith (1998), "Income and wealth heterogeneity in the macroeconomy." *Journal of Political Economy*, 106, 867–896.

Lach, Saul and Daniel Tsiddon (1992), "The behavior of prices and inflation: An empirical analysis of disaggregated price data." *Journal of Political Economy*, 100, 349–389.

Lucas, Robert E., Jr. (2003), "Macroeconomic priorities." *American Economic Review*, 93, 1–14.

Midrigan, Virgiliu (2011), "Menu costs, multi-product firms and aggregate fluctuations." *Econometrica*, 79.

Mills, Frederick (1927), *The Behavior of Prices*. National Bureau of Economic Research, New York.

Nakamura, Emi and Jón Steinsson (2008), "Five facts about prices: A reevaluation of menu cost models." *Quarterly Journal of Economics*, 123, 1415–1464.

Nakamura, Emi and Jón Steinsson (2010), "Monetary-non-neutrality in a multi-sector menu cost model." *Quarterly Journal of Economics*, 125, 961–1013.

Nakamura, Emi and Jon Steinsson (2013), "Price rigidity: Microeconomic evidence and macroeconomic implications." *Annual Review of Economics*, 5, 33–63.

Nevo, Aviv (2001), "Measuring market power in the ready-to-eat cereal industry." *Econometrica*, 69, 307–342.

Newey, Whitney K. and Kenneth D. West (1987), "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix." *Econometrica*, 55, 703–708.

Parks, Richard W. (1978), "Inflation and relative price variability." *Journal of Political Economy*, 86, 79–95.

Pashigian, B. Peter (1988), "Demand uncertainty and sales: A study of fashion and markdown pricing." *American Economic Review*, 78, 936–953.

Reinsdorf, Marshall (1994), "New evidence on the relation between inflation and price dispersion." *Journal of Political Economy*, 84, 720–731.

Schmitt-Grohe, Stephanie and Martin Uribe (2011), "The optimal rate of inflation." In *Handbook of Monetary Economics, Volume 3B* (Benjamin M. Friedman and Michael Woodford, eds.), 653–722, Elsevier, Amsterdam, Holland.

Sheremirov, Viacheslav (2015), "Price dispersion and inflation: New facts and theoretical implications." Working Paper, Federal Reserve Bank of Boston.

Sheshinski, Eytan and Yoram Weiss (1977), "Inflation and costs of price adjustment." *Review of Economic Studies*, 44, 287–303.

Sims, Chrisotpher A. (2003), "Implications of rational inattention." *Journal of Monetary Economics*, 50.

Tauchen, George (1986), "Finite state markov-chain approximation to univariate and vector autoregressions." *Economics Letters*, 20, 177–181.

Van Hoomissen, Theresa Van (1988), "Price dispersion and inflation: Evidence from israel." *Journal of Political Economy*, 96, 1303–1314.

Vavra, Joseph (2014), "Inflation dynamics and time-varying uncertainty: New evidence and an ss interpretation." *Quarterly Journal of Economics*, 129, 215–258.

Vining, Daniel R. and Thomas C. Elwertowski (1976), "The relationship between relative prices and the general price level." *American Economic Review*, 66, 699–708.

Woodford, Michael (2009), "Information-constrained state-dependent pricing." *Journal of Monetary Economics*, 56.

Wulfsberg, Fredrik (2015), "Inflation and price adjustment, evidence from norwegian consumer price data 1975-2004." *American Economic Journal: Macroeconomics*. Forthcoming.

Zbaracki, Mark, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen (2004), "Managerial and customer costs of price adjustment: Direct evidence from industrial markets." *Review of Economics and Statistics*, 86.

Appendix A

Appendix for Chapter 1

A.1 Derivations

A.1.1 Flexible Price Case

The period t profits of intermediate firm i are given by

$$\Pi_{it} = p_{it}y_{it} - W_t L_{it}. \tag{A.1}$$

Using the demand curve the firm faces—equation (1.4)—and the firm's production function—equation (1.8)—we can rewrite the profit function as

$$\Pi_{it} = p_{it}\left(\frac{p_{it}}{P_t}\right)^{-\theta} C_t - \frac{W_t}{A_{it}}\left(\frac{p_{it}}{P_t}\right)^{-\theta} C_t. \tag{A.2}$$

Maximization of this expression as a function of p_{it} yields

$$p_{it} = \frac{\theta}{\theta - 1}\frac{W_t}{A_{it}}. \tag{A.3}$$

Raising the expressions on both sides of this equation to the power $1 - \theta$ and integrating over i yields

$$\int_0^1 p_{it}^{1-\theta} di = \left(\frac{\theta}{\theta - 1} W_t\right)^{1-\theta} \int_0^1 A_{it}^{\theta-1} di. \tag{A.4}$$

Raising both sides of this expression to the power $1/(1 - \theta)$ yields

$$P_t = \frac{\theta}{\theta - 1}\frac{W_t}{A_f}, \tag{A.5}$$

where A_f is given by

$$A_f = \left[\int_0^1 A_{it}^{\theta-1} di\right]^{\frac{1}{\theta-1}}. \tag{A.6}$$

Combining firm i's production function—equation (1.8)—and the demand curve for firm i's output—equation (1.4)—yields

$$L_{it} = \left(\frac{p_{it}}{P_t}\right)^{-\theta} A_{it}^{-1} C_t. \tag{A.7}$$

Integrating over i and using firm i's price setting equation—equation (A.3)—yields

$$\int_0^1 L_{it}\, di = \left(\frac{\theta}{\theta - 1}\frac{W_t}{P_t}\right)^{-\theta} \int_0^1 A_{it}^{\theta-1} di\, C_t. \tag{A.8}$$

Using equation (A.5), this equation can be simplified to yield

$$Y_t = A_f L_t, \tag{A.9}$$

where Y_t denotes aggregate output and we have $Y_t = C_t$.

A.1.2 Sticky Price Case

As in the flexible price case, combining firm i's production function—equation (1.8)—and the demand curve for firm i's output—equation (1.4)—yields

$$L_{it} = \left(\frac{p_{it}}{P_t}\right)^{-\theta} A_{it}^{-1} C_t. \tag{A.10}$$

In this case, however, some labor is potentially used to change prices. Let's denote this by L_t^{pc}. Taking this into account and integrating the above expression over i, we get that aggregate labor supply is

$$L_t = \int_0^1 \left(\frac{p_{it}}{P_t}\right)^{-\theta} A_{it}^{-1} di C_t + L_t^{pc} \tag{A.11}$$

Rearranging this equation and using the fact that $Y_t = C_t$, we get that

$$Y_t = A_t(\bar{\pi})(L_t - L_t^{pc}), \tag{A.12}$$

where

$$A_t(\bar{\pi}) = \left[\int_0^1 \left(\frac{p_{it}}{P_t}\right)^{-\theta} A_{it}^{-1} di\right]^{-1}. \tag{A.13}$$

A.2 Data Appendix

The process of scanning the microfilm cartridges left us with about 600 image folders—one corresponding to each cartridge. Each folder contains roughly 2000 images for a total of about 1 million images. The images are called Price Trend Listings and contain a table of data regarding several products (the rows in the table) over a 12 month period (the columns in the table). All data on each image comes from products in a particular product category (ELI). The left most column on each image contains a location (city) identifier, an outlet identifier, a product identifier, and a version identifier. The top row lists the time periods to which the data refers. The right most column contains the data for the month the images was created in. The remaining 11 columns repeat older data on the same products—i.e., show the "price trend" for each product. We refer to the month the image was created in as the "collection period" for this image. Each of the interior cells in the table is divided into three sub-cells. The top sub-cell reports the price. The middle sub-cell reports a number of flags including a sales flag and a flag related to whether the price is imputed. The bottom sub-cell contains the percentage change of the price since the last collection period for that product. Within each collection period, images are sorted by product category (ELI). Within each image, rows are sorted by the values of the identifiers in the left most column. They are first sorted by the outlet identifier, then by the product identifier, and finally by the version identifier.

As we describe in the main text, we used optical character recognition (OCR) software to convert the scanned images to machine readable form. We worked closely with a software company to create custom software that could convert the Price Trend Listings as accurately as possible. While we were able to find ways to eliminate most common systematic errors that we came across (especially in the price variable), it is inevitable given the current state of OCR technology that there are some random errors that remain in the raw data that results from the OCR process. Fortunately, the format of the Price Trend Listings described above implies that there is a large amount of redundancy in the raw dataset. This redundancy can be used to validate the output of the OCR process. Below we describe the procedure we use to validate the output from the OCR procedure.

A.2.1 Product Categories and Product Identifies

The first step is to validate and improve on the OCR output for category labels and product identifies. We first set to missing all ELI values that do not correspond to one of the values on the list of ELI values in the BLS classification. We then use the fact that the images are ordered by ELI within each collection period to fill in missing ELI information in cases where the last observed ELI value and the next observed ELI value are the same. Finally, in cases where there are large blocks of images that still have missing values for the ELI, we manually review the original scanned images to determine which image separates the different ELIs. By these steps we are able to validate the ELI for 99.7% of the images we have.

We use a similar procedure if there is a missing value of a product-identifier. We use the fact that the images are sorted by ELI and the observations within image are sorted first by location, then by outlet, then by product, and finally by version. First, we set identifiers that are out of order to missing. We then fill in identifiers in cases where the identifier before and after a block of observations with missing values for the identifier are the same.

Errors in reading product-identifies lead to spurious products in our dataset. Whenever such errors occur, an entire row in a single Price Trend Listing image will be associated with this "phantom product" (i.e., this erroneous product code). We will therefore have up to 12 months of price data for these phantom products, which will result in up to 11 months of price change statistics. If the distribution of phantom products is non-uniform across product categories, their presence could bias our results by putting more weight on product categories where there are more phantom products. However, it is unlikely that exactly the same error will occur for multiple images with the same product-month. This implies that product months of phantom products will appear only on a single image (i.e., we won't have more than one replicate for product-months of phantom products). This, in turn, implies that our first algorithm for accepting price observations in to the final dataset (described below) will not accept these observations. We have rerun our results with only data accepted by the first algorithm and they are virtually identical. This makes us confident that the phantom products are not biasing our results.

The presence of the phantom products does artificially inflate the number of observations in the dataset. Whenever the same observation appearing on more than one image is read in different ways from different images, what is suppose to be a single observation, turns into more than one observation. We see signs of this occurring in the early part of our sample and in particular in the latter half of 1979 and the first half of 1980. Over this period, the number of observations per month rises to 170,000 per month. But the number of observations that appear on only a single image also rises very substantially (to over 50,000 per month). We take this as a sign that in this period a substantial number of phantom observations are appearing in the dataset.

A.2.2 Prices

We use two main procedures to validate the OCR output for prices. First, since each Price Trend Listing contains not only the price for that period but also prices for up to 11 earlier month, each product-month observation may appear multiple times in the raw dataset. For example, suppose we consider a product that the BLS collected a price for from October 1979 to Novermber 1980. The October 1979 price will appeal on an image in October 1979 and will then be repeated on images in each of the subsequent 11 months. The October 1979 price of this product will therefore show up 12 times in the dataset. We refer to these 12 instances as 12 replicates of the same product-month observation.

Most product-month observations will have fewer than 12 replicates. Our data is based on monthly micro-film cartridges from May 1977 to December 1980 and bimonthly micro-film cartridges from January 1981 onward. Product-month observations in the later part of our dataset will therefore appear at most 6 times. Also, certain products in certain cities are sampled only bimonthly. These will also only appear at most 6 times. Finally, product-month observations towards the end of a product's life-time in the BLS sample may appear less often. But even the last observation for a product in many cases appears more than once. Figure A.1 plots the distribution of number of replicates in our sample.

We use this redundancy to verify the accuracy of the OCR output for prices. Our rule is to accept the price that we observe most often among the replicates for each product-month as long as we observe this price at least twice. If no price is observed more than once, we don't accept any price using this algorithm (and instead rely on the 2nd algorithm described below). It is very rare that more than one price is observed at least twice. This occurs for

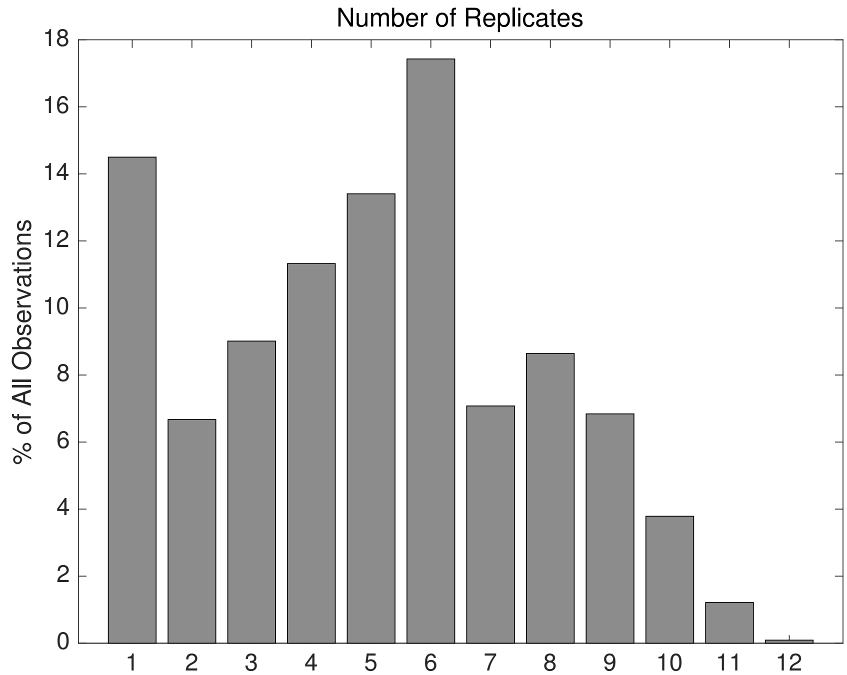

Figure A.1: Distribution of Number of Replicates for Each Product-Month Observation

only 0.04% of product-months. It is even rarer, that there are more than one prices that are observed at least twice and an equal number of time, i.e., that there is a tie as to which price is observed most often. This occurs for only 0.004% of product-months. In these cases, we accept neither price (and again rely on the 2nd algorithm). Overall, prices are accepted using this algorithm for 86.5% of product-months.

Once we have finish running this first validation algorithm, we take all replicates for product-months for which prices have been accepted and we fill in the accepted price. In other words, we "correct any errors" in all the replicates for the product-months for which the first algorithm accepts observations. We do this in order to improve the chances that the second algorithm is able to validate prices for additional product-months.

The second redundancy we make use of is the percentage change variable. We calculate the percentage change in the price for a particular product-month from the price we observe on the image in that month and the price we observe on the image in the previous month, whenever both are present on the same image. We round this calculated percentage change to the next whole percent. If the value we calculate for the percentage change in this way matches the value reported in the percentage change variable on the image, we accept the price for both the product-month in question and the previous product-month (i.e., both price values used to calculate the percentage change).

The second way in which we use the percentage change variable is that when the following conditions are met:

- The percentage change in the price that we calculate from the price observations does not line up with the percentage change variable read directly from the image in a particular product-month and also does not line up for the same product in the next month.

- The percentage change read from the image for the product-month in question and both the next and last month for that product are all equal to zero

we set the price in the product-month to the value read in the previous month and accept this value into our main dataset. This is meant to catch cases where the OCR software made an error in the price variable, but the percentage change variable gives us a strong indication that the price actually stayed constant.

The procedure we describe above that uses the price change variable on the Price Trend Listing images is repeated for each image that a particular product-month observation occurs on. A price for the same product-month can therefore be accepted more than once (in principle as often as a particular product-month occurs on different images). It is even possible that two or more different prices are accepted for the same product month using this procedure. This only occurs for 0.14% of observations. Whenever this occurs, we drop all accepted observations based on the percentage change procedure. Overall, these procedures for using the percentage change variable raises the overall acceptance rate to 98.4% of the observations in our raw dataset. We drop the remaining observations.

A.2.3 Price Flags

As we discuss above, the second sub-cell of each cell in the Price Trend Listings images contains a string of characters which code various information regarding the price in question. This string contains at most seven characters. In some cases, some of these characters are left blank. The first three spots are reserved for characters indicating, among other things whether a product is on sale, whether it was unavailable, whether it is a seasonal item, etc. Typically, at most one of these spots will contain a letter and the others will be left blank. The next three spots give information about which pricing cycle the product belongs to (some products are priced monthly and others bimonthly). The last spot may contain the letter "I" indicating that the price was imputed. If the price was not imputed, this spot is left blank. We had the OCR software convert blank spots to # signs to help us tell which spot of the string each character occurred in. However, this conversion was somewhat imperfect.

Our OCR procedure turned out to be less accurate in converting these price flags, but fortunately the price flags are chosen from a restricted set of characters and the errors follow well-defined patterns. Our interest centers on identifying two pieces of information from this string of characters: 1) whether the product was on sale, 2) whether the price was imputed.

The letter “B” is the character that indicates that the product was on sale. We therefore create a sales flag variable and set it to 1 if the letter “B” occurs in the string of characters. Due to worries about accuracy or the OCR procedure for these flags, we manually compared the OCR output with the corresponding images for a subset of the images. This process revealed that the OCR process sometimes converted the letter “B” in the string to “6”, “8”, “9”, “0”, “O”. We therefore set the sales flag to 1 whenever we observed any of these characters in the string.

As we mention above, the letter “I” in the last spot of the string indicates that the price was imputed. We therefore create an imputation flag variable and set it to 1 if the letter “I” is observed at the end of the string. Our manual comparison of the OCR output and the corresponding images revealed that in some cases the letter “I” was read as “1” by the OCR procedure. The number “1” also appears as a part of the pricing sample part of the string. But our manual inspection indicated that the number “1” appearing at the end of the string and being preceded by another number, gave strong indication that the price was imputed. In these cases, therefore, we set the imputation flag to one.

In addition to this, several of the characters appearing in the first three spots of the string signal that the price was imputed. These include “A” (seasonal item not available), “C” (closeout or clearance sale)[1], “D” (ELI not available), “U” (unable to price), “T” (temporarily unable to price). The first three of these (A, C, D) may also appear in the pricing cycle portion of the string. We therefore set the imputations flag to one whenever we observe these characters in one of the first three spots of the strings. The other three characters

[1]BLS documentation we received indicated that “C” referred to “closeout or clearance sales” but out inspection of a subset of images indicated that these observations were imputed (e.g., they tended to include more than two numbers after the decimal place (fractions of a cent).

should not appear in other parts of the string. We therefore set the imputation flag to one whenever we observe these characters.

As in the case of prices, a price flag should appear in multiple "replicates" due to the structure of the Price Trend Listing, and in principle, these replicates should be the same. In cases where they disagree, we choose the value of the flag corresponding to the majority of replicates. This rule is meant to balance a concern for false positives and false negatives based on our inspection of a subset of cases where is disagreement. The fraction of observations that we drop because of imputation is 9.3%.

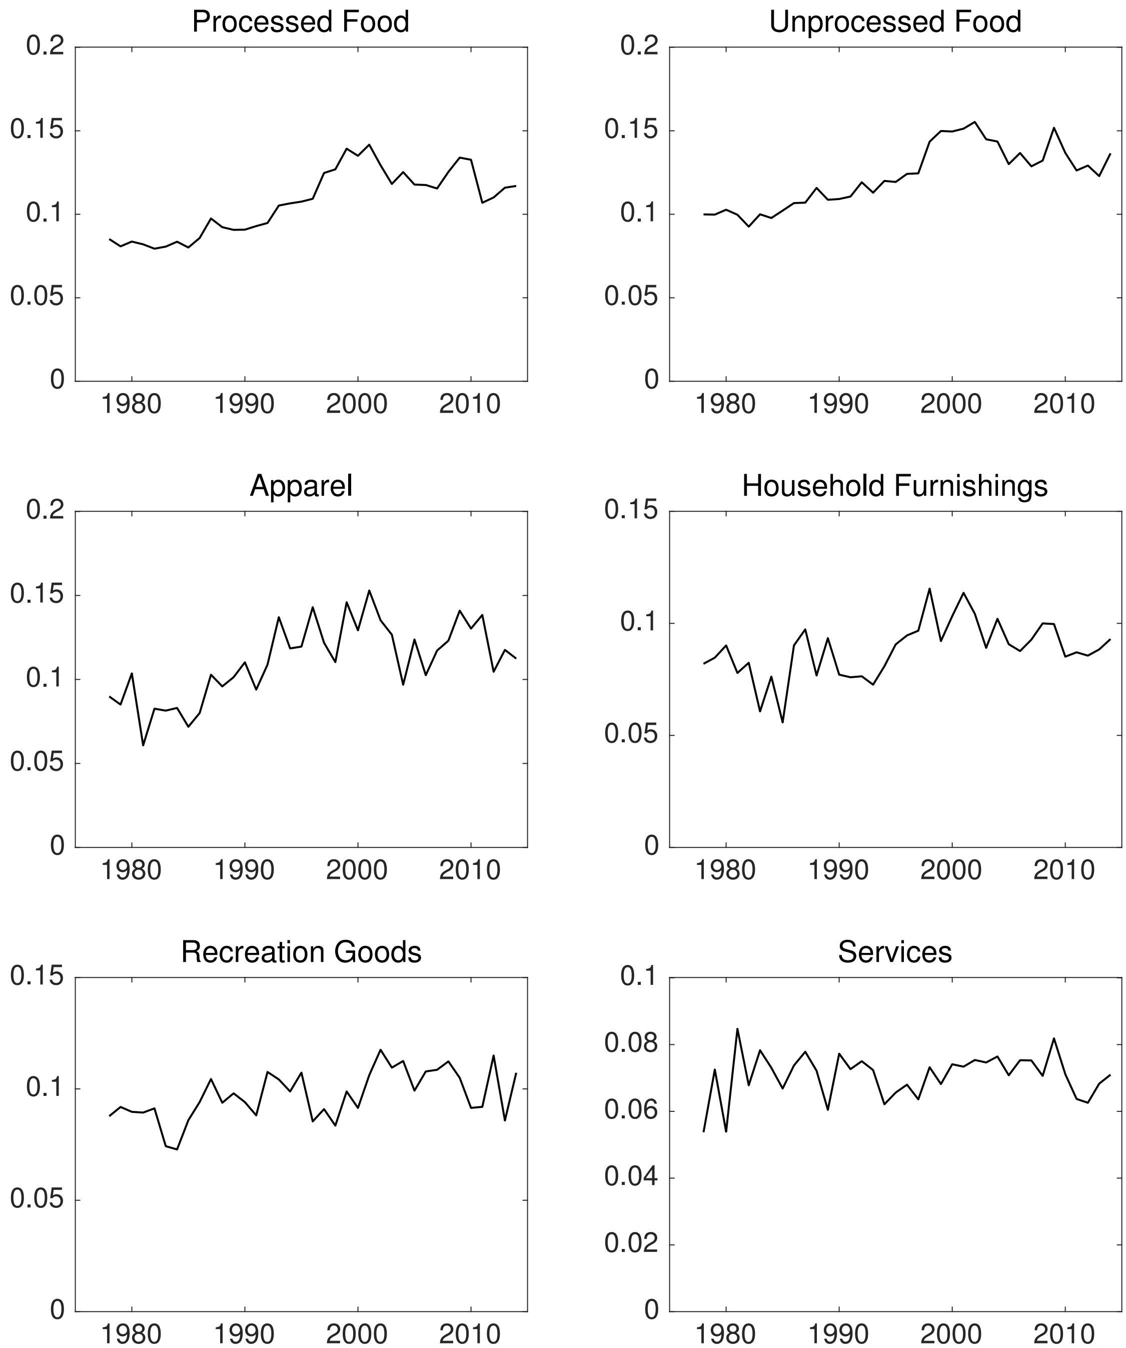

Figure A.1: Mean Absolute Size of Price Changes by Sector

Note: To construct the series plotted in this figure, we first calculate the mean absolute size of price changes in each ELI for each year. We then take the weighted median across ELI's within each sector for each year.

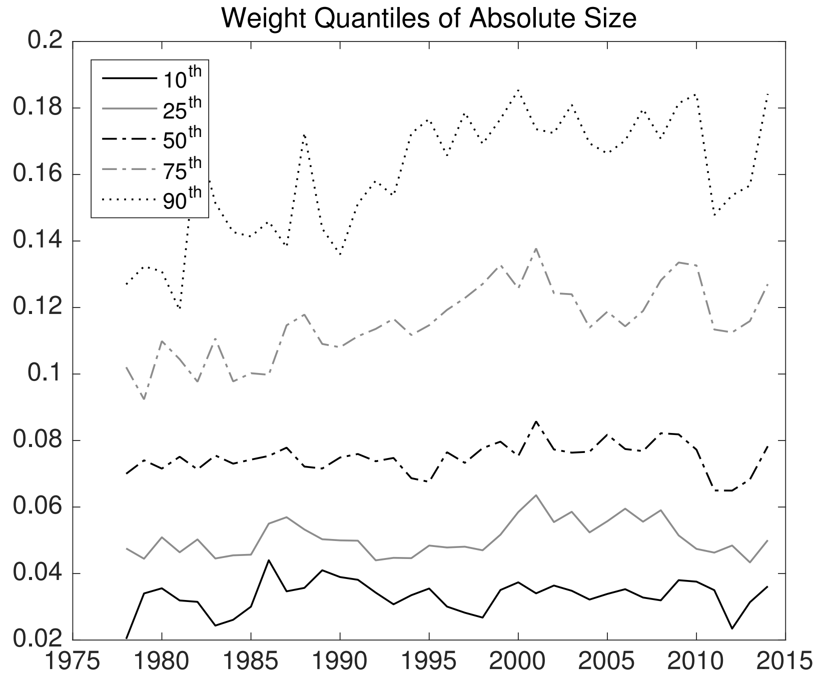

Figure A.2: Quantiles of the Absolute Size of Price Changes
Note: To construct the series plotted in this figure, we first calculate the mean absolute size of price changes in each ELI for each year. We then calculate quantiles of the distribution of the mean absolute size of price changes across ELI's for each year.

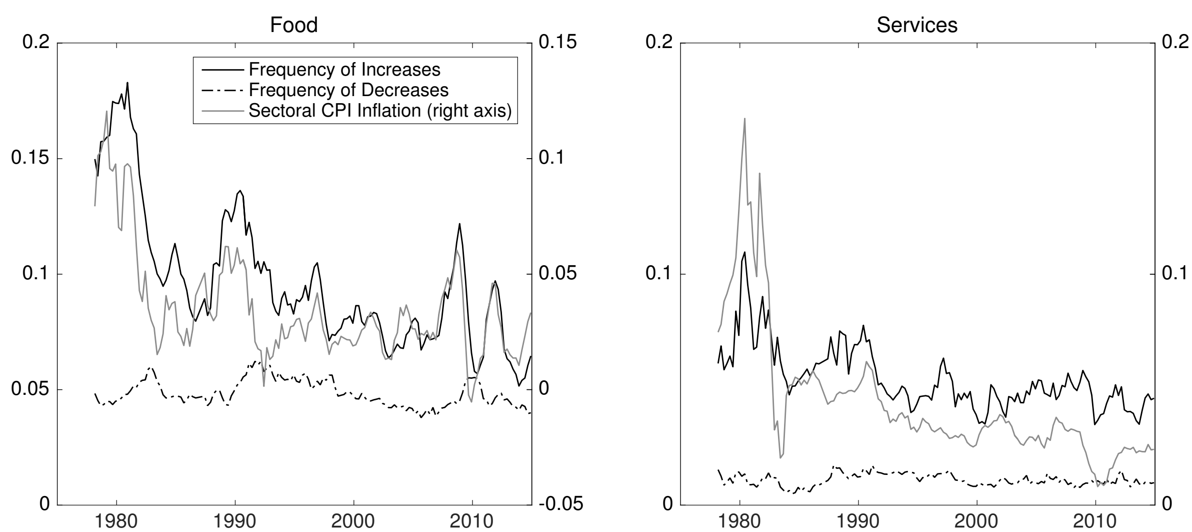

Figure A.3: Frequency of Price Increases and Decreases for Food and Services
Note: To construct the frequency series plotted in this figure, we first calculate the mean frequency of price increases and decreases in each ELI for each month. We then take the weighted median across ELI's for food and services separately and plot these in the two separate panels.

Appendix B

Appendix for Chapter 2

B.1 Data Set and Statistics

As mentioned in the main text, the data set we use for our empirical analysis is the micro data underlying the U.S. CPI for the period 1977-2014, with the previously unavailable period being 1977-1986. I worked intensively in the process of re-constructing this data set from the micro film made available by the Bureau of Labor Statistics. This process is described in detail in Appendix A.2., and it leaves us with a large data set that tracks the prices of individual, narrowly-defined products in a monthly or bi-monthly frequency. We then combine this data set with the existing CPI data (1987 onwards), and that forms the data set for our analysis. Figure A.1 below shows the size of our sample month by month. We plot both the number of non-missing available prices each month, as well as the number of price change observations available. The distinction is important, because we are always interested in price *change* statistics. The number of price observations is greater than the number of price change observations because for the price change to be observed in a particular month, we need both the current month's price, and last month's price. So when a product has a missing price for some month, the price change will be missing for that month and the following month.

Figure A.1: Number of observations by month

The BLS makes a considerable effort to ensure that the prices of individual products are tracked, so that the price changes cannot be attributable to changes in any product characteristics. This conforms with our goals very well, as we are also only interested in price changes of identical products. An individual product could be, for example, a two quart bottle of Diet Coke in a particular supermarket location in New York City, or a specific futon model in a particular furniture store in Los Angeles. The BLS also identifies whenever a product substitution occurs, or when a new "version" of a particular product is introduced. Since a change of version indicates that some characteristic of the product has changed, we treat a new version as an entirely new product, and only compute price changes by comparing price changes within identical versions. We compute price changes as the difference of the log price, or:

$$\Delta p_{it} = log(\frac{P_{it}}{P_{it-1}}).$$

As discussed in Section 2.3, we exclude observations for which there is any indication that the price was not actually observed but imputed, and for which the product was on sale. There are therefore missing observations in the price spells that we use. To compute the price change for any given month, we compare the price for that month to the previous month's price, when it is available. When the previous month's price is not available, we compare the current price to the price from two months before. Without this, we would have to drop a significant amount of data, as many prices are only sampled every two months. Since price changes are relatively infrequent, we believe that it is overwhelmingly likely that if a price changed between any two months, it only changed once, which means that we are observing the true price change, whether it occurred in the current or previous month. This is then not extremely important, as for much of our analysis we combine the price changes by quarter or year.

With the price change observations, we then form distributions of these price changes, keeping only the non-zero changes, for each period (either month, quarter, or year). For the dispersion and skewness statistics, we first separate observations into categories that we label major groups. There are thirteen of these, and table B.1 below provides a list, along with the share of expenditure weight that they represent.

Services represent the lion's share of the weight. We then compute the dispersion and skewness statistics from each major group, and for each time period we then take an

Table B.1: CPI group weight

Major Group	Weight (%)
Processed Food	8.2
Unprocessed Food	5.9
House Furnishings	5.0
Apparel	6.5
Transportation	8.3
Medical Care	1.7
Recreation	3.6
Edu. Supplies	0.5
Miscellaneous	3.2
Services	38.5
Utilities	5.3
Gasoline	5.1
Travel Services	5.5

expenditure-weighted average of the statistics, which represents the value of the statistics that we will use. If, for example, $Skew_{kt}$ is the skewness of the distribution of price changes in major group k and period t, then the value of skewness that we use in our analysis, $Skew_t$, is given by:

$$Skew_t = \sum_k w_k Skew_{kt}.$$

We follow the same method for the dispersion, and thus obtain time series for the skewness and dispersion of price changes. This also applies for the frequency, but there we calculate the frequency first by ELI, which is a much narrower category. That is because the frequency is merely an average of the dummy variable indicating whether a price has changed or not, and it is calculated based on the number of price change observations (zero or non-zero), while the other moments are only calculated based on the non-zero changes (which gives fewer observations). This means that the frequency can be estimated with reasonable precision by ELI. Finally, the expenditure weights that we use are those from the 1998 revision of the CPI, which are the latest ones available. Different weights were used for 1977-1987 and 1988-1997, but we keep the weights constant throughout the sample so that changes in the weights do not induce changes in the statistics that we estimate.

B.2 Computational Procedure and Calibration

We solve the sticky price models in this paper by value function iteration, following the method described in Nakamura and Steinsson (2010). The main difficulty with this method applied to this type of problem is that an important variable entering the firm's profit function is the aggregate price level. Since its future evolution depends on each firm's price, every firm's current state is, in principle, a state variable for all firms, making the problem intractable. To get around this, we follow the example of Krusell and Smith (1998) and approximate the law of motion of the price level with a finite number of moments, as in Nakamura and Steinsson (2010). In particular, we impose that firms perceive future inflation to depend only on future nominal aggregate demand (S_t, which is exogenous), and the current price level:

$$\pi_t \equiv log(\frac{P_t}{P_{t-1}}) = \Gamma(\frac{S_t}{P_{t-1}}).$$

Under this assumption, the state space can be reduced to three dimensions: the firm's idiosyncratic productivity (exogenous), the firm's relative price (choice variable), and real aggregate demand (C_t, which determines the real wage in equilibrium). The latter is endogenously determined, but the probability distribution of its future value is known fully with the law of motion of nominal aggregate demand, and the assumed law of motion of inflation.

The firm's problem can therefore be written recursively with the following Bellman equation:

$$V(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{S_t}{P_t}) = \max_{p_t(z)} \left\{ \Pi_t^R(z) + E_t \left[D_{t,t+1}^R V(A_{t+1}(z), \frac{p_t(z)}{P_{t+1}}, \frac{S_{t+1}}{P_{t+1}}) \right] \right\},$$

where $V(\cdot)$ is firm z's value function, $\Pi_t^R(z)$[1] is firm z's real profits at time t, and $D_{t,t+1}^R$ is the real stochastic discount factor between time t and t+1. Our procedure to solve the model then closely follows Nakamura and Steinsson (2010): First, we discretize the state variables and propose a guess for the function $\Gamma(\frac{S_t}{P_{t-1}})$ on the grid. Then, we solve for the firm's policy function, F[2], by value function iteration, using the proposed $\Gamma(\cdot)$ function, the stochastic processes for the exogenous variables (applied using the Tauchen (1986) method), and the menu cost structure of the firm's problem. We then check whether F and Γ are consistent, by computing the price level (and inflation) implied by F for each value on the $\frac{S_t}{P_{t-1}}$ grid and comparing it to the value given by Γ. If they are consistent, we stop and use F to simulate the models. If they are not consistent, we update Γ and go back to the value function iteration step and continue. To determine whether they are consistent, we compare the inflation values, grid point by grid point, and consider that they are consistent when the difference is smaller the difference in values between grid points.

The method described above applies to all the menu cost models (including the Calvo model). However, the imperfect information models are markedly different in several ways, and therefore require different methods. We solve these models using the same methods and

[1] It can be shown that the profit function under CES preferences and linear production using only labor can be written as $\Pi^R(A, \tilde{p}, C) = C\tilde{p}^{-\theta}[\tilde{p} - \frac{\omega C}{A}]$

[2] Because the value of the menu cost in our general model is stochastic, the policy function is also a function of the menu cost. However, because we assume that the menu costs are iid over time, they are not a state variable.

parameter values used in the original papers (Alvarez et al. (2011b) for the observation costs model; Woodford (2009) for the rational inattention model), and use the policy functions to simulate the models.

As mentioned in Section 2.2, the existing menu cost models and the Calvo model are calibrated to match the median frequency of price change and the median average size of price change in the data. The way we compute these moments is by first calculating the frequency of monthly price changes and the mean absolute value of price change by ELI-year. We then compute the median across the ELI frequencies for each year (to obtain an annual series for the median frequency) and to then take the mean across years. The average frequency that we obtain is 11.3%, and the average size of price change is 8.0%. For the Midrigan model (as well as our random menu cost model), we also target the fraction of price changes that are small (less than 1% in absolute value). We compute this as with the frequency and average size: evaluate fractions by ELI-year, and take weighted medians across ELI's. WE find a value of 12%. Table B.2 below shows the model-implied moments for the Golosov and Lucas, Midrigan, and Calvo models, as well as the random menu cost model from Chapter 3, and compares them to their empirical values:

Table B.2: Model implied moments

Model	Average Frequency (%)	Average Size (%)	Fraction Small (%)
Golosov and Lucas	11.1	8.0	0
Midrigan	11.0	8.0	12.4
Calvo	11.0	7.9	17.3
Random MC	11.3	8.0	12.2
Data	11.3	8.0	12.0

All the models match the frequency and size moments almost exactly, and the Midrigan and random menu cost models match the fraction of small changes very closely. The Calvo and Golosov and Lucas models over- and undershoot the empirical value, respectively, as they do not target it. Table B.3 below shows the parameter values that we choose for these models.

Table B.3: Parameter values for models

Parameter	Golosov and Lucas	Value
χ	Menu cost (as share of steady state revenue)	0.019
σ_ϵ	Std. dev. of idiosyncratic tech. shocks	0.042
	Midrigan	
χ^{High}	Menu cost (when positive)	0.034
σ_ϵ	Std. dev. of idiosyncratic tech. shocks	0.076
p_z	Probability of free price change	0.037
p_ϵ	Probability of receiving idio. shock	0.153
	Calvo	
α	Probability of price change	0.111
σ_ϵ	Std. dev. of idiosyncratic tech. shocks	0.197

For the multi-sector model, we use the same parameter values as in Nakamura and Steinsson (2010), which make the model match the average frequency and size of price change for each of 14 sectors.

B.3 Additional Empirical Results

In Section 2.2, we presented results on the empirical result between inflation and various price change moments, using both scatter plots and regressions. For the scatter plots, the measure of inflation that we used was Core PCE inflation, which excludes food and energy prices that tend to be quite volatile (and that could be influenced by sectoral shocks that we do not consider in the models). In addition, since the PCE index is chained, it tends to yield a lower value for inflation than the CPI. However, for the regressions, we used CPI inflation because we include expected inflation as a control, and the survey of inflation expectations asks about expectations of CPI inflation specifically. We therefore used CPI inflation to make the two variables more comparable. In Figure A.2 below, we plot the twelve month log change for both indexes. They both co-move very strongly, although the peak is much higher for the CPI.

Figure A.2: Inflation

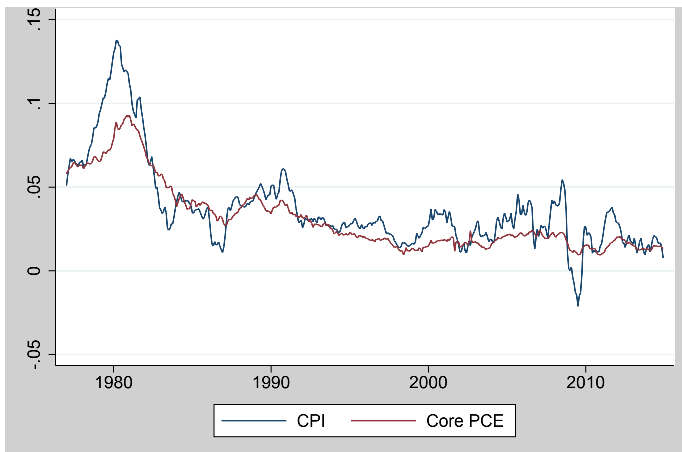

In this section we show that our results do not depend on which inflation measure we use, so we present scatter plots with CPI inflation, and regression results with Core PCE inflation as the regressor. The only difference that this makes is that in the regressions, the absolute value of the coefficients on inflation are slightly larger, because core PCE inflation does not attain as high a value, so the estimated slope of the moments on inflation is smaller. We also present results using series filtered by a moving average smoother and seasonally adjusted by removing quarterly dummies. Again, the the same results hold, but they come out a bit more clearly. For all of these results, we focus on using the quarterly inflation and moment series, although the same results would hold with the monthly and annual series.

Figures A.3-A.6 below present scatter plots of the smoothed moment and inflation series.

Figure A.3: Frequency of price change & inflation smoothed, quarterly

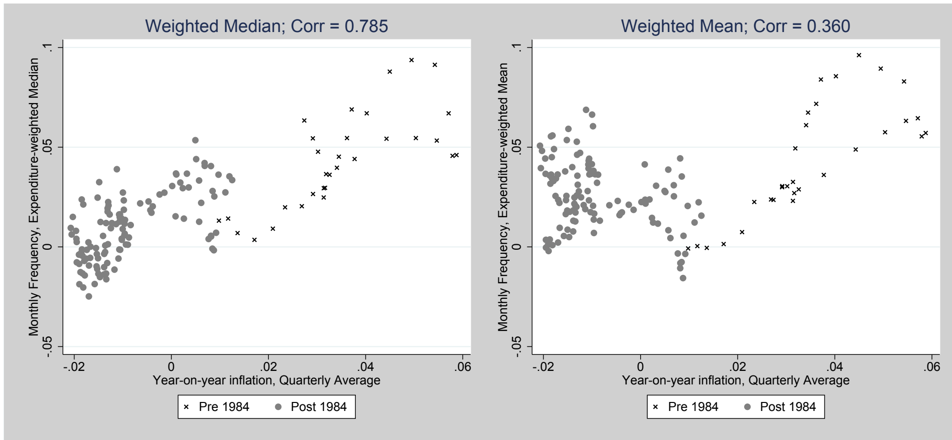

Figure A.4: IQR of price change & inflation smoothed, quarterly

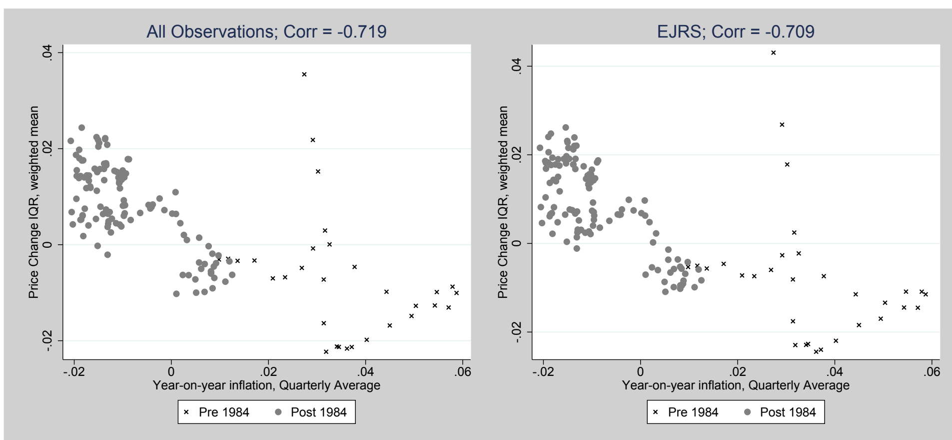

Figures A.7-A.10 are scatter plots using CPI inflation.

The patterns in these scatter plots are the same as in the ones presented in Section 2.3. We further confirm these results with the regression tables below.

What these tables show is that while the size of the coefficients varies somewhat across specifications, the results presented in Section 2.2 still hold: the frequency of price change

Figure A.5: Skewness & inflation smoothed, quarterly

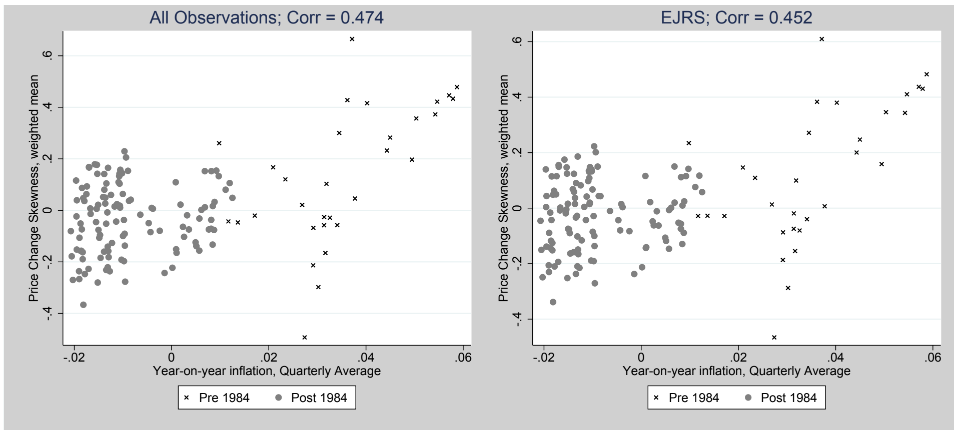

Figure A.6: Kelly skewness & inflation smoothed, quarterly, corr=0.734

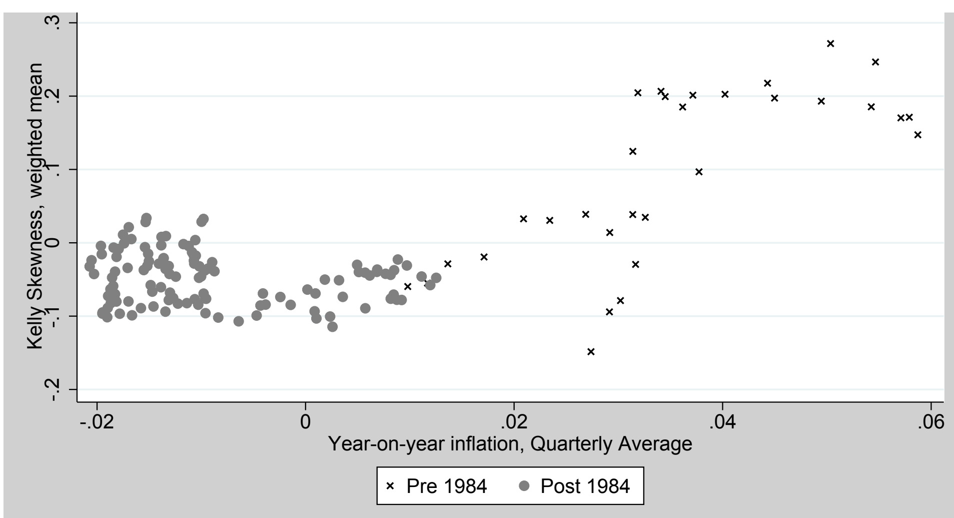

Figure A.7: Frequency of price change & CPI inflation, quarterly

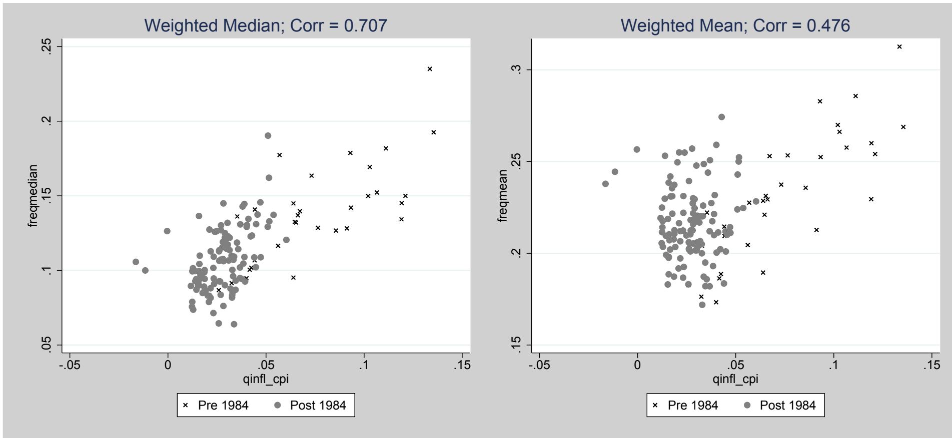

Figure A.8: IQR & CPI inflation, quarterly

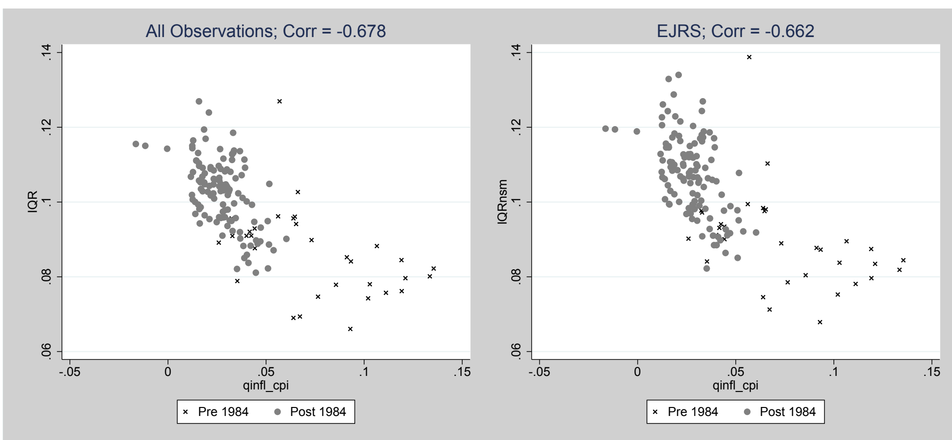

Table B.4: Core inflation as regressor - frequency

Coefficients for Frequency Regressions				
	Weighted Median		Weighted Mean	
Specification	1977-2014	1985-2014	1977-2014	1985-2014
All	0.906***	1.362***	-0.046	-0.231
	(0.271)	(0.313)	(0.244)	(0.305)
Fed Dummies	1.248***	1.503***	0.978***	0.281**
	(0.220)	(0.214)	(0.223)	(0.258)
Inflation Only	0.877***	1.083***	0.374**	-0.580**
	(0.122)	(0.253)	(0.173)	(0.296)

Figure A.9: Skewness & CPI inflation, quarterly

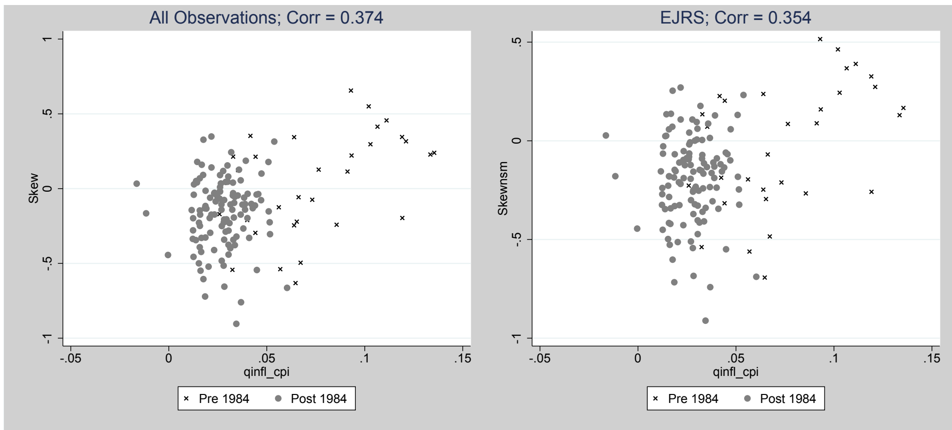

Figure A.10: Kelly skewness & CPI inflation, quarterly, corr=0.674

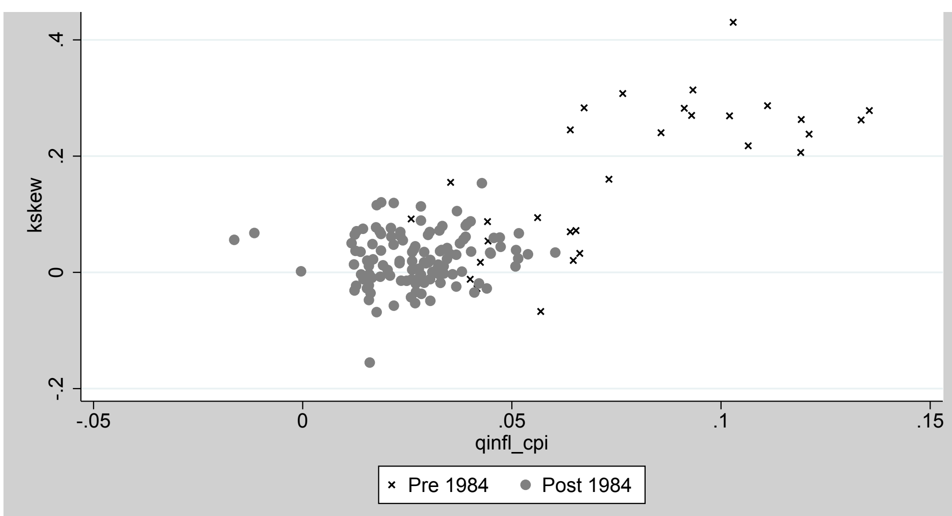

Table B.5: Smoothed and seasonal adjusted series - frequency

Coefficients for Frequency Regressions				
	Weighted Median		Weighted Mean	
Specification	1977-2014	1985-2014	1977-2014	1985-2014
Fed & Expected Infl	0.711***	0.796***	0.462	0.326*
	(0.125)	(0.210)	(0.138)	(0.189)
Fed Dummies	0.778 ***	0.889***	0.723***	0.284*
	(0.075)	(0.207)	(0.109)	(0.163)
Inflation Only	0.716***	0.824***	0.437***	-0.178
	(0.062)	(0.223)	(0.105)	(0.240)

Table B.6: Core inflation as regressor - IQR

Coefficients for IQR Regressions				
	All Observations		EJRS	
Specification	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	-0.412***	-0.676***	-0.461***	-0.803***
	(0.060)	(0.081)	(0.068)	(-0.086)
Fed Dummies	-0.354***	-0.686***	-0.401***	-0.824***
	(0.082)	(0.095)	(0.095)	(0.099)
Fed & Expected Infl	-0.366***	-0.485**	-0.429***	-0.594***
	(0.127)	(0.117)	(0.142)	(0.128)

Table B.7: Smoothed and seasonal adjusted series - IQR

Coefficients for IQR Regressions				
	All Observations		EJRS	
Specification	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	-0.301***	-0.493***	-0.330***	-0.561***
	(0.043)	(0.073)	(0.047)	(0.086)
Fed Dummies	-0.241***	-0.495***	-0.249***	-0.556***
	(0.048)	(0.084)	(0.054)	(0.097)
Fed & Expected Infl	-0.164**	-0.377**	-0.178**	-0.431 ***
	(0.069)	(0.073)	(0.075)	(0.083)

Table B.8: Core inflation as regressor - skewness

Coefficients for Skewness Regressions				
	All Observations		EJRS	
Specification	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	4.537***	2.131	4.315***	1.658
	(1.306)	(2.062)	(1.285)	(1.895)
Fed Dummies	7.546***	3.716	6.997***	3.396
	(1.686)	(2.270)	(1.572)	(2.087)
Fed & Expected Infl	4.683	6.224*	4.039*	5.991
	(2.870)	(3.316)	(2.657)	(3.136)

Table B.9: Smoothed and seasonal adjusted series - skewness

Coefficients for Skewness Regressions				
	All Observations		EJRS	
Specification	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	3.656***	1.208	3.263***	0.699
	(0.776)	(1.222)	(0.776)	(1.148)
Fed Dummies	3.683***	0.925	3.404***	0.688
	(0.689)	(1.349)	(0.680)	(1.245)
Fed & Expected Infl	0.969	0.453	0.785	0.152
	(1.206)	(1.504)	(1.182)	(1.367)

Table B.10: Core inflation as regressor - Kelly skewness

Coefficients for Kelly Skewness Regressions		
	All Observations	
Specification	1977-2014	1985-2014
Inflation Only	2.973***	-0.603
	(0.537)	(0.512)
Fed Dummies	4.035***	0.504
	(0.713)	(0.606)
Fed & Expected Infl	2.066**	0.136*
	(1.047)	(0.721)

Table B.11: Smoothed and seasonal adjusted series - Kelly skewness

Coefficients for Kelly Skewness Regressions		
	All Observations	
Specification	1977-2014	1985-2014
Inflation Only	2.465***	-0.088
	(0.342)	(0.394)
Fed Dummies	2.479***	0.282
	(0.329)	(0.435)
Fed & Expected Infl	1.636**	0.204
	(0.731)	(-0.430)

rises with inflation, the dispersion falls, and the skewness does not fall with inflation (the relationship is positive but not significant in the low inflation period, and positive and mostly significant in the whole sample).