

Asymmetric Inflation Expectations, Downward Rigidity of Wages, and Asymmetric Business Cycles*

David Rezza Baqaee[†]

December 23, 2015

Abstract

Household expectations of the inflation rate are much more sensitive to inflation than to disinflation. To the extent that workers have bargaining power in wage determination, this asymmetry in their beliefs can make wages respond quickly to inflationary forces but sluggishly to deflationary ones. I microfound asymmetric household expectations using ambiguity-aversion: households, who do not know the quality of their information, overweight inflationary news since it reduces their purchasing power, and underweight deflationary news since it increases their purchasing power. I embed asymmetric beliefs into a general equilibrium model and show that, in such a model, monetary policy has asymmetric effects on employment, output, and wage inflation in ways consistent with the data. Although wages are downwardly rigid in this environment, monetary policy need not have a bias towards using inflation to grease the wheels of the labor market.

*I thank Natalie Bau, John Campbell, Raj Chetty, Emmanuel Farhi, Gita Gopinath, Nathaniel Hendren, Cosmin Ilut, Larry Katz, David Laibson, Damjan Pfajfar, Ricardo Nunes, Ricardo Reis, Alp Simsek, Silvana Tenreyro and seminar participants at various conferences for helpful comments. I also thank Ricardo Perez-Truglia for his comments and for sharing his data.

[†]Department of Economics, London School of Economics and Political Science, d.r.baqaee@lse.ac.uk.

1 Introduction

Keynesian macroeconomic theory posits that sticky wages are a crucial feature of labor markets. Rigid wages can cause involuntary unemployment, amplify fluctuations in employment at business cycle frequencies, and break monetary neutrality. In recent years, a large number of empirical papers have shown not only that wages are very sticky, but that there is a clear asymmetry in the way that wages are sticky. In particular, wages appear more flexible when they are rising than when they are falling. Examples include Barattieri et al. (2010), Dickens et al. (2007), Daly and Hobijn (2013), as well as the seminal contribution by Bewley (1999). Furthermore, recent work by Kaur (2014) shows that downward nominal rigidities distort labor market outcomes in rural India.

This paper argues that informational frictions for households can help to explain the asymmetric adjustment of wages during the business cycle, and analyzes the positive and normative implications of these frictions. The main contributions of this paper are: (1) it empirically documents the existence of an asymmetry in how households form their expectations of inflation; households are much better at anticipating accelerations in the inflation rate than decelerations. In a typical model, this asymmetry in household beliefs feeds into asymmetry in wage-setting where wages respond more vigorously to inflationary forces than to disinflationary forces. This makes demand-driven business cycles asymmetric. (2) It micro-founds the source of asymmetric belief formation in an equilibrium model where households are ambiguity-averse and make robust decisions. The advantage of having the microfoundation is that it allows us to investigate policy questions. In particular, with this microfoundation, monetary policy need not have a systematic inflationary bias despite the existence of downward rigidities.

The intuition for the microfoundation is simple. When negotiating their wages, workers observe their nominal wages perfectly, but foresee the real cost of the goods and services that they will consume imperfectly. As in Lucas (1973), workers face a signal extraction problem when trying to determine their purchasing power. Furthermore, households' signals of the general price level are subject to Knightian uncertainty, since households do not precisely know how informative their signals are. Not only is the information households receive of ambiguous quality, but the mapping from aggregate news to each individual's idiosyncratic consumption basket is also unclear. If this Knightian uncertainty induces households to try to minimize their worst-case losses, then they will be more sensitive to inflationary news than disinflationary news because, for a fixed nominally denominated employment contract, inflationary news lowers their purchasing power whereas disinflationary news raises it. This asymmetry of beliefs can then show up in wage-setting since distrustful workers will, to the extent that they can, refuse wage cuts in the presence of deflation, but demand wage increases in the presence of inflation.

This distrustful attitude of workers towards the inflation rate is attested to in many surveys. For example, according to Shiller (1997), the "biggest gripe about inflation" expressed by 77% of the general public is that "inflation hurts my real buying power. It makes me poorer." Interestingly,

only 12% of economists chose this answer.¹ The households' answer makes sense given the problem most households are facing, since over short horizons, households can treat their wages as known and exogenous to the inflation rate.

I embed ambiguity into a general equilibrium model and show that monetary shocks have asymmetric effects on wage inflation and output. In particular, positive monetary shocks result in high wage-inflation and small output booms, since households react strongly to the inflationary signals by demanding wage increases, whereas, negative monetary shocks cause relatively little wage-disinflation but are accompanied with large output drops, since households distrust the disinflationary signals and refuse wage cuts. I show that the model's predictions about the asymmetric impact of monetary policy on wage-inflation is consistent with aggregate time series data for the US.

The model of Knightian uncertainty I use is the one axiomatized by Gilboa and Schmeidler (1989). In this framework, workers have multiple priors about the information content of their signals, and they act according to their worst-case prior when making decisions. A similar modelling device is used by Epstein and Schneider (2008) in the context of asset pricing to model skewness in asset returns, by Ilut (2012) to model departures from uncovered interest rate parity, and Ilut et al. (2014) to model employment decisions of firms. Kuhnen (2012) finds empirical evidence in support of asymmetric learning in the context of financial markets, where agents are overly pessimistic in the loss region.²

The asymmetry implied by this model substantially alters the welfare costs of business cycles when compared to Lucas (1987). Whereas, in the Lucas model, positive shocks cancel out negative shocks so that the welfare cost of fluctuations is second order, in this model, positive demand shocks do not cancel with negative demand shocks, and stabilization policy reaps first order gains. In a paper about external devaluation, Schmitt-Grohé and Uribe (2011) have also recently drawn attention to this point. This harks back to the point made by De Long and Summers (1988) that demand stabilization may fill in the troughs without shaving the peaks. I also investigate the conduct of monetary policy in my model. The received wisdom in the literature, following Akerlof et al. (1996), is that if wages are downwardly rigid, then central banks should have an inflationary bias to "grease the wheels" of the labor market. This way real wage cuts can be masked by a positive inflation rate. Unfortunately, in my model, inflationary biases from the central bank are not helpful since household expectations adjust to take them into account. Any inflationary bias built into central bank policies are undone by endogenously-formed household expectations. In other words, the model predicts an asymmetric equilibrium, but the central bank is powerless to

¹Instead, the most popular reason given by economists was, "inflation makes it hard to compare prices, forces me to hold too much cash, and is inconvenient." Only 7% of households chose this answer.

²Another recent paper that incorporates ambiguity aversion into a macroeconomic model is by Ilut and Schneider (2014). However, my results differ markedly from theirs both in terms of the research question and the set up of the model. Ilut and Schneider (2014) are interested in how ambiguity about the level of productivity affects output, whereas I am interested in how ambiguity about the informativeness of price signals affect wage setting.

do anything about the asymmetry.

Other theoretical treatments of downwardly rigid wages often take the rigidity as given and investigate its consequences (e.g. Daly and Hobijn (2013); Schmitt-Grohé and Uribe (2011); Kaur (2014); Akerlof et al. (1996); Hall (2005); Elsby (2009)). These models are usually motivated by an exogenous fairness norm, and assumptions about the function relating wages to worker effort. Akerlof (1982) is a seminal paper in this strand of the literature. Other attempts to microfound downward wage rigidity are based on implicit contracts, where firms insure their workers against fluctuations by uncoupling the real wage from marginal product of labor; a leading example is Holmstrom (1983), but this literature is focused on real wages and does not bring inflation into the analysis.

The outline of this paper is as follows. In section 2, I set out a basic partial equilibrium model with ambiguity aversion that demonstrates the mechanism. In section 3, I endogenize prices and output, and study the effects of monetary shocks on real and nominal variables. In section 4, I present empirical evidence in favor of asymmetric adjustment of beliefs, and time series evidence that wage inflation responds asymmetrically to positive and negative monetary shocks as predicted by the model. I also discuss the extent to which the model can explain the cross-sectional distribution of wage-changes. In section 5, I investigate a simple monetary policy problem. I summarize and conclude in section 6.

2 Partial Equilibrium Model

Consider the following partial equilibrium model that establishes the intuition for the rest of the paper. Suppose that there is an employer and a worker. The worker has log utility in his real wage, is endowed with a unit of labor, and an exogenous outside option d (I take this to be the utility of leisure). In other words, his preferences are given by

$$u(w_t/p_t, x_t) = \log(w_t/p_t)\mathbf{1}(x_t = 1) + d\mathbf{1}(x_t = 0),$$

where x_t is a binary variable for whether or not he works, w_t is the nominal wage, p_t is the price level in period t and d is an exogenous outside option. The employer makes a nominal wage offer w_t to the worker, who then chooses whether or not to work. If the worker does not work, he receives the exogenous outside option d .

The worker chooses to work if

$$E_t \left(\log \left(\frac{w_t}{p_t} \right) \right) \geq d,$$

where the expectation is taken with respect to the worker's information set. Since the nominal wage is always known with certainty, we can rearrange this expression to get the lowest wage for

which the worker will work

$$w_t = \exp(d + E_t(\log(p_t))). \quad (1)$$

Suppose that workers receive a public signal s_t about the inflation rate. Then we can rewrite (1) as

$$\log(w_t) = d + \log(p_{t-1}) + E(\log(\pi_t)|s_t),$$

where π_t is the inflation rate from period $t - 1$ to t and workers are assumed to know the price level in the previous period. This equation makes clear that the wage inherits the properties of the conditional expectation function when viewed as a function of the signal s_t . If households' expectations of inflation are, for some reason, asymmetric (they rise more quickly than they fall), then the wage will also behave asymmetrically with respect to inflationary pressures.

To demonstrate how ambiguity-aversion can deliver convex conditional expectations, suppose that the price level p_t is given by

$$\log(p_t) = \mu + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2),$$

and both the worker and the employer know μ , but do not know the shock ε_t . Let s_t be a noisy public signal of the price shock ε_t ,

$$s_t = \varepsilon_t + \varepsilon_s, \quad \varepsilon_s \sim \mathcal{N}(0, \sigma_s^2),$$

where ε_s is noise in the signal with variance σ_s^2 . Note that

$$\varepsilon_t | s_t, \mu, \sigma^2, \sigma_s^2 \sim \mathcal{N}\left(\frac{\sigma^2}{\sigma^2 + \sigma_s^2} s_t, \frac{\sigma_s^2 \sigma^2}{\sigma^2 + \sigma_s^2}\right).$$

This means we can rewrite the work condition (1) as

$$w_t \geq \exp\left(d + \mu + \frac{\sigma^2}{\sigma^2 + \sigma_s^2} s_t\right).$$

Now Suppose that there is ambiguity about how informative the signal s_t is. That is, the signal-to-total variance ratio $\frac{\sigma^2}{\sigma^2 + \sigma_s^2}$ is unknown. For example, suppose that the worker knows only that the standard deviation of the noise term lies in some interval $\sigma_s \in [\underline{\sigma}_s, \bar{\sigma}_s]$. Gilboa and Schmeidler (1989) have axiomatized and provided a representation theorem for the preferences of such agents. In particular, such agents follow a minmax procedure, where they make decisions that maximize their worst-case expected utility. In other words, when information quality is ambiguous, expression

(1) becomes

$$\begin{aligned}
 w_t &= \max_{\sigma_s \in [\underline{\sigma}_s, \bar{\sigma}_s]} \exp(d + \mu + E(\log(p_t)|s_t)) \\
 &= \exp\left(d + \mu + \max_{\sigma_s} E(\log(p_t)|s_t)\right) = \exp\left(d + \mu + \tilde{E}(\log(p_t)|s_t)\right). \tag{2}
 \end{aligned}$$

For notational convenience, I denote the expectations taken with respect to the worst-case prior by \tilde{E} . Maximizing worst-case expected outcomes in this way is very similar to the “Robustness” framework proposed by Hansen and Sargent (2008). The specific information structure I use is similar to the one posited by Epstein and Schneider (2008), who use it to study skewness in asset prices.

In this section, the source and nature of the ambiguity is not important. It could be due to Knightian uncertainty about official statistics, or a reduced form representation of the fact that consumers have idiosyncratic consumption baskets and there is ambiguity about the extent to which official statistics are relevant to one’s individual consumption basket. We will return to these issues later, for now, let us take (2) as given.

Expression (2) implies that $\sigma_s = \underline{\sigma}_s$ when $s_t \geq 0$, and $\sigma_s = \bar{\sigma}_s$ when $s_t < 0$. In figure 1, we see an asymmetry in the adjustment of wages to signals of the price level. In particular, wages increase much more rapidly in response to inflationary signals than they fall in response to disinflationary signals. Similar results obtain for the more general constant-relative-risk-aversion utility case, and in the case with ambiguity in the variance of the monetary shock σ instead of ambiguity in the variance of the noise term σ_s . The intuition for this result is that households take inflationary news




Figure 1: Critical wage as a function of s .

very seriously, since the worst case scenario is that bad news is informative. On the other hand, households distrust or ignore disinflationary news, since the worst case scenario is that the good news is noisy.

To motivate this basic mechanism, we can look for evidence of this asymmetry by using inflation expectation surveys of households. I use the Michigan Survey of Consumers. Denote inflation in

period t by π_t and median household inflation expectations of inflation 12 months ahead by $\hat{\pi}_{t+12|t}$.³ In figure 2, expected revisions to the inflation rate $\hat{\pi}_{t+12|t} - \pi_t$ are plotted against actual changes to the inflation rate $\pi_{t+12} - \pi_t$ from the start of the Great Moderation. As expected, we see a steep convexity, indicating that the median household’s expectations of inflation are more responsive to positive rather than negative changes to inflation.

In figure 3, we see that the median forecasts made by professional forecasters, taken from the Survey of Professional Forecasters, do not exhibit this convexity. This suggests that source of the asymmetry, at least in the United States, is in how households process information, rather than in the information itself. Furthermore, the fact that there is an asymmetry in the beliefs of households is unique to this model and would not be found in preference-based theories that rely on loss-aversion or fairness norms. This empirical finding is in the same spirit as the experimental results of Fehr and Tyran (2001) who emphasize that subjects respond weakly to deflationary shocks and strongly to inflationary shocks, although my results are about accelerations and decelerations in the inflation rate.⁴

3 A simple general equilibrium model

Now, we endogenize prices and output, so that we can meaningfully study employment and the impact of monetary shocks. In this section, I use a highly stylized two period model, with the minimum number of ingredients to demonstrate the basic insight. In appendix III, I show how this type of story can be embedded into an otherwise standard New Keynesian model with wage rigidities.

Consider a two period model with a representative firm and two types of households. The first household type are the workers, who supply the firm with labor. The second household type are the owners of the firm, who inelastically supply the firm with capital. The separation of the laborers and capitalists means that workers do not take the profits of the firm or the returns to capital into account when making their labor decisions. In period 1, nature sets money supply M , and there is a noisy public signal s of M . The firm posts a nominal wage W conditional on the signal, and workers decide whether or not to apply for a job. In period 2, M becomes common knowledge, the firm chooses the fraction of the population it wishes to employ, and workers spend the money supply on consuming the output. Intuitively, in period 2, nominal wages are fixed, but the price level changes; this proxies a world where nominal wages are fixed over the length of a contract while prices continue to change.

In period 1, following our earlier discussion, the firm sets the wage to equate the utility of

³I use the median, rather than the mean, forecast since it is less sensitive to the existence of extreme outliers in survey responses.

⁴Both figures 2 and 3 are from after the Volcker disinflation. The period of structural change from 1978-1982 does not show evidence of asymmetries, and I examine this period in more detail in section 4.1.1.




Figure 2: Forecast revisions of the annual inflation rate by the median household in the Michigan Survey of Inflation Expectations from 1983-2015, plotted against realized changes in the annual inflation rate as measured by the CPI.




Figure 3: Forecast revisions by the median professional forecaster in the Michigan Survey of Inflation Expectations from 1983-2012, plotted against realized changes in the annual inflation rate as measured by the CPI, with a nonparametric curve of best fit.

working with the outside option

$$\tilde{E}(u(C^w)|s) = d, \tag{3}$$

where C^w is consumption of workers when employed and d is the exogenous outside option in utility terms. Let households have log utility so that the wage, in period 1, is given by

$$\tilde{E}(\log(C^w)|s) = \tilde{E}(\log(W/P)|s) = d. \quad (4)$$

This makes the households indifferent between working and consuming their outside option. Rearrange this for the wage to get

$$\log(W) = d + \tilde{E}(\log(P)|s). \quad (5)$$

In period 2, the stock of money is revealed, the firm sets marginal product of labor equal to the real wage

$$f_L(K, L) = \frac{W}{P}. \quad (6)$$

To give a role to money, suppose that households have a cash-in-advance constraint, so that their total expenditures have to equal the money supply

$$P(C^w + C^k) = M, \quad (7)$$

where C^k is the consumption of the capitalists. Market clearing for the consumption good implies that

$$C^w + C^k = f(K, L).$$

Let the firm's production technology be given by

$$f(L) = K^{1-\alpha}L^\alpha, \quad \alpha \in (0, 1),$$

then the firm's first order condition (6) implies that

$$L = \left(\frac{W}{\alpha P} \right)^{\frac{1}{\alpha-1}} K.$$

Combine this with market clearing, and (7) to get

$$\frac{M}{P} = C = K^{1-\alpha}L^\alpha = \left(\frac{W}{\alpha P} \right)^{\frac{\alpha}{\alpha-1}} K. \quad (8)$$

Rearrange this to get

$$P = \frac{M^{1-\alpha}W^\alpha}{\alpha^\alpha K}, \quad (9)$$

so the equilibrium price is a geometric average of the money stock and the wage. Substitute this

expression for P into the wage setting equation (5) to get

$$\log(W) = \frac{d}{1-\alpha} + \tilde{\mathbb{E}}(\log M|s) - \frac{\alpha}{1-\alpha} \log(\alpha) - \frac{1}{1-\alpha} \log(K) \quad (10)$$

as the wage in equilibrium. To get equilibrium output, substitute the equilibrium price (9) into (7) to get

$$C = \frac{M}{P} = \frac{\alpha^\alpha M}{M^{1-\alpha} W^\alpha} K = \left(\frac{\alpha M}{W} \right)^\alpha K.$$

Finally, equilibrium labor is given by using the production function $K^{1-\alpha} L^\alpha = C$ to get

$$\log(L) = \log(\alpha) + \log(M) - \log(W) + \log(K) = \text{const} + \log(M) - \tilde{\mathbb{E}}((\log(M)|s)). \quad (11)$$

Equations (10) and (11) show that if, for whatever reason, conditional expectations of the money shock as a function of the signal are more convex than the signal is as a function of the money shock, we should observe asymmetries in wage-setting and in employment fluctuations.

As before, ambiguity about the shocks to money supply can deliver such asymmetric conditional expectations. So, suppose that s is a normal noisy signal of the shock to $\log(M)$. Denoting $\log(M)$ by m ,

$$m|s \sim \mathcal{N} \left(\mu + \frac{\sigma^2}{\sigma^2 + \sigma_s^2} s, \frac{\sigma^2 \sigma_s^2}{\sigma^2 + \sigma_s^2} \right).$$

Denote the signal-to-total variance ratio $\frac{\sigma^2}{\sigma^2 + \sigma_s^2}$ by ψ , and note that $\psi \in [\underline{\psi}, \bar{\psi}]$. Then the equilibrium wage (10) is

$$\log(W) = \frac{d}{1-\alpha} + \mu + \bar{\psi} s \mathbf{1}(s \geq 0) + \underline{\psi} s \mathbf{1}(s < 0) - \frac{\alpha}{1-\alpha} \log(\alpha),$$

and equilibrium employment (11) is

$$\log(L) = \frac{1}{1-\alpha} \log(\alpha) + m - \frac{d}{1-\alpha} - \mu - \bar{\psi} s \mathbf{1}(s \geq 0) - \underline{\psi} s \mathbf{1}(s < 0). \quad (12)$$

In the benchmark case of full information, $\underline{\psi} = \bar{\psi} = 1$, employment is independent of monetary shocks and the wage is a linear function of the size of the monetary shock. This corresponds to the neoclassical case without frictions. In the case with no ambiguity, $0 \leq \underline{\psi} = \bar{\psi} < 1$, the nominal wage and the level of employment are linear in monetary shocks. The intuition here is the same as for the Lucas (1973) islands model. In the case with ambiguity, $0 \leq \underline{\psi} < \bar{\psi} \leq 1$, shown in figure 4, we have asymmetric nominal wage adjustment and employment fluctuations in response to monetary shocks. So we recover the intuition from the partial equilibrium model in section 2, but with additional predictions about the level of employment and the effects of monetary policy (both of which were absent in the partial equilibrium model).

In the section 4, we consider the empirical evidence for the validity of this model's assumptions




Figure 4: The nominal wage and employment as a function of shocks to money supply.

and predictions. Before that however, I discuss the plausibility of some of the model’s assumptions.

3.1 Why is Inflation Bad?

The microfoundation I propose to explain asymmetric expectations relies on the fact that households view unexpected inflation as worse than unexpected disinflation. To macroeconomists and central bankers, this may seem counterintuitive, since inflation, employment, and output tend to comove in the short run. Positive shocks to the money supply should be viewed as better events than negative shocks to the money supply by the household. The reason the households in this model do not behave that way is because they make their decisions in partial equilibrium: their view of inflation holds fixed their nominal wage, because when they make the decision to work or not, they already know the nominal wage offer. The fact that, in general equilibrium, output and inflation comove is irrelevant to them.

To test to see whether this set up has any empirical basis, I use the Michigan Survey of Consumers to check the association between expected income and expected inflation. The Michigan Survey, described in detail in appendix I, asks consumers the following questions:

- (a) By about what percent do you expect your income to (increase/decrease) during the next 12 months?
- (b) By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?

I subtract (b) from (a) to get a measure of household's expectations of their real income growth. I then regress expected real income growth on expected inflation, with time and individual fixed effects. The coefficient on expected inflation in this regression is -0.73 with a t-statistic of 18.58. A coefficient of -1 would suggest that households expect their nominal wages to not respond at all to expected inflation, and a coefficient of 0 would suggest that households expect nominal wage changes to offset inflation completely. An estimate of -0.73 therefore suggests that the effect of inflation expectations on real income expectations are highly significantly negative.

A second reason why inflation could be a boon rather than a bane is if the workers are debtors. If debts are denominated in nominal terms, unexpected inflation is good news since it erodes the real value of debt obligations. The logic here is that a highly indebted worker's net worth can increase as a result of unexpected inflation. However, this ignores the fact that this unexpected inflation would also erode the value of labor income. For a credit constrained worker, even a highly indebted one, labor income could easily be a more important consideration. For instance, consider the extreme case of an impatient credit-constrained worker who is up against a borrowing constraint.⁵ Her budget constraint is given by

$$p_t c_t + R_t = w_t,$$

where p_t is the price of consumption c_t , R_t is debt repayment, and w_t is income. Then consumption is given by

$$c_t = \frac{w_t - R_t}{p_t}.$$

In this case, the gains from lower debt repayments are less significant than the losses from lost labor income. So, if a worker has little debt or is highly credit constrained, then inflation is, holding wages and debt obligations fixed, a net negative.

Ultimately, the question of whether or not workers dislike inflation is an empirical one, and here the survey evidence is overwhelming. As an example, according to the Pew Research Center's survey of 39 countries in May 2013, 94% of Greeks, 84% of Italians, and 69% of Spaniards complain that inflation poses a very big challenge. In the 39 country sample, a median of 33% of respondents wanted their government to focus on rising prices as its top priority.⁶

3.2 Implications for Quit Rates

Another counterintuitive implication of the model presented above may seem to be that households willfully choose to withdraw their labor because wages are not high enough. This suggests an implausible scenario where workers quit their jobs during recessions. However, this is a misinter-

⁵Consider a worker with discount factor β , labor income w_t , facing exogenous real interest rate $r < 1/\beta$, with an exogenous borrowing constraint. In the stationary equilibrium, the worker will borrow up to the borrowing constraint, and use all of her income to roll over the debt and to consume.

⁶<http://www.pewglobal.org/files/2013/05/Pew-Global-Attitudes-Economic-Report-FINAL-May-23-20131.pdf>

pretation of the model. The decision to cut back hiring, in the model, is made by the firm *not* the household. The firm is forced to offer higher wages than efficient due to the information friction, and so it compensates for the higher wages by reducing hirings. I have modelled this by giving the workers an outside option, but it would be easy to augment this with a moral hazard problem. This would not change the model’s predictions, but instead of paying workers their outside option, the firm would need to pay the workers a high enough wage to prevent them from shirking. In the presence of disinflationary pressure, the firm would be forced to lay off workers instead of cutting wages, because households could not commit to not shirk.

4 Empirical Evidence

In this section, I examine this stylized model’s fit with the the data. I present two types of evidence: (1) direct evidence of asymmetry in household beliefs towards inflation; (2) time-series evidence from the United States relating wage and price inflation to monetary shocks. I also discuss how well the model can explain the cross-sectional distribution of wage changes in different countries and different time periods.

4.1 Evidence on Asymmetric Expectations

The basic mechanism of the model is that households place greater weight on inflationary news than disinflationary news. We can try to test for this mechanism directly by using inflation expectation surveys of households. Technically, the expectations of the agents in the model are not unique, since they have multiple priors. So, I assume that individuals report their “worst-case” or “effective” beliefs in surveys – these are the beliefs that would rationalize their behavior if they were Bayesians.

I use individual-level data from the Michigan Survey of Consumers. This data is described in more detail in appendix I. Consider the following reduced-model for household i ’s forecast of 12-month ahead inflation in period t :

$$\hat{\pi}_{t+12|i,t} = A \cdot \text{past info} + B^+ \cdot \text{new inflationary info} + B^- \cdot \text{new disinflationary info} + C_{it},$$

where A , B^+ and B^- are regression coefficients, and C_{it} are additional controls. As discussed by Lancaster (2000), B^+ and B^- can be consistently estimated by least squares.

We can motivate (4.1) as a linear filtering problem where expectations today are some weighted average of past information and new information. My measure of *new information* is the contemporaneous median 12-month ahead inflation forecast from the survey of professional forecasters (SPF) conducted by the Federal Reserve Bank of Philadelphia, and denoted by $\pi_{t+12|t}^e$. The SPF is conducted quarterly starting from the fourth quarter of 1981. My measure of *Past information* is the lagged median SPF forecast, as well as lagged inflation. New information is considered *inflationary* if it is greater than last period’s forecasted inflation rate $\pi_{t+12|t}^e \geq \pi_{t+8|t-4}^e$, otherwise it

is disinflationary. Denote new inflationary and new disinflationary news by

$$expert_t^+ = (\pi_{t+12|t}^e - \pi_{t+8|t-4}^e) \mathbf{1}(\pi_{t+12|t}^e \geq \pi_{t+8|t-4}^e), \quad expert_t^- = (\pi_{t+12|t}^e - \pi_{t+8|t-4}^e) \mathbf{1}(\pi_{t+12|t}^e < \pi_{t+8|t-4}^e).$$

Finally, the control variables C_{it} include individual, and quarter fixed effects. I test the null hypothesis that $B^+ = B^-$. The results of this regression are reported in table 1. In all specifications, we can easily reject the null hypothesis of symmetry at the 5% significance level.

Coibion and Gorodnichenko (2015) emphasize that oil prices play an important role in determining household inflation expectations. One may worry that asymmetric shocks to oil prices, which are more salient to households than other prices, may be driving the observed asymmetry in the regressions. To check for this possibility, I include inflation in crude petroleum prices as an additional control in table 2. This does not change the results of the regression appreciably. However, oil prices themselves are highly significant drivers of household inflation expectations as shown by Coibion and Gorodnichenko (2015). Furthermore, consistent with the thesis of this paper, household inflation expectations react highly asymmetrically to changes in oil prices themselves. In particular, increases and accelerations of increases in oil prices increase inflation expectations by much more than decreases and decelerations in oil prices do.

Malmendier and Nagel (2015) show that household inflation expectations can be affected by their age. In table 3, I report the results for subsamples of the population. The results are fairly heterogeneous, but for every subsample except one (gender, income, education, age, birth-year, region, and homeownership status), the point estimate of inflationary news is at least twice the size of the point estimate for disinflationary news. The only case for which this is not true are for those born before 1937. We may speculate that this might be due to the fact that these individuals experienced the great depression, and therefore fear deflation more than inflation. However, even these individuals respond asymmetrically to oil price shocks. Interestingly, there is no strong relationship between education and asymmetry. Highly educated individuals respond more vigorously to inflationary news, but not to disinflationary news. This favors the ambiguity-aversion interpretation of the data, which casts the asymmetry as the optimal response to ambiguous signals, and not the result of a behavioral bias or educational deficiency.

4.1.1 Time-varying Estimates of Asymmetries and the Volcker Disinflation

In figure 5, I plot the median household inflation forecast and the inflation rate as time series. The pattern I argue for in this paper can be seen with the naked eye: look for positive and negative spikes in the headline inflation rate, and then see if there is a positive or negative spike in household expectations. For instance, the positive spike in mid-2008 is picked up by households but the negative spike in mid-2009 is not; the positive spike in early 2011 is picked up but the decline in 2014 is not; the positive spikes in 1990 are picked up but the negative ones in 1986 are not. The one part of the sample where this pattern does not hold is 1978 to 1982, where households

Table 1: Responsiveness of Household Inflation Expectations to Positive and Negative Shocks in the United States

	(1)	(2)	(3)
	$\hat{\pi}_{t+12 t}$	$\hat{\pi}_{t+12 t}$	$\hat{\pi}_{t+12 t}$
$expert^+$	0.711*** (0.13)	0.538*** (0.14)	0.553*** (0.14)
$expert^-$	0.180 (0.11)	0.032 (0.11)	0.057 (0.12)
$\pi_{t+8 t-4}^e$	0.867*** (0.07)	0.576*** (0.08)	0.632*** (0.09)
π_{t-1}		0.163*** (0.02)	0.152*** (0.02)
Quarter FE	No	No	Yes
Individual FE	Yes	Yes	Yes
Constant	Yes	Yes	Yes
Observations	126,659	126,659	126,659

Standard errors clustered at the individual level in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Columns regress median 12-months ahead inflation expectations of households on positive and negative changes to the inflation rate forecast by the median professional forecaster and other covariates. The inflation expectation data for households comes from the Michigan Survey of Consumers, for experts from the survey of professional forecasters, and the inflation data comes from the BLS. The question households are responding to in the Michigan survey is “During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now? By what percent do you expect prices to go up, on the average, during the next 12 months?”. All columns have individual fixed effects. Column (2) also controls for the lagged inflation rate, and Column (3) controls for the lagged inflation rate and the lagged median household inflation forecast. For column (1) and (2) the p-value that the coefficient on $expert^+$ and $expert^-$ are the same is less than 1%, and for column (3) it is 1.3%. Hats indicate forecasts, subscripts indicate time periods. The sample period is from the fourth quarter of 1981 to the first quarter of 2015, sampled quarterly.

Table 2: Responsiveness of Household Inflation Expectations to Positive and Negative Inflation Forecast and Oil Price Shocks in the United States

	(1)	(2)	(3)	(4)
	$\hat{\pi}_{t+12 t}$	$\hat{\pi}_{t+12 t}$	$\hat{\pi}_{t+12 t}$	$\hat{\pi}_{t+12 t}$
$expert^+$	0.354*** (0.07)	0.354*** (0.07)	0.361*** (0.07)	0.369*** (0.07)
$expert^-$	0.028 (0.06)	0.028 (0.06)	0.026 (0.06)	0.027 (0.06)
$\pi_{t+8 t-4}^e$	0.338*** (0.04)	0.338*** (0.04)	0.345*** (0.04)	0.344*** (0.04)
$\Delta\pi_t^{oil}$	0.181*** (0.06)	0.181*** (0.06)	0.188*** (0.06)	
π_t^{oil}	0.365*** (0.05)	0.365*** (0.05)		
π_{t-4}	0.055*** (0.01)	0.055*** (0.01)	0.053*** (0.01)	0.054*** (0.01)
$\pi_t^{oil} \geq 0$			0.448*** (0.07)	0.387*** (0.08)
$\pi_t^{oil} \leq 0$			0.223** (0.10)	0.248** (0.10)
$\Delta\pi_t^{oil} \geq 0$				0.332*** (0.11)
$\Delta\pi_t^{oil} \leq 0$				0.003 (0.13)
Quarter FE	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes
Constant	Yes	Yes	Yes	Yes
Observations	144, 577	144, 577	144, 577	144, 577

Notes: Columns regress median 12-months ahead inflation expectations of households on positive and negative changes to the inflation rate forecast by the median professional forecaster and other covariates. The inflation expectation data for households comes from the Michigan Survey of Consumers, for experts from the survey of professional forecasters, and the inflation data comes from the BLS. The question households are responding to in the Michigan survey is “During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now? By what percent do you expect prices to go up, on the average, during the next 12 months?”. All columns have individual fixed effects. Column (2) also controls for the lagged inflation rate, and Column (3) controls for the lagged inflation rate and the lagged median household inflation forecast. For column (1) and (2) the p-value that the coefficient on $expert^+$ and $expert^-$ are the same is less than 1%, and for column (3) it is 1.3%. Hats indicate forecasts, subscripts indicate time periods. The sample period is from the fourth quarter of 1981 to the first quarter of 2015, sampled quarterly.

Table 3: Responsiveness of Household Inflation Expectations to Positive and Negative Shocks

	$expert^+$	$expert^-$	p-value	$\pi^{Oil} > 0$	$\pi^{Oil} < 0$	$\Delta\pi^{Oil} > 0$	$\Delta\pi^{Oil} < 0$
Homeowner	0.559 (0.13)	-0.312 (0.11)	0.000	0.350 (0.12)	0.144 (0.15)	0.126 (0.15)	0.051 (0.19)
Renter or NA	0.276 (0.11)	0.096 (0.08)	0.204	0.296 (0.14)	-0.180 (0.15)	0.434 (0.18)	0.018 (0.21)
Low Income	0.005 (0.21)	-0.453 (0.17)	0.110	0.540 (0.21)	0.092 (0.27)	0.171 (0.32)	0.428 (0.38)
Middle Income	0.454 (0.16)	0.133 (0.13)	0.134	0.170 (0.16)	0.210 (0.21)	0.583 (0.23)	-0.389 (0.27)
High Income	0.191 (0.11)	0.008 (0.10)	0.245	0.473 (0.12)	0.031 (0.15)	0.324 (0.17)	0.205 (0.20)
Male	0.439 (0.10)	0.113 (0.08)	0.015	0.474 (0.10)	0.024 (0.13)	0.449 (0.14)	0.009 (0.17)
Female	0.143 (0.11)	-0.215 (0.09)	0.021	0.314 (0.12)	0.050 (0.15)	0.161 (0.17)	0.160 (0.20)
West	0.189 (0.16)	-0.269 (0.14)	0.043	0.303 (0.17)	0.066 (0.22)	0.266 (0.23)	0.068 (0.29)
North Central	0.351 (0.14)	-0.048 (0.11)	0.036	0.282 (0.14)	0.025 (0.18)	0.353 (0.20)	-0.133 (0.24)
North East	0.369 (0.17)	0.084 (0.14)	0.211	0.329 (0.17)	0.168 (0.23)	0.157 (0.26)	0.387 (0.30)
South	0.284 (0.14)	0.008 (0.11)	0.148	0.539 (0.14)	-0.016 (0.17)	0.352 (0.19)	0.044 (0.24)
Low Education	0.122 (0.14)	-0.081 (0.11)	0.274	0.172 (0.14)	-0.023 (0.18)	0.096 (0.21)	-0.361 (0.24)
Middle Education	0.175 (0.16)	-0.080 (0.13)	0.246	0.672 (0.18)	-0.127 (0.21)	0.120 (0.24)	0.367 (0.29)
High Education	0.484 (0.11)	0.024 (0.09)	0.003	0.494 (0.11)	0.123 (0.15)	0.471 (0.16)	0.401 (0.20)
Less than 20 years	0.753 (0.50)	-0.289 (0.38)	0.117	0.673 (0.48)	-1.294 (0.61)	-0.314 (0.74)	0.636 (0.81)
20 to 40 years old	0.283 (0.12)	-0.058 (0.09)	0.035	0.336 (0.13)	0.034 (0.16)	0.359 (0.18)	0.004 (0.22)
40 to 60 years old	0.324 (0.13)	-0.040 (0.11)	0.038	0.335 (0.12)	0.111 (0.16)	0.364 (0.18)	0.154 (0.21)
More than 60 years old	0.272 (0.16)	0.054 (0.14)	0.351	0.410 (0.17)	-0.023 (0.21)	0.173 (0.25)	0.059 (0.29)
Born before 1937	0.114 (0.16)	0.346 (0.12)	0.280	0.179 (0.16)	-0.036 (0.22)	0.270 (0.25)	-0.256 (0.29)
Born between 1937-1960	0.345 (0.10)	-0.144 (0.08)	0.000	0.470 (0.11)	0.142 (0.13)	0.279 (0.15)	0.111 (0.18)
Born after 1960	0.324 (0.16)	-0.327 (0.14)	0.003	0.413 (0.15)	-0.131 (0.19)	0.405 (0.20)	0.199 (0.25)

Notes: Columns report the coefficient on positive and negative changes to that variable for household beliefs. Each row corresponds to a subsample of the population with standard errors in brackets. The third column is the p-value for a t-test that the coefficients on $expert^+$ and $expert^-$ are equal. Standard errors are clustered at the individual level.

are, if anything, underreacting to inflation. Unfortunately, the data is very limited pre-1980. The regressions in table 1 make use of the entire available sample, but inflation forecasts from the SPF only go back to 1981, and the Michigan Survey only goes back to 1978, so we cannot study the period before the Volcker disinflation.




Figure 5: Median expected price change for the next 12 months from the Michigan Survey, and headline annual inflation rate in that period.

In figure 6, I replot figure 2 but include the Volcker disinflation from 1978-1982, and the strong asymmetry we observed disappears. One rationalization for this fact is that the period from 1978-1982 was a period of structural change, with great media attention on Paul Volcker’s attempt to disinflate the US economy. If these extra signals in favor of disinflation swayed household attitudes, then the asymmetry would disappear in this part of the sample.

Ultimately, it is difficult to say exactly what happened during these four years without more data. In constructing the basic motivating figures 2 and 3, I assume that the the inflation rate is stationary, since I take $\pi_{t+12} \geq \pi_t$ as an inflationary news shock and $\pi_{t+12} < \pi_t$ as a disinflationary news shock. In other words, I assume that households expect the inflation rate to be constant before observing new signals. If this assumption was violated during the Volcker disinflation, as it should be during a structural break, then there is no reason to expect a kink at zero according to the theory. I estimate regression (4.1) for different periods: the early 80s, the mid-80s to the late 90s, and the early 2000s to the end of the sample, and display the results in table 6. The latter two

time periods correspond to the first and second halves of the sample period. Since identification is achieved via time-series variation in the aggregate inflation rate, one has to be careful about splitting the sample. Nevertheless, other than the period in 1981-1983, the estimates of B^+ are significantly greater than the estimates of B^- .




Figure 6: Forecast revisions of the annual inflation rate by the median household in the Michigan Survey of Inflation Expectations from 1978-2015, plotted against realized changes in the annual inflation rate as measured by the CPI.

Since the supporting data is from after the Great moderation, one may question the extent to which such asymmetries can persist in countries with high (but stable) inflation rates. Using household expectations data from Argentina, I verify that that higher average inflation does not appear to affect the existence of the asymmetry. To this end, I run the same regression with data from Argentina and present the results in table 5. Unfortunately, in this case, I only have access to the median inflation response from the household survey, rather than individual responses. However, we can still run the same regression on the median response. As predicted by the theory, the point estimate for B^+ is much larger than for B^- . Since we have many fewer observations, the parameters are imprecisely estimated, and we cannot reject the hypothesis that the coefficients are the same. The inflation data used here are from a private consulting firm, taken from Cavallo et al. (2015), and are not official figures from the government (which are widely known to be unreliable).

4.2 Evidence from aggregate time-series

Next, we look at time series evidence of the relationship between wage inflation, price inflation, and monetary policy shocks. The model implies that wage inflation and output should respond

Table 5: Responsiveness of Household Inflation Expectations to Positive and Negative Shocks in Argentina

	(1)	(2)
	$\hat{\pi}_{t t}$	$\hat{\pi}_{t t}$
$\hat{\pi}_{t-1 t-1}$	0.778*** (0.07)	0.911*** (0.04)
π_{t-1}	0.195* (0.08)	
$(\pi_t - \pi_{t-1})\mathbf{1}(\pi_t - \pi_{t-1} \geq 0)$	0.909* (0.37)	1.015** (0.38)
$(\pi_t - \pi_{t-1})\mathbf{1}(\pi_t - \pi_{t-1} < 0)$	0.365 (0.62)	0.185 (0.60)
Constant	0.910 (1.01)	1.843 (1.08)
Observations	79	79

t statistics in parentheses

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

Notes: Columns regress median contemporaneous inflation expectations of households on positive and negative changes to the inflation rate as measured by a private consulting company. The inflation expectation data and the inflation data were kindly shared by Cavallo et al. (2015). Both columns control for the lagged expected inflation rate. Column (2) also controls for the lagged inflation rate. The hypothesis that the coefficients for positive and negative changes are equal in magnitude cannot be rejected. Hats indicate forecasts, subscripts indicate time periods. The sample period is from August 2006 to March 2013. Observations are at the month level.

asymmetrically to expansionary and contractionary monetary shocks. There is already evidence that monetary policy has asymmetric effects on output, for example in De Long and Summers (1988), Cover (1992), Angrist et al. (2013), and Santoro et al. (2014). But this model has novel predictions about asymmetries in the response of wage inflation to monetary shocks. To test these predictions, I use a measure of structural monetary shocks from Romer and Romer (2004). The sample period runs from March 1969 to December 1996. In order to identify the impact of these shocks to price and wage inflation, I use the local projection method from Jordà (2005). This easily accommodates non-linearities in the estimating equations and has the advantage of not imposing dynamic restrictions implicitly embedded in standard VARs.

I estimate the following reduced form model

$$\pi_{t+h}^w = \alpha_0^h + \sum_{j=1}^J \alpha_j^h \pi_{t-j}^w + \beta_h^+ \varepsilon_t^+ + \beta_h^- \varepsilon_t^- + \nu_t,$$

where π_{t+h}^w is monthly wage inflation h periods ahead, ε_t^+ and ε_t^- are positive and negative monetary shocks, and ν_t is the error term. I use three different measures of wage inflation: the change in the log of non-seasonally adjusted average weekly earnings, hourly earnings, and unit labor costs.⁷ For the first two, to account for productivity growth in wages, I use detrended earnings numbers, though this does not affect the results. I estimate $H + 1$ equations for horizons $0, \dots, H$ by estimating a system of seemingly unrelated regressions. I then use the robust standard errors for panel regressions with crosssectional dependence from Driscoll and Kraay (1998).

We are interested in the hypothesis that

$$\sum_k \beta_k^+ = \sum_k \beta_k^-,$$

or that the cumulative effect of a positive shock on wage inflation is the same as the cumulative effect of a negative shock. I use the BIC to select the autoregressive lag length J , although the results are robust to changing the number of lags to be higher (for example, the results are virtually unaffected by using 12 autoregressive lags). The cumulative impact of negative and positive shocks are plotted in figure 7.⁸ For all three measures the point estimates for the expansionary (negative shocks) have a larger effect on wage inflation than contractionary (positive shocks). Although, these differences are not always significant and, as is common with linear projections, the confidence intervals are very wide. The prediction that these shocks should have asymmetric effects on wage inflation is a

⁷Unit labor cost is only available at quarterly frequency, so for the regressions using unit labor costs, the other variables are aggregated to be quarterly.

⁸The impulse responses plot $\sum_k \beta_k^+$ and $\sum_k \beta_k^-$ respectively. Standard monetary models imply that both of these should be negative, since an increase in the interest rate should decrease the wage level. However, the model in this paper implies that we should expect the expansionary effect $\sum_k \beta_k^-$ to be bigger in magnitude than the contractionary effect $\sum_k \beta_k^+$.

purely nominal implication of this model that is not generated by alternative theories of asymmetric business cycles like the ones driven by financial frictions that only bind in recessions. Additionally, the impulse response functions suggest that the impact of expansionary monetary shocks are felt more quickly than the effects of contractionary shocks – this is consistent with the dynamic model presented in appendix III, where the shocks not only have different magnitudes, but also differ in their persistence, as contractionary shocks take longer to be fully incorporated into beliefs.

As a further check on these results, I conduct a similar test by replacing wage inflation with monthly CPI and PCE inflation and plot the results in figure 8. The symmetric results are a replication Romer and Romer (2004), with the exception that I use linear projections rather than single-equation OLS. The results on price inflation mirror the findings of Tenreyro and Thwaites (2015): the evidence for asymmetric responses to price inflation are, if anything, that contractionary shocks have a bigger impact on prices than expansionary shocks. In line with Schweitzer and Hess (2000), these results cast doubt on models that imply a very tight short-run connection between the evolution of wages and prices. Since prices in this model are set by a frictionless price equals marginal cost condition, this suggests there are other important frictions in price-setting.

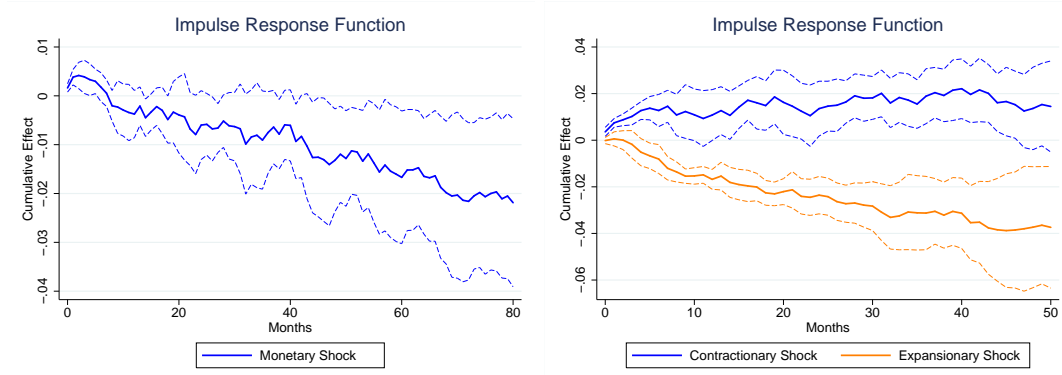
4.3 Cross-sectional distribution of wage changes

Next, we look at the cross-sectional distribution of wages in different countries and time periods. Dickens et al. (2007) demonstrate that the cross-sectional distribution of wage changes is both skewed towards the right, and exhibits bunching at zero.


Dickens et al. (2007) define two forms of wage rigidity. The first, which they call “nominal rigidity”, is a large point mass at zero wage change, the second, which they call “real rigidity”, is an asymmetric distribution of wage changes around the average inflation rate. The type of rigidity predicted by this model is the latter, rather than the former kind of rigidity. An illustration of this difference can be seen in the cross-sectional distributions reported by Dickens et al. (2007). For countries where the inflation rate is relatively high and stable (like the UK in the 1980s), there is not a large point mass at zero wage changes, but the distribution of wage changes is asymmetric around the inflation rate. On the other hand, in countries where the inflation rates are lower and less stable, there is a large point mass at zero. This is also consistent with the histograms in Elsby et al. (2015) where there is an asymmetry around the inflation rate, but there is still a nontrivial fraction of the population who experience nominal wage cuts and the spike at zero is relatively small.

The model presented in section 2 has a degenerate cross-sectional distribution, but it can easily be extended to have cross-sectional heterogeneity. Following Lucas (1973), consider a continuum of islands indexed by elements of the $[0, 1]$ interval, with each island inhabited by a worker and an employer. Worker-employer pair i observe a noisy island specific signal $s_i = \varepsilon + \varepsilon_i$ of the log price level $p = \mu + \varepsilon$, and then write a wage contract. Assume that the worker considers the signal-to-

(a) Weekly Earnings



(b) Unit Labor Costs



(c) Hourly Earnings




Figure 7: Cumulative impact of a negative and positive shock to interest rates on the wage level with 90% Driscoll and Kraay (1998) confidence intervals. The sample is from 1969 to 1996.

noise ratio to be ambiguous. As before, when worker utility is log, the prevailing log wage w in island i is

$$w_i = \mu + \frac{\sigma^2}{\sigma^2 + \sigma_s^2(s_i)} s_i + d.$$

where $\sigma_s(s_i) = \bar{\sigma}_s$ whenever $s_i \leq 0$, and $\underline{\sigma}_s$ otherwise. This means that the cross-sectional distribu-




Figure 8: Cumulative impact of a negative and positive shock to interest rates on the price level with 90% Driscoll and Kraay (1998) confidence intervals.

tion of wages will be discontinuous around the expected price level μ , with higher variance on the right-hand side and bunching on the left-hand side. If the aggregate shock ε is sufficiently large, the discontinuity and bunching disappear. That is, when the unexpected monetary shock is large (high surprise inflation), no asymmetry is observed in the cross-sectional distribution, see figure 9.

Crucially, the model implies that the key point of asymmetry is the ex-ante expected price level (the price level before the signal is observed). This means that, in this model, there is nothing special about zero per se, unless households believe that they are in a very low inflation environment and that absent any signal, prices are not going to change. Conversely, in environments with high and stable inflation, we would observe an asymmetry not at zero, but around expected inflation. Where the asymmetry appears will depend on household expectations in the absence of any new information. It is difficult to infer this from the data available, but the empirical evidence presented in Dickens et al. (2007) is consistent with the idea that the location of asymmetry is higher than zero in environments with persistently high inflation.

A further finding by Dickens et al. (2007), as well as Holden and Wulfsberg (2009), is that the




Figure 9: Theoretical cross-sectional wage-change distributions for different aggregate monetary shocks.

degree of real wage rigidity is strongly correlated with union density. Again, the basic model in section 2 can naturally be extended to account for this finding. In particular, note that the degree of rigidity does not solely depend on the degree of ambiguity, but also on the relative elasticities of labor supply and labor demand – or, loosely speaking, the bargaining power of workers and employers. The kinked beliefs of the workers only affect their wage to the extent that workers can withdraw labor in response to their perceived real wage. In particular, the model in section 2 assumed that the workers’ outside option is exogenous and constant. This makes labor supply completely elastic, since workers effectively make an ultimatum to the the employers and refuse to supply any labor when wages are lower than what they demand. Consider, instead the following log labor supply curve

$$l = \gamma(w - \tilde{E}(p|s) - d),$$

where l is log labor and the elasticity of labor supply is given by γ . When γ tends to infinity, we recover the previous set up. On the other hand, when $\gamma = 0$, labor supply is completely inelastic at fixed supply. Let log labor demand be given by

$$l = \frac{1}{1 - \alpha}(\log(\alpha) + E(p|s) - w),$$

easily derived from profit maximization with a Cobb-Douglas production function. The equilibrium log wage is given by

$$w = \kappa(\log(\alpha) + E(p|s)) + (1 - \kappa) \left(\tilde{E}(p|s) + d \right),$$

where $\kappa = \frac{1}{\gamma(1-\alpha)+1} \in [0, 1]$.

The equilibrium nominal wage is a convex combination of the beliefs of the workers and the beliefs of the employers. In the extreme case of infinitely elastic labor supply, $\gamma = \infty$, only the beliefs of the workers matters and we get maximum rigidity. In the other extreme of completely inelastic labor supply, $\gamma = 0$, only the beliefs of the employer matter. In particular, in the completely inelastic labor supply case, ambiguity has no effect on equilibrium wages. This is intuitive: if workers cannot withdraw their labor in response to the wage offer, then wages are determined solely through competition between employers. If employers are not adversely affected by inflation, then the equilibrium wage will not exhibit a discontinuity around the expected inflation rate even if employers are ambiguity-averse.

So, the degree of bargaining power, captured here by the workers' ability to withdraw labor when wages fall, affects the degree of rigidity in wages. This is consistent with the empirical findings of Dickens et al. (2007) who find that countries where unions have more power exhibit more wage rigidity of the kind generated by this model (i.e. asymmetry around the expected inflation rate).

4.4 Is Ambiguity-Aversion Necessary?

While ambiguity aversion is consistent with the evidence I provided, it is not the only theory that could generate these predictions. In particular, any theory that delivers convex conditional expectations will generate similar predictions. The simplest alternative theory is a Bayesian expected-utility maximizer. Since professional forecasters exhibit no convexity in their beliefs, there must then be something unusual about the priors of the households. In particular, priors that assume the signal is more accurate when there is accelerations of the inflation rate than decelerations will generate very similar results to mine. However, these priors will be incorrect and hard to justify intuitively. Furthermore, the professional forecasters will be objectively doing a better job.

The virtue of ambiguity-aversion, other than its tractability, is that that there is no sense in which households are acting irrationally or making a mistake. If households do not know the precise mapping between aggregate statistics and the prices relevant for them, it is reasonable for them to rely on robustness heuristics – indeed, this was the original motivation for axiomatic theories of ambiguity-aversion.

5 Monetary Policy Problem

Received wisdom in Keynesian economics is that that downward wage rigidity implies that the central bank should have an inflationary bias. Inflation is said to “grease the wheels” of the labor market since it allows wage cuts to take place that would otherwise not have occurred. In this section, I show that this intuition holds in my model if we take the conditional expectation function of the households as exogenous, but fails if we account for the fact that household expectations will

react to the change in policy.

Consider a scenario where the central bank has some, but not complete, control over the distribution of demand shocks that hit the economy. Crucially, suppose that although the central bank can affect the distribution of shocks, it has no control over the distribution of the public signal. In other words, the central bank chooses a distribution of demand shocks to minimize expected losses, taking as given the information content of a noisy public signal (national statistics and professional forecasters).

Most central banks are tasked with maintaining price stability and full employment. In a model like the one in section 3, with only aggregate demand shocks, price stability and deviations from first-best employment are both log-linear functions of the level of the monetary surprise. In particular, if we denote first-best employment by l^{fb} , then

$$l - l^{fb} = \frac{1}{1 - \alpha}(m - \tilde{E}(m|s)),$$

where first-best employment is employment in the perfect information world. On the other hand, $m - \tilde{E}(m|s)$ also captures price instability. So, I assume that the central bank's loss function is given by

$$L(g) := E_g \left(\left(m - \tilde{E}(m|s) \right)^2 \right),$$

where g is the marginal distribution of demand shocks m , and the expectation is taken with respect to g . Note that the central bank's objective function treats employment and unemployment symmetrically, as is standard in the optimal monetary policy literature. This objective is not the welfare relevant objective function of the households in this model, and so the goal is not to investigate optimal policy in this model, but rather to examine the case for inflationary bias using the standard objective function.

Assume that if the central bank takes no action, demand shocks will have a reference distribution q . The central bank chooses the distribution of demand shocks g to minimize losses subject to the requirement that g is not too different from the reference distribution q . I formalize "not too different" using a constraint on the relative entropy between g and q

$$E_q(\log(g)) - E_q(\log(q)) \leq K,$$

where E_q denotes the expectation taken with respect to the density q . Relative entropy is an analytically tractable measure of difference between probability distributions.

As discussed by Hansen and Sargent (2008), there is a natural interpretation for relative entropy constraints. In this context, we can interpret q as the distribution of demand shocks set by nature, and the relative entropy constraint represents the set of alternative models the household cannot distinguish between. In other words, the central bank has some ability to trick agents and to alter

the underlying distribution of demand shocks by changing the money supply. The central bank's ability to do this is constrained by the fact that for large enough changes, the agents will detect the changes in the money supply, and adjust their expectations and prices to take the change into account. So extreme monetary policy changes are detected by the agents, and are therefore neutral.

The analysis is done in two parts. First, I consider the central bank's optimal policy when household conditional expectations are treated as exogenous. In other words, the central bank takes the existence of a convex conditional expectation function as given, and chooses the distribution which minimizes loss. This is where the analysis would stop if we simply took the empirical results of section 4 at face value without a structural model of belief-formation. For this case, we reach a very intuitive policy prescription: the central bank should skew the distribution of demand shocks, and make expansionary shocks much more frequent than contractionary shocks. The intuition is that since households make larger mistakes in response to contractions than expansions, the central bank should minimize the occurrence of contractions.

In the second part of this section, I use endogenize belief formation, and let the central bank take into account the fact that by changing the distribution of demand shocks, it also changes the households' signal extraction problem. This is the advantage of having a structural model for belief formation, and having this structure significantly alters the optimal policy. With endogenous beliefs, the inflationary bias of the central bank disappears. If the central bank makes inflationary shocks much more common, then household conditional expectations account for this change, and the fact that they account for this change diminishes the central bank's motives for causing this change in the first place. The result is that there can be significant asymmetries in equilibrium, but the central bank is powerless to do anything about it.

5.1 Exogenous Belief-Formation

In this section, I consider the problem of a central bank who takes the function mapping signals to beliefs of the household as given, and does not internalize the fact that changing the distribution of demand shocks will affect how signals are mapped to conditional expectations. In other words, the central bank solves the following problem

$$\min_{g(m)} \iint (m - \phi(s))^2 f(s|m)g(m)dsdm$$

such that

$$\int q(m) \log(g(m))dm - \int q(m) \log(q(m))dm \leq K$$

$$\int g(m)dm = 1,$$

where q is the distribution of demand shocks when the central bank is passive, $f(s|m)$ is the density of the signal conditional on the monetary policy shock, and $\phi(s)$ is household expectations of the demand shock conditional on the signal, which the bank takes as exogenous. The first constraint requires that the distribution of demand shocks the bank chooses be sufficiently close to the reference distribution q in relative entropy terms. The second constraint ensures that the chosen density implies a valid probability distribution. The slack non-negativity constraints have been suppressed since they are implied by the first constraint.

The Lagrangian is given by

$$\min_{g(m)} \iint (m - \phi(s))^2 f(s|m)g(m)dsdm - \lambda \int q(m) \log(g(m))dm - \mu \int g(m)dm.$$

The first order condition is given by

$$\begin{aligned} \frac{d}{dt} \iint (m - \phi(s))^2 f(s|m)(g(m) + th(m))dsdm \\ - \lambda \int q(m) \log(g(m) + th(m))dm - \mu \int g(m) + th(m)dm \Big|_{t=0} = 0, \quad \forall h \end{aligned}$$

At the optimum, $g(m)$ solves the following equation

$$\int (m - \phi(s))^2 f(s|m) - \lambda \frac{q(m)}{g(m)} - \mu = 0.$$

Rearrange this to get

$$g(m) = \frac{\lambda q(m)}{\int (m - \tilde{E}(m|s))^2 f(s|m)ds - \mu}, \quad (13)$$

where $\tilde{E}(m|s)$ is substituted for $\phi(s)$.

This first order condition is very intuitive to interpret. Draws of the monetary shock m with large expected squared error in the household's forecast, $E((m - \tilde{E}(m|s))^2|m)$, are less likely to occur relative to the reference distribution q . In other words, if households are more likely to have incorrect beliefs during deflationary episodes than inflationary episodes, then the central bank will reduce the probability of deflationary shocks. This is despite the fact that the central bank's loss function treats under- and over-employment symmetrically.

Since we found household beliefs to be more likely to be incorrect after disinflationary periods than inflationary periods, equation (13) suggests that the central bank should maintain an inflationary bias in policy. This is in keeping with the intuition, and the advice, found in papers like Akerlof et al. (1996) or Kim and Ruge-Murcia (2009), that recommend positive steady state inflation in the presence of downwardly sticky wages. However, in the context of this model, this line of reasoning is susceptible to the Lucas critique if the conditional expectations of households

respond endogenously to the distribution of demand shocks.

Before proceeding to the case with endogenous expectations, let us get a better sense for how the solution behaves with the following numerical example. This example shows that, the naive optimal policy will feature an inflationary bias. Suppose that the signal s is given by

$$s = m + \varepsilon,$$

where ε is a mean-zero normally distributed noise term with variance $\sigma \in [\underline{\sigma}, \bar{\sigma}]$. Then

$$\tilde{E}(m|s) = \max_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \int m \frac{f(s|m, \sigma)g(m)}{f(s|\sigma)} dm, \quad (14)$$

where

$$f(s|m, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(s-m)^2}{2\sigma^2}\right), \quad (15)$$

and

$$f(s|\sigma) = \int f(s|m, \sigma)g(m)dm. \quad (16)$$

Equations (13), (14), (15), and (16) determine the equilibrium of this economy.

Let the reference distribution q be a standard normal distribution with mean 0 and variance 1. By calibrating $\underline{\sigma}$, $\bar{\sigma}$, and λ we can compute the equilibrium distribution of monetary shocks. We can calibrate $\underline{\sigma}$ and $\bar{\sigma}$ by fitting a piecewise linear regression to the expectations data. The slope of the piecewise linear function ψ gives the signal-to-total variance ratio, which in turn pins down $\underline{\sigma}$ and $\bar{\sigma}$. A good estimate seems to be $\underline{\sigma} = 0.7$ and $\bar{\sigma} = 2$. Calibrating λ is harder, so we can plot solutions for a range of λ to get a sense of what the optimal solution looks like. In figure 10, we see that as the constraint on the central bank becomes looser, the distribution of shocks becomes more positive and concentrated. For comparison, figure 11 shows that without ambiguity, the distribution simply becomes more concentrated, but there is no inflationary bias. This lines up with the received wisdom that central banks should have an inflationary bias because of downward wage rigidity.




Figure 10: The marginal distribution of monetary shocks




Figure 11: The marginal distribution of monetary shocks with $\underline{\sigma} = \bar{\sigma} = 1$.

5.2 Endogenous Belief-Formation

The intuitive result in the previous section is in line with other work in recommending inflationary bias in the presence of downward wage rigidity. However, this result depends crucially on the assumption that the function mapping signals to conditional expectations for the households is fixed. If the central bank takes into account the fact that changing the distribution of monetary shocks changes the signal-extraction problem faced by households, then the inflationary bias disappears.

To that end, consider a central bank that faces the following problem:

$$\min_{g(m)} \iint (m - \phi(s))^2 f(s|m)g(m)dsdm$$

such that

$$\begin{aligned} \int q(m) \log(g(m))dm - \int q(m) \log(q(m))dm &\leq K \\ \int f(m)dm &= 1, \\ \phi(s) &= \max_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \int \frac{mf(s|m, \sigma)g(m)}{f(s|\sigma)}dm. \end{aligned}$$

The analytical first order condition for this problem is harder to interpret, since it accounts for a non-linear filtering problem. Instead I plot example solutions using a normal error term and a normal reference distribution in figure 12. Unlike the previous section, we see no inflationary bias in the central bank's optimal response, even though the degree of the asymmetry is very extreme. The reason is that if the central bank attempted to skew the distribution towards more inflationary shocks, conditional expectations of households would take this skew into account when interpreting the signal.

The results of this section do not prove that zero percent inflation is the optimal inflation rate. In fact, in this model, the mean value of the inflation rate, as long as it is known by all agents, has no effect on welfare, since wages and prices are flexible. In practice, there are other reasons why we might want to implement a positive inflation target, ranging from concerns about hitting the zero lower bound to other causes of downward wage rigidity besides the one studied here (for instance a nominal fairness norm).




Figure 12: The optimal distribution of monetary shocks, with $\underline{\sigma} = 1$, $\bar{\sigma} = 10$.




Figure 13: The expectation of the price level conditional on the signal, with $\underline{\sigma} = 1$, $\bar{\sigma} = 10$.

6 Concluding Remarks

In this paper, I document an asymmetry in how household inflation expectations react to inflationary and disinflationary news. I argue that this asymmetry can explain part of the reason why nominal wages can be downwardly rigid. A simple general equilibrium model with asymmetric beliefs implies that nominal and real variables respond asymmetrically to monetary policy shocks.

In particular, contractionary shocks cause larger changes to output than expansionary shocks. On the other hand, contractionary monetary policy shocks cause smaller changes to wage inflation than expansionary shocks, and these predictions are consistent with time series data from the United States.

I argue that information frictions, coupled with ambiguity-aversion, can be a microfoundation for these asymmetric responses. The intuition is that households pay closer attention to and respond more strongly to bad news that their purchasing power might be lower than they thought than to good news that their purchasing power is higher than they thought.

Normatively, the asymmetry induced by ambiguity aversion increases the welfare costs of business cycles. Since positive and negative shocks do not cancel, reductions in variance reap first-order gains. A back of the envelope calculation, in appendix IV, shows that these costs are around 20 times higher than the ones in Lucas (1987).

Furthermore, contrary to popular wisdom, where it is typically assumed that downward wage rigidity imbues central banks with an inflationary bias, in this model, this intuition fails to hold because household's expectations respond to inflationary pressure from the central bank. In other words, the idea that in a world with downward wage rigidity, positive inflation "greases the wheels" of the labor market may be subject to the Lucas critique for reasons similar to the long-run Phillips curve.

References

- Akerlof, G. A. (1982). Labor contracts as partial gift exchange. *The Quarterly Journal of Economics* 97(4), 543–569.
- Akerlof, G. A., W. T. Dickens, G. L. Perry, R. J. Gordon, and N. G. Mankiw (1996). The macroeconomics of low inflation. *Brookings Papers on Economic Activity* 1996(1), pp. 1–76.
- Angeletos, G.-M. and J. La'O (2009). Incomplete information, higher-order beliefs and price inertia. *Journal of Monetary Economics* 56, S19–S37.
- Angrist, J. D., Ò. Jordà, and G. Kuersteiner (2013). Semiparametric estimates of monetary policy effects: string theory revisited. Technical report, National Bureau of Economic Research.
- Barattieri, A., S. Basu, and P. Gottschalk (2010, June). Some evidence on the importance of sticky wages. Working Paper 16130, National Bureau of Economic Research.
- Bewley, T. F. (1999). *Why wages don't fall during a recession*. Harvard University Press.
- Cavallo, A., G. Cruces, and R. Perez-Truglia (2015). Learning from potentially biased statistics. Technical report.

- Coibion, O. and Y. Gorodnichenko (2015). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics* 7(1), 197–232.
- Cover, J. P. (1992). Asymmetric effects of positive and negative money-supply shocks. *The Quarterly Journal of Economics* 107(4), 1261–1282.
- Daly, M. and B. Hobijn (2013). Downward nominal wage rigidities bend the phillips curve. *FRBSF Working Paper* (8).
- De Long, J. B. and L. H. Summers (1988). How does macroeconomic policy affect output? *Brookings Papers on Economic Activity* 1988(2), pp. 433–494.
- Dickens, W. T., L. Goette, E. L. Groshen, S. Holden, J. Messina, M. E. Schweitzer, J. Turunen, and M. E. Ward (2007, Spring). How wages change: Micro evidence from the international wage flexibility project. *Journal of Economic Perspectives* 21(2), 195–214.
- Driscoll, J. C. and A. C. Kraay (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of economics and statistics* 80(4), 549–560.
- Elsby, M. W. (2009). Evaluating the economic significance of downward nominal wage rigidity. *Journal of Monetary Economics* 56(2), 154–169.
- Elsby, M. W., D. Shin, and G. Solon (2015). Wage adjustment in the great recession and other downturns: Evidence from the united states and great britain.
- Epstein, L. G. and M. Schneider (2008). Ambiguity, information quality, and asset pricing. *The Journal of Finance* 63(1), 197–228.
- Fehr, E. and J.-R. Tyran (2001). Does money illusion matter? *American Economic Review*, 1239–1262.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141 – 153.
- Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. *American economic review*, 50–65.
- Hansen, L. P. and T. J. Sargent (2008). *Robustness*. Princeton university press.
- Holden, S. and F. Wulfsberg (2009). Wage rigidity, institutions, and inflation.
- Holmstrom, B. (1983). Equilibrium long-term labor contracts. *The Quarterly Journal of Economics* 98, pp. 23–54.

- Ilut, C. (2012). Ambiguity aversion: Implications for the uncovered interest rate parity puzzle. *American Economic Journal: Macroeconomics*, 33–65.
- Ilut, C., M. Kehrig, and M. Schneider (2014). Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news. Technical report, National Bureau of Economic Research.
- Ilut, C. L. and M. Schneider (2014). Ambiguous business cycles. *The American Economic Review* 104(8), 2368–2399.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, 161–182.
- Kaur, S. (2014). Nominal wage rigidity in village labor markets. Technical report, National Bureau of Economic Research.
- Kim, J. and F. J. Ruge-Murcia (2009). How much inflation is necessary to grease the wheels? *Journal of Monetary Economics* 56(3), 365 – 377.
- Kuhnen, C. M. (2012). Asymmetric learning from financial information.
- Lancaster, T. (2000). The incidental parameter problem since 1948. *Journal of econometrics* 95(2), 391–413.
- Lucas, Robert E, J. (1973, June). Some international evidence on output-inflation tradeoffs. *American Economic Review* 63(3), 326–34.
- Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.
- Malmendier, U. and S. Nagel (2015). Learning from inflation experiences*. *The Quarterly Journal of Economics*.
- Masolo, R., F. Monti, et al. (2015). Ambiguity, monetary policy and trend inflation. Technical report, Bank of England.
- Romer, C. D. (1986). Is the stabilization of the postwar economy a figment of the data? *The American Economic Review* 76(3), pp. 314–334.
- Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 1055–1084.
- Santoro, E., I. Petrella, D. Pfajfar, and E. Gaffeo (2014). Loss aversion and the asymmetric transmission of monetary policy. *Journal of Monetary Economics* 68, 19–36.
- Schmitt-Grohé, S. and M. Uribe (2011). Pegs and pain. Technical report, National Bureau of Economic Research.

- Schweitzer, M. and G. Hess (2000). Does wage inflation cause price inflation. Policy Discussion Paper 00-01, Federal Reserve Bank of Cleveland.
- Shiller, R. J. (1997). Why do people dislike inflation? In *Reducing Inflation: Motivation and Strategy*, pp. 13–70. University of Chicago Press.
- Tenreyro, S. and G. Thwaites (2015). Pushing on a string: US monetary policy is less powerful in recessions.
- Woodford, M. (2003). Imperfect common knowledge and the effects of monetary policy. *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Strother Phelps*, 25.

Appendix I: Data Description

The Survey of Consumer Attitudes and Behavior, conducted by the Survey Research Center at the University of Michigan, is available on a monthly basis from January 1978. The survey is a rolling cross-sectional panel of households designed to be representative of the US population. After a first interview, respondents are randomly selected to be re-interviewed after six months. So, each month, around 40% of respondents are re-interviewed, while 60% are interviewed for the same time. The questions relevant to this paper are: (1) “During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?” with multiple choice answers “Go up”, “Go up (at same rate)”, “Same”, “Go down”, “DK”, and “NA”; (2) “By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?” with numeric answers. I use answers to the second question for household’s expectations, where the number is multiplied by minus one if the household indicated that they expect prices to go down. To account for the severe outliers in the survey, I dropped any responses that were outside of the 90% interquartile range for that time period.

My measure for oil prices is the quarterly inflation rate in the Crude Petroleum category of the producer price index from the BLS.

Appendix II: Robustness Checks

Table 6: Responsiveness of Household Inflation Expectations to Positive and Negative Shocks in the United States

	(1)	(2)	(3)
	$\hat{\pi}_{t+12 t}\mathbf{1}(1981 \leq t \leq 1983)$	$\hat{\pi}_{t+12 t}\mathbf{1}(1983 < t \leq 1999)$	$\hat{\pi}_{t+12 t}\mathbf{1}(1999 < t)$
$expert^+$	-1.478 (1.00)	0.287*** (0.10)	0.379** (0.16)
$expert^-$	-0.083 (0.38)	0.011 (0.10)	-0.510*** (0.13)
$\pi_{t+8 t-4}^e$	-0.037 (0.25)	0.308*** (0.06)	0.047 (0.13)
π_{t-1}	0.269*** (0.10)	0.093*** (0.03)	0.095*** (0.02)
$\Delta\pi_t^{oil}$	4.107 (2.50)	0.125 (0.08)	0.270** (0.11)
π_{t-1}^{oil}	3.689 (2.40)	0.214*** (0.06)	0.273*** (0.09)
Quarter FE	No	No	Yes
Individual FE	Yes	Yes	Yes
Constant	Yes	Yes	Yes

Standard errors clustered at the individual level in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Columns regress median 12-months ahead inflation expectations of households on positive and negative changes to the inflation rate forecast by the median professional forecaster and other covariates. The inflation expectation data for households comes from the Michigan Survey of Consumers, for experts from the survey of professional forecasters, and the inflation data comes from the BLS. π^{oil} is the quarterly inflation rate in the Crude Petroleum prices from the PPI. All columns have individual fixed effects. The sample period is from the fourth quarter of 1981 to the first quarter of 2015, sampled quarterly. Standard errors are clustered at the individual level.

Appendix III: New Keynesian Model with ambiguous sticky wages

In this section, I embed ambiguous information quality into a standard New Keynesian model with sticky wages. Other than showing that our earlier intuitions survive in this context, I show that in a New Keynesian model time-varying ambiguity is observationally equivalent to a supply or cost-push shock. I also show that ambiguity not only causes the amplitude of positive and negative shocks to be asymmetric, but it can also change their persistence.

In a typical New Keynesian model with sticky wages, households are monopolistically-competitive suppliers of their labor. In such a world, it is no longer obviously the case that inflation is bad news and disinflation is good news, since monopolists care about both the relative price and the quantity of what they sell. However, Masolo et al. (2015) show that trend inflation is worse than trend deflation. Therefore, I impose kinked beliefs on the households. In any case, given the empirical support for kinked expectations, it is interesting to put kinked household beliefs into the work-horse New Keynesian model.

Consider a continuum of households indexed by $i \in [0, 1]$. Household seek to maximize

$$\tilde{E}_t \left(\sum_{t=t_0}^{\infty} \beta^t \frac{C_{it}^{1-\gamma}}{1-\gamma} - \frac{L_{it}^{1+\varphi}}{1+\varphi} \right),$$

where \tilde{E}_t represents expectation with respect to kinked beliefs in period t . There is a representative firm that produces the consumption good Y using technology

$$Y_t = A_t L_t,$$

where A_t is a productivity shock and L_t is a CES aggregate of labor inputs

$$L_t = \left(\int_0^1 L_{it}^{1-\frac{1}{\eta}} di \right)^{\frac{\eta}{\eta-1}}.$$

This implies that labor demand is given by

$$l_{it} - l_t = -\eta(w_{it} - w_t), \tag{17}$$

where lower case variables are in logs, and w_t is the log of the CES wage aggregate. Assume that due to free-entry, the firm makes zero profits in equilibrium, therefore

$$w_t - p_t = a_t. \tag{18}$$

Assume that firms are subject to a cash in advance constraint for the labor they purchase

$$w_t + l_t = \theta_t, \tag{19}$$

where θ_t is a stochastic process representing the money supply. Let

$$\theta_t = \theta_{t-1} + v_t, \quad v_t \sim \mathcal{N}(0, \sigma_v^2).$$

Assume that agents receive a noisy public signal x_t of the stance of monetary at the start of the period

$$x_t = v_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2).$$

As before, households do not know the true signal-to-noise ratio. Suppose agents only know that $\sigma_\varepsilon \in [\sigma_\varepsilon - d, \sigma_\varepsilon + d]$, where $d > 0$ is a parameter that captures the amount of ambiguity or Knightian uncertainty.

The timing of this model will be similar to Angeletos and La'O (2009): at the beginning of the period, an exogenous fraction $1 - \lambda$ of households set their wages optimally subject to their information set and Calvo frictions. At the end of the period, v_t becomes common knowledge, and consumption and production take place. I assume that a monopoly tax eliminates the markup.

Proposition 6.1. *In equilibrium, the wage in period t will be given by*

$$w_t = c_1 w_{t-1} + c_2 \psi_t x_t + c_3 \theta_{t-1} + c_4 a_t + c_5, \tag{20}$$

where $\psi_t = \bar{\psi} = \sigma_v^2 / (\sigma_v^2 + \sigma_\varepsilon^2 - d)$ if $x_t \geq 0$, and $\psi_t = \underline{\psi} = \sigma_v^2 / (\sigma_v^2 + \sigma_\varepsilon^2 + d)$ if $x_t < 0$, and c_1, \dots, c_5 are constants.

We can spot two differences between the reset wage (20) and a standard New Keynesian model with sticky wages. First is the presence of the asymmetric response of real variables to monetary shocks. Second, is the presence of the constant term c_5 . This term is an ambiguity premium, and has the same interpretation as ambiguity premia in asset pricing contexts. Ambiguity-averse households try to insure themselves against monetary shocks in the future by setting higher wages than they otherwise would. This results in a steady state level of real wages that is higher than, and output that is lower than, in the case with no ambiguity. This ambiguity premium has exactly the same implications as a mark-up and it generates a distortion of the steady state. Some impulse response functions can be seen in figures 14 and 15. As expected, negative shocks cause larger changes than positive shocks. The persistence of either type of shock, however, is identical, since after the first period, the shock becomes common knowledge.

Proof of Proposition 6.1. The log-linearized optimal reset wage is given by

$$w_{it} = \tilde{E}_t \left[(1 - \beta\lambda) \sum_{k=0}^{\infty} (\beta\lambda)^k (mrs_{i,t+k|t} + p_{t+k}) \right], \quad (21)$$

where $mrs_{i,t+k|t}$ is log marginal rate of substitution for household i at time $t+k$ conditional on the wage being set in t . Observe that

$$\begin{aligned} mrs_{i,t+k|t} &= \gamma c_{it} + \varphi l_{it}, \\ &= \gamma c_t + \varphi l_{it}, \\ &= \gamma y_t + \varphi (l_t - \eta(w_{it} - w_t)), \\ &= \gamma a_t + (\gamma + \varphi) l_t - \varphi \eta w_{it} + \varphi \eta w_t, \end{aligned} \quad (22)$$

where the second equality follows from complete insurance markets. Use (22), (18), and (19) to get

$$mrs_{i,t+k|t} + p_{t+k} = \underbrace{\frac{\gamma + \varphi}{1 + \varphi \eta}}_{\alpha} \theta_t + \underbrace{\frac{\varphi \eta - \gamma - \varphi + 1}{1 + \varphi \eta}}_{1-\alpha} w_t + \underbrace{\frac{\gamma - 1}{1 + \varphi \eta}}_{\xi_t} a_t. \quad (23)$$

Note that

$$\tilde{E}_t \theta_t = \theta_{t-1} + \psi_t x_t,$$

where $\psi_t = \bar{\psi} = \sigma_v^2 / (\sigma_v^2 + \sigma_\varepsilon^2 - d)$ if $x_t \geq 0$, and $\psi_t = \underline{\psi} = \sigma_v^2 / (\sigma_v^2 + \sigma_\varepsilon^2 + d)$ if $x_t < 0$.

Conjecture an equilibrium where

$$w_{it} = b_1 w_{t-1} + b_2 \psi_t x_t + b_3 \theta_{t-1} + b_4 \xi_t + b_5. \quad (24)$$

Observe that, from aggregation,

$$\begin{aligned} w_t &= \lambda w_{t-1} + (1 - \lambda) w_{it}, \\ &= \lambda w_{t-1} + (1 - \lambda) (b_1 w_{t-1} + b_2 \psi_t x_t + b_3 \theta_{t-1} + b_4 \xi_t + b_5). \end{aligned} \quad (25)$$

Note that we can write (21) recursively

$$\begin{aligned} w_{it} &= (1 - \beta\lambda) \tilde{E}_t (mrs_{it|t} + p_t) + \beta\lambda \tilde{E}_t (w_{i,t+1}). \\ &= (1 - \beta\lambda) \tilde{E}_t (\alpha \theta_t + (1 - \alpha) w_t + \xi_t) + \beta\lambda \tilde{E}_t (w_{i,t+1}), \\ &= (1 - \beta\lambda) \tilde{E}_t (\alpha \theta_t + (1 - \alpha) w_t + \xi_t) + \beta\lambda \tilde{E}_t (b_1 w_t + b_2 \psi_{t+1} x_{t+1} + b_3 \theta_t + b_4 \xi_{t+1} + b_5). \end{aligned} \quad (26)$$

Combine (26) and (25) and, by matching coefficients, derive expressions for b_1, b_2, b_3, b_4 , and b_5 .

This requires noting that

$$\tilde{E}_t \xi_t = \xi_t, \quad E_t(\psi_{t+1}x_{t+1}) = \frac{\sigma_x}{\sqrt{2\pi}}(\bar{\psi} - \underline{\psi}),$$

where σ_x denotes the time-varying variance of x_{t+1} . Matching coefficients gives

$$\begin{aligned} b_1 &= (1 - \alpha)(1 - \beta\lambda)(\lambda + (1 - \lambda)b_1) + \beta\lambda b_1((1 - \lambda)b_1 + \lambda), \\ b_2 &= (1 - \beta\lambda)\alpha + b_2(1 - \alpha)(1 - \beta\lambda)(1 - \lambda) + b_2\beta\lambda b_1(1 - \lambda) + \beta\lambda b_3, \\ b_3 &= (1 - \beta\lambda)\alpha + b_3(1 - \alpha)(1 - \beta\lambda)(1 - \lambda) + b_3\beta\lambda b_1(1 - \lambda) + \beta\lambda b_3, \\ b_4 &= (1 - \beta\lambda) + b_4(1 - \alpha)(1 - \beta\lambda)(1 - \lambda) + b_4\beta\lambda b_1(1 - \lambda), \\ b_5 &= b_5(1 - \alpha)(1 - \beta\lambda)(1 - \lambda) + b_5\beta\lambda b_1(1 - \lambda) + \beta\lambda b_5 + \beta\lambda \frac{\sigma_x(\bar{\psi} - \underline{\psi})}{\sqrt{2\pi}}. \end{aligned}$$

■




Figure 14: The nominal wage and employment as a function of permanent positive shock to money supply.




Figure 15: The nominal wage and employment as a function of permanent negative shock to money supply.

6.1 Time-varying ambiguity

Consider a world where the degree of Knightian uncertainty d is a time-varying quantity, suppose for example that d is a random walk. Then in equilibrium,

$$w_t = c_1 w_{t-1} + c_2 \psi_t x_t + c_3 \theta_{t-1} + c_4 a_t + c_5 (d_t), \quad (27)$$

where b_5 is increasing in d_t . This results in cost-push shocks, which increase wages and reduce output, giving a new microfoundation for the existence of a meaningful policy tradeoff between output and inflation for the central bank and violation of the divine coincidence.

6.2 New Keynesian Model with imperfect information and no Calvo frictions

If we eliminate the Calvo friction, the dynamics of the model become degenerate since there is perfect information at the end of each period. An alternative way of endowing the model with some persistence is to follow Woodford (2003). In this set up, ambiguity aversion not only makes the shocks asymmetric on impact, but it also changes their persistence. In particular, disinflationary signals take longer to be incorporated into agents' beliefs with the result that recessions are not just deeper, but also longer-lived, than booms.

Each period, agents receive a public signal x_t as before, but now, instead of the true state being

revealed after one period, the true state is never revealed. On the other hand, we dispense with the Calvo friction so that wages can be reset every period. Now

$$\tilde{E}_t(\theta_t) = \sum_{j=0}^{\infty} \left(\prod_{i=0}^{j-1} \psi_{t-i} \right) (1 - \psi_{t-j}) x_{t-j},$$

where ψ_t corresponds to the Kalman gain coefficient under worst case beliefs. We can substitute this into (25) to get

$$w_t = \frac{\alpha}{1 - \alpha} \left[\sum_{j=0}^{\infty} \left(\prod_{i=0}^{j-1} \psi_{t-i} \right) (1 - \psi_{t-j}) x_{t-j} \right] + \frac{1}{1 - \alpha} \xi_t.$$

We see that now, the shock not only affects the magnitude, but also the persistence of the shocks. In particular, negative shocks will on average take longer to be incorporated into the price, which in turn will result in more persistent declines in output. This increases the welfare costs of negative demand shocks.

Appendix IV: Costs of Business Cycles

In this environment, demand shocks are more costly to welfare than in standard models of business cycles. Lucas (1987), in a highly influential study, performs a back-of-the-envelope calculation that implies that the welfare costs of ordinary business cycles, measured in units of life-time consumption, are extremely small (around one-twentieth of one percent). The basic intuition underlying this result is that negative shocks are cancelled out by positive shocks, resulting in second order gains from demand-management policies. However, as pointed out by Schmitt-Grohé and Uribe (2011), in a world with asymmetric rigidities, such calculations need not be true. In the present environment, as seen in figure 4, negative shocks cause far larger drops than positive shocks – therefore, demand-management policy can reap first-order gains.

To formalize this intuition, we can replicate the calculation in Lucas (1987) for the present model. Let

$$c_t^{det} = c_0 e^{gt}$$

represent a deterministic consumption path growing at rate g starting from c_0 . Let the stochastic consumption stream be

$$c_t = c_0 e^{gt} \exp(\bar{\kappa} \mathbf{1}(\varepsilon_t \leq 0) \varepsilon_t + \underline{\kappa} \mathbf{1}(\varepsilon_t > 0) \varepsilon_t),$$

where ε_t is a Gaussian- $(0, \sigma^2)$ demand shock, and $\bar{\kappa} > \underline{\kappa}$ corresponds to the piecewise-linear slopes of (12). This is equilibrium consumption in a model like the one presented above that also features

deterministic growth. To measure the welfare costs of demand shocks in permanent consumption units, set

$$\sum_{t=0}^{\infty} \beta^t u(\lambda c_t^{det}) = E \left(\sum_{t=0}^{\infty} \beta^t u(c_t) \right),$$

and solve for λ . Note that the expectation on the right-hand side features no ambiguity-aversion, but is the objective probability distribution of $\{c_t\}_{t=0}^{\infty}$ given normally distributed demand shocks. Following Lucas (1987), assume CRRA utility with risk aversion parameter γ . Then we can derive the following analytical expression for λ

$$(1 + \lambda)^{1-\gamma} = \frac{1}{2} \left\{ \exp((\gamma - 1)\bar{\kappa}(1 + (\gamma - 1)\bar{\kappa})\sigma^2) \left[1 + \Phi \left(\frac{1 + 2(\gamma - 1)\bar{\kappa}\sigma}{2\sqrt{2}} \right) \right] + \exp((\gamma - 1)\underline{\kappa}(1 + (\gamma - 1)\underline{\kappa})\sigma^2) \left[1 - \Phi \left(\frac{1 + 2(\gamma - 1)\underline{\kappa}\sigma}{2\sqrt{2}} \right) \right] \right\},$$

where Φ is the CDF of a standard normal distribution. Note that if we set $\bar{\kappa} = \underline{\kappa} = 1$, we recover the original calculation done by Lucas. To get a sense for how large this cost is, Lucas calibrates his model by letting $\gamma \in [1, 4]$ and $\sigma = 0.032$. The standard deviation of the shocks are taken from the residuals of a linear regression. The present model implies that such a procedure would underestimate the variance of the true underlying shocks. This harks back to the debate between Romer (1986) and De Long and Summers (1988) about whether macroeconomic policy reduces the variance of shocks or fills in the troughs without shaving the peaks. In the table below, I take $\sigma = 0.032$ and $\bar{\kappa} = 1$, to make my results directly comparable with those of Lucas (1987), with the caveat, that re-estimating σ would result in even larger differences. In particular, the variance of the consumption process I specify is $(1 + \underline{\kappa}^2)\sigma^2/2$ which is strictly less than the variance used by Lucas as long as $\underline{\kappa} < 1$.

		γ			
		1	2	3	4
	0	25.28	25.62	25.98	26.34
	0.1	22.86	23.24	23.63	24.03
	0.2	20.45	20.87	21.3	21.74
	0.3	18.03	18.5	18.98	19.47
	0.4	15.6	16.13	16.67	17.21
$\underline{\kappa}$	0.5	13.18	13.77	14.36	14.97
	0.6	10.75	11.41	12.07	12.74
	0.7	8.315	9.048	9.787	10.53
	0.8	5.88	6.695	7.514	8.337
	0.9	3.441	4.345	5.251	6.159

Table 7: The ratio of λ in this model to that in Lucas (1987).

The results of the calibration are in table 7. We see that even a modest amount of asymmetry can substantially increase the welfare costs of demand shocks for the US. In particular, if we calibrate κ using table 1, then the welfare cost of demand-driven business cycles are approximately 1% percent of life-time consumption, or about 20 times the cost found by Lucas (1987). There are reasons to believe that these estimates are a lower bound on the welfare costs even in the context of a complete markets, representative consumer economy. First, according to the model, the variance of the underlying shocks in the data is larger than what one would estimate from the residuals of a least squares regression. Second, as we will see in the next section, in a dynamic model, ambiguous information can cause distortions in the deterministic steady-state of a linearized model and make negative shocks more persistent than positive shocks, further increasing the welfare costs of shocks in the model.