INFLATION DYNAMICS AND TIME-VARYING VOLATILITY: NEW EVIDENCE AND AN SS INTERPRETATION

Joseph Vavra[*]
University of Chicago[†]

9/12/2012

Abstract

Is monetary policy less effective at increasing real output during recessions than during normal times? In this paper, I argue that volatility rises during recessions, which leads to an increase in aggregate price flexibility so that nominal stimulus mostly generates inflation rather than output growth. To do this, I estimate price-setting models with "volatility shocks" and show these models match new facts in CPI microdata that standard price-setting models miss. Furthermore, these models imply that output responds dramatically less to nominal stimulus during volatile recessions. For example, the estimated output response to additional nominal stimulus in October 2001, a time of high volatility, is less than half of the response in September 1995, a time of low volatility. I confirm this prediction by estimating forward-looking Phillips curves that show a greater inflation-output tradeoff during volatile periods.

JEL Classification: E10, E30, E31, E50, D8
Keywords: Volatility, uncertainty, Ss model, fixed cost, menu cost, time-varying impulse response, monetary policy

[*]I would like to thank Eduardo Engel for invaluable advice as well as John Leahy, my discussant at the NY Fed Monetary Workshop. I would also like to thank Rudi Bachmann, David Berger, Nick Bloom, Martin Eichenbaum, Francois Gourio, Amy Meek, Giuseppe Moscarini, Guillermo Ordonez, and Tony Smith. I am also grateful for comments from seminar participants at various universities and conferences. This research received generous support from the Society for Computing in Economics and Finance student paper prize.

[†]Correspondence: Email: joseph.vavra@chicagobooth.edu. Phone: 773-834-0959. Address: 5807 S. Woodlawn Ave Chicago, IL 60637

1 Introduction

The recent recession has led to renewed interest over whether the ability of monetary policy to stimulate the economy varies across time. If monetary policy becomes less effective during recessions, then its use as the primary stablizing mechanism for the economy might be called into question. While there are many features of the most recent recession that might be special,[1] I argue that all recent U.S. recessions have been associated with increased cross-sectional volatility and that this volatility leads to a reduction in the effectiveness of monetary policy at precisely the time it is needed most. In particular, in price-setting models that are consistent with micro evidence, greater volatility leads to greater aggregate price flexibility so that nominal stimulus generates mostly inflation rather than output growth.

An explosion of microdata has led to a large literature studying aggregate price-flexibility through the lens of micro price-setting models. Midrigan (2011) and Nakamura and Steinsson (2010) show that Ss price-setting models where firms face a fixed cost of price adjustment can capture a variety of micro facts while generating significant aggregate monetary non-neutrality.[2] Furthermore, these models endogenize the frequency of price adjustment and thus have scope to generate time-varying aggregate price flexibility. However, it is well-known since Caplin and Spulber (1987) that the frequency of adjustment is not enough to pin down aggregate price-flexibility in Ss price-setting models, and Golosov and Lucas (2007) show this in a more quantitative setting. In addition to the frequency of adjustment, it is necessary to track higher moments of the distribution of firms' desired price changes in order to determine the aggregate price response to changing aggregate conditions.

Towards this end, I use the BLS micro data that underlies the CPI to directly examine how the distribution of price changes moves over the business cycle. In particular, I document two new facts: 1) The cross-sectional standard deviation of price changes is strongly countercyclical: price changes become substantially more

[1] Much attention has been devoted to the effects of financial frictions and the zero lower bound for monetary policy in the most recent recession.

[2] Eichenbaum, Jaimovich, and Rebelo (2009) provides more direct evidence that price-setters engage in Ss like behavior, adjusting prices when their markup drifts out of a target region.

disperse during recessions, when other measures of volatility typically rise. 2) The standard deviation of price changes comoves strongly with the frequency of price adjustment in the economy. That is, the dispersion of price changes (conditional on adjustment) is high when more products are changing prices.

I next assess the ability of standard price-setting models with only aggregate first moment shocks[3] to match these new empirical facts. In particular, I focus on Ss price-setting models because they endogenize both the frequency of adjustment and the distribution of price changes.[4] In these models, firms face idiosyncratic productivity shocks as well as aggregate nominal shocks and must pay a fixed cost to adjust their nominal prices. This fixed cost induces firms to follow Ss pricing strategies: they only adjust to their desired price when their current price falls outside of some threshold (S,s) region. Within this region, it is not worth paying the adjustment cost and firms maintain a constant price. While these models have been shown to do a good job of capturing a wide range of micro pricing facts, I show that they get the new empirical facts wrong. Ss models with only first moment shocks imply a counterfactual negative correlation between price change dispersion and the frequency of adjustment, and they generate procyclical price change dispersion.

I then show that the addition of countercyclical volatility shocks, which increase the cross-sectional variance of firm level idiosyncratic productivity, improves the model fit dramatically.[5] Many recent papers, including Bloom (2009); Gilchrist, Sim, and Zakrajsek (2010); Arellano, Bai, and Kehoe (2010); Basu and Bundick (2011); and Bloom et al. (2012) have advanced the view that increases in volatility and uncertainty may cause recessions. In an opposing literature, Bachmann and Moscarini (2011) argues that the causality runs in the opposite direction, with business cycles causing additional volatility, and Bachmann and Bayer (2011b) argues that volatility shocks are too small to be significant drivers of the business cycle. Despite this disagreement, there is growing consensus that the variance of firm-level

[3] First moment shocks, such as changes in nominal output, shift all firms' flexible price by the same amount.

[4] Other popular price-setting models miss the correlation between frequency and price change dispersion for trivial reasons: the Calvo model features a constant exogenous frequency of adjustment, and most pure information processing models imply that all prices are adjusted every period so the frequency of price adjustment is constant at 100%.

[5] Keeping with Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012) and Fernandez-Villaverde, Guerro, Rubio-Ramirez, and Uribe (2011), I do not attempt to model the sources of second moment shocks.

productivity is countercyclical even if the source of that countercyclical volatility and its implications for business cycles remains unsettled. In this paper, I show that even modest countercyclical volatility has important implications for micro price-setting behavior and the monetary transmission mechanism.[6]

In the model, "second moment" shocks that increase idiosyncratic volatility have two effects: First, an increase in volatility has a direct effect that pushes more firms to adjust for a given region of inaction. However, greater volatility also increases the option value of waiting, which leads to a "wait-and-see" effect that widens the size of the inaction region and decreases price adjustment. Assessing the relative strength of these effects is a quantitative question, and I find that in the estimated models the volatility effect strongly dominates so that the frequency of adjustment rises during times of high volatility. Furthermore, price change dispersion also rises as the variance of the shocks rises so that there is a positive correlation between the frequency of adjustment and dispersion.

More importantly, as the region of inaction grows but more firms are pushed to adjust by larger shocks, the price level becomes more flexible so that the price response to nominal shocks is greater in times of high volatility. This greater price response then implies that real output responds much less to nominal shocks in times of high volatility so that the inflation-output tradeoff worsens. The quantitative importance of this effect is large. For example, the total response of real output to nominal stimulus in October 2001, a highly volatile time, is between one-half and one-third of the response in September 1995, a time of very low volatility. In contrast, in Ss models with only first moment shocks, the real response to nominal shocks does not vary cyclically.

Why does greater volatility lead to greater aggregate price responsiveness? In Ss models, the price response to nominal shocks can be decomposed into two margins: the intensive and extensive margin. In response to a positive nominal shock, the intensive margin is given by the extra amount that firms who were already adjusting now raise their prices. The extensive margin is given by the change in the mix of adjusters. When there is a positive nominal shock, some firms that would otherwise have left their prices constant now raise prices while some firms that would have lowered prices now leave them constant. When volatility increases, both margins

[6] Furthermore, these results do not depend on any causal link between volatility and the business cycle. It is not important whether volatility drives recessions or vice versa.

become more important as there are more firms adjusting prices, and more firms are pushed near the margin of adjustment by the more volatile shock process. Overall, the price response on impact to a nominal shock is 120% larger at the 90th percentile of volatility than it is at the 10th percentile of volatility, and two-thirds of this increase is driven by a more responsive extensive margin.

For the majority of the paper, I focus on a version of the Golosov and Lucas (2007) model. I report results from this relatively simple model for two main reasons: 1) It illustrates the effects of volatility in a transparent manner. 2) In the first part of the paper, I argue that Ss models with first moment shocks are inconsistent with micro evidence while Ss models with second moment shocks can match this evidence. In order to state this conclusively and show that these are robust implications, it is important to estimate standard errors for the model estimates. This is computationally burdensome and using the simple Golosov and Lucas (2007) version of the model drastically reduces the computing time required.[7] Nevertheless, this model has several weaknesses. Most importantly, it is well-known that the model exhibits little non-neutrality and thus might seem to be an uninteresting departure point for the study of time-varying non-neutrality. In addition, as emphasized by Midrigan (2011), the model does a poor job of matching the average distribution of price changes, and it also features a particularly simple aggregate shock process.

Nevertheless, I show that the importance of volatility shocks documented for the Golosov and Lucas (2007) model survives after a variety of empirical enrichments. I introduce random menu costs as in Dotsey, King, and Wolman (1999), which generates results for price-setting behavior that are similar to the multiproduct firms of Midrigan (2011). Furthermore, I introduce leptokurtic shocks to match the excess kurtosis of price-changes observed in CPI data, which again moves the model closer to Midrigan (2011) or Gertler and Leahy (2008). Finally, I introduce more realistic aggregate shocks into the model. These extensions capture the important features of the "second-generation state-dependent" models that Klenow and Kryvtsov (2008) argue are consistent with empirical patterns of price-setting. In particular, these extensions substantially improve the model's ability to fit the distribution of price changes, amplify non-neutrality and generate hump-shaped impulse response func-

[7] Estimating standard errors for the Golosov and Lucas (2007) takes approximately two weeks. Each of these extensions adds at least one state-variable and increases computational time by an order of magnitude.

tions. Thus, they have substantial effects on the average response of the model to nominal shocks. However, I show that the basic insights of the Golosov and Lucas (2007) model remain: all of the models with only first moment shocks remain inconsistent with micro evidence while the models with second moment shocks are consistent with this evidence. More importantly, aggregate price flexibility continues to rise substantially with volatility. Thus, countercyclical volatility improves the empirical fit of an extremely wide class of price-setting models and has important implications for the monetary transmission mechanism.

The general conclusion of the model, that the price level becomes more flexible so that the inflation-output tradeoff worsens during volatile recessions, is thus robust to a number of important model extensions. However, I also look for additional evidence of this mechanism that does not rely on a particular structural model.[8] Using aggregate inflation data, I directly test for a time-varying inflation-output tradeoff by estimating forward looking Phillips curve as in Gali and Gertler (1999) but allowing the response of inflation to marginal cost to vary with volatility.[9] Consistent with the model, I find that inflation responds much more strongly to marginal cost during times of high volatility. The short-run Phillips curve estimated during times of high volatility has a significantly greater slope than the Phillips curve estimated during times of low volatility. This time-variation is predicted by the Ss models with volatility shocks while it is at odds with older Keynesian models that predict that the slope of the Phillips curve should be procyclical. It is also at odds with Ss models with only first moment shocks as well as with the Calvo price-setting model that imply an acyclical inflation-output tradeoff. Thus, in addition to providing a better fit to micro data, the Ss models with volatility shocks also better match aggregate inflation dynamics.

The remainder of the paper is organized as follows: Section 2 contains the empirical findings. Section 3 shows that Ss models with only first moment shocks get the

[8]Several recent papers such as Bachmann and Sims (2011) and Auerbach and Gorodnichenko (2012) have looked for reduced form evidence that the fiscal policy transmission mechanism varies over the business cycle.

[9]While it may seem more natural to estimate a VAR that varies with the state of volatility, my model is about the response of inflation and real output to changes in nominal output rather than to federal fund rate shocks, per se. Changes across time in VAR estimates could reflect time-varying responses to nominal output shocks or time-varying responses of nominal output to interest rate shocks. The Phillips curve evidence is thus a more direct test of my model. Nevertheless, in previous versions of this paper I estimated such VARs and found that the estimated responses were also consistent with my model.

empirical findings wrong. I first analyze a simple analytical Ss model and show that it generates a counterfactual negative correlation between price change dispersion and the frequency of adjustment. I then estimate a fully-specified quantitative Ss model and show that the analytical results still hold, and that furthermore, in contrast to the empirical evidence, the model implies procyclical price dispersion. Section 4 adds second moment shocks to the model and shows that the model's fit is significantly improved. Section 5 discusses policy implications of the second moment shocks. Section 6 shows results for a number of empirical extensions of the benchmark model. Section 7 presents the reduced form evidence that price flexibility rises during times of high volatility and Section 8 concludes.

2 Empirical Results

The restricted access *CPI research database* collected by the Bureau of Labor Statistics (BLS) contains price observations for the thousands of non-shelter items underlying the CPI from January 1988 through January 2012. Prices are only collected monthly for the entire sample period in New York, Los Angeles and Chicago so my main analysis is restricted to these cities.[10] The database contains thousands of individual "quote-lines" with price observations for many months. Quote-lines correspond to an individual item at a particular outlet. For example, the type of quote-line collected in the research database might be 2-liter "Brand-X" cola at a particular Chicago outlet. See Nakamura and Steinsson (2008) for additional description.

This database has received great attention in recent years, beginning with Bils and Klenow (2004). While initial studies of this micro pricing data focused on static first moments of the data such as the average frequency and size of price changes, only recently have more dynamic features of the data begun to receive attention. (See Klenow and Malin (2011) for detailed summaries of the recent literature utilizing this data). However, despite the widespread attention this data has received, analysis of higher moments of the price change distribution has begun only recently. Klenow and Malin (2011) provides brief evidence of the relationship between first and higher moments of inflation and calls for additional attention to this topic.

In this paper, I focus on the business cycle properties of the distribution of price changes rather than the relationship between the distribution of price changes and

[10] Using the full sample does not qualitatively affect the results. See the data appendix.

inflation.[11] In particular, I will show that price change dispersion (the second moment of the price change distribution)[12] has robust, dynamic patterns with strong implications for models of price setting. Let $dp_{i,t} = \log \frac{p_{i,t}}{p_{i,t-1}}$ be the log price change observed for item i at time t. Then, using aggregation weights provided by the BLS, it is straightforward to compute the cross-sectional dispersion of log price changes for each month and investigate how it varies over time. Figure 1 shows the interquartile range of price changes[13] as well as the median frequency of adjustment. It is clear from this figure that there is a large increase in the dispersion of price changes during recessions. Furthermore, there is a more modest rise in the frequency of adjustment during recessions.[14] Since both the frequency of adjustment and the the interquartile range exhibit high-frequency noise and low-frequency trends, I next bandpass filter the two series to leave only the variation at business cycle frequencies. Figure 2 shows that at business cycle frequencies, two robust relationships emerge. 1) The dispersion of price changes is countercyclical. 2) There is strong positive comovement between price change dispersion and the frequency of adjustment.[15] Table I confirms these relationships numerically.

In principle, the aggregate relationships reported in Table I may be driven by compositional changes in the make up of price changes over the business cycle. Appendix 1 explores this in much more detail and argues that this is not the case. In particular, these results hold both within as well as across sectors, and they hold for both price increases and decreases. I also find similar results using a balanced panel of prices spanning recessions so that the results are unlikely to be driven by product entry and exit. This suggests that the countercyclicality of price dispersion and the positive relationship between price dispersion and the frequency of adjustment are robust facts that are common to the majority of items and price changes in the U.S. economy.

Appendix 1 discusses the construction of these statistics in more detail as well

[11] The empirical results in this paper extend and supplant those in Berger and Vavra (2011). I am grateful for David Berger's support and advice in investigating these empirical questions.

[12] See the empirical appendix for discussion of higher moments.

[13] All dispersion numbers exclude zeros. This is not an important restriction since jointly matching the frequency of adjustment and the distribution of price changes excluding zeros necessarily implies matching the distribution of price changes including zeros.

[14] A regression on recession dummies shows that during recessions, the interquartile range is 38% higher and frequency is 10% higher.

[15] It should also be noted that while the price dispersion relationships are most dramatic for the most recent recession, both facts remain after the exclusion of these dates.

as providing a battery of robustness checks. In particular, while the above facts are computed using bandpass filtered data, similar qualitative results obtain using alternative filtering.[16] In addition, the results are robust to the inclusion or exclusion of zero price changes, as well as to the inclusion of sales and product substitutions. Finally, different procedures to control for outliers and measurement error do not significantly alter the results.

The remainder of the paper then takes these empirical facts as given and assesses the extent to which they can be generated by Ss pricing models. In particular, I will argue that when viewed through the lens of Ss models, these price-setting facts strongly suggest that second moment shocks are an important feature of the economic environment affecting firms' pricing decisions.

3 Ss Models with First Moment Shocks

In this section, I present evidence that Ss models with only first moment shocks imply a negative correlation between the dispersion of price changes and the frequency of adjustment, in contrast to the empirical evidence. I begin with simple intuition for why Ss models with only first moment shocks generate a negative correlation between frequency and price change dispersion in a discrete time setting. I next prove that the intuition holds in an analytical continuous time model, and I then move to a fully-specified quantitative model and show that the result from the analytical model still holds and that, in addition, price dispersion is procyclical.

3.1 Intuition for Negative Correlation Between Frequency and Price Change Dispersion

Ss models with only first moment shocks imply a negative correlation between the frequency of price changes and the dispersion of those price changes. In Ss models, firms face a fixed cost of adjusting prices so that a firm will only choose to adjust its price if its current price is far from its desired price. In order to increase the frequency of adjustment, more firms must be pushed out of the inaction region. However, aggregate first moment shocks will, by definition, affect all firms' desired price changes

[16] Ashley and Verbrugge (2007) argues that two-sided bandpass filters are misspecified and instead argue for the use of a one-sided bandpass filter. Recomputing statistics using this one-sided filter, if anything, strengthened the above facts.

in the same way. Thus, firms must all be pushed out of the inaction in the same direction. While this leads to an increase in the frequency of adjustment, more price changes are then in the same direction, which leads to a decrease in price change dispersion. This is illustrated in Figure 3.

The upper panel shows the distribution of "price gaps", the difference between a firm's current price and the price it would choose if it adjusts.[17] After a first moment shock that increases firms' desired prices, the distribution of price gaps shifts from the blue to the green distribution. The lower panels show the effect on the distribution of actual price changes. After the first moment shock, the frequency of adjustment increases since the additional mass of firms raising prices is greater than the reduction in mass of firms decreasing prices. At the same time, the variance of price changes falls as we move from the blue to the green distribution. Thus, there is a negative correlation between the frequency of adjustment and price change dispersion.

While this figure is just to illustrate the intuition, in Appendix 2 I prove that this intuition holds in the continuous time limit of a simple two-sided Ss model. While this analytical result is quite strong, it relies on a model with several simplifying assumptions. It is partial equilibrium, imposes a quadratic loss function, and it has a very simple process for the evolution of the price gap, which is not derived from micro foundations. Thus, I now move to a more quantitative Ss model. This general equilibrium model will have fully specified micro foundations and its quantitative fit to the empirical data will be evaluated more formally. Nevertheless, I will show that the result from the analytical model still holds: the Ss model with only first moment shocks implies a negative correlation between price change frequency and dispersion.

3.2 Quantitative Equilibrium Ss model

The baseline quantitative model closely follows Golosov and Lucas (2007), and I incorporate aggregate shocks using methods first explored in Midrigan (2011). I focus on this model for two main reasons. First, I want to argue that Ss models with and without second moment shocks have very different implications for the relationship between price change dispersion and frequency. Making this statement formally requires estimating standard errors for model parameters and moments, which is computationally burdensome and only feasible in this simple environment.

[17] The distribution will be skewed in the presence of any positive trend inflation.

Second, I want to illustrate the importance of time-varying volatility for aggregate price flexibility in a transparent manner before enriching the model. Nevertheless, in Section 6, I will enrich the model along a number of dimensions emphasized in recent literature, and I will argue that similar results obtain both for the models' implications for micro price-setting facts as well as for time-variation in aggregate price flexibility.

The baseline economy is composed of a representative household and a continuum of monopolistically competitive firms. I first discuss the household problem, and I then present the firm problem and define equilibrium.

3.2.1 Households

Households allocate income and labor to maximize a Dixit-Stiglitz consumption aggregate subject to indivisible labor supply

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \left[\log C_t - \omega n_t\right],$$

subject to

$$\int_0^1 p_t^i c_t^i di + E_t\left[q_{t,t+1} B_{t+1}\right] \leq B_t + W_t n_t + \int_0^1 \Pi_t^i di,$$

where

$$C_t = \left(\int_0^1 \left(c_t^i\right)^{\frac{\theta-1}{\theta}} di\right)^{\frac{\theta}{\theta-1}}$$

is a Dixit-Stiglitz aggregator of consumption goods c_t^i, p_t^i is the price of good i, n_t is the household's labor supply, ω is the disutility of labor, W_t is the nominal wage, Π_t^i is nominal profits the household receives from owning firm i, and θ is the elasticity of substitution. A complete set of Arrow-Debreau state-contigent claims are traded in the economy so that B_{t+1} is a random variable that delivers payoffs in period $t+1$ from financial assets purchased in period t and $q_{t,t+1} = \beta \frac{C_t}{C_{t+1}}$ is the stochastic discount factor used to price these claims.

3.2.2 Firms

Firms produce output using a linear technology in labor

$$y_t^i = z_t^i a_t l_t^i,$$

where firm i's idiosyncratic productivity z_t^i evolves according to

$$\log z_t^i = \rho_z \log z_{t-1}^i + \sigma_z \varepsilon_t^i; \quad \varepsilon_t^i \sim N(0,1),$$

aggregate productivity a_t evolves according to

$$\log a_t = \rho_a \log a_{t-1} + \sigma_a \varepsilon_t^a; \quad \varepsilon_t^a \sim N(0,1)$$

and l_t^i is labor rented by firm i. After choosing prices, firms fulfill all of the resulting consumer demand:

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

where P_t is the Dixit-Stiglitz price index

$$P_t = \left(\int_0^1 \left(p_t^i\right)^{1-\theta} di\right)^{\frac{1}{1-\theta}}.$$

I assume that nominal aggregate spending $S_t = P_t C_t$ follows a random walk with drift in logs:[18]

$$\log S_t = \mu + \log S_{t-1} + \sigma_s \varepsilon_t^s; \quad \varepsilon_t^s \sim N(0,1).$$

Firms must pay a fixed cost f in units labor in order to adjust their nominal price. Given these constraints, firm i's problem is then to choose prices to maximize discounted profits

$$\max_{p_t^i} E_t \sum_{t=0}^{\infty} q_{t,t+1} \pi_t^i,$$

where firm profits are given by

$$\pi_t^i = \left(\frac{p_t^i}{P_t} - \frac{W_t}{z_t^i a_t P_t}\right)\left(\frac{p_t^i}{P_t}\right)^{-\theta} C_t - f\frac{W_t}{P_t} I_{p_t^i \neq p_{t-1}^i},$$

[18] This is a computational simplification that reduces the aggregate state-space by one dimension. In Section 6 I relax this assumption.

and $I_{p_t^i \neq p_{t-1}^i}$ is an indicator function for nominal price changes.

3.2.3 Computing the Equilibrium

In order to bound the state-space of the problem, all nominal variables are normalized by current nominal spending in the economy. The idiosyncratic states of the economy are given by the firm's previous nominal price p_{t-1}^i and its current level of productivity z_t^i. The aggregate state of the economy can be summarized by the current level of nominal spending S_t, the value of aggregate productivity a_t, and the joint distribution of idiosyncratic states $\phi\left(p_{t-1}^i, z_t^i\right)$. Since the evolution of aggregate state variables depends on this joint distribution, the state space of the problem is thus infinite dimensional. Following Krusell and Smith (1998) and its application to Ss models in Midrigan (2011), I conjecture that the decomposition of changes in S_t into changes in P_t is given by the following forecasting rule:

$$\log \frac{P_t}{S_t} = \gamma_0 + \gamma_1 \log a_t + \left[\gamma_2 + \gamma_3 \log a_t\right] \chi_{1,t}$$

where $\chi_{1,t} = \log \frac{P_{-1}}{S_t} + \log a_t$.[19] Given this conjecture, I then search for a value of the transition coefficients, $\boldsymbol{\gamma}$ so that the true law of motion in the economy is well approximated by the conjectured law of motion. At this point, a regression of the actual law of motion on the conjectured law of motion gives R^2 in excess of 99%. Furthermore, adding an additional moment (the cross-sectional variance of price gaps) to the forecasting rule did not change the qualitative conclusions.[20] Finally, rather than comparing the conjectured law of motion to the actual law of motion period-by-period as is implied by the linear regression, a series of aggregate variables can be computed entirely from the conjectured law of motion and compared to results computed directly from the simulated model as suggested by Den Haan (2010). The approximation errors remain extremely small.

Given the conjectured law of motion, the firm problem can be written recursively as

$$V\left(\frac{p_{-1}^i}{S}, z^i; \chi_1, a\right) = \max\left[V^N\left(\frac{p_{-1}^i}{S}, z^i; \chi_1, a\right), V^A(z^i; \chi_1, a)\right]$$

[19] This is a sufficient statistic in a Calvo environment. It essentially captures the average "price gap" across firms.

[20] These effects are essentially captured by a_t.

where the value of not adjusting and adjusting are given respectively by

$$V^N\left(\frac{p^i_{-1}}{S}, z^i; \chi_1, a\right) = \pi\left(\frac{p^i_{-1}}{S}, z^i; \chi_1, a\right) + \beta E \frac{\frac{S}{P}}{\frac{S'}{P'}} V\left(\exp\left[\log\frac{p^i_{-1}}{S} - (\mu + \varepsilon^s)\right], z^{i\prime}; \chi_1', a'\right)$$

and

$$V^A(z^i; \chi_1, a) = -f\omega\frac{S}{P} + \max_{\log p^i/S}\left[\pi\left(\frac{p^i}{S}, z^i; \chi_1, a\right) + \beta E\frac{\frac{S}{P}}{\frac{S'}{P'}} V\left(\exp\left[\log\frac{p^i}{S} - (\mu + \varepsilon^s)\right], z^{i\prime}; \chi_1', a'\right)\right],$$

and firm flow profits can be written as

$$\pi\left(\log\frac{p^i}{S}, z^i; \chi_1, a\right) = \left(\frac{p^i}{S} - \frac{\omega}{az^i}\right)\left(\frac{p^i}{S}\right)^{-\theta}\left(\frac{P}{S}\right)^{\theta-2}.$$

For more details on this representation as well as for expanded expressions for the law of motion for these variables see Appendix 3.

3.2.4 Estimation and Results

The model period is one month, so I set the discount factor $\beta = .997$. The calibration of the nominal shock process follows Nakamura and Steinsson (2010). Since there is no long-run real growth in the model economy, I set $\mu = .002$ to match the mean growth rate of nominal GDP minus real GDP, and I set $\sigma_s = .0037$ to match the standard deviation of nominal GDP growth, over the period 1998-2012.[21] The production function is linear in labor, the sole factor of production, so I calibrate the aggregate productivity process with $\rho_a = .91$ and $\sigma_a = .006$ so that the model matches the quarterly persistence and standard deviation of average labor productivity.[22]

The remaining parameters are estimated using simulated method of moments. There are four model parameters: the elasticity of substitution θ, the persistence and standard deviation of idiosyncratic productivity, ρ_z and σ_z, and the fixed cost f. These parameters are selected to fit four moments from the data: the average frequency of adjustment, the average size of increases, the average size of decreases, and the fraction of price changes that are increases. For more details on the estimation

[21] This yields inflation that matches the behavior of U.S. inflation on average.

[22] As measured by non-farm business output per hour. Alternatively, calibrating the productivity process in the model to match TFP would imply higher persistence. Increasing the persistence of productivity in the model did not change the qualitative conclusions.

procedure as well as a discussion of alternative moments and estimation schemes, see Appendix 3.

Table II shows the model's estimated parameters and best fit moments. The estimated parameters of the model are in line with recent literature. The elasticity of substitution of 6.9 is in between the values used by Golosov and Lucas (2007) and Nakamura and Steinsson (2010). The fixed cost implies that total adjustment costs in the economy represent just over 0.3% of steady-state monthly revenues.[23] The estimated persistence of productivity is relatively low, with a monthly persistence of 0.63. This low persistence is largely driven by matching the large and relatively frequent price changes observed in the data. Again, the productivity parameters are roughly in line with previous estimates in the menu cost literature.

Unsurprisingly, the model does a good job of matching the frequency of adjustment, the size of increases and decreases, and the fraction of price changes that are increases. A main reason for using Ss models has been their ability to match these micro moments. In contrast, as predicted by the analytical model, the Ss model with first moment shocks generates an extremely strong negative correlation between the frequency of adjustment and the cross-sectional standard deviation of price changes. This is in contrast to the strong empirical positive correlation.[24]

A final empirical failure of the model with only first moment shocks is the business cycle behavior of price dispersion. Empirically, there is a negative correlation of -.41 between output and the standard deviation of price changes while in the model there is a positive correlation of 0.19.

4 Second Moment Shocks

The previous section argues that Ss models with first moment shocks are unable to match the empirical evidence on price dispersion. In this section, I argue that the addition of time-varying volatility or second moment shocks improves the model fit dramatically. I focus on this mechanism because there is mounting evidence

[23]This measure of the fixed cost is given by $f * freq * \frac{\theta-1}{\theta} / Y_{ss}$. The cost conditional on adjustment is around 3% of revenues, which is roughly in line with the estimates in Zbaracki, Ritson, Levy, Dutta, and Bergen (2004).

[24]Previous versions of this paper attempted to explicitly target the empirical correlation, and no configuration of parameters is able to generate a positive correlation for the model with only first moment shocks.

that volatility is countercyclical. Furthermore, it can be shown that several other plausible mechanisms for variation over the business cycle are inconsistent with a positive comovement between frequency and price change dispersion. For example, time-varying costs of price adjustment or market power would still yield negative comovement.

While there are various notions of volatility, in the model, an increase in volatility will be represented by a common increase in the standard deviation of firms' idiosyncratic productivity. Since it is idiosyncratic productivity differences across firms that generate price dispersion through variation in markups, it is natural to increase the standard deviation of productivity in order to increase price change dispersion. Nevertheless, any mechanism that generates countercyclical markup dispersion so that firms have greater desire to adjust prices during recessions should generate similar results.[25]

With time-varying volatility, the negative correlation between frequency and price dispersion implied by standard Ss models can potentially be broken. If more dispersed shocks induce more firms to adjust, then price dispersion and the frequency of adjustment can comove.

4.1 The Model

The model differs from the model in Section 3 in only one dimension. A firm's idiosyncratic productivity now evolves as

$$\log z_t^i = \rho_z \log z_{t-1}^i + d_t \sigma_z \varepsilon_t^i; \quad \varepsilon_t^i \sim N(0,1)$$

where the standard deviation of firm level shocks d_t itself evolves as

$$\log d_t = \rho_d \log d_{t-1} + \sigma_d \varepsilon_t^d; \quad \varepsilon_t^d \sim N(0,1).$$

That is, firms face idiosyncratic shocks with common standard deviation d_t, and this standard deviation is itself time-varying. For computational simplicity, I make the assumption that aggregate productivity and d_t are perfectly negatively correlated.

[25] For example Bachmann and Moscarini (2011) or countercyclical volatility of demand.

That is,

$$\begin{aligned}\rho_a &= \rho_d \\ \varepsilon_t^a &= -\varepsilon_t^d.\end{aligned}$$

While this is a strong assumption, it provides computational advantages by reducing the state-space by one dimension so that the model can again be estimated. Under the assumption that a_t and s_t are perfectly negatively correlated, the firm problem can be characterized in terms of the same state variables as in the problem with only first moment shocks. The only difference in the firm problem is that the standard deviation of idiosyncratic firm shocks now depends on the aggregate state of the economy and firms' expectations must account for this time-varying standard deviation. The form of the Krusell and Smith (1998) transition rule remains unchanged, and at a fixed point, the conjectured law of motion again provides an extremely accurate forecast for the true law of motion of the economy. In Section 6 I relax this perfect negative correlation and show that the qualitative results of the model remain unchanged.

4.2 Estimation and Results

The estimation procedure remains the same as in the model with only first moment shocks. The model is estimated to match the frequency, size of increases and decreases. There is now one additional parameter, σ_d, which I calibrate using external evidence on the size of shocks to idiosyncratic volatility. In particular, I pick the standard deviation of idiosyncratic volatility to match U.S. census data in Bloom et al. (2012). Bloom et al. (2012) computes the cross-sectional standard deviation of firm level TFP for each year from 1972-2009, and I target the time-series variation in this measure of productivity dispersion. For the years corresponding to the BLS pricing sample, Bloom et al. (2012) finds an annual coefficient of variation for the standard deviation[26] of firm level TFP equal to 1.01, and this is the value I calibrate my model to match.

However, it should be noted that my data and model are about the dispersion of *product level* shocks for a broad swathe of the U.S. economy while the data in Bloom

[26] The coefficient of variation for the IQR is slightly larger. After detrending, all values are slightly smaller.

et al. (2012) comes from *plant level* manufacturing data. Since the two measures of volatility may not coincide, and since Bachmann and Bayer (2011b) also argue for smaller values using German data, I also investigate the sensitivity of my analysis to changes in the size of volatility shocks, results of which are described in Section 6.[27]

Once I calibrate the size of volatility shocks, estimation of the remaining parameters proceeds as before. In general, the estimated parameters for the model with second moment shocks, shown in Table III, are similar to those for the model with only first moment shocks.

As can be seen from Table IV, the model fit with first and second moment shocks is a dramatic improvement over the model with only first moment shocks. The model now implies a correlation between the frequency and standard deviation of price changes that is positive and closely in line with the empirical data instead of the strongly negative correlation implied by the model with only first moment shocks.

Why does time-varying volatility imply a positive correlation between the frequency of price changes and the cross-sectional standard deviation of those changes? As emphasized by Bloom et al. (2012), in the presence of fixed adjustment costs, an increase in volatility has two effects: 1) There is a "wait-and-see" effect. Greater volatility makes the option value of waiting increase as it is not worth paying adjustment costs today if a firm will want to reverse its decision tomorrow. This makes firms' inaction regions widen so that price adjustment is less likely. 2) There is a "volatility" effect. When shocks have a greater standard deviation, more firms will hit adjustment bands of any given width so that price adjustment rises.

The two effects work in opposite directions, but quantitatively, the volatility effect dominates.[28] Figure 4 shows how the density of firm price gaps as well as the adjustment hazard respond to an increase in volatility. The wait-and-see effect

[27] Previous versions of this paper instead estimated the level of volatility to match time-series variation in the dispersion of price changes. This yielded similar results.

[28] Why does the volatility effect dominate the wait-and-see effect? This is the opposite of the result obtained in the investment model of Bloom et al. (2012), and it is because the estimated fixed costs of adjustment in my model of pricing are substantially lower than those estimated for models of firm investment. In Bloom et al. (2012), the investment fixed cost is estimated to be 1.5% of annual sales with additional irreversible investment resale losses of 33.9%. In contrast, the estimated fixed cost of adjustment in my pricing model with first and second moment shocks is only 0.46% of monthly sales and there are no irreversibilities. When adjustment costs are relatively small, the option value of saving adjustment costs is also relatively small so it has little effect on firm pricing. If fixed costs in the model are increased by more than twenty times then the wait-and-see effect can be made to dominate, but this cost is implausibly large and the model then misses along many other moments.

is clear, as the adjustment hazard widens in the state of the world with increased volatility. Similarly, the density of price gaps also spreads out so that the density of firms with a high probability of adjustment rises. On net, the second effect dominates so that the frequency of adjustment rises from 7% to 12% per month when volatility moves from the 25th to the 75th percentile.

When volatility rises, adjustment bands widen so that the difference between the average price increase and the average price decrease grows. At the same time, the fraction of price changes that are increases is little affected, so the standard deviation of price changes grows. And since the volatility effect dominates, the frequency of adjustment also grows. Thus, the model implies a positive correlation between the frequency of adjustment and the cross-sectional standard deviation of price changes. It is also interesting to note that the interaction between the volatility effect and the wait-and-see effect implies that the frequency of large (absolute) price changes grows with volatility while the frequency of small changes declines. I find support for this in the BLS CPI microdata. While the overall frequency of adjustment is countercyclical, this hides important compositional differences: the frequency of large price changes is strongly countercyclical while the frequency of small price changes is actually procyclical.[29]

5 Policy Implications

I now show that time-varying volatility has striking implications for the transmission of nominal shocks to the real economy. In times of high volatility, the real effect of nominal shocks is substantially reduced so that the economy exhibits state-dependent impulse responses. Table V shows the first element of the real output impulse response to a one-month doubling of the nominal output growth rate,[30] computed for different states of the economy[31].

The real effect on impact of the same nominal shock is nearly halved when moving from the 10th percentile of volatility to the 90th percentile of volatility. Differences in the cumulative impulse response functions are even larger.

[29] Correlation with bandpass filtered industrial production is -0.39 and 0.4, respectively, while the corresponding correlation for the total frequency is -0.28. For this calculation I used a price change of 5% as the cutoff between small and large.

[30] From .002 to .004. (This is roughly a 0.5 standard deviation shock).

[31] IRFs are calculated using the ergodic distribution.

Since $S = PY$, the fraction of the nominal shock that does not generate real output growth instead generates inflation. This means that the tradeoff between inflation and output worsens dramatically when volatility increases, since getting the same increase in real output requires much larger increases in nominal output and thus inflation.

Why does aggregate price flexibility increase during times of high volatility? I will first describe what determines the price response to a nominal shock in Ss models and then discuss how this interacts with volatility. As previously mentioned, in Ss models the price impulse response to a positive nominal shock can be decomposed into two components. The first component is the intensive margin: conditional on adjustment, all firms will raise prices more[32] after a positive nominal shock. The second component is the extensive margin: some firms close to raising prices will be pushed into action by a positive nominal shock, and some firms who previously would have lowered prices instead choose inaction after the shock.

When will each of these margins be more important? The intensive margin response to a nominal shock increases with the frequency of adjustment. The more firms that will be adjusting independent of the additional nominal shock, the larger is the aggregate price response to that shock on the intensive margin. Intuitively, as the frequency of adjustment rises, a greater fraction of aggregate price adjustment will occur along the intensive margin.[33]

The extensive margin response to a nominal shock will be more important when more firms' adjustment decisions are changed by the shock and when the difference between adjustment and non-adjustment is large. Thus, the extensive margin grows with the number of firms near the margin of adjustment and with the width of the inaction region (holding the fraction of firms near the boundary of the inaction region constant). If there are lots of firms that are on the margin of adjusting, then the mix of firms that choose to adjust will vary more in response to a nominal shock and the extensive margin will be stronger. This is the classic "selection effect" emphasized by Golosov and Lucas (2007). Furthermore, the wider the inaction region, the more

[32] Or lower less.

[33] Here it is important to note that a separate literature (such as Klenow and Kryvtsov (2008)) refers to the frequency of adjustment as the extensive margin. This is distinct from my use of the extensive margin terminology, and it is because they are decomposing inflation at a point in time rather than the response of inflation to a new aggregate shock. See Gagnon, Lopez-Salido, and Vincent (2012) for additional discussion.

this effect is amplified to give the total effect of the extensive margin on the price level: If the difference between inaction and action is small, then shifting mass from inaction to action will have less effect on the overall price level than if the difference between inaction and action is large. See Caballero and Engel (2007) for a more detailed discussion of these two margins.

From Figure 4 it is clear that both margins become more important in times of high volatility. There are more firms in the adjustment region so that the intensive margin grows. In addition, there are more firms near the adjustment bands, and the bands are of greater width, both of which increase the importance of the extensive margin. Thus, the aggregate price response along both the intensive and extensive margin is greater during times of high volatility, and as the price level becomes more flexible, the real response to nominal shocks necessarily declines. Table VI shows the contribution of the intensive and extensive margin to price flexibility at the 10th and 90th percentile of volatility. Clearly both margins become more important as volatility increases, however, the increase in the extensive margin is substantially larger than the increase in the intensive margin and accounts for approximately two-thirds of the overall increase in price flexibility. Here it is worth noting that a Calvo model calibrated to match time-variation in the frequency of adjustment would capture time-variation in the intensive margin, but it would miss time-variation in the extensive margin that is quantitatively more important.

Second moment volatility shocks are the driving feature of this time-varying response. In the model with only first moment shocks, there is no relationship between the real impact of nominal shocks and the cycle: for all values of aggregate productivity, the average output impulse response on impact is equal to roughly 55% of the nominal shock.[34] In contrast, if I solve a version of the model with second moment shocks but no aggregate productivity shocks, time-varying policy responses remain, with price flexibility still rising dramatically with volatility. However, it is important to note that when there are only second moment shocks in the model, price dispersion

[34] I find small effects of non-linearities on time-varying responsiveness in the model with first moment shocks in contrast to Bachmann, Caballero, and Engel (2010) for several reasons: 1) They study investment where the frequency of adjustment is lower and estimated adjustment costs are higher. 2) Aggregate nominal and real shocks in my model have opposite effects on firms' desired price changes. Thus there is little change in the distribution of firms' desired price changes over the business cycle, and movements in this distribution are what generate time-varying responsiveness in their model.

becomes procyclical.[35]

As a final exercise, following Bachmann et al. (2010), I can back out aggregate shocks from the model to fit U.S. economic data and then compute how the output response to nominal shocks varies across time. In order to compute the sequence of shocks[36] that best explains the observed data, I begin from the ergodic distribution and pick the value of the nominal shock as well as the value of aggregate productivity in order to match CPI inflation and industrial production growth in each month.[37]

Given the sequence of aggregate shocks, I can then calculate the output impulse response at each date from 1988-2012. Figure 5 shows the output impulse response on impact across time, as a measure of the responsiveness of real output to nominal shocks. The responsiveness index is clearly procyclical and plunges in recessions. Thus, the model implies that the inflation-output tradeoff is substantially worse during recessions.

Figure 6 shows the entire output impulse response in September 1995 and October 2001. These are times of very low, and very high volatility, respectively. The model with second moment shocks implies that the total response of real output to a nominal shock in October 2001 is approximately one-third of the response in September 1995. Clearly, there is no such difference for the model with only first moment shocks.

6 Empirical Extensions

While the Golosov and Lucas (2007) model illustrates the quantitative importance of volatility in a simple setting, the model is not without some empirical failings. While the model endogenously generates a number of facts from price-setting data like the frequency of adjustment and the average size of price increases and decreases, the model does a poor job of matching the overall distribution of price changes for the U.S. In particular, in the BLS CPI data, there are a large fraction of small price changes, and the price change distribution also exhibits excess kurtosis. The Golosov and Lucas (2007) model is inconsistent with these facts, and Midrigan (2011) argues that these moments are particularly significant for aggregate non-neutrality. In particular, he builds a model that matches these facts and shows that it generates

[35] This is due to Jensen's inequality.

[36] Shocks are picked period-by-period, so the numerical procedure is straightforward.

[37] The results are insensitive to various detrending methods of the aggregate data. This data matching procedure exactly matches inflation and quadratically detrended output growth.

significantly more monetary non-neutrality than the Golosov and Lucas (2007) model. In this section, I explore the robustness of my results along several dimensions. First, I introduce several extensions to the model to better match the distribution of price changes. I then introduce more realistic aggregate shocks into the model.

In total, these extensions capture the important features of the "second-generation state-dependent" models that Klenow and Kryvtsov (2008) argue are consistent with empirical patterns of price-setting, and I show that my conclusions are robust to each of these extensions. In particular, each of these models continues to generate a counterfactual negative correlation between the dispersion of price changes and the frequency of price changes when only subject to aggregate first moment shocks. In contrast, these models are able to match the positive correlation observed empirically with the addition of realistic second moment shocks. More importantly, in each of these extensions, aggregate price flexibility is substantially larger during times of high volatility.

6.1 Matching the Distribution of Price Changes

It is well-known that the Golosov and Lucas (2007) model generates a distribution of price changes that is too bimodal relative to the data. Furthermore, the empirical distribution of price changes is fat tailed. In this section I explore the implications of matching these facts.

In the Golosov and Lucas (2007) model, firms must pay a fixed cost to adjust prices. Instead, as in Dotsey et al. (1999) or Nakamura and Steinsson (2010), I assume that firms draw a random cost of adjusting. In particular, I assume that firms can adjust their prices for free with some positive probability, which is picked to match the fraction of small price changes in the CPI.[38] This extension flattens the adjustment hazard as a function of the price gap and means that some firms with small price gaps will still choose to adjust their prices. For an individual product, this is essentially a reduced form for the multiproduct firms in Midrigan (2011). In that model, multiproduct firms can pay one menu cost to simultaneously adjust the prices of all of their products, which means that some items will change prices even if their current price gap is small. A flatter adjustment hazard weakens the strength of the extensive margin and moves the model closer to Calvo.

[38] I target 20% of price changes to be smaller than 1/4 the mean price change.

In addition, I introduce leptokurtic shocks to firms' desired price changes. In particular, following Midrigan (2011), I now assume that

$$\log z_t^i = \begin{cases} \rho_z \log z_{t-1}^i + d_t \sigma_z \varepsilon_t^i; \quad \varepsilon_t^i \sim N(0,1) \text{ with probability } \gamma \\ \log z_{t-1}^i \text{ with probability 1-}\gamma \end{cases}$$

where I pick $\gamma = 0.13$ to target the excess kurtosis of price changes in the CPI and adjust the other parameters of the productivity process to again target the average size of price changes. As in Gertler and Leahy (2008) and Midrigan (2011), this extension reduces the strength of the extensive margin and increases non-neutrality as most firms hit with the γ shock will tend to adjust, independent of the aggregate nominal shock. (In addition to excess kurtosis, Appendix 1 shows that empirically, the kurtosis of price changes in the CPI database is procyclical. Allowing for a time-varying γ to match this fact would strengthen my results as it would generate additional countercyclical price-flexibility). Table VII shows the results of this model, with and without second moment shocks.

As before, the model with first moment shocks generates a counterfactual negative correlation between the standard deviation of price changes and the frequency of adjustment, in contrast to the model with second moment shocks. More importantly, aggregate price flexibility continues to rise with volatility as shown in Table VIII and Figure 7. The cumulative output impulse response estimated for October 2001 remains more than fifty percent larger than that for September 1995. The effects of time-varying price volatility on time-varying price flexibility are dampened somewhat relative to the Golosov and Lucas (2007) model, since random menu costs move the model towards Calvo, which features a constant impulse response. However, even after these model extensions, quantitative differences in price flexibility over the business cycle remain quite large. Thus, these extensions significantly amplify non-neutrality, so that they have a large effect on the average response of output to nominal shocks, but they do not change the conclusion that volatility matters for aggregate price flexibility. This is because this model features a weaker extensive margin on average, but there is still large variation in the extensive margin across time.[39]

Finally, this model is a substantially better fit for the distribution of price changes,

[39] See Gagnon et al. (2012) for separate work showing the extensive margin is important for explaining the inflation response to several important events.

both on average as well as in recessions, and it does a better job of matching variation in the frequency of adjustment over the business cycle. See Appendix 1 for additional discussion.

6.2 More Realistic Aggregate Shocks

The benchmark model features cross-sectional volatility that is perfectly negatively correlated with aggregate productivity. While there is now substantial evidence that volatility is countercyclical, there is not a perfect negative correlation. Furthermore, the baseline model generates price change dispersion that is too countercyclical. In this section, I relax the assumption that aggregate productivity is perfectly negatively correlated with volatility. I also show how my results are affected by changes in the persistence of volatility as well as the size of volatility shocks. Finally, I investigate a specification of the model with autocorrelated nominal output growth shocks. This is a more empirically realistic specification, and it allows the model to generate hump-shaped impulse response functions.

Table IX shows that while more realistic specifications for the volatility process make the inflation-output tradeoff less countercyclical, price flexibility still rises substantially during recessions. Furthermore, the relationship between volatility and price flexibility is little affected. This is unsurprising since the mechanism that generates countercyclical price flexibility in the model is countercyclical volatility rather than the business cycle per se. This is consistent with reduced form time-series evidence in the next section. Finally, given the debate in the literature over the empirical size of countercyclical volatility, in the last column of Table IX, I show that even halving the size of volatility shocks leaves significant room for volatility to matter for aggregate price flexibility.

The final model extension I investigate is a more realistic process for nominal output growth. In particular, as in Midrigan (2011), I assume that the growth rate of nominal spending follows an autoregressive process which I calibrate to have persistence 0.61 and standard deviation .0037. Again I find that the model with second moment shocks is able to match the empirical correlation between price change dispersion and frequency while the model with only first moment shocks implies a strong negative correlation. Furthermore, autocorrelated nominal output shocks actually amplify the importance of volatility. This is because when volatility increases, the wait-and-see effect reduces price flexibility on impact, and it takes one to two months for the volatility effect to push more firms to adjust. With autocorrelated nominal output shocks, the effects of an additional impulse to nominal output build up over several months so that the peak response of nominal output occurs exactly when the volatility effect is strongest. Figure 8 shows how the impulse response function varies with volatility. At the 90th percentile of volatility, the output impulse response is less than half that at the 10th percentile of volatility.

7 Time-Varying Phillips Curves

My model with second moment shocks generates countercyclical price dispersion and it is able to match the positive correlation between price dispersion and the frequency of adjustment. More importantly, it implies that the price level becomes more flexible in times of high volatility. That is, the slope of the short-run Phillips curve rises substantially so that increases in output require larger increases in inflation when volatility is high. Furthermore, this time-varying inflation-output survives under a range of model extensions.

Nevertheless, a natural question is the extent to which less structural treatments of the data imply results similar to my model.[40] In this section, I estimate the slope of the short-run Phillips curve to test the model's implications directly. In particular, I ask whether increases in marginal cost lead to greater increases in inflation during times of high volatility. To do this, I estimate forward looking Phillips curves using aggregate data as in Gali and Gertler (1999), and I indeed find that the slope of the estimated Phillips curve rises substantially with volatility.

It is well-known that a wide variety of price-setting models lead to an inflation equation of the form

$$\pi_t = \lambda mc_t + \beta E_t \left\{ \pi_{t+1} \right\},$$

and Gertler and Leahy (2008) derive a Phillips curve with this functional form in an Ss model. In the Ss model with volatility shocks, the inflation response to shocks to marginal cost varies across time, so the corresponding inflation relationship is given

[40] One can imagine a number of reduced form tests for time-varying price flexibility. Previous drafts of this paper provided evidence that federal funds rate shocks have greater effects on prices using a time-varying FAVAR. In addition, I also showed that import price pass-through is significantly larger during times of high volatility. Finally, inflation volatility is larger during periods with high cross-sectional volatility. However, all of this evidence is a less direct test of the model.

by[41]

$$\pi_t = \lambda_t mc_t + \beta E_t \{\pi_{t+1}\}. \tag{1}$$

The structural Ss model predicts that λ_t rises with volatility so that inflation responds more to marginal cost during volatile times. I test this prediction by estimating Equation 1 separately for times of high volatility and low volatility. I pick the periods of high and low volatility by first ranking all periods by the interquartile range of plant-level TFP from Bloom et al. (2012). I then call the one-third of periods with the highest value for the interquartile range high volatility periods and the one-third of periods with the lowest value for the interquartile range low volatility periods. The Bloom et al. (2012) measure of cross-sectional volatility is based on census data from 1972-2009, so my estimates cover these dates.[42]

Under rational expectations, the error in the forecast of π_{t+1} is uncorrelated with information dated t and earlier,[43] so Equation 1 can be estimated using GMM after instrumenting for $E_t \{\pi_{t+1}\}$.[44] As in Gali et al. (2005), I use the labor share of income in the nonfarm business sector to proxy for marginal cost, and I instrument for $E_t \{\pi_{t+1}\}$ using four lags of inflation and two lags of the labor income share, the output gap and wage inflation.[45] I then estimate Equation 1 using GMM separately for high and low volatility periods. Table X displays the results.

Consistent with the structural model, I find that the short-run Phillips curve is substantially steeper during times of high volatility. That is, λ is significantly greater than zero when estimated during high volatility periods while it is not significantly different from zero during periods of low volatility, and the point estimate is even slightly negative. The estimates of β are within two standard deviations of a standard value of 0.99, and reestimating the model imposing that β is equal in the high and

[41] Given the Krusell-Smith specification for price forecasts, this inflation relationship holds numerically. However, this functional form also arises (approximately) in a wide variety of pricing models including a Calvo model with time-varying probabilities of adjustment, so that it is agnostic about the underlying model of price-setting.

[42] Previous versions of this paper used an earlier "Uncertainty Index" from an older version of Bloom et al. (2012) that was available for a longer sample. It produced similar results.

[43] This will be true even if (as predicted by the model) there is serial correlation in λ_t as long as agents know λ_t when making their forecasts of future inflation.

[44] A number of papers have been written criticizing GMM in this context. See Gali, Gertler, and Lopez-Salido (2005) for a rejoinder that argues that GMM is an appropriate econometric choice.

[45] I measure inflation using the GDP deflator, wages using nominal compensation per hour in the non-farm business sector and I measure the output gap as the deviation from a quadratic GDP trend. Alternative measures did not substantively affect the results.

low volatility regimes did not substantially affect the estimates of λ (the difference in estimated slopes was more significant). While standard errors are relatively large given the short data sample (48 observations for each regression), the slope of the Phillips curve is significantly larger during the high volatility regime.

Since the CPI micro data is only available beginning in 1988, it is possible that the price dispersion facts documented in Section 2 do not hold for prior data, so that the model with volatility shocks may not be a good description of that data. I investigate this possibility by reestimating Equation 1 using only data from 1988:1-2009:4, and I find similar results. In addition to this robustness check, I have also investigated additional instruments as well as additional lags when instrumenting for expected inflation and results were similar. In addition, I have estimated backward looking Phillips curves and again the results were little affected. Splitting the sample using different thresholds of the Bloom et al. (2012) volatility measure or instead using the dispersion of price changes in CPI data as the measure of volatility also did not change the conclusions.

Since volatility is countercyclical, it is also natural to ask whether time-variation in the inflation-output tradeoff is coming from time-variation in volatility or if it is coming from other effects of the business cycle. While volatility is countercyclical, there is not a perfect negative correlation between volatility and output, so I redo the analysis conditioning on GDP rather than volatility. As before, I divide the sample into the periods with the highest one-third of the GDP gap and the lowest one-third of the GDP gap and estimate Equation 1 separately for these two samples. While the point estimate for λ is higher when GDP is below trend, the estimates are no longer statistically different. Thus, just as in the structural models, it appears that it is indeed time-varying volatility that drives time-variation in the reduced form Phillips curve rather than the business cycle, per se.

Overall, this time-series evidence shows that increases in marginal cost lead to greater increases in inflation during the volatile times that accompany recessions. This is predicted by the Ss model with volatility shocks while it is at odds with older Keynesian models that predict that the slope of the Phillips curve should be procyclical. It is also at odds with Ss models with only first moment shocks as well as with the Calvo price-setting model that imply an acyclical inflation output tradeoff. Thus, in addition to providing a better fit to microdata, the Ss model with volatility shocks also better matches aggregate inflation dynamics.

8 Conclusions

There is mounting empirical evidence that volatility rises during recessions. In the presence of adjustment frictions, this can have important implications for the transmission of aggregate shocks. Bloom et al. (2012) argues that a fall in investment following an increase in volatility may be an important source of business cycle fluctuations. Furthermore, stabilization policies to stimulate investment may have reduced effectiveness when volatility is large.

While there is a growing literature that studies the effects of volatility in real business cycle models, the implications for monetary policy have received less attention. In this paper, I argue that countercyclical volatility can explain a number of micro price-setting facts that standard price-setting models otherwise miss. Furthermore, fixed costs of price adjustment and volatility have important interactions that generate time-varying real responses to nominal shocks. During times of high volatility, firms have greater desired price changes, which in turn lead the aggregate price level to become more responsive (and output less responsive) to nominal stimulus. This means that achieving a given increase in real output requires a greater increase in inflation during times of high volatility, and I test for this directly by estimating forward looking Phillips curves using aggregate data. Consistent with the model, I find that the estimated inflation-output tradeoff is substantially worse during times of high volatility.

Increases in volatility mean that firms become more responsive on the price margin so that monetary policy of normal magnitude becomes less effective at influencing output. Since volatility rises during recessions, monetary policy faces less desirable tradeoffs at precisely the time when it is needed most. Even in the most conservative model estimates, at the height of volatility during recent recessions, nominal stimulus would have less than half the impact on real output as during the relative calm of the mid-nineties.

References

Arellano, C., Y. Bai, and P. Kehoe (2010). Financial markets and fluctuations in uncertainty. *Minneapolis Fed Paper*.

Ashley, R. and R. J. Verbrugge (2007). Mis-specification in phillips curve regressions: Quantifying frequency dependence in this relationship while allowing for feedback. *Mimeo*.

Auerbach, A. and Y. Gorodnichenko (2012). Fiscal multipliers in recession and expansion. *Fiscal Policy after the Financial Crisis*.

Bachmann, R. and C. Bayer (2011a). Investment dispersion and the business cycle. *Nber Working Paper 16861*.

Bachmann, R. and C. Bayer (2011b). Uncertainty business cycles - really? *NBER Working Paper 16862*.

Bachmann, R., R. J. Caballero, and E. M. Engel (2010). Aggregate implications of lumpy investment: New evidence and a DSGE model. *Cowles Foundation Discussion Paper*.

Bachmann, R. and G. Moscarini (2011). Business cycles and endogenous uncertainty.

Bachmann, R. and E. R. Sims (2011). Confidence and the transmission of government spending shocks. *NBER Working Paper 17063*.

Basu, S. and B. Bundick (2011). Uncertainty shocks in a model of effective demand. *Mimeo*.

Berger, D. and J. Vavra (2011). Dynamics of the U.S. price distribution. *Yale Mimeo*.

Bertola, G. and R. J. Caballero (1990). Kinked adjustment costs and aggregate dynamics. *NBER Macroeconomics Annual 5*.

Bils, M. (2009). Do higher prices for new goods reflect quality growth or inflation? *Quarterly Journal of Economics 124*(2).

Bils, M. and P. Klenow (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy 112*(5).

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica 77*(3).

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. Terry (2012). Really uncertain business cycles. *NBER Working Paper 18245*.

Caballero, R. J. and E. M. Engel (2007, September). Price stickiness in ss models: New interpretations of old results. *Journal of Monetary Economics 54*(Supp 1).

Caplin, A. S. and D. F. Spulber (1987). Menu costs and the neutrality of money. *The Quarterly Journal of Economics 102*(4).

Danziger, L. (1999). A dynamic economy with costly price adjustments. *American Economic Review 89*(4).

Den Haan, W. J. (2010). Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents. *Journal of Economic Dynamics and Control 34*(1).

Dotsey, M., R. King, and A. Wolman (1999). State-dependent pricing and the general equilibrium dynamics of money and output. *Quarterly Journal of Economics 114*(2).

Eichenbaum, M., N. Jaimovich, and S. Rebelo (2009). Reference prices, costs and nominal rigidities. *American Economic Review*.

Fernandez-Villaverde, J., J. Guerro, J. Rubio-Ramirez, and M. Uribe (2011). Risk matters: The real effects of volatility shocks. *American Economic Review 101*(6).

Gagnon, E., D. Lopez-Salido, and N. Vincent (2012). Individual price adjustment along the extensive margin. *International Finance Discussion Papers Number 1052*.

Gali, J. and M. Gertler (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics 44*(2).

Gali, J., M. Gertler, and D. Lopez-Salido (2005). Robustness of the estimates of the hybrid new keynesian phillips curve. *Journal of Monetary Economics 52*(6).

Gertler, M. and J. Leahy (2008). A phillips curve with an ss foundation. *Journal of Political Economy 116*(3).

Gilchrist, S., J. Sim, and E. Zakrajsek (2010). Uncertainty, financial frictions, and investment dynamics. *Mimeo*.

Golosov, M. and R. E. Lucas (2007). Menu costs and phillips curves. *Journal of Political Economy 115*(2).

Guimaraes, B. and K. Sheedy (2011). Sales and monetary policy. *American Economic Review 101*(2).

Hassler, J. (1996). Variations in risk and fluctuations in demand: A theoretical model. *Journal of Economic Dynamics and Control 20*(6).

Kehoe, P. J. and V. Midrigan (2008, October). Temporary price changes and the real effects of monetary policy. *NBER Working Paper 14392*.

Khan, A. and J. Thomas (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica 76*(2).

Klenow, P. and O. Kryvtsov (2008). State-dependent or time-dependent pricing: Does it matter for recent U.S. inflation? *The Quarterly Journal of Economics 123*(3).

Klenow, P. and B. Malin (2011). Microeconomic evidence on price-setting. *Handbook of Monetary Economics 3A*.

Knotek, E. S. K. and S. Terry (2008). Alternative methods of solving state-dependent pricing models. *Federal Reserve Bank of Kansas City Working Paper No. 08-10*.

Krusell, P. and A. A. Smith (1998). Income and wealth heterogeneity in the macroeconomy. *The Journal of Political Economy 106*(5).

Midrigan, V. (2011). Menu costs, multi-product firms and aggregate fluctuations. *Econometrica 79*(4).

Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics 123*(4).

Nakamura, E. and J. Steinsson (2010, August). Monetary non-neutrality in a multisector menu cost model. *Quarterly Journal of Economics 154*(4).

Stokey, N. L. (2009). *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton University Press.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters 20*(2).

Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004). Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *The Review of Economics and Statistics 86*(2).

9 Appendix 1: Empirical Results

This appendix discusses a number of robustness checks for the main empirical results of the paper and also discusses additional dynamic features of the distribution of price changes in the data and models. All correlations reported in Table I were computed using Baxter-King (18,96,33) bandpass filtered, seasonally adjusted data. The bandpass filter was chosen because it eliminates high-frequency noise in the price dispersion and frequency data. However, despite their widespread use, Ashley and Verbrugge (2007) argue that two-sided bandpass filters may produce inconsistent estimates of the frequency components of interest. Recomputing statistics using their alternative one-sided bandpass filter nevertheless produced similar results. In addition, similar results were obtained when using a moving average smoothed version of the series as well as when comparing raw correlations of series' growth rates.[46] The top-left panel of Figure 9 shows the overall distribution of price changes conditional on the state of the economy. The recession distribution shows only price changes during NBER recession months. The non-recession distribution shows the distribution of price changes for months 3-9 before and after these recessions. It is clear even in this raw, unadjusted data that there is a spreading of the distribution of price changes during recessions.

The benchmark series excludes sales and product substitutions for several reasons. Many recent papers[47] argue that the behavior of prices after excluding sales is likely to be more relevant for monetary policy. Furthermore, Bils (2009) argues that product substitutions induce a quantitatively significant source of measurement error into price series. In addition, the benchmark analysis focuses on the monthly price series collected in NY, Los Angeles and Chicago. The monthly series more accurately measures the timing of price changes relative to monthly data. Seasonal adjustment was computed using deviations from month dummies as in Klenow and Malin (2011). Finally, my benchmark results exclude zeros from the measures of price change dispersion. This is not a strong restriction since I am interested in jointly matching the frequency of adjustment and distribution of price-changes conditional on adjusting. Matching these two series implies that I will also match the distribution of price changes including the zeros. Nevertheless, the top half of Table A-I shows that my

[46] And as reported in the body of the text, similar results are delivered with regressions of these statistics on NBER Recession dummies.

[47] E.g. Eichenbaum et al. (2009), Kehoe and Midrigan (2008) and Guimaraes and Sheedy (2011).

conclusions are robust to all of these restrictions. This table shows the correlations between various statistics of interest for different sample restrictions.[48]

Both the median and mean frequency are positively correlated with the standard deviation and interquartile range of price changes across all specifications. In addition, both measures of price-change dispersion are countercyclical. The frequency of adjustment is moderately countercyclical across all specifications but it moves less than the dispersion of price changes (I will return to the implications of the model for the frequency of adjustment at the end of this section).

In principle, these results could be driven by compositional changes in the mix of items adjusting prices over the business cycle rather than being characteristic of the behavior of individual price-setters. I explore this issue in a number of ways. The bottom panel of Table A-I recomputes statistics for a number of subsamples. First, the dispersion of both price increases and price decreases is countercyclical. Thus, it does not appear that countercyclical price change dispersion is driven by a switch from price increases (which are smaller in absolute value on average) to price decreases over the business cycle.[49]

The next rows of Table A-I show that for the vast majority of product categories, price change dispersion is countercyclical and comoves with the frequency of adjustment. Across the 13 listed sub-samples, 2 measures of frequency and 2 measures of price dispersion, there are 39 significant positive relationships between dispersion and frequency out of a possible 52. In addition, there are another 10 positive but insignificant dispersion-frequency relationships. Only one product category (raw food) exhibits significant negative correlations between frequency and price change dispersion (and only for one measure of price change dispersion). In addition, using the 2 measures of price change dispersion and the 13 subsamples, there are a total of 24 out of a possible 26 significant negative correlations between dispersion and industrial production. The remaining two point estimates are also negative although only marginally significant. The frequency of adjustment within product categories

[48] In all cases I trim price changes larger than 500% as these are likely to be due to measurement error resulting from transcription errors or miscoded prices. The results are not particularly sensitive to this restriction.

[49] Ss models do not have strong implications for the relationship between frequency and dispersion after conditioning on the sign of the price change, so the fact that there is not a strong relationship between these statistics in the data is not a concern. The strong implications of Ss models are for the total frequency and total dispersion of price changes, rather than separate relationships for increases and decreases.

is also mildly countercyclical with 17 out of 26 correlations significantly negative, another 8 exhibiting an insignificant negative correlation and only one exhibiting an insignificant positive correlation. Thus, it appears that the empirical facts are not driven by compositional changes in the product categories over the business cycle and hold within various sectors.

A final concern is that product entry and exit over the business cycle may potentially be driving my conclusions (for example if there is more product entry during booms and if new products exhibit less frequent price changes). The structure of the CPI sampling limits the ability to recompute my main statistics using a balanced sample since products are rotated out of the database after a maximum of roughly four years. Nevertheless, it is possible to get a rough sense for whether this is likely to be a concern by restricting the analysis to only include items that span at least one recession. I can then compare the distribution of price changes for this balanced panel during recession months to non-recession months. More precisely, I include only items which are observed in the database from 9 months before a recession until 9 months after a recession. I then compare the distribution of price changes during the recession to the distribution 3-9 months before or after the recession. The upper right panel of Figure 9 shows that restricting to the balanced panel does not change the qualitative conclusions, so my results do not appear to be driven by product entry and exit.

While I focus mostly on the second moment of price changes, it is straightforward to compute higher moments of the distribution of price changes. Table Table A-I shows that in contrast to the second moment of price changes, there is little robust relationship between the skewness of price changes and the business cycle. There are many subsamples with significant positive relationships and many others with significant negative relationships, and a substantial fraction exhibit no significant relationship. In contrast, there does appear to be a relationship between the kurtosis (4th moment) of price changes and the business cycle. The vast majority of samples exhibit a positive correlation between kurtosis and the business cycle: there is more excess kurtosis during booms than during recessions. The model in Section 6.1 introduces poisson shocks to match the average kurtosis observed in the data, but I do not impose any exogenous time-variation in this poisson probability. Nevertheless, the model exhibits mildly procyclical kurtosis: during times of high volatility, firms that are adjusting are less likely to have a desired price change extremely close to zero. At the same time, the low persistence of idiosyncratic shocks limits the model's ability to generate very fat tails even during times of high volatility. However, this effect is quantitatively small and matching the kurtosis observed in the data would require the introduction of shocks with time-varying kurtosis. For simplicity I do not explore this extension, but it is likely to reinforce my results. Midrigan (2011) shows that greater kurtosis leads to a reduction in price flexibility. In the data, kurtosis is procyclical, so matching this fact would make price flexibility even more countercyclical.

Overall the model in 6.1 does a good job of matching the distribution of price changes, both on average and across recessions. The bottom panel of Figure 9 shows that the distribution of price changes for the model lines up well with that in the data.

Finally, I briefly explore the implications of my model for time-series variation in the frequency of adjustment and for the decompositions of inflation in Klenow and Kryvtsov (2008). Overall, the Golosov and Lucas (2007) model exhibits excess frequency volatility. Empirically, the correlation between frequency and the business cycle is only -0.28. In the benchmark Golosov and Lucas (2007) model with volatility shocks, the frequency of adjustment is too countercyclical with a correlation of -0.7. Furthermore, in the Golosov and Lucas (2007) model, the frequency of adjustment rises from 7% to 13% when moving from non-recessions to recessions while the data only exhibits a 1-2% increase in the frequency of adjustment.

While the Golosov and Lucas (2007) model implies too much variation in frequency across time, the performance of the model with random menu costs is much better. This model exhibits a correlation of only -0.23 between the frequency of adjustment and output, and during the average recession, the frequency only rises slightly. This is because the flatter adjustment hazards induced by random menu costs mean that even when the distribution of firms' desired price changes is spread out by greater volatility, there is only a small additional mass of firms that is pushed into action. Nevertheless, more firms are pushed into the upward sloping region of the hazard function, and this increases the strength of the extensive margin and raises aggregate price flexibility as shown in the body of the paper.

Klenow and Kryvtsov (2008) decomposes the variance of inflation into intensive and extensive margin contributions. Using their terminology, the extensive margin contribution to inflation is the contribution of changes across time in the frequency of adjustment to changes in inflation (which is different from the extensive margin response to an aggregate shock, as described and used in previous sections). While the latter concept is unobservable and requires a model, the former is observable, and we can decompose inflation using a Taylor expansion

$$\begin{aligned} var\left(\pi_t\right) &= var\left(dp_t fr_t\right) \\ &= \underbrace{var\left(dp_t\right) * \overline{fr}}_{IM} + \underbrace{var(fr_t) * \overline{dp}^2 + 2\overline{fr} * \overline{dp} * cov(fr_t, dp_t) + O_t}_{EM}. \end{aligned}$$

I find that in the model with random menu costs and leptokurtic shocks, the IM contributes 83% of the variance of inflation, which is in line with the IM contribution of 91% reported by Klenow and Kryvtsov (2008). (In the benchmark Golosov and Lucas (2007)) the IM contribution falls to just under 75% as volatility shocks contribute to frequency variations that are moderately larger than those in the data).

In summary, across a variety of empirical specifications, I find that countercyclical price change dispersion and a positive comovement between price change dispersion and frequency are robust features of the data. In addition to being able to match this fact, the empirical extensions of the model are able to also match higher moments of the distribution of price changes as well as the decomposition across time of inflation into changes in frequency and changes in the size of adjustment. More importantly, all the models still prescribe an important role for volatility shocks, both for matching micro price facts and for their role in the monetary transmission mechanism.

10 Appendix 2: Analytical Model

The analytical model is a standard two-sided Ss model. (See Bertola and Caballero (1990) for a review of several early examples). Time evolves continuously and firms discount payoffs at rate r. Let p_t be a firm's log nominal price at time t. p_t^* is a firm's optimal price if there are no adjustment frictions, $z_t \equiv p_t - p_t^*$ is the firm's "price gap" and p_t^* follows a Brownian motion with drift $\pi > 0$ and variance $\sigma = \sigma_A + \sigma_I$ so that the variance of the desired price has both an aggregate and idiosyncratic component. Firms pick their nominal price at time t to minimize a quadratic loss function subject to the constraint that adjusting the nominal price entails paying a fixed cost $F > 0$.

A firm's optimal policy is a two-sided Ss rule: firms raise prices by L when z_t

reaches a lower threshold $L < 0$ and lower prices by U when z_t reaches an upper threshold $U > 0$. This continuous time environment yields the following result:

Theorem 1 *Assume firms face a quadratic loss function in their deviation from the optimal price z, and that the price gap follows a Brownian motion with variance σ^2 and drift π and face fixed cost F of price changes. Then the variance of price changes (conditional on adjustment) and the frequency of adjustment are negatively correlated.*

Proof. First, note that when there is no drift, a quadratic loss function and no variable cost, the frequency of increases will be equal to the frequency of decreases. Now, as inflation increases, the frequency of increases must rise while the frequency of decreases must fall. Thus, $f_{up} > f_{down}$ with positive inflation and a constant fixed cost of price changes. Now, let our optimal policy be described by thresholds U and L with $U > 0$ and $L < 0$. Firms raise prices when their price gap $z = p - p^*$ reaches L and lower prices when z reaches U. Without loss of generality, we can normalize the optimal reset point to 0.

We now wish to show that if $f_{up} > f_{down}$ then the variance of price changes and the frequency of price changes are negatively correlated in response to first moment shocks. If we hold optimal policy constant, then without loss of generality, we can compute the response to an aggregate shock by computing how the distribution varies in response to changes in π.[50] In order to do this, we will take an optimal policy U, L as given and then compute how the frequency and variance of prices respond to inflation by computing changes in the ergodic density in response to these changes. It can be shown[51] that the ergodic density is given by:

$$f(z) = \begin{cases} A + Be^{\alpha z} & 0 < z < U \\ C + De^{\alpha z} & L < z \leq 0 \end{cases}$$

[50] In continuous time, shocking inflation while holding policy functions constant is the analogue of the discrete time first moment shock considered in Section 3.

[51] The invariant distribution must satisfy $\alpha f'(z) = f''(z) + O(\Delta_t)$ and at the boundaries we must have $f(U) = f(L) = 0$ and $\int f = 1$. Together these conditions imply the given density. See e.g. Stokey (2009) for a more formal discussion.

with

$$
\begin{aligned}
A &= \left(1-e^{\alpha L}\right) e^{\alpha U}/K \\
B &= -\left(1-e^{\alpha L}\right)/K \\
C &= -e^{\alpha L}\left(e^{\alpha U}-1\right)/K \\
D &= -\left(1-e^{\alpha U}\right)/K \\
K &= Ue^{\alpha U}(1-e^{\alpha L}) - Le^{\alpha L}(1-e^{\alpha U}),
\end{aligned}
$$

where $\alpha = -2\pi/\sigma^2$. This implies that the frequency of increases and decreases is given by

$$
\begin{aligned}
f_{up} &= \frac{\sigma^2}{2} f'(L) = -\frac{\sigma^2}{2}\alpha\left(1-e^{\alpha U}\right)e^{\alpha L}/K && (2) \\
f_{down} &= -\frac{\sigma^2}{2} f'(U) = \frac{\sigma^2}{2}\alpha\left(1-e^{\alpha L}\right)e^{\alpha U}/K. && (3)
\end{aligned}
$$

so that the total frequency of adjustment is equal to $\frac{\sigma^2}{2}\frac{\alpha}{K}\left[e^{\alpha U}-e^{\alpha L}\right]$. Differentiating with respect to α gives that $\frac{\partial(f_{up}+f_{down})}{\partial\alpha} < 0$ if $e^{\alpha U}e^{\alpha L} < \frac{e^{\alpha U}+e^{\alpha L}}{2}$. Now if $f_{up} > f_{down}$ then (2) and (3) imply $e^{\alpha U}e^{\alpha L} < \frac{e^{\alpha U}+e^{\alpha L}}{2}$. Thus if $f_{up} > f_{down}$ then $\frac{\partial(f_{up}+f_{down})}{\partial\alpha} < 0$. Since an increase in inflation decreases α, we can thus conclude that $\frac{\partial(f_{up}+f_{down})}{\partial\pi} > 0$.

It can be shown that the variance of price changes is given by

$$
\frac{f_{up}f_{down}}{\left(f_{up}+f_{down}\right)^2}\left[L-U\right]^2.
$$

For given L and U with $f_{up} > f_{down}$, this is clearly decreasing in f_{up} and increasing in f_{down}. Since agn a positive inflation realization increases f_{up} and decreases f_{down} while holding L and U constant, the variance of price changes is thus decreasing in inflation: $\frac{\partial var}{\partial\pi} < 0$. This completes the proof that the frequency of adjustment and the variance of price changes negatively comove in response to inflation shocks, in contrast to the empirical evidence. ∎

11 Appendix 3: Computational Procedure and Estimation

11.1 Computing the model

Let p be a firm's nominal price after adjustment, P be the price level, ω be the disutility of labor, C be aggregate real demand, z be a firm's productivity and θ be the elasticity of substitution. Then current real profits are given by[52]

$$\begin{aligned}\pi(p,z;\chi,a) &= \left(\frac{p}{P}-\frac{\omega C}{az}\right)\left(\frac{p}{P}\right)^{-\theta}C \\ &= \left(\frac{p/S}{P/S}-\frac{\omega C}{az}\right)\left(\frac{p/S}{P/S}\right)^{-\theta}C\end{aligned}$$

Now, note that by assumption $S=PC$. In general, the price level will depend on the current value of the aggregate shocks and the joint distribution of idiosyncratic firm states, but I conjecture that

$$\log\frac{P}{S}=\gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\,\chi,$$

with the mean price flexible price gap: $\chi_1\equiv\log\frac{P_{-1}}{S}+\log a$. This implies that

$$C=\frac{S}{P}=e^{-(\gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\chi_1)}.$$

Substituting into the profit function gives

$$\pi(p,z;\chi_1)=\left(p/S-\frac{\omega}{az}\right)(p/S)^{-\theta}\,e^{(\gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\chi_1)(\theta-2)},$$

so if we take S as given, then instead of p as the state, we can write real profits as

$$\pi(p/S,z;\chi_1)=\left(p/S-\frac{\omega}{az}\right)(p/S)^{-\theta}\,e^{(\gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\chi_1)(\theta-2)}.$$

Finally, it is straightforward to calculate transition rules for these variables. Since S follows a random walk in logs we get

[52] Note that the household labor supply problem implies that the real wage is equal to ωC.

$$\log\frac{p'}{S'}=\log\frac{p}{S}-(\mu+\varepsilon^{s}).$$

By assumption,

$$\log z'=\rho_z\log z+d_t\sigma_z\varepsilon^{z},$$

and

$$\log a'=\rho_a\log a+\sigma_a\varepsilon^{a},$$

and

$$\chi_1'=\gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\,\chi_1-(\mu+\varepsilon^{s})+\log a'.$$

Thus, we can write the firm i's value function as

$$V\left(\frac{p_{-1}}{S},z;\chi_1,a\right)=\max\left[V^{N}\left(\frac{p_{-1}}{S},z;\chi_1,a\right),V^{A}(z;\chi_1,a)\right],$$

with

$$\begin{aligned}V^{N}\left(\log\frac{p_{-1}}{S},\log z;\chi_1,\log a\right) &= \pi\left(\frac{p_{-1}}{S},z;\chi_1,a\right)\\ &\quad +E_{\varepsilon^{z},\varepsilon^{a},\varepsilon^{s}}\,QV\begin{pmatrix}\log\frac{p_{-1}}{S}-(\mu+\varepsilon^{s}),\rho_z\log z+d\,(a)\,\sigma_z\varepsilon^{z},\\ \gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\,\chi_1\\ -(\mu+\varepsilon^{s})+\rho_a\log a+\sigma_a\varepsilon^{a},\\ \rho_a\log a+\sigma_a\varepsilon^{a}\end{pmatrix}\end{aligned}$$

and

$$\begin{aligned}V^{A}(\log z;\chi_1,\log a) &= -f\omega e^{-(\gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\chi_1)}\\ &\quad+\max_{p/S}\left[\pi\left(\frac{p}{S},z;\chi_1\right)+E_{\varepsilon^{z},\varepsilon^{a},\varepsilon^{s}}\,QV\begin{pmatrix}\log\frac{p}{S}-(\mu+\varepsilon_M),\rho_z\log z+d\,(a)\,\sigma_z\varepsilon^{z};\\ \gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\,\chi_1\\ -(\mu+\varepsilon^{s})+\rho_a\log a+\sigma_a\varepsilon^{a},\\ \rho_a\log a+\sigma_a\varepsilon^{a}\end{pmatrix}\right]\end{aligned}$$

where $Q=\beta\frac{e^{-(\gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\chi_1)}}{e^{-(\gamma_0+\gamma_1\log a'+\{\gamma_2+\gamma_3\log a'\}\{\gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\chi_1-(\mu+\varepsilon^{s})+\log a'\})}}$ is the stochastic discount factor and $\omega e^{-(\gamma_0+\gamma_1\log a+[\gamma_2+\gamma_3\log a]\chi_1)}$ is the real wage.

Given this recursive representation, I then solve the problem using value function iteration on a grid. Knotek and Terry (2008) argues that discretizing fixed adjustment cost models has robustness advantages versus collocation or other interpolation methods. Nevertheless, earlier versions of my model were solved using cubic spline interpolation and the results were unchanged. The random variables are discretized using the method of Tauchen (1986). In the benchmark analysis, 171 grid points were used for the pricing grid, 21 grid points were used for the idiosyncratic productivity grid, 14 grid points were used for the χ_1 grid and 5 grid points were used for the aggregate productivity grid. Although not a state, expectations must be computed for ε^s, and it was discretized using 7 grid points. Results were unchanged when more grid points were added.

Once the model is solved for a given conjecture for $\boldsymbol{\gamma}$, a panel of 5000 firms[53] is simulated for 14,400 months[54] with a 100 month burnin. The conjectured law of motion

$$\log \frac{P}{S} = \gamma_0 + \gamma_1 \log a + [\gamma_2 + \gamma_3 \log a]\, \chi_1$$

is then updated by regressing these variables on the simulated data. The solution and simulation is then repeated until convergence. In the benchmark analysis, the standard for convergence is a less than 1% change in any of the $\boldsymbol{\gamma}$ coefficients across iterations. Higher standards of convergence did not change the qualitative results.

In addition, at the best fit parameters, I recomputed a version of the model with significantly greater precision and more thoroughly tested the accuracy of aggregate transition rules. Using the method proposed by Den Haan (2010), I computed the maximum error between the conjectured and simulated law of motion over 10,000 periods. Even over this extremely long time frame the maximum difference between aggregate variables computed using only simulation and those computed only using the conjectured law of motion is less than 0.1%, and the average error is much lower. Results suggest that forecasting errors can be made arbitrarily small by increasing grid sizes and simulations. Finally, errors in the forecasting equation are unrelated to output and to volatility in the model. None of the qualitative conclusions of the model are changed when precision is increased from the benchmark analysis.

[53] I investigated panels of up to 500,000 firms. Results were unchanged.

[54] 14,400 is 50 replications of the length of the empirical sample window.

11.2 Estimating the Model

The model is estimated using simulated method of moments. For a given set of parameters, the model is simulated for 50 replications of the same length observed in the data, and all statistics are computed by using the same procedure applied to the empirical data. Let M_i be the vector of moments for replication i. I perform a search over parameters to minimize the log squared deviation between $\sum \frac{M_i}{50}$ and M_{data}.

Once the best fit pr of parameters was identified, I then used bootstrapping to calculate standard errors and model goodness of fit. 100 bootstrap replications of length 288 were computed from the model with best fit parameters to generate $\widehat{M}_{data,1}, ..., \widehat{M}_{data,100}$. The model was then reestimated on this "fake data" to generate a new set of best fit moments and parameters, which directly yield confidence intervals for the original model.

Previous versions of this paper explored alternative estimation schemes, including estimating the size of volatility shocks by targeting time-series variation in price change dispersion. I also explored an indirect inference procedure using a GARCH model for inflation. Ultimately, all results were similar.

Table I

Dispersion Measure	Corr w/ Production	Corr w/ Ave Freq	Corr w/ Med Freq
IQR	-0.39***	0.66***	0.68***
Standard Deviation	-0.41***	0.51***	0.61***

Zeros are excluded when computing dispersion. Excluding sales and product substitutions. Log data is seasonally adjusted using 12 monthly dummies and bandpass filtered using a Baxter King (18,96,33) filter. n=222. ***=at least 1% significance.

Table II

Model with First Moment Shocks

Parameter	Estimated Value	
Elasticity of Substitution	6.8	(6.2,7.4)
Productivity Persistence	.63	(.60,.69)
Productivity Standard Deviation	.04	(.039,.041)
Fixed Cost $\times 10^3$	3.2	(2.9, 3.5)

Moment	Data	Model	
Frequency	.10	.10	(.099,.104)
Fraction Up	.65	.65	(.64,.66)
Size Up	.08	.076	(.075,.078)
Size Down	.10	.088	(.087,.089)
Correlation Dispersion and Frequency	.61	-.67	(-.69,-.62)

Bootstrapped 95% Confidence intervals in parentheses. Fixed cost is average fraction of monthly revenues paid by all firms

Table III

Model with Second Moment Shocks

Parameter	Estimated Value	
Elasticity of Substitution	8.0	(7.7, 8.6)
Productivity Persistence	.66	(.63,.72)
Productivity Standard Deviation	.0425	(.04,.045)
Fixed Cost $\times 10^3$	4.6	(3.8,5.5)

Bootstrapped 95% Confidence intervals in parentheses. Fixed cost is average fraction of monthly revenues paid by all firms

Table IV

Model Fit Comparison

Moment	Data	Only First		First + Second	
Frequency	.10	.10	(.099,.104)	.10	(.089, .117)
Fraction Up	.65	.65	(.64,.66)	.65	(.62,.67)
Size Up	.08	.076	(.075,.078)	.082	(.078, .088)
Size Down	.10	.088	(.087,.089)	.096	(.093, .101)
Correlation Dispersion and Frequency	.61	-.67	(-.69,-.62)	.66	(.50,.79)

Bootstrapped 95% Confidence intervals in parentheses

Table V

Output Impulse Response: Baseline Model

Volatility	Output IRF
10th percentile	70%
25th percentile	64%
50th percentile	56%
75th percentile	49%
90th percentile	43%

Output Impulse on impact as a percent of total nominal shock. The nominal shock is a 1 month doubling of nominal output growth from .002 to .004

Table VI

Price Impulse Response: Baseline Model

Volatility	Intensive Margin	Extensive Margin
10th percentile	5%	25%
90th percentile	15%	42%

Price Impulse contributions on impact as a percent of total nominal shock

Table VII

Model Fit Comparison: With Random Menu Cost and Kurtosis

Moment	Data	Only First	First + Second
Frequency	.10	.11	.11
Fraction Up	.65	.67	.66
Size Up	.08	.055	.055
Size Down	.10	.061	.064
Correlation Dispersion and Frequency	.61	-.58	.61

Table VIII

Output Impulse Response: With Random Menu Cost and Kurtosis

Volatility	Output IRF
10th percentile	72%
25th percentile	69%
50th percentile	67%
75th percentile	65%
90th percentile	63%

Output Impulse on impact as a percent of total nominal shock. The nominal shock is a 1 month doubling of nominal output growth from .002 to .004.

Table IX

Output Impulse Response on Impact

Business Cycle	Baseline	Match Corr	Match Corr-Pers	Small Volatility
10th Percentile Output	44%	50%	49%	50%
90th Percentile Output	68%	62%	65%	62%
90th Percentile Volatility	43%	46%	46%	49%
10th Percentile Volatility	70%	66%	66%	62%

Output Impulse on impact as a percent of total nominal shock. The nominal shock is a 1 month doubling of nominal output growth. Baseline model is the model with volatility shocks in Section 5. Match Corr keeps the persistence of volatility the same as TFP but matches the cyclicality of price dispersion. Match Corr-Pers lowers the persistence of volatility to match that of price dispersion. Small volatility halves the size of the volatility shocks in the baseline model. The first 3 columns all feature volatility shocks of the same size.

Table X

Estimated Phillips Curves

	λ	β	$\lambda_{high} - \lambda_{low}$
(1) High volatility: Full Sample	0.025**	1.07***	0.034**
	(0.011)	(0.05)	(0.0145)
(2) Low volatility: Full Sample	-0.009	1.001***	
	(0.009)	(0.02)	
(3) High volatility: Post 88	0.015*	1.09***	0.025*
	(0.0088)	(0.074)	(0.013)
(4) Low volatility: Post 88	-0.010	1.02***	
	(0.01)	(0.019)	
(5) Low GDP Gap	0.007	1.046***	0.004
	(0.0045)	(0.012)	(0.008)
(6) High GDP Gap	0.003	0.965***	
	(0.006)	(0.023)	

This table reports GMM estimates of parameters of Eq. 1. The first two rows show results for 1970:1-2009:4 while the next 2 rows show results for 1988:1-2009:4. Rows 5 and 6 split the sample using the GDP gap instead of volatility Instruments used are four lags of GDP deflator inflation, and two lags of labor income share, quadratic detrended GDP and wage inflation. A Newey-West covariance matrix with optimal lag length was used. Standard errors in parentheses.

Table A-I

Series	(1) Freq,Y	(2) Med,Y	(3) XSD,Y	(4) IQR,Y	(5) Skew,Y	(6) Kurt,Y	(7) Freq,XSD	(8) Freq,IQR	(9) Med,XSD	(10) Med,IQR
Benchmark	-.28***	-.27***	-.41***	-.39***	.26***	.28***	.51***	.66***	.61***	.68***
Bimonthly	-.06	-.10	-.47***	-.43***	.16**	.28***	.20***	.38***	.47***	.51***
In-SaleSub	-.38***	-.55***	-.54***	-.44***	-.35***	.40***	.13*	.44***	.17**	.35***
No-Seasonal	-.28***	-.29***	-.40***	-.38***	.26***	.28***	.50***	.67***	.62***	.71***
In-Zeros	-.28***	-.27***	-.36***	-.36***	.41***	.12*	.72***	.84***	.74***	.73***
Decreases	-.33***	-.58***	-.33***	-.33***	-.50***	.53***	.11	.08	.12*	.28***
Increases	.25***	.38***	-.15**	-.14**	.18***	.14**	-.20***	-.16**	.24***	.32***
Core	-.24***	-.32***	-.13*	-.43***	.16**	.24***	.10	.24***	.32***	.36***
Durable	-.09	-.24***	-.19***	-.28***	.01	.42***	.20***	.13**	.15**	.11*
Non-Durable	-.11	-.04	-.33***	-.40***	-.02	.37***	.45***	.60***	.60***	.57***
Services	-.20***	.01	-.42***	-.17**	.11	-.05	.20***	.22***	.22***	.30***
Process. Food	-.08	-.04	-.22***	-.45***	.19***	.34***	.50***	.10	.27***	.03
Raw Food	-.29***	-.26***	-.54***	-.47***	-.52***	.21***	.05	-.25***	.17**	-.31***
HH Furnish.	-.20***	-.17***	-.42***	-.30***	.23***	.31***	.23***	-.05	.46***	.38***
Apparel	-.06	-.22***	-.14**	-.32***	-.02	.43***	.32***	.18***	.59***	.71***
Transportation	-.26***	-.25***	-.11*	-.34***	.20***	.07	.32***	.14**	.33***	.16**
Recreation	-.32***	-.17**	-.22***	-.21***	-.31***	.15**	.15**	.08	.19***	.10
Other	-.28***	-.23***	-.21***	-.13**	-.12*	.12*	.29***	.05	.36***	.27***
Vehicle Fuel	-.03	-.06	-.31***	-.36***	.13*	-.03	.46***	.15**	.52***	.19***
Travel	-.30***	-.24***	-.29***	-.28***	-.12*	.09	.07	.20***	.10	.23***

Freq is average frequency of adjustment, Med is median frequency, XSD is std deviation of price changes, IQR is interquartile range, Skew is skewness, Kurt is kurtosis, Y is industrial production. Benchmark uses monthly data excluding sales, substitutions, zeros. It is seasonally adjusted using 12 monthly dummies. All-Series are similar except where noted. Bimonthly uses the full bimonthly sample. In-SaleSub includes sales and product substitutions. No-Seasonal has no seasonal adjustment. In-Zeros includes zeros. Decreases and Increases include only price changes of neg and pos sign. Core restricts to items in core CPI. Durable restricts to durable goods, non-durable restricts to non-durable goods. Sample period 1988m1-2012m1. All series are bandpass filtered using a Baxter-King(18,96,33) filter.

***=1%, **=5%, *=10%.

Figure 1: Price Changes Across Time

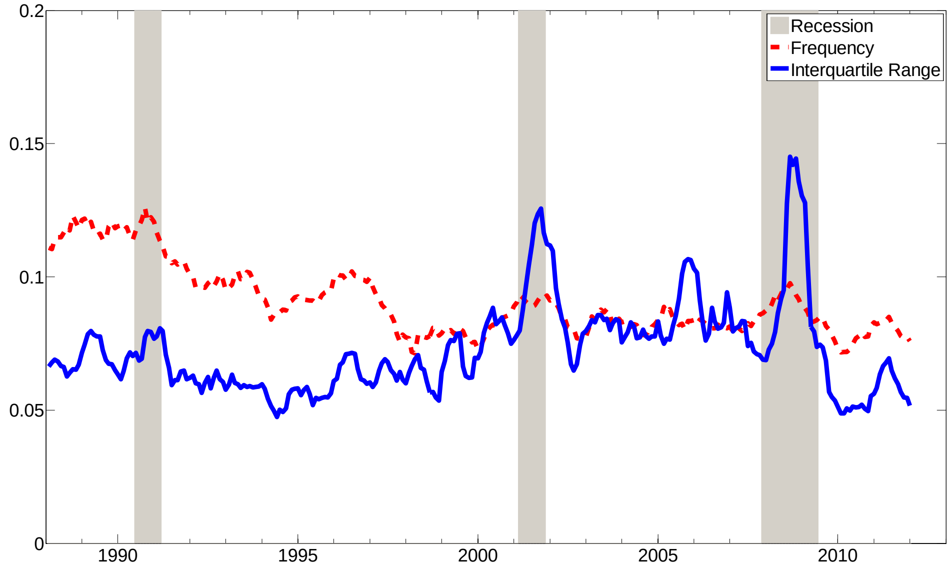

Data is seasonally adjusted using 12 monthly dummies and smoothed with a 6 month moving average.
Frequency is the Median Frequency of Adjustment.

Figure 2: Price Changes Over the Business Cycle

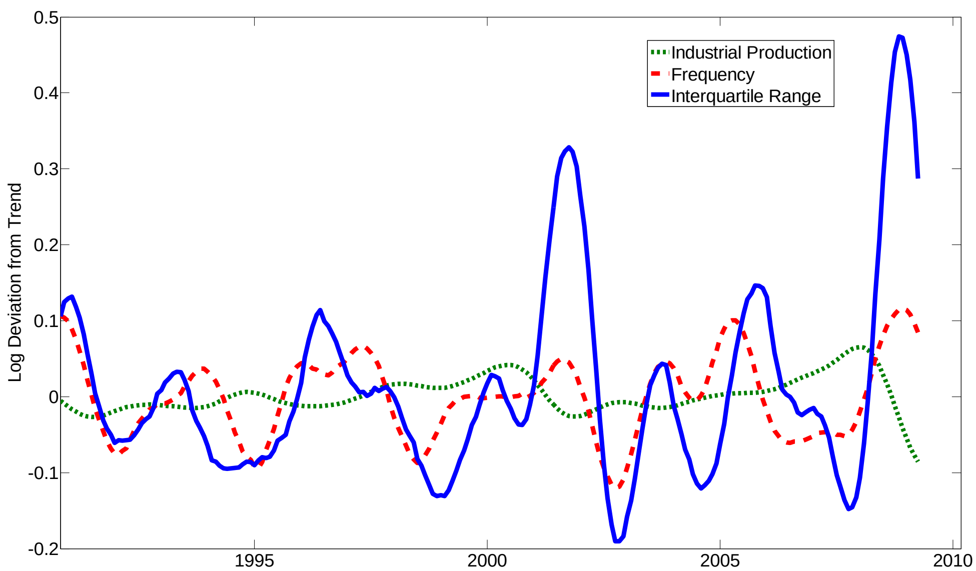

Log frequency and interquartile range are seasonally adjusted using monthly dummies. All series are bandpass filtered with a Baxter King (18,96,33) filter.

Figure 3: Negative Correlation Between Frequency and Price Change Dispersion

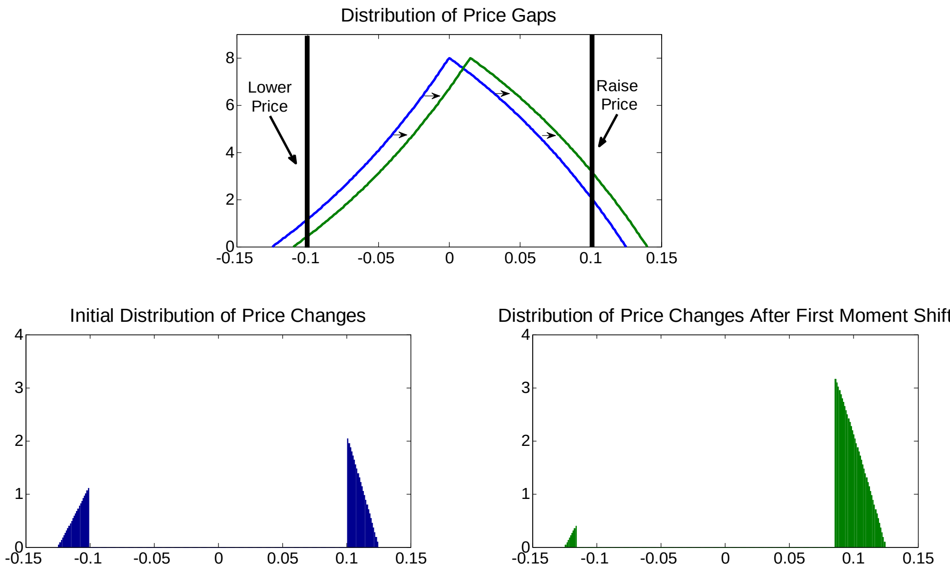

Figure 4: Response to an Increase in Volatility

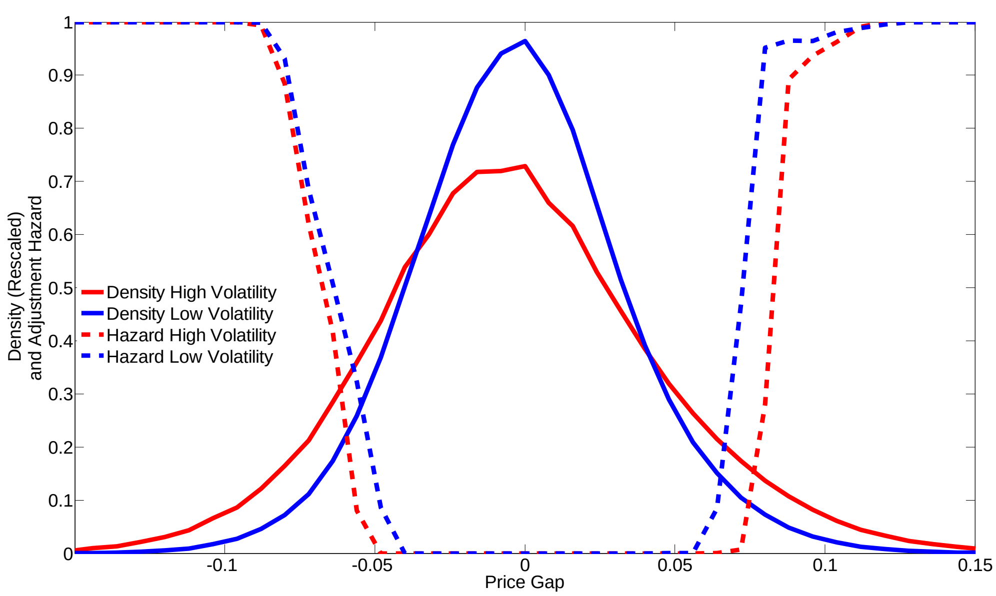

Figure 5: Output Impulse Response on Impact

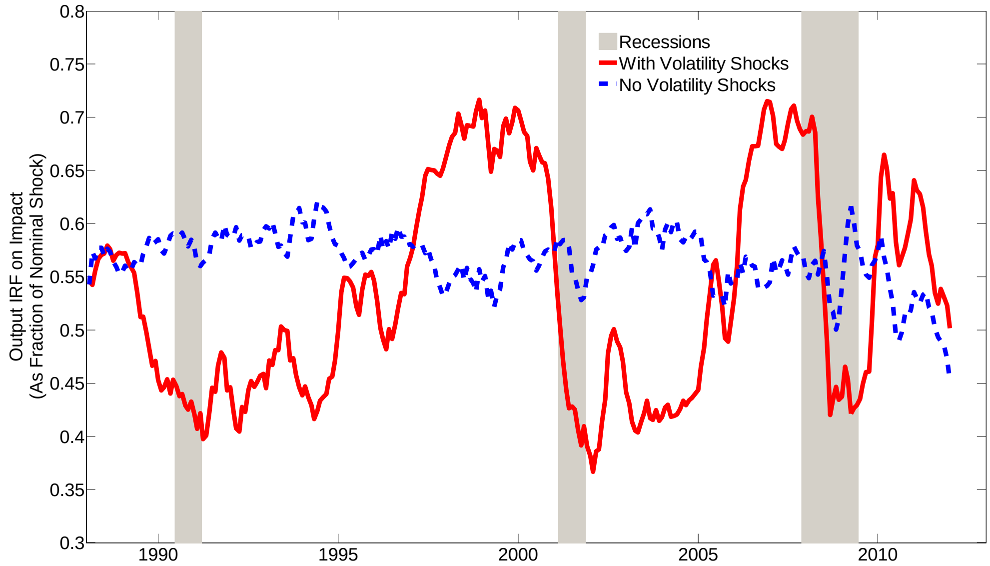

Figure 6: Real Output Impulse Response to Nominal Shock

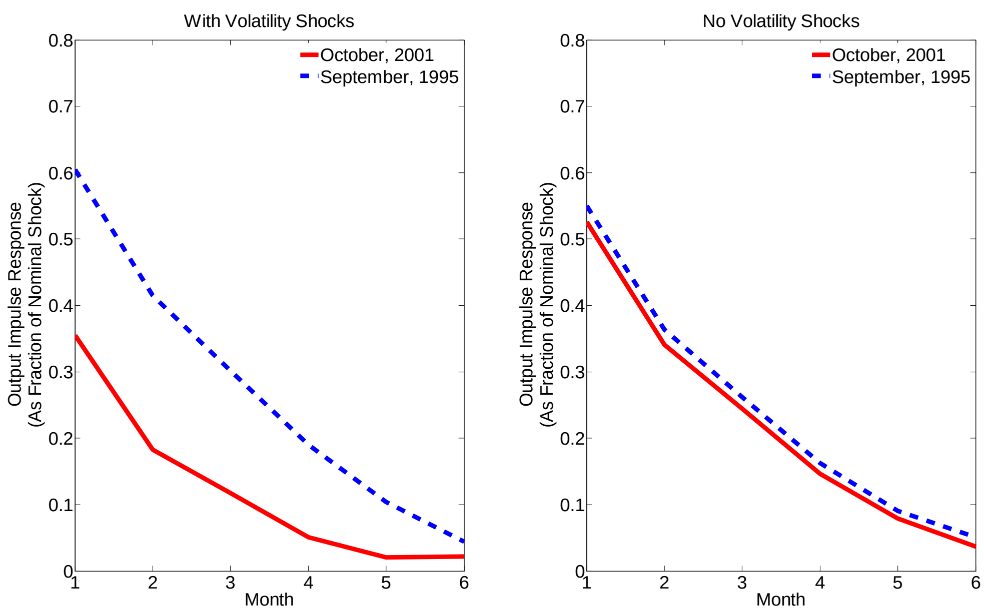

Figure 7: Real Output Impulse Response to Nominal Shock (Model with Random Fixed Cost and Leptokurtic Shocks)

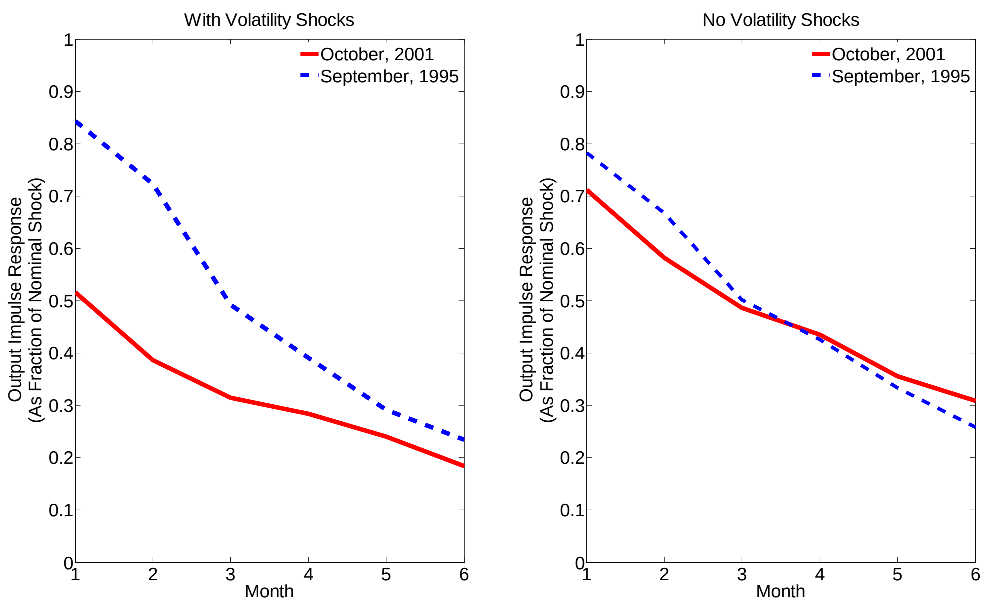

Figure 8: Real Output Response to Autocorrelated Nominal Shock

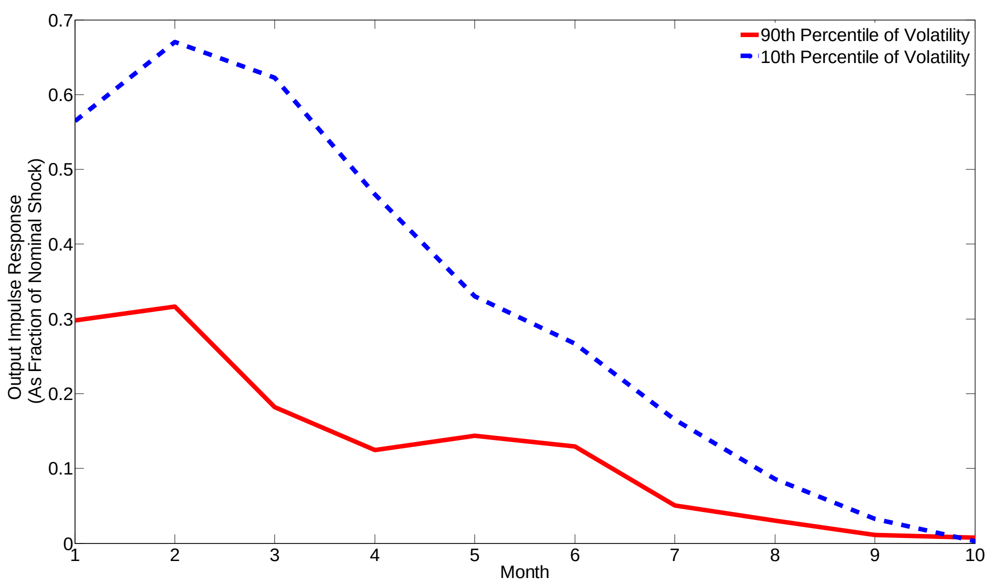

Figure 9: Price Change Distributions