

NBER WORKING PAPER SERIES

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

Joseph S. Vavra

Working Paper 19148
<http://www.nber.org/papers/w19148>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2013

I would like to thank Eduardo Engel for invaluable advice. I also thank the editor Robert Barro, three anonymous referees and my discussants John Leahy and Martin Eichenbaum. I would also like to thank Rudi Bachmann, David Berger, Nick Bloom, Francois Gourio, Erik Hurst, 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. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Joseph S. Vavra. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation
Joseph S. Vavra
NBER Working Paper No. 19148
June 2013
JEL No. D8,E10,E30,E31,E50

ABSTRACT

Is monetary policy less effective at increasing real output during periods of high volatility than during normal times? In this paper, I argue that greater volatility leads to an increase in aggregate price flexibility so that nominal stimulus mostly generates inflation rather than output growth. To do this, I construct price-setting models with "volatility shocks" and show these models match new facts in CPI micro data that standard price-setting models miss. I then show that these models imply output responds less to nominal stimulus during times of high volatility. Furthermore, since volatility is countercyclical, this implies that nominal stimulus has smaller real effects during downturns. For example, the estimated output response to additional nominal stimulus in September 1995, a time of low volatility, is 55 percent larger than the response in October 2001, a time of high volatility.

Joseph S. Vavra
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
joseph.vavra@chicagobooth.edu

1 Introduction

Is there a link between volatility and the ability of monetary policy to stimulate the economy? I argue that in price-setting models that are consistent with micro data, greater volatility leads to greater aggregate price flexibility so that nominal stimulus generates mostly inflation rather than real output growth. In addition, volatility is strongly countercyclical. This in turn leads to a reduction in the ability of monetary policy to stimulate the real economy at precisely the time it is needed most.

An explosion of micro data 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. 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 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 shocks¹ to match these new empirical facts. In particular, I focus on Ss

¹First moment shocks, such as changes in nominal output, shift all firms' flexible price by the

price-setting models because they endogenize both the frequency of adjustment and the distribution of price changes. While the simplest versions of these models are easy to reject, more realistic versions such as Midrigan (2011) have had much more empirical success.

In this paper, I show that my new facts are completely at odds with even sophisticated Ss models that only feature first moment shocks. This is because Ss models with only first moment shocks imply an extremely robust counterfactual negative correlation between price change dispersion and the frequency of adjustment. I provide intuition for this result using the analytical framework first developed in Barro (1972) and then show that this counterfactual negative correlation also holds in more realistic, quantitative versions of the model. Ss models induce an inaction region within which it is not worth paying the adjustment cost to change prices. I show that while aggregate first moment shocks can push more firms out of the inaction region and increase the frequency of adjustment, they must push all firms out of the inaction region in the same direction, which lowers the dispersion of price changes. A similar negative relationship between price change dispersion and frequency also arises in response to time-varying fixed costs or market power, so it is a quite general implication of these models. However, this robust implication is easily rejected by my new empirical evidence.²

While Ss models with aggregate first moment shocks are inconsistent with the empirical evidence, I show that a different and empirically plausible type of shock can reconcile Ss models with the data. In particular, increases in firm level volatility can simultaneously increase both the frequency of adjustment and price change dispersion. Furthermore, such time-variation in volatility is consistent with a growing empirical literature documenting that volatility rises during recessions.³

Why do "second moment" shocks that increase idiosyncratic volatility lead to positive comovement between frequency and price change dispersion? Increases in volatility have two effects: 1) An increase in volatility has a direct effect that pushes

same amount.

²In contrast, many existing tests of Ss models such as their implications for the duration hazard of price adjustment are theoretically ambiguous and have led to conflicting empirical evidence. It's also worth noting that my evidence rejects other price-setting models for more trivial reasons: the Calvo model features a constant frequency of adjustment and most pure information processing models imply that all prices are adjusted every period so the frequency of adjustment is constant at 100%.

³See Bloom (2009), Gilchrist, Sim, and Zakrajsek (2010), Bachmann and Bayer (2011b), and Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012).

more firms to adjust for a given region of inaction. 2) An increase in volatility increases the option value of waiting, which widens the size of the inaction region and decreases price adjustment. It is well-known that in response to permanent increases in volatility, the first effect typically dominates so that both frequency and price change dispersion grow. In this paper, I show that this "steady-state" intuition holds in a more realistic environment with transitory increases in volatility. Thus, Ss price-setting models with volatility shocks are able to match the new empirical evidence that standard price-setting models miss.

Beyond matching the new facts, countercyclical volatility has important implications for the nominal transmission mechanism. I show that the aggregate price level becomes much more responsive to nominal shocks during times of high volatility. This increase in price flexibility 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 very large. For example, the total response of real output to nominal stimulus in September 1995, a time of very low volatility, is 55% larger than the response in October 2001, a time of very high volatility. In contrast, in Ss models with only first moment shocks, the real response to nominal shocks is acyclical.

Why does greater volatility lead to greater aggregate price flexibility? Caballero and Engel (2007) show that 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 become more important. The frequency of adjustment rises, 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 36% larger at the 90th percentile of volatility than it is at the 10th percentile of volatility. Eighty percent of this increase is driven by a more responsive extensive margin rather than by increases in the frequency of adjustment. Thus the model delivers large time-variation in price flexibility while remaining consistent with empirically small movements in the frequency of adjustment over the business cycle.

For the majority of the quantitative sections of the paper, I focus on a version of the Golosov and Lucas (2007) model with random menu costs as in Dotsey, King, and Wolman (1999) and leptokurtic shocks as in Midrigan (2011). These model features are important for two reasons. First, Golosov and Lucas (2007) shows that the model without these features exhibits little non-neutrality. Second, as emphasized by Midrigan (2011), the Golosov and Lucas (2007) model does a poor job of matching the average distribution of price changes. Thus, my benchmark model captures the important microeconomic features of the "second-generation state-dependent" models that Klenow and Kryvtsov (2008) argue are consistent with empirical patterns of price-setting.

Nevertheless, the benchmark model is not without fault. In particular, it assumes a particularly simple aggregate shock process, which means the model generates price change dispersion that is moderately too countercyclical and fails to generate hump-shaped impulse response functions. However, I enrich the aggregate shock process and show that the insights of the benchmark 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 a wide class of price-setting models and has important implications for the monetary transmission mechanism.

The remainder of the paper is organized as follows: Section 2 contains the empirical findings. Section 3 provides context for the empirical results using the menu cost model popularized by Barro (1972) and Dixit (1991). I show that nearly all potential shocks in this model induce a counterfactual negative correlation between price change dispersion and the frequency of adjustment. Only changes in volatility are consistent with the positive correlation observed in the data. However, while this simple model provides useful intuition, it nevertheless relies on a number of very strong assumptions. Section 4 shows that the results from the analytical model remain in an empirically richer quantitative Ss model. In particular, the quantitative model with second moment volatility shocks matches the empirical positive correlation between price change dispersion and frequency while the model with only first moment shocks continues to generate a counterfactual negative correlation. Section 5 discusses policy implications of the second moment shocks. Section 6 shows results for extensions of the benchmark model that allow for richer aggregate shocks, and Section 7 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.⁴ 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 a great deal of attention, beginning with the work of Bils and Klenow (2004). While initial studies of this micro pricing data focused on static first moments such as the average frequency and size of price changes, only recently have more dynamic features of the data begun to receive attention. 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. However, there remains little empirical evidence on price-setting behavior across time.⁵

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 inflation. In particular, I will show that price change dispersion (the second moment of the price change distribution) 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 .⁶ Then, using CPI 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.

For concreteness, all dispersion numbers condition on price adjustment and thus exclude zeros. This is not an important restriction, since jointly matching the fre-

⁴Using the full sample increase the sample size from around 15,000 prices per month to around 85,000 prices per month but does not qualitatively affect the results. See Appendix 1.

⁵See Klenow and Malin (2011) for more general summaries of the recent literature utilizing this data.

⁶Since there are occasional transcription errors in decimal places, I exclude price changes of greater than 500%. In addition, the unit of price collection occasionally changes (e.g. One can vs. a six pack) which can result in spurious price changes due to rounding. I thus exclude price changes less than 0.1%. Finally, I exclude items with imputed prices and prices which are collected but not used in the construction of the CPI. My results are not sensitive to these choices.

quency of adjustment and the distribution of price changes excluding zeros necessarily implies matching the distribution of price changes including zeros. In addition, in my benchmark results I exclude sales and product substitutions.⁷ Many recent papers argue that the behavior of prices after excluding sales is likely to be more relevant for monetary policy.⁸ Furthermore, *Bils (2009)* argues that product substitutions induce a quantitatively significant source of measurement error into price series. Nevertheless, Appendix I discusses the construction of these statistics in more detail and shows that all of my results are robust to various alternative treatments of the data.

Figure 1 shows the interquartile range of price changes 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 modest rise in the frequency of adjustment during recessions. Since both series exhibit high-frequency noise and low-frequency trends, I bandpass filter the two series to leave only the variation at business cycle frequencies.⁹ Figure 2 shows that at business cycle frequencies, two clear patterns emerge. 1) The dispersion of price changes is countercyclical. 2) There is strong positive comovement between price change dispersion and the frequency of adjustment.

Tables I and II assess these patterns more formally. Table I shows the relationship between price-setting statistics and different measures of the business cycle. As suggested by Figures 1 and 2, there is a strong and significant negative correlation between price change dispersion and the business cycle. (For simplicity I report Newey-West standard errors to account for autocorrelation. However, other methods for computing standard errors produced similar results).¹⁰ The first row of Table 1 shows that the standard deviation of price changes as well as the interquartile range both exhibit large negative correlations with the growth rate of industrial production. The second row of Table I shows that this result is even stronger at business cycle frequencies. Finally, the third row of Table I regresses price-setting statistics on recession dummies. This regression shows that the standard deviation of price changes

⁷See Appendix 1 for a discussion of the sales filter.

⁸E.g. *Eichenbaum, Jaimovich, and Rebelo (2009)*, *Kehoe and Midrigan (2008)* and *Guimaraes and Sheedy (2011)*.

⁹See Appendix 1 for a discussion of alternative filtering procedures.

¹⁰Newey-West standard errors may be understated for the business cycle results because they do not account for the presence of errors induced by the bandpass filter. They can also have bad small-sample properties. Nevertheless, I have explored various bootstrapping procedures to adjust standard errors for this concern, and they did not significantly alter my results. *Prais-Winsten* and *Cochrane-Orcutt* estimation also produce similar results.

increases by almost 25% and the interquartile range by more than 40% during NBER recessions relative to the non-recession means displayed in row four.

The frequency of adjustment exhibits similar but weaker cyclical patterns. Both the average and median frequency of adjustment are modestly countercyclical although the significance varies across specifications. Finally, the kurtosis of price changes is procyclical. Figure 3 provides a visual complement to the statistical results in Table I. It compares the overall distribution of price changes during recessions to the distribution during non-recessions. Evident in this figure is the substantial reduction in small price changes during recessions so that the overall distribution is more disperse.

The final column of Table I shows that in addition to being more countercyclical, the dispersion of price changes rises more in recessions than the frequency of adjustment.¹¹ In addition, measures of price change dispersion exhibit larger coefficients of variation, so they move more across time than the frequency of adjustment. This implies that in price-setting models, movements in price flexibility should arise primarily from movements in price change dispersion rather than from movements in the frequency of adjustment.

Table II confirms the positive relationship between price change dispersion and the frequency of adjustment observed in Figure 2. Since my theoretical motivation focuses on the business cycle dynamics of monetary policy it is worth noting that while there is significant positive comovement in the raw data, the relationship is much stronger at business cycle frequencies.

In principle, countercyclical price change dispersion and positive comovement between dispersion and frequency could be driven by compositional changes in the make-up of price changes over the business cycle. There are many potential sources of compositional change over the business cycle: Price decreases are more common during recessions, and decreases are also larger and more disperse than price increases. The sectoral composition of price changes might also vary over the business cycle, and it is well known that there is pervasive heterogeneity in average price-setting behavior across sectors. Finally, the patterns I document could potentially be driven by product entry and exit across time. For brevity I leave the detailed results for Appendix I, but I do not find evidence that the aggregate relationships are driven by compositional changes.

¹¹For brevity I only report results for IQR/Med. Point estimates for SD/Med, SD/Freq and IQR/Freq are similar, but the results for SD/Med are insignificant or marginally significant depending on the particular business cycle measure.

In particular, Appendix I shows that while the exact magnitude and statistical significance varies somewhat, the aggregate results also hold within nearly all sectors, and they hold separately for both price increases and decreases. Only raw food exhibits statistically significant patterns that move in the opposite direction from the aggregate results, and even then, not all specifications are reversed.¹² In addition to these individual sectoral results, I have also run panel regressions with sector level fixed effects, and they produced similar results. In addition to these sector level results, following Klenow and Kryvtsov (2008) I have recomputed all results using "standardized price changes" that subtract the mean price change within each ELI and divide by the average standard deviation of price changes within an ELI. These results eliminate the possibility that my results are driven by a few sectors with particularly volatile prices.

Finally, 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 change dispersion and the positive relationship between price change 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.

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 Analytical Results

In this section, I use the classic analytical menu cost model popularized by Barro (1972) and Dixit (1991) to frame the empirical results in the previous section. In particular, I will show that this model has strong predictions for the relationship between the frequency of adjustment and the dispersion of price changes, and that only changes in volatility are consistent with the empirical results.

Firms face a dynamic control problem with state variable z . Departures of z from

¹²Furthermore, as I will show in the following sections, even a "zero" correlation between price change dispersion and frequency would still be strong evidence against Ss models with only first moment shocks. These models will predict a strong negative correlation, for which there is no empirical support.

zero entail a flow cost of bz^2 and firms discount payoffs at rate r . When not adjusting, z follows a Brownian motion. Assume for now that there is no drift so that $dz = \sigma dw$, where dw is the standard Wiener process. Firms must pay fixed cost k to adjust the value of z . This well-studied environment gives rise to a simple, symmetric Ss rule.

Proposition 1 *The optimal policy is no adjustment when $|z| < S$ where $S = \left(\frac{6k\sigma^2}{b}\right)^{1/4}$, and adjusting z to zero when $|z| \geq S$. In addition, the frequency of adjustment is given by $\left(\frac{b}{6k}\right)^{1/2} \sigma$ and the standard deviation of price changes is equal to S .*

Proof. See Barro (1972) for the first proof of these results (in an environment with monopolistic competition rather than a quadratic loss function). Dixit (1991) and Alvarez and Lippi (2012) provide proofs for the simplified quadratic loss function. The outline of the proofs is reproduced in Appendix 2. ■

What does this imply for the relationship between the frequency and standard deviation of price changes? In response to increases in the fixed cost k , the frequency of adjustment is falling while the standard deviation of price changes is rising. Thus, changes in fixed costs induce a counterfactual negative relationship between price change dispersion and frequency. Similarly, declines in the cost of deviating from the optimal price, b , also induce a fall in frequency and a rise in the dispersion of price changes.¹³ Both parameter changes result in a larger "inaction" region which increases the dispersion of price changes but reduces their frequency.

In contrast, increasing the volatility of shocks, σ , leads to a simultaneous increase in both the standard deviation of price changes and the frequency of adjustment, as in the data. In the presence of fixed adjustment costs, an increase in volatility increases the size of the inaction region through a real options effect, which leads to an increase in the standard deviation of price changes and a decline in the frequency of adjustment. However, the direct effect of greater volatility means that firms will hit adjustment bands of a given width more frequently. The two effects work in opposite direction, but Proposition 1 shows that in steady-state, the direct effect of greater volatility dominates so that frequency rises. A key question of the quantitative exercises that follow will be assessing whether this simple result holds out of steady-state in an environment with transitory shocks to volatility.

To provide additional intuition, Figure 4 shows the effects of an increase in volatility in a discrete time version of the model. The upper panel shows the distribution

¹³This parameter can capture, e.g. firm market power.

of "price gaps", the difference between a firm's current price and the price it would choose if it adjusts as well as the firms' inaction region. The bottom panel shows the implied distribution of observed price changes. Increases in volatility spread out the distribution of price gaps and also increase the size of the inaction region. Figure 4 shows a case where the increase in the inaction region is small relative to the increase in volatility so that more firms adjust and both dispersion and frequency rise. However, if the effect on the inaction region was stronger, then the dispersion of price changes would still rise but frequency might fall. Again, which of these two effects dominates in response to transitory changes in volatility is a quantitative question.

The above analysis shows that among the free parameters of the model, only volatility is a promising candidate for explaining the positive empirical relationship between frequency and price change dispersion. However, this model is explicitly built to explain how firms respond to shocks to z . Can aggregate shocks to z generate a positive relationship between variance and frequency? I now show that they cannot.

Proposition 2 *In an environment with zero inflation, small aggregate shocks to z do not change either the frequency or standard deviation of price changes. In an environment with positive trend inflation, small aggregate increases in z raise frequency and decrease the standard deviation of price changes.*

Proof. See Appendix 2. ■

I leave the proof to Appendix 2 and instead provide intuition with Figure 5, which shows the effects of positive first moment shock in a discrete time setting. In order to increase the frequency of adjustment in an Ss model, more firms must be pushed out of the inaction region. However, aggregate first moment shocks will, by definition, affect all firms' desired price changes in the same way. Thus, firms must all be pushed out of the inaction region 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.

The upper panel of Figure 5 shows the distribution of price gaps.¹⁴ After a first moment shock that increases firms' desired prices, 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

¹⁴The distribution will be skewed in the presence of any positive trend inflation.

of firms decreasing prices. At the same time, the variance of price changes falls. Thus, there is a negative correlation between the frequency of adjustment and price change dispersion.

While the analytical results are quite strong, they rely on a model with many simplifying assumptions. Most importantly, the comparative static result that greater volatility leads to greater frequency of adjustment need only hold in steady-state after a permanent increase in volatility. It is not obvious if this result will carry forward to a more realistic environment with transitory shocks to volatility. In addition, in the simple model, all price changes are of size $\pm S$, in contrast to the disperse distribution of price changes in the data. Finally, since the simple model is partial equilibrium, it is also less suited for studying time-variation in aggregate price flexibility.¹⁵ I now move to a more empirically realistic Ss model, which I can take seriously to micro data. Nevertheless, I will show that the result from the analytical model still holds: Ss models with only first moment shocks imply a counterfactual negative correlation between price change frequency and dispersion while Ss models with second moment shocks are able to match the empirical positive correlation.

4 Quantitative Ss models

The baseline quantitative model follows Golosov and Lucas (2007), with the addition of several features emphasized by Midrigan (2011). The 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.

4.1 Households

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

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

¹⁵Capturing interactions between endogenous aggregate price flexibility and individual firm's pricing decisions requires general equilibrium.

subject to

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

where

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

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-Debreu state-contingent claims is 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.

4.2 Firms

Firms produce output using a linear technology in labor

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

I discuss the evolution of idiosyncratic productivity z_t^i in the following sub-section. 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 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

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

Nominal aggregate spending $S_t = P_t C_t$ follows a random walk with drift in logs:¹⁶

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

Firms must pay a fixed cost f in units of labor in order to adjust their nominal price. To match the presence of small price changes, I assume that $f = 0$ with some *iid* probability each period.¹⁷ 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 flow profits are given by

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

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

4.3 Idiosyncratic Productivity

Firm i 's idiosyncratic productivity z_t^i evolves according to an AR process in logs, but following Gertler and Leahy (2008) and Midrigan (2011), I assume that shocks to firm i 's idiosyncratic productivity z_t^i arrive infrequently, according to a Poisson process

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

This specification nests that in Golosov and Lucas (2007) but allows the model to better match the fat-tailed empirical distribution of price changes. Motivated by the analytical model in Section 3 as well as mounting empirical evidence that volatility moves across time, the standard deviation of firm level idiosyncratic productivity shocks d_t itself evolves as

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

¹⁶This is a computational simplification that reduces the aggregate state-space by one dimension. In Section 6 I relax this assumption.

¹⁷This is basically a reduced form for the multiproduct firms in Midrigan (2011).

Thus, firms face idiosyncratic shocks with common standard deviation d_t , and this standard deviation is itself time-varying. 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.¹⁸

For computational simplicity, I initially make the assumption that aggregate productivity and d_t are perfectly negatively correlated. 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. The true correlation between first and second moment shocks is likely to be negative but not perfectly so. For example, using German micro data, Bachmann and Bayer (2011b) finds that the cross-sectional standard deviation of firm-level Solow residuals is strongly counter-cyclical with a correlation of -.48 with detrended output. In addition, a number of recent papers such as Bloom et al. (2012) and Gilchrist et al. (2010) build models where increases in volatility endogenously lead to declines in aggregate productivity.

The assumption that aggregate productivity and volatility are perfectly negatively correlated captures this empirical and theoretical relationship in the simplest way possible. In addition, in Section 6 I show that more empirically realistic specifications for the relationship between productivity and volatility barely change my conclusions, and they further complicate the model.

4.4 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 firm's idiosyncratic states are given by its previous nominal price p_{t-1}^i and its current level of productivity z_t^i . Under the simplifying assumption that volatility is perfectly negatively correlated with aggregate productivity, the aggregate state of the economy can be summarized by the current

¹⁸For example Bachmann and Moscarini (2011) or countercyclical volatility of demand.

level of nominal spending S_t , the value of aggregate productivity a_t , and the joint distribution of idiosyncratic states $\phi(p_{t-1}^i, z_t^i, f^i)$. 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:

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

where $\chi_{1,t} = \log \frac{P_{t-1}}{S_t} + \log a_t$. Given this conjecture, I then search for a value of the transition coefficients, γ 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. 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

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

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

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

and

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

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.

4.5 Calibration 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. 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.¹⁹

I calibrate the size of volatility shocks σ_d using external evidence on movements in 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 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 et al. (2012) comes from *plant level* manufacturing data. Since the two measures of volatility may not coincide, and since Bachmann and Bayer (2011b) argues 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 5.

The remaining parameters are calibrated to match various micro data moments.

¹⁹As 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.

There are six remaining parameters: the elasticity of substitution θ , the persistence and standard deviation of idiosyncratic productivity, ρ_z and σ_z , the fixed cost f , the probability of being hit with a productivity shock p^z , and the probability of drawing a fixed cost of zero. These parameters are selected to fit six 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, the fraction of small price changes²⁰, and the excess kurtosis of the price change distribution. For more details on the calibration procedure as well as a discussion of alternative moments and estimation schemes, see Appendix 3.

Table III shows the model's best fit parameters and Table IV the resulting moments. The parameters of the model are in line with recent literature. The elasticity of substitution of 6.8 is between the values used by Golosov and Lucas (2007) and Nakamura and Steinsson (2010). The specification for the fixed cost implies that total adjustment costs in the economy represent just under 0.5% of steady-state monthly revenues.²¹ The persistence of productivity is relatively low, with a monthly persistence of 0.62. 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 can match the frequency of adjustment and fraction of price changes that are increases. The model also does a good job of matching the distribution of price changes, although it slightly underpredicts the average size of price changes.²² A main reason for using Ss models (with the features emphasized by Midrigan (2011)) is their ability to endogenously match these micro moments.

Most importantly, the model with second moment shocks implies a correlation between the frequency and standard deviation of price changes that is positive and closely in line with the empirical data. This confirms that the intuition from the steady-state model in Section 3 extends to more realistic transitory shocks to volatility. Even with these transitory shocks, the direct effect of greater volatility dominates the real options effect so that the frequency of adjustment rises with volatility.

²⁰I define a small price change as $|dp_{i,t}| < \frac{1}{2} \text{mean}|dp|$

²¹This measure of the fixed cost is given by $f * (\text{freq} - p_{f=0}) * \frac{\theta-1}{\theta} / Y_{ss}$. The cost conditional on adjustment is around 4% of revenues, which is in line with the estimates in Zbaracki, Ritson, Levy, Dutta, and Bergen (2004).

²²Differences between average size in the model and data are largely driven by extreme outliers in the data. The model does a good job of replicating all but the most extreme percentiles of the price change distribution.

In addition, the model also does a good job of replicating the fact that price change dispersion moves more with volatility than the frequency of adjustment. This is driven by the behavior of two sets of firms. 1) Firms with very large price gaps will always adjust, regardless of the level of volatility. However, as volatility rises, the variance of these price changes rises. Thus, these firms contribute to greater price change dispersion but not to greater frequency. 2) Similarly, the random fixed costs of adjustment mean that some firms will adjust regardless of their price gap. Greater volatility increases the variance of these price changes without increasing their frequency. Appendix I and Figure 3 provide additional evidence that the model is a good fit for the distribution of price changes more generally across time.

Finally, it is interesting to note that the interaction between the volatility effect and the real options 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 micro data. The overall frequency of adjustment is countercyclical, but this hides compositional differences: the frequency of large price changes is strongly countercyclical while the frequency of small price changes is actually procyclical.²³

I next show that the model without second moment shocks fails to match the empirical evidence. The third column of Table IV shows results for the model after setting $\sigma_d \equiv 0$. 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 dispersion of price changes. This is in contrast to the strong empirical positive correlation.²⁴

In addition, the model with only first moment shocks generates essentially no correlation between other moments of the distribution of price changes and the business cycle. This is because recessions in the model tend to occur when both productivity and nominal demand fall, but these aggregate shocks have offsetting effects on firms' desired prices so the distribution of price changes is essentially acyclical. While it is possible to get greater movements in the distribution of price changes by amplifying the size of one of the first moment shocks, this comes at the cost of an even worse fit for the correlation between frequency and dispersion. The inability to match the pos-

²³I use a cutoff of 5% as the threshold between large and small and get the following numbers: the correlation of large changes with industrial production is -0.39 while the corresponding correlation with small price changes is 0.4. The correlation of total frequency with industrial production is -0.28. I arrive at similar numbers using a threshold of "small" price changes of 2.5% and 1%.

²⁴Previous versions of this paper explicitly targeted the empirical correlation, and no configuration of parameters generates a positive correlation for the model with only first moment shocks.

itive correlation between dispersion and variance is an extremely robust implication of versions of the model with only first moment shocks.

5 Policy Implications

I now show that time-varying volatility has important implications for the transmission of nominal shocks to the real economy. In times of high volatility, the real effects of nominal shocks are substantially reduced. Table V shows the response of log real output and inflation to a permanent increase in log nominal output of .002, computed for different levels of volatility.²⁵ This shock corresponds to a one-month doubling of the nominal output growth rate.

At the 90th percentile of volatility, a permanent increase in log nominal output of .002 leads to an increase in log real output *on impact* of .00121. Thus, 60.7% of the nominal shock goes into output on impact. At the 10th percentile of volatility, the real output response to the same nominal shock is .00142 so that 70.1% of the shock goes into real output on impact. This implies that when volatility is reduced from the 90th percentile to the 10th percentile the real effect on impact of the same nominal shock is increased by roughly 20%.

In addition to these differences on impact, the real effects of nominal shocks become more persistent when volatility falls: the half-life of the real response is increased by about 20% when reducing volatility from the 90th to the 10th percentile. This means that the cumulative effects on real output are more sensitive to volatility than are the impact effects. At the 90th percentile of volatility, the cumulative increase in log real output to a shock to log nominal output of size .002 is .0065 so that real output is increased by 325% more than the nominal shock. At the 10th percentile of volatility, the cumulative increase in log real output to the same nominal shock is .00836 so that real output is increased by 418% more than the nominal shock. Thus, the cumulative output impulse response increases by 29% when moving from the 90th to the 10th percentile of volatility.

²⁵In an Ss environment, the size of the nominal shock is not innocuous. Large changes in S will cause more firms to adjust, increasing $\frac{\Delta\pi}{\Delta S}$. Thus, the level of the output IRF cannot be defined independently of the size of the change in S . However, the random menu costs in my model move it closer towards Calvo and reduce the importance of shock size. Within a reasonable range, changing the size of ΔS has only modest effects on $\frac{\Delta\pi}{\Delta S}$ and does not alter my conclusions about the relationship between $\frac{\Delta\pi}{\Delta S}$ and volatility. For concreteness, all results in this section are reported for $\Delta S = .002$, comparable to a one month doubling of the nominal output growth rate.

It is also informative to look at how the *price response* on impact differs with volatility. Since $S = PY$, the fraction of the nominal shock that does not generate real output growth instead generates inflation, so that the tradeoff between inflation and output worsens as volatility increases. In particular, when volatility is large, getting the same increase in real output requires much greater nominal stimulus, and more of this increase in nominal output goes into P rather than Y . I find that getting the same increase in real output on impact at the 90th percentile of volatility as at the 10th percentile of volatility requires 60% more inflation.²⁶

Why does aggregate price flexibility increase with volatility? To provide intuition for this result, it is helpful to introduce the framework developed in Caballero and Engel (2007). First, I present their result that the price response on impact in Ss models can be decomposed into two distinct margins.²⁷ I then show how these margins interact with volatility.

Let $x_t = \log(p_t^*/p_{t-1})$ be the difference between a firm's current price and the price it will choose if it adjusts. Let the economy-wide distribution of price gaps be given by $f(x)$, and assume that firms have an adjustment probability that is increasing in their price gap $\Lambda(x)$. This implies that inflation will be given by

$$\pi = \int x \Lambda(x) f(x) dx.$$

If there is a positive shock ΔS to firms' desired prices, inflation will be given by

$$\pi(\Delta S) = \int (x + \Delta S) \Lambda(x + \Delta S) f(x) dx.$$

Taking a first-order Taylor approximation of $\pi(\Delta S)$ around $\Delta S = 0$, rearranging and taking the limit as $\Delta S \rightarrow 0$ gives that the price response on impact is

$$\lim_{\Delta S \rightarrow 0} \frac{\Delta \pi}{\Delta S} = \underbrace{\int \Lambda(x) f(x) dx}_{\text{Intensive Margin} \equiv \text{Freq}} + \underbrace{\int x \Lambda'(x) f(x) dx}_{\text{Extensive Margin}}.$$

This price response on impact is the sum of two components: 1) The intensive

²⁶In an Ss environment, the required inflation change is very mildly sensitive to the size of the baseline shock (see previous footnote). But the required change in inflation is increasing in the size of the nominal shock over quantitatively relevant ranges, so the reported results are conservative.

²⁷While their analytical environment is not identical to mine, it is quite similar and I have verified that these formulas provide extremely close approximations to my quantitative results.

margin gives the inflation contribution of items whose prices would have adjusted without the aggregate shock. These firms adjust to the aggregate shock by changing the size of their adjustment. 2) The extensive margin gives the additional inflation contribution of firms whose decision to adjust is either triggered or canceled by the aggregate shock.

When will each of these margins be more important? Inspecting the expression for the intensive margin shows that this component is equal to the frequency of adjustment. The more firms that are adjusting absent the aggregate shock, the greater the aggregate price response to that shock through the intensive margin.

The extensive margin grows with the number of firms near the margin of adjustment (firms with large $\Lambda'(x)$). In addition, the extensive margin is amplified if firms near the margin of adjustment also have large values of $|x|$: if the difference between adjusting and not adjusting grows, then triggering firms to switch their adjustment decisions will have a bigger effect on the overall price level.

Table VI decomposes the price responses shown in Table V into the intensive and extensive margin contributions and shows that during times of high volatility, both margins become more important on impact.²⁸ Adding both margins we again see that the price response increases from 28.9% to 39.3% as volatility increases. The intensive margin contribution grows because greater volatility increases the frequency of adjustment. The extensive margin grows because greater volatility simultaneously leads to a wider inaction region and pushes more firms near the margin of adjustment. That is, greater volatility means there are more firms with both large $\Lambda'(x)$ and large desired price gaps. While both margins become more important as volatility rises, the increase in the extensive margin is substantially larger than the increase in the intensive margin and accounts for four-fifths of the overall increase in price flexibility. Thus, even with small and empirically plausible movements in the frequency of adjustment, the model still delivers large variation in price flexibility across time.²⁹

Second moment volatility shocks are the driving feature of this time-varying re-

²⁸While the Caballero and Engel (2007) results provides an extremely good approximation for the price response on impact, the cumulative price response is substantially more complicated and cannot be characterized analytically. Nevertheless, similar intuition should apply in periods after the initial nominal shock.

²⁹Since the intensive margin is equal to the frequency of adjustment, it is easy to measure empirically. Unfortunately, no such empirical counterpart exists for the extensive margin. Nevertheless, the model is consistent with empirical movements in the distribution of price changes across time, which suggests that it should produce reasonable movements in the extensive margin.

sponse. In the model with only first moment shocks, there is no relationship between the real impact of nominal shocks and the business cycle. 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 continuing to rise substantially with volatility.³⁰

Finally, it is instructive to map the model mechanism more directly to specific episodes in the U.S. economy. Following Bachmann, Caballero, and Engel (2010), I can back out aggregate shocks from the model to fit U.S. time-series data and then compute how the output response to nominal shocks varies across time.³¹ Given the sequence of aggregate shocks, I can then calculate the output impulse response for the simulated U.S. economy to the same .002 log nominal output shock at each date from 1988-2012. Figure 6 shows the 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 September 1995 is approximately 55% larger than the response in October 2001. In contrast, there are essentially no differences for the model with only first moment shocks.

6 More Realistic Aggregate Shocks

The benchmark model illustrates the quantitative importance of volatility in a setting with simple aggregate shocks. This section argues that my results are robust to more empirically realistic specifications. I relax the strong assumptions on the evolution of volatility and then consider a more realistic process for nominal output shocks.³²

My baseline model features cross-sectional volatility that is perfectly negatively correlated with aggregate productivity. While there is now substantial evidence

³⁰It is important to note that when there are only second moment shocks in the model, price change dispersion becomes procyclical. However, as noted in the previous section, this specification is inconsistent with the strong empirical and theoretical negative relationship between productivity and volatility.

³¹In order to compute the sequence of shocks 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.

³²For brevity, I concentrate on the implications of more realistic aggregate shocks for time-varying price flexibility. In all cases, the models without second moment shocks continue to generate counterfactual negative relationships between dispersion and frequency while the models with second moment shocks generate more realistic positive correlations.

that volatility is countercyclical, there is not a perfect negative correlation in the data. Furthermore, the baseline model generates price change dispersion that is too countercyclical. In particular, the benchmark model implies a correlation between price change dispersion and output of -0.6 while in the data this correlation is only -0.41. In this section, I relax the assumption that aggregate productivity is perfectly negatively correlated with volatility to assess the robustness of my results. Instead I assume that volatility is the sum of two components, the first of which is perfectly negatively correlated with productivity, and the second of which is orthogonal to productivity. By picking the relative variance of these two components, I can match the empirical cyclical volatility. In addition, once I relax the perfect negative correlation with productivity, I can also study the effects of volatility persistence.

Table VII summarizes how price flexibility varies with the business cycle and with volatility in these models. The first row shows how much price flexibility increases during recessions (as we move from the 90th to the 10th percentile of output). The second row shows how much price flexibility increases as we move from the 10th percentile of volatility to the 90th percentile of volatility. For example, for the benchmark model the earlier results in Table V show that the price response on impact is 5.78×10^{-4} ($= .289 \times .002$) at the 10th percentile of volatility and 7.86×10^{-4} ($= .393 \times .002$) at the 90th percentile of volatility. Thus, there is a 36% increase in the price response on impact as volatility moves from the 10th to the 90th percentile, which is summarized in Table VII.

The first thing to note in Table VII is that in all cases, price flexibility varies more with volatility than it does with the business cycle. This is because it is time-varying volatility that generates greater price flexibility rather than the business cycle per se. That the model generates countercyclical price flexibility is only because volatility is negatively correlated with productivity.³³

Matching the empirical correlation between volatility and productivity moderately reduces cyclical price flexibility, but significant variation remains. Similarly, reducing the persistence of volatility only mildly dampens time-variation in price flexibility. Finally, given the debate in the literature over the empirical size of volatility shocks, in the last column of Table VII, I show that time-variation in price flexibility is roughly linear in the size of volatility shocks. Halving the size of volatility shocks roughly

³³Note that even in the baseline case when volatility is perfectly negatively correlated with productivity, volatility is not perfectly countercyclical since there are also aggregate nominal shocks.

halves the time-variation in price flexibility. But even at this overly conservative lower bound, I still find substantial time-variation in the inflation-output tradeoff. Generating the same short-run increase in output at the 90th percentile of volatility requires 25% more inflation than at the 10th percentile of volatility.

In addition to volatility that is perfectly negatively correlated with productivity, my benchmark model assumes that nominal output follows a random walk with drift. I now extend the model to allow for autocorrelated nominal output growth shocks. This is a more empirically realistic specification, and it allows the model to generate hump-shaped impulse response functions. 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.

Figure 7 shows that the model with autocorrelated nominal output growth shocks is able to generate an empirically realistic hump-shaped output IRF. Furthermore, autocorrelated nominal output shocks amplify the importance of volatility. This is because when volatility increases, the real options effect reduces price flexibility on impact as the inaction region widens, and it takes one to two months for the volatility effect to push more firms to adjust. With autocorrelated nominal output shocks, the effect of an additional impulse to nominal output builds up over several months so that the peak response of nominal output occurs exactly when the volatility effect is strongest. The cumulative output IRF is now 45% larger at the 10th percentile of volatility than at the 90th percentile of volatility relative to the 29% increase in the baseline model.

Thus, more realistic nominal output shocks amplify time-varying price flexibility relative to the baseline model while more realistic volatility shocks moderately dampen this time-variation. This suggests that a model with both aggregate shock extensions would deliver large time-variation similar to the baseline model.³⁴

7 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

³⁴Solving the model with both aggregate shocks simultaneously would two state variables to the baseline problem and complicate the Krusell-Smith forecasting rule, so solving such a model would be substantially more complicated.

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.

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 much worse tradeoffs at precisely the time when it is needed most. Nominal stimulus would have 50% greater impacts on real output during the calm of the mid-nineties as at the height of volatility during recent recessions.

References

- Alvarez, F., M. Gonzalez-Rozada, A. Neumeyer, and M. Beraja (2011). From hyperinflation to stable prices: Argentina's evidence on menu cost models.
- Alvarez, F., L. Guiso, and F. Lippi (2011). Optimal price setting with observation and menu costs. *Quarterly Journal of Economics* 126(4).
- Alvarez, F. and F. Lippi (2012). Price setting with menu cost for multi-product firms.
- Ashley, R. and R. J. Verbrugge (2007). Mis-specification in phillips curve regressions: Quantifying frequency dependence in this relationship while allowing for feedback. *Mimeo*.
- 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.
- Barro, R. (1972). A theory of monopolistic price adjustment. *Review of Economic Studies* 39(1).
- 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).
- Dixit, A. (1991). Analytic approximations in models of hysteresis. *Review of Economic Studies* 58(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*.
- 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 multi-sector 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).

8 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. The correlations in row 2 of 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 and regressions on recession dummies (the latter two relationships are shown in Table 1).

In addition to these statistical relationships, the top-left panel of Figure 3 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 3-9 months before and 3-9 months 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.

For the reasons addressed in the body of the paper, I exclude sales and product substitutions from my benchmark results. To identify sales I use a 2-step approach as in Nakamura and Steinsson (2008). I first exclude all price changes marked as sales by BLS price collectors. However, many sales prices are not flagged by price collectors, so I then also exclude "v-shaped" price changes of less than 3 months. While I exclude these price changes in my benchmark results, the fourth row of Table A-I shows that all of my results are robust to using posted prices.

In my benchmark results I also focus on the monthly data available in the top 3 metro areas. Using monthly data reduces measurement error in the timing of price changes across time. Nevertheless, Table A-I shows that my results are robust to using the full sample. While the magnitude of the results is slightly weakened, this is likely due to misattributing the timing of price changes within these two month periods, which will attenuate time-series correlations.

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

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.³⁵

The row labeled "standardized" recomputes my benchmark statistics after standardizing each price change. That is, following Klenow and Kryvtsov (2008) before computing the distribution of price changes I set $dp = \frac{dp - \mu_{dp,ELI}}{XSD_{ELI}}$ where $\mu_{dp,ELI}$ is the average price change within an ELI and XSD_{ELI} is the standard deviation of these price changes. This standardization reduces the possibility that my results are driven by a few sectors with particularly volatile prices. While sectors with particularly volatile prices would receive outsized weight in the raw statistics, this is not the case when using standardized price changes. Table A-I shows that standardizing the distribution of price changes does not qualitatively affect my conclusions.

The next rows of Table A-I show that for the vast majority of product categories,

³⁵Ss 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.

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 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. In addition to these sector level results, I have also run panel regressions with sector level fixed effects. Since there are only 11-major groups, I am unable to cluster by major group, but my results were essentially unchanged when using various other procedures for correcting standard errors. 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, although within major sectors the results are necessarily noisier as sample sizes are reduced.

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 3 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 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. My model 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 Section 3 does a good job of matching the distribution of price changes, both on average and across recessions. The bottom panel of Figure 3 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 the decompositions of inflation in Klenow and Kryvtsov (2008). 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

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

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 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.

9 Appendix 2: Analytical Model

Firms face a dynamic control problem with a single state variable z . At each point in continuous time, the ideal value of z is zero and departures from zero entail a flow cost of bz^2 and firms discount payoffs at rate r . When not adjusting, z follows a Brownian motion. Assume for now that there is no drift so that $dz = \sigma dw$, where dw is the standard Wiener process. (Note that for the study of aggregate shocks, it is convenient to break shocks into some idiosyncratic and some aggregate component $\sigma = \sigma_I + \sigma_A$, which is without loss of generality in partial equilibrium). Firms can adjust the value of z , subject to a fixed cost k . This environment has been well-studied, and it gives rise to a simple, symmetric *Ss* rule with a closed form solution. In particular,

Proposition 3 *Firms' optimal policy is to not adjust when $|z| < S$ where $S = \left(\frac{6k\sigma^2}{b}\right)^{1/4}$, and to adjust z to zero when $|z| \geq S$. In addition, the frequency of adjustment is given by $\left(\frac{b}{6k}\right)^{1/2} \sigma$ and the standard deviation of price changes is given by S .*

Proof. A version of this problem was first solved by Barro (1972), and the simplified version solved here is first described in Dixit (1991). As the solution method is now standard, I only sketch the basic outline. This firm problem yields a Bellman equation of the form $rV(z) = \frac{1}{2}\sigma^2V''(z) + bz^2$ when not adjusting together with value matching conditions at the adjustment threshold $V(S) - V(0) = k$ and smooth pasting condition $V'(S) = 0$. The solution to this differential yields a value for the threshold that (for small k) is given by $S = \left(\frac{6k\sigma^2}{b}\right)^{1/4}$. The expected time to hit the boundary of the inaction region starting from zero is given by $T = \frac{S^2}{\sigma^2}$, so that the average number of price adjustments is given by $1/T = \left(\frac{b}{6k}\right)^{1/2}\sigma$. Finally, the symmetry of the problem delivers a standard deviation equal to $\sqrt{\frac{1}{2}(S - 0)^2 + \frac{1}{2}(-S - 0)^2} = S$. ■

Proposition 4 *In an environment with zero inflation, small aggregate shocks to z do not change either the frequency or standard deviation of price changes. In an environment with positive trend inflation, small aggregate increases in z raise frequency and decrease the standard deviation of price changes.*

Proof. For the first part of the proof, see Alvarez and Lippi (2012). They show that when $dz = -\mu t + \sigma dw$ then $\frac{\partial freq}{\partial \mu}|_{\mu=0} = \frac{\partial var}{\partial \mu}|_{\mu=0} \equiv 0$. At the zero inflation steady-state, marginal increases in μ do not affect the frequency or variance of price changes. Since the model is partial equilibrium, changes in μ are the same as any other aggregate shock to all firms' z . Now consider an environment with positive trend inflation. First note that with strictly positive trend inflation firms' optimal price-setting policies imply $P_{up} > P_{down}$. Furthermore, since firms' optimal adjustment thresholds already account for aggregate shocks to z ($\sigma = \sigma_I + \sigma_A$), these thresholds will not change in response to aggregate shocks, and only the distribution of z within these thresholds will change. Let the optimal thresholds and return point when adjusting be given by \tilde{L}, \tilde{U}, C . By subtracting C from both thresholds neither the frequency of adjustment nor variance of price changes will be affected, so without loss of generality we can renormalize the optimal policy to be given by L, U with firms raising prices by L when they reach the left threshold and lowering prices by U when they reach the right threshold. Given the stochastic process for z and adjustment thresholds L, U Stokey (2009) Chapter 5 provides formulas for the ergodic distribution of z as well as the relative probability of price increases and the implied frequency of

adjustment. In particular, formula 5.16 gives probabilities:

$$P_{down} = \frac{1 - e^{\alpha L}}{e^{\alpha U} - e^{\alpha L}}$$

$$P_{up} = \frac{e^{\alpha U} - 1}{e^{\alpha U} - e^{\alpha L}}$$

and formula 5.23 gives the frequency:

$$freq = -\frac{\alpha\sigma^2}{2} \frac{[e^{\alpha U} - e^{\alpha L}]}{(U - L) - e^{\alpha L}U + e^{\alpha U}L},$$

with $\alpha = 2\mu/\sigma^2$. Using formula 5.16 it is straightforward to calculate that the variance of price changes is given by

$$var = (1 - P_{up}) P_{up} [U - L].$$

We are then interested in how the frequency and variance of price changes respond to aggregate shocks to z . Positive shocks to z are equivalent to negative shocks to μ , and vice versa, so we can calculate the response of the frequency and variance to a first moment shock by taking their derivative with respect to μ . (Here it is again worth noting that these derivatives will be taken holding U, L constant, which will be true for aggregate shocks but not for permanent changes in μ since permanent changes in μ will shift the optimal policy). Taking these derivatives, it is straightforward to show that $\frac{\partial P_{up}}{\partial \mu} > 0$. Since the variance is maximized at $P_{up} = 1/2$ (which will hold when $\mu = 0$), the variance of price changes is falling in μ . Taking the derivative of frequency yields that $\frac{\partial freq}{\partial \mu} > 0$ if $e^{\alpha U} + e^{\alpha L} > 2$. Furthermore, $e^{\alpha U} + e^{\alpha L} > 2 \iff P_{up} > P_{down}$, which again will hold in an environment with positive inflation. Thus, we get that frequency and variance move in opposite directions in response to first moment shocks.

Finally, it is of interest to compute how frequency and variance move in response to steady-state changes in μ rather than to shocks to μ . (That is, allowing for optimal policy to change in response to changes in μ). To my knowledge, closed form solutions to the system of differential equations that determine the optimal thresholds do not exist for $\mu \neq 0$. However, it is straightforward to solve the system of differential equations numerically. Although I do not have a proof, these results suggest that the above proposition also holds when U, L change in response to μ . That is, starting

from positive inflation, increases to trend inflation raise the frequency of adjustment and lower the variance of price changes. ■

10 Appendix 3: Computational Procedure

10.1 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 by³⁶

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

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 and using p/S as the idiosyncratic price 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

³⁶Note 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,

$$\log z' = \begin{cases} \rho_z \log z + d_t \sigma_z \varepsilon^z & \text{with probability } p^z \\ \log z & \text{with probability } 1-p^z. \end{cases}$$

and

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

and

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

In addition, each period firms draw iid cost of adjustment f with probability p^f and 0 with probability $1 - p^f$. Thus, we can write the firm i 's value function as

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

with

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

and

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

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 I used 171 grid points for the pricing grid, 21 grid points for the idiosyncratic productivity grid, 14 grid points for the χ_1 grid and 5 grid points 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 γ , a panel of 5000 firms³⁷ is simulated for 14,400 months³⁸ with a 100 month burnin. The 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 γ 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.

³⁷I investigated panels of up to 500,000 firms. Results were unchanged.

³⁸14,400 is 50 replications of the length of the empirical sample window.

While this version of the paper calibrates the parameters of the model, previous versions of this paper estimated simpler versions of the quantitative model more formally and all of the qualitative conclusions were similar. I explored identification of the micro parameters fairly extensively in this simpler model. The relative size of increases and decreases identifies the elasticity of substitution, the fraction increases identifies the persistence of productivity, the average size and frequency identify the standard deviation of productivity and the size of fixed costs. As mentioned in the body of the text, in the version of the model with random fixed costs and leptokurtic shocks I use data on the frequency of small price changes and the kurtosis of price changes to identify these additional parameters.

Table I: Business Cycle Properties of Price-Setting

Dependent Variable	S.D.	IQR	Freq	Med	Skew	Kurt	IQR/Med
(1) Cor. IP growth	-0.367*** (0.063)	-0.328*** (0.083)	-0.108 (0.068)	-0.235*** (0.061)	0.051 (.056)	0.111** (0.045)	-0.243*** (0.069)
(2) Cor. IP (Bandpass)	-0.410*** (0.143)	-0.393*** (0.119)	-0.282*** (0.098)	-0.268 (0.165)	0.251 (0.200)	0.279*** (0.093)	-0.345** (0.156)
(3) Recession	0.017*** (0.004)	0.030*** (0.008)	0.012*** (0.004)	0.013*** (0.004)	-0.122 (0.092)	-1.129*** (0.396)	0.174** (0.078)
Mean of Dep. Var. Non-Recession:	0.073	0.069	0.148	0.107	0.133	6.549	0.648
Mean of Dep. Var.:	0.075	0.073	0.150	0.109	0.117	6.403	0.671
Coefficient of Variation	0.231	0.354	0.108	0.120	6.72	0.361	0.331

Each column reports a time-series correlation of a price-setting statistic with a measure of the business cycle. Row 3 is a regression on recession dummies. For comparison, Mean of Dep. Var. Non-Recession shows the mean of these variables outside of recessions. Mean of Dep. Var displays the overall variable means including both recessions and non-recessions. Zeros are excluded when computing dispersion. All statistics exclude sales and product substitutions. All data is seasonally adjusted using 12 monthly dummies. Regressions in rows (1) and (3) include a quadratic time-trend. All data for regressions in row 2 are bandpass filtered using a Baxter-King (18,96,33) filter. IP=Industrial Production, Recessions use NBER data, S.D.=Standard Deviation, IQR=Interquartile Range, Freq=Average frequency of adjustment, Med=Median Frequency of Adjustment, IQR/Med Freq is the ratio of those statistics. Skew=Skewness, Kurt-Kurtosis.

n=222 for row 2, n=288 for rows 1 and 3. ***=at least 1% significance. **=5% significance (Newey-West standard errors in parentheses)

Table II: Correlation Between Frequency and Dispersion

Dep. Variable	1. S.D.	2. IQR	3. S.D. (Bandpass)	4. IQR (Bandpass)
Freq	0.134* (0.063)	0.162*** (0.051)	0.506*** (0.106)	0.660*** (0.085)
Med	0.244*** (0.058)	0.276*** (0.057)	0.610*** (0.112)	0.676*** (0.080)

This table reports correlations between measures of frequency and price change dispersion. (Newey-West standard errors in parentheses). Zeros are excluded when computing dispersion. Excluding sales and product substitutions. All data is seasonally adjusted using 12 monthly dummies. Regressions in columns (1) and (2) include a quadratic time-trend. All data for regressions in columns 3 and 4 are bandpass filtered using a Baxter King (18,96,33) filter. S.D.=Standard Deviation, IQR=Interquartile Range, Freq=Average frequency of adjustment, Med=Median frequency of Adjustment. n=288 for columns 1 and 2. n=222 for columns 3 and 4. ***=at least 1% significance.

Table III

Model Parameter Values	
Parameter	Value
Elasticity of Substitution	6.8
Productivity Persistence	.62
Productivity Standard Deviation	.08
Poisson Probability of Prod. Shock	.13
Fixed Cost (If positive)	.09
Probability Zero Adjustment Cost	.09

Table IV

Moments Targeted	Model Fit		
	Data	First+Second	Only First
Frequency	.11	.11	.11
Fraction Up	.65	.66	.67
Size Up	.07	.055	.055
Size Down	.09	.064	.061
Fraction Small	.33	.34	.34
Kurtosis	6.4	7.2	7.3
Moments Not Targeted			
Correlation Dispersion and Frequency	.61	.62	-.58
Correlation Dispersion and Y	-.41	-.60	.14
Correlation Freq and Y	-.27	-.24	-.07
Correlation Dispersion/Freq and Y	-.30	-.39	.11

Table V

Volatility	Impulse Response: Baseline Model		
	Price IRF on Impact	Output IRF on Impact	Total Output IRF
10th percentile	28.9%	71.1%	418.0%
50th percentile	35.0%	65.0%	360.8%
90th percentile	39.3%	60.7%	325.0%

Output Impulse as a percent of 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	10.2%	18.7%
90th percentile	12.2%	27.1%

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

The shock doubles nominal output growth from .002 to .004 for one month

Table VII

Increase in Price Impulse Response on Impact

Business Cycle	Baseline	Match Corr	Match Corr-Pers	Small Volatility
90th to 10th Percentile of Output	33%	25%	22%	13%
10th to 90th Percentile of Volatility	36%	30%	25%	17%

90th to 10th percentile of output shows percentage increase in price response as output falls. 10th to 90th percentile of volatility shows percentage increase in price response as volatility rises. The nominal shock is a 1 month doubling of nominal output growth. Baseline model is the model with volatility shocks in Section 3. Match Corr keeps the persistence of volatility the same as TFP but matches the cyclical 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. The impulse response is an increase of the nominal output growth rate from .002 to .004 for one month

Table A-I

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Series	Freq,Y	Med,Y	XSD,Y	IQR,Y	Skew,Y	Kurt,Y	Freq,XSD	Freq,IQR	Med,XSD	Med,IQR
Benchmark	-.28***	-.27***	-.41***	-.39***	.25	.28***	.51***	.66***	.61***	.68***
Standardized	-.28***	-.27***	-.33***	-.42***	-.02	.39***	.35***	.51***	.43***	.51***
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. Standardized subtracts ELI means from each price change and divides by ELI standard deviation. 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

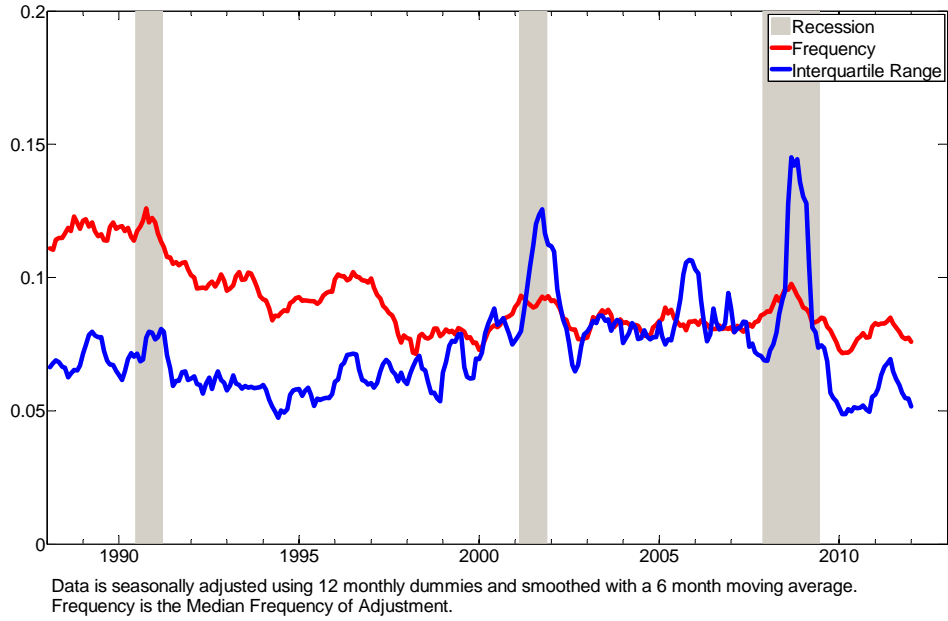


Figure 2: Price Changes Over the Business Cycle




Figure 3: Price Change Distributions

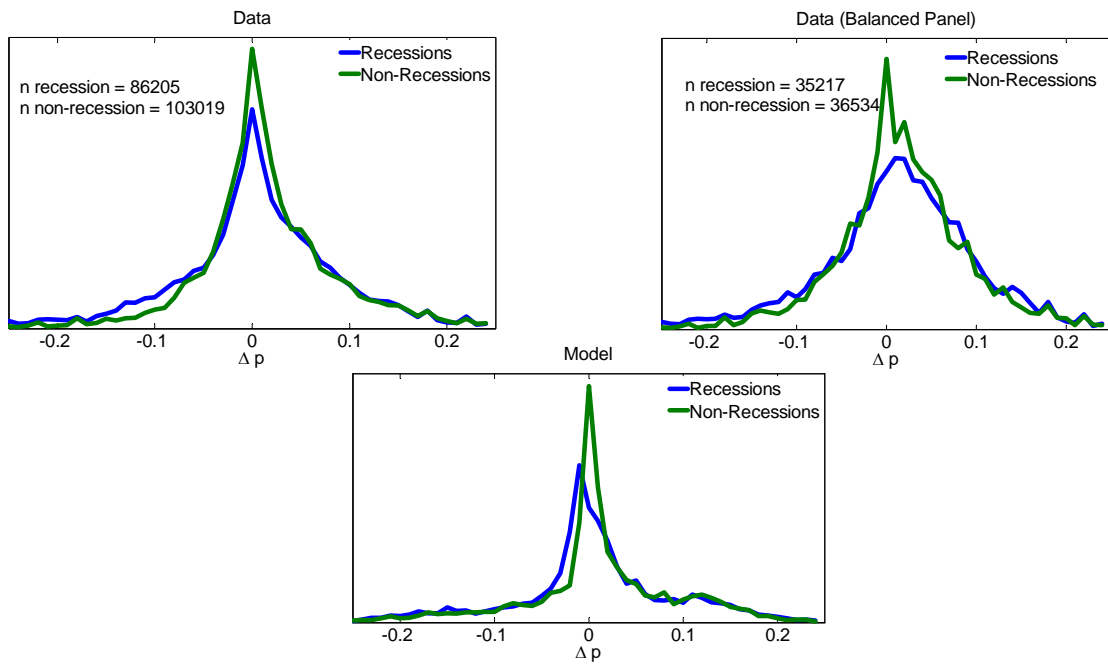


Figure 4: Greater Volatility Can Induce Greater Frequency and Price Change Dispersion

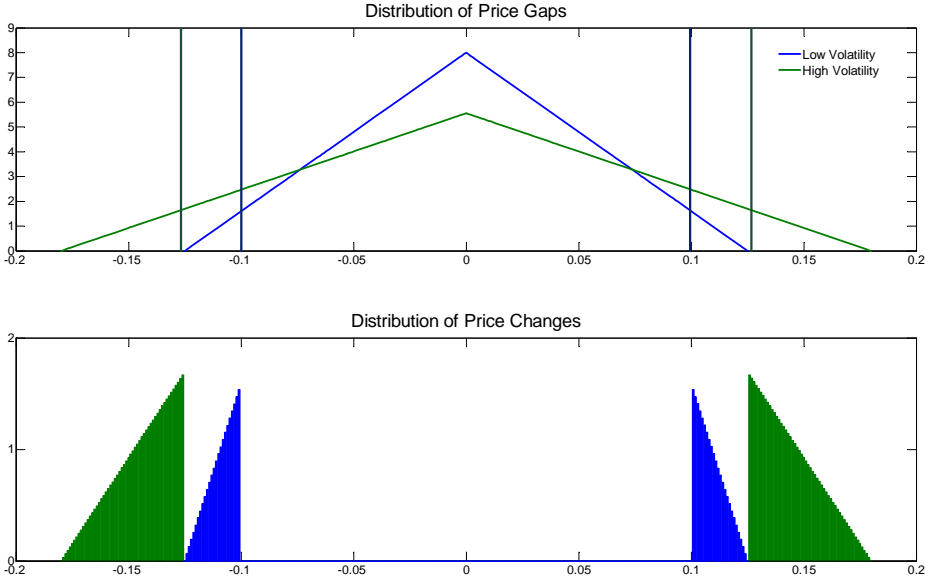


Figure 5: Negative Correlation Between Frequency and Price Change Dispersion

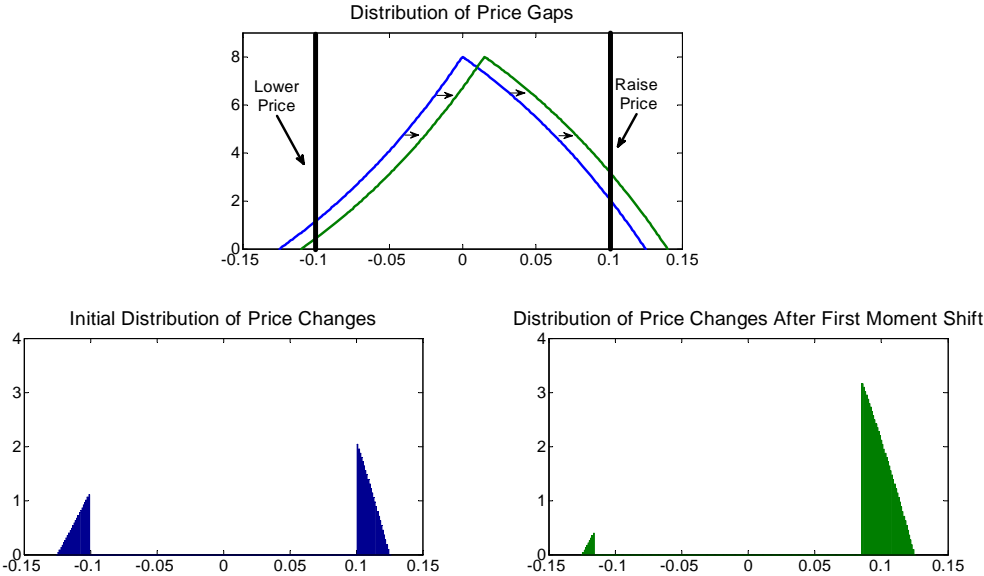


Figure 6: Real Output Impulse Response to Nominal Shock

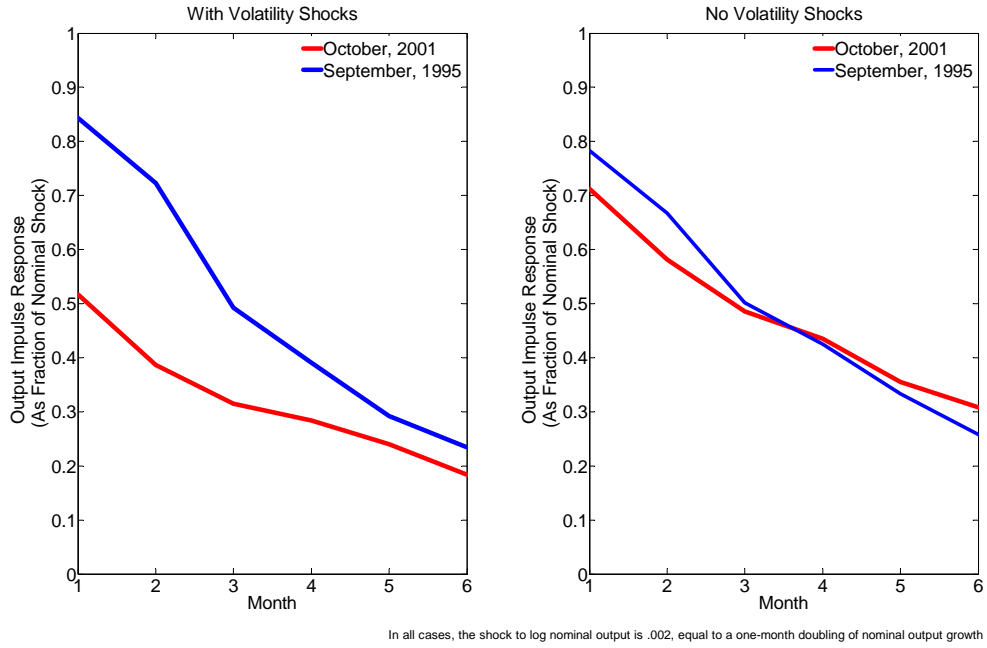


Figure 7: Real Output Response to Autocorrelated Nominal Shock

