Inflation Dynamics and Time-Varying Uncertainty: New Evidence and an Ss Interpretation

Joseph Vavra*
Yale University

10/22/2011

Abstract

Is monetary policy less effective during recessions? In this paper, I investigate whether the transmission of monetary policy shocks to the real economy varies across time in micro founded Ss pricing models. I find that standard Ss models do not imply time-varying real effects of monetary policy, but these models have counterfactual implications for the dynamics of the distribution of price changes in BLS micro data. I then introduce counter-cyclical "uncertainty shocks" to match these new micro price facts and show that the price level becomes more flexible in times of high uncertainty, which in turn decreases the response of real output to nominal shocks. The estimated response of real output to a nominal shock in October of 2001, a highly uncertain time, is one fourth of the response in September of 1995, a time of low uncertainty. U.S. time-series evidence provides additional support for my model: I show that during times of high uncertainty, exchange rate pass-through to import prices is significantly higher. In addition, the price level responds more strongly to monetary policy shocks in a FAVAR during such periods of uncertainty.

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

1

1 Introduction

The recent recession has led to renewed interest over whether the effectiveness of monetary policy may vary across time. If monetary policy becomes less effective during recessions, then its use as the primary stablizing mechanism for the economy might be called into question. In this paper, I investigate whether the transmission of monetary policy shocks to real output varies across time, in micro founded Ss pricing models. I find that standard Ss models do not imply time-varying real effects of monetary policy, but that is because these models miss important dynamic facts about the distribution of price changes in the U.S. In particular, these models have counterfactual implications for the second moment of price changes. I then introduce second moment "uncertainty shocks" into the model as in Bloom [2009] to match these micro price facts and show that the price level becomes more flexible in times of high uncertainty, which in turn decreases the response of real output to nominal shocks. Since uncertainty is countercyclical, this implies that nominal shocks have little real effect during uncertain recessions.

Recent work, including Bloom [2009], Bloom, Floetotto, and Jaimovich [2009], and Gilchrist, Sim, and Zakrajsek [2010], has brought renewed attention to the view that increases in uncertainty may cause recessions. While these papers have focused on investment and the real economy, increases in uncertainty should also manifest themselves in firms' prices. In this paper, I argue that micro price data indeed supports the view that uncertainty shocks are an important feature of the economy and that these shocks have important implications for monetary policy.

I begin with empirical results using the BLS micro data that underlies the CPI. I provide evidence that the cross-sectional standard deviation of price changes is strongly countercyclical: price changes become substantially more disperse during recessions, when other measures of uncertainty typically rise. Furthermore, the standard deviation of price changes comoves strongly with the average frequency of price adjustment in the economy. That is, the dispersion of price changes (conditional on adjustment) is high when the frequency of adjustment is high.

I next assess the ability of Ss models with only first moment shocks to match these new empirical facts. While there are a variety of models of firm pricing, I focus on Ss models, as they provide a simple framework for endogenizing both the frequency and size of price changes.[1] In these models, firms face idiosyncratic productivity shocks as well as aggregate productivity and nominal shocks and must pay a fixed cost to adjust their nominal prices. This fixed cost induces firms to follow Ss pricing strategies: they only adjust to their desired price when their current price falls outside of some threshold (S,s) region. Within this region, it is not worth paying the adjustment cost, and firms maintain a constant price. While these models have been shown to do a good job of capturing a wide range of micro pricing facts, I show that they get the new empirical facts wrong. Ss models with only first moment shocks imply a counterfactual negative correlation between price dispersion and the frequency of adjustment, and they generate procyclical price dispersion.

I then show that the addition of second moment shocks, that increase the variance of firm level idiosyncratic productivity, improves the model fit dramatically. Furthermore, the model with second moment shocks generates very different implications for the transmission of nominal shocks.

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

As the region of inaction grows but more firms are pushed to adjust by larger shocks, the price level becomes more flexible so that the price response to nominal shocks is greater in times of high uncertainty. Equivalently, real output responds much less to nominal shocks in times of high uncertainty. In contrast, in the Ss model with only first moment shocks, the real response to nominal shocks is not time-varying.

---

[1] Eichenbaum, Jaimovich, and Rebelo [2009] provides evidence that price-setters engage in Ss like behavior, adjusting
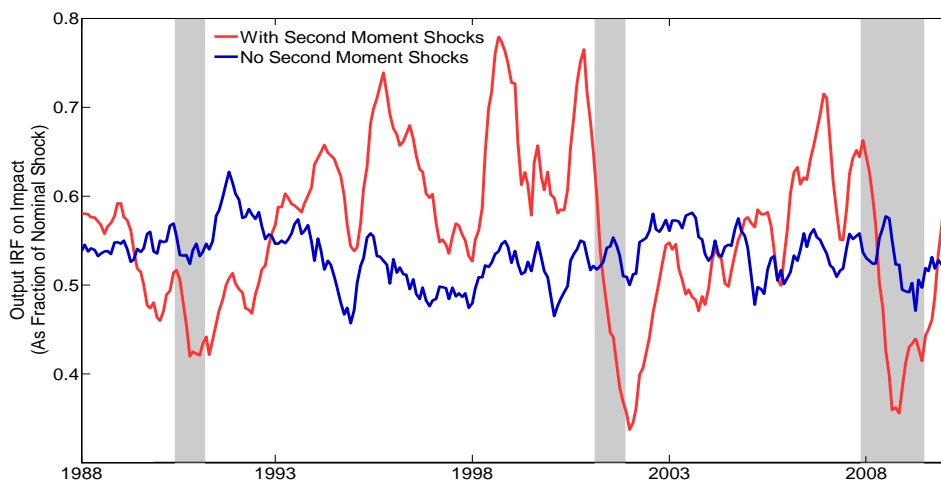
Figure 1: Output Impulse Response on Impact

Figure 1 displays the estimated real output impulse response on impact to a one standard deviation nominal shock, computed for the U.S. economy from 1988-2010. Figure 2 compares the full impulse response in 1995 to that in 2001.

In the model with second moment shocks, the total response of real output to a nominal shock in October of 2001, a highly uncertain time, is one quarter of the response in September of 1995, a time of very low uncertainty. Clearly, there is no such difference for the model with only first moment shocks.

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 uncertainty increases, both margins become more important as there are more firms adjusting prices, and more firms are pushed near the

prices when their markup drifts out of a target region.

4

Figure 2: Real Output Impulse Response to Nominal Shock
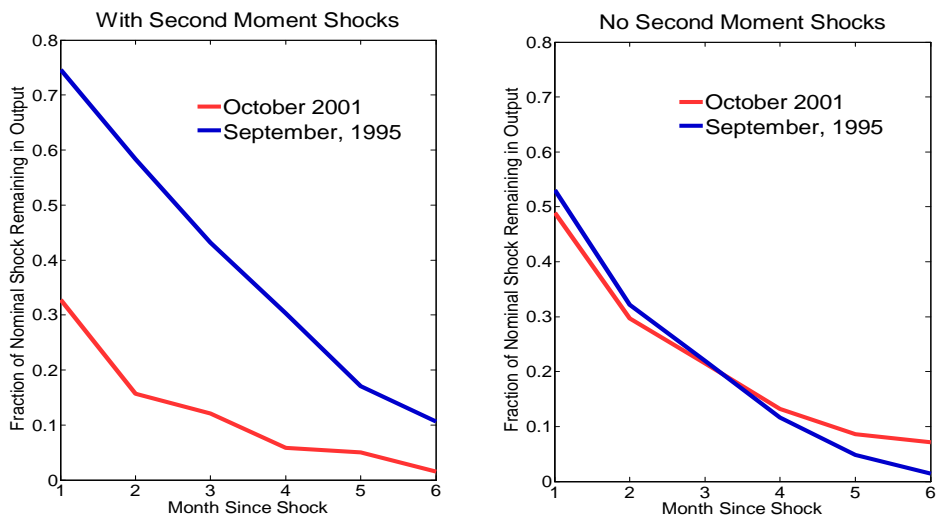
margin of adjustment by the more volatile shock process. Overall, the price response on impact to a nominal shock is 130% larger at the 90th percentile of uncertainty than it is at the 10th percentile of uncertainty, and two thirds of this increase is driven by a more responsive extensive margin.

The general conclusion of the model, that the price response to nominal shocks is dramatically larger during uncertain recessions, is robust to a number of model extensions. Nevertheless, I also present additional empirical evidence that the price level is more flexible during times of high uncertainty, that does not rely on a particular structural model. First, I estimate otherwise standard exchange rate pass-through equations,[2] allowing pass-through to vary with uncertainty. I find that pass-through of nominal exchange rate changes to import prices is substantially larger during times of high uncertainty. Second, I estimate a regime-specific factor augmented VAR, allowing factor loadings and VAR coefficients to vary across high and low uncertainty regimes. Again, I find that the results are consistent with my quantitative model: The estimated price response to identified monetary policy shocks is significantly greater during the high uncertainty regime.

The remainder of the paper is organized as follows: Section 2 contains the empirical findings.

_____

[2] See Campa and Goldberg [2005] and Marazzi, Sheets, Vigfusson, Faust, Gagnon, Marquez, Martin, Reeve, and Rogers [2005].

Section 3 shows that Ss models with only first moment shocks get the empirical findings wrong. I first analyze a simple analytical Ss model and show that it generates a counterfactual negative correlation between price change dispersion and the frequency of adjustment. I then estimate a fully-specified quantitative Ss model and show that the analytical results still hold and that furthermore, in contrast to the empirical evidence, the model implies procyclical price dispersion. Section 4 adds second moment shocks to the model and shows that the model's fit is improved dramatically. Section 5 discusses policy implications of the second moment shocks. Section 6 presents the reduced form evidence that price flexibility rises during times of high uncertainty 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 March, 2010. Prices are only collected monthly for the entire sample period in New York, Los Angeles and Chicago so my analysis is restricted to these cities.[3] The database contains thousands of individual "quote-lines" with price observations for many months. Quote-lines are the greatest level of disaggregation possible and correspond to an individual item at a particular outlet. An example of a quote-line collected in the research database is 2-liter Coke at a particular Chicago outlet.

These quote-lines are then classified into various product categories called "Entry Level Items" or ELIs. The ELIs can then be grouped into several levels of more aggregated product categories finishing with eight major expenditure groups: Apparel, Education and Communication, Food, Other Goods and Services, Housing, Medical Care, Recreation, and Transportion. For more details on the structure of the database see Nakamura and Steinsson [2008].

This database has received great attention in recent years, beginning with Bils and Klenow [2004]. While initial studies of this micro pricing data focused on static first moments of the data such as the average frequency and size of price changes, only recently have more dynamic features of the

---

[3]Using only monthly data reduces the sample from approximately 85,000 items to approximately 15,000 items per month, so it is a significant restriction. However, using the full sample does not qualitatively affect the results. See the data appendix for additional discussion.

data begun to receive attention. (See Klenow and Malin [2010] for detailed summaries of the recent literature utilizing this data). However, despite the widespread attention this data has received, analysis of higher moments of the price change distribution has begun only recently. Klenow and Malin [2010] provides brief evidence of the relationship between first and higher moments of inflation and calls for additional attention to this topic. Berger and Vavra [2011] analyzes the dynamic, business cycle behavior of the distribution of price changes. In that paper, we provide evidence that price dispersion (the second moment of the price change distribution) has many robust, dynamic patterns which may help to distinguish models of price setting.

In this paper, I focus on two facts from Berger and Vavra [2011] that I will argue are of particular significance for models of firm pricing. Let $dp_{i,t} = \log \frac{p_{i,t}}{p_{i,t-1}}$ be the log price change observed for item i at time t. Then, using aggregation weights provided by the BLS,[4] it is straightforward to compute the cross-sectional dispersion of log price changes[5] for each month and investigate how it varies over time.

Figure 3 shows the relationship between the cross-sectional interquartile range of price changes (excluding non-adjusters[6]) and industrial production.[7] This figure clearly shows Fact 1: price change dispersion is countercyclical. This fact is confirmed numerically in Table 1, which shows that there is a strong negative correlation between various measures of price change dispersion and industrial production.[8]

---

[4] These weights are used for aggregating individual price series to create the CPI. The weights describe how individual quote-lines are weighted within ELIs and how ELIs are aggregated into overall expenditure. The ELI aggregation step is based on the Consumer Expenditure Survey and is updated by the BLS periodically.
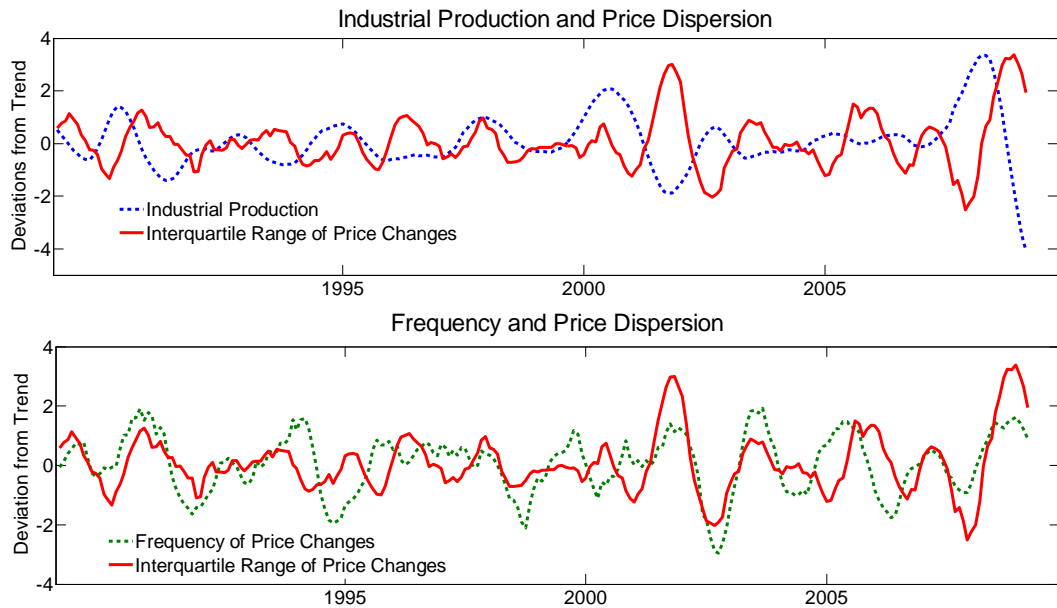
[5] We also computed several alternative measures of the size of price changes that are more robust to outliers. Results were qualitatively similar.

[6] Even stronger results obtain for price dispersion including zeros, but the models that I will be investigating have stronger implications for dispersion excluding zeros.

[7] Industrial production is used as it is available monthly and so can be directly compared to monthly IQR numbers. Similar qualitative results obtain if dispersion is aggregated quarterly and compared to GDP.

[8] From Figure 3 it appears that there may be some lead-lag structure to the relationship. Indeed, I find that price dispersion mildly leads output by 1-2 months. This is consistent with the view in Bloom, Floetotto, and Jaimovich [2009] that uncertainty shocks may be an important driving force for business cycles. There is no such lead-lag relationship for the relationship between frequency and price dispersion, which is consistent with the timing of uncertainty realizations in my model.

Figure 3: Time-Series Relationships with Price Dispersion

Table 1

Dispersion Measure	Correlation with Production	Correlation with Frequency
IQR	-0.40***	0.52***
Standard Deviation	-0.45***	0.42***

Zeros are excluded when computing dispersion. Excluding sales and product substitutions.

Data is seasonally adjusted and bandpass filtered. n=242. ***=at least 1% significance.

Figure 3 also shows the relationship between the interquartile range and the frequency of adjust-ment, which leads to Fact 2: price dispersion is positively correlated with the frequency of adjustment. Again, Table 1 numerically confirms this qualitative result.[9]

---

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

In principle, the aggregate relationships reported in Table 1 may be driven by compositional changes in the make up of price changes over the business cycle. However, we find that these facts hold both within as well as across sectors so that they are unlikely to be driven by any particular items.[10] Furthermore, the dispersion of price increases, the dispersion of price decreases, and the difference between the mean increase and the mean decrease are all countercyclical so that the dispersion of all components of price dispersion rise during recessions.[11] This suggests that the countercyclicality of price dispersion and the positive relationship between price dispersion and the frequency of adjustment are robust facts that are common to the majority of items and price changes in the U.S. economy.

The data appendix discusses the construction of these dispersion numbers in more detail as well as providing a number of robustness checks. In particular, while the above facts are computed using bandpass filtered data, similar qualitative results obtain using alternative filtering.[12] In addition, the results are robust to the inclusion or exclusion of zero price changes, as well as to the inclusion of sales and product substitutions. Finally, different procedures to control for outliers and measurement error do not significantly alter the results.

The remainder of the paper 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 empirical facts strongly suggest that second moment shocks are an important feature of the economic environment affecting firms' pricing decisions in the United States.

3   Ss Models with First Moment Shocks

In this section, I present evidence that Ss models with only first moment shocks imply a negative correlation between the dispersion of price changes and the frequency of adjustment and a positive correlation between dispersion and output, in contrast to the empirical evidence. I begin with a

---

[10] We investigated a variety of sector definitions: Major Groups, ELIs, durables and non-durables, and core and non-core.

[11] The relationships are most dramatic for price decreases.

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

simple analytical model that provides intuition for the negative correlation between price dispersion and frequency. I then move to a fully-specified quantitative model and show that the result from the analytical model still holds and that, in addition, price dispersion is procyclical.

3.1   Analytical Ss model

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

A firm's optimal policy is a two-sided Ss rule: firms raise prices by L when z_t reaches a lower threshold $L < 0$ and lower prices by U when z_t reaches an upper threshold $U > 0$.[13] This continuous time environment yields the following result:

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

Proof. See Appendix 2. ∎

The formal proof consists of calculating the frequency and variance of price changes under the ergodic distribution analytically. The response of these variables to an aggregate shock can then be calculated as $\frac{\partial freq}{\partial \pi}$ and $\frac{\partial var}{\partial \pi}$ as an increase in π is equivalent to a long-sequence of positive shocks in the ergodic distribution. While the proof simply consists of signing these derivatives, the intuition can largely be expressed graphically.

---

[13] That the optimal reset point is 0 is without loss of generality, as any other constant target can be subsumed by renormalizing U and L.

With inflation close to zero, and $U = -L$, the ergodic distribution is nearly symmetric about zero and the frequency of price increases is equal to the frequency of price decreases. Figure 4 displays this case. Now as inflation grows, the ergodic density becomes skewed as positive inflation pushes

Figure 4: Zero Inflation
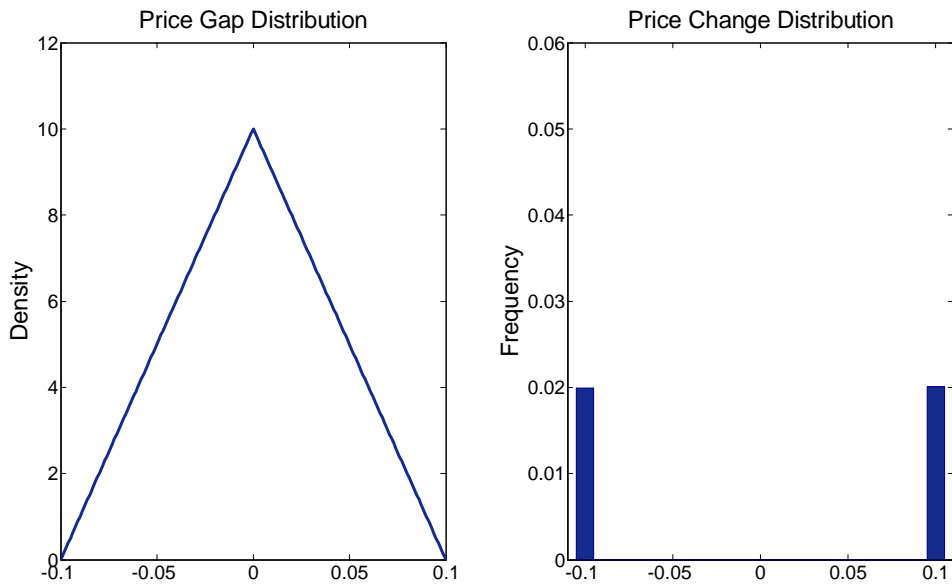
a greater mass of firms towards raising their prices[14] as shown in Figure 5. This implies that the frequency of increases rises and the frequency of decreases falls. However, what is gained in price increases is larger than what is lost in price decreases as the total mass adjusting is now increased. This also implies that the price change distribution is now less symmetric: there are many more price increases than decreases. Since the variance of the price change distribution (conditional on adjustment) can be written as the variance of a linear transformation of a Bernoulli distribution, the variance necessarily falls as the price change distribution becomes less symmetric. Thus, we get that the frequency of adjustment and the variance of price changes move in opposite directions. This is

---

[14] Note that we can interpret the change in frequency and variance after an increase in inflation as the result of a long sequence of positive shocks under the ergodic distribution rather than a permanent increase in inflation. The former implies that the optimal policy is unchanged.

Figure 5: Positive Inflation

Price Gap Distribution

Price Change Distribution
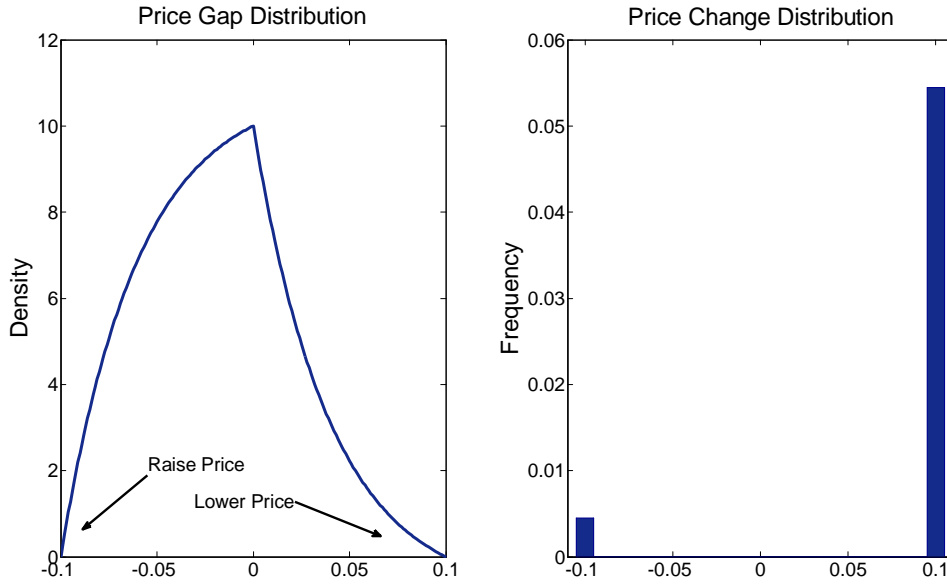
in contrast to the strong positive correlation that is observed empirically.

While this analytical result is quite strong, it relies on a model with several simplifying assumptions. It is partial equilibrium, imposes a quadratic loss function, and it has a very simple process for the evolution of the price gap, which is not derived from micro foundations. Thus, I now move to a more quantitative Ss model. This general equilibrium model will have fully specified micro foundations and its quantitative fit to the empirical data will be evaluated more formally. Nevertheless, I will show that the result from the analytical model still holds: the Ss model with only first moment shocks implies a negative correlation between price change frequency and dispersion. Furthermore, the quantitative model's fit can also be investigated along other dimensions, and I will show that it implies procyclical price dispersion, in contrast to the empirical evidence.

3.2 Quantitative Equilibrium Ss model

The quantitative model closely follows Golosov and Lucas [2007], and I incorporate aggregate shocks using methods first explored in Midrigan [2011]. The economy is composed of a representative household and a continuum of monopolistically competitive firms. I first discuss the household problem. I then present the firm problem and define equilibrium.

3.2.1 Households

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

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

subject to

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

where

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

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

3.2.2 Firms

Firms produce output using a linear technology in labor

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

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

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

aggregate productivity a_t evolves according to

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

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

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

where P_t is the Dixit-Stiglitz price index

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

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

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

where firm profits are given by

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

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

14

3.2.3 Computing the Equilibrium

I assume that nominal aggregate spending $S_t = P_t C_t$ follows a random walk with drift in logs:[15]

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

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

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

where $\chi_{1,t} = mean\left(\log\left[\frac{p_{t-1}^i}{S_t}a_t z_t^i\right]\right) = \log\frac{P_{-1}}{S_t} + \log a_t$ up to a constant. Given this conjecture,[16] I then search for a value of the transition coefficients, $\boldsymbol{\gamma}$ so that the true law of motion in the economy is well approximated by the conjectured law of motion. At this point, a regression of the actual law of motion on the conjectured law of motion gives R^2 in excess of 99%. Furthermore, adding one additional moment from the joint distribution does not change the qualitative conclusions. Finally, rather than comparing the conjectured law of motion to the actual law of motion point-by-point 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

---

[15] I have solved a version of the model with autocorrelated nominal spending shocks. With autocorrelated nominal spending shocks, the model generates a hump-shaped output IRF, in line with VAR evidence. The other qualitative results of the model (both with and without uncertainty shocks) were unchanged.

[16] This conjecture gives a linear relationship between the current real price level and the old price level relative to the new money stock, with slope and intercept that can vary with the aggregate real state of the economy. Other specifications arrived at similar qualitative results.

suggested by Den Haan [2010]. The approximation errors remain extremely small.

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

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

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

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

and

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

and firm flow profits can be written as

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

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

3.2.4 Estimation and Results

The model period is one month, so I set the discount factor $\beta = .997$. The calibration of the nominal shock process follows Nakamura and Steinsson [2010]. Since there is no long-run real growth in the model economy, I set $\mu = .002$ to match the mean growth rate of nominal GDP minus real GDP, and I set $\sigma_s = .0037$ to match the standard deviation of nominal GDP growth, over the period 1998-2010. 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.[17]

---

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

The remaining parameters are estimated using simulated method of moments. There are four model parameters: the elasticity of substitution θ, the persistence and standard deviation of idiosyncratic productivity, ρ_z and σ_z, and the fixed cost f. These parameters are selected to fit six moments from the data: the average frequency of adjustment, the average size of increases, the average size of decreases, the fraction of price changes that are increases, the correlation between frequency and the standard deviation of price changes, and the time-series standard deviation of cross-sectional price dispersion. The first four moments are standard while the latter two parameters capture salient features of the empirical data presented in Section 2. For more details on the estimation procedure as well as a discussion of alternative moments and estimation schemes, see Appendix 3.

Table 2

Model with First Moment Shocks

Parameter	Estimated Value	
Elasticity of Substitution	6.9	(6.4,7.5)
Productivity Persistence	.63	(.59,.68)
Productivity Standard Deviation	.0395	(.0363,.043)
Fixed Cost $\times 10^3$	2.3	(2.1, 2.6)

Moment	Data	Model	
Frequency	.10	.11	(.10,.13)
Fraction Up	.65	.64	(.62,.66)
Size Up	.08	.075	(.071,.081)
Size Down	.10	.086	(.080,.091)
Correlation Dispersion and Frequency	.42	-.68	(-.78,-.59)
Standard Deviation of Dispersion	.081	.050	(.037,.062)

Bootstrapped 95% Confidence intervals in parantheses. Fixed cost is average

fraction of steady-state revenues paid by all firms

---

change the qualitative conclusions.

Table 2 shows the model's estimated parameters and best fit moments. The estimated parameters of the model are in line with recent literature. The elasticity of substitution of 6.9 is in between the values used by Golosov and Lucas [2007] and Nakamura and Steinsson [2010]. The fixed cost implies that total adjustment costs in the economy represent just over 0.2% of steady-state monthly revenues.[18] The estimated persistence of productivity is relatively low, with a monthly persistence of 0.63. This low persistence is largely driven by matching the large and relatively frequent price changes observed in the data. Again, the productivity parameters are roughly in line with previous estimates in the menu cost literature.

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

Furthermore, the Ss model with only first moment shocks generates too little time-series variation in the standard deviation of price changes. In the estimation procedure I only include price dispersion excluding non-adjusters. However, while not used in estimation, the lack of time-series variation is even more apparent for the standard deviation of price changes *including* zeros. The time-series standard deviation of price dispersion including zeros is 0.018 in the model, while it is 0.11 in the data. This is clearly because in the data, the positive correlation between the frequency of adjustment and price dispersion excluding zeros makes the time-series variation of dispersion including zeros larger than the time-series variation of dispersion excluding zeros. In contrast, in the Ss model with first moment shocks, movements in the frequency of adjustment offset increases in price dispersion excluding zeros so that the time-series variation of dispersion excluding zeros is larger than that including zeros.

A final empirical failure of the model with only first moment shocks is the business cycle behavior of

---

[18] This measure of the fixed cost is given by $f * freq * \frac{\theta-1}{\theta}/Y_{ss}$. The cost conditional on adjustment is around 2% of revenues, which is roughly in line with the estimates in Zbaracki, Ritson, Levy, Dutta, and Bergen [2004].

price dispersion. Empirically, there is a negative correlation of -.45 between output and the standard deviation of price changes while in the model there is a positive correlation of 0.24.

4   Second Moment Shocks

The previous section argues that Ss models with first moment shocks are unable to match the empirical evidence on price dispersion. In this section, I argue that the addition of time-varying uncertainty or second moment shocks improves the model fit dramatically. While there are various notions of uncertainty, in the model, an increase in uncertainty will be modeled as a common increase in the standard deviation of firm's idiosyncratic productivity. Since it is idiosyncratic productivity differences across firms that generate price dispersion, it is natural to increase the standard deviation of this productivity in order to increase price dispersion. With time-varying uncertainty, the negative correlation between frequency and price dispersion implied by Ss models can potentially be broken. If more dispersed shocks induce more firms to adjust, then dispersion and the frequency of adjustment can comove.

4.1   The Model

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

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

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

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

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

productivity and d_t are perfectly negatively correlated. That is,

$$\rho_a = \rho_d$$
$$\varepsilon_t^a = -\varepsilon_t^d.$$

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 [2009] find that the cross-sectional standard deviation of firm-level Solow residuals are strongly counter-cyclical with a correlation of -.48 with detrended output.

Furthermore, Bloom, Floetotto, and Jaimovich [2009] and Gilchrist, Sim, and Zakrajsek [2010] show that in the presence of fixed costs of capital adjustment or financial market imperfections, second moment shocks alone can generate falls in aggregate productivity. Thus, my assumption that aggregate productivity and uncertainty are negatively correlated may also be justified as a reduced form for the interactions between capital and uncertainty. Admitedly, this interaction is not present in my model since labor is the sole factor of production.

In any event, I have solved a version of the model without this perfect negative correlation. While the model is more computationally burdensome, and as such, it cannot be fully estimated, the qualitative conclusions appear essentially unchanged.[19]

Under the assumption that a_t and s_t are perfectly negatively correlated, the firm problem can be characterized in terms of the same state variables as in the problem with only first moment shocks. The only difference in the firm problem is that the standard deviation of idiosyncratic firm shocks now depends on the aggregate state of the economy and firms' expectations must account for this time-varying standard deviation. The form of the Krusell and Smith [1998] transition rule remains unchanged, and at a fixed point, the conjectured law of motion again provides an extremely accurate forecast for the true law of motion of the economy.

---

[19]Furthermore, when the perfect negative correlation between productivity and uncertainty is relaxed, I can also investigate the effects of changes in the persistence of uncertainty. The results are not particularly sensitive to decreasing the persistence of uncertainty to match the empirical persistence of price dispersion.

4.2 Estimation and Results

The estimation procedure remains the same as in the model with only first moment shocks. The model is estimated to match the frequency, size of increases and decreases, the correlation between frequency and price dispersion and the time-series standard deviation of dispersion. There is now one additional parameter to be estimated. The standard deviation of uncertainty, σ_d, must be estimated in addition to the elasticity of substitution, persistence and standard deviation of idiosyncratic productivity, and the cost of price adjustment. Again, identification as well as more details on the estimation procedure and alternative estimation schemes are discussed in Appendix 3. In general, the estimated parameters for the model with second moment shocks are similar to those for the model with only first moment shocks.

Table 3

Model with Second Moment Shocks

Parameter	Estimated Value	
Elasticity of Substitution	7.9	(7.3, 8.5)
Productivity Persistence	.640	(.585,.705)
Productivity Standard Deviation	.0425	(.0394,.0457)
Fixed Cost $\times 10^3$	4.0	(3.4,5.0)
Standard Deviation Uncertainty	.091	(.068,.115)

Bootstrapped 95% Confidence intervals in parantheses. Fixed cost is average

fraction of steady-state revenues paid by all firms

Table 4

Model Fit Comparison

Moment	Data	Only First		First + Second	
Frequency	.10	.110	(.100,.130)	.095	(.078, .114)
Fraction Up	.65	.64	(.62,.66)	.66	(.62,.70)
Size Up	.08	.075	(.071,.081)	.081	(.076, .091)
Size Down	.10	.086	(.080,.091)	.097	(.092, .105)
Correlation Dispersion and Frequency	.42	-.68	(-.78,-.59)	.43	(.17,.61)
Standard Deviation of Dispersion	.081	.050	(.037,.062)	.080	(.050, .120)

Bootstrapped 95% Confidence intervals in parantheses

As can be seen from Table 4, the model fit with first and second moment shocks is a dramatic improvement over the model with only first moment shocks. The model now implies a correlation between the frequency and standard deviation of price changes that is positive and closely in line with the empirical data instead of the strongly negative correlation implied by the model with only first moment shocks. Furthermore, the time-series standard deviation of price dispersion now matches the data instead of being too small. This improvement in model fit is more dramatic for the time-series variation of dispersion including zeros, which was not used in estimation. Empirically, the standard deviation of dispersion including zeros is 0.11. This same value is generated by the model with first and second moment shocks, while, as mentioned previously, a standard deviation of only 0.018 is generated by the model with only first moment shocks. Finally, the model with first and second moment shocks implies that price dispersion is strongly countercyclical: the model generated correlation[20] is -0.82 while the empirical correlation is -0.45.
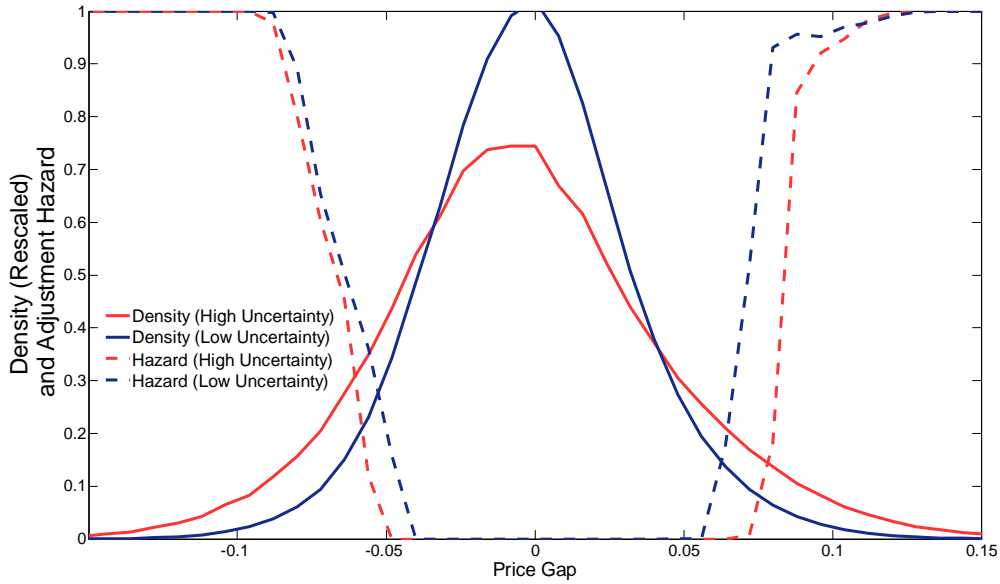
Why does time-varying uncertainty imply a positive correlation between the frequency of price changes and the cross-sectional standard deviation of price changes? As emphasized by Bloom [2009], in the presence of fixed adjustment costs, an increase in uncertainty has two effects: 1. There is a "wait-and-see" effect. Greater uncertainty makes the option value of waiting increase as it is

---

[20] As mentioned in the previous section, I have solved a version of the model where productivity and uncertainty have a less than perfect negative correlation. In this case, the empirical countercyclicality of dispersion can be matched exactly, and other results are qualitatively unchanged.

not worth paying adjustment costs today if a firm will want to reverse its decision tomorrow. This wait-and-see effect makes firms inaction region widen so that price adjustment is less likely. 2. There is a "volatility" effect. When shocks have a greater standard deviation, more firms will hit adjustment bands of any given width so that price adjustment rises.

The two effects work in opposite directions, but quantitatively, the volatility effect dominates. Figure 6 displays how the density of firm price gaps as well as the adjustment hazard for a given price gap respond to an increase in uncertainty. The wait-and-see effect is clear, as the adjustment hazard widens in the state of the world with increased uncertainty. Similarly, the density of price gaps also spreads out so that the density of firms with a high probability of adjustment rises. On net, the second effect dominates so that the frequency of adjustment rises from 7% to 12% per month when uncertainty moves from the 25th to the 75th percentile.

Figure 6: Response to an Increase in Uncertainty

When uncertainty rises, adjustment bands widen so that the difference between the average price increase and the average price decrease grows. At the same time, the fraction of price changes

that are increases is little affected, so the standard deviation of price changes grows. And since the volatility effect dominates, the frequency of adjustment also grows. Thus, the model implies a positive correlation between the frequency of adjustment and the cross-sectional standard deviation of price changes.

Since the model with no uncertainty shocks generates a negative correlation between these variables, it is clear that the magnitude of uncertainty shocks is a critical input for matching the empirical relationship. How large are the uncertainty shocks required to generate a positive correlation between frequency and dispersion in line with the empirical evidence? The estimated standard deviation of (log) uncertainty shocks is 0.091. This implies that at the 95th percentile of uncertainty, firms' standard deviation of productivity is roughly 35% above average. The annual coefficient of variation of the cross-sectional standard deviation of productivity is 15%. This falls in between the empirical coefficient of variation of 2.67% found by Bachmann and Bayer [2009] using German data, and the coefficient of variation of 17% estimated in Bloom [2009].[21] Thus, the size of the shocks estimated in my model is roughly in-line with previous estimates on the magnitude of uncertainty shocks.

5 Policy Implications

I now show that time-varying uncertainty has striking implications for the transmission of nominal shocks to the real economy. In times of high uncertainty, the real effect of nominal shocks is substantially reduced. Table 5 shows the first element of the real output impulse response, as a percentage of the nominal shock.

---

[21]Both of these estimates are in a context with capital, which is not present in my model.

Table 5

Output Impulse Response

Uncertainty	Output IRF
10th percentile	75%
25th percentile	66%
50th percentile	57%
75th percentile	49%
90th percentile	43%

Output Impulse on impact as a percent of total nominal

shock. The nominal shock is a 1 month doubling of

nominal output growth.

Since $S = PY$, the remainder of the nominal shock goes into inflation. Thus, the real effect of a nominal shock, on impact, is nearly doubled when moving from the 90th percentile of uncertainty to the 10th percentile of uncertainty.[22] The cumulative real effect of a nominal shock at the 90th percentile of uncertainty is nearly tripled relative to that at the 10th percentile of uncertainty.[23]

For another perspective on the relationship between uncertainty and the real effects of nominal shocks, I compute the size of the nominal shock required to generate the same real output response under different degrees of uncertainty. At the 90th percentile of uncertainty, the increase in nominal output must be 80 percent larger than at the 10th percentile of uncertainty in order to generate the same real effect. This requires increasing inflation by four times as much in the high uncertainty state.

---

[22] While the model with second moment shocks implies substantial time-variation in the output IRF, it should be noted that the average IRF with and without second moment shocks is similar, so the addition of second moment shocks does not affect the result in Golosov and Lucas [2007] that on average, menu cost models generate substantially smaller real effects of nominal shocks than time-dependent models. I investigated a version of the model with random fixed costs as in Dotsey, King, and Wolman [1999], calibrated to match the fraction of small price changes in the CPI. This reduces the importance of the extensive margin effect and increases the average output IRF. The other qualitative features of the model (with and without second moment shocks) were unchanged. In the model with random fixed costs and uncertainty shocks, high uncertainty continues to imply small real effects of nominal shocks.

[23] IRFs are calculated using the ergodic distribution, so these numbers represent the average response to a nominal shock at different percentiles of the uncertainty distribution. The actual response depends on the previous sequence of shocks as well as the pre-shock realization of nominal output.

Why does nominal output have much smaller real effects in times of high uncertainty? As mentioned earlier, in Ss models the price impulse response to a positive nominal shock can be decomposed into two components. The first component is the intensive margin: conditional on adjustment, all firms will raise prices more (or lower less) after a positive nominal shock. The second component is the extensive margin: firms close to raising prices will be pushed into action by a positive nominal shock, and some firms who previously would have lowered prices are pushed into inaction by the shock.

The intensive margin response to a nominal shock is a function of the frequency of adjustment and of the expected change in the price level. The more firms that will be adjusting independent of the nominal shock, and the greater those firms' response to the nominal shock, the larger is the intensive margin. The extensive margin response to a nominal shock is determined by how many firms are near the margin of adjustment and by the width of the adjustment bands. If there are lots of firms that are on the margin of adjusting, then the mix of firms that choose to adjust will vary more in response to a nominal shock and the extensive margin will be stronger. This is the classic "selection effect" emphasized by Golosov and Lucas [2007]. Furthermore, the wider the adjustment bands, the more this effect is amplified to give the total effect of the extensive margin on the price level. If the difference between the average price increase and decrease is small, then shifting mass from price decreases to price increases will have less effect on the overall price level than if the difference between the average price increase and decrease is large. See Caballero and Engel [2007] for a more detailed discussion of these two margins.

From Figure 6 it is clear that both margins become more important in times of high uncertainty. There are more firms in the adjustment region so that the intensive margin grows. In addition, there are more firms near the adjustment bands, and the bands are of greater width, both of which increase the importance of the extensive margin. Thus, both the intensive and extensive margin price response increase in times of high uncertainty, and as the price level becomes more flexible, the real effect of nominal shocks necessarily falls. Table 6 shows the contribution of the intensive and extensive margin to price flexibility at the 10th and 90th percentile of uncertainty. Clearly both margins become more important as uncertainty increases, however, the increase in the extensive margin is substantially

larger than the increase in the intensive margin and accounts for approximately two thirds of the overall increase in price flexibility.

Table 6

Price Impulse Response

Uncertainty	Intensive Margin	Extensive Margin
10th percentile	4%	21%
90th percentile	15%	42%

Price Impulse contributions as a percent of total nominal shock

The relative importance of the wait-and-see effect and the volatility effect can be investigated quantitatively by recomputing impulse responses when these effects are turned off. The direct contribution of the wait-and-see effect can be computed by assuming that firms' bands widen but that volatility does not actually increase. Similarly, the direct contribution of the volatility effect can be computed by holding firm policies constant as uncertainty varies. Finally, there is an interaction between the volatility effect and the wait-and-see effect. If the mass near the adjustment bands is held constant as the adjustment bands widen, the extensive margin becomes more important. The total increase in the price impulse response when moving from the 10th percentile of uncertainty to the 90th percentile of uncertainty can then be decomposed into the contribution of the volatility effect, the contribution of the wait-and-see effect, and an interaction component. Combining all three effects, I find that the volatiltiy effect accounts for 114% of the total increase in the price IRF, the wait-and-see effect accounts for -31%, and the interaction effect accounts for 17%.

Thus, I find that prices are substantially more responsive in times of high uncertainty, and that this is because the volatiltiy effect dominates the wait-and-see effect. Second moment shocks are the driving feature of this time-varying response. Recall that since I assumed that first and second moment shocks are perfectly negatively correlated, I can turn off the second moment shocks and investigate how the output impulse response varies with the cycle. When there are no second moment shocks, at the tenth percentile of aggregate productivity (which corresponds to the 90th percentile of uncertainty in the model with second moment shocks), the output impulse response is identical to that at the 90th percentile of aggregate productivity. With only first moment shocks, there is no relationship between

the real impact of nominal shocks and the cycle: for all values of aggregate productivity the average output impulse response on impact is equal to 60% of the nominal shock. In contrast, if first moment shocks are turned off instead of second moment shocks, time-varying policy responses remain. With only second moment shocks, the output impulse response on impact ranges from 45% at the 90th percentile of uncertainty to 72% at the 10th percentile of uncertainty. Thus, the relationship between uncertainty and the transmission of nominal shocks remains even when there are no first moment shocks. However, it is important to note that, as mentioned in the previous section, when there are only second moment shocks in the model, price dispersion becomes procyclical.

Why does the volatility effect dominate the wait-and-see effect? This is the opposite of the result obtained in the investment model of Bloom [2009], and it is because the estimated fixed costs of adjustment in my model of pricing are substantially lower than those estimated for models of firm investment. In Bloom [2009], the investment fixed cost is estimated to be 1.5% of annual sales with additional irreversible investment resale losses of 33.9%. In contrast, the estimated fixed cost of adjustment in my pricing model with first and second moment shocks is only 0.46% of monthly sales and there are no irreversibilities. When adjustment costs are relatively small, the option value of saving adjustment costs is also relatively small so it has little effect on firm pricing. This can be seen by increasing the fixed costs of adjustment in the model. If fixed costs are multiplied by more than twenty-fold, the policy implications eventually reverse[24] so that the price level is modestly less flexible when uncertainty is large. However, with this reparameterization, the implied frequency of adjustment falls to under 1%, nearly 100% of price changes are increases, and the correlation between frequency and price dispersion becomes negative. This implies that while the wait-and-see effect can be made to dominate with large enough fixed costs, the other predictions of the model are then grossly counterfactual. Thus, the implication that nominal shocks have reduced real effects in times of high uncertainty appears to hold for all plausible parameter values.[25]

---

[24] Changing other parameters lowers the required increase in fixed costs. Aggregate productivity shocks dampen the effects of increases in fixed costs on the wait-and-see motive, as decreases in productivity increase the fraction of firms adjusting as uncertainty rises. A lower elasticity of substitution also decreases the costs of deviating from the optimal price. Thus, lowering the elasticity of substitution and turning off aggregate productivity shocks reduces the increase in fixed costs required. If these parameters are not varied, the required increase in fixed costs is closer to 100-fold.

[25] While the policy implications of uncertainty can be reduced for high enough fixed costs, the model then implies a negative correlation between frequency and dispersion. One might wonder if matching the positive correlation between frequency and dispersion mechanically implies a positive relationship between dispersion and price flexibility. If only the

It should be noted that nothing about the volatility effect dominating in my model is inconsistent with the wait-and-see effect dominating in Bloom [2009]. There is no reason to believe that fixed costs of investment should be similar to fixed costs of price adjustment. Furthermore, it's also interesting to note that due to the difference in our modeling environments I find policy implications similar to Bloom, Floetotto, and Jaimovich [2009] despite finding that different effects of uncertainty shocks dominate. This is because in pricing models, nominal shocks have reduced real effects when price adjustment rises due to uncertainty. In contrast, in Bloom, Floetotto, and Jaimovich [2009], fewer firms adjust investment during uncertain times. However, their model considers real shocks to the cost of investment, which have greater effect when more firms adjust.

As a final exercise, following Bachmann, Caballero, and Engel [2010], I can back out aggregate shocks from the model to fit U.S. economic data and then compute how the output response to nominal shocks varies across time. In order to compute the sequence of shocks that best explains the observed data, I begin from the ergodic distribution and then 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.[26] The implied uncertainty series generated by my model can be compared to uncertainty indices constructed in Bloom, Floetotto, and Jaimovich [2009]. Figure 7 shows that the uncertainty series implied by my model tracks empirical measures of uncertainty, despite being computed from disparate models and data sources.[27]

---

intensive margin response to nominal shocks was operative this would necessarily be the case. However, the extensive margin depends on both the width of the bands and the fraction of firms near the margin of adjustment. Since, in general, the latter may fall when uncertainty rises, there is no mechanical relationship between uncertainty and price flexibility, even if we match the positive relationship between uncertainty and the frequency of adjustment.

[26] The results are insensitive to various detrending methods of the raw empirical data. This data matching procedure exactly matches inflation and output growth. Since the model assumes that uncertainty and TFP are perfectly negatively correlated, a different procedure picks the aggregate shocks to match inflation and dispersion instead of inflation and output growth. Finally, in the extension of the model without perfectly correlated uncertainty and TFP, I can pick the 3 aggregate shocks to exactly match these three series. In all cases, the results were fairly similar.

[27] The uncertainty index in Bloom, Floetotto, and Jaimovich [2009] is the average of several empirical measures of uncertainty in the United States. The idiosyncratic uncertainty index is adapted from Bloom, Floetotto, and Jaimovich [2009] and excludes measures of aggregate uncertainty, since my model only includes idiosyncratic uncertainty. My uncertainty measure is that backed out of my structural model.

Figure 7: Uncertainty Comparisons

Idiosyncratic Uncertainty Index Includes Census Based Measures and Excludes Measures of Aggregate Uncertainty
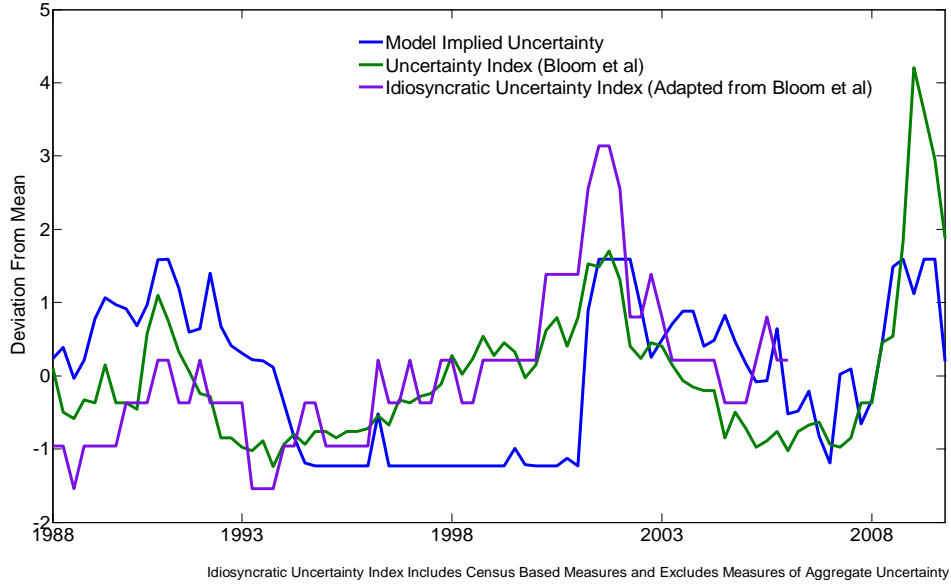
Given the sequence of aggregate shocks, I can then calculate the output impulse response at each date from 1988-2010. As first discussed in the introduction, Figure 1 shows the output impulse response on impact as a measure of the responsiveness of real output to nominal shocks. The responsiveness index is clearly procyclical and plunges in recessions. Thus, the model implies that the inflation-output tradeoff is substantially worse during recessions.

Again, Figure 2 shows the implied output impulse response in September 1995 and October 2001. These are times of very low, and very high uncertainty, respectively. The model with second moment shocks implies that the total response of real output to a nominal shock in October of 2001 is one-quarter of the response in September of 1995. Clearly, there is no such difference for the model with only first moment shocks.

6 Additional Evidence that Price-Flexibility Rises with Uncertainty

My model with second moment shocks generates countercyclical price dispersion and it is able to match the positive correlation between price dispersion and the frequency of adjustment. More importantly, it implies that the price level becomes more flexible in times of high uncertainty. This time-varying price flexibility survives a number of model extensions.[28] Nevertheless, a natural question is the extent to which less structural treatments of the data imply results similar to my model.

I begin by assessing the extent to which the response of import prices to nominal exchange rates varies with uncertainty. Since nominal exchange rate shocks are observable, they are frequently used to test models of price-setting. Since an increase in the nominal exchange rate should increase all importing firms' desired prices, I can test whether this response varies with uncertainty. A typical aggregate pass-through regression estimates

$$\Delta p_t = \alpha + \sum_{i=0}^{2} a_i \Delta e_{t-i} + \sum_{i=0}^{2} b_i \Delta w_{t-i} + \sum_{i=0}^{2} c_i \Delta d_{t-i} + \varepsilon_t$$

where p_t is an index of import prices, e_t is a trade-weighted nominal exchange rate, w_t is a measure of foreign production costs and d_t is a measure of domestic demand.[29] Short-run pass-through is then given by a_0 while long-run[30] pass-through is given by $\sum_{i=0}^{2} a_i$. See Campa and Goldberg [2005] for additional discussion. Since my model predicts that prices should respond more to aggregate nominal exchange rate shocks during times of high uncertainty, I augment this specification by interacting pass-through with indicators for high and low uncertainty regimes.[31] The high uncertainty regime is the one-third of quarters with highest uncertainty according to the Bloom, Floetotto, and Jaimovich

---

[28] See previous footnotes. Extensions include random fixed costs and autocorrelated money growth shocks in order to match the persistence and hump-shaped response to monetary shocks observed in VARs, as well as the addition of uncertainty shocks that are not perfectly negatively correlated with aggregate productivity.

[29] I use quarterly data from 1985-2010. p_t is the BLS import price index excluding petroleum. e_t is the trade-weighted nominal exchange rate constructed by the Federal Reserve, w_t is the OECD (ex US) trade-weighted PPI and d_t is real chained GDP - exports + imports, taken from the BEA.

[30] I have investigated additional lags in the long-run pass-through equation, and it did not substantively affect the results.

[31] Marazzi, Sheets, Vigfusson, Faust, Gagnon, Marquez, Martin, Reeve, and Rogers [2005] argues that exchange rate pass-through has fallen throughout time. Since uncertainty may also have secular (or seasonal) patterns, I also add interactions with a linear time trend as well as a quarter dummy.

[2009] uncertainty index, while the low uncertainty regime is the one-third of quarters with the lowest uncertainty.[32]

Table 7

Short-Run and Long-Run Pass-Through Coefficients

Average SR	Average LR	High - Low Unc SR	High - Low Unc LR
0.18	0.35	0.19	0.46
(.06)	(.07)	(.08)	(.14)

Standard Errors in Parantheses. See footnote 29 for data description. n=104.

Table 7 shows that moving from low uncertainty to high uncertainty regimes increases pass-through by 19 percentage points on impact and by 45 percentage points in the long-run. This is relative to an average pass-through on impact of 18 percent, and a long-run pass-through of 35 percent, so as predicted by my model, there are large effects of uncertainty on price flexibility.   Of course, aggregate pass-through equations of this sort may be driven by changes in composition across time. I have estimated pass-through equations at the 1 and 2 digit industry level and found similar results. Nevertheless, there may be cyclical compositional changes in the mixture of imports at higher levels of disaggregation. Such an investigation requires the micro data underlying the BLS import price indices, and I leave this to future work.

Nominal exchange rate changes are observable and plausibly exogenous from the perspective of individual price setters, and so they are particularly convenient as a test of my model's implications. However, my model and micro facts are about retail pricing rather than import pricing, so I also look at the relationship between uncertainty and monetary policy shocks identified using monetary policy VARs.[33]   Towards that end, I estimate a Factor Augmented VAR as in Bernanke, Boivin, and Eliasz [2005], but I allow factor loadings and VAR coefficients to vary between a high uncertainty and low

---

[32] I also investigated a specification where pass-through was interacted with a continuous uncertainty variable instead of a high/low dummy and found similar large increases in pass-through as uncertainty rises.

[33] Subject to the usual caveats about identification.

uncertainty regime.[34] I choose to estimate a FAVAR as opposed to a simpler VAR to allow for the possibility that the Fed conditions on different information during times of high uncertainty.[35]

Figure 8 shows the impulse response to a loosening of monetary policy and that, consistent with my quantitative model, the estimated price response is significantly greater during high uncertainty regimes than during low uncertainty regimes.   While 90 percent confidence intervals for the overall

Figure 8: Regime-Specific FAVAR

Impulse Response of CPI (Ex Shelter) to 25 Basis Point Decrease in FFR
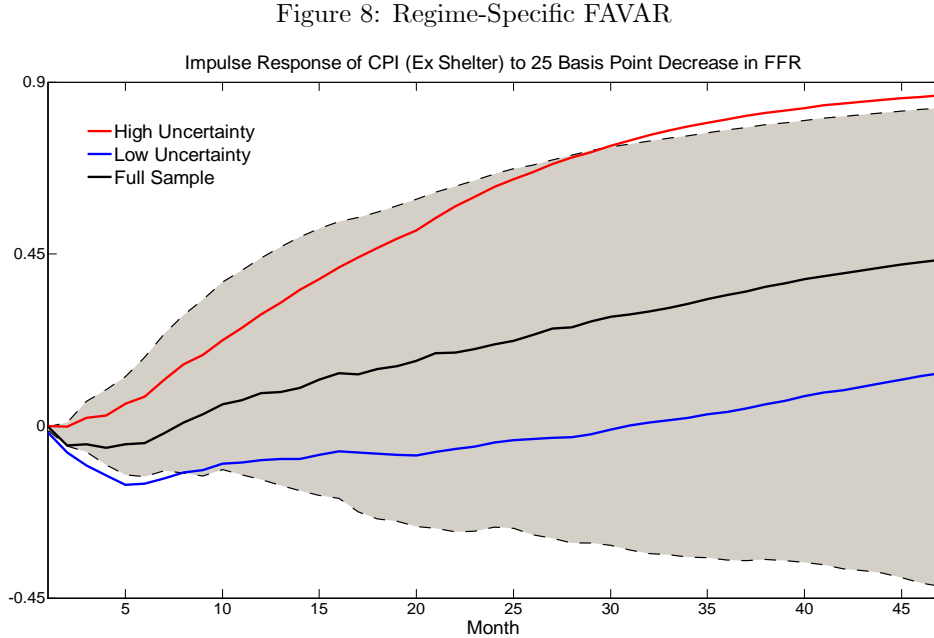
FAVAR are large, the hypothesis that the regime-specific VARs do not differ from the standard linear VAR can be rejected.   These results are consistent with those in Zhang [2011].   She estimates

---

[34]Specifically I split the 1976-2005 sample into thirds using the Bloom, Floetotto, and Jaimovich [2009] uncertainty index.   I then estimate factor loadings separately on the high and low uncertainty regime.   I allow structural VAR coefficients to vary with regimes and estimate these VARs using standard techniques under the identifying assumptions in Bernanke, Boivin, and Eliasz [2005]. The figure in the text assumes 9 factors and 12 lags in the estimating FAVAR. The qualitative results are unchanged with changes in these assumptions although the significance of the results is somewhat sensitive to these assumptions.   The overall non-linear VAR strategy is similar to that in Zhang [2011] and Bachmann and Sims [2011].

[35]Nevertheless, I have found similar results when using a Christiano, Eichenbaum, and Evans [1999] type VAR augmented with housing starts and the purchasing managers index.   The use of these leading indicators is necessary to eliminate the substantial price puzzle present in VARs using data from 1995-present, even after including commodity prices.   While the Christiano, Eichenbaum, and Evans [1999] approach is more restrictive than the FAVAR model, it has the advantage of generating tighter confidence intervals so that equal price responses during times of high and low uncertainty can be more strongly rejected.

similar regime-specific VARs but splits the sample by business cycle state rather than uncertainty and concentrates on the price response of disaggregated sectors. She shows that cross-sector price dispersion is larger during recessions and that disaggregate prices respond more quickly to monetary policy shocks. This VAR evidence again confirms my structural model's implication that the price level should respond more to monetary policy shocks during times of high uncertainty.

7 Conclusions

There is mounting empirical evidence that uncertainty rises during recessions. I estimate uncertainty shocks in a Ss model of price setting to match the empirical behavior of price dispersion documented in Berger and Vavra [2011], and I find that the uncertainty series estimated in my model is broadly similar to the empirical uncertainty index in Bloom, Floetotto, and Jaimovich [2009]. Uncertainty is strongly countercyclical, falling during booms and spiking during recessions. In the presence of adjustment frictions, this can have important implications for the transmission of aggregate shocks. Bloom, Floetotto, and Jaimovich [2009] argue that a fall in investment following an increase in uncertainty may be an important source of business cycle fluctuations. Furthermore, stabilization policies to stimulate investment may have reduced effectiveness when uncertainty is large.

While there is a growing literature that studies the effects of uncertainty in real business cycle models, the implications for monetary policy have received little attention. In this paper, I argue that countercyclical price dispersion can be explained by a Ss model with uncertainty shocks. Furthermore, fixed costs of price adjustment and uncertainty have important interactions that generate time-varying real responses to nominal shocks. I find that greater uncertainty pushes more firms near the margin of adjustment, so that the price level becomes more responsive to nominal shocks and the output effect of these shocks is substantially dampened.

Together, these papers point towards similar policy responses to uncertainty. When uncertainty rises, firms no longer invest, so that policies that affect the real price of investment have little effect. Furthermore, firms adjust more rapidly on the price margin so that monetary policy of normal magnitude becomes less effective. This increase in uncertainty during recessions means that monetary

policy becomes less effective at precisely the time when it is needed most. My model implies that at the height of uncertainty during the most recent recession, nominal shocks had one quarter of the impact on real output as during the relative calm of the mid-nineties.

References

ALTONJI, J. G., AND L. M. SEGAL (1996): "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business & Economic Statistics*, 14(3).

AMIRAULT, D., C. KWAN, AND G. WILKINSON (2006): "Survey of Price-Setting Behaviour of Canadian Companies," Discussion paper, Bank of Canada.

ARELLANO, C., Y. BAI, AND P. KEHOE (2010): "Financial Markets and Fluctuations in Uncertainty," *Minneapolis Fed Paper*.

ASHLEY, R., AND R. J. VERBRUGGE (2007): "Mis-Specification in Phillips Curve Regressions: Quantifying Frequency Dependence in This Relationship While Allowing for Feedback," *Mimeo*.

BACHMANN, R., AND C. BAYER (2009): "Firm-Specific Productivity Risk over the Business Cycle: Facts and Aggregate Implications," *CESifo Working Paper Series 2844*.

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 E. R. SIMS (2011): "Confidence and the Transmission of Government Spending Shocks," Working Paper 17063, National Bureau of Economic Research.

BERGER, D., AND J. VAVRA (2011): "Dynamics of the U.S. Price Distribution," *Yale Mimeo*.

BERNANKE, B., J. BOIVIN, AND P. ELIASZ (2005): "Measuring the Effects of Monetary Policy: A Factor Augmented Vector Autoregressive (FAVAR) Approach," *Quarterly Journal of Economics*, 120(1).

BERTOLA, G., AND R. J. CABALLERO (1990): "Kinked Adjustment Costs and Aggregate Dynamics," *NBER Macroeconomics Annual*, 5.

BILS, M. (2009): "Do Higher Prices for New Goods Reflect Quality Growth or Inflation?," *Quarterly Journal of Economics*, 124(2).

BILS, M., AND P. KLENOW (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112(5).

BILS, M., P. KLENOW, AND B. MALIN (2009): "Reset Price Inflation and the Impact of Monetary Policy Shocks," *National Bureau of Economic Research Working Paper Series*, No. 14787.

BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3).

BLOOM, N., M. FLOETOTTO, AND N. JAIMOVICH (2009): "Really Uncertain Business Cycles," *Mimeo*.

BURSTEIN, A., AND C. HELLWIG (2007): "Prices and Market Shares in a Menu Cost Model," *National Bureau of Economic Research Working Paper Series*, No. 13455.

CABALLERO, R. J. (1993): "Durable Goods: An Explanation for Their Slow Adjustment," *The Journal of Political Economy*, 101(2).

CABALLERO, R. J., AND E. M. ENGEL (2007): "Price stickiness in Ss models: New interpretations of old results," *Journal of Monetary Economics*, 54(Supplement 1).

CAMPA, J. M., AND L. S. GOLDBERG (2005): "Exchange Rate Pass-Through Into Import Prices," *The Review of Economics and Statistics*, 87(4).

CAPLIN, A. S., AND D. F. SPULBER (1987): "Menu Costs and the Neutrality of Money," *The Quarterly Journal of Economics*, 102(4).

CHEVALIER, J., A. KASHYAP, AND P. ROSSI (2003): "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review*, 93(1).

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): "Chapter 2 Monetary policy shocks: What have we learned and to what end?," in *Handbook of Macroeconomics*, vol. 1. Elsevier, 1 edn.

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).

DHYNE, E., L. J. ALVAREZ, H. L. BIHAN, G. VERONESE, D. DIAS, J. HOFFMANN, N. JONKER, P. LUNNEMANN, F. RUMLER, AND VILMUNEN (2005): "Price setting in the euro area: Some stylized facts from Individual Consumer Price Data," Discussion paper, Banco de Espana.

DIXIT, A. (1991): "A simplified treatment of the theory of optimal regulation of Brownian motion," *Journal of Economic Dynamics and Control*, 15(4).

DOTSEY, M., R. KING, AND A. WOLMAN (1999): "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics*, 114(2).

EICHENBAUM, M., N. JAIMOVICH, AND S. REBELO (2009): "Reference Prices, Costs and Nominal Rigidities," *American Economic Review*.

FERNANDEZ-VILLAVERDE, J., J. GUERRO, J. RUBIO-RAMIREZ, AND M. URIBE (2009): "Risk Matters: The Real Effects of Volatility Shocks," *Penn Mimeo*.

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).

JUSTINIANO, A., AND G. PRIMICERI (2008): "Time Varying Volatility of Macroeconomics Fluctuations," *American Economic Review*, 98(3).

KEHOE, P. J., AND V. MIDRIGAN (2008): "Temporary Price Changes and the Real Effects of Monetary Policy," *National Bureau of Economic Research Working Paper Series*, No. 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 (2010): "Microeconomic Evidence on Price-Setting," *National Bureau of Economic Research Working Paper Series*, 15826.

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).

MARAZZI, M., N. SHEETS, R. VIGFUSSON, J. FAUST, J. GAGNON, J. MARQUEZ, R. MARTIN, T. REEVE, AND J. ROGERS (2005): "Exchange Rate Pass-Through to U.S. Import Prices: Some New Evidence," *International Finance Discussion Papers*.

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).

——— (2010): "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).

WOODFORD, M. (2009): "Information-Constrained State-Dependent Pricing," *Journal of Monetary Economics*, 56(1).

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).

ZHANG, F. (2011): "Rational Inattention in Uncertain Business Cycles," *Mimeo*.

8 Appendix 1: Empirical Results

This appendix briefly discusses robustness of the empirical results. See Berger and Vavra [2011] for more details.

The model estimation focuses on the the cross-sectional standard deviation of price changes, conditional on adjustment and its relationship to other statistics. The benchmark empirical standard deviation number is constructed after trimming the top and bottom 2.5% of price changes, as the database contains a large fraction of outliers that are likely to be the result of measurement error. Such errors can occur due to transcription errors or miscoded prices. Nevertheless, similar results were obtained when trimming only the top and bottom 1% or after only excluding price changes of more than 500%, which are likely to be due to measurement error. The interquartile range, which is more robust to outliers, also delivers similar results. While the interquartile range is more robust to outliers and measurement error and is thus perhaps a better empirical measure of price dispersion, it can be problematic in Ss models given that the distribution is highly bimodal so that small changes in the distribution can lead to large jumps in the interquartile range. Thus, the estimation instead uses the standard deviation of price changes.

All correlations reported in Table 1 were computed using Baxter-King bandpass filtered data, seasonally adjusted data. The bandpass filter was chosen because it eliminates high-frequency noise in the price dispersion 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. Seasonal adjustment was computed using deviations from month dummies as in Klenow and Malin [2010].

The benchmark series excludes sales and product substitutions for several reasons. Many recent papers[36] argue that the behavior of prices after excluding sales is likely to be more relevant for

---

[36]E.g. Eichenbaum, Jaimovich, and Rebelo [2009], Kehoe and Midrigan [2008] and Guimaraes and Sheedy [2011].

monetary policy. Furthermore, Bils [2009] argues that product substitutions induce a quantitatively significant source of measurement error into price series. Nevertheless, results are similar for price series which also include product substitutions and sales.

Dispersion numbers reported in the body of the text are conditional on price adjustment so that zeros are excluded. This is because, as shown in the analytical section, Ss models with only first moment shocks imply a negative correlation between the frequency of adjustment and the standard deviation of price changes (excluding zeros). In contrast, the implications of Ss models for the relationship between frequency and the standard deviation of price changes (including zeros) are less stark and cannot be characterized analytically.[37]

Nevertheless, the empirical behavior of price dispersion including zeros is similar to price dispersion excluding zeros. Dispersion including zeros is strongly countercyclical, and unsurprisingly, it has an extremely high correlation with the frequency of adjustment. Reestimating models to match this fact did not qualitatively alter the conclusions.

The positive correlation between the frequency of adjustment and price dispersion reported in Table 1 is computed using the average frequency of adjustment. It is well known that the median frequency of adjustment differs from the mean frequency of adjustment due to heterogeneity across items. Recomputing the correlation between the median frequency of price adjustment and the dispersion of price changes produces an even stronger positive correlation.

Finally, the benchmark empirical results focus on monthly data, which restricts the analysis to New York, Los Angeles and Chicago. If prices in between observations are imputed to remain constant then the analysis can be performed on the entire sample. While using the entire sample improves representativeness, it likely introduces additional measurement error in attributing price changes across time. Nevertheless, results were again similar when using the full sample.

---

[37] In the estimated model, the correlation between dispersion (including zeros) and frequency is strongly negative.

9 Appendix 2: Analytical Model

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

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

We now wish to show that if $f_{up} > f_{down}$ then the variance of price changes and the frequency of price changes are negatively correlated in response to shocks to inflation. In order to do this, we will take an optimal policy U, L as given. We can then compute how the frequency and variance of prices will change after a positive shock to inflation by holding the optimal policy constant and computing changes in the ergodic distribution in response to inflation.

It can then be shown[38] that the ergodic density is given by:

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

---

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

with

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

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

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

so that the total frequency of adjustment is equal to $\frac{\sigma^2}{2}\frac{\alpha}{K}\left[e^{\alpha U} - e^{\alpha L}\right]$. Differentiating with respect to α gives

$$
\begin{aligned}
\frac{\partial\left(f_{up} + f_{down}\right)}{\partial\alpha} &\propto \left(e^{\alpha U} - e^{\alpha L}\right)\left(1 - \frac{\alpha}{K}\frac{\partial K}{\partial\alpha}\right) + \alpha U e^{\alpha U} - e^{\alpha L}\alpha L \\
&= \left(e^{\alpha U} - e^{\alpha L}\right)\left(1 - \frac{(\alpha U)^2 e^{\alpha U}(1 - e^{\alpha L}) - (\alpha L)^2 e^{\alpha L}(1 - e^{\alpha U})}{\alpha U e^{\alpha U}(1 - e^{\alpha L}) - \alpha L e^{\alpha L}(1 - e^{\alpha U})}\right) + \alpha U e^{\alpha U} - e^{\alpha L}\alpha L \\
&< 0 \text{ for } e^{\alpha U} e^{\alpha L} < \frac{e^{\alpha U} + e^{\alpha L}}{2}.
\end{aligned}
$$

Now if $f_{up} > f_{down}$ then

$$
-\frac{\sigma^2}{2}\alpha\left(1 - e^{\alpha U}\right) e^{\alpha L}/K > \frac{\sigma^2}{2}\alpha\left(1 - e^{\alpha L}\right) e^{\alpha U}/K
$$

$$
\Rightarrow
$$

$$
e^{\alpha U} e^{\alpha L} < \frac{e^{\alpha U} + e^{\alpha L}}{2}
$$

since $\alpha < 0$ and $K > 0$. Thus if, $f_{up} > f_{down}$ then $\frac{\partial(f_{up} + f_{down})}{\partial\alpha} < 0$. Since an increase in inflation

decreases α, we can thus conclude that $\frac{\partial (f_{up} + f_{down})}{\partial \pi} > 0$.

The variance of price changes is given by[39]

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

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

10 Appendix 3: Computational Procedure and Estimation

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 by[40]

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

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, which given S, we can write as: $\chi \left(\log \frac{P_{-1}}{S}, \log z, \log a \right)$, but I conjecture that firms can forecast aggregate variables with a log-

---

[39] Note that as mentioned in the body of the text, the simplified expression is essentially a Bernoulli variance with $p = frac\ up$.
[40] Note that the household labor supply problem implies that the real wage is equal to ωC.

linear rule in the aggregate state and the first moment of the distribution $\chi_1 = mean \left(\log \left[\frac{P_{-1}}{S} za \right] \right)$. Using the law of large numbers, $\chi_1 = \log \frac{P_{-1}}{S} + \log a$ up to a constant. That implies that firms' forecasting rule is:

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

Since firms will need to forecast all aggregate variables, note that

$$\begin{aligned} \chi_1' &= \log P - \log S' + \log a' \\ &= \log \frac{P}{S} + \log S - \log S' + \log a' \\ &= \gamma_0 + \gamma_1 \log a + [\gamma_2 + \gamma_3 \log a] \chi_1 - (\mu + \varepsilon_M) + \log a', \end{aligned}$$

that

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

and

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

Substituting into the profit function gives

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

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

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

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

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

By assumption,

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

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

and

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

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

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

with

$$V^N\left(\log\frac{p_{-1}}{S}, \log z; \chi_1, \log a\right) = \pi\left(\frac{p_{-1}}{S}, z; \chi_1, a\right)$$

$$+E_{\varepsilon^z, \varepsilon^a, \varepsilon^s} QV\left(\begin{array}{c}\log\frac{p_{-1}}{S} - (\mu + \varepsilon^s), \rho_z \log z + d(a)\sigma_z\varepsilon^z, \\ \gamma_0 + \gamma_1 \log a + [\gamma_2 + \gamma_3 \log a]\chi_1 \\ -(\mu + \varepsilon^s) + \rho_a \log a + \sigma_a\varepsilon^a, \\ \rho_a \log a + \sigma_a\varepsilon^a\end{array}\right)$$

and

$$V^A(\log z; \chi_1, \log a) = -f\omega e^{-(\gamma_0 + \gamma_1 \log a + [\gamma_2 + \gamma_3 \log a]\chi_1)}$$

$$+\max_{p/S}\left[\pi\left(\frac{p}{S}, z; \chi_1\right) + E_{\varepsilon^z, \varepsilon^a, \varepsilon^s} QV\left(\begin{array}{c}\log\frac{p}{S} - (\mu + \varepsilon_M), \rho_z \log z + d(a)\sigma_z\varepsilon^z; \\ \gamma_0 + \gamma_1 \log a + [\gamma_2 + \gamma_3 \log a]\chi_1 \\ -(\mu + \varepsilon^s) + \rho_a \log a + \sigma_a\varepsilon^a, \\ \rho_a \log a + \sigma_a\varepsilon^a\end{array}\right)\right]$$

47

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

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

Once the model is solved for a given conjecture for $\boldsymbol{\gamma}$, a panel of 5000 firms[41] is simulated for 13,300 months[42] with a 100 month burnin. The conjectured law of motion

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

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

In addition, at the best fit parameters, I recomputed a version of the model with significantly greater precision and more thoroughly tested the accuracy of aggregate transition rules. Using the method proposed by Den Haan [2010], I computed the maximum error between the conjectured and simulated law of motion over 10,000 periods. Even over this extremely long time frame the maximum difference between aggregate variables computed using only simulation and those computed only using the conjectured law of motion is less than 0.1%, and the average error is much lower. Results suggest

---

[41] I investigated panels of up to 500,000 firms. Results were unchanged.
[42] 13,300 is 50 replications of the length of the empirical sample window.

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 uncertainty in the model. None of the qualitative conclusions of the model are changed when precision is increased from the benchmark analysis.

10.2 Estimating the Model

The model is estimated using simulated method of moments. For a given set of parameters, the model is simulated for 50 replications of the same length observed in the data, and all statistics are computed by using the same procedure applied to the empirical data. Let M_i be the vector of moments for replication i. I perform a grid search[43] over parameters to minimize the log squared deviation between $\sum \frac{M_i}{50}$ and M_{data} with identity weight matrix. Using the inverse of the variance-covariance matrix of the model moments did not substantitively alter the results, but equal weighting produced more stable numerical results.[44]

Once the best fit pair of parameters was identified, I then used bootstrapping to calculate standard errors and model goodness of fit. 100 bootstrap replications of length 266 were computed from the model with best fit parameters to generate $\widehat{M}_{data,1}, ..., \widehat{M}_{data,100}$. The model was then reestimated on this "fake data" to generate a new set of best fit moments and parameters, which directly yield confidence intervals for the original model. The model with only first moment shocks is strongly rejected. The bootstrapped 95th percentile for the one-sided χ^2 goodness of fit loss statistic is 0.63 while the loss from the best fit parameters is 7.07. Thus, there is no chance that departures of the model with only first moment shocks from the data can be explained by random sampling error from the model. In contrast, the 95th percentile of the χ^2 goodness of fit for the model with second moment shocks is 0.12 while the loss function associated with the estimated parameters is given by 0.0053. Thus, the model with second moment shocks cannot be rejected with 95% confidence.

The model with uncertainty shocks has 6 moments and 5 parameters, so it appears to be over-

---

[43]While other optimization methods have clear advantages over a grid search, the grid search method has the advantage of being parallelizable in the bootstrapping routine.

[44]In some applications, equal weighting also appears to provide better small sample properties. See Altonji and Segal [1996].

identified, but there may be model dependence across moments. I assess this by examining how the moments vary as I search over the parameter space for the optimum. The matrix of model simulated moments is of full rank. While there is some correlation of model moments across parameters, none of the correlations is large enough to raise concern about non-identification.[45] Furthermore, all of the moments appear informative, with substantial variation as parameters are varied.

The six moments chosen in the benchmark analysis were picked because they directly inform the parameters of the model. Nevertheless, there is nothing that precludes the inclusion of additional moments. An alternative estimation procedure included two additional moments: the time-series standard deviation of price dispersion including zeros, and the correlation between price dispersion and output. This did not change the qualitative results. Finally, I also estimated the model using an indirect inference approach. Under this procedure, an auxiliary model is used to capture reduced form relationships in the empirical data. Structural parameters of the model are then calculated to minimize the difference between the auxiliary model parameters calculated using empirical data and auxiliary model parameters calculated using simulated data. The auxiliary model I used was a an ARCH time series model for inflation where I allowed the variance of inflation residuals to vary with cross-sectional price dispersion:

$$\pi_t = \phi\pi_{t-1} + \sigma_t e_t$$

$$\sigma_t = \alpha + \eta d_t$$

$$e_t \sim \text{iid, mean zero, unit variance}$$

$$d_t = \text{Cross-Sectional Std Dev of Prices.}$$

This reduced form model was motivated by the structural model's prediction that inflation should respond more strongly to shocks in times of high uncertainty and thus price dispersion. In addition to the auxiliary model, I also included standard moments to estimate the Ss models. Again, the qualitative predictions were unchanged. The empirical data implies η significantly greater than zero. The model with only first moment shocks implies no relationship while the model with second moment

---

[45] At least locally. As usual, this is not a proof of global identification.

shocks generates $\eta > 0$. The estimated structural parameters are similar when using the indirect inference procedure, and the policy implications of the two models are qualitatively unchanged.