

# Time-Varying Phillips Curves

Joseph Vavra\*

University of Chicago and NBER<sup>†</sup>

December 2013

## Abstract

A growing theoretical literature argues that aggregate price flexibility and the inflation-output tradeoff faced by central banks should rise with microeconomic price change dispersion. However, there is little empirical work testing this prediction. I fill this gap by estimating time-varying forward looking New-Keynesian Phillips Curves (NKPC). I reject a NKPC with constant inflation-output tradeoff in favor of a slope that increases with microeconomic volatility. In contrast, there is no evidence that the inflation-output tradeoff varies with aggregate volatility or the business cycle more generally. Furthermore, I show that greater volatility does not affect price flexibility purely through increases in frequency.

**JEL Classification:** E10, E30, E31, E50

**Keywords:** New Keynesian Phillips Curves, volatility, uncertainty, time-varying price flexibility, menu cost models

---

\*I would like to thank Priyanka Anand, David Berger, Eduardo Engel, John Leahy, Alejandro Justiniano and Marty Eichenbaum for useful suggestions. This is a greatly expanded version of work which first appeared in the first chapter of my dissertation.

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

# 1 Introduction

A growing empirical literature documents that the dispersion of microeconomic price changes as well as other firm level variables varies dramatically across time. Does this variation in microeconomic price dispersion have implications for aggregate price flexibility? A growing theoretical literature answers this question affirmatively. A wide variety of price-setting models imply that increases in microeconomic price churning should be associated with increases in aggregate price flexibility. This in turn typically implies that the trade-off between inflation and output facing central banks worsens: in order to generate the same increase in real output, the central bank must be willing to tolerate higher inflation.

Despite this growing theoretical literature, there is little empirical work testing this key implication. In this paper I provide evidence in support of the theoretical relationship between increases in microeconomic volatility and worsening of the inflation-output tradeoff. It is well-known that a wide variety of price-setting models lead to New Keynesian Phillips Curve (NKPC) of the form:

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

The slope  $\lambda$  of this NKPC directly captures the relationship between real economic activity and inflation. Greater  $\lambda$  implies that a given increase in real activity should be associated with a greater increase in inflation. In a classic paper, Gali and Gertler (1999) provide an empirical methodology for directly estimating this slope.

In this paper, I extend their analysis by exploring whether this slope varies across time.<sup>1</sup> In particular, I estimate a regime-specific NKPC with regimes defined in various different ways. I find that, as predicted by economic theory, I can reject a constant  $\lambda$  in favor of a slope that varies with volatility. However, I show that this is not true for all types of volatility. In particular, I find a strong empirical

---

<sup>1</sup>Stock and Watson (2010) present some evidence in favor of a non-linear relationship between inflation and unemployment. While they do not estimate NKPC and do not look at the relationship with volatility, their results are broadly consistent with mine. However I find that volatility is more important for the slope of the NKPC than the business cycle.

relationship between  $\lambda$  and various forms of *idiosyncratic* volatility, yet I find no relationship between  $\lambda$  and *aggregate* volatility. In addition, I find no significant relationship between  $\lambda$  and the business cycle. Thus, it appears that microeconomic price-churning appears to be of particular significance for aggregate price flexibility.

This is interesting because it is exactly what is predicted by various theoretical papers. The first paper to theoretically link time-varying microeconomic price change dispersion to time-varying aggregate price flexibility is Vavra (2013). He documents a positive relationship between the interquartile range of price changes and the frequency of adjustment, and argues that idiosyncratic volatility shocks can match this relationship in an otherwise standard menu cost model. More importantly, he shows that increases in idiosyncratic volatility lead to increases in aggregate price flexibility. This is partially driven by an increase in the frequency of adjustment, but is mostly driven by a change in the particular firms that choose to adjust. Firms with the largest price changes contribute disproportionately to aggregate inflation and increases in idiosyncratic volatility increase the mass of such firms. Vavra (2013) shows that this "extensive margin" effect accounts for the bulk of the increase in aggregate price flexibility that arises with increases in idiosyncratic volatility. In contrast, increases in the volatility of aggregate shocks are not consistent with microeconomic evidence and have little effect on aggregate price flexibility. In addition, the time-varying inflation-output trade-off in his model is driven by variation in volatility rather than by the business cycle per se. After controlling for the level of idiosyncratic volatility, his model delivers essentially no relationship between business cycles and price flexibility.

While Vavra (2013) is the first theoretical paper to link microeconomic price dispersion with macroeconomic price flexibility, a number of recent papers have reached similar conclusions. In Baley and Blanco (2013), an information friction prevents firms from fully knowing their nominal costs. When firms are very uncertain, they adjust prices strongly in response to new information. In contrast, when firms are more certain of their nominal costs, they are less likely to respond to (noisy) new information. This mechanism generates a positive relationship between the response of prices to shocks and uncertainty. During times of high uncertainty, aggregate price flexibility rises and monetary shocks have smaller real effects.

Berger and Vavra (2013) document a positive relationship between reduced form "pass-through" regressions and microeconomic price change dispersion. Based on this

empirical result they argue that matching the behavior of adjusting prices requires variable markup shocks. These shocks change how "responsive" firms are to changes in their marginal cost. They then show that when changes in variable markups increase responsiveness, firms adjust more strongly to movements in both exchange rates and idiosyncratic costs so that they can match the patterns in the data. In addition, when this variable markup channel delivers greater responsiveness it does so for all changes in nominal cost which means that aggregate price flexibility rises.

Bachmann, Born, Elstner, and Grimme (2013) use German micro data to estimate the response of individual price adjustment to actual measures of uncertainty for individual firms. They find a modest positive relationship between the frequency of price adjustment and this micro level uncertainty. They then build a Calvo model that mechanically links microeconomic uncertainty to the frequency of adjustment and find that it delivers a small but positive relationship between microeconomic price dispersion and aggregate price flexibility. However, their model assumes that uncertainty only affects aggregate price flexibility through its affect on the frequency of adjustment. This is an assumption that can be tested directly by measuring the relationship between uncertainty, frequency and the slope of the Phillips curve.

In the data, to what extent do the aggregate effects of volatility operate through frequency rather than through alternative channels like size? Overall, the data suggests that volatility has important effects not captured by frequency. I show this in two ways: 1) I estimate NKPC that that vary with frequency in addition to varying with volatility. While there is a relationship between frequency and the slope of the Phillips curve, the relationship with volatility is much stronger. 2) I impose additional structure and assume that all changes in the slope of the Phillips curve are driven by variation in frequency as in a Calvo model I then estimate how large the variation in frequency must be to match the variation in the slope of the Phillips curve. After estimating this implied variation in frequency I can then compare it to actual variation in frequency computed from BLS micro data. Overall I find that the implied variation in frequency estimated through this structural specification is substantially larger than the actual variation in frequency in the data. This means that in order to rationalize the degree of time-varying price flexibility in the data, a Calvo model requires counterfactually large variation in frequency. That is, price flexibility varies for reasons not captured by measured frequency. This is consistent with menu cost models that imply a large role for the extensive margin and a more

limited role for the intensive margin (which is proportional to frequency) in explaining price flexibility.

The remainder of the paper proceeds as follows: Section 2 introduces the theoretical motivation and describes the basic estimation strategy. Section 3 presents a variety of results estimating the "reduced form" slope of the NKPC. Section 4 estimates "structural" parameters under the additional assumption that the NKPC arises from a standard Calvo specification and relates the results to alternative structural models. Section 5 concludes.

## 2 Empirical Methodology

### 2.1 Theoretical Motivation

Before discussing time-varying New Keynesian Phillips Curves I briefly review the existing literature deriving and estimating these inflation relationships. A wide variety of theoretical price-setting models give rise to a NKPC. For example, in a Calvo model each firm has some probability  $1 - \theta$  of being able to adjust its price while its price must remain fixed with probability  $\theta$ . Under standard assumptions<sup>2</sup> the aggregate price level (expressed as a percentage deviation around a zero inflation steady-state) evolves as

$$p_t = \theta p_{t-1} + (1 - \theta) p_t^*, \quad (1)$$

where  $p_t^*$  is the optimal reset price of an adjusting firm.<sup>3</sup> If a firm's marginal cost is given by  $mc_t$  and it discounts profits with discount factor  $\beta$  then its optimal reset price is given by

$$p_t^* = (1 - \beta\theta) \sum_{k=0}^{\infty} (\theta\beta)^k E_t \{ mc_{t+k}^n \}. \quad (2)$$

Combining (1) and (2) yields the NKPC:

$$\pi_t = \lambda mc_t + \beta E_t \{ \pi_{t+1} \} \quad (3)$$

---

<sup>2</sup>This discussion follows Galí and Gertler (1999). See Woodford (2003) for much more detail.

<sup>3</sup>With no fundamental heterogeneity, all adjusting firms choose the same reset price.

with

$$\lambda \equiv \frac{(1 - \theta)(1 - \beta\theta)}{\theta}.$$

While this NKPC arises naturally in a Calvo environment, it also arises as a reduced form for various alternative price-setting models. In this sense, it is quite general and agnostic about the underlying mechanism that gives rise to firms' price-setting patterns. For example, Gertler and Leahy (2008) derive a Phillips curve with this functional form in an Ss model.

Vavra (2013) shows that in Ss models with time-varying volatility, a similar NKPC can be derived but with time-varying  $\lambda$ . That is, theoretical models with time-varying microeconomic dispersion predict a NKPC of the form:

$$\pi_t = \lambda_t mc_t + \beta E_t \{ \pi_{t+1} \}. \quad (4)$$

These models predict that  $\lambda_t$  rises with microeconomic volatility. As volatility rises, there is an increase in aggregate price flexibility so that inflation responds more to changes in marginal cost. The remainder of the paper tests whether the constant inflation-output tradeoff in (3) can be rejected in favor of an inflation-output tradeoff that varies with volatility as in (4).

## 2.2 Estimating the NKPC

In a highly influential paper, Gali and Gertler (1999) provide a methodology for estimating (3). Under rational expectations, the error in the forecast of  $\pi_{t+1}$  is uncorrelated with information dated  $t$  and earlier, so equation (3) can be estimated using GMM after instrumenting for  $E_t \{ \pi_{t+1} \}$ .<sup>4</sup> As in Gali et al. (2005), I use the labor share of income in the nonfarm business sector to proxy for marginal cost, and I instrument for  $E_t \{ \pi_{t+1} \}$  using four lags of inflation and two lags of the labor income share, the output gap and wage inflation.<sup>5</sup>

---

<sup>4</sup>A number of papers have been written criticizing GMM in this context. See Gali, Gertler, and Lopez-Salido (2005) for a rejoinder that argues that GMM is an appropriate econometric choice.

<sup>5</sup>I measure inflation using the GDP deflator, wages using nominal compensation per hour in the non-farm business sector and I measure the output gap as the deviation from a quadratic GDP trend. In the benchmark results, I compute deviations in the labor share relative to quadratic time-trend to account for the fact that there has been a trend decline in labor share over the last twenty years. (See e.g. Karabarbounis and Neiman (2013) for a discussion of changing labor share trends and Cespedes, Ochoa, and Soto (2005) for discussion of NKPC estimation in non-stationary environments). Alternative measures, many of which are shown in Table 2, did not substantively

Models with time-varying volatility predict a NKPC of the form (4) so that implementing GMM requires

$$E_t \{(\pi_t - \lambda_t m c_t - \beta \pi_{t+1}) z_t\} = 0,$$

where  $z_t$  is a vector of instruments dated  $t$  and earlier.<sup>6</sup> If  $\lambda_t$  is exogenous and is known by agents when they are forecasting future inflation, this is not particularly problematic. This assumption holds for example in the structural model of Vavra (2013) since  $\lambda_t$  moves one-for-one with exogenous and observable volatility. However, it is easy to think of environments that would violate this moment restriction. However, this is not particularly problematic for my empirical strategy because I am interested in testing specification (4) against the null hypothesis that  $\lambda$  is constant as in (3). Under the null hypothesis that  $\lambda$  is constant across time, the moment conditions in (4) reduce to the standard moment conditions in Gali and Gertler (1999). Thus, under the null hypothesis GMM is valid. Rejecting this null hypothesis does not require taking a stand on the validity of moment conditions in alternative environments.

Testing for a  $\lambda_t$  that varies continuously is not feasible, so I proceed by estimating equation (4) separately for times of high and low volatility. In my benchmark results I pick the periods of high and low volatility by first ranking all periods by the interquartile range of plant-level TFP from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012) although I will show extensive robustness checks with respect to alternative measures of volatility. I then call the one-third of periods with the highest value for the interquartile range high volatility periods and the one-third of periods with the lowest value for the interquartile range low volatility periods. The Bloom et al. (2012) measure of cross-sectional volatility is based on census data from 1972-2009, so my estimates cover these dates.

---

affect the results.

<sup>6</sup>As emphasized by e.g. Stock, Wright, and Yogo (2002), reliable estimation can be hampered when  $z_t$  is a weak instrument. However, this weak instrument problem is more important when estimating "hybrid" NKPC with backward looking terms than when estimating pure forward looking Phillips curves. In hybrid specifications, in order to be a strong instrument,  $z_t$  must provide substantial marginal predictive content for  $\pi_{t+1}$  after controlling for  $m c_t$ ,  $\pi_t$  and  $\pi_{t-1}$ . In an empirical model with a purely forward looking specification,  $z_t$  need only satisfy the condition that it improve on a specification without backward looking components. In my applications, there is no evidence for a weak instrument problem. For example, Stock and Yogo (2005) tests very strongly reject weak instrument nulls.

After splitting the sample into a high and low regime I estimate equation (3) separately for each regime and test the hypothesis that  $\lambda_{high} = \lambda_{low}$ . Again, under the null hypothesis that  $\lambda$  is constant across time, estimating (3) in subsamples is just as valid as estimating (3) over the entire sample and should deliver the same  $\lambda$ , while rejecting  $\lambda_{high} = \lambda_{low}$  is evidence against a constant inflation-output tradeoff.

### 3 Results (Reduced Form)

Table 1 shows results for the regime-specific NKPC under various alternative assumptions. Specification (1) estimates the NKPC separately for the highest and lowest one-third of months by the IQR of plant-level TFP shocks allowing both  $\lambda$  and  $\beta$  to vary between the two regimes. The main object of interest is  $\lambda_{high} - \lambda_{low}$ , and the data strongly rejects the hypothesis of constant  $\lambda$ . The estimate of  $\lambda$  is significantly larger (both economically and statistically) during the higher volatility periods.

While specification (1) is the least restrictive specification, a number of models imply additional restrictions on the relationship between parameters in the two different regimes. For example, many models imply that  $\beta$  should be constant in the two regimes. Specification (2) imposes this restriction and shows that it does not change the conclusion, and neither does imposing that  $\beta = 1$ . Specifications (4)-(6) show that estimates of  $\lambda$  continue to be significantly greater during the high volatility regime when using different thresholds to identify high and low volatility.<sup>7</sup> While the point estimates for  $\lambda$  are somewhat different under the two different threshold definitions, this is not particularly surprising given the small sample sizes. In addition, these differences in point estimates when splitting the sample into fourths rather than thirds are not statistically significant. The important point is that  $\lambda_{high} - \lambda_{low}$  continues to be positive and statistically significant under both threshold definitions.

In Section 4 I impose additional model structure on the NKPC and discuss implications of these differences for more fundamental parameters. However, it is worth noting that when using the baseline results, I find estimates of  $\beta$  that are close to conventional values of 0.99. The estimates of  $\lambda_{low}$  are typically negative, which may seem puzzling. However, the estimates are often insignificant, and in Section 4 I show that the exact level of the estimate of  $\lambda_{low}$  is somewhat sensitive to the exact moment conditions used for estimation. In particular, estimating the more struc-

---

<sup>7</sup>For brevity I only report these two thresholds but similar results hold for alternative definitions.



tured specification produces reasonable values for the frequency of adjustment as well as moderately positive estimates of  $\lambda_{low}$ . Thus, across estimation specifications, the average estimated slope of the NKPC is not particularly stable but that is not the focus of my empirical investigation. In contrast, the estimated difference between high and low regimes remains highly significant and robust across all specifications.

Table 1  
Estimated Phillips Curves

	$\lambda_{low}$	$\lambda_{high} - \lambda_{low}$	$\beta_{low}$	$\beta_{high} - \beta_{low}$	$n$
Split thirds by IQR TFP					
(1) No restrictions	-0.082** (0.041)	0.121*** (0.042)	0.998*** (0.028)	0.015 (0.038)	91
(2) $\beta_{high} = \beta_{low}$	-0.082* (0.042)	0.121*** (0.043)	1.003*** (0.017)		91
(3) $\beta_{high} = \beta_{low} = 1$	-0.082* (0.043)	0.121*** (0.044)			91
Split fourths by IQR TFP					
(4) No restrictions	-0.015 (0.015)	0.046*** (0.017)	1.034*** (0.019)	-0.018 (0.033)	71
(5) $\beta_{high} = \beta_{low}$	-0.015 (0.015)	0.045*** (0.017)	1.030*** (0.015)		71
(6) $\beta_{high} = \beta_{low} = 1$	-0.032* (0.019)	0.061*** (0.022)			71

This table reports GMM estimates of parameters of Eq. 4. IQR TFP is the cross-sectional distribution of TFP shocks computed in Bloom et al 2013. Instruments used are four lags of GDP deflator inflation, and two lags of labor income share, quadratic detrended GDP and wage inflation. A Newey-West covariance matrix with 12 lags was used to compute standard errors. The sample period is quarterly from 1972-2009.

"No restrictions" allows for different discount rates in the high and low regime.

How sensitive are these results to various econometric specifications? Table 2 shows a number of robustness checks. Except where specifically indicated, all results use the baseline sample split and instruments. For simplicity I report only results

under the assumption that  $\beta_{high} = \beta_{low}$ , but similar results obtain when allowing  $\beta$  to vary with the regimes or when restricting  $\beta$  to equal one.

**Table 2**  
Estimated Phillips Curves-Alternative Specifications

	$\lambda_{low}$	$\lambda_{high} - \lambda_{low}$	$\beta$	$n$
(1) NFB Inflation	-0.056** (0.022)	0.124*** (0.026)	0.945*** (0.014)	91
(2) No NFBSLS detrending	-.005 (0.009)	0.029** (0.012)	1.033*** (0.017)	91
(3) Additional Instrument Lags	-0.058* (0.033)	0.105*** (0.032)	0.991*** (0.012)	91
(4) Additional Instruments	-0.032 (0.027)	0.081*** (0.026)	1.000*** (0.012)	91
(5) Regime Specific Inflation	-0.083** (0.042)	0.118*** (0.042)	0.991*** (0.028)	91

---

This table reports GMM estimates of parameters of Eq. 4. All specifications split IQR TFP in thirds. "NFB Inflation" uses the non-farm business deflator instead of the GDP deflator.

"No NFBSLS detrending" uses the level of the labor share as the driving variable instead of linear detrending. "Additional Instrument lags" includes 4 lags of all instruments instead of 2.

"Additional instruments" includes commodity price inflation and long-short spreads as additional instruments. "Regime Specific Inflation" allows for regimes to differ permanently in the level of inflation. A Newey-West covariance matrix with 12 lags was used to compute standard errors.

Line (1) of Table 2 shows that the baseline result is not sensitive to the particular measure of aggregate inflation. Using the non-farm business deflator instead of the overall GDP deflator, if anything, delivers stronger results. As mentioned in Section 2, I remove very low frequency trends in the labor share since there has been a trend decline in this variable in recent years, which makes the raw series problematic as a measure of cyclical shocks to marginal cost.<sup>8</sup> Line (2) shows that the results continue

---

<sup>8</sup>See Cespedes et al. (2005) for related discussion of non-stationary NKPC.

to hold using the raw series, although they are somewhat weaker.<sup>9</sup> In Line (3) I show results when using four lags instead of two lags for all instruments, and in line (4) I include two lags of commodity price inflation and long-short interest rate spreads as additional instruments, as in Gali and Gertler (1999). Finally, in line (5) I allow for a regime-specific inflation trend.<sup>10</sup> While this specification does not satisfy the basic specification in equation (3), there is some concern that if equation (3) is misspecified there may be differences in the average level of inflation across regimes. Not allowing for such differences might spuriously drive differences in  $\lambda$  across regimes. However, allowing for level differences in inflation across regimes does not change the results.

To this point, all results have proxied for microeconomic volatility using the interquartile range of plant-level TFP shocks computed in Bloom et al. (2012). I use this series as the benchmark because it is available for a relatively long period of time and closely proxies idiosyncratic volatility since it measures cross-sectional differences. Nevertheless, this series is not without problems. First, theoretical models predict a positive relationship between the dispersion of microeconomic price changes and aggregate price flexibility. Existing evidence suggests that there is a strong positive relationship between microeconomic TFP dispersion and microeconomic price change dispersion, but the relationship is not perfect. In addition, the Bloom et al. (2012) measure uses manufacturing data from the census. This implies that this measure focuses on a narrow subset of the economy. In addition, the nature of the data means that the series can only be constructed annually. With these concerns in mind, Table 3 redoes the baseline analysis using alternative measures of volatility.

In row (1), I directly compute a time-series for the interquartile range of price changes in BLS micro data. While this might appear to be a more natural benchmark, this series has several issues. Producer Price micro data is only available beginning in 1981, Consumer Price data in 1988 and Import Price data in 1993. Thus the overall series is substantially shorter, and constructing the longest series requires merging data from several sources. Nevertheless, the baseline results are unchanged: during periods of time with a high interquartile range of price changes, the estimated slope of the NKPC is significantly larger. In row (2), I again use data from Bloom et al. (2012) but use the interquartile range of firm sales rather than the interquartile range

---

<sup>9</sup>Overall the NKPC fits the data less well over the last 20 years when not accounting for the trend in labor share.

<sup>10</sup>That is, I estimate equation (3) with the addition of a regime-specific constant.

of TFP shocks. Finally, in row (3), I use the dispersion of NAICS industry growth rates as the measure of volatility. That is, I compute the interquartile range of various sectoral indices. This measure of dispersion captures only differences across sectors and not differences across firms within sectors, but it generates similar results.

**Table 3**  
Alternative Measures of Volatility

	$\lambda_{low}$	$\lambda_{high} - \lambda_{low}$	$\beta$	$n$
Idiosyncratic:				
(1) Price IQR	-0.054*** (0.016)	0.091*** (0.022)	0.966*** (0.020)	77
(2) IQR Sales	-0.037 (0.027)	0.058** (0.029)	1.013*** (0.014)	91
(3) Industry IQR	-0.021** (0.010)	0.057*** (0.014)	0.987*** (0.013)	100
Aggregate:				
(4) GDP Volatility	-0.001 (0.010)	0.009 (0.017)	1.010*** (0.011)	127
(5) Stock Market Vol	0.0178 (0.014)	-0.012 (0.019)	1.007*** (0.010)	121
(6) Policy Uncertainty	0.003 (0.009)	0.015 (0.012)	1.006 (0.012)	131

This table reports GMM estimates of parameters of Eq. 4. All specifications split IQR TFP in thirds. Price IQR uses the average interquartile range in BLS data: 1981-2011. IQR sales uses the interquartile range of firm sales computed by Bloom et al (2012):1972-2009. Industry IQR uses the dispersion of NAICS industry growth rates 1972-2011. GDP volatility is the rolling standard deviation of GDP growth in previous 8 quarters: 1960-2011. Stock market volatility taken from Bloom 2009, 1962-2008. Policy uncertainty taken from Bloom et al (2013): 1960-2011. All samples split into the highest and lowest one-third quarters. A Newey-West covariance matrix with 12 lags was used to compute standard errors.

In the second half of Table 3 I show results that focus on aggregate volatility instead of idiosyncratic volatility. In particular, I use three proxies for aggregate volatility. The first measure of aggregate volatility is the standard deviation of real GDP growth in an 8 quarter moving window. Second, I use the stock market volatility measure in Bloom (2009).<sup>11</sup> Finally, I use the measure of policy uncertainty created by Baker, Bloom, and Davis (2013). Rows 4-6 show that none of these measures of aggregate volatility have significant effects on estimates of the slope of the NKPC. This is consistent with the theoretical model in Vavra (2013). In that model, aggregate price flexibility is primarily determined by the cross-sectional dispersion of firms' desired price changes. Increases in idiosyncratic volatility spread out this distribution and drive a positive relationship between microeconomic dispersion and price flexibility. In contrast, increases in aggregate volatility have quantitatively small effects on the distribution of firms' desired price changes and thus have little effect on aggregate price flexibility.

While I have shown a strong relationship between idiosyncratic volatility and the slope of the Phillips curve, perhaps idiosyncratic volatility is proxying for some alternative channel. Using CPI micro data, Vavra (2013) shows that the dispersion of price changes is countercyclical and is positively correlated with the frequency of adjustment. In addition, Stock and Watson (2010) argue that unemployment is more useful for forecasting inflation during recessions.<sup>12</sup> Perhaps the positive relationship between the slope of the Phillips curve and volatility that I document is actually capturing something about the business cycle or about the frequency of adjustment. Table 4 shows results when defining regimes based on measures of the business cycle as well as for the frequency of adjustment. While the point estimates for the GDP gap and GDP growth are negative (showing that  $\lambda$  tends to rise in recessions), this relationship is not significant. Thus, just as in the structural models, it appears that it is indeed time-varying volatility that drives time-variation in the reduced form Phillips curve rather than the business cycle, per se.

---

<sup>11</sup>Since this is the volatility of the overall stock market, it should primarily reflect aggregate volatility rather than idiosyncratic volatility.

<sup>12</sup>Phrased in terms of Phillips curves, their results imply a steeper Phillips curve during recessions.

### Table 4

Alternative Sample Splits: Capturing Something Besides Volatility?

	$\lambda_{low}$	$\lambda_{high} - \lambda_{low}$	$\beta$	$n$
(1) GDP Gap	0.011 (0.012)	-0.006 (0.014)	1.001*** (0.009)	97
(2) GDP Growth	0.028* (0.016)	-0.008 (0.020)	1.016*** (0.013)	97
(3) Frequency	-0.028*** (0.010)	0.046*** (0.013)	1.003*** (0.015)	77
(4) IQR TFP ex ZLB	-0.082** (0.038)	0.109*** (0.041)	0.996 (0.016)	85

This table reports GMM estimates of parameters of Eq. 4. All specifications split sample into thirds. "GDP gap" splits according to the value of the gdp gap, "GDP growth" splits by the growth rate of real GDP. "Frequency" splits the sample using the frequency of adjustment in BLS data, and "IQR TFP ex ZLB" redoes the benchmark calculation excluding 2008q2-present. A Newey-West covariance matrix with 12 lags was used to compute standard errors.

There is a positive relationship between frequency and the slope of the Phillips curve, but this effect is substantially smaller than the effect of volatility. Finally, the fourth specification in Table 4 excludes the time period from 2008q2 to the present. A number of models suggest that inflation at the ZLB should be different than during normal times. Since the financial crisis was also a time of extremely large volatility, it is important to show that my results are not driven by this period, and indeed they are not.

Finally, theoretical models with time-varying volatility predict a forward looking NKPC with time-varying slope as in equation (3), but a number of papers studying average inflation dynamics have argued for the importance of backward looking inflation terms. With this in mind, I now estimate

$$\pi_t = \lambda_t mc_t + \gamma_t^{forward} E_t \{ \pi_{t+1} \} + \gamma^{backward} \pi_{t-1}. \quad (5)$$

To my knowledge, no one has explored a theoretical model with time-varying volatility and both forward and backward looking agents, so it is not obvious if such

a specification would be generated by simple theories. Nevertheless, Table 5 shows that allowing for backward looking inflation terms does not change my basic results. As in the benchmark results, specification (1) allows for different effects of inflation expectations in the two-different regimes and specification (2) restricts these forward looking terms to be identical in both regimes, and in both cases  $\lambda_{high} - \lambda_{low} > 0$ /

Table 5  
Hybrid Specifications

	$\lambda_{low}$	$\lambda_{high} - \lambda_{low}$	$\gamma_{low}^{forward}$	$\gamma_{low}^{forward}$ $-\gamma_{low}^{forward}$	$\gamma_{low}^{backward}$	$n$
(1) No restrictions	-0.049* (0.029)	0.054* (0.030)	0.774*** (0.099)	-0.016 (0.037)	0.271*** (0.092)	91
(2) $\gamma_{low}^{forward} = \gamma_{low}^{backward}$	-0.061* (0.032)	0.074** (0.032)	0.843*** (0.086)		0.180*** (0.085)	91

This table reports GMM estimates of parameters of Eq. 5. Both regressions split regimes by thirds of IQR TFP. Instrument four lags of GDP deflator inflation, and two lags of labor income share, quadratic detrended GDP and wage inflation.

A Newey-West covariance matrix with 12 lags was used to compute standard errors. The sample period is quarterly from 1972-2009.

"No restrictions" allows for different discount rates for forward looking agents in the high and low regime.

## 4 Results (Structural)

To now I have concentrated on estimating the reduced form NKPC

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

However, by imposing additional model structure I can estimate more interpretable parameters. For example, in the Calvo model:

$$\lambda \equiv \frac{(1 - \theta)(1 - \beta\theta)}{\theta},$$

where  $1 - \theta$  is the probability that a firm can adjust its price. In the Ss model of Gertler and Leahy (2008), an identical functional form arises with the probability of a firm receiving a Poisson shock,  $1 - \alpha$ , replacing the probability that a firm exogenously gets to adjust its price  $1 - \theta$ . With these models in mind I now estimate

$$\pi_t = \frac{(1 - \theta)(1 - \beta\theta)}{\theta} mc_t + \beta E_t \{ \pi_{t+1} \}, \quad (6)$$

again splitting the sample into different regimes. Table 6 shows the results with regimes defined using the interquartile range of TFP from Bloom et al. (2012), and the interquartile range of price changes and frequency of adjustment calculating using BLS micro data. For ease of interpretation I focus on the case with constant  $\beta$  but allowing for variation in  $\beta$  did not change the results. The main object of interest is  $(1 - \theta)_{high-low}$ , the increase in the quarterly frequency of adjustment when moving from the low regime to the high regime. In all cases the increase in the implied quarterly frequency of adjustment during the high volatility or frequency regime is large and significant. The quarterly frequency of adjustment rises by 20-30% (corresponding to a decline in the average duration of prices of several quarters).

Table 6

	Structural Specifications					$n$
	$(1 - \theta)_{low}$	$(1 - \theta)_{high}$	$(1 - \theta)_{high-low}$	$\beta$	Freq <sub>high-low</sub>	
(1) TFP IQR	0.079 (0.110)	0.333*** (0.092)	0.254*** (0.024)	1.032*** (0.014)		91
(2) Price IQR	0.194*** (0.039)	0.486*** (0.047)	0.291*** (0.039)	1.070*** (0.030)	0.027*** (0.02)	77
(3) Frequency	0.070 (0.092)	0.277*** (0.072)	0.207*** (0.047)	1.059*** (0.018)	0.168*** (0.012)	77

This table reports GMM estimates of parameters of Eq. 6. All specifications split sample into thirds. See previous tables for variable definitions. A Newey-West covariance matrix with 12 lags was used to compute standard errors.

Thus, when viewed through the lens of a Calvo model, the movements in frequency



implied across time implied by the NKPC are very large. Is this time-variation implied by the Calvo model a good description of reality? One can get a sense of this by comparing the implied frequency needed to rationalize variation in the NKPC to actual variation in frequency. (Note that this is only possible to compute for the sample window corresponding to BLS data). Clearly when using the price change IQR to define regimes, the variation in frequency in a Calvo model needed to rationalize the time-variation in  $\lambda$  is unrealistically large. When regimes are defined using the observed frequency of adjustment, the implied variation in frequency is still larger than actual variation in frequency although this difference is less extreme.

What does this mean? It implies that the variation in aggregate price flexibility across time cannot be explained solely by variation in the frequency of adjustment. This is particularly true when sorting regimes by the interquartile range of price changes. In order to generate the increase in price flexibility observed in the high volatility regime in the data, a Calvo model would require a 30% increase in the frequency of adjustment while the true increase is only 3%. This result is again consistent with Ss models of price-setting. Since the Gertler and Leahy (2008) model has an identical functional form to equation (6) with  $\alpha$  replacing  $\theta$ , all of the estimates of  $\theta$  in Table 6 can be reinterpreted as estimates of  $\alpha$ . That is, increases in the probability of receiving Poisson shocks during high volatility periods can lead to increases in the slope of the Phillips curve without large increases in the frequency of adjustment. This is also consistent with the quantitative results in Vavra (2013). He argues that only around 20% of the time-varying price flexibility in his model is driven by movements in the frequency of adjustment, with the remainder being driven by variation in the mix of which firms choose to adjust.

This implies that the empirical frequency of adjustment is not a good summary statistic for aggregate price flexibility. This in turn implies that models that only allow volatility to affect aggregate price flexibility through its affect on frequency are likely to provide a lower bound on the actual importance of volatility in the data. For example, Bachmann et al. (2013) use micro data to estimate the elasticity of firm-level frequency with respect to firm-level uncertainty. After estimating this elasticity, they calibrate a Calvo model with exogenous frequency to match this elasticity and conclude that variation in price flexibility is modest. However, volatility in this model can only affect price flexibility through its effect on frequency.

## 5 Conclusion

A growing theoretical literature argues that increases in microeconomic price change dispersion should lead to increases in aggregate price flexibility and worsening of the inflation-output tradeoff. In this paper I estimate regime-specific NKPC to assess this theoretical prediction. The results are highly supportive of this theoretical channel. Overall, the time-series evidence shows that increases in marginal cost lead to greater increases in inflation during periods of idiosyncratic volatility. This is predicted by Ss or imperfect information models with volatility shocks while it is at odds with older Keynesian models that predict that the slope of the Phillips curve should be procyclical. It is also at odds with Calvo price-setting models which assume a constant frequency of adjustment. Since the Calvo model is log-linear, increases in volatility have no effects on the inflation-output tradeoff unless they affect frequency.

I also provide more structural results to assess the role of frequency variation in explaining aggregate price flexibility. My results suggest that volatility has effects on aggregate price flexibility that are not captured purely by changes in frequency. Together these results suggest that policy recommendations based on a constant inflation-output tradeoff are unlikely to be particularly informative. The inflation-output tradeoff tends to worsen during the times when central banks are more likely to want to stimulate the economy.

## References

- Bachmann, R., B. Born, S. Elstner, and C. Grimme (2013). Time-varying business volatility, price setting, and the real effects of monetary policy. *Mimeo*.
- Baker, S., N. Bloom, and S. Davis (2013). Measuring economic policy uncertainty. *Mimeo*.
- Baley, I. and J. Blanco (2013). Learning to price. *NYU Mimeo*.

- Berger, D. and J. Vavra (2013). Volatility and pass-through. *NBER Working Paper 19651*.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica* 77(3).
- Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. Terry (2012). Really uncertain business cycles. *NBER Working Paper 18245*.
- Céspedes, L., M. Ochoa, and C. Soto (2005). The new keynesian phillips curve in an emerging market economy: The case of chile.
- Gali, J. and M. Gertler (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics* 44(2).
- Gali, J., M. Gertler, and D. Lopez-Salido (2005). Robustness of the estimates of the hybrid new keynesian phillips curve. *Journal of Monetary Economics* 52(6).
- Gertler, M. and J. Leahy (2008). A phillips curve with an ss foundation. *Journal of Political Economy* 116(3).
- Karabarbounis, L. and B. Neiman (2013). The global decline of the labor share. *NBER Working Paper 19136*.
- Stock, J. and M. Watson (2010). Modeling inflation after the crisis. *NBER Working Paper 16488*.
- Stock, J., J. Wright, and M. Yogo (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics*.
- Stock, J. and M. Yogo (2005). Testing for weak instruments in linear IV regression. *Identification and Inference for Econometric Models*.
- Vavra, J. (2013). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation.
- Woodford, M. (2003, August). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.