A "Local" Model of the Firm: Sticky prices and the Phillips Curve

Clayton Daley

11. July 2007

A "Local" Model of the Firm: Sticky Prices and the Phillips Curve
by Clayton Daley

Introduction

While this paper has been constructed "forward" from theory to reality, it must be honestly stated that the discovery of its mechanism was, in fact, primarily the result of working backwards through the problem and identification of the final piece of the puzzle was indeed fortuitous. However, as the logic was carried to its completion and compared with other literature on price stickiness and the Phillips curve, the ability of a single model to capture the themes of other approaches was striking and undeniable.

Once recognized, this kernel of knowledge proved even more difficult to communicate. Many of the features of the model require complex dynamic decision making and would perhaps be applied, in real life, using specialized and complex decision rules. As this level of complexity is unnecessary to recognize the basic truth of the argument and explore its implications, it became desirable to selectively simplify the model. Unfortunately, each time a major complexity was eliminated, it took with it one of the elements of behavior that make the explanation so appealing.

Consequently, this document will start with the underlying principle of the final mechanism and a basic outline of that mechanism. Once the principle has been communicated as plainly as possible, a model will be constructed and the effects will be illustrated. In many cases, complexity will be avoided by looking at the extremes of a simpler model and explaining why the reality should, in fact, be somewhere in between. Finally, the congruence between this model and existing sticky price explanations will be explored with the hope that an economist who is gifted in complex dynamics will recognize the deep links and explore the idea and its empirical application more thoroughly.

Background

A comment made in passing with the CFO of a Fortune 500 firm captured perhaps the critical problem with typical macro-economic models and assumptions. Faced with an increase in aggregate demand, it would seem natural for an economist to assume that the supply curve slopes upwards and that, given their awareness of this information, firms would behave as if this were true. However, this CFO was adamant that this assumption would never be used as a decision-rule in a business, and indeed that this basic assumption often isn't even the case for the firm.

This might strike a casual economist as a shocking statement and without further thought, could be dismissed as an oddity or perhaps a misunderstanding. Digging more deeply into the sentiment, however, offers a much different perspective on the critical issues that a business faces. Despite the relative ease and limited expense of obtaining macroeconomic conditions, business can rarely if ever safely use this information. This is in stark contrast to the attempt, by economists, to use these variables to "inform" businesses.

This is not itself a novel idea, with special emphasis on Lucas (1972), but Gordon (1981) captures its limitation as, "A crucial weakness of the Lucas (1972) and Barro (1976) papers is that there is no device to generate persistence of output movements as observed in real-world business cycles." Meltzer (1995) offers alternative explanations for the same end, "As in Bomhoff (1983), a principal difficulty in interpreting information is uncertainty about how long changes will persist. This is the central idea developed in Brunner, Cukierman and Meltzer (1983), but we took the idea from Muth's (1961) seminal paper on rational expectations. In Meltzer (1982), I used these ideas to discuss price setting." While this offers temporary rigidity, it is again difficult to explain the persistence factor.

If the proper mechanism were found, however, the possibility that local shocks dominate business behavior (regardless of the price reasoning) is and remains a striking justification for the sentiment offered.[1] Or in Gordon's (1981) words, "For the models of the new classicists to generate a business cycle, each agent must be equipped with a pair of blinders that arbitrarily cuts him off from information printed in the daily newspaper on aggregate variables like interest rates, price indexes, and the money supply. But to accept the non-market clearing framework we must recognize impediments to price adjustment that prevent agents from promptly responding to aggregate disturbances."

The Foundation

Given that a firm cannot consistently rely on any macroeconomic data to determine its ongoing pricing policies and supply/demand expectations, it is reasonable to identify those few factors that are truly local and truly relevant. On the demand side, it seems reasonable to expect that a company needs to consider changing its policies when the demand for its product shifts (and rarely otherwise). Complicating this decision is the fact that firms face demand uncertainty when the underlying levels of demand have not actually changed. Once detected and assuming (momentarily) an upward sloping supply curve, this would imply a need to slide the firm's offerings up or down the demand curve to ensure that the company optimizes its markup while balancing supply and demand.[2] With this in mind, the place to look for a business's basic decision rules is within operations and inventory control.

While no specific model was adequate (and simple enough) to address the broad economic question, the search was fruitful. Specifically, it illustrated that the typical inventory control approach is a complex version of statistical hypothesis testing involving

---

[1] Perhaps the best and simple example would be to consider that many modern firms are subject to local labor costs/shocks and shifts in tastes that are independent of macroeconomic variables. All together, it is nearly inconceivable that a firm's demand will remain stable subject to only real variables and sorting out the difference between real and monetary variables, being difficult enough for economists, is impractical and perhaps impossible for an operating business and its employees.

[2] To ensure that we are focusing strictly on the mechanism in question, we will force our firms to optimize markup. Typical "menu costs" concerns will also be deferred and addressed in force later.

the mean of a certain characteristic, demand in the economic problem. The approach in hypothesis testing is to select an alpha, a confidence level, and either select or reject the hypothesis, a very black and white answer.

But businesses are less concerned about a definitive answer and more worried about being as close as possible to reality. If we reverse the typical statistical process, we can instead determine the level of confidence (alpha) of the marginal decision point (where a hypothesis is on the edge of accept and reject). While it may not be statistically accurate, it is close enough to conceive of this percentage as the probability that the particular demand (or sequence of demands) indicates a *change* in the underlying level of demand. If demand changes, a business will, naturally, change its production levels (and associated prices) to restore equilibrium.

The interesting part about this approach is that one-period demand values very close to the old mean produce very low confidence levels of a change, while values very distant from the model produce high confidence levels of a change. If we load this probability into the price change (or change in expected demand), we find that *real* price stickiness (or real expectations stickiness) arises from this uncertainty. Also important to this problem is the reverse, that a larger change will cause a business to detect the change and adjust prices quickly (as the confidence rises quickly) while small changes will delay substantial detection and adjustment for long periods of time.[3] Capturing this complex dynamic mix becomes the first sticking point of the model.

This inventory sensitive approach is so compelling that the typical approach to marginal cost became a problem again. Firms that were completely unsure of the changes in demand are all of the sudden oracles when it comes to the shape and cost of their labor supply. Further, this assumption of upward sloping supply curves violated our very reasonable sentiment (and concern) about shifts in the real equilibrium. If a labor shock has shifted the equilibrium to a higher quantity and a lower price, the last thing we want to do is to choke off demand (or give it to our competitors) by assuming an upward sloping supply curve.

Again, the behavior of the inventory model was so appealing, both in its output and simple handing of real uncertainty, that it seemed reasonable to look for a comparable approach to labor. A smaller, but not insurmountable leap is to view the individuals attracted to an interview as an uncertain look into the real level of supply (as demand experienced in a single period is an uncertain view into underlying demand). The number of candidates, and more importantly the number of good, qualified candidates, is indicative of the job market; however, not every posting attracts a representative number of each of these different groups. This looks very similar to our model of uncertain demand and a comparable model allows us to produce inventory-like real stickiness in labor markets without introducing additional complexity. When these inventories fall for

---

[3] The basic MV=PQ suggests that deviations in output caused by monetary policy must eventually work their way out through inflation. Since inflation rarely deviates from expectations by more than 1% to 1½% *per y ear*, it's reasonable to suggest that monetary shocks are relatively small, particularly in comparison to real (often local) shocks.

extended periods of time (just as with our product inventory), the firm must assume an actual shift in labor supply and shift prices (both wages and product prices) appropriately.

Follow the process through its steps, and we can see how stickiness and lag accumulates:

1. Demand increase
2. Lag to recognized increased demand
3. Attempt to increase supply
4. Lag to post position and get candidates
5. Lag to recognized low labor supply from candidate counts
6. Increase in supply prices
7. Lag to get new product prices posted

All together, if a company is unable to use macroeconomic factors to make decisions, it will introduce significant real stickiness to its behavior. If each firm faces the same concerns, a shift in macro variables will drive the average firm in the right direction (some faster than others), but this stickiness will not be overcome by any special features of the macroeconomic force or any approach to aggregation.

With the basic framework explained, we will attempt to illustrate the model in greater detail and mathematical rigor.

The Firm

The most extreme version of local information is to limit a firm's knowledge to the actual demand levels and inventory of the firm. This assumption has been selected to explore the nature of this limitation rather than to suggest that it is the sole factor in the economy. In order to explore this element of the process, the firm has been assumed to have certain characteristics:

1. The firm will be the price setting body in both product and labor markets, requiring a level of imperfect competition. To avoid confusion with alternative explanations of stickiness, the firm will maintain a fixed markup.
2. The firm will only be given information about its own inventory values. For product inventory this is strait forward. For labor inventory, this would reflect the number of applicants for a particular job. This model will be agnostic as to whether the target inventory level for labor is required to find an adequately skilled employee or whether it represents an efficiency wage story.
3. In both markets, the firm will face uncertainty. Again, for demand, this is a natural concept. For labor, this model will assume that an interview attracts an uncertain percentage of the available labor, normally distributed around some mean. Consequently, firms will be uncertain about the actual levels of available labor despite obtaining a measurable number of applicants for a particular job. The single firm model will be simplified by assuming that the prototype firm always draws the expected value from that distribution. Once multiple firms are introduced, the qualitative consequences of actual uncertainty will be discussed.
4. The firm can estimate the price elasticity of its labor supply and product demand.

5. The firm experiences no hiring lag. The firm responds to changes in demand by immediately testing the labor market. It then reprices both its wages and products based on labor market signals. This simplifies the problem into a single inventory control and pricing problem (labor). Relaxing this assumption would introduce additional lag into the process, but not alter the basic conclusions.
6. Finally, the firm uses simple decision rules that are represented by rules of thumb taken from basic statistical hypothesis testing.

For the sake of exploring the basic characteristics of the model, we will assume:

- Real economic conditions will be held constant. This means that all demand shifts are monetary and the economy is at full (NAIRU) employment so any additional hiring will force companies to increase their wages.
- The standard deviation (σ) in both markets will be held constant and known to simplify statistical testing and eliminate the need to use more complicated t-tests.

Pricing Policies

As a starting point, the firm with be endowed with belief about the level of demand $E(D_t)$ that it will experience as a result of its current pricing policy, P_t. In each subsequent period, a firm's sales come in as D_t. Aware that it faces uncertain demand, a firm must determine the best way to update its price based on this limited information.

If we hold standard deviation constant, the question during a given period is how reality deviated (or did not deviate) from its expectations. The most direct way to answer this question is to execute a hypothesis test that $D_t = E(D_t)$. Fortunately, holding standard deviation (σ) constant, a single period test can be approximated by a simple test of the z-score,[4] $(D-D^*) / \sigma$

The typical approach to a means test is to select an alpha (α) as a confidence threshold. While this is a historically popular approach, it means that a business would not change its pricing when it's 89% confident but would when its 90% confident. Rather than deal with an arbitrary confidence level, a firm would be better off asking the probability that this demand was drawn from the expected distribution. Since we have assumed that companies draw the expected value of the new demand distribution, we can plot this probability against the z-score of the particular mean shift:

[4] Familiarity with the z-score is used as means of simplifying the model. In reality, this model would require data from multiple periods and a two-sample pooled t-test to address unknown σs.

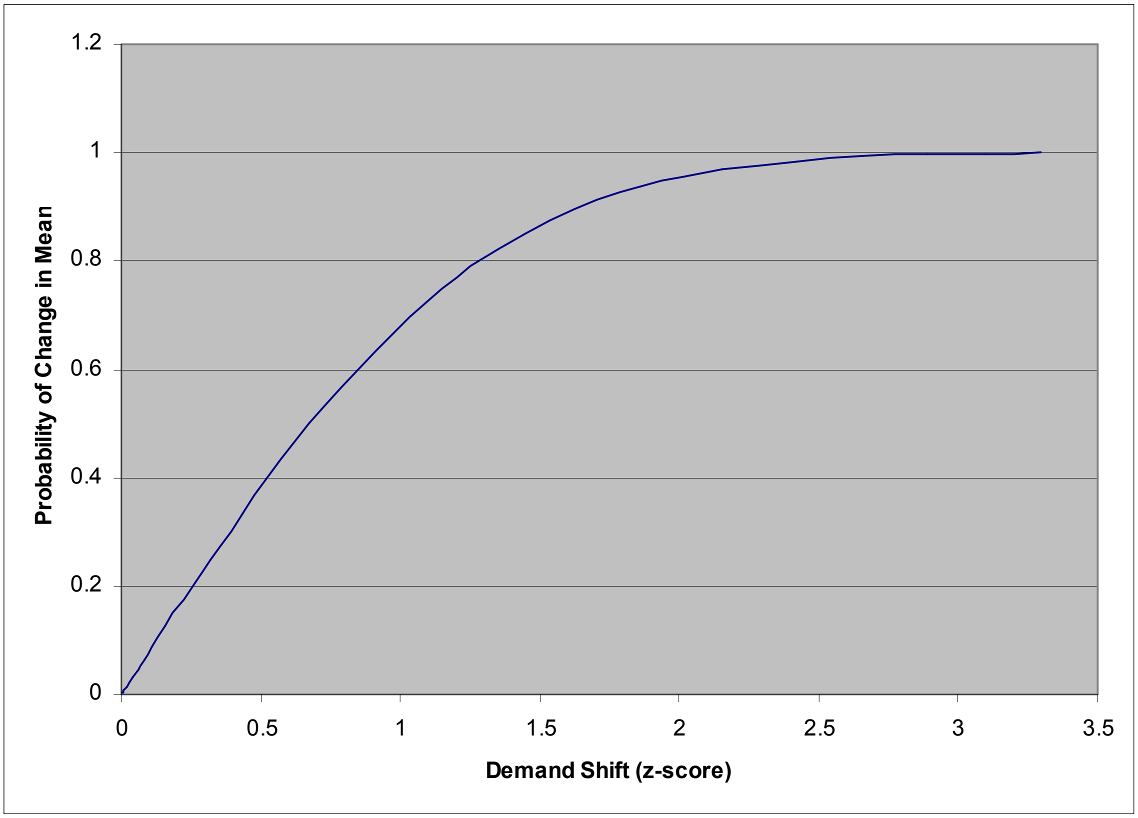

If the z-score indicates a high probability that demand has shifted, the company's best estimate of the new demand level is D_t. Consequently a firm could estimate that it has a probability (1- α) that demand is actually at $E(D_t)$ and a probability (α) that demand has changed to D_t, or:

$$E(D_{t+1}) = E(D_t) * (1 - \alpha) + D * (\alpha)$$

Naturally, this additional demand must be addressed by hiring labor. Consequently, the firm goes out and tests the market for labor. Since we have assumed full employment, attempting to hire additional labor will drive the labor market below equilibrium inventories. Consequently, the firm will have to raise its wages to attract a sufficient number of applicants.[5]

This new price would be based on the elasticity of demand and labor. Assuming the firm faces constant elasticity in both markets and holds its markup constant, the two elasticities can be combined into a single elasticity of magnitude γ. Consequently the company's new pricing policy would be:

$$P_{t+1} = P_t + \gamma * [E(D_{t+1}) - E(D_t)]$$

[5] We remain agnostic as to whether this is required to attract skilled candidates or restore efficiency wages

As this model does not depend on menu costs, it functions fully with very small price adjustments. While the typical posted price is rarely so fluid, Meltzer (1995) gives us the escape required to restore this model as we can, “Assume that firms initially respond to changes in cost and demand by adjusting deliveries, advertising allowances, discounts, and so on while leaving quoted prices unchanged for several months or longer. If managers are uncertain about the duration of changes in demand or cost, they can change other components of the price vector to test the market's response. By changing delivery terms, or offering or removing discounts, firms can change their revenues or the buyers cost without changing the quoted price. This pricing model can be used to rationalize the familiar Keynesian supply curve--a reverse L--when quoted prices are distinguished from other terms in the price vector. Equally, the model can explain the difference in response to the questions in Blinder's [1991] survey.”

Since our initial conditions have real variables set at equilibrium, long run quantity theory would suggest that an increase in demand of a given percentage should be offset by a comparable shift in prices. If we set our elasticity (γ) so that the company would select the long-run equilibrium price at high levels of confidence, we would expect a chart like:

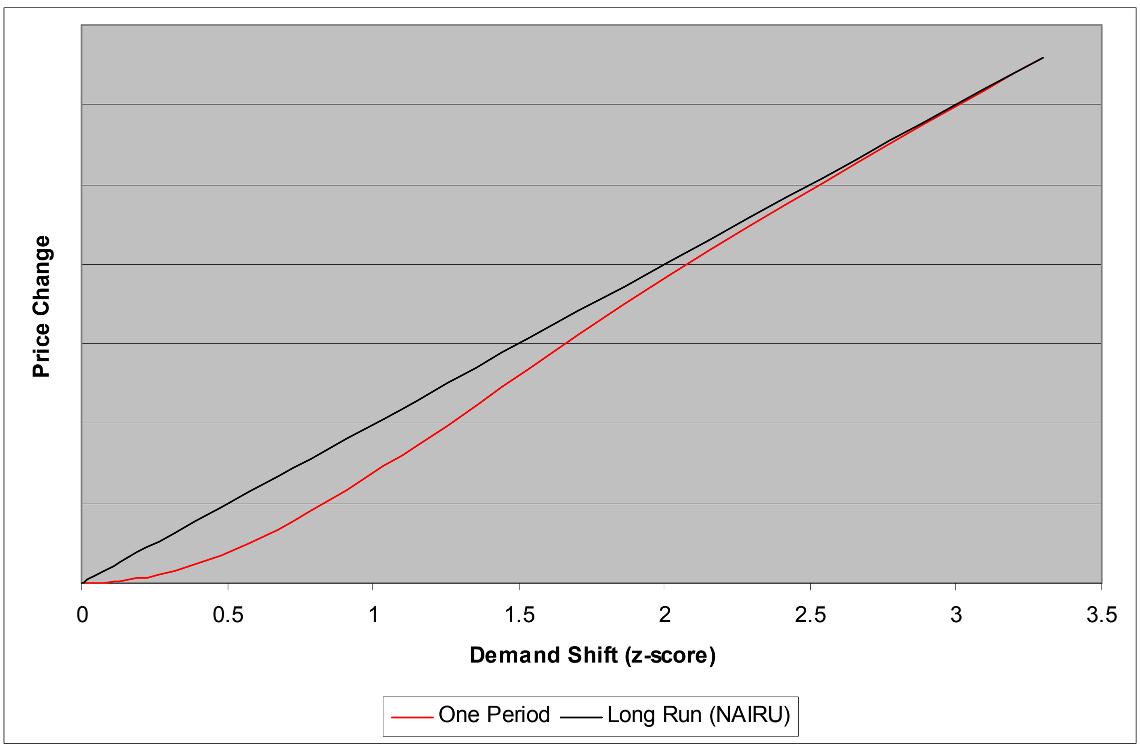

The units of the price change have been suppressed because they should be measured in percentage so that a shift of n% in demand is offset by a shift of n% in price and thus cannot be accurately mapped against a raw z-score.

In a competitive market, elasticity of demand is considerably higher than the aggregate demand budget line as companies that raise their prices transfer additional demand to their competitors. The empirical behavior of the Phillips Curve agrees that the actual elasticity is higher than long-run equilibrium behavior. If we increase the elasticity of the model, we get a chart more like:

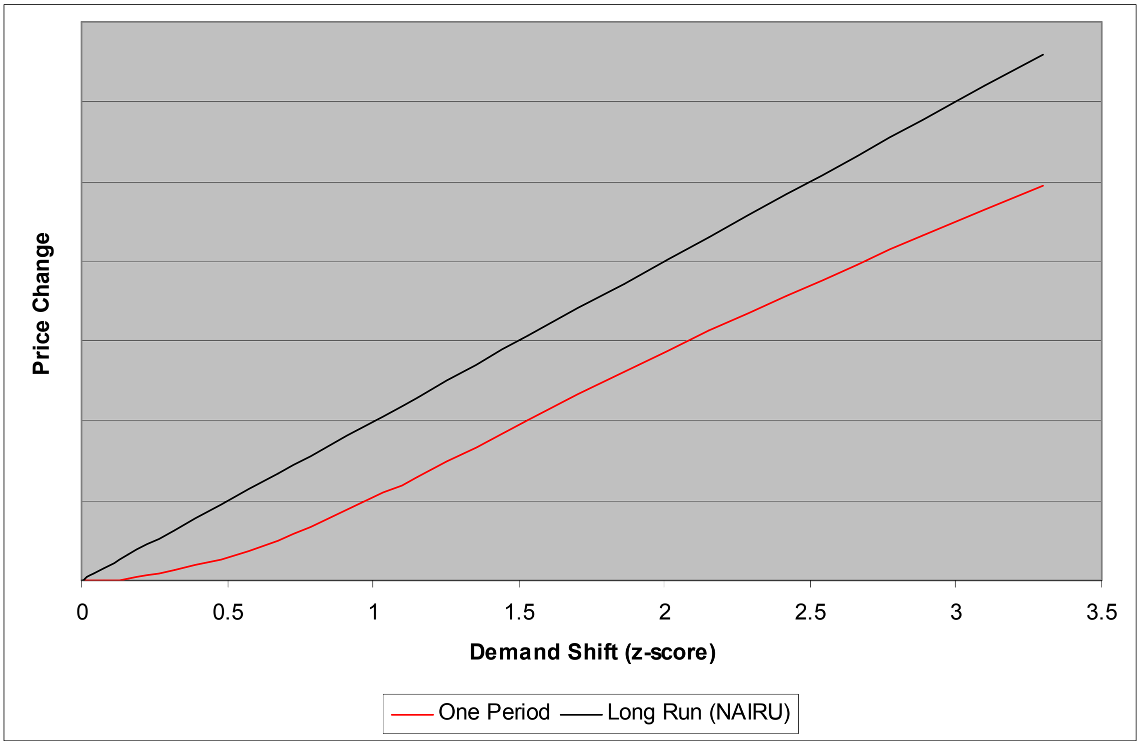

Timing of Price Changes[6]

Naturally, the simplest version of this model delays the price change by a single period, so we plot D_t (the cause) against P_{t+1} (the effect). However, we recognize that changes will occur continuously in this model (as demand shifts are increasingly confident) so we get a second piece of good information by plotting D_t (the effect on demand) against P_{t+1} (the effect on price) if no additional monetary shocks occur:

[6] For simplicity, all graphs except the Phillips curve use the Y axis as NAIRU so positive X values equate to excess demand for labor and negative X values represent an output gap.

As the actual demand for the period should fall between these two curves (as the adjustments occur gradually), the resulting Phillips curve is most compelling if 75% of the final adjustment in price is absorbed into price and 25% of the final adjustment in price is absorbed by quantity. Naturally, these numbers are otherwise arbitrary in a static model, but could be made precise using a dynamic model with one of several reasonable decision rules including a continuous adjustment process. Incidentally, this estimate is also similar to a graph of D_t and P_{t+1} when the elasticity of demand is higher than the elasticity imposed by the monetary (or potentially aggregate demand) budget constraint:

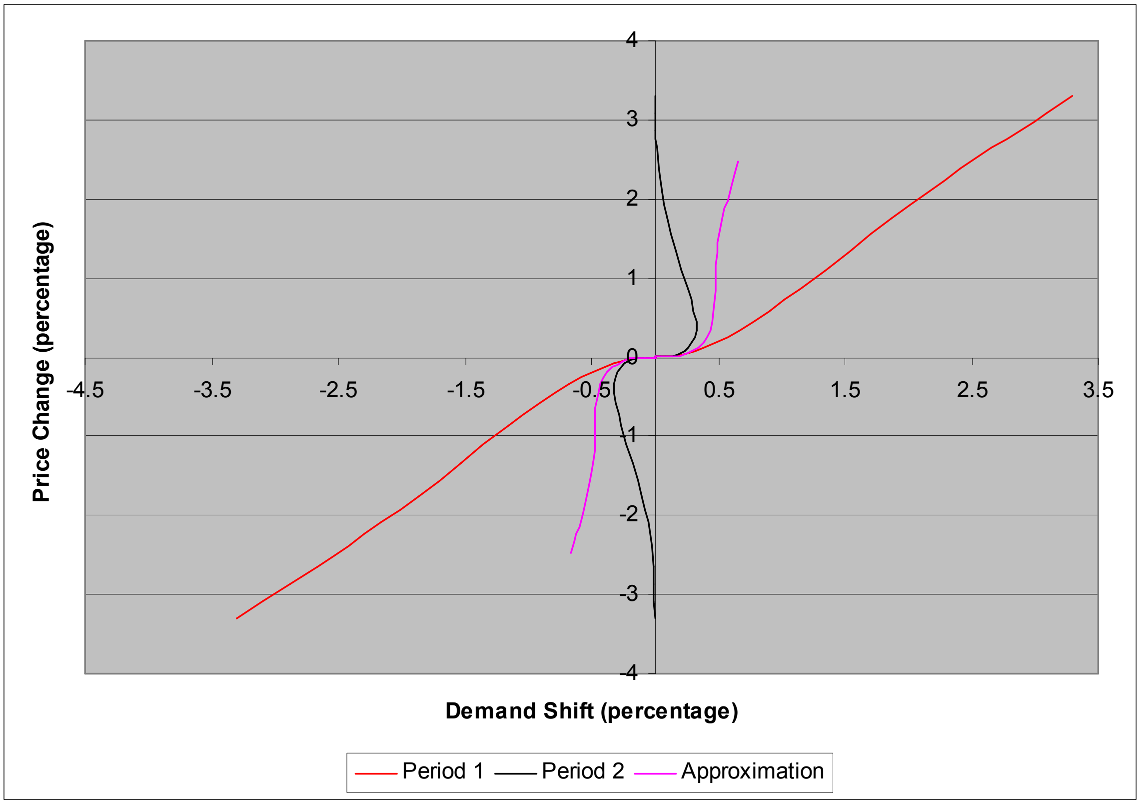

If we lay various budget curves across this approximation, we get:

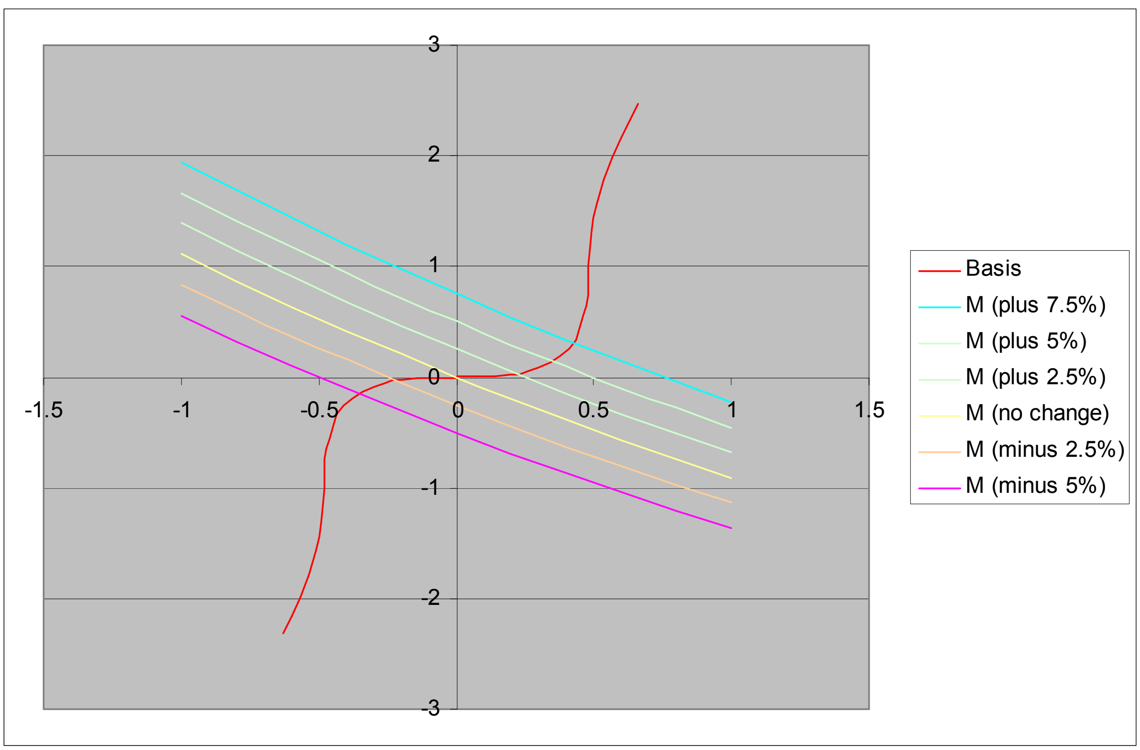

If we calculate the result of various monetary policies, we can derive a Phillips Curve based on ΔP and ΔQ, by defining some arbitrary level of NAIRU unemployment (10% in this image):

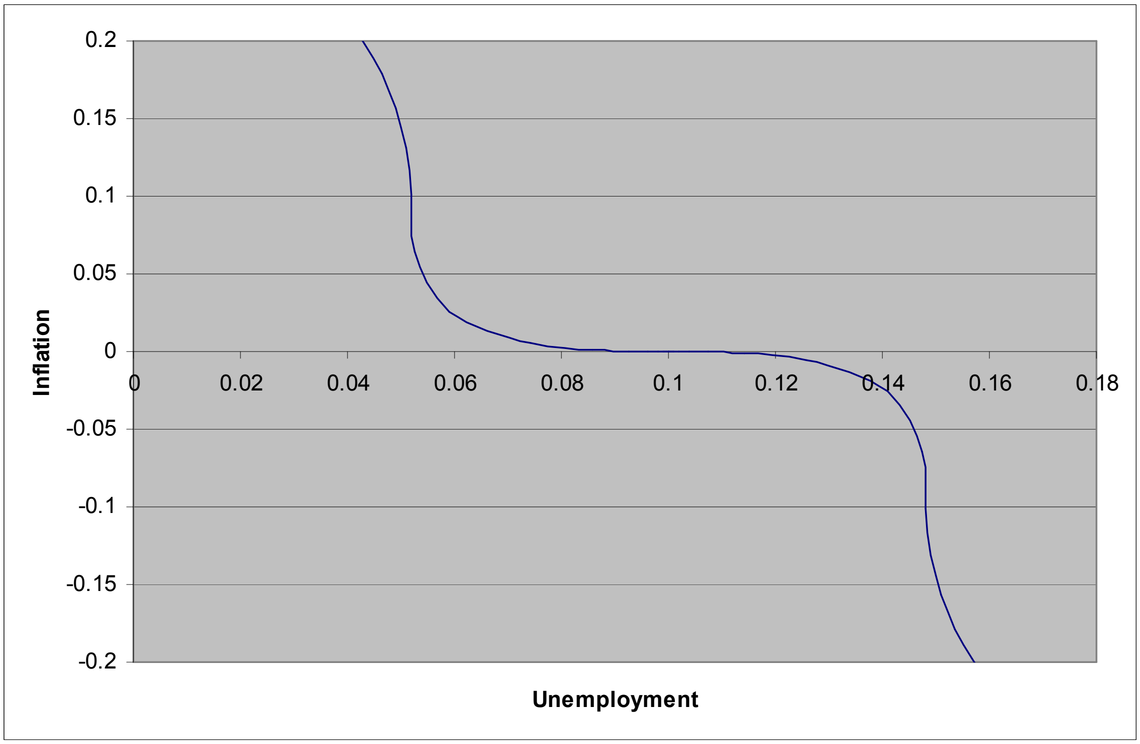

Naturally, with as many simplifications as have already been introduced, the performance of this model is encouraging. Naturally adjustments to the actual performance of the "approximation" would vary this curve only in the details and not in the basic trend. Also, as time passes, the confidence improves and the curve shifts towards NAIRU:

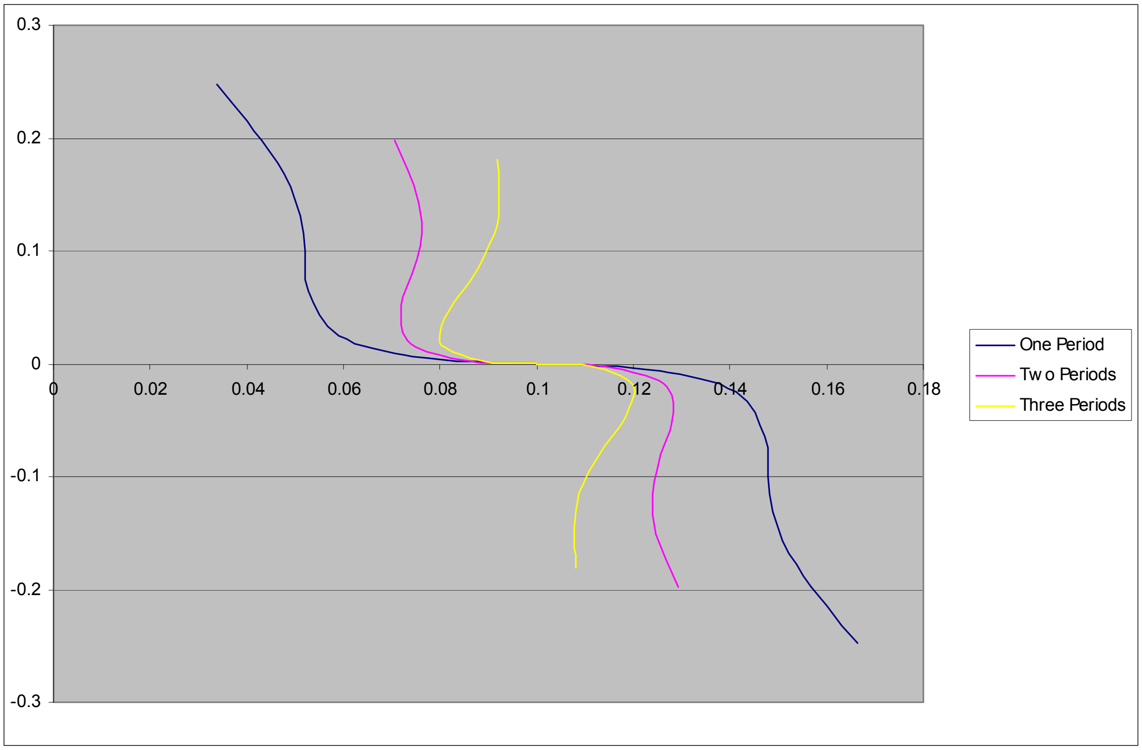

Notice the extremely high rigidity around the NAIRU equilibrium point despite the passage of time that would have, in other models, allowed for firms to accumulated information or change prices. While a dynamic model would eliminate some persistence through extreme draws (from demand uncertainty) and through higher elasticity around average prices (driving laggards more quickly towards equilibrium), this model offers an extremely credible and persistent rigidity.

Expectations

If we introduce expectations regarding the future price level, this would be reflected by an upward shift the curve on the P-Q plot, making the company sticky around its new expectations (until the evidence accumulated):

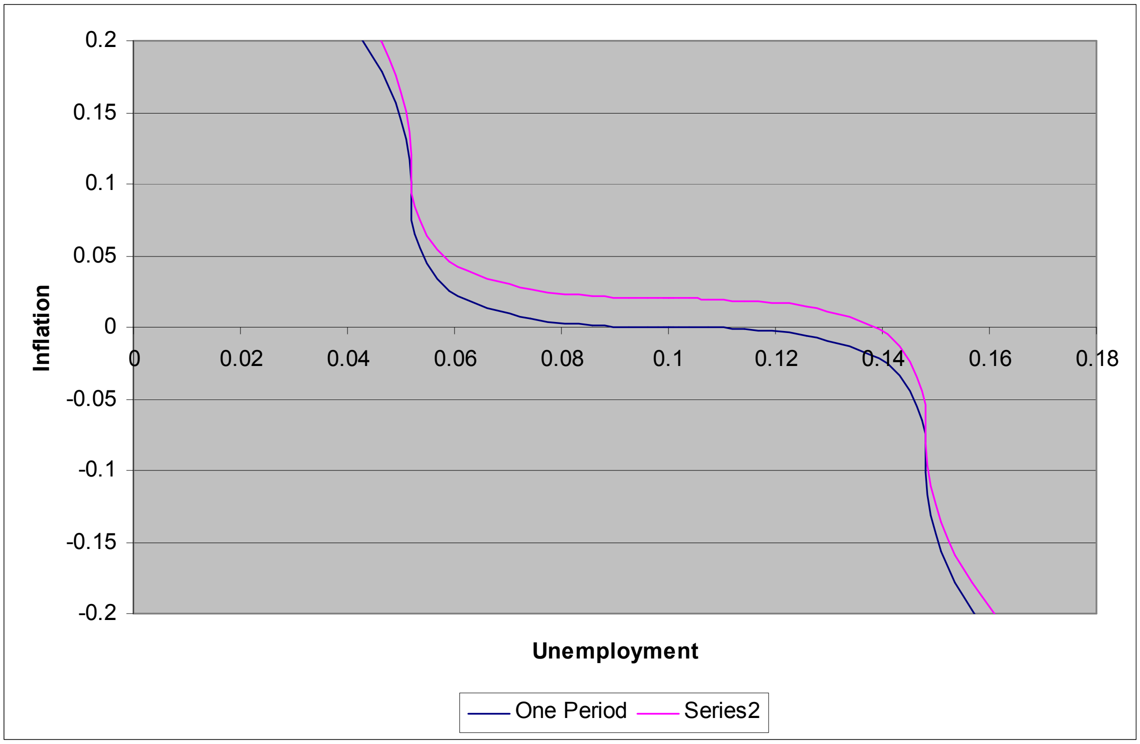

Naturally this is captured by shifting up E(D):

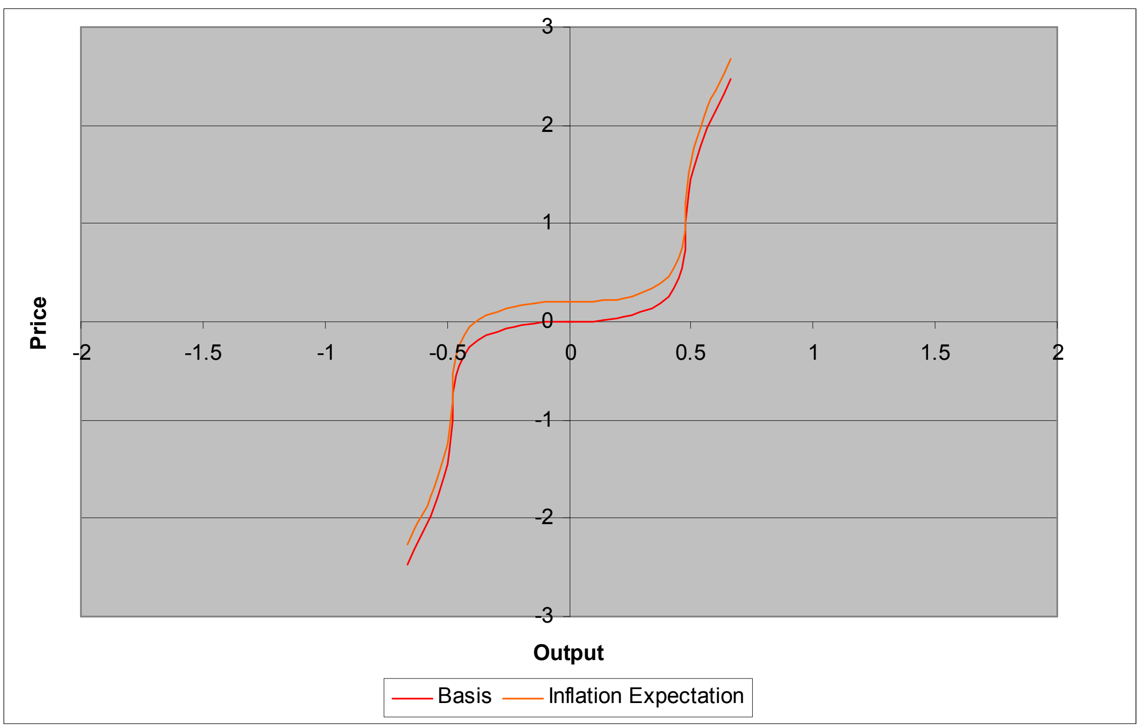

And each output gap occurs at different monetary policies, where equilibrium is represented by a shift in monetary policy equal to the shift in expectations:

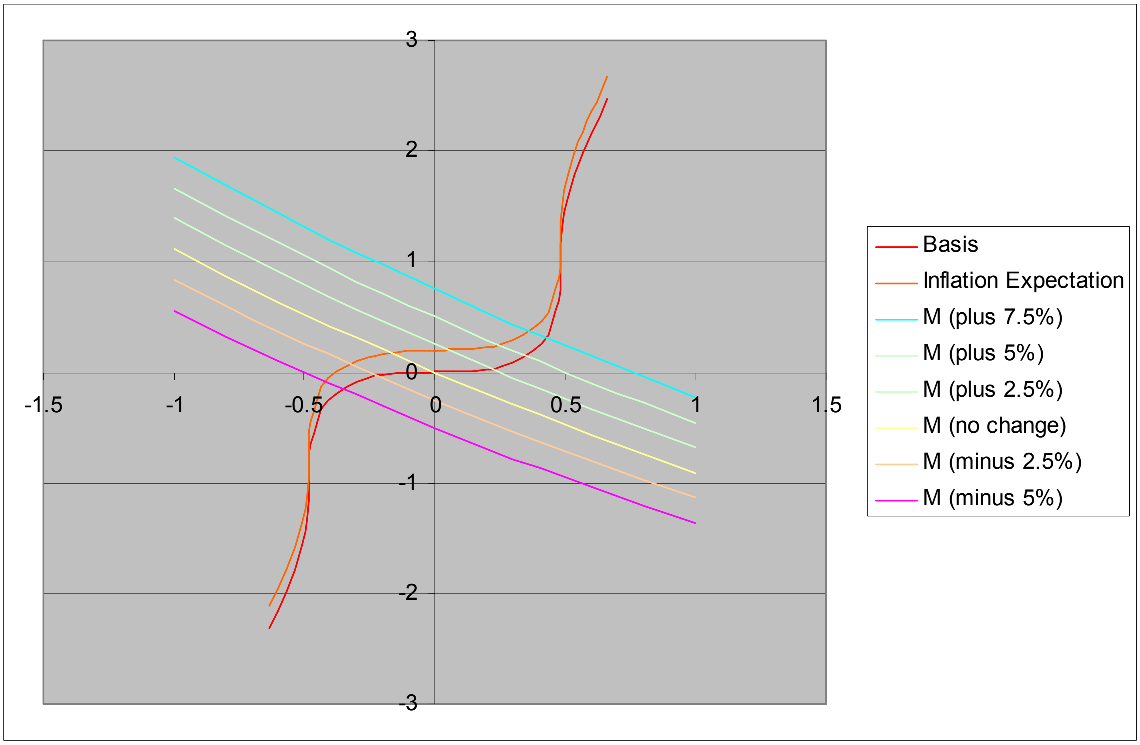

Dynamic Considerations

The single firm behavior is striking, but we need to go further and address the implications of multiple firms. Typically, a company that experience demand towards one of the extremes of the graph will make price adjustments to address the gap. We would expect a number of small shifts, and many in opposite directions. Indeed, this expectation is consistent with the findings of Carlton (1987) that, "Even for what appear to be homogeneous commodities, the correlation of price changes across buyers is very low."

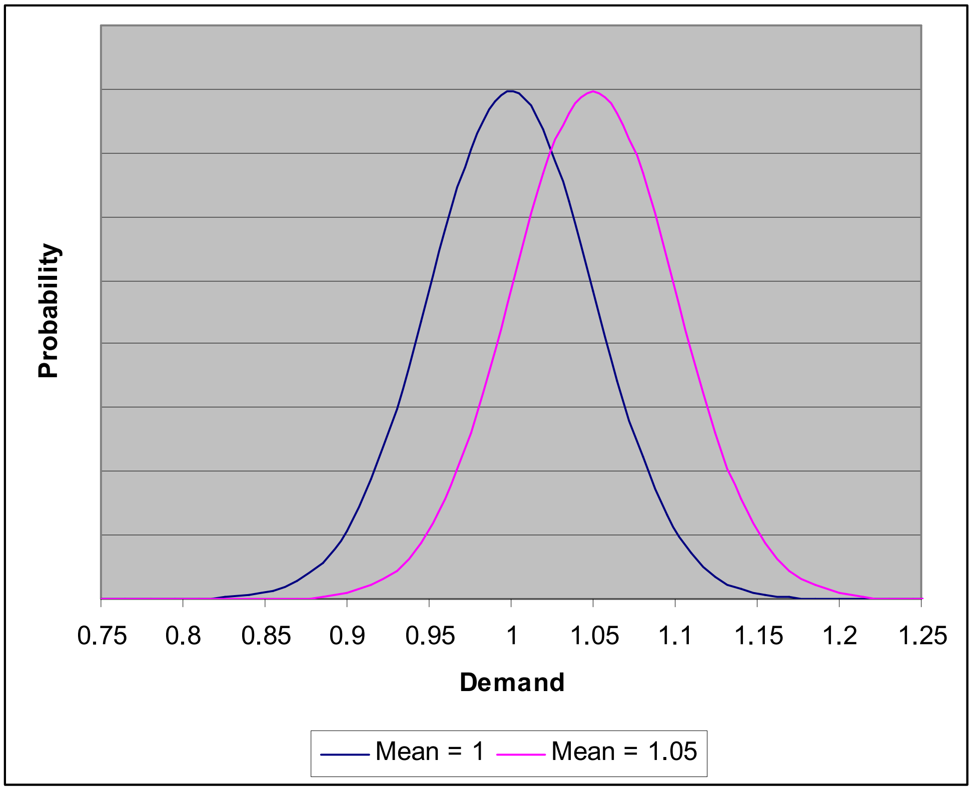

As we can see, with a shift of one full standard deviation, 16% of the new demand distribution still falls below the previous mean, potentially signaling the need for lower prices. Of course, each subsequent shift in an individual company's policies will transfer a disproportionate amount of that impulse on to other companies that have not yet detected the underlying shift (or gotten the opposite indicators). [7]

Consequently, the price adjustment impulse will be progressively passed from firm to firm. During the majority of periods, the average firm will continue to get impulses to shift prices in the correct direction. However, in the intervening periods, their price adjustments will drive them past the average and they will receive a signal to cease (or sometimes even retrench) price shifts.

Once the new equilibrium has been reached, companies would continue to fluctuate around this equilibrium, depending both on irrelevant noise as well as relevant local shocks such as labor shortages near production facilities.

Comparables

[7] It is worthwhile to point out that we would expect a negative autocorrelation within any particular demand series as a population who, by sheer chance, all shows up to purchase goods in a particular period, are unlikely to return in the subsequent period to continue to produce high, albeit misleading, levels of demand. Indeed, we might expect a company in this circumstance to suddenly recognize a drastic fall in demand and respond with measures even more extreme than their competitors who received signals better reflecting real equilibrium.

Now that we have a model of real price stickiness, it is valuable to consider existing stickiness theories and recognize why these theories function adequately.

Menu Cost

The first and easiest to address is the menu cost approach. Gordon (1990) offers, "The new menu-cost literature owes its origins to a paper by Barro (1972) on the S,s approach to price adjustment by a profit-maximizing monopolist who faces a lump-sum cost of adjusting prices… The basic S,s result is derived for a monopolist facing a stochastic additive shift in its demand curve taking the form of a random walk without drift. The optimal strategy for the monopolist is shown to be the selection of "floor" and "ceiling" bands, with the price remaining constant when the shift is within the bands and changing fully to the new desired level when the shift is outside the bands."

If we applied the hypothesis test in its standard threshold form, this is precisely the behavior we might expect. We would suppose therefore that, rather than capture an actual menu cost, this theory captures a threshold version of the inventory theory. Since individual businesses often work off threshold inventory control models, we are unsurprised that this thesis is consistent even in more complex dynamic models. However, pure inventory theory would expect a hard threshold to break down during aggregation and the price adjustment supposed by menu cost theory would be slightly different than an inventory model's confidence weighted shifts.

Staggered Contracts

Again, the literature on the topic offers a behavior that arises naturally in the inventory model. When we consider the overlapping normal curves produced by mapping uncertainty around $E(D_t)$ with the actual uncertain demand experienced around D_t, we realize that many companies will realize little or no shift (particularly when confidence weighted) and will delay price changes while others will update their prices immediately. On average, the shift would be in the appropriate direction, as staggered contracts would suggest. However, major prices changes are now staggered on the basis of real expectations rather than arbitrary periodic changes.

Gordon's (1990) offers on this model, "Ball and Romer (1989) show that staggering is a stable equilibrium if there are firm specific shocks that arrive at different times fir different firms. However, they show that synchronization can also be an equilibrium: 'Multiple equilibria are possible because there is an incentive for synchronized price setters to remain bunched, but not for staggered price setters to move toward synchronization' (1989, p. 193)." Indeed, an inventory model would define the firm-specific uncertainty of demand as the critical firm specific-shock that justifies the unstable equilibrium.

Adaptive Learning

We can also attribute adaptive learning models to pricing rules similar to our expected demand rule:

$$E(D_{t+1}) = E(D_t) * (1 - \alpha) + D * (\alpha)$$

combined with the price rule:

$$P_{t+1} = P_t + \gamma * [E(D_{t+1}) - E(D_t)]$$

Each period, the past beliefs are updated by a new factor including an alpha. However, unlike adaptive learning models, our alpha will vary based on the relationship between the demand shift and the standard deviation of demand. Consequently, adaptive learning models capture the general trend of the pricing rule, but without adjusting for different shift-confidence.

Sticky Information

Sticky information, one of the most recent macroeconomic developments attributed primarily to Sims (1998, 2003), Woodford (2003), and Mankiw and Reis (2002, 2006), is also the closest to our basic thesis. Unlike typical stick information models that depend on slow or costly learning of aggregate shifts, our model depends on the slow development of *confidence* in the local relevance of aggregate shifts. Naturally, we would expect average prices to adjust to macroeconomic shifts even though the information must be passed through its local effects.

These models are relatively new, but have gained some credence from see Klenow & Willis, 2006 and have detractors such as Coibion (2006).

Input Output Table

Finally, an inventory model that takes raw materials into consideration begins to take on the character of the input-output table argument. Gordon (1990) sums up situation as, "To be a credible explanation of real price rigidity, the distinction among local and aggregate cost and demand shocks must be embedded in a world with many heterogeneous firms interacting within a complex input-output table. With only two firms, each supplying the other, information would be cheap enough to permit both firms to disentangle the local versus aggregate component of their costs. But with thousands of firms buying thousands of components, containing ingredients from many other firms, the typical firm has no idea of the identity of its full set of suppliers when all the indirect links within the input-output table are considered." Again, a hard and fast justification for a local model and a natural aggregator.

Work Cited

Ball, Laurence and Romer, David, (1989) "The Equilibrium and Optimal Timing of Price Changes," Rev. Econ. Stud., 56(2), pp. 179-98.

Barro, Robert J. (1972) "A Theory of Monopolistic Price Adjustment," Rev. Econ. Stud., Jan. 1972, 39(1), pp: 17-26.

Barro, Robert J. (1976) "Rational Expectations and the Role of Monetary Policy," . Monet. Econ., Jan. 1976, 2(1), pp.1-32.

Blinder, Alan S. (1991) "Why Are Prices Sticky? Preliminary Results from an Interview Study," The American Economic Review (May 1991), pp. 89-96.

Bomhoff, Edward J. (1983) Monetary Uncertainty. North Holland.

Brunner, Karl, Alex Cukierman and Allan H. Meltzer (1983) "Money and Economic Activity: Inventories and Business Cycles," Journal of Monetary Economics (May 1983), pp. 281-319.

Carlton, Dennis W. (1986) "The Rigidity of Prices," American Economic Review,Vol . 76, No. 4, (September 1986), pp. 637-658.

Coibion, Olivier (2006) "Testing the Sticky Information Phillips Curve." Job market working paper.

Gordon, Rober J. (1981) "Output Fluctuations and Gradual Price Adjustment," J. Econ. Lit, June 1981, 19(2), pp. 493-530

Gordon, Robert J. (1991) "What Is New-Keynesian Economics?," NBER Reprints 1513, National Bureau of Economic Research, Inc.

Klenow, Peter J. and Willis, Jonathan L. (2006) "Sticky information and sticky prices," Research Working Paper RWP 06-13, Federal Reserve Bank of Kansas City.

Mankiw, N. Gregory, and Reis, Ricardo (2002) Sticky information versus sticky prices: A proposal to replace the new Keynesian Phillips curve. Quarterly Journal of Economics 117, 1295-1328.

Mankiw, Gregory N. and Ricardo Reis (2006) “Pervasive Stickiness,” American Economic Review, May 2006, 96 (2), 164–69.

Meltzer, Allan H. (1982) "Rational Expectations, Risk, Uncertainty and Market Responses," in Paul Wachtel, ed., Crises in the Economic and Financial Structure. Lexington Books, 1982, pp. 3-22.

Meltzer, Allan H. (1995) "Information, sticky prices and macroeconomic foundations," Proceedings, Federal Reserve Bank of St. Louis, issue May, pages 101-118.

Muth, John F. (1961) "Rational Expectations and the Theory of Price Movements," Econometrica (July 1961), pp. 315-35.

Sims, Christopher (2003) “Implications of Rational Inattention,” Journal of Monetary Economics, April 2003, 50 (3), 665-90.

Woodford, Michael. (2003) “Imperfect Common Knowledge and the Effects of Monetary Policy” in Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, Princeton University Press.