Seasonality in a Menu Cost Model

Aaron Popp
3/14/2010
The Ohio State University
Preliminary

Abstract

A baseline menu cost model cannot generate substantial monetary nonneutrality without introducing sources of real rigidity, and there are many potential, effective sources of real rigidity. I construct a menu cost model that features seasonal fluctuations in a labor market rigidity, which potentially generates more monetary nonneutrality than the equivalent model without seasonal fluctuations, a large seasonal cycle on the order of 7.5% from peak to trough, and a seasonally varying response of the economy to a monetary shock. Seasonal fluctuations in the economy provide firms a moving target for their optimal prices for which the firms do not fully adjust because of the menu cost and other sources of real rigidity. Real rigidities force firms to cluster their prices closer together, and the seasonally moving target for firms increases the penalty for firms that have prices that differ from the aggregate price level, which amplifies the clustering motive. The effect of seasonality also depends on the persistence of the idiosyncratic fluctuations. As the persistence of idiosyncratic factors increases and becomes more important for the firms, the clustering effect of seasonality weakens.

Email address: popp.22@osu.edu. I thank Paul Evans, Aubhik Khan, Julia Thomas, Joe Kaboski, Fang Zhang, Jing Han, and OSU Macro Lunch participants for comments and suggestions.

1 Introduction

Economics advances when fundamental assumptions are relaxed. Models with imperfect information relax the assumption of rational expectations. Time dependent and state dependent prices relax the assumption of perfectly flexible prices. As economists relax fundamental assumptions, economists either explain phenomena that the earlier lineage of models cannot or find that relaxing restrictive assumptions makes expected, desirable results disappear. In

[1]Email address: popp.22@osu.edu. I thank Paul Evans, Aubhik Khan, Julia Thomas, Joe Kaboski, Fang Zhang, Jing Han, and OSU Macro Lunch participants for comments and suggestions.

this paper, I show that an otherwise standard menu cost model with a seasonally fluctuating labor market friction should generate more monetary nonneutrality in response to a monetary shock than the same menu cost model without seasonal fluctuations in the labor market friction. Relaxing the common assumption that seasonal fluctuations do not matter may help state dependent pricing models generate a large amount of monetary nonneutrality.

A menu cost model must include real rigidities or strategic complementarity in price setting (Ball and Romer, 1990). Equivalently, firms must strongly want to have an optimal price that is near the price of the majority of other firms in the economy, even with differing idiosyncratic effects. If the real rigidity is strong enough, few firms will deviate far from the aggregate price level. When a monetary shock occurs, the firms that would want to adjust are the ones with either a large idiosyncratic shock or the ones that are the furthest from the center of the distribution of relative prices. If fewer firms are on the tails of the distribution of relative prices, fewer firms will change their prices in response to the shock. While the price changes of the firms that adjust will be large, most firms in the economy will still be near the aggregate price level before the shock with a strong rigidity. The aggregate price level responds more sluggishly to the shock, and thus output responds more strongly to the shock. Seasonality gives firms an stronger incentive to set their prices near the prices of other firms. The firms' optimal price changes with the seasons, but they will not be able to completely adjust to the changing seasons. Hence, rather than being near an optimal price when near the center of the distribution of firms across states, as in a model without seasonality, firms are frequently operating off their target price. If they respond more to exogenous idiosyncratic or aggregate motives, then they will suffer greater profit losses. Price changes are generally smaller in the model with seasonality than the model without seasonality, and firms decide to keep their prices between the extremes of the seasonal fluctuation in their optimal, costless adjustment price. Seasonal fluctuations in my model not only generate a seasonal cycle, they are important for the business cycle as well.

My model shows that a relatively small fluctuation in the real rigidity of the economy can generate a sizeable seasonal cycle in the economy. Consumption fluctuates by around 7.5% from peak to trough in my baseline specification, which is approximately the size of the seasonal cycle found by Barsky and Miron (1989). There are several competing explanations for the existence of the seasonal cycle. Olivei and Tenreyro (2007) point to time varying labor market fluctuations, such as staggered labor contracts, as an explanation for the seasonally varying response of the economy to monetary shocks that they find. Nakamura and Steinsson (2008) suggest that seasonal fluctuations in demand could be responsible for the seasonal fluctuations in the micro price data, but neither study conclusively disproves the other. My model is more similar to Olivei and Tenreyro's concept; the source of seasonality in my model is a generic, time-varying labor market friction. Seasonal demand shocks in the model could still play an important role in fitting the micro price data about price changes over the seasons, but it does not appear to be exclusively vital to obtain sizeable seasonal fluctuations in the economy.

The response of my model to monetary shocks should be seasonally dependent as well. Olivei and Tenreyo found that output responds much more to monetary shocks at the beginning of the year than at the end of the year, and my findings are qualitatively consistent with their result. Midrigan (2008) shows that one way to boost the monetary nonneutrality of a menu cost model is to increase the kurtosis of the distribution of price changes, conditional on the fact that they actually adjust. Midrigan chooses to use leptokurtotic cost shocks, which forces firms to cluster more around the median firm's price and produces mostly small price changes. The kurtosis (raw, not excess) in my model fluctuates by season from a low of 1.43 to 3.15 from about 1.4 in the model without seasonality. While the overall kurtosis of the model is still too low to account for the amount of monetary nonneutrality that a similarly calibrated time-dependent model is able to generate, it is possible to add other sources of real rigidity to my simple model to boost the monetary nonneutrality produced in my model to approximately the level produced by a time-dependent model.

A separate, methodological contribution of my paper is how I handle the solution of the menu cost model. My model is the first DSGE model with menu costs, heterogeneous agents, and seasonal fluctuations. Others have solved models featuring a seasonally fluctuating economy (see Braun and Evans (1995), Liu (2000), and Olivei and Tenreyro (2007)), but those models do not feature extensive firm heterogeneity. To solve for the steady state of a menu cost model, one must solve for the stationary distribution of the heterogeneous firms across states. Ignoring seasonal fluctuations, I solve for the stationary distribution using a method similar to solving for the stationary distribution of a Markov chain. I then generalize the method to accommodate the differing pricing rules and distributions of firms across states across the seasons. This cycling distribution of firms across time is termed a cyclostationary distribution, a concept that has been rarely applied in macroeconomics, and I call my seasonal equilibrium the cyclostationary equilibrium. The method could be useful to other similar models. For example, recent models of international trade emphasize that trade is lumpy at the firm level is lumpy, and trade for certain goods is seasonal. If demand for certain goods is seasonal, then the lumpiness could be partially due to seasonality. Kaneda and Mehrez (1998) argue that seasonal fluctuations are nontrivial when modelling international trade and that there are large seasonal fluctuations in the trade of some disaggregated goods and aggregate measures of trade. If firms and traders are heterogenous, one could apply my method to help find the steady state and cyclostationary distribution of firms in a seasonally dependent model of international trade.

Section 1 continues with a review of the related literature. I introduce the model in Section 2. Section 3, Appendix A, and Appendix B explain the algorithm that I use to solve the model. Section 4 shows the main results of the model. Section 5 concludes.

1.1 Literature Review

1.1.1 Other Models with Seasonality

Seasonal fluctuations in the macroeconomy are larger than business cycle fluctuations in the short run. Barsky and Miron (1989) develop some basic facts about the seasonal cycle of the economy, which allow them to establish the existence of the seasonal cycle in the economy. Seasonal fluctuations dominate in the short run in the economy. Regressing on detrended data, Barsky and Miron find that output on average fluctuates by about 8% over the course of the year from the peak in the 4th quarter of the year to the trough in the first half of the year. The price level moves in the opposite direction as output, but the fluctuations in the overall price level are much smaller, on the order of less than half of a percent. Also, the seasonal cycle and the business cycle share qualitative characteristics, which may imply that the fluctuations have some common causes. Their results are empirical, and economic models that incorporate seasonality are left for future research.

Braun and Evans (1995) construct a DSGE real business cycle model with seasonal fluctuations in preferences, technology, and government purchases to attempt to explain the seasonal cycle as laid out by Barsky and Miron. To explain the size of the seasonal cycle, their model requires technology to rise at a 24% annual rate in the fourth quarter of the year and to fall at a 28% annual rate in the first quarter of the year, which they consider evidence that the typical production technology in RBC models is misspecified.

While Braun and Evans (1995) create a model that can explain seasonal fluctuations in real variables, it could not explain fluctuations in nominal variables. Liu (2000) creates a two sector monetary model with seasonal fluctuations to determine whether and how the Federal Reserve should respond to the fluctuations in the seasonal cycle versus the business cycle. Historically, the Federal Reserve smoothed nominal interest rates over the seasonal cycle but not the business cycle. The real-bills doctrine suggests that a central bank should smooth fluctuations in nominal interest rates over both cycles while the quantity theory of money suggests that a central bank should smooth fluctuations in the money supply over the cycles. Liu finds, comparing the historical policy, the real-bills policy, and the quantity-theory of money policy, that the historical policy is closest to the optimal policy prescribed by Friedman's Rule. Notably, his model is the first to combine both business cycle fluctuations and seasonal fluctuations in the model economy.

Olivei and Tenreyro (2007) also combine business cycle and seasonal fluctuations in a DSGE model to explain apparent seasonal fluctuations in the response of the economy to monetary shocks. Using a seasonal VAR, they show that the response of output to monetary policy shocks varies depending on the quarter in which the shock occurs. In the first half of the year, prices respond little initially and output responds quickly to a monetary shock. In the second half of the year, prices respond quickly to the shock, and output does not respond much at all. Olivei and Tenreyro suggest that seasonally varying labor market frictions are the source of the seasonal fluctuations in the response of the economy to monetary shocks. They find evidence that wage contracts are typically signed in the later half of the year, especially during the fourth quarter. If a monetary shock occurs in the latter half of the year, then the firms can price the shock into their wages. Wages respond quickly to the shock, and thus most firms' marginal costs respond quickly to the shock. Prices respond quickly to the shock, which limits the real effects of the shock. In the beginning the year, many firms will have just renegotiated their wages, so marginal costs are roughly constant for those firms. If the monetary shock occurs in the beginning of the year, firms are less willing to respond to the shock since their marginal costs respond sluggishly. They are more willing to change their production than prices, so output responds quickly and robustly to the shock. They create a model of staggered wage contracts in which firms face Calvo timing in wage setting. Every period, firms have a seasonally dependent probability of being able to negotiate their wages. They set the probabilities of adjustment so that firms are much more likely to renegotiate their wages at the end of the year. Their model fits the general pattern that they observe in the VAR analysis.

While seasonality in economics is the subject of a large literature, there are few empirical applications of cyclostationarity in economics, and what applications there are tend to be for time series analysis of macroeconomic or financial data. Broszkiewicz-Suwaj et al (2004) apply cyclostationarity to financial data to find correlation in financial data. Franses (1996, 2004) and Pargano and Parzan (1979) look at the concept of cyclostationarity in time series data and various time series models. Leskow (2001) applies an econometric test featuring cyclostationarity to asset volatility data as an alternative to the ARCH and GARCH approaches to modelling the variance of asset returns. Serpedin et al (2005) is an extensive bibliography on cyclostationarity across disciplines and a good reference for other uses of cyclostationarity.

1.1.2 Menu Cost Models

State-dependent pricing models allow economists to evaluate the assumptions underlying time-dependent models, and they allow economists to better model firm behavior based on the ability to fit the micro data on prices. Menu cost models are a subset of state-dependent pricing models in which firms must pay a fixed cost to change their prices. Some economists do not think that menu cost models can generate sufficient monetary nonneutrality in response to a monetary shock to be useful. These nonneutrality results, including Caplin and Spulber (1987) and Golosov and Lucas (2006), challenge the view that small nominal rigidities could be realistically amplified to explain large monetary nonneutrality in the macroeconomy. If money is neutral in a menu cost model, then either something fundamental about the model is misspecified or the underlying assumptions about firms' price changes in time-dependent pricing models are incorrect. If the latter is true, then the result challenges the validity of a large literature of work that builds upon models that assume time-dependent pricing. On the other hand, others find that there are mechanisms by which small menu costs have substantial real effects, and these models can also fit what economists know about firms' pricing motives. It is possible to develop these models further as a potential replacement for time-dependent models.

Mehrez (1998) explores the implications of the menu cost mechanism in an economy with seasonal fluctuations and hence is the paper most similar in spirit to my paper. Mehrez explores the adjustment incentives of firms, without the now typical DSGE framework, by assuming that firms desire to maximize the flow of profits with a simple penalty from deviating much from a seasonally varying target price. His framework is a sufficient framework to analyze firms' adjustment incentives, but it is not sufficient to pin down the aggregate implications of seasonality. He finds that the observed seasonal fluctuations in the economy are less than the true seasonal fluctuations, because firms cannot adjust completely to the seasonal fluctuations. Also, as inflation rises in his economy, the amplitude of seasonal fluctuations rises as well. Both theoretical results will occur in my model as well. Matching Israeli price data from the 1980s, he finds that the adjustment behavior of the firms is roughly consistent with his model of menu costs coupled with seasonal fluctuations in the firms' desired price.

Caplin and Spulber (1987) create a menu cost model in which money is completely neutral. While firms face a nontrivial menu cost, firms are distributed uniformly across prices by assumption. To generate monetary nonneutrality, the distribution of firms across relative prices must change in response to a shock. If the distribution of firms across relative prices does not change, then the only reaction of firms to the shock will be for the firms with the lowest relative prices to have large nominal price changes. The distribution of firms across nominal prices would simply shift upwards by the amount of the monetary shock, and aggregate prices would fully accommodate the shock. Individual firms will not fully and immediately adjust to the shock, but the firms that do adjust will adjust enough to force monetary neutrality.

Dotsey, King, and Wolman (1999) create a menu cost model in which firms are subject to a random menu cost, which creates heterogeneity in the firm's price adjustment decision. The technical problems of the model increase with this assumption since the steady state distribution of the firms across prices is unknown initially. The problem is tractable, however, and they find that the model produces significant monetary nonneutrality. Their nonstandard distribution of the menu cost shocks drives their results.

Golosov and Lucas (2006) create a continuous time menu cost model that builds upon the previous menu cost literature by matching the price change moments implied by the model to micro data. Specifically, Golosov and Lucas match the mean and variance of nonsale price changes and inflation from Klenow and Kryvstov (2008). Previous models did not match the micro data about prices explicitly. In Dotsey, King, and Wolman (1999), all firms adjust their prices to a common price since the random component of the model is the menu cost, which does not affect the optimal chosen price of the firm. In Golosov and Lucas, firms face an idiosyncratic productivity shock that affects whether and how the firms adjust their prices, which allows them to match the micro data on prices. With a model that matches the micro facts about price changes, they then evaluate how a monetary shock affects the path of output and prices in the model and compare the model to a baseline Calvo model. They find that the model produces little monetary nonneutrality as a result of the monetary shock, consistent with the findings of Caplin and Spulber, which is not surprising since the only difference between Golosov and Lucas' problem and Caplin and Spulber's problem, to a log linear approximation, is the existence of idiosyncratic shocks.

Midrigan (2009) views the Golosov and Lucas and Caplin and Spulber nonneutrality result as a problem of the underlying assumption that the idiosyncratic shocks are normally distributed, the effect of which he calls the selection effect. The selection effect occurs when the firms that change their prices after a monetary shock are the firms that have large price changes. One way to minimize the selection effect is to minimize the number of firms that have large price changes by forcing as many firms as possible to have small price changes. Since large idiosyncratic shocks give firms an incentive to change their prices by a large amount, one way to mitigate the selection effect is to choose an idiosyncratic distribution for the shock that groups most firms around the mean shock value. A leptokurtotic distribution of firms across prices, one in which the tails of the distribution are fat but a large mass of firms is clustered near the mean of the distribution, gives the desired grouping of firms around the median price. Midrigan successfully applies a leptokurtotic distribution of cost shocks to create a menu cost model that can generate substantial real rigidity in response to a monetary shock. The consequence of the leptokurtotic distribution of shocks is a leptokurtotic distribution of desired and actual price changes, which Midrigan finds is consistent with firm level scanner data.

Another contribution of Midrigan (2009) is a method by which to account for small price changes in the data. Midrigan introduces economies of scope in price setting in firms. Firms produce two related goods, and when the firm chooses to change the price of one of its goods, it gets to change the other price for free. Since the firms cluster around the median price and the firms will most often change the price of one of its goods when it receives a large idiosyncratic shock in that good, the free price change for the other good is often quite small. Combined with the leptokurtotic shocks, Midrigan's model is able to account for 80% as much monetary nonneutrality as a time dependent model and account for most of the micro facts about price changes in his data.

Nakamura and Steinsson (2009) create the Calvo-Plus model, a hybrid state and time dependent pricing model, to fit both the micro data and provide an alternative solution to the neutrality result of Golosov and Lucas. The model is a hybrid in the sense that firms face a certain probability of having the either a high or nearly zero menu cost, which is a state dependent pricing mechanism that has elements of (time dependent) Calvo timing. By introducing sector heterogeneity and intermediate inputs into the model, they are able to generate roughly nine times the amount of monetary nonneutrality as a baseline menu cost model. The intuition behind the importance of sector heterogeneity in the model is the idea that the first price change by a firm after a monetary shock contributes the most towards the economy's adjustment to a shock. The second time a firm adjusts its price after a monetary shock, it contributes little to the adjustment, because it has already priced in most of the shock. They find evidence that the frequency and size of price changes varies by sector in micro level pricing data. If a significant number of price changes occur within relatively few firms in the economy, then the firms that frequently change their prices will account for many of the price changes but contribute little to the monetary neutrality of the model. Firms that are in sectors in which prices change infrequently will contribute few price changes, and they will change their prices more sluggishly to the monetary shock but account for much of the monetary nonneutrality of the model. Overall, firms tend to change their prices less in response to exogenous aggregate shocks, which boosts the monetary nonneutrality produced by the model. Forcing some firms in the model to account for a large number of price changes triples the monetary nonneutrality produced by their standard menu cost model.

Intermediate inputs are particularly important for my paper as it is one possible labor market friction that can justify my generic source of real rigidity. The intermediate inputs mechanism of Nakamura and Steinsson is based on the roundabout production method of Basu (1995) where all goods are final goods and inputs for all other goods in the market. A larger share of intermediate inputs in production slows down the rate at which firms change their prices since firms' marginal costs respond more sluggishly to a monetary shock. Firms do not respond to the monetary shock per se; they respond to the changes in their marginal costs. In a standard menu cost model, aggregate wages will respond immediately and completely in response to a shock, which means that firms have a more immediate incentive to change their prices. With intermediate inputs, wages depend on the other firms' chosen prices, because the prices of firms are essentially the marginal costs of firms. The more incomplete the adjustment of marginal costs to a monetary shock, the more likely firms will not respond to the shock, and even if they adjust, adjustment will be incomplete. Hence, monetary nonneutrality increases significantly with intermediate inputs in the model.

Burstein and Hellwig (2007) investigate the size and importance of aggregate and firm level sources of real rigidity in a standard menu cost model. The aggregate real rigidity in their model takes the form of a generic labor market friction. Essentially, the wage that firms face in the model is the geometric weighted average of the money supply, which is the monetary shock and thus adjusts immediately to a monetary shock, and the price level, which responds sluggishly. The aggregate real rigidity is controlled by the weighted average parameter; as more weight is put on the price level, the slower firms will respond (in terms of changing their prices) to the monetary shock so long as other firms have an incentive to respond sluggishly to the shock as well. The decreasing returns to scale production function in labor is the firm level rigidity in the model. Burstein and Hellwig find that the firm level rigidity cannot generate much monetary nonneutrality with reasonable implications on the micro facts about prices. Basically, as decreasing returns to scale become stronger, firms cannot adjust by changing their production as much in response to a monetary shock, so they will be more likely to change their prices. This additional desire to change prices limits the additional rigidity that the model can produce. The bulk of the monetary nonneutrality produced by the model is from the aggregate real rigidity, which is consistent with the findings of Nakamura and Steinsson.

Gorodnichenko (2009) develops a menu cost model with imperfect information about a nominal demand shock but perfect information about the aggregate price level. Firms must choose whether or not to change their prices and whether or not to purchase a better signal about the nominal demand shock. Firms are hesitant to change their prices in response to a monetary shock because their price change contains their information about the demand shock and the state of the economy. The aggregate price level acts as a public signal about the state of the shock, and when enough firms change their prices, enough information is released about the shocks into the public signal to reach a tipping point, and many more firms change their prices. The hesitation on the part of firms to change their prices in response to any shock provides a novel source of real rigidity. Gorodnichenko's model, while highly stylized, can generate not only a large amount of monetary nonneutrality in response to a monetary shock, but a delayed and smoothed hump-shaped response of prices (and thus a hump-shaped response of output) as well.

Kehoe and Midrigan (2008) create a menu cost model in which firms can elect to make not only a typical, permanent price change but a temporary price change as well. The temporary price change accommodates sales and other markdowns, and can be thought of as a price rental. Firms have both a regular price and a desired sales price. If the firm wants to change its regular price, it must pay the full menu cost, but it maintains that new price permanently. If the firm instead wants to use a sales price, then it will "rent" as price change at a lower cost than a regular price change. It will change its price for a period, and the price will revert to its regular price in the next period. Their model is able to account for a reasonably large amount of monetary nonneutrality. They also show that models calibrated to micro data without sales overstate the response of the economy to the shocks, and models calibrated to price data with sales cannot account for much monetary nonneutrality. Instead, a reasonable, imperfect compromise for a menu cost model without temporary price changes is to try to match the percentage of time prices are at their annual mode.

1.1.3 Studies of Price Setting and Seasonality in Price Setting

Seasonality in price setting has been of interest to economists for some time, and there has been a debate about the motives that firms have to adjust their prices. Most papers that look at the micro price data tend to look at price changes in the aggregate rather than season by season, because most practical uses for this data require seasonally adjusted data. Researchers in this literature also tend to remove sales from the data using various algorithms to smooth from the data what they think is unimportant noise for models that focus on sources of aggregate fluctuations. Also, the frequency of price changes, especially sales, in high frequency data would imply that prices change very frequently in menu cost models, which are typically specified as either monthly or quarterly models. In the menu cost literature, other than Kehoe and Midrigan (2008), the models do not have a way to accommodate sales, and my model is the first DSGE models to accommodate explicitly seasonal price changes. I will eventually need to expand upon the literature presented below, especially the work of Nakamura and Steinsson (2008), to find additional facts about the seasonal aspect of price changes.

Bils and Klenow (2004) examine BLS data used to compute the CPI to find some basic facts about micro price changes. The BLS collects monthly data for 70,000-80,000 goods from 22,000 stores in 88 geographic regions across the United States. Bils and Klenow find considerable heterogeneity in price changes across goods, which implies that standard time dependent sticky price models do not do a good job modeling prices at the good level, particularly for goods that have less frequent price changes. Also, they find that the average time per price change is quite low; for half of the goods in the sample, the average duration of a price is less than 5.5 months. The median duration of nonsale price changes is about 4.3 months. They do not mention much about seasonality in price changes in their paper other to say that seasonally adjusted inflation features the lowest persistence and highest volatility, so only seasonal sales would not explain their findings.

Klenow and Kryvtsov (2008) find similar facts as Bils and Klenow in the same data. With sales, prices change on average every four months in the median category, and without sales, prices in the median category change on average every seven months. Price changes are typically large in absolute terms, about 10%, but there are many small price changes on the order of 5% or less. They also find that variance of aggregate inflation is mostly from changes in the intensive margin (the size) of price changes, rather than the extensive margin (the fraction of items) of price changes. They find that the 2004 vintage (and previous) of state dependent and time dependent models have a difficult time fitting all of the facts that they have at the same time. The Golosov and Lucas model does not generate enough small price changes, and the model of Dotsey, King, and Wolman does not generate enough large price changes. Time dependent models predict that older prices will feature larger changes when they do change, and they suggest an incorrect hazard of price changes. They do acknowledge that some of the newer models, such as Midrigan's model, can generate reasonable price changes, however.

Nakamura and Steinsson (2008) work with the same data as Bils and Klenow and Klenow and Kryvtsov and establish what they term are the five facts about prices. Notably for my model, the fourth fact is that price changes are highly seasonal. In particular, they found a monotonically decreasing trend of price changes over the four quarters of the year, after removing sales from the data. Within a quarter, the frequency of price adjustment decreases monotonically as well. They do not provide information about the size of price changes over different seasons, however. Disaggregating the price adjustments into price increases and price decreases, they find that the frequency of price decreases stays steady throughout the year while the frequency of price increases tends to fall monotonically through the year. Also, the hazard of price adjustment for some goods, which gives the probability that a firm will adjust its price given the length of time since its last price change, has a spike for some goods at 12 months, which implies that some prices change annually.

They also find other facts that are important for calibrating a menu cost model. The median frequency of non-sale price changes is about 9%-12% a month, half of what it is leaving in sales. Sales comprise a large percentage of price changes. Also, not all price changes, excluding sales, are price decreases. About one third of nonsale price changes are price decreases. Finally, the slope of the hazard function for individual prices tends to be downward sloping, which means that firms are more likely to adjust in consecutive periods than in any other pair of periods.

Looking at scanner data from Dominick's, a grocery store, in the Chicago area, Kehoe and Midrigan (2008) find that the price changes in that data, including sales, differ considerably from the BLS data that Bils and Klenow and Nakamura and Steinsson use. First, they find that prices change rapidly. In their sample, about 1/3 of prices change every week. While the size of relative price changes is high but volatile, firms spend most of the year (about 60%) at the mode, or regular, price. If a product deviates from its regular price, it is typically for a sale; 30% of the time, a good's price is below its regular price. Hence, sales account for most (83%) price changes in their data, and sales are highly transient; a product's price returns to its regular price the next week with about a 50% chance when the product is currently on sale.

One of the key pieces of the puzzle of creating a menu cost model with seasonality is to fit the high monetary nonneutrality, the high frequency of price changes in the first quarter of the year, and the tendency for seasonal sales to occur in some goods in the fourth quarter of the year (which also produces a slight decline in the CPI as documented by Barsky and Miron (1989)). If seasonality is indeed a motive for firms to adjust their prices, it would have to give firms a strong motive to change prices in the first quarter of the year, and the price changes in the fourth quarter of the year would have to be disproportionately price decreases. Seasonally varying demand shocks could be a culprit, or some other aggregate shock that boosts demand in the second half of the year could cause the fluctuations as well. But why do firms want to lower prices in periods in which demand for their goods is high? There are several competing theories.

Warner and Barsky (1995) evaluate the statistical properties of daily pricing data for eight goods from different stores around Ann Arbor, Michigan between November 1987 and February 1988. Their theory about why prices decrease when demand for the good is high has to do with economies of scale in search. For consumers, the cost of researching goods and travelling between sites is roughly a fixed cost, so when demand is high, consumers optimally search and travel more. Hence, consumers are more sensitive to prices during periods of high demand; they find and purchase goods at lower prices, which lowers the average price for the good.

Chevalier, Kashyap, and Rossi (2003) provide evidence that firms lower prices on certain goods in periods of peak seasonal demand to attract customers into their stores. They provide support for the loss leader model of price setting and advertising, formalized by Lal and Matutes (1994), in which the firms' optimal pricing decision is to lower prices on goods that are in the most demand. Consumers do not know all of the prices at a store before they travel there, but they do know the prices of goods that the firms' advertise. The firms then compete on the basis of a subset of advertised goods, in particular, goods for which there is a high demand. While firms do not earn high margins on the advertised goods in high demand, firms make up for the losses by charging relatively high prices on unadvertised goods. Consumers acquiesce to the high prices because of the inconvenience and uncertainty of travelling to a different store. The loss leader theory does not rely on aggregate demand fluctuations to generate sales; rather, idiosyncratic demand fluctuations for goods are sufficient to generate sales in that good. Chevalier, Kashyap, and Rossi document that prices for certain goods that face idiosyncratic seasonal demand shocks, such as tuna during Lent, decrease during the periods of higher demand. Prices decrease for other goods during periods of high aggregate demand as well, such as beer, cheese, soup, and crackers at Christmas, but the loss leader theory produces seasonal fluctuations regardless of their source.

Nevo and Hatzitaskos (2006) provide an alternative theory of seasonal price fluctuations that can account for sales during holidays and other periods of high demand. During peak periods of seasonal demand, regardless of whether it is an idiosyncratic or aggregate fluctuation, more consumers enter the market for certain goods. These consumers may differ from the consumers who typically purchase the good, and they may be more sensitive to the price of a good. For example, during Lent, the additional consumers who buy tuna may not care as much about the quality of tuna purchased as tuna connoisseurs who purchase tuna throughout the year. Thus, the new entrants would tend to purchase cheaper tuna than the regular tuna purchasers. Alternatively, during holidays, people would tend to enjoy beer at parties or other events, boosting the demand for beer. As Nevo and Hatzitaskos put it (p. 2-3), "after a few beers, it is hard to distinguish between different brands," so consumers of beer may decide to buy cheaper varieties of beer. As the frugal consumers enter the market, the average price of beer sold in the market may decline regardless of any sales in the market. Even if there are sales, they may not account for as much of the decline in average prices as the product substitution in markets. Nevo and Hatzitaskos reestimate the results of Chevalier, Kashyap, and Rossi using fixed weights price indices for the goods, which a change in the market share of brands will not affect. They find that the seasonal fluctuations in prices resulting from sales rather than substitutions are much less than implied by Chevalier, Kashyap, and Rossi. Also, for tuna specifically, demand for higher quality white tuna does not increase much during Lent; the additional demand for tuna is predominantly for the cheaper, light tuna, and the prices of the two varieties of tuna that gain the most market share do not decrease. Both of these facts are inconsistent with the loss leader model and are consistent with the idea that product substitutions account for a significant portion of fluctuations in prices.

The loss leader theory and the substitution theory can account for seasonality in prices from both idiosyncratic and aggregate demand sources, but the search theory and the countercyclical markup theory can only account for seasonality in prices from fluctuations in aggregate demand. Hence, an important issue is the primary source of seasonality in prices. If idiosyncratic sources of seasonality only affect a small subset of goods, then the debate in the theory is not necessarily particularly important for the menu cost literature. Aggregate demand fluctuations would sufficiently account for the seasonality motive in price setting for firms. Bryan and Cecchetti (1995) examine seasonality in CPI components from 1982 to 1993 and find that seasonality from idiosyncratic sources dominates seasonality from a common aggregate source. They perform a linear decomposition of individual price movements into an average seasonally adjusted price movement, an average aggregate seasonal price movement, an idiosyncratic seasonal price movement, and measurement error. The ratio of the idiosyncratic seasonal variance to the aggregate seasonal variance across the goods provides some insight about the relative importance of the two sources of fluctuations. They find that only two out of the thirty-six goods in their sample, auto repair and food away from home, have a variance ratio of less than one, which implies that few goods have small idiosyncratic contributions to overall seasonal price fluctuations. Furthermore, several goods feature extremely large variance contributions from the idiosyncratic seasonal component. The variance ratio for motor fuel, fruits, gas and electricity, fuel oil, and women's apparel all exceed 100. Their analysis ignores the fact that the elasticity of demand and supply for goods is not necessarily the same, however, which could understate the importance of aggregate price changes on goods with a high elasticity of demand or inelasticity of supply. Another fact that they point out is that the source of seasonal fluctuations in the CPI goods does not necessarily come from a single source. Regressing the prices of goods on deterministic seasonal trends, they find that the months in which seasonality in price setting appears varies between the goods and may identify the source of the price fluctuations. For example, cereal and fruit prices tend to decline in autumn when supply is abundant, but they tend to increase in January when fresh supply is scarce. Public transportation, natural gas, and electricity prices have on average large January price increases, which could be an effect of regulation. Fluctuations from demand may not be the only reason for fluctuations in prices.

2 Model

A continuum of firms, indexed by i on $[0,1]$, hire labor from households to produce goods in a monopolistically competitive market. The production function is constant returns to scale in labor:

$$c_t(i) = a_t(i) l_t(i)$$

Firms operate along their demand curves, so consumption of good i, $c_t(i)$, is equal to production. Firm i hires labor, $l_t(i)$, and faces a productivity shock $a_t(i)$ that takes the following form:

$$\log a_t(i) = \rho_a \log a_{t-1}(i) + \vartheta_t$$

The innovation ϑ_t is iid normally distributed with mean 0 and variance σ_a^2. The firms current period nominal profit is then:

$$\pi(p) = p_t(i) c_t(i) - w_t l_t(i)$$

Consumers in the economy gain utility from consumption and leisure. They provide labor l_t to the firms at a nominal wage W_t and earn dividends from equal ownership shares of the firms in the economy. Since consumers are identical in the economy, the consumers' problem is summarized by a representative consumer's problem. The representative consumer seeks to maximize the lifetime discounted sum of utility defined over the aggregated consumption good and labor. The consumer's problem is:

$$\max E_t \sum_{t=0}^{\infty} \beta^t U(C_t, l_t)$$

$$\textit{subject to}$$

$$\int_0^1 p_t(i) c_t(i) di \leq W_t l_t + \pi_t$$

The aggregate consumption good is C_t, the aggregate price level is P_t, the nominal wage is W_t, the share of profits that the firm receives from ownership of the firms is π_t, and the labor supply of the worker is l_t. Aggregate consumption and the aggregate price level come from the standard CES aggregators:

$$P_t = \left[\int_0^1 p_t(i)^{1-\epsilon} di\right]^{\frac{1}{1-\epsilon}}$$

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

The consumer discounts the future at a rate of β. Utility in the model will be time separable and separable in the consumption good as follows:

$$U(C_t, l_t) = \frac{C_t^{\gamma_s}}{\gamma_s} - \chi l_t$$

The labor-leisure condition from the firm's problem is:

$$w_t = \chi C_t^{1-\gamma_s} \tag{1}$$

Note that w_t is the real wage rate, defined as the nominal wage over the aggregate price level. Consumer demand is determined by the cost minimization motive of the consumer across all consumption bundles:

$$\min \int_0^1 p_t(i) c_t(i) di$$

$$\textit{subject to}$$

$$U^* \leq \frac{C_t^{\gamma_s}}{\gamma_s} - \chi l_t$$

Consumer demand is then simply:

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

The representative consumer's demand for good i is $c_t(i)$, and the price of good i at time t is $p_t(i)$. Note that the form for consumer demand is independent of the curvature of the utility function for the simple power/logarithmic forms of utility, so the curvature of the utility function only affects the representative consumer's labor-leisure condition.

I introduce money by letting it equal nominal demand.

$$M_t = P_t C_t \tag{2}$$

Effectively, money in the model is the sum of labor income and profits of the firm as suggested by the binding budget constraint of the consumers. Money grows exogenously in the model following the process:

$$M_t = \mu_t M_{t-1}$$

The money growth rate follows the exogenous stochastic process:

$$\log(\mu_t) = \bar{\mu} + \rho_u \log(\mu_{t-1}) + \eta_t$$

Here, η_t is a monetary disturbance.

The model presented here is similar to the other models in the literature, most notably Burstein and Hellwig (2007). Relative to Burstein and Hellwig, I omit the demand shock and the decreasing returns to scale production function, and I explicitly derive the representative consumer's problem. Recall from the representative consumer's labor-leisure condition (1) that the real wage is related to aggregate consumption, the disutility of labor, and the curvature of the utility function. Substituting out C_t from (1) using (2):

$$w_t = \chi \left(\frac{M_t}{P_t} \right)^{1-\gamma_s}$$
$$W_t = \chi M_t^{1-\gamma_s} P_t^{\gamma_s}$$

Let $\chi = 1$, and thus:

$$W_t = M_t^{1-\gamma_s} P_t^{\gamma_s} \tag{3}$$

The model, with the restriction that $0 \leq \gamma_s < 1$, implies that the nominal wage is the geometric average of the money stock and the aggregate price level, which is a fundamental assumption in Burstein and Hellwig (2007). They assume this form for the law of motion for the nominal wage noting that it can be derived a result of increased curvature in the utility function relative to the standard logarithmic preferences or the result of including intermediate inputs into the model. Further, it would be difficult to argue that the curvature of the utility function is quarterly dependent for the representative consumer in the economy. The representative consumer's problem here is thus more pragmatic than realistic.

The firm's intertemporal problem is to maximize its lifetime discounted sum of normalized profits. No analytical solution for this problem exists because of the menu cost, denoted ξ.

The system of functional equations that define the steady state of the model is:

$$V_a(e^{\eta(i)}, s) = \max_{\tilde{p}^*(i)} \left(\pi \left(\tilde{p}^*(i), e^{\eta(i)}, s\right) + \hat{\beta} \int_0^1 V \left(\frac{\tilde{p}^*(i)}{e^{\mu}}, e^{\eta(i)\prime}, s' \right) dF(\delta, \eta) \right) \quad (4)$$

$$V_n(\tilde{p}_{-1}(i), e^{\eta(i)}, s) = \pi(\tilde{p}_{-1}(i), e^{\eta(i)}, s) + \hat{\beta} \int_0^1 V \left(\frac{\tilde{p}_{-1}(i)}{e^{\mu}}, e^{\eta(i)\prime}, s' \right) dF(\delta, \eta) \quad (5)$$

$$V\left(\tilde{p}_{-1}(i), e^{\eta(i)}, s\right) = \max\left(V_a\left(e^{\eta(i)}, s\right) - \tilde{w}\xi, V_n\left(\tilde{p}_{-1}(i), e^{\eta(i)}, s\right)\right) \quad (6)$$

The firm's stochastic discount factor is $\hat{\beta}$. The significance of γ_s becomes clear once the adjustment cost is omitted. The first order condition of that problem implies:

$$\log\left(\tilde{p}^*(i)\right) = \log\left(\frac{\epsilon}{\epsilon - 1}\right) + \gamma_s \log\left(\tilde{P}\right) - \log\left(a(i)\right) \quad (7)$$

Since the prices are normalized, the money supply does not enter directly into (7). Note that as γ_s increases, the firm is more pressed to keep its price close to the prices of other firms (the aggregate price level). Firms will be hesitant to make the sorts of price adjustments necessary to generate quick adjustment of prices to the monetary shock. Prices will adjust less quickly to the monetary shocks in the economy. Furthermore, the size of price adjustment should be lower on average since the firm will also want to keep its relative price low. Therefore, γ_s is the main source of real rigidity in the economy.

2.1 Planner Equilibrium

First, I seek the steady state equilibrium of the economy, which simplifies to finding the steady state of the equivalent planner economy of the model. The planner economy is the economy in which there is no aggregate uncertainty from monetary shocks, and thus the money supply growth rate is exogenous and known. Hence, the persistence of the money supply does not matter, and thus $\log(\mu_t) = \bar{\mu}$. A more detailed summary of my solution method is in Appendix A and Appendix B, but a summary of the equilibrium is below.

The steady state of the model with seasonality, which I will call the cyclostationary equilibrium, is a generalization of the steady state of the model without seasonality. The crux of the model without seasonality is the steady state normalized wage, which satisfies (3). The normalized wage implies pricing rules for the firms for each possible state after substituting the wage into (4), (5), and (6) and noting that $q = q'$. These pricing rules imply a stationary distribution of firms across the states in the economy. The pricing decision rules and the distribution of firms across prices determine the normalized price level through the standard CES price aggregator, which needs to be consistent with (3). Note that to find the aggregate consumption in the economy, I use (2).

The cyclostationary equilibrium is defined by the distribution of firms across prices and the normalized wages for each season within the cycle. The stationary distribution of firms in the nonseasonal case is no longer technically stationary since it is not the same every period. Rather, it also must be cyclostationary; the distribution of firms across states must be the same every s periods. Also, the pricing rules and the value functions for every season in the cycle now enter into the other seasons' problems. To determine the distribution of firms across prices for the same season in the next cycle, the firms in the current season must progress through every other seasons' pricing rules. The expected value function for the next season explicitly enters into the current season's problem. Other than those technical complications, equilibrium is basically the same as in the case without seasonality. The crux is the set of steady state normalized wages, which imply the pricing rules with (4), (5), and (6) for each period. Equation (2) still provides aggregate consumption. The cyclostationary distribution is still a consequence of the pricing rules.

2.2 Parameterization

The model parameters that I do not calibrate are largely from Burstein and Hellwig (2007). Since the models are very similar, the parameters that best fit Burstein and Hellwig's model without the demand shock should provide a good first approximation of the parameters necessary for my model to fit the data.

Table 1: Baseline Parameterization

Parameter	Use	Value
ϵ	elasticity of substitution	4
β	discount factor	$0.935^{\frac{1}{12}}$
ρ_u	persistence of money shock	0
$\overline{\mu}$	money growth rate	$\log(1.021^{\frac{1}{12}})$
ξ	menu cost*	0.0047322
$\overline{\gamma_s}$	wage rigidity*	0.81786
ρ_a	persistence of technology shock*	0.38355
σ_a	standard deviation of technology shock*	0.074838

*: calibrated

In a baseline case without seasonality, I calibrate ρ_a, $\overline{\gamma_s}$, ξ, and σ_a by minimizing, using a Nelder-Mead algorithm, the sum of weighted squared deviations of the nonseasonal model's moments away from moments obtained from the Dominick's scanner price data set by Midrigan (2009). I choose to match the mean price change (0.1%), the mean absolute price change (12%), the standard deviation of the absolute value of price changes (0.09), and the frequency of price changes (24%). The nonseasonal model's fit to the data is not spectacular; the mean absolute price change is too high (16%), the mean price change is too high (0.7%); the standard deviation of the absolute value of the price changes is too low (0.05); and the frequency of price changes is too high (27%). Menu cost models have difficulty matching the standard deviation of the absolute value of price changes because they lack small price changes, and the other moments are within a couple of percent of the desired fit, which should be good enough for an initial fit.

In the seasonal case, I use the same parameters as the fit above, subject to the seasonal fluctuation discussed below. Surprisingly, the moments averaged across seasons in the seasonal model fit the desired moments better than the nonseasonal model, particularly for the mean absolute price change (14%).

The baseline calibrated parameters are quite consistent with the existing literature. The key parameters in the model are γ_s and ρ_a because they will control the magnitude of the response to seasonal fluctuations. The calibrated value for γ_s fits nicely into the existing literature about labor market rigidities. Burstein and Hellwig (2007) estimate a value for a firm specific rigidity and an aggregate (labor market) rigidity. The true real rigidity of the model is a function of both parameters. Their estimates of the true real rigidity and the aggregate real rigidity are between 0.75 and 0.85. Nakamura and Steinsson's (2008) materials share of intermediate inputs is similar to the real rigidity in this model. They chose a materials share of 0.7. Rotemburg and Woodford (1997) estimate a firm specific rigidity and an aggregate rigidity of around 0.8. My calibrated technology shock persistence is slightly lower than what other authors find but is still reasonable. For example, Burstein and Hellwig assume a persistence of 0.5, Midrigan finds a persistence of 0.483 by calibrating his model to similar moments, and Nakamura and Steinsson choose a persistence of 0.7.

The second problem is to choose the spread between the highest γ_s and the lowest γ_s, which will determine the amplitude of the seasonal cycle. Barsky and Miron (1989) find an approximately 8% fluctuation trough to peak in detrended GDP from 1948 to 1985 with the peak of the consumption in the fourth quarter of the year. Table 2 summarizes their results.

Table 2: Seasonal Patterns 1948-85

Abridged Table 2 from Barsky and Miron (1989)

(Percent Deviations from Trend)

	Quarter			
Variable	1	2	3	4
GDP	-3.76	-0.02	-0.53	4.31
Price Level	-0.15	-0.01	0.16	0
Nominal Wage	-0.01	-0.09	0.01	0.1
Real Wage	0.13	-0.17	-.013	0.17

Their regression results and the theoretical results of Mehrez (1998) imply that a fairly large spread in γ_s will be necessary to account for the seasonal cycle, but I find that a relatively small fluctuation in the nominal value of the parameter sufficiently generates a large seasonal fluctuation in real variables. I choose values for the maximum and minimum of γ_s within 0.0075 of the of baseline, because if the gap between the pricing decisions of the model is too large, too many firms will change their prices in the transition period between the extreme values of γ_s. For example, all firms adjust in every period if the difference between γ_s each period is 0.2 or more. I choose lower values of γ_s at the end of the year, consistent with both the idea that there is a boom in the later half of the year in the typical seasonal cycle and Olivei and Tenreyro (2007), whose staggered wage adjustment mechanism implies more rigidity at the beginning of the year. For simplicity, I assign each γ_s in my twelve month model using $\gamma_s = 0.0075. * \sin\left(\frac{2\pi(s-2)}{12}\right) + \overline{\gamma_s}$ for $s = 1, 2, ..., 12$. As it turns out, small fluctuations in γ_s are sufficient to generate an impressive seasonal cycle, though the simple form that I use for the fluctuation does not capture the fact that consumption booms mainly in the 4th quarter and decreases at varying rates relative to the trend in other quarters. Also, note that I am trying to match the seasonal fluctuations in the real part of the economy with this model. The miniscule nominal fluctuations imply that I need to introduce money into the model in a little more detailed way than the constant velocity that I have assumed here.

3 Results

Figure 1 below shows, for γ_s reported in Table 1, aggregate consumption for the baseline cyclostationary equilibrium, the cyclostationary equilibria with progressively lower idiosyncratic persistence, and the individual steady states assuming no seasonal dependence across months.

Figure 1

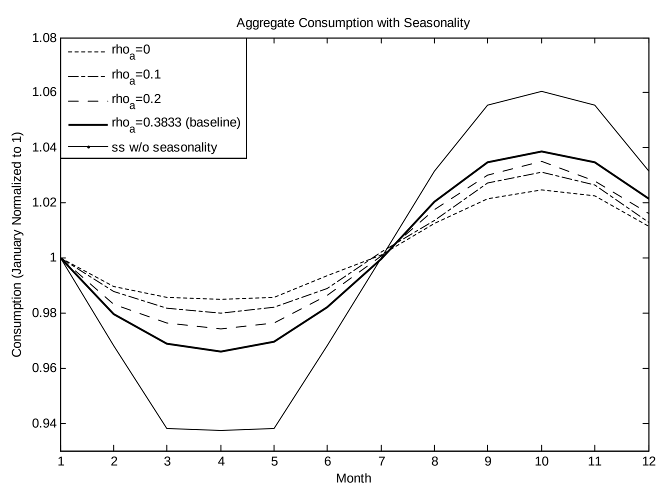

One of the most striking features of Figure 1 is the conflict between the firms' desire to adjust to its idiosyncratic shock and their desire to adjust to the aggregate fluctuation. One could think of the steady states without seasonality as maximizing the effect of the firms' idiosyncratic shocks and of the cyclostationary equilibrium where the firms' idiosyncratic shocks are iid as maximizing the effects of the aggregate seasonal fluctuation. In the individual steady states, the firms face neither a fluctuating future aggregate price nor an aggregate fluctuation in their marginal costs, and hence, in the sum, they fully adjust for the seasonal wage fluctuation. When the firms' idiosyncratic shock is completely transient, the firm optimally focus more on the aggregate fluctuation in the economy. Since γ_s, the source of real rigidity in the model, is positive, the firms are penalized for setting prices different than the aggregate price level. Only if a firm's transient idiosyncratic component is extreme relative to the other firms will it adjust primarily considering its idiosyncratic motive. Otherwise, the firms will primarily adjust to stay near the aggregate price level, which will fluctuate according changes in γ_s, but since adjustment to aggregate fluctuations is incomplete, the firms will not fully adjust to the seasonal fluctuation in the aggregate. As the persistence of the productivity shock increases, the relative weight that the firms place on their productivity shock when making pricing decisions must increase, because the shock is more likely to affect the firms for a longer time. Hence, firms adjust relatively less to stay near the aggregate prices for the firms, and the effect of the seasonal fluctuation in γ_s is greater in the aggregate. With the persistence set at the baseline value, the seasonal fluctuation moves prices more than in the iid case, but the firms still do not fully adjust for the seasonal fluctuations. When the persistence is even higher than the baseline value (not reported), the idiosyncratic component matters even more, and consumption fluctuates more than the baseline case across the seasonal cycle.

Figure 2 shows the seasonal distribution of price changes conditional on adjustment for January, a representative month for comparing the seasonal and nonseasonal models. The top panel is the model that considers seasonal fluctuations. The bottom panel shows the model without seasonal fluctuations (a standard menu cost model).

Figure 2

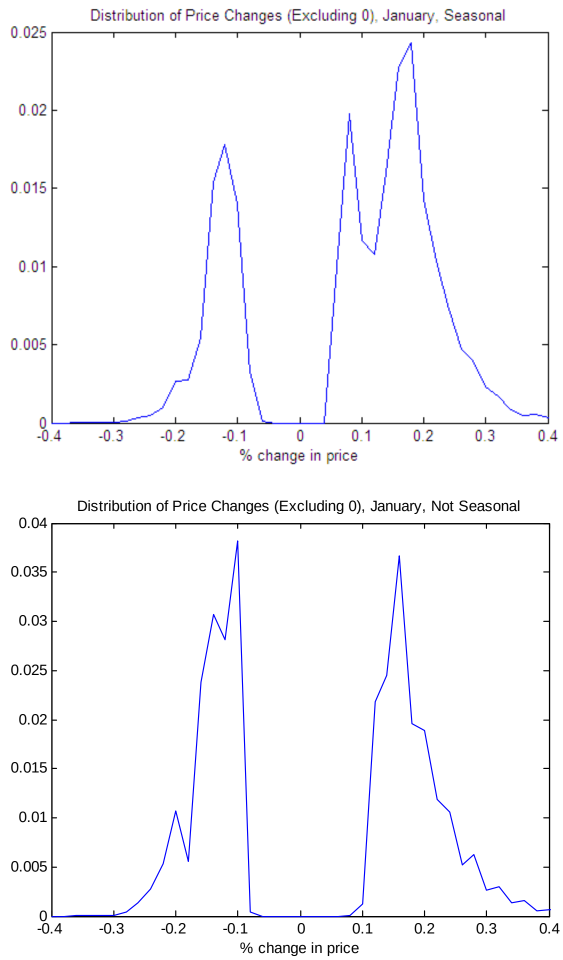

Comparing the two panels provides a couple of observations. First, some of the firms have smaller price changes in the seasonal model, which is indicative of the clustering motive that firms have when facing the added challenge of adjusting to the moving, seasonally dependent aggregate price. Firms face a penalty for a price that is not near the aggregate price level because of the real rigidity, ceteris paribus, so they must balance the needs of adjusting for current conditions and adjusting such that they do not need to change their prices as urgently in the future. This desire will change depending on the season. Second, there is a potential reduction of the selection effect to a monetary shock. The mass of firms that is adjusting to seasonal motives has smaller price changes, which shifts firms out of the tails of the distribution of price changes. Given a small monetary shock, the selection effect implies that a margin of firms with more extreme shocks or prices far away from equilibrium will adjust in addition to the firms that would otherwise adjust. If there is a sufficient mass of firms that adjust to the shock, the response of prices to the monetary shock will be large, thus limiting the ability of the model to explain a large real response to the shock.

At the adjustment margin, where price changes are larger, there are fewer firms in the case with seasonality. Some additional firms could adjust for the seasonal motive, but the fact that their price changes are smaller and that they target the aggregate price level should significantly mitigate their response to a monetary shock. On the whole, the selection effect should be mitigated in the seasonal model.

A potential concern about the model is how it aggregates, especially since the price changes of the firms vary widely month to month. Specifically, if the model does not yield price changes aggregated across the entire year that match the micro price data and other models, then it is not necessarily a good approximation of the behavior of firms. Table 3 shows the vital moments of the distribution of price changes month by month, in the aggregate, and in the steady state without seasonal dependence.

Table 3: Monthly Moments of Seasonal Price Changes

Month	γ	$\% \triangle P$	$\% \triangle P \uparrow$	$\%\overline{\triangle P}$	$\%\overline{abs(\triangle P)}$
January	0.8179	29.2	23.5	9.2	14.4
February	0.8216	27.4	22.1	9.4	14.1
March	0.8243	24.3	17.2	6.7	13.6
April	0.8253	22.9	13.4	3.4	13.3
May	0.8243	23.3	11.4	0.7	12.9
June	0.8216	24.4	8.7	-2.8	12.7
July	0.8179	25.2	7.8	-4.5	12.9
August	0.8141	26.3	7.4	-5.3	13.0
September	0.8114	26.8	9.7	-3.0	13.0
October	0.8104	26.5	13.1	0.7	13.0
November	0.8114	26.1	15.9	3.8	13.6
December	0.8141	26.2	19.0	7.2	14.4
Aggregate	-	25.7	14.1	2.2	13.4
Mean Steady State	-	26.8	14.4	2.6	15.6

Table 3 (cont.): Monthly Moments of Seasonal Price Changes

Month	$std(\triangle P)$	$std(abs(\triangle P))$	$kurt(\triangle P)$
January	0.12	0.05	2.86
February	0.11	0.05	3.15
March	0.13	0.04	2.08
April	0.13	0.04	1.51
May	0.13	0.04	1.42
June	0.13	0.04	1.80
July	0.13	0.04	2.11
August	0.12	0.04	2.33
September	0.13	0.04	1.75
October	0.14	0.04	1.43
November	0.14	0.04	1.57
December	0.13	0.04	2.18
Aggregate	0.13	0.04	2.01
Mean Steady State	0.16	0.05	1.39

Comparing the seasonal and nonseasonal aggregate moments, the clustering motive of the firms in response to the seasonal fluctuation is still evident in a decrease in the standard deviation of price changes and a slight decrease in the mean price change. Also, the kurtosis of the distribution of price changes increases significantly in the seasonal case, which could be a sign of a mitigated selection effect. The strongest effects in Table 3 are the effects on the mean price changes and the extensive margin of price changes. The mean price changes closely follow the seasonal fluctuation in aggregate prices, so seasonal fluctuations in them are expected. Except for the final quarter of the year, the number of price increases in the model decreases monotonically, which is consistent with the seasonal micro price findings of Nakamura and Steinsson (2008). The behavior of all price changes and price decreases are inconsistent with their findings. The number of price changes in the model fluctuate much less than price increases, and the number of price decreases increases every quarter. Nakamura and Steinsson find that the number of price decreases does not change much across the year. Also, they find that one third of nonsale price changes are price decreases, while about 44% of price changes are price decreases in my baseline seasonal model. Whether the behavior of prices in my model is inconsistent with Nakamura and Steinsson, however, is an open question, because they filter out v-shaped sales, which could eliminate seasonal decreases in prices in particular. Sales, therefore, may not be trivial and disposable in a menu cost model. Since I do not presently have access to the confidential BLS database that

Nakamura and Steinsson use, I cannot do more than conject that the micro price moments generated by my model are consistent with the micro data.

It appears that the seasonal motive to change prices in the seasonal model could help a standard menu cost model partially account for the small price changes observed in the data. While the model without seasonality does not have any price changes less than $\frac{1}{2}$ of the mean of price changes, the model with seasonality can account for some of the price changes below $\frac{1}{2}$ of the mean price change. Neither model can account for the smallest price changes that occur in the data, however. Alone, the seasonal model cannot account for all small price changes in the data, but the model feature could supplement other features, such as the economies of scope mechanism in Midrigan (2009) or the high/low menu cost mechanism of Nakamura and Steinsson (2009).

4 Conclusion

Economists have typically viewed the seasonal cycle as secondary to the business cycle, but the results of my model indicate more than a casual connection, particularly in sticky price models with strong real rigidities. The seasonal cycle may provide important information about business cycle fluctuations.

I construct a menu cost model with a seasonally fluctuating labor market rigidity that has the potential to fit the micro data about prices, the seasonal cycle of the economy, and the seasonally varying response of the economy to monetary shocks. The model can readily explain a 7.5% gap between the peak and trough of the seasonal cycle with a minimal change in the labor market rigidity. Also, the pattern in the kurtosis of the conditional distribution of price changes indicates that the model will produce more monetary nonneutrality in the first half of the year than in the second half of the year, and the kurtosis of the conditional aggregated distribution of firms indicates that the model with seasonality produces more monetary nonneutrality in response to a shock than the baseline model without seasonal fluctuations. Essentially, firms have a greater incentive to cluster their prices near when their optimal prices are seasonally fluctuating unless they receive a large idiosyncratic shock. In that case, the firms will adjust as if they will adjust again in the next period. This clustering effect mitigates the selection effect by amplifying the real rigidity that the model produces.

The key determinants of the effects of seasonality are the size of the labor market rigidity and the persistence of the technology shock. As the size of the labor market (real) rigidity increases in the model, the profit penalty for firms that deviate far from the aggregate price level increases. Since seasonality moves the aggregate price level, firms would increasingly cluster near the aggregate price level to avoid lost profit from fluctuations. The aggregate price level is endogenously determined by the firms' behavior, so limitations in the firms' response to fluctuations limit the fluctuation of aggregate variables. The persistence of the technology shock demonstrates the tradeoff that the firms face between adjusting for the idiosyncratic technology shock and the seasonal fluctuation in the economy. As the persistence of the technology shock increases, any extreme idiosyncratic shock that the firm receives is less transient and has more of an effect on firm behavior. In response, aggregate variables and firms' prices respond more to the seasonal fluctuation. Simply, the idiosyncratic component becomes more important to the firms, which offsets the seasonal clustering motive.

There are some extensions to the paper that I will pursue in the near future. First, I could add seasonal demand shocks to the model to test whether they could generate the same seasonal fluctuations as the model here. This feature may sort out whether labor market frictions, seasonal demand shocks, or a combination of the two are relevant for seasonal fluctuations in prices. Second, I need to solve for more than the steady state of the model; I need to estimate the effects of a monetary shock. I will extend the calculations once it is computationally feasible. Third, I could change the type of model that I am using to help the model generate additional monetary nonneutrality in response to a monetary shock. Burstein and Hellwig's model is simple, and there are a variety of extensions that I could undertake based on other research. The sector heterogeneity of Nakamura and Steinsson (2009) is a particularly appealing feature, because the timing and amount of seasonality in prices differs by sector. If a minority of firms account for the majority of price changes, then they should also account for most of the seasonal fluctuations, which would increase the pressure placed on other firms to cluster together, thus further amplifying the real rigidity in the model. If the persistence of idiosyncratic shocks varies by sector and the logical effects of seasonality hold in the new economy, then the model should also be able to explain the finding of Boivin, Giannoni, and Mihov (2009), who find that monetary shocks affect sectors with less volatile and less persistent idiosyncratic shocks more slowly.

References

Ball, Laurence and Romer, David (1990). "Real Rigidities and the Non-neutrality of Money," Review of Economic Studies, 57(2), 183-203

Barsky, Robert and Miron, Jeffrey, (1989). "The Seasonal Cycle and the Business Cycle," Journal of Political Economy, 97(3), 503-34

Basu, Shantanu (1995). "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," American Economic Review, 85(3): 512-531

Boivin, Jean, Giannoni, Marc, and Mihov, Ilian (2009). "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data," American Economic Review, 99(1), 350-384

Broszkiewicz-Suwaj, Ewa, Makagon, Andrzej, Weron, Rafal, and Wylomanska, Agnieszka (2004). "On Detecting and Modeling Periodic Correlation in Financial Data," Physica A: Statistical and Theoretical Physics, 336, 196–205

Bryan, Michael and Cecchetti, Stephen (1995). "The Seasonality of Consumer Prices," NBER Working Paper 5173.
http://www.clevelandfed.org/Research/review/1995/95-q2-bryan.pdf

Bils, Mark and Klenow, Peter (2004). "Some Evidence on the Importance of Sticky Prices," Journal of Political Economy, 112, 947-985

Boivin, Jean, Giannoni, Marc, and Mihov, Ilian (2009). "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data." American Economic Review, 99(1), 350-384.

Braun, Anton and Evans, Charles (1995). "Seasonality and Equilibrium Business Cycle Theories," Journal of Economic Dynamics and Control, 19(3), 503-531

Burstein, Ariel and Hellwig, Christian (2007). "Prices and Market Shares in a Menu Cost Model," Working Paper.
http://www.econ.ucla.edu/people/papers/Hellwig/Hellwig415.pdf

Caplin, Andrew and Spulber, Daniel (1987). "Menu Costs and the Neutrality of Money," Quarterly Journal of Economics, 102(4): 703-725

Chevalier, Judith, Kashyap, Anil, and Rossi, Peter (2003). "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data." American Economic Review, 93(1): 15-37

Dotsey, Michael, King, Robert, and Wolman, Alexander (1999). "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," The Quarterly Journal of Economics, 114(2), 655-690

Dotsey, Michael and King, Robert (2005). "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models," Journal of Monetary Economics, 52(1), 213-242

Eichenbaum, Martin, Jaimovich, Nir, and Rebelo, Sergio (2008). "Reference Prices and Nominal Rigidities," NBER Working Paper 13829.
http://www.nber.org/papers/w13829

Franses, Philip Hans (1996). Periodicity and Stochastic Trends in Economic Time Series, Oxford University Press, Oxford

Franses, Philip Hans, Paap, Richard. Periodic Time Series Models, Oxford University Press, Oxford, 2004.

Gertler, Mark and Leahy, John (2006). "A Phillips Curve with an Ss Foundation," NBER Working Paper 11971

Golosov, Mikhail and Lucas, Robert E. Jr (2007). "Menu Costs and Phillips Curves," Journal of Political Economy, 115(2), pp 171-199

Gorodnichenko, Yuriy (2009). "Endogenous Information, Menu Costs and Inflation Persistence," NBER Working Paper W14184.
http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1165512#

Kaneda, Mitsuhiro and Mehrez, Gil (1998). "Seasonal Fluctuations and International Trade," Working Paper.
http://129.3.20.41/eps/it/papers/9809/9809001.pdf

Kehoe, Patrick and Midrigan, Virgiliu (2008). "Temporary Price Changes and the Real Effects of Monetary Policy," Working Paper.
http://homepages.nyu.edu/~vm50/research.html

Khan, Aubhik and Thomas, Julia (2003). "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?" Journal of Monetary Economics, 50(2), 331-360

Klenow, Peter and Kryvtsov, Oleksiy (2008). "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" Quarterly Journal of Economics, 123(3), 863-904

Kryvtsov, Oleksiy and Midrigan, Virgiliu (2008). "Inventories, Markups, and Real Rigidities in Menu Cost Models," Working Paper.
http://homepages.nyu.edu/~vm50/inventoryApril2008.pdf

Lal, Rajiv and Matutes, Carmen (1994). "Retail Pricing and Advertising Strategies," The Journal of Business, 67(3), 345-370

Leskow, Jacek (2001). "The Impact of Stationarity Assessment on Studies of Volatility and Value-at-risk," Stable Non-Gaussian Models in Finance and Econometrics, Mathematical and Computer Modelling 34 1213–1222

Liu, Zheng (2000). "Seasonal Cycles, Business Cycles, and Monetary Policy," Journal of Monetary Economics, 46(2), 441-464

Mehrez, Gil (1998). "State Dependent Adjustment in an Economy with Seasonal Fluctuations," Working Paper.

http://ssrn.com/abstract=3471

Midrigan, Virgiliu (2008). “Menu Costs, Multi-Product Firms, and Aggregate Fluctuations,” Working Paper.

http://homepages.nyu.edu/~vm50/Main_RR.pdf

Nakamura, Emi and Steinsson, Jón (2008). “Five Facts about Prices: A Reevaluation of Menu Cost Models,” Quarterly Journal of Economics, 123(4), 1415-1464

Nakamura, Emi and Steinsson, Jón (2009) “Monetary Non-Neutrality in a Multi-Sector Menu Cost Model,” Working Paper.

http://www.columbia.edu/~en2198/papers/heterogeneity.pdf

Nevo, Aviv and Hatzitaskos, Konstantinos (2006). “Why Does the Average Price Paid Fall During High Demand Periods?” Working paper.

http://faculty.wcas.northwestern.edu/~ane686/research/tuna3.pdf

Olivei, Giovanni and Tenreyro, Silvana (2007). “The Timing of Monetary Policy Shocks,” American Economic Review, 97(3), 636-663

Pagano, Marcello, Parzen, Emanuel (1979). "An Approach to Modeling Seasonally Stationary Time Series," Journal of Econometrics 9(1-2), 137–153.

Rotemberg, Julio, and Woodford, Michael (1997). “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy,” 1997 NBER Macroeconomics Annual, 297-346

Serpedin, Erchin, Panduru, Flaviu, Sari, Ilkay, and Giannakis, Georgios (2005). "Bibliography on Cyclostationarity," Signal Processing, 85, 2233-2303

Tauchen, George (1986). "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." Economic Letters, 20, 177-181.

Warner, Elizabeth and Barsky, Robert (1995). “The Timing and Magnitude and Retail Store Markdowns: Evidence from Weekends and Holidays,” Quarterly Journal of Economics 110(2), 321–52

Van Loan, Charles (1997). Introduction to Scientific Computing: A Matrix-Vector Approach Using Matlab. Prentice Hall, Upper Saddle River, NJ

A Solving the Model

Figure A1 outlines the computational algorithm that I employ to solve for the steady state of the model without seasonality. The objective of the algorithm is to solve for (6).

Figure A1

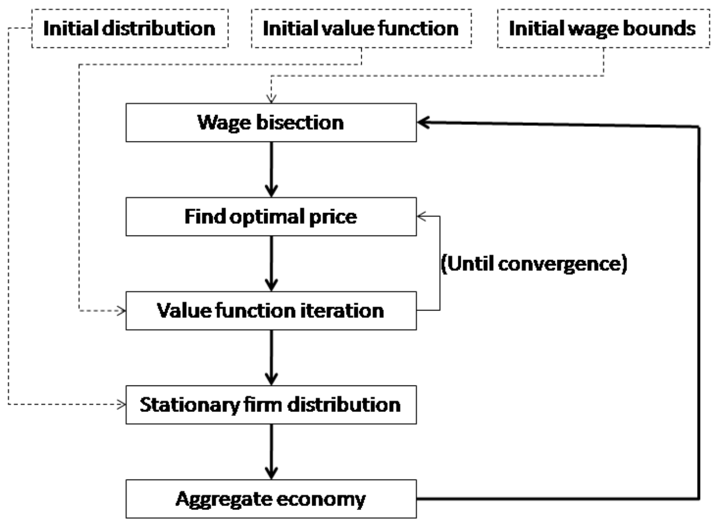

The computational algorithm is similar to the method used in Khan and Thomas (2003). First, I must guess an initial distribution of firms across prices, an initial value function, and initial bounds on the possible equilibrium wage. Also, I must discretize the productivity shock and the normalized prices. The algorithm solves for the steady state wage using a bisection algorithm starting from the initial wage bounds. The wage bisection algorithm provides a guess for the steady state wage in the economy. Using the model equations, all other endogenous variables, except for the firm's optimal price, can be expressed in terms of this wage guess, so for each state, I compute the firms optimal price using a golden section search algorithm on (4). To approximate the expected value function, I employ a cubic interpolation spline. Once I know the optimal price for each state, I know the value function implied by the wage and those prices. However, this value function is not necessarily the same as the one used for the spline approximate, so I perform value function iteration, repeatedly solving for optimal prices across states given the current approximation of the value function and for the implied value function until the value function converges. Then, I have the value function and optimal pricing decision for firms across all states. I must aggregate the economy by finding the stationary equilibrium of firms across the states, however, to find the wage implied by the value function and decision rules. To find this distribution, I employ a technique similar to solving for the steady state of a Markov chain, which is described in more detail in Appendix B. With the distribution of firms across states in hand, I aggregate the firms' individual prices to compute the aggregate price level using the CES aggregators.

The aggregate normalized price level and aggregate consumption are inversely related as implied by (2), and the consumer's labor-leisure condition then provides the wage implied by firms' optimization. The implied wage informs of which bound in the bisection algorithm should be changed, and the algorithm continues with new guesses for the steady state wage until convergence. Once the wage converges, I have found the steady state.

Adding seasonality to the problem means that I will need to find the cyclostationary equilibrium and the cyclostationary distribution of firms across prices. To find the cyclostationary equilibrium of the model, I use an algorithm similar to the one above, except I solve for the cyclostationary normalized wages across seasons in a cycle until the aggregate variables in the model converge to some tolerable level across iterations. The solution to the above algorithm for each value of γ_s is a good initial guess for initial values needed for the algorithm. Instead of using the current season's value function as the approximate for the next period's value function (as in the steady state), I use the next period's value function in the cycle. When solving for the stationary distribution of firms across prices in the current period, I use the stationary distribution of firms across states from the other periods to pin down the distribution of firms across states after a complete seasonal cycle (see Appendix B). The information necessary to solve the current period problem from the other periods' problems is made available to the algorithm in the current cycle. Otherwise, the logic of the algorithm, essentially to find the market clearing wages, remains the same.

Note that I need to discretize the grid of prices and shocks on which I solve for the value function and the distribution of firms across states. For computational speed, I start with a large number of grid points for each state and decrease the number of grid points until the aggregate and micro moments of the data are affected significantly. In practice, I use 60 log-spaced grid points for prices (more than what is probably necessary, but extra grid points in the price dimension affect the computational speed of the model much less than extra grid points for the shocks) and 15 grid points for the productivity shock.

B Computing the Cyclostationary Distribution

This appendix provides the computational algorithm for the stationary or cyclostationary distribution of firms across states. The spirit of the solution method follows Tauchen (1986) in that I need to find the transition matrix between states in the model across periods. The following method finds the stationary distribution across firms of non-iid states $\overline{\omega_t}$ in the case without seasonality:

Step 1: Sort the states into the iid shocks and the non-iid state variables (including persistent shocks). The states in ω_t are only the non-iid (discretized) states, and the overall distribution of firms across states is the joint distribution of the distribution of firms across the non-iid states and the independent distribution of firms across the iid shocks. If there are multiple non-iid state variables, the variables in ω_t are the (vectorized) Cartesian power of the states.

Step 2: Initialize a mass 1 of firms at every state in ω_t and then find the implied mass of firms across all states, including the iid shocks. The mass of firms at any state will then be equal to the mass of firms at the state implied by the joint distribution of the iid shocks. If there are no iid shocks, then the mass of firms at all states is simply 1.

Step 3: Apply the decision rules for the firms across all states once and track the mass of firms from each non-iid state in the present period to each non-iid state after the effect of inflation in the next period. Tracking these firms gives the elements of the transition matrix A in the relationship $\omega_{t+1} = A\omega_t$. The element $A_{i,j}$ is the mass of firms that began at time t at state j and ended up at the beginning of time t+1 at state i.

Step 4: Create an initial distribution of firms across the states ω^0.

Step 5: Note that $\omega_{t+1} = \omega_t$ in equilibrium, and iterate $\omega^{i+1} = A\omega^i$ until convergence.

The rationale for this approach follows from the solution method of the menu cost model and the problem of finding the stationary distribution of a Markov chain. For computational purposes in the menu cost model, one must discretize the states. It would be impossible to pin down the firms' optimal price choices numerically given an infinite number of states. Hence, the distributions of the firms across the states must be discretized as well. The stationary distribution of continuous random variables is approximated by the stationary distribution of discrete random variables. If the states have the Markov property–that given the present state, future and past states are independent–then the problem of finding the stationary distribution simplifies to the problem of finding the transition matrix between states and then solving for the steady state of the Markov chain.

The first step to find is to find the transition matrix between the firms' state now and their state in the next period. In a textbook Markov chain problem, the transition matrix is known, but in a computational setting, one must find the transition matrix implied by the firms' decision rules at different states.

To solve for the cyclostationary distribution in the case with seasonality, one would need the transition matrices $A_1, ..., A_s$ where s is the number of seasons in a year and where $\omega_{t+q} = A_q\omega_{t+q-1}$ for $1 \leq q \leq s$. The cyclostationary distribution satisfies the properties $\omega_{t+s} = A^*\omega_t$ and $\omega_{t+s} = \omega_t$ for all t. Substituting in the relationship $\omega_{t+q} = A_q\omega_{t+q-1}$ simplifies the problem of pinning down A^*:

$$\omega_{t+s} = A_s\omega_{t+s-1}$$

$$\omega_{t+s} = A_sA_{s-1}\omega_{t+s-2}$$

$$\omega_{t+s} = A_sA_{s-1}...A_1\omega_t$$

Thus, the transition matrix A^* for a given period t is a function of the period by period transition matrices:

$$A^* = A_s A_{s-1} ... A_1$$

Solving for the cyclostationary distribution then uses the same methods as solving for the simple stationary distribution case given the transition matrix (from step 4 on).

Step 1: Sort the states into the iid shocks and the non-iid state variables (including persistent shocks). The states in ω_t are only the non-iid (discretized) states, and the overall distribution of firms across states is the joint distribution of the distribution of firms across the non-iid states and the independent distribution of firms across the iid shocks. If there are multiple non-iid state variables, the variables in ω_t are the (vectorized) Cartesian product of states.

Step 2: Initialize a mass 1 of firms at every state in ω_t and then find the implied mass of firms across all states, including the iid shocks. The mass of firms at any state will then be equal to the mass of firms at the state implied by the joint distribution of the iid shocks. If there are no iid shocks, then the mass of firms at all states is simply 1.

Step 3: Apply the decision rules for the firms across all states once and track the mass of firms from each non-iid state in the present period to each non-iid state after the effect of inflation in the next period. Tracking these firms gives the elements of the transition matrix A_1 in the relationship $\omega_{t+1} = A_1 \omega_t$. The element $A_{1(i,j)}$ is the mass of firms that began at time t at state j and ended up at the beginning of time t+1 at state i.

Step 4: Repeat steps 1-3 to find $A_2, ..., A_s$. Then compute $A^* = A_s A_{s-1} ... A_1$.

Step 5: Create an initial distribution of firms across the states ω^0.

Step 6: Note that $\omega_{t+s} = \omega_t$ in equilibrium, and iterate $\omega^{i+1} = A^* \omega^i$ until convergence.