Business Cycles and Endogenous Uncertainty*

Rüdiger Bachmann
U. of Michigan and NBER

Giuseppe Moscarini
Yale U. and NBER

December 2010 (in progress)

Abstract

Recessions are times of increased uncertainty and volatility at the micro level. This widely documented empirical pattern has been interpreted as the effect of uncertainty shocks, mediated by various frictions, on aggregate economic activity. We explore the hypothesis that the causation runs the opposite way: first moment shocks induce risky behavior, which in turn raises observed cross-sectional dispersion and time series volatility of individual economic outcomes. Specifically, we study the cyclical pattern of the dispersion of price changes, and the resulting changes in sales and employment at the firm-level. We formulate an imperfect information version of the standard monopolistically competitive model. The elasticity of demand differs across products, and firms are not sure about the elasticity of the demand they face, but learn it from their volume of sales. Due to a fixed operation cost, information is valuable to decide whether to exit the market. Idiosyncratic demand shocks impair learning. The model is fully microfounded and can be aggregated to study general equilibrium. Deviations from average prices are costly to the firm in terms of forgone profits, but the response of sale volumes is informative about market power. Bad economic times are the best times to price-experiment, as the opportunity cost of price mistakes is lower and exit looms large. Following a negative aggregate shock to nominal spending, firms at first keep their prices relatively high, in order to learn whether the demand for their product is sufficiently inelastic and resilient to survive, and then change them further by an amount that depends on what they have learned.

*We thank Veronica Guerrieri for a very helpful discussion and Eduardo Engel, Mikhail Golosov, and Joseph Vavra for comments. The usual disclaimer applies.

1 Introduction

Bad economic times are also times of increased uncertainty. This common refrain resonated again throughout the recent Great Recession, and finds substantial support in empirical evidence from previous recessions. The finance literature has long recognized that markets are much more volatile when average asset prices suffer, typically in bad economic times. In a recent influential paper, Bloom (2009) shows that the VXO index of stock market volatility covaries positively with cross-sectional dispersion in stock returns and negatively with aggregate economic activity. Bloom, Floetotto and Jaimovich (2009) document that the cross-sectional dispersion of shipments across manufacturing establishments, of growth rates of sales across listed companies and across all industries are strongly countercyclical.

The predominant view in the economics literature is that causality runs from uncertainty to aggregate economic activity. Indeed, Bloom finds that the VXO index spikes in recessions, and its innovations in a VAR identified by Choleski-ordering can account for movements in industrial production. The precise mechanism through which uncertainty shocks cause recessions relies on the type of friction that impedes continuos microeconomic adjustment. The sectoral reallocation hypothesis, going back to Lilien (1982), emphasizes frictions in the reallocation of employment across sectors. An increased dispersion in the size of sector-specific shocks causes temporary unemployment, as workers slowly retrain and reallocate to expanding industries. Brainard and Cutler (1993) measure sectoral shocks by means of excess returns on industry stock indexes, after controlling for industry-specific betas, and indeed find evidence that the dispersion of excess returns moves over time and strongly correlates with unemployment. Its explanatory power for unemployment, however, is quantitatively weak. Davis, Haltiwanger and Schuh (1996, Figure 5.5) find that the dispersion of employment growth rates across manufacturing establishments is significantly larger in 1982 (recession) than in 1978 (expansion).

The real option literature suggests that increased uncertainty discourages and delays irreversible actions, such as investing in new capital, hiring new workers, or changing prices in the presence of non convex adjustment costs. This prediction is not a foregone outcome. Larger shocks increase inertia by moving optimal adjustment boundaries further apart, to take advantage of the larger option value of waiting, but also push more agents against adjustment boundaries, so the net effect is ambiguous. Bloom (2009) in partial equilibrium, and Bloom, Floetotto and Jaimovich (2009) in general equilibrium, investigate this mechanism, allowing for fixed and variable adjustment costs to labor and capital. A calibration of a standard S,s investment model, with garden variety aggregate shocks in mean productivity and less traditional innovations in the volatility of firm-specific shocks, shows that the latter can independently cause a drop in aggregate investment and hiring, and even a recession of plausible magnitude, followed by a sharp rebound as uncertainty is resolved. One can think of the stock market crash on 1987 or the Asian crisis of 1997 as examples of increased uncertainty that hit investment but not hard enough to cause a recession.

Bloom (2009)'s article spurred significant interest in this topic. Bachmann and Bayer (2009a, 2009b) use a long panel of German firms and show that the dispersion in firm-level innovations in TFP, sales and real value added is countercyclical, although the dispersion in investment rates is procyclical. Unlike much of the previous literature, they rely on direct measures of performance for a sample of private and public companies spanning all sectors. While increased volatility in stock prices may be due to enhanced alertness and faster response to news by stock traders, this direct evidence indeed points to more news in firms' fundamentals during recessions. Using this evidence to calibrate a neoclassical economy with non convex adjustment costs, Bachmann and Bayer conclude that shocks to the size of firm-level TFP innovations, if any, mildly amplify first-moment aggregate shocks, but are not an independent source of fluctuations. Bachmann, Elstner and Sims (2010) compare measures of individual business expectations, from surveys of decision makers inside the firms, with ex post business performance, to extract a measure of genuine uncertainty. They replace stock market volatility with mean square error in business forecasts, and repeat Bloom (2009)'s VAR exercise. They find that uncertainty shocks have a negative, but very persistent, effect on industrial output. The real option approach predicts a sharp but short-lived effect. Changing the identification strategy of structural uncertainty shocks by requiring that they have no long-run effects on aggregate economic activity, Bachmann et al (2010) find that they have no discernible impact even in the short run. They actually observe the opposite: various measures of uncertainty significantly increase in response to an identified long-run negative shock to aggregate economic activity.

Arellano, Bai and Kehoe (2009) again find that the dispersion in firm-level sales growth in Compustat is countercyclical. They interpret this fact as uncertainty shocks and argue that they generate, through borrowing constraints, business cycles of plausible magnitude. Berger and Vavra (2010) use the micro-data underlying the Consumer Price Index and document that the dispersion in log price changes is strongly countercyclical, due mostly to an increase in the frequency of large price cuts, so skewness also changes with variance. Vavra (2010) shows that countercyclical innovations to the volatility of firm-specific idiosyncratic TFP shocks can account for this fact, but significantly diminish the impact of monetary policy in an otherwise standard calibrated menu cost equilibrium model.

In this paper, we reverse the direction of causality. We explore possible mechanisms through which a decline in mean economic activity, a first moment shock, may induce agents to undertake riskier activities, that result in more dispersed outcomes. Our view is that negative first-moment shocks to a firm's profitability, whether idiosyncratic or aggregate, force the firm to review its *modus operandi* and to change its strategy to survive. If our view is correct, then the root cause of aggregate fluctuations are first moment shocks, and rising uncertainty serves, at best, as an amplification device. Our goal here is not to diminish the importance of uncertainty shocks, but rather to explore the viability of alternative explanations for countercyclical uncertainty. The mechanism that Bloom stressed may still be operative: faced with fixed adjustment costs to capital or employment, firms that willingly undertake risky actions on other margins, such as pricing or supplier relationships, anticipate receiving news, good or bad, so they have an incentive to wait for those news before investing or hiring. The direction of causality, however, matters for economic policy.If agents are creating idiosyncratic risk, the scope for insurance is much diminished.

Observed comovement between uncertainty measures and aggregate economic activity tends to be contemporaneous. Hence, we focus on choice margins that firms can activate quickly in response to first moment shocks. For this reason, in this paper we focus on prices, which can probably be changed faster than production inputs, the nature of projects/business ventures, relationship with suppliers and customers, and other margins.

We explore and model the following hypothesis. Firms have some market power, but are not sure how much of it. They do not know with certainty the price elasticity of the demand that each of them faces. In good times, they choose prices from which they do not want to deviate often, as the scale of demand is high and the opportunity cost of making pricing mistakes is large. So they all choose similar prices, generate small dispersion in sales, and do not learn much about their demand. Depending on the true demand that they face, their sales vary over time, and some firms are driven out of the market. When a firm observes a string of poor sales, it becomes increasingly pessimistic about its own market power, and starts contemplating the exit decision. At that point, the returns to change prices in order to learn the demand curve increase. The firm makes one last desperate attempt in order to understand whether its product sells well and it is worth continuing. When the economy is hit by a negative aggregate spending shock, many more firms find themselves in this situation. As they try to learn whether they are viable, they deviate more from previous pricing behavior, generate more variable and informative sales, and further adjust prices to what they observe. So the dispersion of price changes rises.

The information friction that we focus on, imperfect information about the slope of the demand curve, implies a specific time pattern of price dynamics, which differ markedly from that of time- or state-dependent pricing models. Specifically, a price change, and particularly a large one away from competitors' prices, brings the price into a region of more active experimentation, where more is learned about the demand curve. Therefore, earlier and larger subsequent changes are likely, as more information accrues. The hazard rate of price changes should be declining over price duration. And large price changes away from competitors' prices should be likely to be followed soon by more large price changes. That is, firms should be more likely to change prices a lot when close to, and away from, competitors', and then change again soon. This is exactly what Campbell and Eden (2010) find in scanner data for groceries from five US cities. They conclude: "Taken together, these results suggest to us that sellers extensively experiment with their prices." In contrast, time- and state-dependent price adjustment models imply a constant or increasing hazard rate.

We also offer a modeling contribution. We formulate the (to the best of our knowledge) first incomplete-information version of a standard model of monopolistic competition, where a firm is uncertain about the constant price elasticity of the demand it faces, and sales are only partially informative because of random shocks to the scale of nominal demand for each good. Crucially, prices affect sales depending on the true price elasticity, while spending shocks are independent of prices quoted. Thus, the signal-noise ratio of sales is larger the more a firm deviates from the average prices quoted by competitors. This makes the price a natural tool for experimentation. A fixed operation cost generates endogenous exit. This cost is large enough that firms that are pretty sure to have low market power leave the market.

We hit this model economy with aggregate, first moment shocks to nominal spending. In bad economic times, experimentation is more common for two reasons: the general level of demand is lower, and the likelihood of exit is higher, raising the option-related value of risk to the firm. If the firm makes drastic price changes, it can learn a lot, and then still exit if the news are disappointing. For now, we present partial equilibrium numerical illustrations: we fix price indexes of different types of goods in each aggregate state so that money is not neutral, and verify when the dispersion of price changes indeed comoves inversely with aggregate economic activity. We go, however, to great lengths to formulate our model so that it is fully microfounded, preserves nice aggregation properties, and can thus be solved in general equilibrium, which is the next step on our agenda. In particular, the structure of our economic model has a Kalman filtering flavor and can easily be extended in several directions. A major one that we envision in future drafts is "quality obsolescence": the introduction of new products stochastically increases the price elasticity of demand of pre-existing products, in a way that the firm cannot directly observe and can learn only from declining profits. The microeconomic literature on the "ignorant monopolist", who learns its own demand function through pricing decisions much like in a multi-armed bandit problem, offers several tractable but rather ad hoc models, which posit a demand function. See Trefler (1993) for an example. To the best of our knowledge, none of these models is amenable to aggregate general equilibrium analysis like ours.

One lesson we learn from our exercise is that price experimentation imparts an upward bias in the level of prices. The reason is that a firm can learn by deviating from competitors' prices both up and down, and then observing the reaction of the customers. High and low prices are both informative. But a higher price has the additional benefit of reducing sales, thus production costs, while bolstering revenues through the unit price. This seems a quite general principle that goes beyond the isoelastic utility framework. Since experimentation is more valuable when the firm is at risk of bankruptcy, that is in recessions, the upward bias is countercyclical, prices tend to fall less with nominal spending, further increasing its real effects.

2 The economy

Time is discrete. A large measure of firms produce varieties of a perishable consumption good (or service), using labor only according to a linear production function. The basket of all varieties is the numeraire. Production entails a fixed "operation" cost Ψ per period. Each firm produces one and only variety $j \in [0,1]$, so WLOG the firm can be indexed by j, hiring homogeneous labor on a competitive market at wage w. Firms maximize the present discounted value of flow profits Π discounted at factor $\beta \in (0,1)$.

Firms are owned by a mass Ω of identical, fully diversified households. Each representative household comprises a unit measure of individuals. Each individual has preferences for consumption and work time h expressed by

$$u\left(\{c_j\}\right) - \chi \frac{h^{1+\varepsilon}}{1+\varepsilon}$$

where

$$u\left(\{c_j\}\right) = \frac{1}{2}\log C_1 + \frac{1}{2}\log C_2$$

each C_i, $i = 1, 2$ is a CES aggregator of varieties

$$C_i = \left(\int_{A_i} c_j^{\frac{\sigma_i - 1}{\sigma_i}}\, dj\right)^{\frac{\sigma_i}{\sigma_i - 1}},$$

c_j is consumption of variety j, $\{A_1, A_2\}$ is measurable partition of the unit interval, $0 < \sigma_1 < \sigma_2 < 1$, $\varepsilon > 0$, $\chi > 0$. So there are two sets of varieties, one more generic and easily substitutable, $j \in A_2$ the other one made of "specialties", less easily substitutable, $j \in A_1$.

Members of the household are indexed by the positive real numbers $k \geq 0$, where k is the Pareto weight in the household's preference aggregator. Across members, k has a log-normal distribution with cdf F, mean 1 and variance s_k^2. The household assigns labor supply to each member, pools the resulting income from all members, and redistributes it for spending. Each member of the household of type k receives income Y_k and visits a random but representative sample of firms of each type $i = 1, 2$. The household objective is to maximize the $\beta-$present value of a utilitarian aggregator

$$\begin{aligned} U &= \max_{\{Y_k,h_k\}_{k\geq 0}} \int_0^\infty k \left(\max_{\{c_{j,k}\}_{j\in[0,1]}} u\left(\{c_{j,k}\}\right) - \chi \frac{h_k^{1+\varepsilon}}{1+\varepsilon} \right) dF\left(k\right) \\ \text{s.t. } Y &= \Pi + w \int_0^\infty h_k dF\left(k\right) \\ Y &= \int_0^\infty Y_k dF\left(k\right) \\ Y_k &= \int_0^1 P_j c_{j,k} dj \end{aligned}$$

where P_j is the price of variety j, $c_{j,k}$ is consumption of variety j by consumer k, the first constraint is the budget constraint of each shopper, the second one is the budget constraint of the household, and the third one says that the household income equals profits from all firms plus labor earnings.

The novel element in our model is the information structure. Each firm meets exactly one shopper per period, drawn independently across time and households from F, and quotes a price. The shopper can tell which set A_i each firm belongs to, that is, how easily substitutable in her own preferences each variety j is. But the firm does not know its own identity j, so it does not know whether it is a high- or low-market power firm, nor it observes the type k of the shopper that shows up at the door, so the income that she has available to spend. For example, consumption good is fruit, the units of all varieties are pounds of fruit, firm j produces kiwis, but does not know whether consumers see kiwis as similar to such staples as apples and oranges, difficult to substitute, or as another kind of berries, which are more price sensitive. The firm's prior belief λ that it is of type 1 is the Lebesgue measure of set A_1, namely the proportion of all firms that are in fact of type 1. Different firms, even of the same type, will have different posterior beliefs, depending on their sales history. So they will quote different prices. By the time the consumer draws these prices, she cannot go back and redistribute income to take advantage of the prices she found.

Over time, the firm observes the prices it quotes and the resulting volumes of sales, and learns from them the price elasticity of the demand it faces. The price affects sales through the elasticity of demand, but the arrival of heterogeneous shoppers with unobserved budgets effectively creates a preference shock, which hides the true elasticity. Because spending shocks faced by firms are iid cross-sectionally and over time, there is nothing to learn from observing aggregate variables.

3 Household optimization

First, we solve the problem of household member k who is given income Y_k to spend on varieties. As all goods are perishable, each period the consumer solves a static optimization problem.

$$\max_{\{c_{j,k}\}_{j\in[0,1]}} \frac{1}{2}\log\left(\int_{A_1} c_j^{\frac{\sigma_1-1}{\sigma_1}}dj\right)^{\frac{\sigma_1}{\sigma_1-1}} + \frac{1}{2}\log\left(\int_{A_2} c_j^{\frac{\sigma_2-1}{\sigma_2}}dj\right)^{\frac{\sigma_2}{\sigma_2-1}}$$

$$\text{s.t. } Y_k = \int_0^1 P_j c_{j,k} dj$$

This problem has a well-known solution. The consumer first allocates her income to the two baskets with equal shares

$$\bar{P}_1 C_{1,k} = \bar{P}_2 C_{2,k} = \frac{1}{2}Y_k$$

where the price index of basket i is

$$\bar{P}_i = \left(\int_{A_i} P_j^{1-\sigma_i}dj\right)^{\frac{1}{1-\sigma_i}}$$

and the elasticity of demand is σ_i. Then the consumer allocates each income share $Y_k/2$ to varieties in each basket, according to

$$c_{j,k|j\in A_i} = \frac{Y_k}{2}\frac{P_j^{-\sigma_i}}{\bar{P}_i^{1-\sigma_i}}$$

so that the budget constraint is indeed satisfied.

$$\int_0^1 P_j c_{j,k} dj = \frac{Y_k}{2}\sum_{i=1}^{2}\int_{A_i}\left(\frac{P_j}{\bar{P}_i}\right)^{1-\sigma_i}dj = \frac{Y_k}{2}\sum_{i=1}^{2}\int_{A_i}\frac{P_j^{1-\sigma_i}}{\int_{A_i}P_{j'}^{1-\sigma_i}dj'}dj = \frac{Y_k}{2}2 = Y_k$$

The indirect utility from consumption of consumer of type k who receives income k is

$$\begin{aligned}
u^* &= \frac{1}{2}\sum_{i=1}^{2}\log\left(\int_{A_i} c_{j,k}^{\frac{\sigma_i-1}{\sigma_i}}dj\right)^{\frac{\sigma_i}{\sigma_i-1}} \\
&= \frac{1}{2}\sum_{i=1}^{2}\log\left(\int_{A_i}\left(\frac{Y_k}{2}\frac{P_j^{-\sigma_i}}{\bar{P}_1^{1-\sigma i}}\right)^{\frac{\sigma_i-1}{\sigma_i}}dj\right)^{\frac{\sigma_i}{\sigma_i-1}} \\
&= \log Y_k - \log 2 + \frac{1}{2}\sum_{i=1}^{2}\log\left(\int_{A_i}\left(\frac{P_j^{-\sigma_i}}{\bar{P}_1^{1-\sigma i}}\right)^{\frac{\sigma_i-1}{\sigma_i}}dj\right)^{\frac{\sigma_i}{\sigma_i-1}}
\end{aligned}$$

Since firms cannot observe spending power Y_k of their customers, their price quotes cannot depend on it. Hence, the household member's marginal utility of income is

$$\frac{du^*}{dY_k} = \frac{1}{Y_k}.$$

Equating the marginal utility of income between any two household members k and k', taking Pareto weights into account, we find that household income is allocated optimally in proportion to the Pareto weight k. Since $E[k] = 1$, we obtain

$$k\frac{du^*}{dY_k} = k'\frac{du^*}{dY_k} \Rightarrow \frac{k'}{k} = \frac{Y_{k'}}{Y_k} \Rightarrow Y_k = kY$$

Each member is infinitesimally small and has a zero measure of income, but in the aggregate, the household budget constraint holds:

$$\int_0^\infty Y_k dF(k) = \int_0^\infty kY dF(k) = Y\int_0^\infty k dF(k) = Y.$$

Income is then distributed across members by the log-normal cdf F.

The ideal price index $\bar{P}$ equates indirect utility from consumption to income Y deflated by $\bar{P}$. We first rescale utility again by adding $\log 4$, an innocuous affine transformation, to express the ideal price index $\bar{P}$ in the simplest units:

$$\frac{Y}{\bar{P}} = \log 4C_1C_2 = \log 4\frac{Y}{2\bar{P}_1}\frac{Y}{2\bar{P}_2}$$

so

$$\bar{P} = \sqrt{\bar{P}_1\bar{P}_2}.$$

This is the numeraire, so we normalize $\bar{P} = 1$. Note that each $\bar{P}_i$ is scale-free, because it is units of the good in the basket per units of the good's varieties in basket 1, for example pounds of fruits of all kinds per pound of staple fruit.

Finally, the household chooses labor supply h_k of each member:

$$\chi k h_k^\varepsilon = \nu w = \frac{\partial U}{\partial Y_k}w = k\frac{du^*}{dY_k}w = \frac{w}{Y} = \frac{w}{wH + \Pi}$$

where ν is the Lagrange multiplier on the household's budget constraint. Aggregating across members, this implicitly defines the household's aggregate labor supply:

$$H = \left(\frac{w}{wH + \Pi}\right)^{\frac{1}{\varepsilon}} \int_0^\infty (\chi k)^{-\frac{1}{\varepsilon}} dF(k).$$

The household controls the labor earnings part of its income by choosing labor supply, but takes as given the wage rate w and profits, because it owns only a small share of each firm and of total labor supply.

4 Bayesian updating

The firm has only a belief λ that it is of type 1 (low demand elasticity, high mark-up), and observes the price P_j it quotes and the quantity Q_j it sells at that price to the consumer k who shows up. But the firm knows neither the customer's income (k) nor its own demand elasticity (σ_i.) Let

$$\begin{aligned} q_j &= \log Q_j \\ p_j &= \log P_j \\ \bar{p}_i &= \log \bar{P}_i \\ \mu_i &= (\sigma_i - 1)\,\bar{p}_i - \log 2 \\ y_k &= \log Y_k \end{aligned}$$

Taking logs in the demand function, a firm that is really of type i and meets shopper k (but does not know either i or k) observes physical sales of

$$q_j = -\sigma_i p_j + y_k + \mu_i. \tag{1}$$

Since k is log normal, $Y_k = kY$ is log normal, so $y_k = \log Y_k$ is normal $N\left(m_y, s_y^2\right)$ across people. Using the properties of the log normal distribution and the moments of the underlying distribution of Pareto weights $E\,[k] = 1, V\,[k] = s_k^2$, it can be shown that:

$$\begin{aligned} s_y^2 &= \log\left(1 + s_k^2\right) \\ m_y &= \log Y - \frac{s_y^2}{2} \end{aligned}$$

From now on we drop the variety index j and member identity k for convenience. So the observation equation is

$$q = -\sigma p + y + \mu$$

where the firm knows (log) sales q and price p but not its own type (σ, thus μ) nor the spending power of the consumer (y). Here $y = y_k$ serves as iid Gaussian noise to mask the true elasticity of demand. As the price quote p affects demand through the elasticity but not through income, it affects the speed of learning. There is room for active experimentation.

Now consider the likelihood of these sales in each state (basket of goods):

$$\begin{aligned} \Pr\left(q \leq \tilde{q} | A_i\right) &= \Pr\left(-\sigma p + y + \mu \leq \tilde{q} | A_i\right) \\ &= \Pr\left(-\sigma_i p + y + \mu_i \leq \tilde{q}\right) \\ &= \Phi\left(\tilde{q} - \mu_i + \sigma_i p\right) \end{aligned}$$

where Φ is the Gaussian cdf of $y \sim N\left(m_y, s_y^2\right)$, the consumer's log expenditure.

Using the expression for the Gaussian density, a firm that believes to be of type 1 with probability λ, quotes (log) price p and observes (log) sales q and aggregate income Y, updates its belief to

$$\lambda' = B\left(\lambda, p, q, Y\right) = \left[1 + \frac{1-\lambda}{\lambda} \exp\left\{\frac{\left(q - \mu_1 + \sigma_1 p - m_y\right)^2 - \left(q - \mu_2 + \sigma_2 p - m_y\right)^2}{2s_y^2}\right\}\right]^{-1}. \tag{2}$$

Dynamic decision-making requires to calculate in advance the evolution of beliefs conditional on a given state i, where true sales are $q = \mu_i - \sigma_i p + y$. Replacing this expression into (2), given an idiosyncratic spending shock y, after some algebra, beliefs are the random variable

$$\begin{aligned} b_i\left(\lambda, p, y, Y\right) &= B\left(\lambda, p, \mu_i - \sigma_i p + y, Y\right) \\ &= \left[1 + \frac{1-\lambda}{\lambda} \exp\left\{\frac{z^2 - \left(\Delta\sigma p - \Delta\mu + z\right)^2}{2s_y^2}\right\} (-1)^{\mathbb{I}(i=1)}\right]^{-1} \end{aligned}$$

where $\mathbb{I}$ is the indicator function. As log expenditure y enters only as deviation from its mean m_y, which is public information, we can replace expectations over y with expectations over the simpler variable $z = y - m_y \sim N\left(0, s_y^2\right)$. Also, as long as the dispersion of Pareto weights s_k^2 does not depend on aggregate output Y, neither does the dispersion of income s_y^2, that we can write simply as s^2. This is a deliberate choice, to avoid introducing an exogenous source of cyclical uncertainty, which is the whole point of our analysis. Thus the posterior belief b_i in state i is independent of Y and we can write it as a function of prior beliefs, price quoted, and noise z as follows:

$$b_i\left(\lambda, p, z\right) = \left[1 + \frac{1-\lambda}{\lambda} \exp\left\{\left(-\frac{1}{2}\right)^{\mathbb{I}(i=1)} \left[\left(\frac{z}{s}\right)^2 - \left(p\frac{\Delta\sigma}{s} - \frac{\Delta\mu}{s} + \frac{z}{s}\right)\right]\right\}\right]^{-1} \tag{3}$$

It can be shown that

$$E_z\left[b_2\left(\lambda, p, z\right)\right] < \lambda = \lambda E_z\left[b_1\left(\lambda, p, z\right)\right] + \left(1-\lambda\right) E_z\left[b_2\left(\lambda, p, z\right)\right] < E_z\left[b_1\left(\lambda, p, z\right)\right]$$

The first inequality says that the belief $b_2\left(\lambda, p, z\right)$ that the state is $i = 1$ when the true state is 2 has a negative drift, and the last inequality that the belief $b_1\left(\lambda, p, z\right)$ that the state is $i = 1$ when the true state is 1 (when it is correct) has a positive drift. The middle equality says that, unconditionally on the state, as usual beliefs are a martingale.

The speed of learning, that is how slack are the last two inequalities and how much posterior beliefs react to sale volume, is controlled by the price through the signal/noise ratio $\Delta\sigma/s$: the farther apart the two possible values of the elasticity, relative to the noise in demand, the more informative are sales. Note also that

$$p\frac{\Delta\sigma}{s}-\frac{\Delta\mu}{s}=\frac{(p-\bar{p}_1)(\sigma_1-1)-(p-\bar{p}_2)(\sigma_2-1)}{s}$$

so a very large $|p-\bar{p}_i|$, either positive or negative, provides much information. In order to learn, firms have to price differently than potential competitors in either basket, and the more pronounced the difference in elasticities the more pronounced the difference in sales for any amount of noise z.

A final remark on the firm's inference. If y_k is interpreted as a location-specific income shock, then one would expect large, multi-stores, geographically diversified firms or products to learn faster, as they would be able to average out more noise in the scale of demand. The (potentially testable) prediction is that the volatility of price changes of a product is inversely related to volume, even conditional on age of the product. It is not difficult, though, to identify possible additional sources of noise in demand, specific to the product and correlated across locations. For example, varieties within a basket may be vertically differentiated and the mean of y_k may contain a variety-specific and uncertain quality factor, which does not average out acrss markets.

5 Firm optimization

The firm is owned by households, hence maximizes profits in units of the numeraire consumer basket of all varieties, that all consumers buy in the same proportions. Firm j quotes a price P, which is the number of units of the consumer basket that the buyer has to pay to obtain one unit of variety j (recall that all varieties are types of the same good, so they have the same units, such as pounds of fruit) The consumer compares P with the price of the basket of its immediate substitutes $\bar{P}_i$, which is the number of consumer baskets of all varieties that can be exchanged, by visiting different firms, with one unit of the basket of varieties of type i. Profits equal

$$\frac{Y}{2}\frac{P^{-\sigma_i}(P-w)}{\bar{P}_i^{1-\sigma_i}}=\frac{Y}{2}\left(\frac{P}{\bar{P}_i}\right)^{-\sigma}\left(\frac{P}{\bar{P}_i}-\frac{w}{\bar{P}_i}\right).$$

Here $P/\bar{P}_i$ is the number of units of the basket i that can be exchanged with one unit of the variety by the buyer, and $w/\bar{P}_i$ is the number of units of the basket i that can be exchanged with one unit of the variety by the worker, who produces one unit of variety per unit time worked. The implicit assumption is that the firm physically appropriates the output produced by each worker, who cannot give up his wages, retain the good and sell it on the open market. Due to market power in the product market, the firm can earn a profits on the difference between the units of consumer baskets it charges the buyer per unit of its own variety and the units of consumer baskets the worker charges the firm to produce that unit. The firm trades units of its variety for units of the consumer basket, but the final amount of the latter it returns to its owners as profits depends also on $\bar{P}_i$. Simply put, the firm cannot tell whether the buyer spends many units of the consumer basket because the relevant competing goods $\bar{P}_i$ are expensive or because the buyer has high income (many units of the consumer basket) to spend. Since all goods are in the same units, say pounds of fruit, the ratios $P/\bar{P}_i$ and $w/\bar{P}_i$ are scale-free numbers. Hence, profits are in units of income Y, which is expressed in units of the numeraire consumer basket.

To fix ideas, we first consider two special cases. First, the firm knows its demand elasticity. Second, it learns demand elasticity from sales but is myopic ($\beta = 0$). Finally, we move to dynamics with learning.

5.1 Perfect information

If the firms knows to be in the basket A_i, hence to face demand elasticity σ_i, then it also knows the demand it faces $Y_k P_j^{-\sigma_i}/\bar{P}_1^{1-\sigma_i}$ up to the possibly unknown expenditure shock Y_k. By a standard argument, and using $E[Y_k] = Y$, the firm solves

$$\max_P \left\langle 0, \frac{Y}{2} \frac{P^{-\sigma_i}(P-w)}{\bar{P}_i^{1-\sigma_i}} - \Psi \right\rangle$$

so it chooses a constant markup

$$P_i^* = \frac{\sigma_i}{\sigma_i - 1} w.$$

maximized profits, before deciding whether to exit or not, are then

$$\Pi_i^* = \frac{Y}{2}(\sigma_i)^{-\sigma_i} \left(\frac{w}{\bar{P}_i(\sigma_i - 1)} \right)^{1-\sigma_i} - \Psi.$$

Since $\sigma_1 < \sigma_2$, so the firm enjoys more market power if it is in set A_1, and makes larger profits. To make the problem interesting and learning valuable, from now on we assume

$$\Pi_2^* < 0 < \Pi_1^*$$

for at least some relevant values of aggregate spending Y.

5.2 Imperfect information: static optimization

Now the firm believes to be in set A_1 with probability λ but is fully myopic and does not care about learning and experimentation. The firm then maximizes per period expected profits

$$\max_P \left\langle 0, (P-w) \frac{Y}{2} \left[\lambda \frac{P^{-\sigma_1}}{\bar{P}_1^{1-\sigma_1}} + (1-\lambda) \frac{P^{-\sigma_2}}{\bar{P}_2^{1-\sigma_2}} \right] - \Psi \right\rangle$$

The NFOC is now a more complicated implicit function, which gives an optimal policy $P^{**}(\lambda)$. We know that perfect information prices are ranked $P_1^* > P_2^*$, because $\sigma_1 < \sigma_2$. $P^{**}(\cdot)$ is monotonically (proof to be added) increasing from $P_2^* = P^{**}(0)$ to $P_1^* = P^{**}(1) > P^{**}(0)$: the more optimistic the firm is to own a special variety that consumers really like, the higher the price it charges.

5.3 Imperfect information: dynamic optimization in steady state

The firm now maximizes the expected present value of profits discounted with factor $\beta \in (0,1)$. Then its value is

$$V(\lambda) = \max \langle 0, W(\lambda) \rangle$$

where the first option is exit and the second is continuation:

$$\begin{aligned} W(\lambda) &= \max_P (P - w) \frac{Y}{2} \left[\lambda \frac{P^{-\sigma_1}}{\bar{P}_1^{1-\sigma_1}} + (1-\lambda) \frac{P^{-\sigma_2}}{\bar{P}_2^{1-\sigma_2}} \right] - \Psi \\ &\quad + \beta \int_{-\infty}^{+\infty} \left[\lambda V(b_1(\lambda, \log P, z)) + (1-\lambda) V(b_2(\lambda, \log P, z)) \right] \frac{e^{-\frac{z^2}{2s^2}}}{s\sqrt{2\pi}} dz \end{aligned} \tag{4}$$

To calculate beliefs b_i in each state i, we use the stochastic law of motion (3) derived earlier. The firm maximizes expected flow profits, net of operation costs Ψ, plus the discounted continuation value, which allows for learning from current sales and possibly exit from the market.

Since the firm enjoys more market power if it is in set A_1, we should expect V to be increasing. Then exit occurs at an interior belief $\underline{\lambda} \in (0,1)$ such that value matching holds

$$V(\underline{\lambda}) = 0 \tag{5}$$

and the firm continues as long as the belief remains above $\underline{\lambda}$. At $\lambda = 1$ there is perfect information and the firm marks up over the wage the usual way: P_1^*. At the other extreme, $V(0) = 0 > W(0)$ as the firm exits before being sure to have low market power, due to the operation cost Ψ which makes $\Pi_2^* < 0$.

6 General equilibrium in a stochastic economy

We assume

$$Y = \frac{M}{\hat{P}}$$

where M is the quantity of money in dollars, and $\hat{P}$ is the price in dollars of the consumer basket, so nominal income $\hat{P}Y$ is determined by money supply, and velocity is constant. There are N aggregate states of money M, or nominal demand $\hat{P}Y$, with transition probabilities $\tau_{MM'}$.from M to M'.

Clearly, there is always an equilibrium where money is neutral: if all prices change proportionally to M, then real quantities and beliefs dynamics are unchanged, because they depend only on relative prices. We conjecture that there exists another equilibrium where money is non-neutral (more on this later), so $\hat{P}$ varies less than 1:1 with money and real income Y varies with money. We also assume that each state M is associated to a unique combination of $\hat{P}$ and Y, so we define transitions over values of Y: $\tau_{YY'}$. there is an equilibrium where $\bar{P}_i$ and w take the same values $\bar{P}_{iY}, w_Y$ whenever Y is realized, so that the aggregate state can be taken to be Y WLOG (this will fail in general, and we will use standard Krusell-Smith approximation techniques to find the equilibrium evolution of price indexes.) The firm is owned by the consumer, hence maximizes profits in units of the numeraire consumer basket. The firm's problem becomes

$$\begin{aligned}V(\lambda, Y) = \max\Bigg\langle 0, \max_{P\geq 0} \frac{Y}{2}(P - w_Y)\left[\lambda \frac{P^{-\sigma_1}}{\bar{P}_{1Y}^{1-\sigma_1}} + (1-\lambda)\frac{P^{-\sigma_2}}{\bar{P}_{2Y}^{1-\sigma_2}}\right] - \Psi \\ +\beta \sum_{Y'} \tau_{YY'} \int_{-\infty}^{\infty} \left[\lambda V\left(b_1(\lambda, \log P, z), Y'\right) + (1-\lambda) V\left(b_2(\lambda, \log P, z), Y'\right)\right] \frac{e^{-\frac{z^2}{2s^2}}}{s\sqrt{2\pi}} dz \Bigg\rangle\end{aligned}$$

Finally, aggregate labor supply equals labor demand, which is the same as output:

$$\Omega \int_0^\infty h^k dF(k) = \int_0^1 c_j dj$$

Individual output of each variety depends on prices, which in turn depend on the distribution of beliefs.

7 Price experimentation

7.1 The incentives to experiment

To understand the firm's incentives to experiment in equilibrium, note that the firm trades off the desire to price at the static optimal mark-up P^{**} against the goal of making posterior beliefs move and learning the truth. Standard in Bayesian learning problems, V is convex. Beliefs are a martingale, and more information cannot hurt, as it can always be ignored:

$$\begin{aligned} & V\left(E_{Y'}\left[E_y\left[\lambda b_1(\lambda, p, z, y - m_{Y'}) + (1-\lambda) b_2(\lambda, p, z, y - m_{Y'})\right]\right]\right) \\ = \; & V(\lambda) \\ \leq \; & E_{Y'}\left\{E_y\left[\lambda V\left(b_1(\lambda, \log P, y - m_{Y'}), Y'\right) + (1-\lambda) V\left(b_2(\lambda, \log P, y - m_{Y'}), Y'\right)\right]\right\} \end{aligned}$$

So the firm likes a mean preserving spread in posterior beliefs, that it can achieve by deviating from competitor's prices, as sales are then very informative. The identity of the relevant competitors is unknown, but the firm has beliefs about it.

The gains from experimentation increase with:

- future expected real income Y'; if the Y process is persistent, the gain is higher the higher the current Y, as information will be put to good use in a more favorable environment;
- the discount factor β, as the future is more important;
- the signal/noise ratio $\Delta\sigma/s$, as experimentation yields more information bang for the buck;
- the operation cost Ψ, as information helps the firm to exit and stop paying it when not worth it.

The cost of experimentation is a possible loss in current profits. This loss increases with:

- current average income Y, which scales current profits;
- the level of the elasticities σ_i, as a given price deviation to attain a certain amount of learning has a more dramatic impact on profits the more elastic is demand;
- the operation cost Ψ, as it gets paid for sure today no matter what price one sets, and so is the fixed cost of not leaving the market to trying and learn.

Note that the current state of demand Y affects both the gains and the costs of experimentation. So experimentation is maximized when the expected increase in Y to Y' is maximal: the opportunity cost of poor sales today is low relative to future potential sales. Temporarily bad times are a motive to experiment, even without operation costs Ψ and exit. The role of Ψ is to make experimentation more valuable for firms in bad shape, especially in bad aggregate economic conditions, so that the effect of fundamentals on the incentives to experiment work both for aggregate and idiosyncratic shocks.

As mentioned, the firm can learn by choosing a large deviation of its own price from price indexes $\bar{p}_i$. Given the structure of preferences, under perfect information a firm of type 1 only cares about $\bar{p}_1$ (and vice versa): only $\bar{p}_1$ is relevant to the firm's decisions, any change in the other price index $\bar{p}_2$ will only reduce proportionally the real income spent on basket 2, leaving nominal spending on each basket unchanged. With imperfect information, the firm knows both values of $\bar{p}_1$ and $\bar{p}_2$ in equilibrium, but not which one is relevant to its sales. Therefore, it can learn by quoting a price that is very high or very low relative to *either* price index $\bar{p}_i$, unless $\bar{p}_1$ and $\bar{p}_2$ are very different. To a first approximation, the effect of a drastic price change on learning only depends on its size and not on its sign: note that the leading term in (3) is $(p\Delta\sigma/s)^2$ for $|p|$ large. A high price ($p >> 0$), though, enjoys a unique advantage: it reduces sales, thus production costs. A low price ($p << 0$) must raise revenues much faster than a higher price reduces them, in order to offset the increased costs. Whether this happens or not depends on the demand system. With isoelastic demand, this never happens. So experimentation always takes the form of a price exceeding the myopic optimum under uncertainty $P^{**}(\lambda)$. Imperfect information introduces a positive bias in prices. This upward pressure on prices is stronger the more pronounced experimentation. In our view, and in our numerical examples that follow, experimentation is more common in bad economic times, when nominal and real spending fall, for reasons explained above. We observe a form of endogenous downward nominal price rigidity. Firms keep prices relatively high to see how they do, and then decide in which direction to change them for good.

7.2 Numerical illustration

The model period is one quarter. We choose the following parameter values.

β	σ_1	σ_2	s_k^2
0.98	8	10	0.1

The discount factor $\beta = 0.98$ is justified by an unmodelled risk premium. A factor much closer to 1 leads firms to set a virtually infinite price, as the benefits from experimentation then offset the complete loss of current sales (revenues go to zero at different rates depending on the true σ_i).

For aggregate dynamics we normalize mean output to 1, calibrate to real values and run a partial equilibrium exercise. We assume:

real output Y	real wage w_Y	Prob. Y switches	oper. cost Ψ
0.95	0.998	0.1	0.058
1.05	1.002	0.1	0.058

Then we solve for the optimal pricing policy, and simulate the economy to verify how the dispersion of price changes comoves with the aggregate output.

To justify these business cycle fluctuations through money shocks, we rely on the endogenous downward price rigidity introduced by imperfect information. Because price experimentation is more useful in bad economic times, and because it tends to make prices higher than the static mark-up, due to the cost-saving benefit of a high-price experimentation strategy, nominal prices

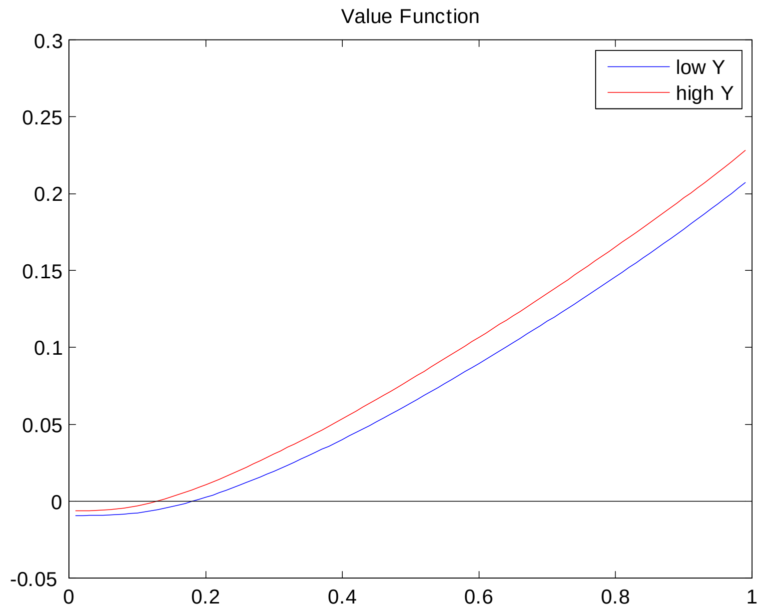

Figure 1: Value function of the firm in aggregate stochastic environment

tend to fall less than money supply. Nominal wages fall faster than prices, so that the real wage declines; absent any productivity shocks, real profits increase; the resulting income effects on labor supply leads the household to work less, rationalizing the fall in the real wage. The map from real wage to output movements is pinned down by an appropriate choice of labor supply disutility χ and elasticity ε.

To solve for the optimal pricing policy, we need to guess the change in the relative price of the two baskets across aggregate states. We report the optimal pricing function in money units (so average prices are not normalized to 1). In those units, our guess is that the overall price level in dollars changes by $\pm 1.8\%$, which is made of a $\pm 2\%$ assumed change in the nominal wage, off a base of 2, minus the additional $\pm 0.2\%$ fluctuation in the real wage to which we calibrate. So the quantity of money fluctuates $\pm 6.8\%$. Due to endogenous price rigidity, the price level implied by the simulations fluctuates by about $\pm 1.6\%$, less than our guess $\pm 1.8\%$. [COMPLETE TO CONVERGE TO GENERAL EQUILIBRIUM]

Figure 1 shows the value function $V(\lambda, Y)$ of the firm's problem as a function of the belief λ that the variety has low elasticity σ_1, when the current state of aggregate economic activity Y takes each of its two possible values. As mentioned, the value is convex, the upper envelope of a smooth increasing and convex function W and the zero outside option when choosing exit, which occurs where the $W = 0$.

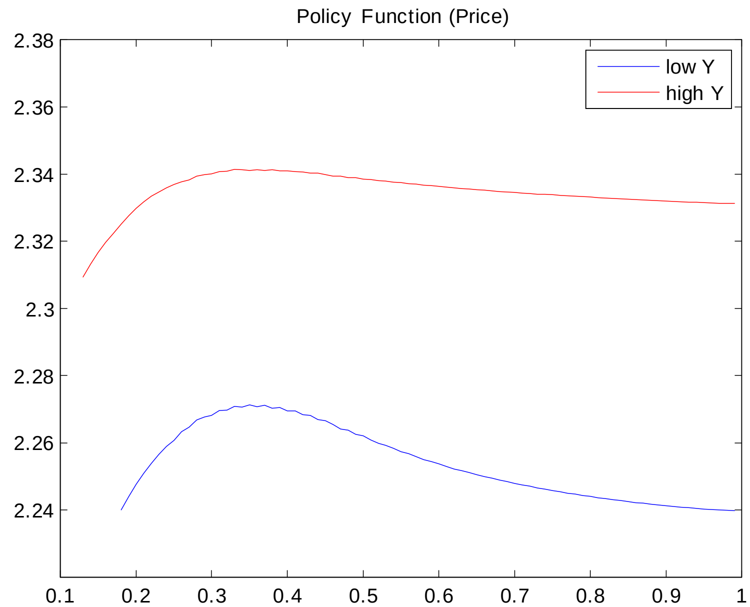

Figure 2: Pricing policy of the firm in aggregate stochastic environment

Figure 2 is the key picture in our exercise and illustrates the optimal policy function (price), again as a function of the belief λ that the variety has low elasticity σ_1, when the current state of aggregate economic activity Y takes each of its two possible values. The myopic pricing function P^* (not shown) is increasing. In contrast, when approaching exit, the firm experiments by pricing high, to read demand elasticity while saving on production costs. This generates the "hump" in the pricing function. Since the operations cost Ψ is inelastic to changes in aggregate demand Y, it is more onerous, and exit is more likely, when Y is low. Correspondingly, the hump is more pronounced in bad times. When Y is high, firms stick to a fairly uniform price, so that learning requires deviating from both, experiment a bit and then cut prices as evidence that the product does not sell well accumulates, and the belief λ falls. When Y is low, experimentation creates the pronounced hump. Since type 1 firms tend to have beliefs closer to 1, and type 2 firms closer to the exit cutoff, the prices of the two baskets $\bar{P}_1$ and $\bar{P}_2$ now differ significantly. Possibly, the inferior varieties of type 2 may have a higher average price, as firms experiment. Thus, it is easier to learn by changing prices around, and the benefits from doing so justify the cost in terms of lost flow profits.

When Y switches, there is a jump in all prices, caused by the discreteness of the aggregate state. The effect, though, is not symmetric, When the economy expands, firms that were closer to the exit cutoff raise their prices by less, as the pressure to experiment diminishes. When the economy slump, those same firms (low beliefs) let their prices fall less, to observe how resilient sale volumes are, then further adjust the price depending on what they observe. Once a low Y is in place, prices at the low end of the belief distribution tend to fall further, as beliefs are a martingale and the pricing function is concave, hence prices are a supermartingale.

In order to simulate this economy, we need to specify the fraction λ_0 of type 1 firm, and to replenish exit. We assume that upon endogenous exit by each firm at $\underline{\lambda}$ a new firm enters the market, and is assigned to basket 1 with chance λ_0. We set $\lambda_0 = 0.5$, so only half of the varieties introduced should survive. We simulate 50,000 firms over 500 model periods. We then calculate in each period the cross-sectional standard deviation of price changes. The choice of numeraire (consumer basket or dollar units of money) does not affect the cross-sectional dispersion, as changes in the dollar price of a basket affects all prices proportionally. Price changes, and particularly price cuts sliding "off the hump" on either side of the peak, are more frequent in bad economic times. We find that this measure of price change dispersion has a correlation of about -0.38 with aggregate economic activity Y. Berger and Vavra (2010) report correlations around -0.4 between the interquartile range of the monthly changes in CPI components and detrended industrial production. In our model, revenues and output are proportional to prices, so a standard deviation in % price changes translates directly into one of sales and output changes.

[TO BE CONTINUED]

8 Extensions

The model is tractable enough to lend itself to plenty of alternative or additional specifications.

It is relatively straightforward to solve the model where demand elasticity σ changes stochastically, unobserved by the firm (other than through noisy sales). We envision the following process. New varieties are introduced each period through some costly R&D process. Some of them "displace" some of the low-elasticity incumbents. For example, the iPhone, when first introduced, may have been considered similar to a Blackberry. Sales (and stock market) then revealed that this is not the case, and the iPhone was quite a unique device. New smartphones introduced later, however, may or may not be close substitutes for the iPhone. Only sales data will tell. Therefore the true elasticity σ has a positive drift, it may rise each period with positive probability, which may depend on the amount of entry. Given the Kalman-filter-like structure of our observation equation (1), this extension is straightforward to implement. It describes well a product life cycle.

The income shock y_k which generates an interesting inference problem for the firm can be rationalized through preference shocks of a representative agent. In particular, we can introduce a scaling factor of c_j in the utility function, which captures vertical differentiation (quality ladders). As long as the firms do not know these factors, which change iid over time across varieties, the optimal pricing problem has the same structure as in this paper.

References

Arellano, C., Y. Bai, P. Kehoe, 2010, "Financial Markets and Fluctuations in Uncertainty", Minneapolis FED.

Bachmann, R. and C. Bayer, 2009a, "Firm-Specific Productivity Risk over the Business Cycle: Facts and Aggregate Implications", mimeo.

Bachmann, R. and C. Bayer, 2009b, "The Cross-section of Firms over the Business Cycle: New Facts and a DSGE Exploration", mimeo.

Bachmann, R., S. Elstner, and E. Sims, 2010, "Uncertainty and Economic Activity: Evidence from Business Survey Data", NBER WP 16143.

Berger, D. and J. Vavra, 2010, "Dynamics of the U.S. Price Distribution", mimeo Yale, https://sites.google.com/site/jvavra/research

Bloom, N., 2009, "The Impact of Uncertainty Shocks", *Econometrica*, May.

Bloom, N., M. Floetotto and N. Jaimovich, 2009, "Really Uncertain Business Cycles", mimeo.

Brainard, L. and D. Cutler, 1993, "Sectoral Shifts and Cyclical Unemployment Revisted", *Quarterly Journal of Economics.*

Campbell, J. and B. Eden, 2010, "Rigid Prices: Evidence from U.S. Scanner Data", July 2010, WP 2005-08

Davis, S., J. Haltiwanger and S. Schuh, 1996, *Job Creation and Destruction*, MIT Press

Lilien, D.M., 1982, "Sectoral shifts and cyclical unemployment", *Journal of Political Economy.*

Trefler, D., 1993, "The Ignorant Monopolist: Optimal Learning with Endogenous Information", *International Economic Review.*

Vavra, J., 2010, "Inflation Dynamics and Time-Varying Uncertainty: New Evidence and an Ss Interpretation", mimeo Yale, https://sites.google.com/site/jvavra/research