Watching each other: inducing rigidity in the midst of change

Primary fields: Economic theory, Monetary economics

Candidate Number: 64257 **Supervisor:** Andrea Prat

Word Count: 7,950

27 May 2008

1 Introduction

One of the central questions in macroeconomics regards the extent and nature of nominal rigidities in relation to prices. New Keynesian models rely upon rigidity in prices and/or wages as a means of inducing movements of variables in calibrated simulations that broadly match those of actual variables over time. In particular, the presence of nominal rigidities in aggregate prices allows for monetary policy to be effective, while if aggregate prices are (perfectly) flexible, monetary policy cannot be effective as changes to costs induced by increases in aggregate demand will be immediately reflected in price changes.

Initial work suggested that most prices change around once per year.[1] However, in their seminal paper, Bils and Klenow (2004) analysed monthly CPI data from the Bureau of Labor Statistics (BLS) and observed that the median duration of prices was 4.3 months, a frequency almost three times higher than previously thought. In response, a debate has ensued in the literature over what represents a price movement that macroeconomists should be concerned with. In particular, the debate has centred on what constitutes a "sale" price as opposed to a "regular" price and how rigid prices are when sales are "backed out" of the data. For example, Nakamura and Steinsson (2007) use a more detailed dataset, also from the BLS, to argue that when sales are excluded, average price duration is between 8 and 11 months. They identify sales according to the flag set by BLS field agents when taking the survey, typically in response to a "sale" sign next to the price being recorded. In contrast, Kehoe and Midrigan (2007) employ an algorithmic approach to identifying sales when looking at weekly supermarket data and observe an average price duration of 4.5 months when sales are excluded and only 3 weeks when they are included.

Eichenbaum, Jaimovich and Rebelo (2008) offer another approach, arguing that firms appear to employ "reference pricing," whereby a reference price is set infrequently (in their data, the figure is about one year), but prices fluctuate around the reference price at a high frequency (roughly once every two weeks). They define a reference price to be the most common price for a good over a period of time (in their analysis they use quarters) and note that non-reference prices do not necessarily correspond to sales, with 25% of non-reference prices in their dataset being above their reference price. Using this framework and employing a dataset comprised of weekly scanner data covering prices, quantities and costs of over 60,000 products in 1,000 stores over the space of three years, they note six regularities in the data as stylised descriptions of price dynamics for the individual firm:

[1] See, for example, Taylor (1999).

1. Prices rarely change unless there is a change in cost.
2. Nevertheless, prices do not always change when costs change.
3. When reference prices do change, they completely pass through all changes in reference costs since the reference price last changed.
4. There are many small changes, both in weekly and in reference prices.
5. Prices are more volatile than marginal costs.
6. The probability of a change in reference prices is significantly associated with a deviation of hypothetical or expected mark-ups from the average markup, suggesting the presence of state dependence in the timing of reference price changes.

A variety of proposals for explaining and modelling price rigidity have been put forward over the last few decades. In a highly influential paper, Calvo (1983) suggested a model wherein all firms faced an exogenous and constant probability of changing their prices in any given period. This idea, which simply imposes price stickiness rather than presenting a microfoundation for it, has become one of the primary work horses of macroeconomic modelling. Proposed explanations for price stickiness include menu cost models, explicit or implicit contracts, the belief that shocks may be temporary, coordination failure and most recently, rational inattention on the part of firms, possibly in response to costly information processing. While each of these ideas are intuitively appealing in their own ways and are variously recognised by firms as reasons for nominal rigidities in surveys (see next section), none of them succeed in capturing all of the stylised facts that we now know about price dynamics.

The model that I present here attempts a micro-founded explanation of firms' pricing decisions, both in timing and in magnitude, on the basis of *watching other firms* in their own pricing behaviour. As a means of explaining price stickiness, it relates most closely to theories of coordination failure and costly information. It is consistent with all of the above stylised facts about price movements for individual firms and with observations of rigidity in aggregate prices. It also represents a framework for the transmission of inflation (and hence, its persistence) across industries or geographies and not simply along production chains.

The remainder of this paper is laid out as follows. Section 2 presents a brief summary of the survey evidence for firms setting prices on the basis of watching other firms. Section 3 gives an overview of models with rational inattention, the basis of the model presented here. Section 4 details the model, section 5 its consequences for aggregate prices and section 6 concludes.

2 Price-setting Surveys

Starting with the work of Blinder (1991) and Blinder *et al.* (1998) in the United States, a variety of surveys have been conducted in an increasingly popular attempt to shed light on precisely how firms set prices. These include work in the UK (Hall, Walsh and Yates, 1997), Sweden (Apel, Friberg and Hallsten, 2005), Japan (Nakagawa, Hattori and Takagawa, 2000), Canada (Amirault, Kwan and Wilkinson, 2006) and nine euro area countries (Fabiani *et al.*, 2005).[2]

The surveys are unanimous in observing that larger firms tend both to review and update their prices more frequently than smaller firms, as do those who feel they are exposed to more intense competition, whether this is from a few other firms that compete aggressively or from a higher number of competitors. The surveys also suggest that both state- and time-dependent pricing rules are used in determining the *timing* of price changes. Amirault, Kwan and Wilkinson (2006) report that 66% of Canadian firms surveyed exclusively use time-dependent pricing rules. This accords with earlier work by Hall, Walsh and Yates (1997) in the UK (79%) and by Blinder *et al.* (1998) in the USA (60%). By contrast, only 34% of firms in eurozone countries reported following purely time-dependent pricing rules, while 20% used state-dependent rules only and 46% of firms used a combination of the two (Fabiani *et al.*, 2005).

When looking at those firms following partially or completely state-based pricing, Canadian firms listed price changes by competitors as the most important cause in triggering an adjustment (Amirault, Kwan and Wilkinson, 2006), as did those in Sweden (Apel, Friberg and Hallsten, 2005). 53% of Spanish firms reported that competitors' price movements were important factors in triggering their own price changes (Alvarez and Hernando, 2005).

In considering the *magnitude* of price changes, 25% of surveyed UK firms reported basing their prices on those of their competitors (Hall, Walsh and Yates, 1997). This figure agreed with the 27% of surveyed eurozone firms reporting the same (Fabiani *et al.*, 2005), although this ranged from 13% in Portugal (Martins, 2005) to 38% in France (Loupias and Ricart, 2004). In the Netherlands, where the

[2] Countries included were: Austria (Kwapil, Baumgartner and Scharler, 2005), Belgium (Aucremanne and Druant, 2005), France (Loupias and Ricart, 2004), Germany (Stahl, 2005), Luxembourg (Lunnemann and Mathä, 2006), The Netherlands (Hoeberichts and Stokman, 2006), Portugal (Martins, 2005) and Spain (Alvarez and Hernando, 2005).

survey was unique in including very small firms among those polled, this figure was 21.6% overall but rose sharply to 34.1% for firms employing only one worker (Hoeberichts and Stokman, 2006).

In investigating the reasons for price stickiness, arguably the primary purpose of the surveys, four theories stand out as being significant: implicit contracts, explicit contracts, cost-based pricing and coordination failure. All of these were among the top five recognised reasons in all 14 surveys when they were included in the options put to surveyed firms. In stark contrast, menu costs and its more recent variant, information costs, were among the least supported ideas, being in the bottom three reasons for most European surveys and Canada (Fabiani *et al.*, 2005). Only in America and Austria were these costs placed in the middle of the group, menu costs being cited as the sixth most proximate cause of price rigidity in the United States and seventh in Austria and information costs coming sixth in Austria (Kwapil, Baumgartner and Scharler, 2005).

Surveys of firms have, for a variety of reasons, occasionally been questioned over their validity as a source of information in building knowledge about firm behaviour. Firstly, survey questions may be poorly worded or may mischaracterise the idea they are trying to present. Secondly, even if survey questions are worded clearly and correctly, they limit the response to a finite set of answers, none of which may be true. Thirdly, firms may (entirely honestly) misrepresent themselves in their responses for reasons analogous to the difference between stated preference and revealed preference in consumer theory. As an example of the first concern, in the Canadian survey the idea of sticky information was put to interviewed firms as follows:

> The information used to review (and ultimately change) prices is available infrequently. Therefore, prices may be slow to adjust to new conditions. (Amirault, Kwan and Wilkinson, 2006 p. 46)

This is a mischaracterisation of the rational inattention models of price stickiness. All versions of these models, both those characterised by Sims (2003) and Reis (2006), suppose that information is always available, but that there are reasons why firms would choose not to avail themselves of it. Nevertheless, the remaining surveys present the costly information ideas reasonably well, so the overall result still needs explaining.

The second and third reasons for being wary of survey results are worth dwelling on for a moment. The model presented in this paper, which takes costly information as its base, produces highly frequent price changes at the level of the individual firm while still inducing nominal rigidities at the aggregate level. It is entirely conceivable that, to the extent that this model is correct, firms may perceive their own pricing decisions as being constrained by coordination failure (one of the most highly ranked reasons for rigidity) and may see the low cost of viewing their competitors' prices as implying that the typical costs of information are low. This is not to suggest that survey evidence should be disregarded. Its use in exploring the validity of past research and offering a guide for future work is invaluable, but it should not be treated as objective data.

3 Sticky information and rational inattention

Starting with the work of Sims (2003), the most recent idea in explaining price rigidity has been that of rational inattention, whereby firms rationally choose to not stay abreast of changes in material trading conditions. While Sims (2003) originally envisaged this as a capacity constraint on firms' information absorption rate,[3] the most widely recognised version of this genre of explanations comes from Greg Mankiw and Ricardo Reis in what they refer to as "sticky information."

In their work, Mankiw and Reis (2002, 2006, 2007) argue that in addition to the costs of physically implementing a change in price (menu costs), there are costs involved in determining the optimal price. Reis (2006) provided a micro-foundation for this work, noting:

> [I]t is costly to acquire, absorb, and process information. It is costly to acquire information in the sense of collecting all the pieces of information that are relevant to assess the current state of the world. It is costly to absorb information in the sense of compiling this information into the relevant sufficient statistics needed to make optimal decisions. And it is costly to process information in the sense of coming up with the optimal action and implementing it. (Reis 2006: p. 795)

In this scenario, the rational producer chooses to update their private information – and hence, their prices – only occasionally and remains inattentive in between adjustments, even when new information is available. In this regard it would therefore appear to observation from the level of the macroeconomy that information itself was "sticky."

Reis (2006) defined the function $G(\mathbf{s}, t) : \mathbb{R}^{S+1} \to \mathbb{R}$, where $\mathbf{s}$ is an $S \times 1$ stochastic Markov process, as "the expected difference between profits earned with full information and profits earned

[3] For examples of models using this idea, see Moscarini (2004) or Mackowiak and Wiederholt (2007).

while following a pre-chosen plan." With $K(\mathbf{s}) : \mathbb{R}^S \rightarrow \mathbb{R}$ as the cost of updating the firm's information and (therefore) optimal plan, he showed that an approximation of the optimal period of inattentiveness when perturbing around $K(\cdot) = 0$ is given by:

$$d^*(\mathbf{s}) = \sqrt{\frac{2K(\mathbf{s})}{G_t(\mathbf{s}, 0)}}$$

Reis identified demand or production volatility and a high elasticity of profits with respect to price or quantity as being possible determinants of inattentiveness through $G_t(\cdot)$. It is worth noting that a further factor will be the size of the firm. $G_t(\cdot)$ clearly increases with firm size, which, in a world of monopolistic competition, can be thought of as the scale of the demand function faced by the firm. So long as $K(\cdot)$ does not increase with firm size, this unambiguously means that larger firms will seek to update their full-information plans more frequently. The intuition is simple enough: the average cost of planning per unit of output is less for larger firms, so they have an incentive to do so more often.

In his later exploration of iso-elastic demand, Reis (2006) assumed instead that planning costs are proportional to firms' profits, which serves to counter this incentive. While this assumption is clearly analytically useful, it is not entirely clear that it is reasonable. Intuitively, full-information planning involves acquiring complete knowledge of the production technology, factor input prices and the shape and scale of demand. There seems little reason to believe that the costs of acquiring information or the amount of information involved (and therefore the costs of processing it) would increase with the size of the firm when investigating the demand or factor input prices it faces. It might conceivably be argued that understanding the productive technology of a firm may prove more difficult as the number of employees increases, but at the same time, it should be noted that larger firms tend to sell multiple goods and services and it seems reasonable to assume at least some economies of scale in investigating product lines simultaneously. Qualitatively, all surveys of price-setting behaviour have also found that larger firms do indeed update their prices more frequently than smaller firms. To the extent that price stickiness is driven by the cost of collecting and processing information, it therefore seems most reasonable to assume a fixed exogenous cost in doing so.

Despite its intuitive appeal, the sticky-information model as presented by Reis (2006), just like its antecedents in standard menu-cost models, counter-factually implies that there should not be many small price changes, nor that prices should be more volatile than marginal costs. Combined with the fact that few firms recognised it as a reason for price stickiness in the various price-setting surveys (Fabiani *et al.*, 2005), this has lead some observers to question the value of sticky-information models.

Two further assumptions of Reis (2006) from the perspective of this paper are that firms set their prices in isolation from each other and that they do so after drawing on accurate sources of information.[4] The first of these implies that price adjustments will be uniformly distributed across geographies and industries and that a constant and exogenously-determined fraction of firms will update their prices in any given period. The second assumption implies that when they do occur, every price adjustment at every firm is always optimal given trading conditions at that instant. These characterisations do not allow for any event-driven updating of information or prices, nor peer-driven updating of the same, nor indeed the possibility of drawing information from one's peers (i.e. with potential bias). They admit no concept of momentum in price adjustments and conclude that price shocks in one sector of the economy spread only by changing the material trading conditions in another.

4 The model

4.1 An informal description

In contrast to Reis (2006), I assume that in each period firms face a choice between three options: continuing with their existing prices, updating their full-information price plan for a significant up-front cost, or adjusting their prices by looking at those of their competitors for a small cost for each firm observed. Note that a rational firm will never perform a full-information update and look at other firms in the same period because the full-information update determines the optimal price by definition; watching other firms would add costs, but gain nothing in return. As will be explained below, this third option will typically lead firms to update their prices more frequently, but to update their full-information price plans less so. It also introduces the idea of a *network* of firms, with links indicating which other firms they each observe. This, in turn, allows us to explore the dynamics of price movements across an economy, looking at how a shock to one firm or industry may lead to price rises in another.

In general, there are two main reasons for watching the price movements of other firms. Firstly, a firm may already know what price it wishes to set, but look to other firms for the timing of that change. This coordination of price movements, presumably arising from concerns about a loss of business when being the first to move, implies that firms producing close substitutes will be more likely to watch each

[4] It should be noted that these assumptions are hardly unique to Reis (2006). They are common to all popular models of price-setting.

other. That is, the magnitude of (positive) cross-price elasticity should be a good indicator of whether two firms watch each other's price movements. Secondly, a firm may look at other firms for guidance on both the timing and the magnitude of a price change. In this scenario, firms will wish to observe those other firms with whom their shocks to production and demand are highly correlated and for whom their price movements are strong indicators of those shocks. It is on this second motive that I focus in this paper.

It may be asked how a firm can be expected to know how closely their shocks covary with those of other firms. It is hypothesised, although not empirically tested here, that the cross-price elasticity of two firms is a good indicator of how closely correlated are their idiosyncratic shocks. This makes intuitive sense on the basis that cross-price elasticity increases with the functional similarity of the two products and decreases with the (geographic) distance between the firms. For a Thai restaurant, the news that the Indian restaurant next door has changed their prices contains more information than identical news about the book store over the road because the two restaurants can be assumed to employ more similar production technologies. Nevertheless, news about price movements in the book store may be more useful than news about another Thai restaurant in an entirely different city because the book store still faces the same factor input prices and shares movements in demand specific to that location.[5]

Viewing the price movement of a close competitor therefore allows a firm to make a more precise estimate of their own shock, which will lower the opportunity cost of not performing a full-information update. However, the precision gained will decrease as the correlation between the two firms' shocks lessens. With a constant or increasing marginal cost of looking at another competitor, each firm will therefore be able to derive a threshold level of cross-price elasticity and only watch competitors above that level. Since the distribution of cross-price elasticities faced by each firm is different, the number of competitors to watch and the threshold cross-price elasticity of those competitors will be specific to each firm. We can then use those thresholds to construct a network linking each firm in the economy.

For example, suppose that there are four firms: two selling bicycles, one selling motorised scooters (like Vespas) and the fourth a gymnasium selling memberships. Scooters and gym memberships may not be considered close substitutes, but scooters and bicycles are substitutes in terms of transportation, while bicycles and gym memberships are substitutes in terms of personal exercise. We might conceivably

[5]For example, the first Thai restaurant and the book store will likely face the same wage demands and (land) rental costs. Likewise, both will experience an increase in demand if the area becomes more popular with high-income consumers.

then have the following cross-price elasticities and thresholds for each firm (own-price elasticities are not shown for simplicity):

	Threshold	(A)	(B)	(C)	(D)
(A) Scooters	0.4		**0.5**	0.2	0.01
(B) Bicycles	0.2	**0.5**		**0.7**	**0.3**
(C) Bicycles in another town	0.3	0.2	**0.7**		0.1
(D) Gym memberships	0.1	0.01	**0.3**	0.05	

From this table, we can then construct the associated graph:

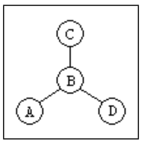

Figure 1: A sample network

We can then see that the speed with which a shock is accounted for across the economy depends critically on which firm first updates their prices and the structure of the network. In the example above, if firms A, C or D were to be the first to update their price in response to an aggregate shock, it may take three periods before all firms had moved, but if firm B were to be the first to move, it may only take two periods.

As will be discussed below, when competitors' prices are watched with a lag - that is, when a firm determining today's price can only see yesterday's prices from their competitors - the number of competitors that each firm chooses to watch depends critically on how long it has been since they last performed a full-information update. In particular, since competitors' prices can only inform a firm about its shocks in the previous period, it is optimal to not watch any other firms in the period immediately following a full-information update. In subsequent periods, the optimal number of firms to watch will increase as the relevance of the last full-information set declines.

4.2 The formal model

4.2.1 The basic setting

I assume a world of monopolistic competition in discrete time, with firms facing iso-elastic demand and constant marginal costs that are subject to persistent, multiplicative shocks. I regard each firm as identical except for the path of shocks it experiences and the correlations between those shocks and those of other firms.

$$
\begin{aligned}
Q_{it} &= e^{b_{it}} P_{it}^{-\gamma} \quad \text{with } \gamma > 1 \\
b_{it} &\sim N\left(\bar{b}, \sigma_b^2\right) \\
C\left(Q_{it}\right) &= c e^{s_{it}} Q_{it} \\
s_{it} &\sim AR(1)
\end{aligned} \tag{1}
$$

In this scenario, the optimal price is independent of shocks to demand:

$$
P_{it}^* = \left(\frac{\gamma}{\gamma - 1}\right) c e^{s_{it}} \tag{2}
$$

For this reason, I henceforth assume for simplicity that there are no shocks to demand ($b_{it} = \bar{b} \ \forall t$), so that all variations in a firm's optimal profits come from shocks to marginal cost. This also means that we have built in Eichenbaum, Jaimovich and Rebelo's (2008) first stylised fact of price dynamics - that prices tend only to change in response to a change in costs - by assumption. Optimal per-period profit is then given by:

$$
\pi_{it}^* = \lambda e^{-(\gamma-1)s_{it}} \qquad \text{with } \lambda \equiv \gamma^{-\gamma} \left(\frac{\gamma - 1}{c}\right)^{\gamma - 1} e^{\bar{b}} \tag{3}
$$

Since I assume that firms have identical marginal costs and own-price elasticities (or at the least that even when different they are universally known) observing a firm's price allows one to directly calculate the value they used for s_{it}. Provided that their shocks are positively correlated, a firm may then estimate their optimal price by observing that of a competitor and making an assumption about whether that other firm has updated their price optimally.[6]

In constructing this estimate, there are two broad approaches that the firm may take. In a naive approach, they assume that whenever a competitor updates their price it is a result of full-information

[6] If competitors' own-price elasticities or marginal costs were not known, firms could still make inferences about their own shocks by looking at the percentage change in other firms' prices. This is because with iso-elastic demand, a firm's optimal price is proportional to $e^{s_{it}}$, so that

$$
\Delta s_{it} \equiv s_{it} - s_{it-1} = \ln\left(P_{it}^*\right) - \ln\left(P_{it-1}^*\right)
$$

This means that by observing a competitor's percentage change in price and assuming that they set their prices optimally, a firm can infer the change in that competitor's shock. With their shocks being positively correlated, this in turn allows the firm to estimate the change in their own shock and, on the basis of an earlier full-information update, the level of their current shock. The analysis in this scenario is considerably more complex and is left for future work.

optimisation, while with more sophisticated updating they take into account the fact that other firms are watching each other too. In what follows, I assume that firms undertake naive updating, meaning that they believe each price update from their competitors contains entirely new information.

Confirming whether firms who set their prices by watching each other do so under a naive or sophisticated approach (or somewhere in between) is clearly a matter for future research. However, I can offer the following as a single datum of anecdotal evidence: Australia's three largest cities are Sydney (population 4.3 million), Melbourne (3.8 million) and Brisbane (1.9 million). By contrast, my home town of Warwick, which is closest to Brisbane, has a population of 12,000 (ABS 2007). I discussed price-setting behaviour with the manager of the sole equipment hire business in Warwick. He stated that when setting prices, he anonymously called a number of hire businesses in Brisbane to obtain their prices and then sets his prices as an average of them, subject to what he referred to as a "sanity check" - in essence, his priors. When asked if he ever took into account the fact that the Brisbane firms might be watching each other or other businesses in Sydney, his response was negative.

Demarzo *et al.* (2003) refer to this failure to properly adjust for the possible repetition of information as "persuasion bias." They observe that the presence of such a bias is consistent with psychological evidence relating repetition of statements and belief in their validity, and argue that it helps explain a variety of phenomena in political discourse, court trials and marketing. The mechanism for belief dynamics that they employ, which has a direct parallel to the price updating in this paper, supposes that agents cannot determine the source of all of the information that they use in updating their beliefs. They suggest that agents maintain a "running tally" summarising their beliefs (in effect, a weighted average of information received), which they update with a rule that would be optimal if all information obtained were entirely new.

It is important to be clear about the timing of price changes and of firms' updating of their internal beliefs. I assume that all firms observe the following process in each period:

1. Decide whether to update prices via full-information optimisation or by observing other firms' prices
2. If setting price on the basis of competitors' prices, observe their prices from the previous period and determine the new price
3. This period's shock is realised
4. If setting price on the basis of full-information, determine the optimal price

5. Update the price and this period's profit is realised

The key point is that firms can only observe their competitors' prices from the *previous* period, meaning that the persistence of firms' shocks plays a key role in determining how much information any observations can contain. In particular, it also implies that it will be optimal to not observe the prices of any competitors in the period immediately following a full-information optimisation (since to do so could only shed light on what is already known), but to start in the period after that.

4.2.2 The error structure

The overall shock faced by a firm in a given period is comprised of an aggregate component common to all firms and an idiosyncratic component specific to that firm. I assume that all shocks follow the same, weakly stationary, AR(1) process and that innovations are normally distributed around a mean of zero.[7]

$$\begin{aligned} s_{it} &= \mu_t + \varepsilon_{it} \\ \mu_t &= \phi\mu_{t-1} + \theta_t \qquad \theta_t \sim N\left(0, \sigma_\theta^2\right) \\ \varepsilon_{it} &= \phi\varepsilon_{it-1} + e_{it} \qquad e_{it} \sim N\left(0, \sigma_e^2\right) \qquad |\phi| < 1 \;\; \forall i, t \end{aligned} \tag{4}$$

I further assume that while aggregate and idiosyncratic shocks are independent of each other, idiosyncratic shocks are correlated across firms within the same period and that this covariance is universally known.

$$\begin{aligned} Cov\left(\theta_t, e_{kt}\right) &= 0 && \forall k, t \\ Cov\left(e_{it}, e_{it+s}\right) &= 0 && \forall t, s \\ Cov\left(e_{it}, e_{jt}\right) &= \sigma_{e_i e_j} \neq 0 && \forall t \end{aligned} \tag{5}$$

This means that the shock to firm j in a given period can be thought of as an imperfect, but unbiased signal about the shock to firm i in the same period.

$$s_{jt} = s_{it} + v_{ijt} \tag{6}$$

[7] Using an identical AR structure for all firms (i.e. the same value of ϕ for both aggregate and all firms' idiosyncratic shocks and a common value for σ_e^2) is clearly quite a strong assumption. It is employed here solely for its tractability.

The error term is defined as $v_{ijt} \equiv \varepsilon_{jt} - \varepsilon_{it}$ and is distributed normally with mean zero and variance given by:

$$Var\left(v_{ijt}\right) \equiv \sigma_{v_{ij}}^2 = \frac{2}{1-\phi^2}\left(\sigma_e^2 - \sigma_{e_i e_j}\right) \tag{7}$$

Note that the error terms for different competitors are *not* independent, their covariances being given by:

$$Cov\left(v_{ijt}, v_{ikt}\right) \equiv \sigma_{v_{ij} v_{ik}} = \frac{1}{1-\phi^2}\left(\sigma_e^2 + \sigma_{e_j e_k} - \sigma_{e_i e_j} - \sigma_{e_i e_k}\right) \tag{8}$$

This then provides the basis for a firm to update its belief regarding its optimal price by observing the prices of competitor firms.[8]

4.2.3 Forming beliefs

Because firms can only observe their competitors' prices from the previous period, beliefs will be formed with a lag. Let $\psi_{it+q}^{n_q}$ be the belief of firm i regarding s_{it+q}, formed in period $t+q+1$, when they last performed a full-information update in period t and they have viewed n_q competitors' prices from period $t+q$. Because of the Gaussian assumption regarding all firms' shocks and the linear relationships between them, $\psi_{it+q}^{n_q}$ will be a normal distribution. Let $m_{it+q}^{n_q}$ and $\sigma_{it+q}^{2^{n_q}}$ represent the expectation and variance respectively of that belief. In a given period, ψ_{it+q}^{0} will therefore describe firm i's prior belief before observing their competitors' prices.

For $q = 1$, when no previous estimates have been made, the prior belief will be based directly on the evolution from the full-information update, so that $\psi_{it+1}^{0} \sim N\left(m_{it+1}^{0}, \sigma_{it+1}^{2^0}\right)$ where $m_{it+1}^{0} = \phi s_{it}$ and $\sigma_{it+1}^{2^0} = \sigma_\theta^2 + \sigma_e^2$.

For $q \geq 2$, the prior belief for each period will be based on the posterior belief of the previous period, so that $\psi_{it+q+1}^{0} \sim N\left(m_{it+q+1}^{0}, \sigma_{it+q+1}^{2^0}\right)$ where $m_{it+q+1}^{0} = \phi m_{it+q}^{n_q}$ and $\sigma_{it+q+1}^{2^0} = \sigma_\theta^2 + \sigma_e^2 + \phi^2 \sigma_{it+q}^{2^{n_q}}$. This is based on the fact that $s_{it+q+1} = \phi s_{it+q} + \theta_{t+q+1} + e_{it+q+1}$.

[8] Note that since the signal errors are not independent across competitor firms, we *cannot* use the commonly used simplification of Bayes' Law for Gaussian-but-independent signals as presented by Chamley (2004), where priors of $N\left(m_{it}, \sigma_i^2\right)$ are updated to posteriors of $N\left(m_{it}', \sigma_i'^2\right)$ by:

$$\begin{aligned}
\sigma_i'^2 &= \frac{\sigma_i^2 \sigma_{s_j}^2}{\sigma_i^2 + \sigma_{s_j}^2} \\
m_{it}' &= \alpha s_{jt} + (1-\alpha) m_{it} \quad \text{with } \alpha \equiv \frac{\sigma_i'^2}{\sigma_{s_j}^2} = \frac{\sigma_i^2}{\sigma_i^2 + \sigma_{s_j}^2}
\end{aligned}$$

As should be intuitively clear, the more closely another firm's shock covaries with that of firm i, the more precision that firm i can expect to gain by watching that other firm's price. Therefore, without loss of generality, I construct an index specific to firm i, placing firm i at zero (0) and arranging the other firms in decreasing order of covariance with firm i, so that $\sigma_{e_i e_{j_i}} \geq \sigma_{e_i e_{j_i+1}}\ \forall j_i$. This in turn implies that $\sigma^2_{v_{ij_i}} \leq \sigma^2_{v_{ij_i+1}}\ \forall j_i$.

Define the vector $\boldsymbol{\beta}_{iq} = \begin{bmatrix} \psi^0_{it+q} & s_{1t+q} & s_{2t+q} & s_{3t+q} & \dots \end{bmatrix}'$ as the information available to firm i regarding period $t+q$ (remember that it will only be available in period $t+q+1$). The variance-covariance matrix of $\boldsymbol{\beta}_{iq}$ is given by:

$$VCov\left(\boldsymbol{\beta}_{iq}\right) = \begin{bmatrix} \sigma^{2^0}_{it+q} & & \boldsymbol{\Sigma}^{0\prime}_{iq} & & \\ & \sigma^2_{v_{i1}} & \sigma_{v_{i1}v_{i2}} & \sigma_{v_{i1}v_{i3}} & \cdots \\ \boldsymbol{\Sigma}^0_{iq} & \sigma_{v_{i1}v_{i2}} & \sigma^2_{v_{i2}} & \sigma_{v_{i2}v_{i3}} & \\ & \sigma_{v_{i1}v_{i3}} & \sigma_{v_{i2}v_{i3}} & \sigma^2_{v_{i3}} & \\ & \vdots & & & \ddots \end{bmatrix} \tag{9}$$

where $\left[\boldsymbol{\Sigma}^0_{iq}\right]_{j_i} = Cov\left(\psi^0_{it+q}, s_{j_i t+q}\right)$, the covariance between the prior belief and the signals of the other firms. This covariance clearly depends on q, the length of time since the last full-information update. For $q=1$, it is given by $\left[\boldsymbol{\Sigma}^0_{i1}\right]_{j_i} = \sigma^2_\theta + \sigma_{e_i e_{j_i}}$,[9] but for later periods when the prior belief is based on the previous period's posterior, the result will be considerably more complex.

Watching the first n_q other firms' prices, firm i's belief is therefore updated to the posterior of $\psi^{n_q}_{it+q} = \psi^0_{it+q} | \left\{s_{j_i t+q}\right\}^{n_q}_{j_i=1} \sim N\left(m^{n_q}_{it+q}, \sigma^{2^{n_q}}_{it+q}\right)$ using the standard formulaes for jointly normal variables. For example, looking at only one other firm's price gives the following:

$$\begin{aligned} m^1_{it+q} &= m^0_{it+q} + \frac{\left[\boldsymbol{\Sigma}^0_{iq}\right]_1}{\sigma^2_{v_{i1}}}\left(s_{1t+q} - s_{it+q}\right) \\ \sigma^{2^1}_{it+q} &= \sigma^{2^0}_{it+q} - \frac{\left(\left[\boldsymbol{\Sigma}^0_{iq}\right]_1\right)^2}{\sigma^2_{v_{i1}}} \end{aligned} \tag{10}$$

[9]

$$\begin{aligned} \left[\Sigma^0_{i1}\right]_{j_i} &= Cov\left(\psi^0_{it+1}, s_{j_i t+1}\right) \\ &= Cov\left(\phi s_{it} + \theta_{t+1} + e_{it+1}, \phi s_{j_i t} + \theta_{t+1} + e_{jt+1}\right) \\ &= Cov\left(\theta_{t+1} + e_{it+1}, \theta_{t+1} + e_{jt+1}\right) \\ &= \sigma^2_\theta + \sigma_{e_i e_{j_i}} \end{aligned}$$

Note that the larger is the covariance in idiosyncratic shocks between the competitor and firm i, the larger will be $\left[\mathbf{\Sigma}_{iq}^{0}\right]_{1}$, the smaller will be $\sigma_{v_{i1}}^{2}$ and so the smaller will be the posterior variance. When looking at two or more firms, deriving a closed-form solution rapidly becomes too complex to be done by hand, although in principal it should be possible to obtain a closed-form solution (not just a numerical estimate) by using software such as Matlab. Completing this analysis has been left for future work.

4.2.4 How many firms to watch

When firm i last performed a full-information update in period t, the price they set for period $t+q+1$ will depend on their estimate for s_{it+q+1}, given by the prior belief for that period, $\psi_{it+q+1}^{0}\left(n_q\right)$, with $m_{it+q+1}^{0} = \phi m_{it+q}^{n_q}$ and $\sigma_{it+q+1}^{2^0} = \sigma_\theta^2 + \sigma_e^2 + \phi^2 \sigma_{it+q}^{2^{n_q}}$. If observing the n_q other firms imposes an additional cost $D\left(n_q\right)$, then the expected period profit for period $t+q+1$ will be:

$$
\begin{aligned}
E_{t+q}\widetilde{\pi}_{it+q+1} &= E_{t+q}\lambda\left(\gamma e^{-(\gamma-1)\psi_{it+q+1}^{0}(n_q)} - (\gamma-1)\, e^{s_{it+q+1}-\gamma\psi_{it+q+1}^{0}(n_q)}\right) - D\left(n_q\right) \qquad (11)\\
&= \lambda\left(\begin{array}{c}\gamma\exp\left[\begin{array}{c}(1-\gamma)\,\phi m_{it+q}^{n_q}\\ +\frac{1}{2}(\gamma-1)^2\left(\sigma_\theta^2+\sigma_e^2+\phi^2\sigma_{it+q}^{2^{n_q}}\right)\end{array}\right]\\ -(\gamma-1)\exp\left[\begin{array}{c}\phi^{q+1}s_{it}-\gamma\phi m_{it+q}^{n_q}\\ +\frac{1}{2}\left(\sigma_\theta^2+\sigma_e^2\right)\sum\limits_{s=0}^{q}\phi^{2s}\\ +\frac{1}{2}\gamma^2\left(\sigma_\theta^2+\sigma_e^2+\phi^2\sigma_{it+q}^{2^{n_q}}\right)\\ -\gamma Cov\left(s_{it+q+1},\psi_{it+q+1}^{0}\left(n_q\right)\right)\end{array}\right]\end{array}\right) - D\left(n_q\right)
\end{aligned}
$$

Since ψ_{it+q+1}^{0} and s_{it+q+1} are both normally distributed, it is clear that the expected profit excluding observation costs will increase as the variance in these terms falls, so that there is a positive but decreasing marginal benefit to looking at another firm (decreasing because the marginal contribution to precision of each firm will be less than the previous one). With a non-decreasing cost of viewing competitors' prices ($D'\left(n_q\right) \geq 0$), this will implicitly define the optimal number of firms to look at in each period, n_q^*. It is postulated, but not proved in this paper, that the marginal benefit of looking at another firm will be increasing in q, meaning that $n_{q+1}^* \geq n_q^*$.

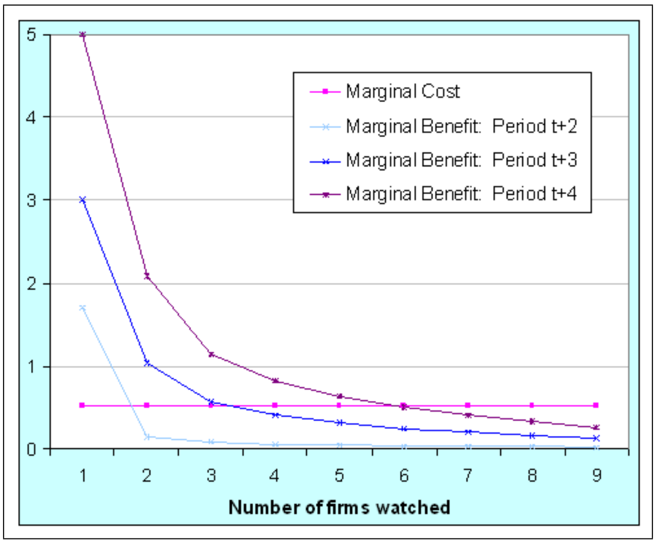

Figure 2: Stylised example: the anticipated marginal benefits from viewing other firms' prices

4.2.5 Frequency of full-information updates

For a firm that last performed a full-information update in period t, let $G(q+1, n_q)$ be the period $t+q+1$ expected difference between full-information profits and those obtained after observing the prices of n_q other firms from the $(t+q)^{th}$ period. Provided that the information obtained from each observed competitor is sufficiently small (e.g. $\sigma_{e_i e_j}$ is small relative to $\sigma_\theta^2 + \sigma_e^2$) so that the marginal benefits are not too high, we would expect $G\left(q+1, n_q^*\right)$ to be increasing in q. That is, that the gap between full-information profits and the highest profit available from watching other firms will grow over time.

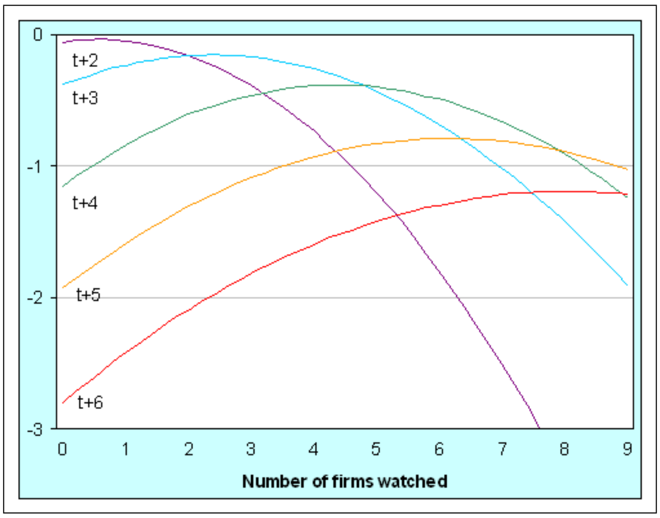

Figure 3: Stylised example: expected profit relative to full-information profit

The firm will therefore perform another full-information update when that gap exceeds the cost of that update. We can then define $\bar{q}$ to be the smallest value of q such that $G\left(\bar{q}+1, n_{\bar{q}}^{*}\right) \geq K(\cdot)$. The number of competitors watched by the firm would thus cycle over time as $\{0, 0, n_1, n_2, \ldots, n_{\bar{q}-1}, 0, 0, n_1, \ldots\}$, where the first zero in each iteration of the sequence is a period in which they perform a full-information update, the second zero is the period immediately afterwards when nothing can be gained from watching other firms and a (weakly) positive number of firms are watched each period thereafter. Full-information updates will be performed every $\bar{q}+1$ periods. This will typically be a longer delay than would occur without the option to view competitors' prices. For example, in the above diagram, if the cost of a full-information update was 1, the firm would update their information set once every four periods if they watched no other firms and only once every six periods if they observed their competitors. Of course, $G(q+1, n_q)$ will be a function of the various parameters of the model, so if these change between full-information updates, $\bar{q}$ may vary from cycle to cycle.

More generally, we can extend the analysis of Reis (2006) to incorporate this model. Let $G(\mathbf{s}, q, n) : \mathbb{R}^{S+1} \rightarrow \mathbb{R}$ be the expected difference between full-information profits and the situation in which the firm watches the prices of n other firms. Reis' formulation is then a nested scenario with $n = 0$, but in the more general case we will have

$$d^{*}(\mathbf{s}, n) = \sqrt{\frac{2K(\mathbf{s})}{G_t(\mathbf{s}, 0, n)}}$$

with the proof being directly analogous to that for proposition 4 in Reis (2006). With the firm choosing n to maximise the profit without full-information, it follows that $G_t(\mathbf{s}, 0, n^{*}) \leq G_t(\mathbf{s}, 0, 0)$. In other words, the additional option to watch and emulate competitors' prices will tend to lengthen the delay between full-information updating.

5 Network structure and aggregate effects

The network between firms resulting from this model is a directed graph. We can represent this as a matrix $\mathbf{G}_t = [g_{ijt}]$, where $g_{ijt} \in \{0, 1\}$ and $g_{ijt} = 1$ if firm i looks at the price of firm j when estimating their shock for period t (this will take place in period $t+1$). One way to think of the above results is to argue that the network will be highly dynamic, with links between firms (g_{ijt}) changing every period, albeit in a predictable manner. Perhaps a better characterisation would be to identify the network as being defined by union of the maximal observation groups for each of the firms. Let $S(i, q)$ be the set of firms that firm i considers when estimating their shock for the q^{th} period after they last performed a full-information update. Then we can define the network as $g_{ijt} = \mathbf{1}\left(j \in S\left(i, \bar{q}_i\right)\right)$ where $\mathbf{1}\left(\cdot\right)$ is the indicator function. This makes the network endogenous, but time-invariant, so that $\mathbf{G}_t = \mathbf{G}\ \forall t$.

Because of the assumption of universal Gaussianity, the expectations of firms' beliefs in each period will be weighted averages of their priors and the signals they receive by watching other firms. We can denote these weights by $\mathbf{W}_t = [w_{ijt}]$, where $w_{ijt} = 0$ if $g_{ijt} = 0$, $w_{ijt} \in [0,1]$ when $g_{ij} = 1$ and $\sum_j w_{ijt} = 1\ \forall i, t$. This means that agents' expectations regarding their shocks will evolve as

$$\mathbf{m}_{t+1} = \mathbf{W}_t \mathbf{m}_t$$

with the additional constraint that every $\bar{q}_i + 1$ periods, firm i will take a draw from the truth, so that $m_{it} = s_{it}$. In their analogous model, DeMarzo *et al.* (2003) suppose that $\mathbf{W}_t$ evolves over time according to the process

$$\mathbf{W}_t = \left(1 - \lambda_t\right)\mathbf{I} + \lambda_t \mathbf{W}_{t-1}$$

which involves common and simultaneous movements in weights for all agents over time. In the model presented above, however, $\mathbf{W}_t$ will cycle through a fixed series of values and the changes from one period to the next could be very different for each agent. For example, suppose that we had four firms and all of them had the same frequency of full-information updates of once every four periods (i.e. $\bar{q}_i = 3\ \forall i$), each staggered by one period. In that case, we would have $\{\mathbf{W}_t\} = \left\{\mathbf{W}^{(1)}, \mathbf{W}^{(2)}, \mathbf{W}^{(3)}, \mathbf{W}^{(4)}, \mathbf{W}^{(1)}, \mathbf{W}^{(2)}, \ldots\right\}$.

We can still say something about the evolution of $\mathbf{W}_t$ over this cycle even without closed-form solutions. The weight firm i places on their own priors, w_{iit}, will be zero in a period of full-information updating, will equal one in the subsequent period and will decrease (weakly) monotonically thereafter. The weight placed on other firm's signals will vary depending on the time since the last full-information update, the variance-covariance matrix of idiosyncratic shocks, the cost of watching other firms and the cost of a full-information update. However we would expect that, holding the other parameters constant, the weight should be weakly increasing in the covariance of the two firms' idiosyncratic shocks: $\frac{\partial w_{ijt}}{\partial \sigma_{e_i e_j}} \geq 0$.

A key result of this is that when the network is highly connected - i.e. when firms watch many other firms - the aggregate price level can display a remarkable degree of rigidity, lasting long after all firms have performed a full-information optimisation. To illustrate this, consider the following stylised example. Suppose that there are four firms with the following matrix of correlations in the shocks to their errors ($\rho_{ij} = \frac{Cov(\theta_t + e_{it}, \theta_t + e_{jt})}{Var(\theta_t + e_{it})} = \frac{\sigma_\theta^2 + \sigma_{e_i e_j}}{\sigma_\theta^2 + \sigma_e^2}$):

	A	B	C	D
A	1.0	0.7	0.6	0.5
B		1.0	0.8	0.6
C			1.0	0.5
D				1.0

With a complete solution, this table of primitives, together with the cost of watching other firms or performing a full-information update, would be sufficient to identify the results. For exposition purposes here, however, suppose further that all four firms choose to operate on four-period cycles, staggered by one period, and all choose to watch all of their competitors when it is efficient to watch any. This means that the number of other firms that are watched in each period will obey the following cycle, with firm A performing a full-information update in period 1, firm B in period 2 and so on:

	(1)	(2)	(3)	(4)
A	0	0	3	3
B	3	0	0	3
C	3	3	0	0
D	0	3	3	0

The following period-specific weights would then be consistent with the characterisation presented above:

Period 1

	A	B	C	D
A	0	0	0	0
B	0.22	0.33	0.26	0.19
C	0.10	0.14	0.67	0.09
D	0	0	0	1.00

Period 2

	A	B	C	D
A	1.00	0	0	0
B	0	0	0	0
C	0.21	0.28	0.33	0.18
D	0.10	0.12	0.10	0.67

Period 3

	A	B	C	D
A	0.67	0.13	0.11	0.09
B	0	1.00	0	0
C	0	0	0	0
D	0.21	0.25	0.21	0.33

Period 4

	A	B	C	D
A	0.33	0.26	0.22	0.19
B	0.11	0.67	0.13	0.09
C	0	0	1.00	0
D	0	0	0	0

In what follows, I treat aggregate prices as an unweighted, linear average of prices, which are calculated as $P_{it} = e^{m_{it}}$ for each firm (i.e. I assume that $\lambda = 1$).

I first consider the effect of an aggregate shock, common to all firms, so that $\theta_0 = 1$, $\theta_t = 0\ \forall t \geq 1$ and $e_{is} = 0\ \forall i, s$. Despite the fact that all four firms perform full-information optimisations within the first four periods, an aggregate shock leads to a delay of up to 11 periods before the aggregate price level reaches its optimal level.[10] The time taken for the aggregate price level to reach its optimal path increases in the persistence of the shock, ϕ.

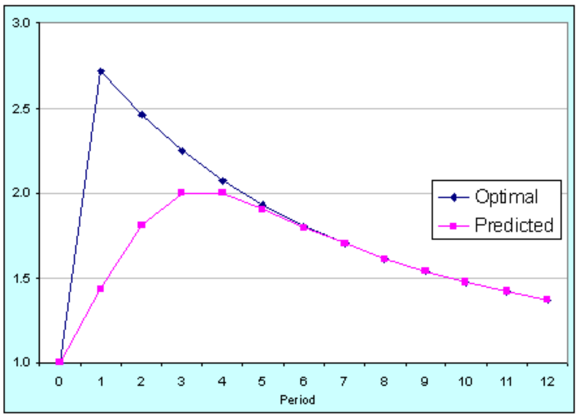

Figure 4a: $\phi = 0.9$.

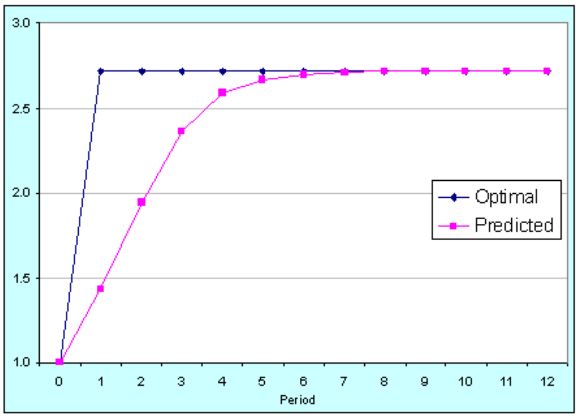

Figure 4b: $\phi = 1.0$.

Figure 4: Aggregate price level in response to an aggregate shock to marginal costs

I next consider the response to a non-uniform shock. Specifically, I assume that $\theta_t = 0\ \forall t$ and set $e_{A0} = 1$, $e_{B0} = 0.7$, $e_{C0} = 0.6$ and $e_{D0} = 0.5$ (in line with the correlation of idiosyncratic shocks with firm A), with $e_{it} = 0\ \forall i, t \geq 1$. Interestingly, a non-uniform shock causes the aggregate price level to oscillate around its optimal level. This oscillation does not decrease over time relative to the shock, but instead represents the introduction of a permanent dynamic to the aggregate price level. When the shocks are transitory, however, the magnitude of this oscillation does decreases along with the shock. Once again, despite all firms having accessed the "truth" within four periods, it takes more than that before the aggregate price level to crosses the optimum level.

The reason for the excess nominal rigidity is that by watching other firms and being subject to persuasion bias, firms are lead to placing too much weight on the prices of their competitors relative to their own priors. In the case of an aggregate shock, this causes all firms to believe that their prices are too high and lower them below the optimal level until they next perform a full-information update. In

[10] Strictly speaking, the aggregate price level never reaches the optimal level, but the difference between the two decays towards zero over time.

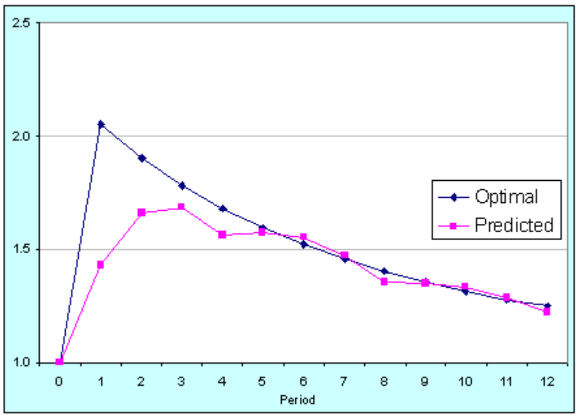

Figure 5a: $\phi = 0.9$.

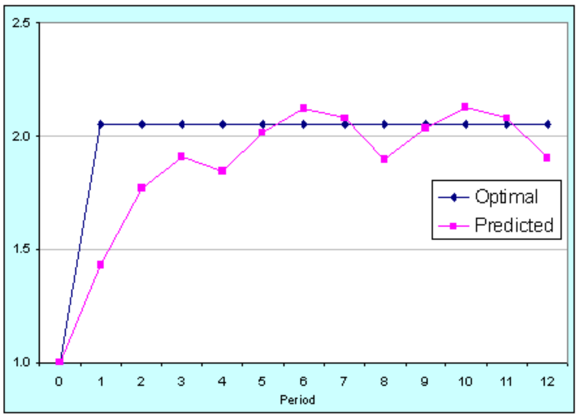

Figure 5b: $\phi = 1.0$.

Figure 5: Aggregate price level in response to a non-uniform shock to marginal costs

the case of a non-uniform shock, firms with the larger shocks will believe that their prices are too high and lower them, while firms who experience the smaller shocks will believe that their prices are too low and raise them.

The fact that firms' cycles are not synchronised is therefore the cause of the oscillations observed in the presence of non-uniform shocks. In this scenario, it is enlightening to view the oscillations of individual firms' prices.

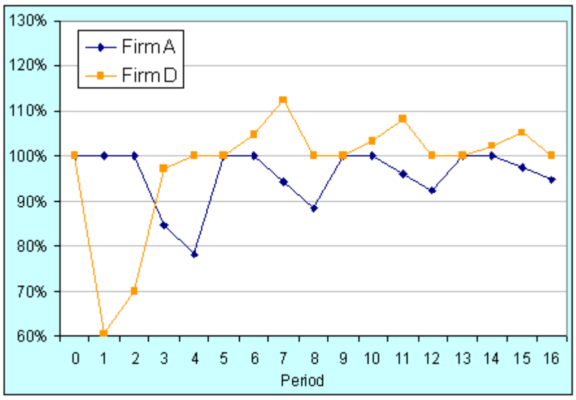

Figure 6a: $e_{A0} = 1$ (high), $e_{D0} = 0.5$ (low) and $\phi = 0.9$.

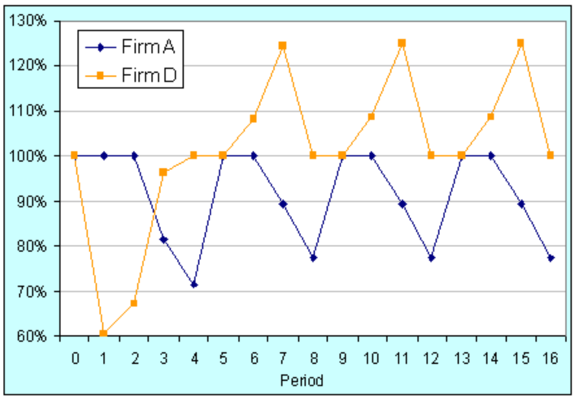

Figure 6b: $e_{A0} = 1$ (high), $e_{D0} = 0.5$ (low) and $\phi = 1.0$.

Figure 6: Prices as a fraction of their optimal level

Note that in figure 6a (non-uniform, non-permanent shocks) it appears that the deviations from optimum may degrade over time, but they do stabilise to a steady pattern of oscillation analogous to those when the shocks are permanent (figure 6b). When the shock is an aggregate (i.e. uniform) one, all firms behave as firm A does in the above figures, but the percentage deviations from the optimum price *do* degrade to zero over time. An objection might be raised that a firm that repeatedly updated to its optimal price and later moved away from it because its competitors hadn't yet moved would eventually realise that to do so was suboptimal, but this misses the point of the persuasion bias. The firm does not move away from its optimal price because its competitors have not yet moved, but because it (incorrectly) believes that its underlying costs have changed.

It is worth noting the striking visual similarity between these stylised graphs of firms' price movements and those of Eichenbaum, Jaimovich and Rebelo (2008) in their observation of reference prices in the data.

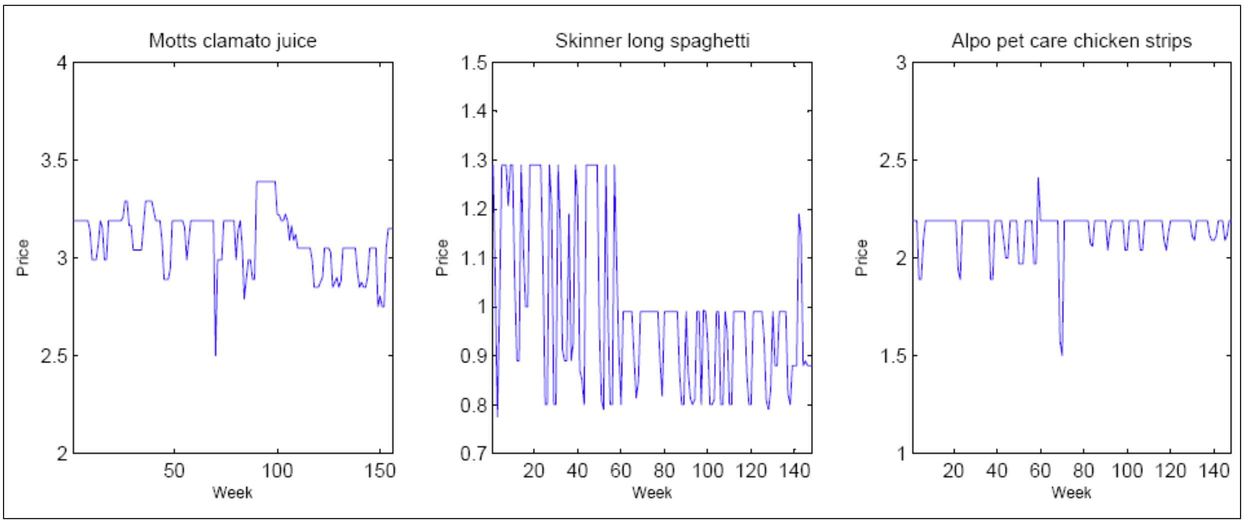

Figure 7: Time-series observations on the prices of three goods in the Eichenbaum *et al.* dataset (Eichenbaum, Jaimovich and Rebelo, 2008: Figure 1).

This suggests that some of the deviations from reference prices observed by Eichenbaum *et al.* may be the result of firms suboptimally adjusting their prices in response to the prices of their competitors. The stylised movements of firms' prices here also qualitatively match the six facts observed by Eichenbaum, Jaimovich and Rebelo (2008), particularly in the case of permanent, non-uniform shocks to firms' costs.[11]

It should be stressed that the above example rested on the strongly connected nature of the network. If the four firms competed in information isolation, prices would have wholly and correctly adjusted

[11] Of course, the first fact - that prices rarely change unless there is (believed to be) a change in cost - was expressly built in by assumption.

to any shock within four periods. This highlights the importance of network topology in exploring the dynamics of aggregate prices. For example, the presence of clustering, where linked nodes tend to have common neighbours, "fat tails" in distributions of node degrees, the number of links to or from each node, and assortative linking, where high-degree nodes tend to link to other high-degree nodes and low-degree nodes to other low-degree nodes - three attributes common to many social networks - would here lead to significant differences in responses to shocks across the network. In particular, one might expect that areas of densely populated, highly connected firms - cities, in other words - would exhibit greater nominal rigidity in aggregate prices while simultaneously invoking more frequent and smaller price changes at the level of the individual firm.

6 Conclusion

This paper proposes a model of firms' pricing decisions based partially on a cost of determining optimal prices in the style of Reis (2006) and partially on firms watching the prices of their competitors. This latter component is motivated by the observation that when surveyed, a large fraction of firms across North America and Europe admit to looking to other firms in deciding both the timing and the magnitude of price changes. When the marginal cost of looking at a competitor is low and the cost of a full-information optimisation is high, firms will rationally choose to perform a full-information update less frequently than in standard sticky information models, but to nevertheless change their prices more frequently as a result of mimicking the movement of their competitors.

The presence of persuasion bias, in which firms believe that their competitors' price changes are always the result of full-information optimisation, induces firms to value their prior beliefs regarding their optimal price too little and the prices of their competitors too much. This leads firms to deviate from their profit-maximising prices when it has been a long time since they last performed a full-information update. In the event of common shocks to firms' costs, this results in nominal rigidities in aggregate prices despite individual firms' prices changing almost every period. In the event of non-uniform shocks to firms' costs, in addition to nominal rigidities in aggregate prices it also introduces indefinite oscillations to both individual and aggregate prices. If the non-uniform shocks are temporary, these oscillations decrease in magnitude with the shocks themselves, but if they are permanent, the oscillations remain in the system for ever.

This model, particularly in the event of permanent, non-uniform shocks, appears to be consistent with a variety of stylised facts about price dynamics as presented by Eichenbaum, Jaimovich and Rebelo

(2008). Prices change more frequently than marginal costs, temporarily deviating from their optimal levels by small amounts relative to the size of the initial shock. When a full-information update occurs, prices changes fully take any change in marginal cost into account and the timing of each full-information update is based on the expected difference between profits gained by continuing in the current process and optimal profits (which are proportional to markup). In addition, the model also represents a framework for the transmission of inflation (and hence, its persistence) across industries or geographies and not simply along production chains, thereby opening a window for the further study of the effects of monetary policy.

There are nevertheless a variety of shortcomings in the model as presented. The most obvious of these is the absence of a closed-form solution to the firm's problem of how many competitors to watch in each period. Solving the problem when firms do not know each other's marginal costs (and so have to rely on percentage changes in prices) and when they are not wholly subject to persuasion bias are also promising areas for future work. Further study on the dynamics of prices across the network is also clearly called for. Possible avenues in this area include providing a measure of the social welfare impact of the persuasion bias (in the style of Calvo-Armengol and de Marti (2007) or DeMarzo, Vayanos and Zwiebel (2003)) and characterising the effects of various network attributes such as assortative linking and network diameter on aggregate price levels. The work of Galeotti and Goyal (2007) may be of benefit in this regard. Finally, a study of aggregation - whether clusters of interlinked firms can be thought of as a single firm for the purposes of analysing the dynamics of aggregate prices - would represent the groundwork for the model to be brought into the realm of Dynamic Stochastic General Equilibrium (DSGE) models and other forms of macroeconomic modelling.

Empirically, the work as it currently stands clearly calls for verification of the presence of persuasion bias on the part of price setters and for analysis on whether cross-price elasticity is a good predictor of the correlation between two firms' shocks. If these two assumptions are confirmed, a general data-gathering exercise to obtain cross-price elasticities for a wide sample of firms would be necessary to construct a reasonable model of a localised economy before testing its ability to predict aggregate price changes across that economy.

7 Bibliography

ABS (2007). Regional Population Growth, Australia, 2006-07, *Australian Bureau of Statistics*.

Akerlof, G. (2007). "The Missing Motivation in Macroeconomics." *American Economic Review* **97**(1): pp. 5-36.

Allcott, H., D. Karlan, *et al.* (2007). "Community Size and Network Closure." *American Economic Review* **97**(2): pp. 80-85.

Álvarez, L. and I. Hernando (2005). "The Price Setting Behaviour of Spanish Firms: Evidence from Survey Data." *ECB* Working Paper No. 538.

Amirault, D., C. Kwan, *et al.* (2006). "Survey of Price-Setting Behaviour of Canadian Companies." *Bank of Canada* Working Paper No. 2006-35.

Apel, M., R. Friberg, *et al.* (2005). "Microfoundations of Macroeconomic Price Adjustment: Survey Evidence from Swedish Firms." *Journal of Money, Credit, and Banking* **37**(2): pp. 313-338.

Aucremanne, L. and M. Druant (2005). "Price-setting Behaviour in Belgium: What can be Learned from an Ad Hoc Survey?" *ECB* Working Paper No. 448.

Bils, M. and P. Klenow (2004). "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* **112**(5): pp. 947-985.

Blinder, A. (1991). "Why are Prices Sticky? Preliminary Results from an Interview Survey." *American Economic Review* **81**(2): pp. 89-96.

Bramoullé, Y. and R. Kranton (2007). "Public Goods in Networks." *Journal of Economic Theory* **135**(1): pp. 478-494.

Calvo-Armengol, A. and M. Jackson (2004). "The Effects of Social Networks on Employment and Inequality." *American Economic Review* **94**(3): pp. 426-454.

Calvo-Armengol, A. and J. d. Marti (2007). "Communication Networks: Knowledge and Decisions." *American Economic Review* **97**(2): pp. 86-91.

Calvo, G. (1983). "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* **12**(3): pp. 383-398.

DeMarzo, P., D. Vayanos, *et al.* (2003). "Persuasion Bias, Social Influence and Unidimensional Opinions." *Quarterly Journal of Economics* **118**(3): pp. 909-968.

Eichenbaum, M., N. Jaimovich, *et al.* (2008). "Reference Prices and Nominal Rigidities." *NBER* Working Paper No. 13829.

Fabiani, S., M. Druant, et al. (2005). "The Pricing Behaviour of Firms in the Euro Area: New Survey Evidence." *ECB* Working Paper No. 535.

Fabiani, S., A. Gattulli, et al. (2004). "The Pricing Behaviour of Italian Firms: New Survey Evidence on Price Stickiness." *ECB* Working Paper No. 333.

Fafchamps, M. and F. Gubert (2007). "Risk Sharing and Network Formation." *American Economic Review* 97(2): pp. 75-79.

Galeotti, A. and S. Goyal (2007). "Network multipliers and the optimality of indirect communication", Essex university, *mimeo.*

Hall, S., M. Walsh, et al. (1997). "How do UK companies set prices?" *Bank of England* Working Paper No. 67.

Hoeberichts, M. and A. Stokman (2006). "Price Setting Behaviour in the Netherlands: Results of a Survey." *ECB* Working Paper No. 607.

Jackson, M. and B. Rogers (2007). "Meeting Strangers and Friends of Friends: How Random Are Social Networks?" *American Economic Review* **97**(3): pp. 890-915.

Jackson, M. and A. Wolinsky (1996). "A Strategic Model of Social and Economic Networks." *Journal of Economic Theory* **71**(1): pp. 44-74.

Jackson, M. and L. Yariv (2007). "Diffusion of Behavior and Equilibrium Properties in Network Games." *American Economic Review* **97**(2): pp. 92-98.

Kehoe, P. and V. Midrigan (2007). "Sales and the Real Effects of Monetary Policy." *Federal Reserve Bank of Minneapolis* Working Paper No. 652.

Kwapil, C., J. Baumgartner, et al. (2005). "The Price-Setting Behavior of Austrian Firms: Some Survey Evidence." *ECB* Working Paper No. 464.

Loupias, C. and R. Ricart (2004). "Price Setting in France: New Evidence from Survey Data." *ECB* Working Paper No. 423.

Lunnemann, P. and T. Mathä (2006). "New Survey Evidence on the Pricing Behaviour of Luxembourg Firms." *ECB* Working Paper No. 617.

Mackowiak, B. and M. Wiederholt (2007). "Optimal Sticky Prices under Rational Inattention." Northwestern University, *mimeo.*

Mankiw, G. and R. Reis (2002). "Sticky information versus sticky prices: A proposal to replace the New Keynsian Phillips curve." The *Quarterly Journal of Economics* **117**(4): pp. 1295-1328.

Mankiw, G. and R. Reis (2006). "Pervasive Stickiness." *American Economic Review* **96**(2): pp. 164-169.

Mankiw, G. and R. Reis (2007). "Sticky Information in General Equilibrium." *Journal of the European Economic Association* **5**(2-3): pp. 603-613

Martins, F. (2005). "The Price Setting Behaviour of Portuguese Firms: Evidence from Survey Data." *ECB* Working Paper No. 562.

Moscarini, G. (2004). "Limited Information Capacity as a Source of Inertia." *Journal of Economic Dynamics and Control* **28**(10): pp. 2003-2035.

Nakagawa, S., R. Hattori, et al. (2000). "Price-Setting Behavior of Japanese Companies." *Bank of Japan* Research Papers No. 2000-09.

Nakamura, E. and J. Steinsson (2008). "Five Facts About Prices: A Reevaluation of Menu Cost Models", Columbia University, *mimeo.*

Reis, R. (2006). "Inattentive Consumers." *Journal of Monetary Economics* **53**(8): pp. 1761-1800.

Reis, R. (2006). "Inattentive Producers." *Review of Economic Studies* **73**(3): pp. 793-821.

Sims, C. (2003). "Implications of Rational Inattention." *Journal of Monetary Economics* **50**(3): pp. 665-690.

Stahl, H. (2005). "Price Setting in German Manufacturing: New Evidence from New Survey Data." *ECB* Working Paper No. 561.

Taylor, J. (1999). "Staggered price and wage setting in macroeconomics" in *Handbook of Macroeconomics.* Eds. J. Taylor and M. Woodford, Elsevier: pp. 1009-1050.