**DEPARTMENT OF ECONOMICS
DISCUSSION PAPER SERIES**

**ACTIVELY LEARNING BY PRICING: A MODEL OF
AN EXPERIMENTING SELLER**

Tim Willems

Number 687
December 2013
Revised February 2014

Manor Road Building, Manor Road, Oxford OX1 3UQ

Actively Learning by Pricing: A Model of an Experimenting Seller*

Tim Willems[†]

January 28, 2014

Abstract

This paper presents a model of a rational seller who is actively learning the slope of his demand curve via his pricing strategy. Consequently, this seller optimally experiments with his price. Resulting price patterns show a lot of discreteness (as observed in the data), which has proved to be a major challenge to most price setting models. The model's learning dynamics are able to reconcile individual price flexibility with aggregate price sluggishness, while the experimentation motive can explain the presence of many idiosyncratic price changes in the data as well as the observation that prices are more volatile than costs.

JEL-classification: D21, D83, E31

Key words: Active learning, Price experimentation, Nominal rigidities, Sales, Discrete pricing

1 Introduction

This paper investigates the extent to which the active learning motive of a seller who faces uncertainty on the slope of his demand curve can reconcile the volatile, discrete pattern followed by individual prices, with the sluggishness observed in the aggregate price level.

*I am especially grateful to Wouter den Haan and Sweder van Wijnbergen for their comments, guidance and support. I also thank Hans Amman, Martin Eichenbaum, Martin Ellison, Xavier Gabaix, Paolo Giordani, Liam Graham, David Kendrick, John Leahy, Peter Leeflang, Jean-Paul L'Huillier, Bartosz Maćkowiak, Albert Marcet, Filip Matějka, Giuseppe Moscarini, Kevin Sheedy, Christian Stoltenberg, Lars Svensson, Volker Wieland, Stephen Wright, and many seminar and conference audiences for stimulating comments and discussions. Part of this research was carried out at the Sveriges Riksbank and I thank them for their hospitality and feedback.

[†]Nuffield College, and Department of Economics, University of Oxford, Manor Road Building, Manor Road, Oxford OX1 3UQ, UK. Also a member of the Centre for Macroeconomics. E-mail: tim.willems@economics.ox.ac.uk.

Recent papers analyzing higher-frequency data on retail prices (like Bils and Klenow (2004)), point out that they are quite flexible at the micro-level. In particular, Kehoe and Midrigan (2010) and Eichenbaum, Jaimovich and Rebelo (2011), who both looked at supermarket scanner data, report that posted prices tend to change about once every two to three weeks. Figure 1, which shows some prices posted by a large US retailer at the weekly frequency, illustrates this.

Simultaneously, there is also evidence that the aggregate price level is sluggish in responding to shocks (*cf.* Christiano, Eichenbaum and Evans (1999)). This has raised the question whether the co-existence of these two observations is actually for real, and if so, why the flexibility at the micro-level does not translate into macro price flexibility.

Lately, the profession has made considerable progress in bringing these two (seemingly conflicting) observations together. Papers like Gertler and Leahy (2008), Nakamura and Steinsson (2010) and Kehoe and Midrigan (2010) manage to replicate these features by refining the standard menu cost model. Others have pointed out that informational imperfections may also play a role: Nimark (2008) for example shows that higher order expectations can reconcile individual price flexibility with aggregate price stickiness, while Maćkowiak and Wiederholt (2009) and Matějka (2010a,b) can match these observations by modeling attention as a scarce resource (following Sims (2003)).

Higher-frequency price data sets have also shown that individual prices are characterized by at least one other interesting feature: next to their flexibility, they simultaneously show a rigidity. This rigidity however does not take the form of price stickiness in the traditional sense, but rather implies that prices tend to bounce back and forth between only a few rigid values - thereby showing a lot of discreteness (see Figure 1). Matching this dimension of the data with "continuous" models (*i.e.* models without an exogenous preference towards discreteness), turns out to be more challenging. In fact, from the aforementioned papers that manage to reconcile individual price flexibility with aggregate price stickiness, only the ones by Matějka (2010a,b) yield discrete pricing patterns from a continuous setup.¹

The present paper therefore develops a different explanation for the micro-macro conflict in pricing behavior - one that is also able to generate individual price series that show a lot of discreteness. It starts by relaxing the strong (but standard) assumption that firms can observe their entire demand curve. Instead, firms in this paper are uncertain on their demand curve's slope.

¹Gabaix (2011) also obtains discreteness, but he does not analyze what his model predicts with respect to the behavior of the aggregate price level. Other pricing models (such as the Calvo one) only generate discreteness by assumption, not from a continuous setup without a built-in preference for discreteness.

Figure 1: Prices posted by a large US retailer for various goods (source: Eichenbaum, Jaimovich and Rebelo (2008)).

Given this uncertainty, firms then try to learn the true value of this parameter - thereby endogenously generating an estimate of their demand curve. The way in which this learning process is modeled, however differs from the passive approach that Sargent (1999) for example used to analyze the Fed's learning process on the Phillips curve. Under that strategy, the estimation stage is separated from the control stage as a result of which the resulting learning rule is not optimal.

This paper on the other hand does optimize the learning rule by taking the links between estimation and control into account. It lets firms learn *actively* about their demand curve. Under active learning a seller realizes that he is learning from self-generated observations (*i.e.* that he is like an econometrician who has the luxury of being able to generate his own data points). Consequently, he recognizes that there is a link between his actions, and the amount of information that is generated on the parameters he is trying to learn. As a result, this seller becomes willing to take somewhat more extreme actions that are suboptimal in the short run, because he realizes that these actions generate information - thereby enabling him to achieve higher pay-offs in the future.

Through this channel, the present paper provides a rationale for the occurrence of

sales. This phenomenon has lately received a lot of attention from macroeconomists, as their prevalence has raised the question whether sales are important for the dynamics of the aggregate price level (Nakamura and Steinsson, 2008). Recent papers have explained the existence of sales by referring to the formation of implicit contracts between buyers and sellers (Nakamura and Steinsson, 2011), the presence of consumers with different price elasticities (Guimarães and Sheedy, 2011), information processing constraints (Matějka, 2010a,b) or by assuming that price adjustment costs for a sale are lower than those accompanying a regular price change (Kehoe and Midrigan, 2010). The (complementary) motive proposed in this paper is that sales result from the desire of firms to obtain more information on the slope of their demand curve (which enables them to achieve higher profits in the future).

If this learning motive indeed plays a role for putting items on sale in reality, one would expect that the frequency of sales is increasing in the uncertainty on demand. And as documented by Pashigian and Bowen (1991), this actually is the case. Next to that, there is also more direct evidence that sellers experiment with their price. Einav *et al.* (2011) for example find evidence that experimentation is extremely common on eBay: of the 100 million listings appearing on eBay on a given day, more than half will reappear again as a separate listing, often with modified sale parameters (such as price). Consequently, they state: "Particularly interesting from an economic viewpoint is the possibility that sellers engage in active experimentation to improve their business practices" (p.25). Campbell and Eden (2010) analyze supermarket scanner data and report that grocers deliberately seem to select extreme prices which they then quickly abandon. This suggests to them that "sellers extensively experiment with their prices" (p.1). Their hypothesis is confirmed by survey evidence from the marketing literature: in a questionnaire held among 32 large US retailers, 90 percent of them say to price-experiment (Gaur and Fisher, 2005).² The fact that an increasing number of sellers has started to use pricing optimization software (algorithms that tend to take the experimentation motive explicitly into account - very much like in this paper's model), only adds to this.³

²In September 2000, Amazon.com actually made the headlines with its experimental behavior: users found out that the same DVDs were simultaneously offered to different buyers at different prices. Amazon stated that these differences were due to price experimentation, but critics suggested that they were based on consumer profiling. Since the latter is perceived as being unfair, Amazon promised to refrain from cross-sectional price experimentation in the future - leaving only the time-series dimension to do so (which is going to be the focus in the remainder of this paper). See "Have They Got a Deal For You" in *The Washington Post* of June 19, 2005 for more on this case.

³According to the 2004 "Retail Horizons"-survey conducted among 300 US retailers, 36 percent of them used pricing optimization software (notable users include Walmart and 7-Eleven). In their 2002 survey, this number was only 12 percent. If this trend (along with the increasing importance of online selling) continues, the active learning process of sellers is likely to become even more important in the future as

More generally, the marketing literature tells us that retailers are well aware of the challenges associated with optimal price setting in the face of demand uncertainty. Retailers themselves for example name "measuring the impact of pricing decisions" as their top operational business challenge,⁴ in response to which the revenue management literature has produced a large number of papers that looks at exactly this problem (see Araman and Caldentey (2009) for an example).

In this paper, I show that the price paths resulting from such an experimentation process contain a lot of discreteness: despite the fact that the model fully consists of continuous equations, prices bounce back and forth between only a few rigid values, just as in the data (recall Figure 1). The reason is that the price setting rules under active learning take the form of a step function: there is a low price platform (associated with low demand) and a high one (for high demand), separated by a sharp kink (which serves to avoid the posting of "uninformative prices" - a concept described below). In simulations, the resulting price paths look very much like those generated by Matějka's (2010a,b) rationally inattentive seller and consumer models, as they also manage to produce discrete pricing patterns without having an exogenous preference towards it.⁵

In sum, this paper shows that modeling the seller as an optimal experimenter is able to replicate at least three aspects of the data: the experimentation motive generates volatile pricing patterns showing a lot of discreteness at the individual level, while the fact that there still is some learning going on makes the aggregate price level respond sluggishly to shocks. So once one considers a model that optimizes the learning process itself, it turns out that learning is not only able to match the aggregate dimension of the data (which is sort of well-known by now), but that it is also capable of reproducing important micro-elements of it. Do note that this result is *not* driven by some form of irrationality: sellers in the model are just responding optimally to the fact that they face a demand curve with an unobserved, time-varying slope.

The remainder of this paper is structured as follows. First, Section 2 provides a non-exhaustive overview on the use of active learning in economics (see Kendrick, Amman and Tucci (2011) for a broader survey). Section 3 then describes the model, after which Section 4 explains the solution method. Section 5 contains the calibration and model results, after which a discussion (Section 6) and conclusion (Section 7) complete the paper.

both developments greatly facilitate the practical implementation of optimal price experimentation.

⁴See the RSR-report "Getting Back to Good: Retail Pricing 2010".

⁵As I will explain later on, the discreteness that arises in the present paper is not exact, but looks very much like it. In the rational inattention papers of Matějka (2010a,b) on the other hand, the discreteness is exact. See Matějka and Sims (2010) for a general discussion of discreteness in models with rationally inattentive agents.

2 Relation to the literature

Active learning was first introduced in the economics literature by Prescott (1972).⁶ He considered general optimal control problems in the presence of an experimentation motive. Subsequently, active learning has for example been applied to experimental drug consumption (*cf.* Grossman, Kihlstrom and Mirman (1977)), investment and growth under uncertainty (Bertocchi and Spagat, 1998) and to optimal monetary policy (see Wieland (2000b; 2006), Ellison and Valla (2001) and Svensson and Williams (2007)).

The first application of the active learning-concept to a monopolist who tries to learn his demand curve can be found in Rothschild (1974). He framed the learning process as a two-armed bandit problem and noted that this problem contains a trade-off between long and short run profit maximization (the former calling for price experimentation, while the latter goal is achieved by posting the price that present information indicates is most profitable). Subsequently, several papers have derived important analytical results in this context: Aghion *et al.* (1991) characterize limit actions and beliefs, Treffler (1993) establishes results that determine the direction of experimentation, while Keller and Rady (1999) investigate how the experimentation incentives vary with the discount rate and the persistence of the demand curve.⁷

Following advances in computational power, the present paper is able to take a more practical approach that visualizes the experimentation motive. In this sense, it also relates to Kiefer (1989), who solved Rothschild's (1974) two-armed bandit version of the problem numerically and analyzed how the properties of the value and policy functions varied with the discount factor. In contrast to the present paper, he however did not analyze what this implies for the paths followed by individual and aggregate prices, nor did he establish the link with sales.

In line with most of the aforementioned contributions, this paper will assume that firms are using their pricing strategy to learn about the slope of their demand curve only. In sharp contrast, there is assumed to be no learning motive on the intercept term.⁸ Balvers and Cosimano (1990), which also forms an important predecessor to the present

⁶The concept of active learning is also referred to as "optimal experimentation" (Wieland (2000a), Cosimano (2008)) or "optimal learning" (Wieland (2000b), Aghion *et al.* (1991)). It originated in the engineering literature, where it is known as "dual control".

⁷Contemporaneous work by Bachmann and Moscarini (2012) also analyzes a model in which sellers have an experimentation motive, but they mainly use their model to show that macroeconomic fluctuations can generate fluctuations in uncertainty (as measured by price dispersion). I will return to the relevance of Bachmann and Moscarini's work for the present paper in Section 6.

⁸Due to the dimensionality curse associated with the solution method, solving the full problem in which firms are learning on both slope and intercept is very difficult. See Section 6 for a further discussion of this issue and an idea on how the addition of a learning motive for the intercept would affect the results.

paper, consider the other polar case in which firms are only trying to learn the intercept of their demand curve in an active manner. They show that this provides firms with an incentive to mute their price adjustments in response to changes in demand. Under their specification, active learning thus implies *the absence of price changes*. Consequently, their model can replicate the delayed response of the aggregate price level in response to shifts in demand, but is unable to match the recently uncovered observation that individual prices change so frequently. This paper's specification on the other hand, turns out to be able to reproduce both dimensions of the data (see Section 5).

3 Model

Consider a price-setting seller i who operates in a monopolistically competitive market.⁹ He sells his product i in period t at price $p_{i,t}$ and serves whatever quantity is demanded at that price. There are no price adjustment costs, so $p_{i,t}$ is set on a period-by-period basis in a fully flexible way. The seller faces a marginal production cost that is equal to c , which is assumed to be known and constant over time. Indicating period t demand for his product by $q_{i,t}$ and using δ to denote the discount factor, his profit maximization problem reads:

$$\max_{p_{i,t}} \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \delta^t [p_{i,t} - c] q_{i,t} \right\} \quad (1)$$

He believes that demand for his product follows a linear specification.¹⁰ In particular:

$$\begin{aligned} q_{i,t} &= \alpha_t - \beta_{i,t} [p_{i,t} - \bar{p}_t] + \gamma q_{i,t-1} \\ &\equiv \alpha_t - \beta_{i,t} r_{i,t} + \gamma q_{i,t-1} \end{aligned} \quad (2)$$

⁹The retailer I consider only sells his product at one location. If he would possess multiple selling points, he could in principle also use the cross-sectional dimension to experiment. But if market power differs between various locations (for example due to the fact that store A does have a nearby competitor, while store B does not), it is not clear whether there is much to learn for A from B and vice versa. Moreover, supermarkets in some countries (such as the UK) have made agreements with authorities to follow a "national pricing policy", which implies that the same price is charged at all selling points. In practice, retail chains also do not always share information among all sellers, for example due to the fact that different stores have different owners (franchising). Finally, experimenting cross-sectionally is difficult in an online environment (recall footnote 2).

¹⁰Whether this belief is correct or not is not important (in this paper it will be correct though, so that there is no misspecification bias). The linearity assumption is not essential either (*cf.* Keller and Rady (1999)): the only thing that matters is that the seller estimates a regression equation in his learning step that contains a price sensitivity term which operates over the *relative* price he posts. The linear demand curve (2) satisfies this requirement, as does the CES demand curve (the other popular demand specification).

Here, $\bar{p}_t$ represents the aggregate price level. Its value lies beyond the control of each individual firm. Furthermore, $r_{i,t} \equiv p_{i,t} - \bar{p}_t$ is firm i 's relative price, α_t is the aggregate demand component that is common to all firms, and $\beta_{i,t}$ is the slope of the demand curve for the good under consideration. This slope expresses the degree of product differentiation perceived by consumers (do consumers see any close substitutes to variety i ?); consequently, it operates over firm i 's *relative* price. Following evidence from the marketing literature (*cf.* Gordon, Goldfarb and Li (2013) and the references therein), this parameter will be time-varying (*cf.* equation (4) below).¹¹

There are numerous reasons why the perceived degree of product substitutability can vary over time: it may be due to aggregate conditions (consumers becoming more price sensitive in recessions), store-specific ones (a competitor opening a nearby store, increasing the price sensitivity of shoppers), while brand-specific conditions may also play a role (think of a successful advertising campaign making consumers less price sensitive towards your product, or bad press establishing the reverse). In this paper, all of these factors are captured by the stochastic process for $\beta_{i,t}$.

As both the industrial organization and marketing literature have generated evidence that consumer behavior is subject to some form of brand loyalty (*cf.* Klemperer (1995), Chintagunta, Kyriazidou and Perktold (2001) and the references therein), I allow for the formation of good-specific habits via the parameter $\gamma > 0$.¹² The value of this parameter is assumed to be constant and known to the seller. It should however be stressed that the experimentation motive would still be present without any form of habit formation, so from that point of view its inclusion is not essential (see *e.g.* Kiefer (1989), Wieland (2000a) and Bachmann and Moscarini (2012) who analyze active learning problems without a lagged dependent variable). But, as I will explain in Section 5 of this paper, the presence of good-specific habits does play a role in generating individual price volatility (as this volatility is going to stem from *the interaction* of the active learning motive with the existence of good-specific habits).

The crucial difference with the standard model is the seller's information set. Whereas the standard model makes the rather strong assumption that sellers can perfectly observe their entire demand curve, this paper relaxes this assumption and keeps the intercept

¹¹Even though time-variation in this parameter may exist, Chari, Kehoe and McGrattan (2009) show that the driving shock is open to multiple interpretations and therefore non-structural. Until more evidence is available on the exact origin of the shock, caution should therefore be taken in using such a model for policy analysis. The same proviso applies to the standard New Keynesian model.

¹²To keep the analysis clear, I abstract from the time-consistency issues that are central to Nakamura and Steinsson (2011). As a result, the commitment-channel they emphasize is not at work in the present paper.

α_t and slope $\beta_{i,t}$ unobserved to the seller. Apart from the (p, q) -point he is producing at every period, the seller thus faces uncertainty on the demand curve for his product. Instead, the seller only knows the stochastic processes that α_t and $\beta_{i,t}$ follow.

Although in practice both the intercept and slope of the demand curve are estimated to be persistent, this paper makes the simplifying assumption that aggregate demand α_t follows an i.i.d.-process with known mean $\bar{\alpha}$, hence:

$$\alpha_t = \bar{\alpha} + \mu_t, \text{ with } \mu_t \sim N(0, \sigma_\mu^2) \quad (3)$$

Consequently, there is no learning motive for the intercept term (since it has no persistence its expected value is always equal to $\bar{\alpha}$). This limits the number of state variables to a minimum, thereby facilitating the process of finding a numerical solution to the model (see Section 6 for a more detailed discussion of this assumption).

The unobserved process for $\beta_{i,t}$ is assumed to be persistent. As a result, sellers in the model do have an incentive to learn the actual value of this parameter. I assume that $\beta_{i,t}$ follows an $AR(1)$ around known mean $\bar{\beta}$ (where the latter is common to all sellers, representing aggregate economic conditions):

$$\beta_{i,t} = (1 - \rho_\beta) \bar{\beta} + \rho_\beta \beta_{i,t-1} + \eta_{i,t}, \text{ with } \eta_{i,t} \sim N(0, \sigma_\eta^2) \quad (4)$$

The timing and informational structure of the model are as follows: at the beginning of period t , firms have prior beliefs regarding the unobserved parameter they are trying to learn, $\beta_{i,t}$. These prior beliefs can be described by a conditional normal distribution:

$$p(\beta_{i,t} | \Xi_{i,t-1}) = \mathcal{N}(b_{i,t|t-1}, \Sigma_{i,t|t-1}^b)$$

Here, notation is such that $\mathbb{E}_{t-1}[\beta_{i,t}] \equiv b_{i,t|t-1}$ and $\text{var}_{t|t-1}[\beta_{i,t}] \equiv \Sigma_{i,t|t-1}^b$. From this expression one can see that $\Xi_{i,t-1}$, the vector of state variables describing beliefs, contains two elements: $b_{i,t|t-1}$ and $\Sigma_{i,t|t-1}^b$.

Given these beliefs, firm i sets its relative price $r_{i,t}$, after which all shocks materialize and the firm gets to observe realized demand $q_{i,t}$. With this new information, the firm then updates its beliefs via Bayes' rule, which states that:

$$p(\beta_{i,t} | \Xi_{i,t}) = \frac{p(q_{i,t} | \beta_{i,t}, r_{i,t}, \Xi_{i,t-1}) p(\beta_{i,t} | \Xi_{i,t-1})}{p(q_{i,t} | r_{i,t}, \Xi_{i,t-1})}$$

Next to the actual shocks hitting the economy (μ_t and $\eta_{i,t}$), sellers in practice also have

to deal with the existence of transitory measurement error that is present in real time (*cf.* Bomfim (2001) and Blanchard, L'Huillier and Lorenzoni (2013)). Denoting the variance of the noise associated with α_t by σ_ϵ^2 , and that of the noise associated with $\beta_{i,t}$ by σ_ν^2 , we can write the Bayesian updating and prediction equations as (*cf.* Zellner (1971)):

$$b_{i,t|t} = b_{i,t|t-1} - \frac{\Sigma_{i,t|t-1}^b r_{i,t}}{r_{i,t}^2 \Sigma_{i,t|t-1}^b + \sigma_\mu^2 + \sigma_\epsilon^2} (q_{i,t} - \gamma q_{i,t-1} - \bar{\alpha} + b_{i,t|t-1} r_{i,t}) \quad (5)$$

$$\Sigma_{i,t|t}^b = \Sigma_{i,t|t-1}^b - \frac{\left(\Sigma_{i,t|t-1}^b\right)^2 r_{i,t}^2}{r_{i,t}^2 \Sigma_{i,t|t-1}^b + \sigma_\mu^2 + \sigma_\epsilon^2} \quad (6)$$

$$b_{i,t+1|t} = (1 - \rho_\beta) \bar{b} + \rho_\beta b_{i,t|t} \quad (7)$$

$$\Sigma_{i,t+1|t}^b = \rho_\beta^2 \Sigma_{i,t|t}^b + \sigma_\eta^2 + \sigma_\nu^2 \quad (8)$$

Apart from the fact that the updating equations are non-linear functions of $q_{i,t}$ and $r_{i,t}$, one should also note that they establish a link between different periods since actions taken in the past affect (the quality of) future beliefs. This adds a dynamic dimension to the firm's problem.

Furthermore, active learning endogenizes the model's information structure. In particular, the firm's speed of learning, which is determined by the Kalman gain-like expression $\Sigma_{i,t|t-1}^b r_{i,t} / (r_{i,t}^2 \Sigma_{i,t|t-1}^b + \sigma_\mu^2 + \sigma_\epsilon^2)$ in equation (5), can be affected by the firm's decisions. More precisely, it is a function of the relative price $r_{i,t}$ firm i sets. The reason is that by setting its relative price in an "informative way", a firm can generate more information on the system it is trying to learn about - thereby facilitating the learning process. This contrasts with standard signal extraction problems, such as Lucas (1973), in which the speed of learning is simply a function of the exogenously given signal-to-noise ratio.

In the present setting on the other hand, sellers can choose to generate more precise signals by letting the control $r_{i,t}$ take on more extreme values (*i.e.*: values located further away from zero in absolute value).¹³ After all, if one estimates a regression equation like (2) (which sellers in this model are doing), one would like to have somewhat more extreme observations of $r_{i,t}$ so as to obtain better information on its coefficient $\beta_{i,t}$.¹⁴

¹³Herewith, active learning also has an interesting relationship with rational inattention (Sims, 2003): while both concepts endogenize the signal-to-noise ratio, rational inattention achieves this by endogenizing the amount of *noise* in the system; active learning on the other hand endogenizes the *signal*.

¹⁴To see this, note from (5) and (6) that one would not learn anything about $\beta_{i,t}$ if $r_{i,t}$ would remain constant at zero over time. Equation (8) then implies that uncertainty about the true value of $\beta_{i,t}$ will grow with $(\sigma_\eta^2 + \sigma_\nu^2)$ every period. The reason is that posting the average price is not informative on the perceived degree of product differentiation. More formally, when $r_{i,t} = 0$, $\beta_{i,t}$ ceases to play a role in the demand curve (2) - thereby introducing an identification problem with respect to this parameter.

And since the explanatory variable of the regression also happens to be the control of the system, this can actually be achieved: the experimenting seller realizes that he is like an econometrician who has the luxury of being able to generate his own data points.

4 Solving the model

The seller's problem is given by (1) subject to equations (2)-(8). As shown by Kiefer and Nyarko (1989), this type of problem still admits a contraction mapping argument. In particular, the problem satisfies Blackwell's sufficiency conditions of monotonicity and discounting, as a result of which it has a fixed point (being the value function). Consequently, it can be solved by dynamic programming. The problem has three state variables, $\{q_{i,t-1}, b_{i,t|t-1}, \Sigma_{i,t|t-1}^b\} \equiv \{q_{i,t-1}, \Xi_{i,t-1}\}$, and the Bellman equation can be written as:

$$V(q_{i,t-1}, \Xi_{i,t-1}) = \max_{r_{i,t}} \left[\begin{aligned} &\Pi(q_{i,t-1}, \Xi_{i,t-1}, r_{i,t}) + \delta \int V(q_{i,t}, \Xi_{i,t}) \\ &\times f(q_{i,t}|q_{i,t-1}, \Xi_{i,t-1}, r_{i,t}) dq \end{aligned} \right]$$

This equation nicely shows the trade-off between estimation and control that firms face: on the one hand, they want to maximize current period profits $\Pi_{i,t}$, but on the other hand they also have to take the expected continuation value into account (the second term). The latter contains the potential improvements in future performance that may result from past price experimentation and is therefore increasing in the absolute value of $r_{i,t}$. The fact that price experimentation is costly in terms of current period profits, prevents the seller from posting prices that are "too extreme" (even though such actions would produce a lot of information on the true value of $\beta_{i,t}$). Hereby, the profit maximizing objective of sellers imposes an *endogenous* bound on the precision of the signals they choose to generate. Hence, in this paper's setup, optimizing behavior fulfills a similar role as the exogenous information processing constraint does in rational inattention models.

The model is solved by iterating over the Bellman equation until the value function converges. The exact algorithm is described in Beck and Wieland (2002). It implies that, starting from an initial guess of the value function, one updates the latter by maximizing the right-hand side of the Bellman equation. There, the conditional expectation is evaluated using Gauss-Hermite nodes (since the dependent variable q is normally distributed). To speed up convergence, policy iterations are carried out after each value iteration.

Since these type of control problems are typically characterized by multiple local maxima (see footnote 18), the algorithm starts with a rough grid search, after which a golden section search is carried out to compute the maximum more precisely.

5 Price setting under active learning

We are now in the position to solve the problem set out in Section 3, and see what the introduction of active learning implies for the price setting behavior of firms. In order to do so, Section 5.1 first discusses the model's calibration, after which Section 5.2 describes the model outcomes.

5.1 Calibration

As this paper deals with higher-frequency price movements, I calibrate the model at a weekly frequency. Table 1 summarizes the calibration. The weekly discount factor δ is set to imply a standard value for the quarterly discount factor of 0.99. The parameter capturing good-specific habits (γ) equals the estimate of Ravn, Schmitt-Grohé and Uribe (2006) converted to a weekly frequency. The marginal cost of production c is normalized to 1, while the value of the autoregressive parameter for $\beta_{i,t}$ (ρ_β) is based upon Smets and Wouters (2007), who estimate their equivalent of ρ_β to equal 0.89 on quarterly US data. This leads to a weekly value of 0.99.

I will simplify the analysis by looking at the partial equilibrium case in which the aggregate price level $\bar{p}_t$ is exogenous. In particular, I will just set it equal to a constant, hence $\bar{p}_t = \bar{p} \ \forall t$ (the same assumption is effectively made in Matějka (2010a)).

The actual value of $\bar{p}$, along with the values of the other two parameters that determine the location of the demand curve (the average intercept $\bar{\alpha}$ and average slope $\bar{\beta}$), are however difficult to calibrate as their values cannot be readily estimated. Consequently, these parameters are set in order to ensure two things: i) that sales q_t remain positive in typical simulations and ii) that the average price posted by a firm over time equals $\bar{p}$ (so that individual and aggregate prices are internally consistent in the model). I have found values of $\bar{\alpha} = 0.875$, $\bar{\beta} = 1.85$ and $\bar{p} = 1.25$ to work well for these purposes.

Whereas the actual value picked for $\bar{\beta}$ is not that important, the relative uncertainty that $\bar{\beta}$ is surrounded with, is key. The reason is that this determines the seller's experimentation incentives. This relative uncertainty can for example be measured by the t -statistic on β_t . The latter has recently been estimated in the marketing literature by Gordon, Goldfarb and Li (2013). Using household panel data running from January 2001 up to December 2006, they come up with an average t -statistic of about 6.65 across all 19 grocery categories in their study. Given that $\bar{\beta} = 1.85$ this implies that β_t should get a 0.2782 standard deviation, which requires σ_η^2 to equal 0.00154. The variance of aggregate demand α (σ_μ^2) is set equal to the variance of the natural log of linearly detrended US

nominal GDP from 1947q1 to 2011q1 (taken from the St. Louis Fed).

These estimates of σ_μ^2 and σ_η^2 are however fabricated by researchers ex post. Firms in real time face greater uncertainties as they also have to filter out the transitory measurement error that is present in real time data (Bomfim, 2001). In this respect, Blanchard, L’Huillier and Lorenzoni (2013) estimate a signal-to-noise ratio of about 1 on actual data, which motivates my choice for σ_ϵ^2 and σ_ν^2 (incidentally, 1 was also the signal-to-noise ratio that was used by Bomfim (2001)).

Symbol	Interpretation	Value
$\bar{\alpha}$	Average aggregate demand	0.875
$\bar{\beta}$	Average slope of demand curve	1.85
c	Marginal production cost	1
γ	Habit formation coefficient	0.9875
δ	Discount factor	0.9992
$\bar{p}$	Aggregate price level	1.25
ρ_β	AR-component in β -process	0.99
σ_μ^2	Variance of α -innovation	0.0164
σ_η^2	Variance of β -innovation	0.00154
σ_ϵ^2	Variance of noise in α -process	0.0164
σ_ν^2	Variance of noise in β -process	0.00154

Table 1: Calibration.

5.2 Model results

To visualize the impact of informational structure on the final solution, I will start by reducing the problem described by equations (1)-(8) to two simpler problems.

First, one can derive the policy rule for the standard case in which firms live in a perfect information world where they can observe the true values of α_t and $\beta_{i,t}$ without error. In that case, the problem under consideration reduces to a simpler one and one can show that the profit maximizing price equals:

$$p_{i,t}^* = \bar{p} + \frac{1}{2\beta_{i,t}} [\alpha_t + \beta_{i,t}(c - \bar{p}) + \gamma q_{i,t-1}] - \delta \gamma \mathbb{E}_t \left\{ \frac{q_{i,t+1}}{2\beta_{i,t+1}} \right\} \quad (9)$$

Here, both the forward looking term ($\gamma \mathbb{E}_t \{q_{i,t+1}/2\beta_{i,t+1}\}$) as well as the lagged term ($\gamma q_{i,t-1}/2\beta_{i,t}$) result from the dynamic link that is established by the formation of good-specific habits. In particular, this equation shows the trade-off (described in Klemperer (1995)) that sellers face between setting a low price to lock consumers into their specific

product (which is what the forward looking term captures) and setting a high price to harvest profits by exploiting consumers who are already "addicted" due to past consumption (this is what is expressed by the lagged term).

Next, we can move on to the specification in which α_t and $\beta_{i,t}$ cannot be observed, as a result of which agents have to learn their true values. In that case, firms could follow the so-called "passive" policy rule. It is however key to note that this rule is not optimal as it implies that the seller does not realize that he is learning from self-generated observations. Consequently, he disregards the link between current and future beliefs (*i.e.* in the optimization step he neglects equations (5)-(8)), as a result of which this learning rule lacks any experimentation motives: although passive learners do update their beliefs as more observations arrive over time, they do not actively seek for better information through price experimentation.

For this case, the passive price setting rule can be shown to equal:¹⁵

$$p_{i,t}^\diamond = \bar{p} + \frac{1}{2b_{i,t|t-1}} [\bar{\alpha} + b_{i,t|t-1} (c - \bar{p}) + \gamma q_{i,t-1}] - \delta \gamma \mathbb{E}_t \left\{ \frac{q_{i,t+1}}{2b_{i,t+1|t}} \right\} \quad (10)$$

Deviations of the active learning rule from this passive one, measure the extent of price experimentation.

Unfortunately, it is not possible to obtain an analytical expression for the price setting rule under active learning. We can however characterize an important characteristic of it, namely that the optimal action under active learning will always be more extreme than the one under the passive policy rule. This is done in Proposition 1, which builds upon the following two lemmas. The proofs of these statements can be found in the Appendix.

Lemma 1 *Uncertainty on β is decreasing in the squared deviation of the control variable from zero, that is: $\partial \Sigma_{i,t|t}^b / \partial (r_{i,t}^2) < 0$.*

Lemma 1 formally establishes the previously stated notion that the presence of more extreme observations of the explanatory variable, facilitates the learning process on the unknown coefficient β_t . Note that the direction of the deviation (positive or negative?) does not matter, as $r_{i,t}$ only enters equation (6) in a squared fashion - a feature that I will return to later. Lemma 2 shows that this information is actually valuable.¹⁶

¹⁵Note that the passive policy rule is equal to the certainty equivalent one (*i.e.* the rule that one obtains if one replaces actual values by estimates in equation (9)). Also see Balvers and Cosimano (1990) on this.

¹⁶This lemma does not require differentiability of the value function, which is also not always guaranteed in active learning problems (*cf.* Balvers and Cosimano (1994), upon which the proof of this lemma (and that of Proposition 1) is inspired).

Lemma 2 *The value of the firm is non-increasing in the uncertainty on β , that is: $\Delta V(q_{i,t-1}, b_{i,t|t-1}, \Sigma_{i,t|t-1}^b) / \Delta \Sigma_{i,t|t-1}^b \leq 0$.*

With help of these two lemmas, we can prove Proposition 1:

Proposition 1 *The relative price set under active learning (call it $r_{i,t}^\dagger$) is always more extreme than the price chosen under the passive learning strategy, i.e. $|r_{i,t}^\dagger| > |r_{i,t}^\diamond|$.*

The intuition for this lies in the fact that it can never be optimal to post a price that is smaller in absolute value than the price chosen under passive learning: after all, the price set under the passive learning rule is the one that maximizes profits given current beliefs, while any deviation from that price towards zero produces less information. Since the losses from price experimentation are only second-order in the neighborhood of $r_{i,t}^\diamond$ (through an envelope theorem argument), while price experimentation yields first-order benefits by Lemma's 1 and 2, it follows that the inequality in Proposition 1 is strict.

To gain more insight into the effects of active learning, we can solve for the policy function using the Beck-Wieland algorithm set out in Section 4 of this paper. This was done in order to generate Figure 2. That figure plots the active learning pricing rule and compares it with the passive policy for different combinations of beliefs (b, Σ^b) .¹⁷

The figure illustrates the experimentation motive: while the passive price setting rule is just a linear, increasing function of demand, the rule resulting from the active learning process contains a kink. This non-linearity shows up in that part of the state space where posting the less extreme, passive learning price would be uninformative on the perceived degree of product differentiation, $\beta_{i,t}$. Hence, the discontinuity serves to avoid uninformative actions (that is: it avoids posting a relative price $r_{i,t}$ close to zero).¹⁸

The fact that the policy functions tend to flatten out away from the discontinuity, shows that there is less price experimentation when demand for variety i is either relatively high or low. The reason is that in these extremal parts of the state space the relative price picked by the passive policy (which is the optimal one if one neglects the

¹⁷Note that the perfect information price setting rule (9) equals the passive learning rule (10) for $b_{i,t|t-1} = \beta_{i,t}$ (this is the certainty equivalence property). Consequently, the passive learning and perfect information rules would coincide in Figure 2 as a result of which I have not included the latter explicitly in the figure.

¹⁸More technically, the reason lies in the fact that the value function is twin-peaked. This is due to the fact that it is a combination of a concave function in $r_{i,t}$ (static profits) and a convex one (the continuation value). The continuation value is increasing in the absolute value of $r_{i,t}$ as more extreme relative prices are more informative. This generates two local maxima (one where $r_{i,t} < 0$ and one where $r_{i,t} > 0$), from which the max-operator in the Bellman equation then selects the one yielding the highest value. This leads to an avoidance of the region where the continuation value is minimized, i.e. that region where $r_{i,t}$ is "too close" to zero.

learning motive) has already moved away from zero, due to the fact that demand is so high/low. Consequently, posting the passive price is not that bad anymore from a learning perspective, thereby reducing the need for additional price experimentation (which is costly as it lowers current period profits). As a result, the active learning rule crawls towards the passive one as conditions get more extreme.

Therefore, the optimal price setting rule is shaped like a step function, and prices tend to jump between two rather rigid values: when demand is low, firms post a low price, while an increase in demand above a certain threshold makes them post the higher price. This can be seen as an endogenous form of "local" price stickiness: as long as the model dynamics do not push a seller from the high to the low price platform of his policy function (or vice versa), his optimal price is relatively unresponsive to shocks.

Figure 2: Comparison of active and passive price setting rules for different beliefs.

Consequently, the experimentation motive is able to replicate the empirical observation (displayed in Figure 1) that individual prices follow a rather discrete pattern over time. This is illustrated in the upper panel of Figure 3,¹⁹ which contains a simulated series (spanning a period of approximately three years) for the control variable $r_{i,t}$ under

¹⁹This particular figure was constructed by giving the seller correct initial beliefs about the true value of β_1 , but with an initial uncertainty $\Sigma_{1|0}^b = 0.052$. The typical pattern arising in simulations is robust to choosing different initial beliefs.

active learning. Since the optimal price set under active learning varies a bit with beliefs $(b_{i,t|t-1}, \Sigma_{i,t|t-1}^b)$, and since the upper and lower platforms of the price setting rules are not entirely flat, the resulting price pattern is not exactly discrete. However, it looks very much like it. The reason is that the policy function under active learning has a rather flat slope in all dimensions of the problem. With respect to the $b_{i,t|t-1}$ -dimension, the intuition is similar to that for the $q_{i,t-1}$ -dimension: less conscious price experimentation is needed when conditions are more extreme (*i.e.* when $\beta_{i,t}$ is believed to be either very high or very low). In those circumstances the passive learner already posts a relative price that has moved away from zero (thereby making it more informative) and the active learner does not have to undertake so much costly experimentation on top. With respect to $\Sigma_{i,t|t-1}^b$, the fact that the slope parameter is time-varying (which is motivated by evidence from the marketing literature; *cf.* Gordon, Goldfarb and Li (2013)) implies that uncertainty continues to renew itself and that the learning process never stops. Consequently, uncertainty on the estimate of $\beta_{i,t}$ remains relatively stable over time and the slope in that dimension only gets to play a subordinate role.

Adding a small menu cost would probably increase the degree of discreteness in Figure 3: then the small price changes that can be observed in the current simulation would no longer be worthwhile (a "region of inaction" would arise). Moreover, a menu cost could also be helpful in decreasing the frequency of price changes: in the model, prices tend to change from high to low (or vice versa) about once every two weeks, which is slightly more often than (but close to) the findings of Kehoe and Midrigan (2010) and Eichenbaum, Jaimovich and Rebelo (2011) (both studies report that posted prices tend to change about once every three weeks in their data sets).

The price path arising under active learning bears similarity to the one obtained by Matějka (2010a). Sellers in his model have limited information processing capacity, as a result of which it becomes optimal for them to pay attention to a source of information that provides a small number of different signals only. Such a signal could be the first digit of unit input costs and since a digit can only take a few different values, there are only a few different signals in his framework. This leads to exact discreteness in his setup.²⁰

In the present paper, the (near-)discreteness arises purely from the active learning process: apart from the latter, the model is a rather standard profit maximization problem without price adjustment costs. This can also be seen by looking at the two lower panels of Figure 3. Those show the price paths in response to the same series of shocks under both passive learning as well as under perfect information (the case in which sellers can

²⁰In Matějka (2010b), the discreteness follows from modeling *the consumer* as being rationally inattentive.

simply observe the true values of α_t and $\beta_{i,t}$). As one can see, these series no longer show any tendency for discreteness (a result of the fact that the passive and perfect information policy rules are not kinked).

Figure 3: Simulated paths for the relative price posted under active learning, passive learning and under perfect information.

Also note that passively learning sellers have no problems with posting a relative price that is close to zero, as they do not realize that this obstructs their learning process on the degree of product differentiation, $\beta_{i,t}$. Sellers who are learning in the optimal active manner do realize this, as a result of which they avoid posting a zero relative price (this is what the kink in the policy function takes care of). Consequently, learning is faster in the active specification. This is illustrated by Figure 4, which was generated by giving a seller incorrect initial beliefs on the value of β (namely that its value equals 2.80 instead of 1.85), after which both an active and a passive learner start updating their beliefs. During this learning process, all shocks were shut down to illustrate the process of convergence towards the truth. As can be seen from the figure, learning is faster when it is done in an active manner, since an active learner consciously generates more extreme (and thereby more informative) observations.

Turning back to Figure 3, one can see that the model is also able to replicate the empirical observation that sellers *alternate* between posting the high and the low price. This would not happen if there were no good-specific habits. The reason is that, from

Figure 4: Learning process on β under both active and passive learning. Initial estimate is 2.80, while the true value of β equals 1.85.

an informational point of view, a low price is just as informative as a high one with the same distance from the aggregate price level (recall Lemma 1). As noted by Bachmann and Moscarini (2012), what matters is the *absolute deviation* of $p_{i,t}$ from $\bar{p}$, *not its sign*. In the most standard setting, sellers would therefore just post the high (or low) price and stick to that choice ever after (this happens in Bachmann and Moscarini (2012)). But with good-specific habits, the demand curve contains a lagged dependent variable which introduces an additional dynamic link: a seller who posts an above average price in period t (*i.e.* someone who has set $r_{i,t} > 0$) will sell less units in that period (by the demand curve (2)) as a result of which he faces a lower endogenous state variable (the stock of habits $q_{i,t}$) in the next period. Optimality subsequently requires him to post a lower price, but the active learner then realizes that it is better to post a negative relative price, rather than sticking to a positive one with an absolute value closer to zero (as the latter option produces too little information on $\beta_{i,t}$). It is thus *the interaction* between active learning and the existence of good-specific habits that is able to replicate this aspect of the data.

In addition, Figure 3 illustrates that the experimentation motive introduces additional volatility in the price series (also driven by the aforementioned interaction). This could be an explanation for the observation that prices are more volatile than underlying marginal costs (as documented by Eichenbaum, Jaimovich and Rebelo (2011) and Carlsson and

Skans (2012)). Time-variation in the slope of the demand curve already helps (due to this, prices under perfect information and passive learning are also more volatile than marginal costs (which are constant in this paper)), but the experimentation motive gives this channel an extra kick.

Related to this, the present paper also links to Golosov and Lucas (2007) (and the large literature following that paper), who argue that the relatively large idiosyncratic price movements observed in the data indicate the presence of large idiosyncratic shocks. This paper however shows that it do not necessarily have to be large shocks driving these movements; instead, it could be that much smaller shocks are being amplified by the endogenous, idiosyncratic experimentation process of each individual seller. On this, De Graeve and Walentin (2011) actually provide evidence that idiosyncratic shocks are not as large as previously thought (thereby questioning the Golosov-Lucas calibration), while Nakamura (2008) argues that the large shocks to manufacturers' productivity assumed by Golosov and Lucas (2007) would generate substantial comovement across prices for the same good at different stores, which is not observed in the data.

Finally, one can analyze how the aggregate price level in the active learning model responds to aggregate shocks, and compare this response to the ones obtained under passive learning and perfect information. Following Matějka (2010a), I do this by simulating, and averaging over, the behavior of a large number of agents. Like in Matějka (2010a), this paper's partial equilibrium setup is not ideal to study this issue (as an aggregate shock will change the aggregate price level $\bar{p}$, which is assumed constant here), but such an exercise does seem to capture the core of the underlying learning dynamics at play.²¹

To tackle this question, we need to make a slight modification to the model considered so far. In particular, we need to introduce aggregate risk with respect to the slope of the demand curve. In the benchmark model, aggregate conditions were captured by the common mean value of the process for β (indicated by $\bar{\beta}$). There, this mean was however assumed to be known and kept constant over time, which is going to be relaxed in the remainder of this section.

The most insightful way to do this, and simultaneously illustrate the accompanying learning dynamics that are at work, is to assume that the economy can be in two regimes: a boom or a recession. Associated with these two regimes are two values for the long

²¹If anything, intuition suggests that allowing $\bar{p}$ to change following aggregate shocks only strengthens the conclusions that are to follow if $\bar{p}$ is unobserved as well. The reason is that sellers then have to learn about an extra parameter (and potentially form higher order beliefs), which complicates their learning process - thereby reinforcing the upcoming results (which are driven by the sluggishness stemming from the learning process). See Nimark (2008) for a model along these lines. He indeed finds that the formation of higher order beliefs on the aggregate price level makes the latter respond sluggishly to shocks.

run average of the slope parameter: $\bar{\beta}_R$ (applicable in a recession) and $\bar{\beta}_B$ (applicable in a boom). I will assume that consumers are more price sensitive in recessions (*i.e.* $\bar{\beta}_R > \bar{\beta}_B$), which is intuitive and also line with empirical findings of Gordon, Goldfarb and Li (2013). The actual realizations of $\bar{\beta}_t$ will however remain unobserved to the sellers in the model (making it a two-state hidden Markov problem), so they have to learn which regime the economy is in.

Related to this learning process, the two-state Markov assumption implies that beliefs are conveniently characterized by only one state variable: $\lambda_{i,t|t-1} = \Pr_i(\bar{\beta}_t = \bar{\beta}_R)$, which is seller i 's prior time t belief that the average price elasticity is high. This limits the dimensionality curse associated with the solution method, while still capturing the notion of aggregate risk (see Ellison and Valla (2001) for a similar approach).

The cost of introducing Markov switching is that it is no longer possible to assume that idiosyncratic shocks $\eta_{i,t}$ enter a persistent $AR(1)$ -process. The reason is that Bayesian updating in a mixed $AR(1)$ /Markov switching model no longer admits an analytical solution. Instead, simulation methods such as the Gibbs-sampler must be employed (*cf.* Kim and Nelson (1999)), which is not feasible within this model's framework.

Consequently, I proceed by assuming that all persistence comes through the mean value for β . Since idiosyncratic movements in the slope parameter would cancel out in the aggregate anyway (irrespective of whether they are persistent or not), and since this specification still contains an active learning motive with respect to the slope term, intuition suggests that this assumption does not seriously limit the generality of the aggregate simulation that is to follow.

With respect to the persistence of the mean value for β , the parameter π expresses the conditional probability of staying in the current regime (so $\pi = \Pr(\bar{\beta}_{t+1} = \bar{\beta}_J | \bar{\beta}_t = \bar{\beta}_J)$, $J \in \{R, B\}$). In the simulation I set $\pi = 0.998$ (as in Matějka (2010a)), which implies that there is a 10 percent probability that the aggregate state changes in any given year.

Summarizing, the process for the slope parameter can be described by:

$$\beta_{i,t} = \bar{\beta}_t + \eta_{i,t}, \text{ with } \eta_{i,t} \sim N(0, \sigma_\eta^2) \quad (11)$$

$$\bar{\beta}_t \in \{\bar{\beta}_R, \bar{\beta}_B\} \quad (12)$$

By working on Bayes' rule, one can show that the belief updating process now depends on the relative probability of observing the actual signal in either of the two regimes:

$$\lambda_{i,t|t} = \frac{\lambda_{i,t|t-1} P(q_{i,t} | \bar{\beta}_t = \bar{\beta}_R)}{\lambda_{i,t|t-1} P(q_{i,t} | \bar{\beta}_t = \bar{\beta}_R) + (1 - \lambda_{i,t|t-1}) P(q_{i,t} | \bar{\beta}_t = \bar{\beta}_B)}, \quad (13)$$

with $P(q_{i,t}|\bar{\beta}_t = \bar{\beta}_J) = \exp\left(-0.5[q_{i,t} - \bar{\alpha} + \beta_J r_{i,t} - \gamma q_{i,t-1}]^2 / [\sigma_\mu^2 + \sigma_\epsilon^2]\right)$, $J \in \{R, B\}$.

Since there is a $(1 - \pi)$ probability that the regime switches in going from the end of period t to the beginning of period $t + 1$, we have that:

$$\lambda_{i,t+1|t} = \pi \lambda_{i,t|t} + (1 - \pi) (1 - \lambda_{i,t|t}) \quad (14)$$

Compared to Section 3, the other elements of the setup (as well as the calibration) are left unchanged so I do not repeat them here.

The results of this exercise are shown in Figure 5. That figure shows the posted price averaged over 10,000 sellers. During the simulation, all individual sellers were hit by two types of disturbances: idiosyncratic shocks (all sellers have their individual series for $\beta_{i,t}$, for example representing the popularity of the particular brand they are selling) and aggregate ones (all these individual series for $\beta_{i,t}$ evolve around a common mean slope of $\bar{\beta}$, capturing aggregate economic conditions). For the first 9 weeks, the common mean value of $\beta_{i,t}$ is set equal to 2.25. Then, in week 10, a positive aggregate shock to the economy suddenly makes all consumers less price sensitive, as a result of which $\bar{\beta}$ falls to 1.45.²²

Several things are to be noted from the figure. First, observe that the aggregate price level evolves smoothly in the absence of aggregate shocks. The reason is that all experimentation is idiosyncratic to each individual seller, as a result of which it cancels out in the aggregate. Although this is probably an exaggeration of the amount of idiosyncrasy present in reality, the dataset analyzed in Nakamura (2008) does hint in this direction, as she reports that only 16 percent of all price variation is common to different stores selling an identical product.

However, when the aggregate shock hits in week 10, the price level does respond in all three specifications. In particular, it goes up, which reflects the increase in market power for sellers resulting from the decrease in the average price sensitivity of consumers. Next to that, all three models also produce a hump-shaped response to the shock. This is due to the existence of good-specific habits: consumers are "addicted" to the specific good they have been consuming in the past, as a result of which sellers temporarily obtain the power to charge prices that lie above their new long-run values.

The persistence of this hump however differs between the three models. In particular, the fact that the perfect information line in the upper panel of Figure 4 converges rather

²²The size of the imposed shock is unrealistically large given the fact that it is assumed to take place within a week, but considering a more gradual change from $\bar{\beta}_R$ to $\bar{\beta}_B$ would only increase the sluggishness in the responses of aggregate prices - thereby strengthening the findings that emerge from Figure 5.

Figure 5: Average posted price (top panel) and average beliefs (bottom panel) under active learning, passive learning and under perfect information.

quickly to its new, higher value shows that habit formation in itself (that is: without help of informational imperfections) does not generate much persistence. As the passive learning response shows, assuming that the true values of $\beta_{i,t}$ and $\bar{\beta}_t$ cannot be observed, helps in this respect. The reason is that the learning process implies that firms only gradually find out that the mean value of the slope parameter has changed. Because a regime change is not that likely *a priori*, sellers initially attribute the "anomalous" observations they receive to transitory factors to which they do not wish to respond. Only as they continue to receive observations that are difficult to reconcile with the old regime, they gradually become convinced that a regime switch has occurred.

The passive learning response was however generated under the assumption that learning takes place in the suboptimal passive manner. That is: firms post the certainty equivalent price and update their beliefs with whatever information this pricing strategy happens to generate.

Under active learning in contrast, firms internalize this learning process and consciously use their pricing strategy to this end via experimentation. Consequently, learning is faster compared to the passive case (see the lower panel of Figure 5), as a result of which

the aggregate shock is propagated to a somewhat lesser extent. However, average beliefs still need time to converge to the new value of $\bar{\beta}$, so the aggregate price level continues to move in an inertial way. Hence, one can conclude that the active learning process of sellers who are trying to learn the slope of their demand curve, is not only able to replicate the volatile, discrete pattern followed by individual prices, but is also consistent with the idea that the aggregate price level responds sluggishly to aggregate shocks. We can thus conclude that individual price flexibility does not translate into aggregate price flexibility when individual price changes are driven by the experimentation motive of a seller who is learning the slope of his demand curve.

6 Discussion

Active learning problems rapidly become quite challenging to solve. In order to advance, this paper therefore had to abstract from several interesting issues.

First, the analysis in this paper is only partial equilibrium in nature. In this respect, contemporaneous work by Bachmann and Moscarini (2012) forms an important contribution. In order to generate "endogenous uncertainty", they manage to place the experimenting seller in a dynamic stochastic general equilibrium environment. Although this requires them to simplify in some dimensions (they for example reduce the problem to a discrete two-state setting, where the price elasticity of demand is either high or low), the main idea underlying the present paper should carry over to their setup. Investigating whether it is possible to construct a general equilibrium active learning model that is consistent with the behavior of prices at both the micro and macro level, would therefore be a logical next step along this line of research. Subsequently, one could investigate to what extent the experimenting seller model is able to replicate the quantitative aspects of the data.

Second, this paper simplifies by assuming that firms are using their pricing strategy to learn about only one of the parameters of the demand curve. This avoids the dimensionality curse that is inherent to the solution method.²³ In particular, I have assumed that firms are only using their pricing strategy to facilitate the learning process on their slope parameter. The full problem, in which firms are simultaneously learning on both

²³The problem in which sellers are actively learning about both the intercept and slope of the demand curve would have six state variables: $q_{i,t-1}$, two estimated means, two variances and one covariance estimate. At a weekly frequency (which implies a discount factor close to one), the model becomes very difficult to solve. Potentially, such a bigger problem could be tackled with faster algorithms that approximate the original problem (such as the one in Cosimano (2008)).

slope and intercept, is studied in Wieland (2006) - although in a very different context: he analyzes a central bank that is trying to learn the Phillips curve in an active manner. There he actually finds that the extent of experimentation is *greater* if learning takes place with respect to both intercept and slope, rather than just with respect to the slope (although the differences are small).²⁴ This suggests that the results obtained in the present paper are robust to the more general setting in which firms are learning simultaneously on multiple parameters of the demand curve, although this of course remains to be verified for the particular problem under consideration.

Third, I see this paper's sale motive as only part of the complete story. After all, generating information on demand is unlikely to be the only reason for firms to hold sales. It is however important to note that the experimentation motive is fully complementary to other sale motives offered by the literature (such as Matějka (2010a,b), Guimarães and Sheedy (2011) and Nakamura and Steinsson (2011)), so the channel emphasized by this paper is by no means at odds with these earlier contributions. In addition, this paper abstracts from potential forces at work in reality that make firms want to abstain from price changes. Menu costs are an obvious example, while there may also exist information-related reasons for this (see *e.g.* L'Huillier (2011) where sellers want to delay price adjustments to keep consumers uninformed).

Fourth, consumers behave rather passively in this paper (the consumer side is simply represented by the demand curve (2)). Consequently, consumers do not take into account that sellers have an experimentation motive. If consumers would be aware of this, strategic interactions between the two sides of the market might arise - the analysis of which I leave for future work.

Fifth, I do not consider changes to marginal costs. Eichenbaum, Jaimovich and Rebelo (2011) suggest that the latter are important, as price changes very often coincide with cost changes in their data set. Although there does not seem to be a consensus view on this issue yet,²⁵ adding cost shocks to the model could be an interesting next step. This modification would not eliminate the experimentation motive with respect to the slope term, but if costs are imperfectly observed as well, it does give the seller an incentive to produce more units as that would improve the precision of his cost signal (see Gal-Or

²⁴Because Wieland (2006) is not interested in the higher frequency features of the data, he can calibrate his model at a yearly frequency. This allows him to set the discount factor equal to 0.95, which greatly facilitates the convergence of the value function. Nevertheless, he reports that solving his model still takes 3 days on a 2.21 GHz AMD processor with 1.48 GB of RAM.

²⁵Nakamura's (2008) results for example suggest that cost shocks only play a minor role in explaining high-frequency price movements; similarly, Chevalier and Kashyap (2011) document that prices in Dominick's data set are also often adjusted without movements in costs.

(1988)). In addition, changes in marginal costs might also produce shifts in the pricing band, just as we observe in the data where the high and/or low price occasionally settle at new levels.

Finally, since both this paper’s approach (recall footnote 13) as well as its results are rather similar to that of certain papers in the rational inattention literature, it may be worth pointing out some interesting differences. First, the model developed in this paper does not predict a sharp distinction between a seller’s response to aggregate and idiosyncratic shocks. This contrasts with Maćkowiak and Wiederholt (2009), where sellers allocate most of their attention to idiosyncratic conditions. As a result, sellers in their model quickly react to changes in the latter, while responding rather sluggishly to aggregate developments. Although Boivin, Giannoni and Mihov (2009) and Maćkowiak, Moench and Wiederholt (2009) present evidence that this indeed is the case in reality, the recent papers by De Graeve and Walentin (2011) and Carlsson and Skans (2012) question this. They find that firms respond sluggishly to *both* aggregate *and* idiosyncratic conditions - a finding that is more in line with a learning model in which sellers face uncertainty on the persistence of both shock types. Hopefully, future work can shed more light on these contrasting views.

Second, compared to the rational inattention approach, active learning imposes somewhat more discipline on the model builder: since active learners behave in a fully rational way, the model does not contain a free parameter whose value is difficult to calibrate (like the information processing capacity in rational inattention models). Instead, in this paper’s setup, the costs of more precise signals are determined *endogenously* by the amount of profits a seller has to give up to experiment. Despite the absence of such a degree of freedom, the model is still able to produce rather similar results as the rational inattention models in this literature.

7 Conclusion

Most standard price setting models make the rather strong assumption that sellers can observe their entire demand curve. This paper relaxes this assumption by making the demand curve unobserved and analyzes the optimal pricing strategy of a seller who operates in such an environment. It is shown that when learning is modeled in an optimal way, it can not only match the behavior of aggregate variables (a point that has been made before already), but learning is also able to replicate certain micro-dimensions of the data. In particular, as achieved by Matějka (2010a,b), the model constructed in this

paper manages to replicate all of the following three observations that seem present in reality:

1. Individual prices change frequently.
2. Individual prices tend to move back and forth between only a few rigid values.
3. The aggregate price level responds gradually to aggregate shocks.

Note that the model’s ability to replicate these elements is *not* due to some form of irrationality: all results are driven by the rational behavior of sellers who are responding optimally to the fact that they cannot observe their entire demand curve. Given that more and more retailers are using sophisticated pricing optimization software (of which this paper’s model is essentially a mini-version; also recall footnote 3), rational price setting can not only be thought of as an important optimal benchmark that approximates the real economy and is worth analyzing (like rational expectations), but also as actually materializing to a first degree in practice.²⁶

The model can match observation #1 because the introduction of active learning gives firms an incentive to experiment with their price. Due to the interaction with the existence of good-specific habits, this produces price volatility. Next to the fact that price experimentation is optimal for sellers facing uncertainty on the slope of their demand curve, there is also empirical evidence that this type of behavior plays a role in reality.²⁷ Experimentation could also be the reason why prices are more volatile than marginal costs in the data. It simultaneously nuances the point of view (expressed by Golosov and Lucas (2007), among many others) that the large idiosyncratic variations in prices observed in the data indicate the presence of large idiosyncratic *shocks*; instead, this paper shows that these large movements could also stem from the *endogenous* experimentation motive of individual sellers.

In addition, the experimenting seller model can go a long way in reproducing the discreteness of individual prices observed in reality (observation #2). Replicating this feature of the data has proved to be a major challenge to most price setting models. Although the discreteness arising in this paper is not exact, it looks very much like

²⁶In this respect, one can also take the idea of this paper one step further and develop a new class of macroeconomic models with *actual* pricing algorithms at its core. Since these algorithms determine prices for such a large share of the economy, it may be interesting to investigate how they interact with each other and how the prices they set respond to typical economic shocks.

²⁷See Pashigian and Bowen (1991, who documented a positive impact of demand uncertainty on the frequency of sales), Campbell and Eden (2010, who analyzed US scanner data), Einav *et al.* (2011, who looked at eBay-listings), and Gaur and Fisher (2005, who held a survey among large US retailers).

it. The reason is that the experimentation motive shapes the price setting rules like a step function: they have a high price platform (associated with high demand) and a low one (associated with low demand), separated by a sharp kink. Consequently, the model displays a form of endogenous, local price stickiness (despite the absence of price adjustment costs): as long as a shock does not move a seller from the high to the low platform (or vice versa), it is optimal for him to keep his price rather fixed - even if demand varies locally.

Finally, notwithstanding the flexibility of individual prices under active learning, the aggregate price level still responds sluggishly to shocks. This is due to the fact that average beliefs change only gradually over time. Consequently, the model is also able to capture the inertia found in prices at the aggregate (observation #3).²⁸ To the extent that price changes at the individual level are driven by the experimentation motive of a seller who is actively learning the slope of his demand curve, the flexibility of prices at the micro stage thus does not translate into aggregate price level flexibility.

8 Appendix: proofs

Lemma 1 *Uncertainty on β is decreasing in the squared deviation of the control variable from zero, that is: $\partial \Sigma_{i,t|t}^b / \partial (r_{i,t}^2) < 0$.*

Proof. First define:

$$\begin{aligned} u &\equiv \Sigma_{i,t|t-1}^b - \frac{\left(\Sigma_{i,t|t-1}^b\right)^2 r_{i,t}^2}{r_{i,t}^2 \Sigma_{i,t|t-1}^b + \sigma_\mu^2 + \sigma_\epsilon^2} \\ v &\equiv r_{i,t}^2 \end{aligned}$$

Differentiating these with respect to $r_{i,t}$ yields:

$$\begin{aligned} \frac{\partial u}{\partial r_{i,t}} &= - \frac{2r_{i,t} \left(\Sigma_{i,t|t-1}^b\right)^2 (\sigma_\mu^2 + \sigma_\epsilon^2)}{\left(r_{i,t}^2 \Sigma_{i,t|t-1}^b + \sigma_\mu^2 + \sigma_\epsilon^2\right)^2} \\ \frac{\partial v}{\partial r_{i,t}} &= 2r_{i,t} \end{aligned}$$

²⁸In conversations, I always liked to simplify the message of this paper as follows: this paper presents a model of "active learning". The "active"-part of the setup delivers the volatility and (near-)discreteness of individual prices, while the fact that there is still some "learning" going on, generates the sluggish responses of prices to aggregate shocks.

Via the chain rule we then obtain that:

$$\frac{\partial \Sigma_{i,t|t}^b}{\partial (r_{i,t}^2)} = \frac{\partial u}{\partial v} = \frac{\partial u}{\partial r_{i,t}} \cdot \frac{\partial r_{i,t}}{\partial v} = - \frac{\left(\Sigma_{i,t|t-1}^b \right)^2 (\sigma_\mu^2 + \sigma_\epsilon^2)}{\left(r_{i,t}^2 \Sigma_{i,t|t-1}^b + \sigma_\mu^2 + \sigma_\epsilon^2 \right)^2},$$

which is strictly negative. ■

Lemma 2 *The value of the firm is non-increasing in the uncertainty on β , that is:*
 $\Delta V \left(q_{i,t-1}, b_{i,t|t-1}, \Sigma_{i,t|t-1}^b \right) / \Delta \Sigma_{i,t|t-1}^b \leq 0.$

Proof. Fix the optimal policy sequence $r_{i,t+j}(\Sigma_{i,t|t-1}^b)$ for all $j \geq 0$. Now decrease $\Sigma_{i,t|t-1}^b$ to any $\Sigma^b < \Sigma_{i,t|t-1}^b$ leaving the $r_{i,t+j}$ -sequence unchanged. Call the associated (non-optimized) value function $W \left(q_{i,t-1}, b_{i,t|t-1}, \Sigma^b, r(\Sigma_{i,t|t-1}^b) \right)$ (where the last argument expresses that the $r_{i,t+j}$ -sequence continues to be based upon the old level of uncertainty $\Sigma_{i,t|t-1}^b$). Now compare this value to the optimized value function $V \left(q_{i,t-1}, b_{i,t|t-1}, \Sigma_{i,t|t-1}^b \right)$.

To do so, first observe from the profit function in equation (1) that current period profits are not affected by the change in uncertainty. Furthermore note that the change in uncertainty does not affect future pay-offs either, since the firm continues to set its control based upon the old level of uncertainty $\Sigma_{i,t|t-1}^b$. Consequently, $W \left(q_{i,t-1}, b_{i,t|t-1}, \Sigma^b, r(\Sigma_{i,t|t-1}^b) \right) = V \left(q_{i,t-1}, b_{i,t|t-1}, \Sigma_{i,t|t-1}^b \right)$. In addition, optimization implies that $V \left(q_{i,t-1}, b_{i,t|t-1}, \Sigma^b \right) \geq W \left(q_{i,t-1}, b_{i,t|t-1}, \Sigma^b, r(\Sigma_{i,t|t-1}^b) \right)$. Putting these two results together tells us that $V \left(q_{i,t-1}, b_{i,t|t-1}, \Sigma^b \right) \geq V \left(q_{i,t-1}, b_{i,t|t-1}, \Sigma_{i,t|t-1}^b \right)$, so firm value is non-increasing in uncertainty on β . ■

Proposition 1 *The relative price set under active learning (call it $r_{i,t}^\dagger$) is always more extreme than the price chosen under the passive learning strategy, i.e. $|r_{i,t}^\dagger| > |r_{i,t}^\diamond|$.*

Proof. First, consider the case for which $r_{i,t}^\diamond > 0$. Take any $r_{i,t}^\bullet < r_{i,t}^\diamond$. Note that this cannot be optimal since $r_{i,t}^\diamond$ maximizes firm-value absent learning considerations, while $r_{i,t}^\diamond$ is also more valuable for future periods by Lemmas 1 and 2 (since it is more extreme). Hence, $r_{i,t}^\dagger \geq r_{i,t}^\diamond$ is required for optimality.

By an envelope theorem argument it then follows that $r_{i,t}^\dagger > r_{i,t}^\diamond$. After all, starting from the optimum at $r_{i,t}^\diamond$, variations in $r_{i,t}$ only affect firm-value to a second order, while increasing $r_{i,t}$ increases the value of the firm to a first order by Lemmas 1 and 2.

The proof for the case where $r_{i,t}^\diamond < 0$ follows analogously, *mutatis mutandis*. ■

9 References

- Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien. 1991. "Optimal Learning by Experimentation." *Review of Economic Studies* 58 (4): 621-54.
- Araman, Victor F., and René Caldentey. 2009. "Dynamic Pricing for Nonperishable Products with Demand Learning." *Operations Research* 57 (5): 1169-88.
- Bachmann, Rüdiger, and Giuseppe Moscarini. 2012. "Business Cycles and Endogenous Uncertainty." Mimeo, Yale University.
- Balvers, Ronald J., and Thomas F. Cosimano. 1990. "Actively Learning about Demand and the Dynamics of Price Adjustment." *Economic Journal* 100 (402): 882-98.
- Balvers, Ronald J., and Thomas F. Cosimano. 1994. "Inflation Variability and Gradualist Monetary Policy." *Review of Economic Studies* 61 (4): 721-38.
- Beck, Günter W., and Volker Wieland. 2002. "Learning and Control in a Changing Economic Environment." *Journal of Economic Dynamics and Control* 26 (9-10): 1359-77.
- Bertocchi, Graziella, and Michael Spagat. 1998. "Growth Under Uncertainty With Experimentation." *Journal of Economic Dynamics and Control* 23 (2): 209-31.
- Bils, Mark, and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112 (5): 947-85.
- Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni. 2013. "News, Noise, and Fluctuations: An Empirical Investigation." *American Economic Review* 103 (7): 3045-70.
- Boivin, Jean, Marc P. Giannoni, and Ilian Mihov. 2009. "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data." *American Economic Review* 99 (1): 350-84.
- Bomfim, Antulio N. 2001. "Measurement Error in General Equilibrium: The Aggregate Effects of Noisy Economic Indicators." *Journal of Monetary Economics* 48 (3): 585-603.
- Campbell, Jeffrey R., and Benjamin Eden. 2010. "Rigid Prices: Evidence from US Scanner Data." Federal Reserve Bank of Chicago Working Paper No. 2005-08.
- Carlsson, Mikael, and Oskar N. Skans. 2012. "Evaluating Microfoundations for Aggregate Price Rigidities: Evidence from Matched Firm-Level Data on Product Prices and Unit Labor Cost." *American Economic Review* 102 (4): 1571-95.
- Chari, V.V., Patrick Kehoe, and Ellen McGrattan. 2009. "New Keynesian Models: Not Yet Useful for Policy Analysis." *American Economic Journal: Macroeconomics* 1 (1), 242-66.
- Chevalier, Judith A., and Anil K. Kashyap. 2011. "Best Prices." Chicago Booth

Working Paper No. 11-01.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?", in: John Taylor and Michael Woodford (*eds.*), *Handbook of Macroeconomics*, Vol. 1A, Amsterdam, North-Holland: Elsevier: 65-148.

Cosimano, Thomas F. 2008. "Optimal Experimentation and the Perturbation Method in the Neighborhood of the Augmented Linear Regulator Problem." *Journal of Economic Dynamics and Control* 32 (6): 1857-94.

De Graeve, Ferre, and Karl Walentin. 2011. "Stylized (Arte)Facts on Sectoral Inflation." Sveriges Riksbank Working Paper No. 254.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. 2008. "Reference Prices and Nominal Rigidities." CEPR Working Paper No. 6709.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. 2011. "Reference Prices and Nominal Rigidities." *American Economic Review* 101 (1): 234-62.

Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. 2011. "Learning from Seller Experiments in Online Markets." Mimeo, Stanford University.

Ellison, Martin, and Natacha Valla. 2001. "Learning, Uncertainty and Central Bank Activism in an Economy With Strategic Interactions." *Journal of Monetary Economics* 48 (1): 153-71.

Gabaix, Xavier. 2011. "A Sparsity-Based Model of Bounded Rationality." Mimeo, NYU.

Gal-Or, Esther. 1988. "The Advantages of Imprecise Information." *RAND Journal of Economics* 19 (2): 266-75.

Gaur, Vishal, and Marshall L. Fisher. 2005. "In-Store Experiments to Determine the Impact of Price on Sales." *Production and Operations Management* 14 (4): 377-87.

Gertler, Mark, and John Leahy. 2008. "A Phillips Curve with an Ss Foundation." *Journal of Political Economy* 116 (3): 533-72.

Golosov, Mikhail, and Robert E. Lucas. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115 (2): 171-99.

Gordon, Brett, Avi Goldfarb, and Yang Li. 2013. "Does Price Elasticity Vary with Economic Growth? A Cross-Category Analysis." *Journal of Marketing Research* 50 (1): 4-23.

Grossman, Sanford J., Richard E. Kihlstrom, and Leonard J. Mirman. 1977. "A Bayesian Approach to the Production of Information and Learning By Doing." *Review of Economic Studies* 44 (3): 533-47.

- Guimarães, Bernardo, and Kevin Sheedy. 2011. "Sales and Monetary Policy." *American Economic Review* 101 (2): 844-76.
- Kehoe, Patrick J., and Virgiliu Midrigan. 2010. "Prices Are Sticky After All." NBER Working Paper No. 16364.
- Keller, Godfrey, and Sven Rady. 1999. "Optimal Experimentation in a Changing Environment." *Review of Economic Studies* 66 (3): 475-507.
- Kendrick, David A., Hans M. Amman, and Marco P. Tucci. 2011. "Learning about Learning in Dynamic Economic Models." *Handbook of Computational Economics*, forthcoming.
- Kiefer, Nicholas M. 1989. "A Value Function Arising in the Economics of Information." *Journal of Economic Dynamics and Control* 13 (2): 201-23.
- Kiefer, N.M., and Yaw Nyarko. 1989. "Optimal Control of an Unknown Linear Process with Learning." *International Economic Review* 30 (3): 571-86.
- Kim, Chang J., and Charles R. Nelson. 1999. *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, Cambridge: MIT Press.
- Klemperer, Paul. 1995. "Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade." *Review of Economic Studies* 62 (4): 515-39.
- L'Huillier, Jean-Paul. 2011. "Consumers' Imperfect Information and Price Rigidities." Mimeo, EIEF.
- Lucas, Robert E. 1973. "Some International Evidence on Output-Inflation Trade-offs." *American Economic Review* 63 (3): 326-34.
- Maćkowiak, Bartosz, Emanuel Moench, and Mirko Wiederholt. 2009. "Sectoral Price Data and Models of Price Setting." *Journal of Monetary Economics* 56 (S): 78-99.
- Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99 (3): 769-803.
- Matejka, Filip. 2010a. "Rationally Inattentive Seller: Sales and Discrete Pricing." Mimeo, CERGE-EI.
- Matejka, Filip. 2010b. "Rigid Pricing and Rationally Inattentive Consumer." Mimeo, CERGE-EI.
- Matejka, Filip and Christopher A. Sims. 2010. "Discrete Actions in Information-Constrained Tracking Problems." Mimeo, Princeton University.
- Nakamura, Emi. 2008. "Pass-Through in Retail and Wholesale." *American Economic Review: Papers and Proceedings* 98 (2): 430-37.
- Nakamura, Emi, and Jón Steinsson. 2008. "Five Facts about Prices: A Reevaluation

- of Menu Cost Models." *Quarterly Journal of Economics* 123 (4): 1415-64.
- Nakamura, Emi, and Jón Steinsson. 2010. "Monetary Non-Neutrality in a Multisector Menu Cost Model." *Quarterly Journal of Economics* 125 (3): 961-1013.
- Nakamura, Emi, and Jón Steinsson. 2011. "Price Setting in Forward-Looking Customer Markets." *Journal of Monetary Economics* 58 (3): 220-33.
- Nimark, Kristoffer. 2008. "Dynamic Pricing and Imperfect Common Knowledge." *Journal of Monetary Economics* 55 (2): 365-82.
- Pashigian, B. Peter, and Brian Bowen. 1991. "Why Are Products Sold on Sale?: Explanations of Pricing Regularities." *Quarterly Journal of Economics* 106 (4): 1015-38.
- Prescott, Edward C. 1972. "The Multi-Period Control Problem under Uncertainty." *Econometrica* 40 (6): 1043-58.
- Ravn, Morten O., Stephanie Schmitt-Grohé, and Martin Uribe. 2006. "Deep Habits." *Review of Economic Studies* 73 (1): 195-218.
- Rothschild, Michael. 1974. "A Two-Armed Bandit Theory of Market Pricing." *Journal of Economic Theory* 9 (2): 185-202.
- Sargent, Thomas J. 1999. *The Conquest of American Inflation*. Princeton and Oxford: Princeton University Press.
- Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665-90.
- Smets, Frank, and Raf Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586-606.
- Svensson, Lars E.O., and Noah Williams. 2007. "Bayesian and Adaptive Optimal Policy under Model Uncertainty." Mimeo, Princeton University.
- Trefler, Daniel. 1993. "The Ignorant Monopolist: Optimal Learning with Endogenous Information." *International Economic Review* 34 (3): 565-81.
- Wieland, Volker. 2000a. "Learning by Doing and the Value of Optimal Experimentation." *Journal of Economic Dynamics and Control* 24 (4): 501-34.
- Wieland, Volker. 2000b. "Monetary Policy, Parameter Uncertainty and Optimal Learning." *Journal of Monetary Economics* 46 (1): 199-228.
- Wieland, Volker. 2006. "Monetary Policy under Uncertainty about the Natural Unemployment Rate: Brainard-Style Conservatism versus Experimental Activism." *Advances in Macroeconomics* 6 (1): Article 1.
- Zellner, Arnold. 1971. *Introduction to Bayesian Inference in Econometrics*, Wiley.