

Control costs, rational inattention, and retail price dynamics*

James Costain Anton Nakov
Banco de España ECB and CEPR

March 30, 2021

Abstract

This paper models the adjustment of regular and sales prices in retail microdata under the assumption that precise decision-making is costly. We build on two related approaches: “control costs” and “rational inattention”. We adapt and extend the equivalence results proved by [Steiner et al. \(2017\)](#) to the retail pricing context in order to build a tractable control cost model that approximates the solution of the rational inattention model.

We calibrate the approximate rational inattention model to microdata from supermarkets to study retail price dynamics. Using our simulated model to diagnose the roles of stochastic price discrimination versus the inherent discreteness of the rational inattention solution as drivers of retail prices, we find that the former helps explain the “jumpiness” of retail sales, while the latter helps explain the “stickiness” of nominal prices and of nominal price points. Our model also suggests that limited memory, in addition to limited information processing capability, may play a role in nominal rigidity.

Keywords: State-dependent pricing, retail sales, control costs, rational inattention, limited memory, stochastic price discrimination

JEL Codes: E31, D83, C73

*Thanks to Fernando Álvarez, Marlon Azinović, Luca Dedola, Wouter den Haan, Tetsuya Hoshino, Samuel Hurtado, Bartosz Mackowiak, Lilia Maliar, Filip Matejka, and participants in the ESCB PRISMA network and the CEF2019 and ESWC2020 meetings for helpful comments. We particularly thank Luca Riva and Federico Rodari for assistance with data work and for helpful discussions. Views expressed here are those of the authors and do not necessarily coincide with those of the Bank of Spain, the Eurosystem, the ECB, or the CEPR.

1 Introduction

Retail prices present economists with a paradox: they seem both *sticky* and *jumpy*. On one hand, prices often remain fixed at exactly the same nominal value for extended periods of time, even after economic conditions change; nominal rigidity in the face of monetary policy shocks plays a central role in many theories of monetary economics. On the other hand, price changes are often large, but frequently also transitory, jumping back after a big, brief decrease or increase. Remarkably, stickiness and jumpiness seem linked: retail prices often jump back and forth between two or more “price points” that may each recur many times at exactly the same nominal value (see Sec. 3 for examples). The extensive literature addressing this paradox explains it in two main ways. First, excess short-run volatility may be a response to heterogeneous demand. When demand elasticities are sufficiently diverse across customers, a retailer’s profit function may not be concave, making it more profitable to randomize between high and low prices than to set a single intermediate price (Varian 1980; Guimaraes and Sheedy 2011). Second, the “rational inattention” literature (Sims 2003) has argued that retailers may economize on information costs by randomizing between several discrete price levels instead of adjusting precisely in response to every shock (Matějka 2016; Stevens 2019).¹

This paper seeks to evaluate the importance of these two explanations of retail price dynamics by building a model that incorporates both mechanisms. To do so, it builds on two existing modelling frameworks: the “control costs” framework from game theory (Van Damme 2002), and the “rational inattention” theory of Sims (2003). Steiner et al. (2017) showed that these two approaches are closely related; we extend their results to allow for limited memory as well as limited information processing capacity. Applying our limited-memory framework to the retail context yields a rich but tractable model of price dynamics that extends our previous model of a single sticky price (Costain and Nakov 2019) to allow for jumps between multiple sticky price points, such as “regular” and “sales” prices. Besides shedding light on the economic mechanisms underlying the stickiness and jumpiness of nominal prices, and on their monetary policy implications, our framework also leads to new computational methods for the numerical solution of models of near-rational choice. We now discuss each of these points in turn.

The framework developed in this paper is a natural extension of our previous work, which showed that a model where making precise decisions is costly (a “control cost” model) helps explain why prices are sticky, and how this stickiness affects the macroeconomy. If setting exactly the right price requires costly effort, then retail firms will sometimes leave their prices

1. A third set of papers assumes that price adjustment involves two different types of fixed costs: a small cost for each price change, and another fixed cost for selecting a list of possible future prices from which to choose (e.g. Eichenbaum et al. 2011; Kehoe and Midrigan 2015). While imposing this cost structure can generate sticky and jumpy prices, it does little to explain the underlying economic mechanisms.

unchanged (Costain and Nakov 2015). If choosing exactly the right moment to adjust the price requires costly effort, then firms’ reaction to monetary policy shocks will not be fully synchronized, weakening the “selection effect” and hence strengthening the real effects of the shock (Costain and Nakov 2019). The present paper adds an extra decision step to our earlier model: besides choosing whether or not to adjust its price, the firm also decides whether to revisit a price it remembers setting previously. In this way, it generalizes our earlier model of a sticky scalar, making it a model of a sticky vector. Additionally, in the current paper, we take advantage of the results of Steiner et al. (2017) that formalize the relation between rational inattention and control costs. Rational inattention theory points out that reducing a choice across a large set of alternatives to a choice over a smaller set may economize on information costs (Matějka 2016; Stevens 2019). Therefore, intuitively, it may be efficient for the firm to consider repeating one of the prices it has set recently, before considering other alternatives. Rational inattention theory also informs us about how to parameterize our control cost model, given the empirical frequencies of events that have counterparts in the data (adjusting or not adjusting the price, and returning to a recent price versus introducing a new price point).

Solving rational inattention models, as these are usually defined in existing literature, is typically very challenging. Even though these models impose a constraint on the flow of information to the decision maker (DM), they impose no constraint on the DM’s memory, which makes the DM’s state variable a high-dimensional object. As an alternative that relates directly to the empirical context we are studying, we propose a framework we call *short-term memory rational inattention* (STMRI), in which the DM faces a constraint on information flow, and furthermore, only remembers the signals observed in the last τ periods, for some integer $\tau \geq 0$. We show (Prop. 4) that the solution of an STMRI model takes precisely the same form as that of an RI model, and that the STMRI solution converges to RI as $\tau \rightarrow \infty$. Part of the appeal of the STMRI framework is that limited memory may be more realistic than infinite memory; but regardless, it facilitates computation by reducing the DM’s state variable from an infinite sequence to a finite vector.

Besides its computational advantages, STMRI leads naturally to a theory of multiple sticky price points. Our simulations help distinguish between the economic mechanisms underlying “stickiness” and “jumpiness”: they suggest that the former is explained primarily by costs associated with decision-making, while the latter is primarily driven by heterogeneous demand elasticity. Rational inattention can make prices jump back and forth across discrete values, but if the demand curve has constant elasticity these jumps will roughly track the underlying shocks, making extremely rapid fluctuations like those observed at many retail firms implausible. On the other hand, randomization in response to customer heterogeneity does not imply that prices should repeatedly revisit the same nominal points. Indeed, even rational inattention need not imply that prices revisit exactly the same *nominal* prices. But

rational inattention can indeed explain sticky nominal price points, when we apply it to the appropriate choice set.

To understand what we mean by this, note that if the firm’s choice set at each point in time were simply the set of all possible nonnegative prices, then rational inattention would lead to *real* price stickiness rather than nominal price stickiness (see Example 2 in Sec. 4). But rational inattention models can be applied to choice sets of arbitrary form. Therefore, our model recognizes that leaving the current nominal price unchanged is qualitatively different from setting any other nominal price, because it is the outcome that occurs even if the firm fails to make any choice at all. Therefore, it should be treated as a separate, discrete alternative within the choice set. Likewise, we distinguish the option to return to a remembered price from the alternative of setting any other price. Since these are qualitatively different actions for the DM, their measure relative to other elements of the choice set is a free parameter, which can be deduced from the data (by revealed preference) if the rational inattention model is true. In other words, when modelling detailed dynamics in microdata under the RI assumption, one should not (can not) test RI alone. Instead, RI must be tested jointly with a hypothesis about what choice set is under consideration by the decision-maker. This should not be viewed as a negative conclusion, because the economic context makes it relatively simple to formulate a reasonable conjecture about the nature of the choice set.

We end this introductory section with a brief overview of related literature. In Section 2, we formally define control cost models and rational inattention models. We review some key analytical results about how these models behave, and we define a new class of rational inattention models with limited memory. In Sec. 3, we discuss stylized facts about retail price dynamics, focussing especially on “jumpiness” across recurrent nominal price points. Sec. 4 builds a model of “regular” price dynamics, abstracting from retail sales. Sec. 5 uses our limited-memory framework to model “sales” prices as well as regular prices, and presents numerical results to distinguish between the role of costly decision-making and the role of heterogeneous demand elasticities in explaining price dynamics. Sec. 6 studies the general equilibrium impact of monetary policy shocks ((TO BE COMPLETED)) on price behavior, including sales frequency, and on the real economy. Finally, Sec. 7 concludes.

1.1 Related literature

This paper builds closely on several previous studies that model price stickiness as the result of costly decisions. Firstly, it extends the authors’ previous control cost model (Costain and Nakov 2019) — which studied the intermittent adjustment of a single sticky price — to encompass retail sales too, by assuming that the DM remembers some previously set prices, and entertains the possibility of jumping back to some price it has set before. Another predecessor is Matějka (2016), which points out that rational inattention typically reduces

a large or even infinite choice set to a smaller, discrete and finite set of actions that are actually chosen with positive probability. Matejka offers this fact as a possible explanation of price stickiness, but (as we will show) his model typically generates *real* rigidity, not the *nominal* rigidity that is observed in the data we study. Other than our own, the most closely related paper is [Stevens \(2019\)](#), which builds a rational inattention model to explain nominal stickiness and retail sales. Stevens models *sticky plans*, while we model *sticky price points*; we discuss this distinction in Sec. 3, presenting evidence that the data favor the latter. While Stevens’ approach relies, for numerical tractability, on the coexistence of several different types of information costs, including a lump sum cost of obtaining full information, we instead obtain tractability by assuming limited memory, which allows us to impose a single, uniform information constraint on the firm’s problem. Technically, we build most closely on [Steiner et al. \(2017\)](#), which shows that a dynamic rational inattention model is equivalent to a dynamic control cost model with an optimal benchmark distribution. We extend their result in a simple but useful way, by allowing for finite memory (their model instead assumes infinite memory). This is helpful both because the case with limited memory is easier to compute, while approximating the infinite memory model, and because limited memory leads to a model of sticky price points which appears to match the data well.

These closely related papers are not the only studies arguing that costly decision-making can help explain price rigidity. The idea goes back at least to [Akerlof and Yellen \(1985\)](#), though many early papers made little distinction between menu costs and decision costs. Cite also: [Gorodnichenko \(2010\)](#), [Pastén and Schoenle \(2016\)](#), [Dixon \(2019\)](#). In an abstract setting, [Wilson \(2014\)](#) argues that limited memory can make decisions rigid. [Ilut et al. \(2019\)](#) argue that Knightian uncertainty about demand can cause nominal rigidity. [Knotek \(2011\)](#) argues that cognitive limitations may make prices with fewer digits more convenient, leading to stickiness at round numbers.² [Zbaracki et al. \(2004\)](#) present evidence from a case study showing that most of the costs of price adjustment are related to decision-making or negotiation, rather than menu costs per se. Cite also: [Ellison et al. \(2018\)](#), [Levy et al. \(2002\)](#).

It is frequently argued that sales are largely acyclical, implying that the effects of monetary policy depend primarily on the flexibility of regular prices ([Eichenbaum et al. 2011](#); [Kehoe and Midrigan 2015](#); [Guimaraes and Sheedy 2011](#)). Nonetheless, the cyclicity of sales remains controversial in the empirical literature: [Coibion et al. \(2015\)](#); [Kryvtsov and Vincent \(2014\)](#); [Carvalho and Kryvtsov \(2018\)](#); [Eden et al. \(2019\)](#). Some models of sales are based on dual menu costs: [Kehoe and Midrigan \(2015\)](#); [Eichenbaum et al. \(2011\)](#).

2. Footnote discussing complementarity between Knotek’s mechanism and our own.

2 Models of costly precision

We begin by discussing several modelling frameworks in which it is costly to choose one’s action so that it is precisely optimal given the state of the world. We define “control cost” problems (Mattsson and Weibull 2002) and “rational inattention” problems (Sims 2003), and also briefly mention “rate distortion” problems (Shannon 1959). We review some known results about the relationship between these problems and about their solutions, in preparation for the model of price dynamics we will propose in Section 4. Much of the rational inattention literature has focussed on linear-quadratic Gaussian (LQG) problems (e.g. Mackowiak et al. 2018); here we instead study more general, non-LQG problems, relying on the characterization theorems of Matějka and McKay (2015) and Steiner et al. (2017).

2.1 The control cost framework

Throughout this paper, we think of decision-making as the allocation of probability over a (possibly time-varying) set of feasible actions $a \in \mathcal{A}$.³ For notational and analytical simplicity, we assume the action set \mathcal{A} is discrete, but it could represent a grid-based approximation to a continuous action space.⁴

Control cost (CC) models (Stahl 1990; Van Damme 2002; Mattsson and Weibull 2002) are motivated by the observation that making precise decisions is costly. Many possible cost functions could operationalize this principle in an economic model, because there are many possible measures of precision. Here, we assume that decision costs are proportional to the relative entropy of the chosen probability distribution of actions, relative to an exogenous benchmark distribution of actions. This functional form is convenient for two basic reasons. First, it means we can solve analytically for the policy function, which takes the form of a logit, and also for the value function. Second, entropy measures possess a useful invariance property: if decision costs are measured by relative entropy, then directly selecting an action $a \in \mathcal{A}$ has exactly the same cost as a multi-step decision process that first selects a subset $\tilde{\mathcal{A}} \subset \mathcal{A}$ and then subsequently selects an action $a \in \tilde{\mathcal{A}}$ from that subset. Real-world choice may or may not satisfy this invariance property, but it provides a natural benchmark case against which to compare other, more complex, choice environments.

Before we define CC problems, we briefly define and discuss relative entropy, which is also called Kullback-Leibler divergence. Consider two probability mass functions $\pi_1(a)$ and $\pi_2(a)$ defined over the same support \mathcal{A} . The Kullback-Leibler divergence $\mathcal{D}(\pi_1||\pi_2)$ of distribution

3. See Machina (1985) or Anderson et al. (1992), Chap. 2, for discussion of this assumption.

4. Steiner et al. (2017) conjecture that the equivalence result we use here (Prop. 3) extends to the case of a continuum of possible actions. (MOVE THIS NOTE TO SEC. 2.3??)

π_1 relative to π_2 is:

$$\mathcal{D}(\pi_1||\pi_2) \equiv \sum_{a \in \mathcal{A}} \pi_1(a) \ln \left(\frac{\pi_1(a)}{\pi_2(a)} \right) = E_{\pi_1(a)} \ln \left(\frac{\pi_1(a)}{\pi_2(a)} \right). \quad (1)$$

It is not difficult to show that Kullback-Leibler divergence is a non-negative, convex function (Cover and Thomas 2006, Thms. 2.6.3 and 2.7.2), which equals zero if and only if $\pi_1(a) = \pi_2(a)$ for all $a \in \mathcal{A}$. Intuitively, it is a measure of the difference between distributions π_1 and π_2 . Using it as a decision cost function means that the decision-maker (DM) can costlessly set the action probabilities $\pi_1(a)$ equal to the benchmark distribution $\pi_2(a)$, but to choose any other distribution of actions she must expend some resources (which we will usually interpret as time devoted to cognitive effort) on making the decision. When π_2 is a diffuse distribution over \mathcal{A} , concentrating probabilities on actions near the optimal action a^* implies choosing a π_1 very different from π_2 , implying a high decision cost. This is the sense in which “precision is costly” in the CC models we will consider.

A *control cost problem* simply maximizes expected utility by choosing a distribution π over actions \mathcal{A} subject to a control cost function. In a static example, it takes the form:

$$V(\theta) = \max_{\pi \in \Delta(\mathcal{A})} E_{\pi} u(a, \theta) - \kappa \mathcal{D}(\pi||\eta). \quad (2)$$

Here $u(a, \theta)$ is the utility function (gross of decision costs), which depends on the action a and the state of the world θ . Distribution π is chosen from $\Delta(\mathcal{A})$, the simplex of all possible probability distributions over the action space \mathcal{A} . The distribution η is the benchmark distribution of actions, which is taken as given, and $\kappa > 0$ is a parameter relating to the costs of decision-making.

Note that the expectation in (2) is taken over actions, but not over θ , the state of the world; formally, the CC problem treats θ as known. One way to interpret this setup is to suppose that the DM has enough information to calculate the optimal action a , but that performing the calculation is costly, perhaps in terms of time. Paying a higher cost may permit a more precise decision if considering additional factors, computing higher order terms, or checking calculations takes time and reduces the probability of making big mistakes.

Expanding out the expectation and decision cost terms, problem (2) becomes:

$$\begin{aligned} V(\theta) &= \max_{\pi(a)} \sum_{a \in \mathcal{A}} \pi(a) u(a, \theta) - \kappa \sum_{a \in \mathcal{A}} \pi(a) \ln \left(\frac{\pi(a)}{\eta(a)} \right) \\ \text{s.t.} & \sum_{a \in \mathcal{A}} \pi(a) = 1 \quad \text{and} \quad \pi(a) \geq 0, \quad \forall a \in \mathcal{A}. \end{aligned} \quad (3)$$

Ignoring the inequality constraints $\pi(a) \geq 0$, the first order condition for the probability $\pi(a)$

of a given action a is:

$$u(a, \theta) - \kappa \left(1 + \ln \left(\frac{\pi(a)}{\eta(a)} \right) \right) - \mu = 0, \quad (4)$$

where μ is the multiplier on the constraint that probabilities must sum to one.

Rearranging, and ensuring that the probabilities indeed sum up to one, the optimal distribution of actions in any state θ is a weighted multinomial logit:

$$\pi(a|\theta) = \frac{\eta(a) \exp(\kappa^{-1}u(a, \theta))}{\sum_{a' \in \mathcal{A}} \eta(a') \exp(\kappa^{-1}u(a', \theta))}. \quad (5)$$

Equation (5) shows that $\pi(a|\theta) > 0$ whenever $\eta(a) > 0$, so the inequality constraints in (3) are non-binding as long as $\eta(a) > 0$. In general the probability of taking any given action a increases smoothly with the utility of action a relative to other actions. As $\kappa \rightarrow 0$, so that decisions are costless, the probability of choosing the best option $a^*(\theta) \equiv \arg \max_{a \in \mathcal{A}} u(a, \theta)$ approaches one. If instead $\kappa \rightarrow \infty$, making decisions extremely expensive, then $\pi(a|\theta)$ approaches $\eta(a)$: the distribution of actions converges to the benchmark distribution $\eta(a)$, regardless of the state of the world θ .

Finally, taking logs in (5) and plugging the result into the objective of problem (3), we can also find an analytical formula for the value function:

$$V(\theta) = \kappa \ln \left(\sum_{a \in \mathcal{A}} \eta(a) \exp(\kappa^{-1}u(a, \theta)) \right). \quad (6)$$

2.1.1 Invariance to the decision sequence

As we mentioned above, a convenient property of relative entropy control costs is their invariance to the structure of the decision sequence. Under appropriate assumptions, a single-step decision (3) across all options $a \in \mathcal{A}$ is equivalent to a multistep procedure which first chooses a subset of the action space, then perhaps chooses an even finer subset, and so forth, until finally a single alternative a is selected. Also, in a dynamic context, entropy control costs are consistent with breaking down the choice of a time series (a_0, a_1, a_2, \dots) into a sequence of Bellman equations that recursively describe the choice of each action a_t at each time t .

We demonstrate the invariance property of relative entropy control costs by way of an example. Suppose there exists some pre-existing setting $a^- \in \mathcal{A}$ in the action space which will occur if the decision-maker “does nothing”. This situation may be relevant, for example, if the DM takes her action by means of a physical mechanism, such as writing a number on a price tag, hitting a sequence of keys on a computer, or controlling a steering wheel. In such cases, there may be a well-defined outcome even if the DM fails to come to a decision and thereby fails to take a deliberate action (for example, even if she falls asleep or otherwise

absents herself). Later, this example will be useful for thinking about nominal price rigidity.

Cases like these may be analyzed as a two-step decision process, choosing first between “do nothing” or “do something”, and then selecting a specific action. We may formalize the first step as a choice between $\mathcal{A}^- \equiv \{a^-\}$ and $\tilde{\mathcal{A}} \equiv \mathcal{A} \setminus \mathcal{A}^-$. At the second step, the DM solves a problem like (2) to choose an action in set $\tilde{\mathcal{A}}$ or in set \mathcal{A}^- :

$$\tilde{V}(\theta) = \max_{\tilde{\pi} \in \Delta(\tilde{\mathcal{A}})} E_{\tilde{\pi}} u(a, \theta) - \kappa \mathcal{D}(\tilde{\pi} || \tilde{\eta}). \quad (7)$$

$$V^-(\theta) = \max_{\pi^- \in \Delta(\mathcal{A}^-)} E_{\pi^-} u(a, \theta) - \kappa \mathcal{D}(\pi^- || \eta^-) = u(a^-, \theta). \quad (8)$$

Problem (8) is trivial because $\Delta(\mathcal{A}^-)$ is a singleton: when a^- is the only feasible action, $\pi^-(a^-) = 1$ is the only feasible distribution. This is a special case of (2) with zero decision costs (because $\eta^-(a^-) = 1$ is the only feasible benchmark), so $V^-(\theta) = u(a^-, \theta)$.

At the first step, the DM chooses across the two possible choice sets:

$$V^0(\theta) = \max_{\lambda \in [0,1]} (1 - \lambda)V^-(\theta) + \lambda\tilde{V}(\theta) - \kappa \left[(1 - \lambda) \ln \left(\frac{1 - \lambda}{1 - \eta_0} \right) + \lambda \ln \left(\frac{\lambda}{\eta_0} \right) \right]. \quad (9)$$

This first-stage problem is again formally equivalent to (2). Note that λ can be interpreted as the probability of choosing to take an action, while $1 - \lambda$ can be interpreted as the probability of doing nothing. In multi-period or continuous-time generalizations of this setup, λ represents the arrival rate of a change of action, as in the CC model of price stickiness developed by [Costain and Nakov \(2019\)](#).

Problems (7)-(9) all have the same form as (2), so they all generate logit solutions: It is easy to verify that appropriate choices of the benchmark distributions $\tilde{\eta}$, η^- , and η_0 make the solutions of (7)-(9) equivalent to (5). Namely, these benchmarks must be constructed so that they represent conditional probabilities derived from the benchmark distribution η that appears in problem (2). This requires $\eta_0 = \eta(a^-)$, $\eta^-(a^-) = 1$, and $\tilde{\eta}(a) = \eta(a|a \in \tilde{\mathcal{A}})$; then (7)-(9) are jointly equivalent to (2).⁵

The example considered above is a special case of a general point: given appropriately-defined benchmark distributions, single-step and multi-step CC problems are equivalent. We state the general result as [Proposition 1](#). The proof is omitted, since it is just a matter of verifying that the logit solutions are equivalent when we plug in benchmark distributions that represent conditional expectations, as in the example outlined above.

Proposition 1 *Let $V(\theta; \mathcal{A}, \eta)$ be the value of (2) when the choice set is \mathcal{A} and the benchmark distribution is η . Let $\mathcal{A}^0 \equiv \{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_n\}$ be a partition of \mathcal{A} . Define $\eta^0(\mathcal{B}_i) \equiv \sum_{a \in \mathcal{B}_i} \eta(a)$,*

5. This invariance property follows from the “grouping axiom” of [Shannon \(1948\)](#); see [Ash \(1990\)](#).

and $\xi_i(a) \equiv \eta(a)/\eta^0(\mathcal{B}_i)$ for each $a \in \mathcal{B}_i$.

Consider the CC problem:

$$V^0(\theta; \mathcal{A}^0, \eta^0) = \max_{\pi^0 \in \Delta(\mathcal{A}^0)} E_{\pi^0} V(\theta; \mathcal{B}, \xi) - \kappa \mathcal{D}(\pi^0 || \eta^0). \quad (10)$$

The values of problems (10) and (2) are the same: $V^0(\theta; \mathcal{A}^0, \eta^0) = V(\theta; \mathcal{A}, \eta)$, and the distribution over actions $a \in \mathcal{A}$ generated by choosing a set $\mathcal{B}_i \in \mathcal{A}^0$ according to (10), and subsequently choosing an action $a \in \mathcal{B}_i$ according to (2), is the same as the distribution (5) generated by choosing an action $a \in \mathcal{A}$ directly according to (2).

2.2 The rational inattention framework

Like the CC framework, rational inattention (RI) assumes that it is costly to choose precisely the action that maximizes the gross utility function conditional on the underlying state of the world. But RI (as defined by Sims 2003) imposes a more specific interpretation on the costs, relating them to the amount of information required to make the decision.

Concretely, the cost function in a rational inattention problem is proportional to the *mutual information* between the state of the world, θ , and the chosen action, a . Hence, before we define RI problems, we briefly discuss mutual information, which is a special case of Kullback-Leibler divergence. For two random variables $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ with joint distribution $\pi(x, y)$, mutual information $\mathcal{I}(X; Y)$ is defined as:

$$\mathcal{I}(X, Y) \equiv \mathcal{D}(p(x, y) || p_X(x)p_Y(y)) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \pi(x, y) \ln \left(\frac{\pi(x, y)}{\pi_X(x)\pi_Y(y)} \right) = E_{\pi(x, y)} \ln \left(\frac{\pi(x, y)}{\pi_X(x)\pi_Y(y)} \right), \quad (11)$$

where $\pi_X(x) \equiv \sum_{y \in \mathcal{Y}} \pi(x, y)$ and $\pi_Y(y) \equiv \sum_{x \in \mathcal{X}} \pi(x, y)$ are the marginal distributions of X and Y , respectively. Note that if X and Y are independent, $\pi(x, y)/(\pi_X(x)\pi_Y(y)) = 1$ for all x and y , so $\mathcal{I}(X, Y) = 0$; otherwise $\mathcal{I}(X, Y)$ is strictly positive. Hence one interpretation of mutual information is that it measures the degree of dependence between X and Y .⁶

To understand optimal choice when information is costly, note that the chosen action may be considered a signal. If it were costless to learn the state of the world and choose an action accordingly, the DM would always choose the optimal action $a^*(\theta)$, so θ could be inferred. In other words, the conditional distribution $\pi(a|\theta)$ would place probability one on $a^*(\theta)$. More generally, when information is costly, the chosen conditional distribution $\pi(a|\theta)$ will randomize across values of a . It is therefore convenient to rewrite mutual information in terms of the conditional distribution $\pi(a|\theta) = \pi(a, \theta)/\pi_\theta(\theta)$, where $\pi_\theta(\theta)$ represents prior

6. Footnote regarding other interpretations, e.g. $\mathcal{I}(X; Y) = H(Y) - H(Y|X) = H(X) - H(X|Y)$.

knowledge about possible states of the world.

$$\mathcal{I}(a, \theta) = \sum_{\theta \in \Theta} \sum_{a \in \mathcal{A}} \pi(a|\theta) \pi_\theta(\theta) \ln \left(\frac{\pi(a|\theta)}{\pi_a(a)} \right) = E_{\pi(a, \theta)} \ln \left(\frac{\pi(a|\theta)}{\pi_a(a)} \right). \quad (12)$$

We now define a static *rational inattention problem* as the choice of a conditional action distribution subject to a limit on the mutual information between the action and the state of the world:

$$U(\pi_\theta(\theta)) = \max_{\pi(a|\theta) \in \Delta(\mathcal{A})} E_{\pi(a, \theta)} u(a, \theta) - \kappa \mathcal{I}(a; \theta) \quad (13)$$

$$= \max_{\pi(a|\theta) \in \Delta(\mathcal{A})} \sum_{\theta \in \Theta} \pi_\theta(\theta) \sum_{a \in \mathcal{A}} \pi(a|\theta) \left[u(a, \theta) - \kappa \ln \left(\frac{\pi(a|\theta)}{\pi_a(a)} \right) \right] \quad \text{s.t.} \quad \pi_a(a) = \sum_{\theta \in \Theta} \pi_\theta(\theta) \pi(a|\theta). \quad (14)$$

Eqs. (13)-(14) take expectations with respect to the joint distribution $\pi(a, \theta)$, factoring it as $\pi(a, \theta) = \pi_\theta(\theta) \pi(a|\theta)$. The problem is solved by choosing a conditional distribution $\pi(a|\theta)$ from the simplex $\Delta(\mathcal{A})$ for each $\theta \in \Theta$.

(ALSO STATE OR PROVE LEMMA 10.8.1 of Cover/Thomas: $\mathcal{I}(X, Y) = \min_r \mathcal{D}(\pi(x, y) || \pi_x(x)r(y))$).

Given these preliminaries, we can now characterize the solution of the static RI problem following [Matějka and McKay \(2015\)](#). They showed that the RI problem is a special case of the CC problem of [Mattsson and Weibull \(2002\)](#), and therefore its solution takes the form of a multinomial logit. More specifically, it is a CC problem in which the benchmark action distribution is chosen optimally. In turn, optimally choosing that distribution means setting it equal to the *marginal* action distribution. Formally, we can state these results as follows:

Proposition 2 (*Matejka and McKay, 2015*). *The rational inattention problem (13) is equivalent to a control cost problem with an optimally-chosen benchmark distribution:*

$$U(\pi_\theta(\theta)) = \max_{q(a) \in \Delta(\mathcal{A})} \sum_{\theta \in \Theta} \pi_\theta(\theta) \max_{\pi(a|\theta) \in \Delta(\mathcal{A})} \sum_{a \in \mathcal{A}} \pi(a|\theta) \left[u(a, \theta) - \kappa \ln \left(\frac{\pi(a|\theta)}{q(a)} \right) \right]. \quad (15)$$

Therefore, the rational inattention problem is solved by the following weighted multinomial logit, in which the weights correspond to the marginal probabilities of each action:

$$\pi(a|\theta) = \frac{q(a) \exp(\kappa^{-1} u(a, \theta))}{\sum_{a' \in \mathcal{A}} q(a') \exp(\kappa^{-1} u(a', \theta))}, \quad (16)$$

$$q(a) = \pi_a(a) = \sum_{\theta \in \Theta} \pi_\theta(\theta) \pi(a|\theta). \quad (17)$$

The value function is:

$$U(\pi_\theta(\theta)) = \kappa \sum_{\theta \in \Theta} \pi_\theta(\theta) \ln \left(\sum_{a \in \mathcal{A}} q(a) \exp(\kappa^{-1} u(a, \theta)) \right). \quad (18)$$

Notice that the inner maximization in (15) represents the value function $V(\theta)$ from the control cost problem (2), given benchmark action probabilities $q(a)$. Therefore (16), which is the first-order condition for $\pi(a|\theta)$, states that conditional action probabilities should take the form of a logit. Equation (17) is the first-order condition for $q(a)$: the optimal benchmark distribution of actions is the *marginal* distribution of actions. Intuitively, optimal information-constrained choice gives the DM a “predisposition” to choose more frequently the actions that are *unconditionally* more likely to be optimal.⁷ Together, (16) and (17) offer a simple iterative algorithm for calculating the optimal RI behavior $\pi(a|\theta)$. Given a utility function $u(a, \theta)$ and the decision cost κ , we can guess any unconditional action distribution $q(a) \in \Delta(\mathcal{A})$ and calculate the corresponding logit distribution of actions from (16). Given the logit $\pi(a|\theta)$, we can construct the corresponding marginals from (17). These two steps are guaranteed to converge, and their fixed point solves the rational inattention problem (13).⁸

Proving that RI is equivalent to a particular CC problem is harder in a dynamic context, but Steiner et al. (2017) have demonstrated that essentially the same arguments apply (Appendix A restates the main arguments of their proof). In a dynamic context, the DM’s information at t consists of the history a^{t-1} of signals a_s that she has observed at previous times $s \leq t$. Hence we can define a *dynamic rational inattention problem* as follows:

$$U(a^0) = \max_{\pi(a_t|\theta^t, a^{t-1}) \in \Delta(\mathcal{A})} E \left[\sum_{t=1}^{\infty} \delta^t (u(a_t, \theta_t) - \kappa \mathcal{I}(a_t, \theta_t | a^{t-1})) \middle| a^0 \right]. \quad (19)$$

We can generalize Prop. 2 to the dynamic case as follows.

Proposition 3 (Steiner, Stewart, and Matejka, 2017).

(i). The dynamic RI problem (19) is equivalent to the following double optimization:

$$U(a^0) = \max_{\pi, q} E \left[\sum_{t=1}^{\infty} \delta^t \left(u(a^t, \theta^t) - \kappa \log \left(\frac{\pi(a_t|\theta^t, a^{t-1})}{q(a_t|a^{t-1})} \right) \right) \middle| a^0 \right]. \quad (20)$$

7. Steiner et al. (2017) call the benchmark action distribution a “predisposition”. As they explain, “actions that are unappealing ex ante can only become appealing through costly updating of beliefs”.

8. This fixed point method is known as the Blahut-Arimoto algorithm. It was proposed by Blahut (1972) as a method to solve *rate distortion problems*, a class of information compression problems proposed by Shannon (1959) which are dual to rational inattention problems. The algorithm can be proved to converge (Csiszár and Tusnády 1984) because it represents two alternating optimization problems over convex sets, as seen in (15): choosing $\pi(a|\theta) \in \Delta(\mathcal{A})$ for each θ , taking as given $q(a)$, and choosing $q(a) \in \Delta(\mathcal{A})$, taking as given $\pi(a|\theta)$. See Cover and Thomas (2006), Fig. 10.9, for discussion.

(ii). Problem (20) represents an expectation across full-information CC problems under an optimal benchmark distribution q :

$$U(a^{t-1}) = \delta \max_{q(a|a^{t-1}) \in \Delta(\mathcal{A})} \sum_{\theta^t} \pi(\theta^t | a^{t-1}) V(\theta^t; a^{t-1}, q). \quad (21)$$

(iii). In (21), $V(\theta^t; a^{t-1}, q)$ is the value of a recursive, full-info CC problem:

$$V(\theta^t; a^{t-1}, q) = \max_{\pi(a_t | \theta^t, a^{t-1}) \in \Delta(\mathcal{A})} \sum_{a_t \in \mathcal{A}} \pi(a_t | \theta^t, a^{t-1}) \left[u(a_t, \theta_t) - \kappa \ln \left(\frac{\pi(a_t | \theta^t, a^{t-1})}{q(a_t | a^{t-1})} \right) + \delta \sum_{\theta^{t+1}} \pi(\theta^{t+1} | \theta^t) V(\theta^{t+1}; a^t, q) \right]. \quad (22)$$

(iv). Hence, (19) and (20) are solved by a weighted multinomial logit:

$$\pi(a_t | \theta^t, a^{t-1}) = \frac{q(a_t | a^{t-1}) \exp(\kappa^{-1} \hat{v}(a_t, \theta^t; a^{t-1}, q))}{\sum_{a' \in \mathcal{A}} q(a' | a^{t-1}) \exp(\kappa^{-1} \hat{v}(a', \theta^t; a^{t-1}, q))}, \quad (23)$$

where

$$\hat{v}(a_t, \theta^t; a^{t-1}, q) \equiv u(a_t, \theta_t) + \delta \sum_{\theta'} \pi(\theta' | \theta^t) V(\theta'; a^t, q). \quad (24)$$

(v). The optimal q is the marginal distribution, conditional on signals observed:

$$q(a_t | a^{t-1}) = \sum_{\theta^t} \pi(a_t | \theta^t, a^{t-1}) \pi(\theta^t | a^{t-1}). \quad (25)$$

Propositions 2 and 3 clarify the analysis of RI problems in several ways. First, they show that RI problems are equivalent to an expectation across full-information CC problems. Therefore, the optimal distribution of actions under RI must be a weighted multinomial logit. Furthermore, the weights are intuitive: they represent the *marginal* probabilities of the actions. That is, while each CC problem takes the state of the world θ_t (or the history θ^t) as known, the RI problem imposes benchmark action probabilities that represent the *average* probabilities with which those actions would be played, given current information (*i.e.* taking an expectation with respect to the current prior over θ_t).

Discuss: this could still be very hard to take to the data: ideally need to know $\pi(a_t | a^{t-1})$ in every possible information set a^{t-1} ! Hence previous work has only computed this solution in simple toy examples. Nonetheless, $\pi(a)$ can be observed *on average*. Therefore data tell us which weights should be imposed on the problem *on average*. Moreover, whether those weights actually vary across information sets can be investigated empirically by observing $\pi(a)$ with coarse conditioning: That is, even if we do not have repeated information on all possible information sets, we can subdivide our data to condition on some previous actions; for example, we can condition on a_{t-1} even if we do not condition on the full history a^{t-1} . To

do this, however, we need to be sure that we define the action set \mathcal{A} in a stationary way, so that we can make repeated, comparable observations of the same situation, in order to infer the unconditional or coarsely conditional distribution $\pi(a)$.

This intuition also suggests an approximation scheme for computing RI models. The computational intractability of the RI problem relates to the fact that it conditions on the full history of previous signals received, a^{t-1} . However, we might consider instead a decision-maker restricted to short-term memory, who recalls only a finite history $\mathcal{B}_\tau^{t-1} \equiv (a_{t-1}, a_{t-2}, \dots, a_{t-\tau})$. Given prior knowledge $\pi_\theta(\theta)$, we can define an RI problem for this DM exactly as we did before. Concretely, a *short-term memory rational inattention (STMRI) problem* is:

$$U(\mathcal{B}_\tau^0) = \max_{\pi(a_t|\theta^t, \mathcal{B}_\tau^{t-1}) \in \Delta(\mathcal{A})} E \left\{ \sum_{t=1}^{\infty} \delta^t (u(a_t, \theta_t) - \kappa \mathcal{I}(a_t, \theta^t | \mathcal{B}_\tau^{t-1})) \middle| \mathcal{B}_\tau^0 \right\}. \quad (26)$$

This problem has the same form as (19), and the same properties apply: it can be rewritten as an expectation across full-information CC problems with optimally-chosen benchmark distributions. Therefore, like the infinite-memory RI problem (19), it gives rise to weighted logit behavior, with optimal weights. Hence, we can generalize Prop. 3 as follows (see Appendix A for the proof).

Proposition 4 (a). *The results of Prop. 3 also hold for the STMRI problem (26), when the infinite signal history $\mathcal{B}_\infty^{t-1} \equiv a^{t-1}$ is replaced by the finite history \mathcal{B}_τ^{t-1} . In particular, (26) is solved by a logit:*

$$\pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1}) = \frac{q(a_t | \mathcal{B}_\tau^{t-1}) \exp(\kappa^{-1} \hat{v}(a_t, \theta^t; \mathcal{B}_\tau^{t-1}, q))}{\sum_{a' \in \mathcal{A}} q(a' | \mathcal{B}_\tau^{t-1}) \exp(\kappa^{-1} \hat{v}(a', \theta^t; \mathcal{B}_\tau^{t-1}, q))}, \quad (27)$$

$$q(a_t | \mathcal{B}_\tau^{t-1}) = \sum_{\theta^t} \pi(\theta^t | \mathcal{B}_\tau^{t-1}) \pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1}). \quad (28)$$

where \hat{v} is derived from the value function V of a recursive CC problem:

$$\hat{v}(a, \theta, \mathcal{B}_\tau^{t-1}, q) \equiv u(a, \theta) + \delta \sum_{\theta'} \pi(\theta' | \theta) V(\theta', \mathcal{B}_\tau^{t-1}; q), \quad (29)$$

$$V(\theta^1, \mathcal{B}_\tau^0; q) = \max_{\pi(a_1 | \theta^1, \mathcal{B}_\tau^0) \in \Delta(\mathcal{A})} \sum_{a_1} \pi(a_1 | \theta^1, \mathcal{B}_\tau^0) \left[u(a_1, \theta_1) - \kappa \ln \left(\frac{\pi(a_1 | \theta^1, \mathcal{B}_\tau^0)}{q(a_1 | \mathcal{B}_\tau^0)} \right) + \delta \sum_{\theta^2 | \theta^1} \pi(\theta^2 | \theta^1) V(\theta^2, \mathcal{B}_\tau^1; q) \right]. \quad (30)$$

(b). *As $\tau \rightarrow \infty$, the probabilities and value functions (27)–(30) that solve the STMRI problem converge to the probabilities (23) and (25) and value functions (24) and (22) that solve the infinite-memory RI problem.*

We should emphasize that this setup imposes a plausible but arbitrary form of limited memory—remembering the most recent signals received (actions taken). *Optimal* use of limited memory would probably compress these memories in some way, but analyzing that problem is beyond the scope of this paper. A useful aspect of this limited memory setup is part (b) of the proposition, which allows us to think of STMRI as an approximation scheme for calculating the RI model: as the size of memory increases, the STMRI model converges to the infinite-memory RI model. We will also see that the STMRI model naturally generates pricing patterns similar to those seen in retail pricing data.

2.3 Discreteness of the RI solution

Even when the underlying shocks $\theta \in \Theta$ and/or the action $a \in \mathcal{A}$ may take a continuum of values, optimal decisions under a rational inattention constraint typically randomize across a smaller, discrete set of actions. This fact has played an important role in the literature on price dynamics under RI (Matějka 2016, Jung et al. 2015, Stevens 2019), where the inherent discreteness of the RI solution has been offered as an explanation of nominal stickiness (the repetition of precisely the same nominal price across multiple points in time) and of transitory price fluctuations that return to the same price point visited earlier, which we colloquially call “sales”. This property was originally proved in the information theory literature by Fix (1978), who showed that rate distortion problems may have discrete solutions even when the signal to be encoded is continuous.

To see why an RI solution may randomize over a small set of isolated points, when the set of actions available is large, or is even a continuum, consider the maximization over the benchmark distribution q in (15), or in (21)-(22), which is formally identical. The constraint $q(a) \in \Delta(\mathcal{A})$ in (15) is an abbreviation for $\sum_{a \in \mathcal{A}} q(a) = 1$ and $q(a) \geq 0, \forall a \in \mathcal{A}$. Let the multiplier on the constraint $\sum_{a \in \mathcal{A}} q(a) = 1$ be ξ . Then the complementary slackness condition for $q(a)$ in (15) is

$$\sum_{\theta \in \Theta} \pi_{\theta}(\theta) \frac{\pi(a|\theta)}{q(a)} \leq \xi, \quad (31)$$

with equality for any a such that $q(a) > 0$ strictly. The first-order condition for $\pi(a|\theta)$, which leads to a logit weighted by $q(a)$, implies that $\pi(a|\theta)$ is strictly positive (zero) if and only if $q(a)$ is strictly positive (zero). Hence, summing over all a such that $q(a) > 0$, (31) implies that $\xi = 1$:

$$1 = \sum_{\theta \in \Theta} \sum_{a \in \mathcal{A}} \pi(a, \theta) = \xi \sum_{a \in \mathcal{A}} q(a) = \xi. \quad (32)$$

If we now plug the logit policy function back into the complementary slackness condition, we obtain the following result.

Proposition 5 (Fix, 1978; Matejka 2016).

(a). The benchmark distribution $q(a) \geq 0$ solves (15) if it satisfies

$$C(a) = \frac{\sum_{\theta \in \Theta} \pi_{\theta}(\theta) \sum_{a' \in \mathcal{A}} \exp(\kappa^{-1}u(a, \theta))}{\sum_{a' \in \mathcal{A}} q(a') \exp(\kappa^{-1}u(a', \theta))} \leq 1, \quad (33)$$

at all $a \in \mathcal{A}$, with equality wherever $q(a) > 0$ strictly.

Likewise, $q(a_t | \mathcal{B}_{\tau}^{t-1})$ solves (26) if (33) holds when $u(a, \theta)$ and $\pi_{\theta}(\theta)$ are replaced by $\hat{v}(a_t, \theta^t; \mathcal{B}_{\tau}^{t-1}, q)$ and $\pi(\theta^t | \mathcal{B}_{\tau}^{t-1})$. This is a solution of (19) in the limiting case $\tau = \infty$.

(b). Suppose u is a quadratic function. Then generically, the set \mathcal{Q} of points such that $q(a) > 0$ is a discrete, finite set.

Prop. 5 shows that RI behavior may concentrate all probability on a strict subset $\mathcal{Q} \subset \mathcal{A}$ of the available actions. Remarkably, \mathcal{Q} may be a finite, discrete set even when \mathcal{A} is an interval (an uncountable set). In fact, Matějka (2016) proves, for the case of a quadratic payoff function, that as long as the set of possible actions \mathcal{A} is bounded, \mathcal{Q} must be a finite, discrete set. The familiar LQG class of RI problems, in which $\mathcal{Q} = \mathcal{A} = \mathcal{R}$ (the real line), turns out to be a very special case: restricting choices to a bounded subset $\mathcal{A} \subset \mathcal{R}$ causes \mathcal{Q} to collapse to a finite, discrete set, and changing the distribution of shocks (Gaussian) may have the same effect.⁹

While the quadratic case is only an example, numerical studies have shown that the discreteness of RI solutions hold much more generally. Therefore, Matějka (2016) argues that this discreteness result may provide an explanation for price stickiness. However, discreteness and stickiness are not exactly the same property. Moreover, the stickiness that Matejka’s model generates may be *nominal* or *real* price stickiness. We will discuss the assumptions needed in Matejka’s framework to obtain *nominal* price stickiness like that seen in microdata when we come to our empirical model, in Sec. 4.

3 Retail pricing data

A cursory glance at microdata from retail stores demonstrates that nominal prices vary frequently between widely differing levels, and that they often return to points they have visited before. Therefore several papers have proposed that retail firms intermittently make “plans” consisting of a set of possible price points, and then subsequently select their prices from within that set. Models of “sticky plans” include Stevens (2019), which derives them from information constraints; but they also include full-information models with two types of menu costs (a cost for making a plan, and a second, lower cost for setting the price), as in Eichen-

9. More precisely, Matějka (2016) shows that \mathcal{Q} must be a finite, discrete set, unless the action set is unbounded and the payoff function is analytic.

baum et al. (2011) or Kehoe and Midrigan (2015). However, the fact that prices sometimes revisit previous values does not necessarily imply that firms intermittently reconsider their entire plans. Actually, a closer look at the data reveals very little evidence of intermittent replanning.

To understand these issues, a taxonomy of different types of price adjustment events is helpful. Consider a series of nominal prices $P_{i,j,t}$ for product i in store j across weeks t . For each t , let $\mathcal{B}_\tau^{t-1} \equiv (P_{i,j,t-1}, P_{i,j,t-2}, \dots, P_{i,j,t-\tau})$ be the vector of prices observed in the last τ periods. Likewise, let $\mathcal{F}_\tau^{t+1} \equiv \{P_{i,j,t+1}, P_{i,j,t+2}, \dots, P_{i,j,t+\tau}\}$ be the prices in the next τ periods. There is a *price change* at t if $P_{i,j,t} \neq P_{i,j,t-1}$. We define three types of price changes, according to their stickiness and recurrence. We say that a price change is *transitory* if $P_{i,j,t}$ is never observed in the backwards or forwards windows:¹⁰

$$P_{i,j,t} \neq P_{i,j,t-1} \text{ and } P_{i,j,t} \notin \mathcal{B}_\tau^{t-1} \text{ and } P_{i,j,t} \notin \mathcal{F}_\tau^{t+1}. \quad (34)$$

A *price recurrence* is a price change that repeats a value observed in the backwards window:

$$P_{i,j,t} \neq P_{i,j,t-1} \text{ and } P_{i,j,t} \in \mathcal{B}_\tau^{t-1}. \quad (35)$$

A *price introduction* is a change that introduces a new price point which later recurs in the forwards window:

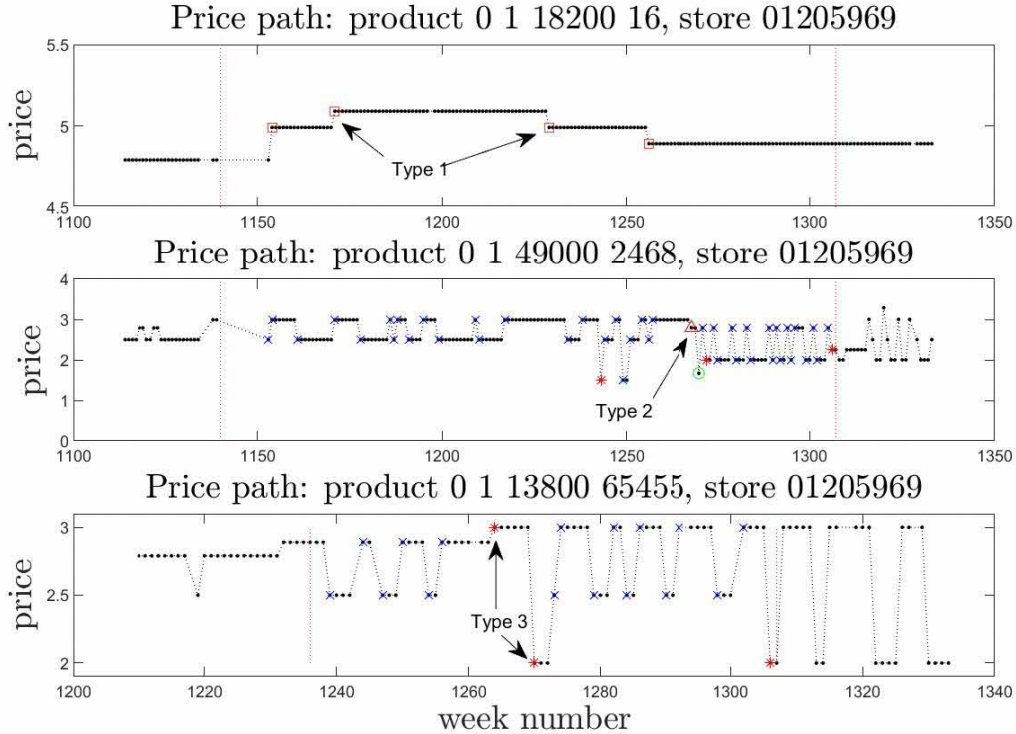
$$P_{i,j,t} \neq P_{i,j,t-1} \text{ and } P_{i,j,t} \notin \mathcal{B}_\tau^{t-1} \text{ and } P_{i,j,t} \in \mathcal{F}_\tau^{t+1}. \quad (36)$$

We will be particularly interested in price introductions, which we further classify into three types. A **type 1** price introduction occurs when no prices coincide between the backwards and forwards windows, and furthermore there is no recurrence within those windows. Type 1 price introductions will be observed if prices are sticky but price points are not – that is, if a product’s nominal price often remains fixed over time, with no tendency to revisit its past values. A **type 2** price introduction occurs when no prices coincide between the backwards and forwards windows, but prices recur within the backwards window or the forwards window. Therefore a type 2 price introduction at time t might represent a *change of plan*: prices do recur, but not at time t , when all the price points change. Finally, we define a **type 3** price introduction as one that is overlapped by the recurrence of one or more other price points. So if $P_{i,j,t} \neq P_{i,j,t-1}$ is a price introduction, and

$$\mathcal{B}_\tau^{t-1} \cap \mathcal{F}_\tau^{t+1} \neq \emptyset \quad (37)$$

10. Although we defined \mathcal{B} as a vector, we will sometimes use the same notation to refer to the set of prices appearing in the vector. Likewise for \mathcal{F} .

Figure 1: Three examples of price sequences in store A



Notes: Weekly time series of nominal prices for three selected products from store A. Green circles indicate transitory price changes. Blue “X” marks price recurrences. Price introductions are marked by red square (type 1), red triangle (type 2), or red star (type 3); the three products shown are selected because they exhibit especially high frequencies of types 1, 2, and 3 introductions (relative to other introductions), respectively. Red vertical lines highlight the end of the first backwards window and the beginning of the last forwards window (using window length $\tau = 26$).

then it is a price introduction of type 3. Thus, observing a type 3 price introduction at time t indicates that there is some stickiness of price points, and that some but not all of those price points are changing at time t .

To illustrate, we study price dynamics at two US stores (to be called A and B) in the IRI retail scanner dataset; the former has a relatively low frequency of price changes (a median frequency of 14.7% per week, across its products) while the latter has a higher frequency (28.1% per week).¹¹ The three panels of Figure 1 show examples of weekly nominal price sequences from store A for products characterized by frequent type 1, type 2, and type 3 price introductions, respectively. The top panel shows the price series for a product that exhibits price stickiness, but shows no evidence of “sales” or of multiple sticky price points. Hence,

11. We restrict our analysis to products for which we have sufficiently long price series (at least XXX weeks). For documentation of the IRI scanner data, see Bronnenberg et al. (2008). Stores A and B have identifiers 01205969 and 01250094 in the dataset; they are chosen to illustrate the cases of low and high adjustment frequencies in the dataset.

the only price changes observed are type 1 price introductions, marked as red squares.

The products shown in the second and third panels both exhibit sales behavior, with two or more sticky price points evident at all times. The second panel gives an example of a type 2 price introduction (red triangle). This product’s price bounces between two or three recurring points both before and after the type 2 introduction, but none of the price points occurring after that introduction coincide with those before it. Therefore this type 2 introduction might represent a change of “plan”. The third panel shows another product with sticky price points, but in this case there is no sign of changing plans; all the observed introductions are type 3 (red stars), being overlapped by the recurrence of other sticky price points. In particular, new high and low price points are introduced shortly after week 1260, but these introductions are both overlapped by a recurring middle price point, so they are classified as type 3.

Table 1 documents the frequencies of different types of price adjustment events.¹² While the data display heterogeneous adjustment behavior across stores and products, two clear conclusions jump out. First, the majority of price changes are recurrences. Across products, the mean fraction of price changes that are recurrences is 60.1% in store A, and 53.8% in store B. The remaining price changes are introductions (31.2% in store A, and 25.5% in store B) or transitory changes (only 8.8% in store A, but substantially higher, at 20.7%, in store B). Second, most price introductions are of type 3. Across products, the mean fraction of price introductions that are of type 3 is 73.1% in store A, and 82.4% in store B. Store A appears to avoid sales for some of its products at some times, as 16.7% of price introductions are type 1, while store B engages in sales for almost all its products, with only 6% of price introductions classified as type 1. “Replanning” events, if any, appear to be infrequent, as only 10.3% (11.6%) of price introductions at store A (B) are of type 2.

Figures 2 and 4 plot histograms that show how the frequencies of the various pricing events vary across products. First, consider the left panel of Fig. 2, which analyzes the frequencies of transitory changes, recurrences, and introductions at store A. To construct the figure, we calculated the frequencies of each type of price change for each product (as percentages of all price changes) and then plotted a histogram with bins counting the number of products with given sample frequencies, over the simplex of possible frequencies. We see again that transitory price changes are rare at store A; most products exhibited introductions or recurrences only. For typical products, around 80% of price changes are recurrences, and the remainder introductions. But for about 70 products, the only observed price changes are price introductions; these may be products without sales behavior.

Next, the histogram in the right panel shows the frequencies of the three types of price

12. There are also a few products that display no price changes at all. For 39 of the 1231 products meeting our sample size criterion at store A, no price changes occur; likewise for 34 out of the 5060 products we analyze at store B. These products are not considered when constructing the statistics in Table 1 and the graphs in Figs. ?? and 4.

Table 1: Nominal price adjustment statistics (medians across products)*

	Store A			Store B		
	Overall	Type 1	Type 3	Overall	Type 1	Type 3
<i>Sample size^a</i>						
Num. of products	1231	140	563	5060	114	1963
Length of sample	220	220	219	218	229	174
Fraction missing	0.10	0.11	0.10	0.03	0.03	0.03
<i>Price changes</i>						
Num. of distinct prices ^b	7	4	7	21	4	19
Freq. price changes (pct weekly)	14.7	2.4	18.8	28.1	1.5	35.6
<i>Classifying price changes^c</i>						
Pct. introductions	25.0	77.5	16.7	21.1	100	16.7
Pct. recurrences	65.6	0	77.8	55.2	0	61.8
Pct. transitory	5.0	16.7	3.6	20.0	0	19.4
Pct. introtype 1	0	100	0	0	100	0
Pct. introtype 2	0	0	0	6.3	0	0
Pct. introtype 3	88.9	0	100	91.3	0	100
<i>Short-run volatility^d</i>						
Ratio $VR^{avg}(\tau)$	2.7	0.56	3.8	2.1	0.45	3.2
Ratio $VR^{diff}(\tau)$	35.3	11.6	38.2	33.3	9.74	37.2
Ratio $VR^{reg}(\tau)$	4.1	1.1	5.2	3.2	1.0	3.7
Ratio $VR^{abs}(\tau)$	18.1	1.9	23.3	18.4	1.41	23.2

Notes: *All statistics, except “number of products”, are medians across products.

^aColumns of “overall” statistics refer to the sample of products with sample length ≥ 100 weeks, amounting to XX% (YY%) of the original number of products at store A (B).

Columns “type 1” (3) refer to the subset of the “overall” sample with type 1 (3) introductions only.

^bTotal number of distinct prices observed over the sample for a given product.

^cIntroductions, recurrences, and transitory changes are reported as percent of all price changes.

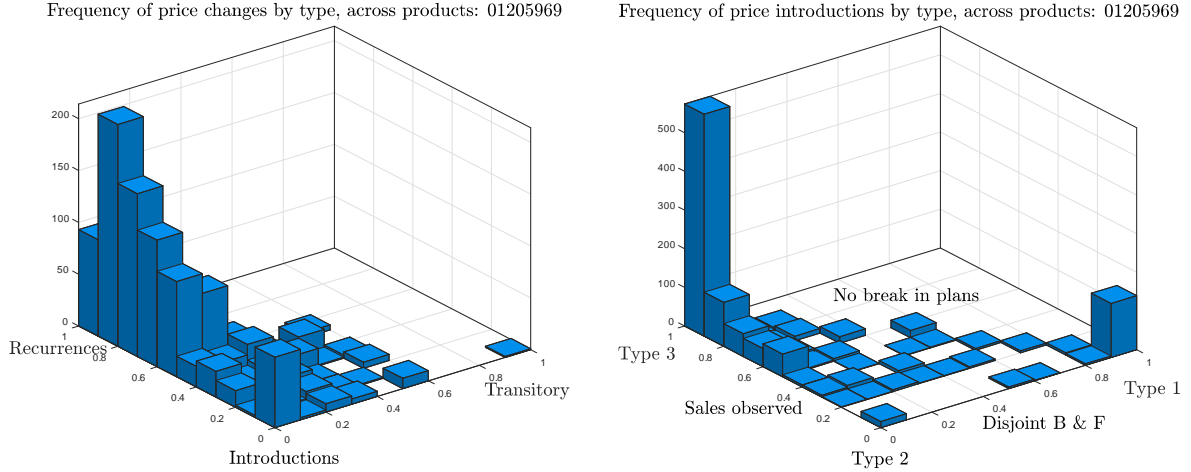
Introductions of types 1, 2, and 3 are reported as percent of all introductions.

^dRatios measuring short-run excess price volatility, as defined in the text.

introductions (excluding products where no introductions occur). Here, the histogram is graphed over the simplex of possible sample frequencies of the three types of introductions. As in Table 1, we see that type 3 price introductions are by far the most common; under the modal behavior (almost 600 products), at least 90% of all introductions are of type 3. A smaller set of products display both type 2 and 3 introductions, with less than half being type 2. Finally, there are 140 products with type 1 price introductions only; these likely include the products from the left panel that exhibited price introductions only (without any other price changes).

Given the cross-product heterogeneity seen in Fig. 2, Table 1 also reports statistics separately for the subsamples of products that display type 1 introductions only (second and fifth data columns) or type 3 introductions only (third and sixth data columns). The 140 products at store A with type 1 introductions only have a much lower frequency of price

Figure 2: Frequencies of pricing events, across products, in store A.



Notes: Figures represent histograms of sample frequencies of different types of price change events, across products, for store A.

Left: Histogram classifying price changes as transitory, recurrences, or introductions. Each point on the simplex represents the sample frequencies of these three types of changes, for a given product. The height of the bars represents the number of observed products in the given bin of sample frequencies.

Right: Histogram classifying price introductions as types 1, 2, or 3. Each point on the simplex represents the sample frequencies of the three types of introductions, for a given product. The height of the bars represents the number of observed products in the given bin of sample frequencies.

adjustment (2.4% per week), and score much lower on several ratios that measure short-run price volatility relative to long-run volatility (discussed further below). In contrast, a total of 563 products, almost half the sample, display type 3 introductions only; their median price adjustment frequency is 18.8% per week, slightly higher than the overall median.

As we saw in Fig. 1, products with type 2 or type 3 price changes can have a highly variable price from week to week, even when a moving average calculated over several weeks would show a stable mean price. In other words, retail sales behavior appears to introduce a form of excess short-run price volatility, relative to longer-term measures of price volatility. Table 1 quantifies this by reporting several measures of excess short-run volatility, which are all higher in the subsample of products displaying type 3 introductions only than they are for the subsample with type 1 introductions only. Since type 3 introductions are far more common than type 1, the ratios calculated from the overall sample are typically close to those from the type 3 subsample.

The excess volatility measures reported in Table 1 are all ratios of short-run to long-run price volatility, defined in several different ways. We define the long-run price either as a moving average, or as a “regular price”, in the sense of Eichenbaum et al. (2011). Let $\mu_{i,j,t}(\tau)$ be the moving average of the log price in the time window from $t - \tau$ to $t + \tau$, and

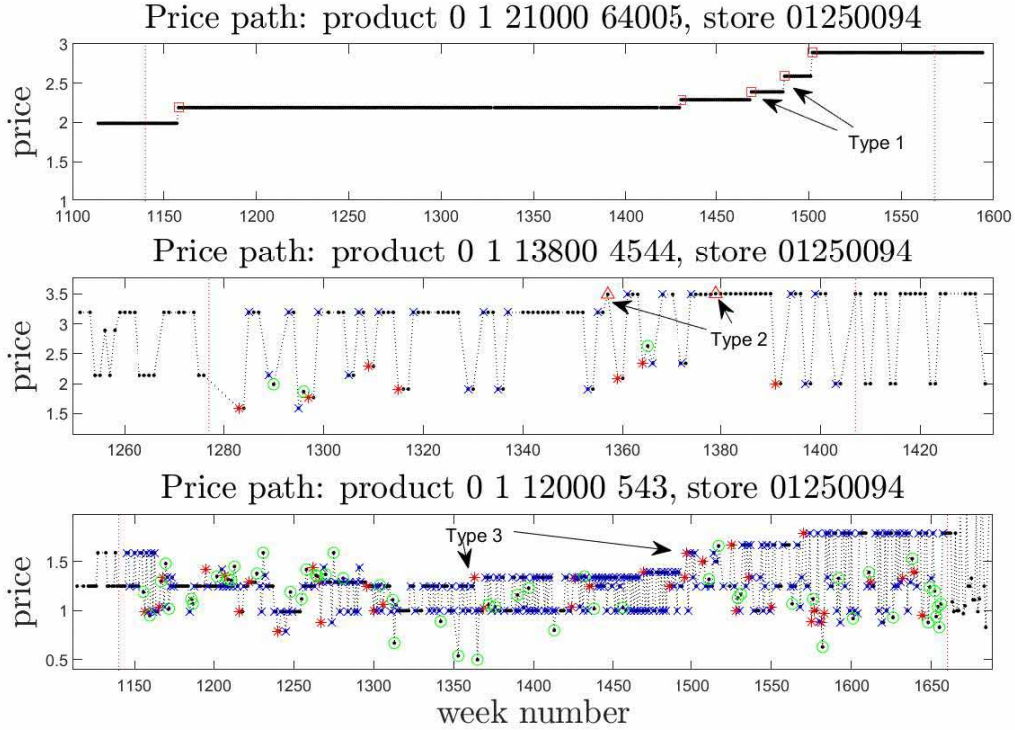
let $\hat{\mu}_{i,j,t}(\tau)$ be the mode of the log price over the same window (i.e. the log of the regular price). As a first measure of excess short-run volatility, we define the difference from the moving average, $p_{i,j,t} - \mu_{i,j,t}(\tau)$, and report the ratio $VR^{avg}(\tau) \equiv \sigma(p_t - \mu_t(\tau))/\sigma(\mu_t(\tau))$ of the standard deviation of the difference to the standard deviation of the moving average.¹³ This ratio is less than one for products in the type 1 subsample, and takes values around 3 in the overall sample, and in the type 3 subsample. Our second measure is the ratio $VR^{diff}(\tau) \equiv \sigma(p_{t+1} - p_t)/\sigma(\mu_{t+1}(\tau) - \mu_t(\tau))$ of the standard deviation of the difference of the log price, to that of the difference of the moving average. This ratio is around 10 in the type 1 subsample, but over 30 for products with type 3 introductions only. Third, we report the $VR^{reg}(\tau) \equiv \sigma(p_{t+1} - p_t)/\sigma(\hat{\mu}_{t+1}(\tau) - \hat{\mu}_t(\tau))$ of the standard deviation of the difference of the log price to that of the difference of the log regular price. CHECK THAT DEFINITION– levels or diff??? This ratio is approximately one in the type 1 subsample, and ranges from 3 to 5 overall, or for products with type 3 introductions only. Finally, $VR^{abs}(\tau) \equiv \sum_t |p_{t+1} - p_t|/\sum_t |\mu_{t+1}(\tau) - \mu_t(\tau)|$ analyzes absolute variation over time: it is the ratio of the sum of absolute differences of the log price to those of the moving average. It is between 1 and 2 for products with type 1 introductions only, but is closer to 20 in the overall sample and in the type 3 subsample.

Since behavior is quite heterogeneous across stores, we next explore pricing patterns from store B, where there are several important quantitative differences (three price series from store B exemplifying type 1, type 2, and type 3 introductions are shown in Fig. 3). First, Table 1 shows that price adjustments are almost twice as frequent at store B, relative to store A. Second, while store B also has some products lacking sales behavior – that is, exhibiting type 1 introductions only – these are an even smaller fraction of the sample than they were in store A. Third, in contrast to store A, transitory price changes are not uncommon at store B. These make up 20.7% of all price changes, and are visible as green circles in the time series in Fig. 3. The fact that transitory changes are relatively common explains why probability mass is shifted right in the histogram in the left panel of Fig. 4. However, previously literature has indicated that many transitory price changes may actually be the result of measurement error; Cavallo (2018) uses daily web-scraped data to demonstrate that apparently transitory prices in weekly scanner data often result from averaging across a mid-week price change. Therefore we will not try to model the transitory changes, which are anyway irrelevant for the patterns we document regarding the stickiness of price points, and the frequency of recurrences and different types of introductions.

In spite of these differences, the key features of the data for our purposes are unchanged between the two stores. In both stores, recurrences are the most common type of price

13. Here, σ indicates standard deviation, and we dropped subscripts i and j for brevity. Each of the volatility ratios is calculated at the product level, and the table then reports the mean across products.

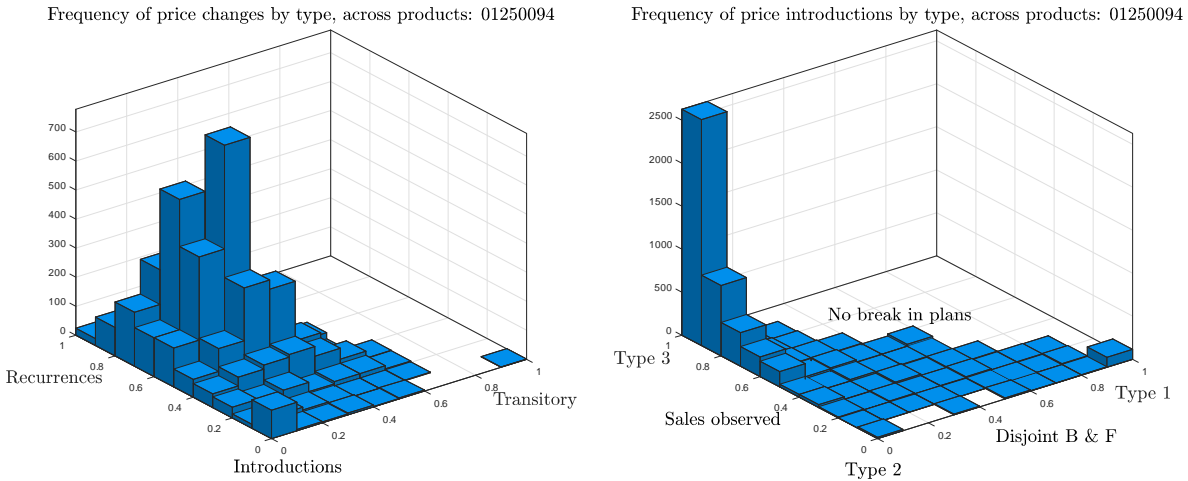
Figure 3: Three examples of price sequences in store B



Notes: Weekly time series of nominal prices for three selected products from store B. Green circles indicate transitory price changes. Blue “X” marks price recurrences. Price introductions are marked by red square (type 1), red triangle (type 2), or red star (type 3); the three products shown are selected because they exhibit especially high frequencies of types 1, 2, and 3 introductions (relative to other introductions), respectively. Red vertical lines highlight the end of the first backwards window and the beginning of the last forwards window (using window length $\tau = 26$).

changes, representing approximately 72% of all non-transitory price changes for the median product. In both stores, for the median product, around 90% of price introductions are type 3. In both stores, type 2 introductions are rare. Strikingly, even in the long series in the third panel of Fig. 3, which shows 36 type 3 introductions, not a single type 2 introduction is observed. Sticky price points are introduced and exit the sample in a staggered fashion, with more than three recurrent price points visible at some times in the series. In summary, while both stores offer a minority of products that feature a single sticky price, with no evidence of sales behavior, the large majority of products display frequent fluctuations across multiple sticky price points. The examples shown in this section are typical of the price series at these stores: the introduction and elimination of price points occurs in a staggered way, and the set of price points is never entirely renewed all at once. While it is possible that some other retailers have different price management strategies, it is clear that the data from these stores are inconsistent with the class of models in which a fixed cost is paid to reset the entire “price

Figure 4: Frequencies of pricing events, across products, in store B.



Notes: Figures represent histograms of sample frequencies of different types of price change events, across products, for store B.

Left: Histogram classifying price changes as changes, recurrences, or introductions. Each point on the simplex represents the sample frequencies of the three types of changes, for a given product. The height of the bars represents the number of observed products in the given bin of sample frequencies.

Right: Histogram classifying price introductions as types 1, 2, or 3. Each point on the simplex represents the sample frequencies of the three types of introductions, for a given product. The height of the bars represents the number of observed products in the given bin of sample frequencies.

plan”. Therefore, we will build an alternative model in which a set of sticky nominal price points is updated in a staggered fashion.

4 An empirical model of sticky regular prices

While most products in our data display multiple sticky price points, a minority of products are instead characterized by a single sticky nominal price which varies intermittently over time. We begin by modelling the dynamics of a single sticky price, following [Costain and Nakov \(2019\)](#).^{14,15} We will define a CC model which is easily mapped into retail price data, but represents the solution of a STMRI model. Then, in the next section, we generalize our framework to address multiple sticky price points:

14. Since the final publication dates might otherwise give a misleading impression of the relation between the two papers, we emphasize that our article appeared in two working paper series (as Banco de España WP 1301 and ECB WP 1693) before the first version of [Steiner et al. \(2017\)](#) circulated (it first appeared in 2015 as CEPR DP 10720).
15. The model developed in this section has two possible applications. First, it may be applied directly to price dynamics for products which are characterized by a single sticky price. Second, one might also filter transitory “sales” out of the data, and then apply this model to the filtered “regular” price data. But to know whether the second procedure gives valid macroeconomic inferences, we should also build a model with multiple sticky price points, and compare its macroeconomic implications to those obtained when applying a model of a single sticky price to the sales-filtered data.

For a given store i , we assume that the time- t nominal price $P_{j,t}$ of good j is chosen from a large, discrete (finite or countably infinite) set of nominal prices, \mathcal{P} . We take as given a system of benchmark distributions $N(\cdot|\mathcal{B})$; that is, for any information set \mathcal{B} , we impose the benchmark distribution $N(\tilde{P}|\mathcal{B}) \in \Delta(\mathcal{P})$ defined over possible nominal prices $\tilde{P} \in \mathcal{P}$. We assume the firm's decision is described by the following control cost problem:

$$V(P, z, \mathcal{B}_\tau, N) = \max_{\pi \in \Delta(\mathcal{P})} \sum_{\tilde{P} \in \mathcal{P}} \pi(\tilde{P}|z, \mathcal{B}_\tau, N) \hat{V}(\tilde{P}, z, \mathcal{B}'_\tau, N) - \kappa W(\mathcal{B}_\tau) \mathcal{D}(\pi || N(\cdot|\mathcal{B}_\tau)) \quad (38)$$

where $\mathcal{B}'_\tau \equiv (\tilde{P}, \mathcal{B}_{\tau-1})$
and $\hat{V}(\tilde{P}, z, \mathcal{B}'_\tau, N) \equiv U(\tilde{P}, z) + E(\delta(z')V(\tilde{P}, z', \mathcal{B}'_\tau, N)|z)$.

This notation supposes that the control variable is the nominal price. The information set contains memories \mathcal{B}_τ of nominal prices (signals received, under an RI interpretation) from the previous τ periods. The post-decision continuation value \hat{V} depends on the nominal price \tilde{P} chosen for the current period, and on the resulting vector of memories \mathcal{B}' , of length τ , which we write as $(\tilde{P}, \mathcal{B}_{\tau-1})$, meaning the concatenation of the price \tilde{P} with the $\tau - 1$ most recent previous prices.¹⁶ Finally, κ parameterizes decision costs in units of time, while $W(\mathcal{B}_\tau)$ represents the DM's nominal value of time.¹⁷ From our previous results, we know that problem (38) is solved by a logit distribution:

$$\pi(\tilde{P}|z, \mathcal{B}_\tau, N) = \frac{N(\tilde{P}|\mathcal{B}_\tau) \exp[\kappa^{-1} W(\mathcal{B}_\tau)^{-1} \hat{V}(\tilde{P}, z, \mathcal{B}'_\tau, N)]}{\sum_{P' \in \mathcal{P}} N(P'|\mathcal{B}_\tau) \exp[\kappa^{-1} W(\mathcal{B}_\tau)^{-1} \hat{V}(P', z, \mathcal{B}'_\tau, N)]}. \quad (39)$$

While the action space \mathcal{P} in (38) is a set of nominal prices, the model can also be described in real terms, which is more convenient for computation and for estimation. Given the nominal price $P_{j,t}$ of a product j at time t , we define its log real price $p_{j,t} \equiv \ln(P_{j,t}/\bar{P}_t)$, where \bar{P}_t is an aggregate price index. We define

$$\mathbf{p}_t \equiv \{p : \bar{P}_t \exp(p) \in \mathcal{P}\} \quad (40)$$

as the set of possible log real prices at time t (the log real prices p corresponding to some

16. We also write the beginning-of-period value V as a function of the beginning-of-period price P , though strictly speaking this is redundant if $\tau > 0$, because P is the first element of \mathcal{B} , so dependence on \mathcal{B} already implies dependence on P .

17. Note that (38) may be regarded as a partial equilibrium problem, or may represent one component of a general equilibrium model. If z is an exogenous random variable affecting a single firm or product, then (38) is a partial equilibrium control cost problem that maximizes the flow of gross nominal payoffs U subject to the shocks z . If instead (38) forms part of a general equilibrium model, then the shock process z must include any shocks that affect the aggregate economy, and W should be written as a function of z too.

nominal price in \mathcal{P}). Similarly, we transform the memories into real terms as follows:

$$\mathbf{b}_\tau^{t-1} \equiv \left(\ln \left(\frac{P_{j,t-1}}{P_t} \right), \ln \left(\frac{P_{j,t-2}}{P_t} \right), \dots, \ln \left(\frac{P_{j,t-\tau}}{P_t} \right) \right). \quad (41)$$

Finally, let w be the real value of time, and let η_t be the benchmark distribution over log real prices $p \in \mathbf{p}_t$ that is equivalent to the nominal benchmark distribution N :

$$w(\mathbf{b}_\tau^{t-1}) \equiv \bar{P}^{-1} W(\mathcal{B}_\tau^{t-1}) \quad (42)$$

$$\eta_t(p|\mathbf{b}_\tau^{t-1}) = N(\bar{P}_t e^p | \mathcal{B}_\tau^{t-1}) \quad (43)$$

We assume the profit function U is homogeneous of degree one (HD1) in nominal variables, which means that the value functions are likewise HD1, and allows us to drop all time subscripts as we rescale these functions into real terms:

$$u(p, z) \equiv \bar{P}^{-1} U(\bar{P} \exp(p), z) \quad (44)$$

$$v(p, z, \mathbf{b}_\tau, \eta) \equiv \bar{P}^{-1} V(\bar{P} \exp(p), z, \mathcal{B}_\tau, N) \quad (45)$$

$$\hat{v}(p, z, \mathbf{b}_\tau, \eta) \equiv \bar{P}^{-1} \hat{V}(\bar{P} \exp(p), z, \mathcal{B}_\tau, N) \quad (46)$$

We can now rewrite the Bellman equation in real terms:

$$v(p, z, \mathbf{b}_\tau, \eta) = \max_{\pi \in \Delta(\mathbf{p})} \sum_{\tilde{p} \in \mathbf{p}} \pi(\tilde{p}|z, \mathbf{b}_\tau, \eta) \hat{v}(\tilde{p}, z, \mathbf{b}'_\tau, \eta) - \kappa w(\mathbf{b}_\tau) \mathcal{D}(\pi || \eta(\cdot | \mathbf{b}_\tau)) \quad (47)$$

$$\text{where } \mathbf{b}'_\tau \equiv (\tilde{p}, \mathbf{b}_{\tau-1})$$

$$\text{and } \hat{v}(\tilde{p}, z, \mathbf{b}'_\tau, \eta) \equiv u(\tilde{p}, z) + E(\delta(z') v(\tilde{p} - i', z', \mathbf{b}'_\tau - i', \eta) | z),$$

where $i' \equiv \ln(\bar{P}'/\bar{P})$ is the inflation rate between the current period and the next. This problem again reflects the fact that the choice variable is a *nominal* price \tilde{P} . Given the *real* value $\tilde{p} \equiv \tilde{P}/\bar{P}$ of this price when it is chosen at time t , its real value at the beginning of $t+1$ is $\tilde{p} - i'$. Likewise, since this problem writes the *nominal* price memories \mathcal{B}_τ in *real* terms as \mathbf{b}_τ , this real value must be updated at the beginning of $t+1$ by adjusting for inflation.¹⁸ The solution to the real control cost problem (47) is formally equivalent to (39), replacing N with η , \hat{V} with \hat{v} , and so forth.

Next, we decompose the decision into a multi-step process that begins with the option of making no change, and then chooses which new price to set, if the firm chooses to make an adjustment. This way of rewriting the model is useful because these choices are easily

18. We have taken some liberties with notation for the sake of brevity. $\mathbf{b}'_\tau \equiv (\tilde{p}, \mathbf{b}_{\tau-1})$ indicates the vector \mathbf{b}'_τ (of length τ) formed by concatenating \tilde{p} at the beginning of the vector $\mathbf{b}_{\tau-1}$ (of length $\tau-1$). The notation $\mathbf{b}'_\tau - i'$ indicates subtracting the scalar i' from each element of the vector \mathbf{b}'_τ .

mapped into observable actions in the data. When we say that the firm's control variable is its nominal price, we mean that if the firm "does nothing", its nominal price remains unchanged. Therefore the first step, in which the firm decides whether to adjust or not, can be written as follows:

$$v(p, z, \mathbf{b}_\tau, \eta) = \max_{\lambda \in [0,1]} (1 - \lambda) \hat{v}(p, z, \mathbf{b}'_\tau, \eta) + \lambda \tilde{v}(p, z, \mathbf{b}_\tau, \eta) - \kappa w(\mathbf{b}_\tau) \mathcal{D}(\lambda | \bar{\lambda}(\mathbf{b}_\tau)) \quad (48)$$

$$\text{where } \bar{\lambda}(\mathbf{b}_\tau) = 1 - \eta(p | \mathbf{b}_\tau), \quad \text{and } \mathbf{b}'_\tau \equiv (p, \mathbf{b}_{\tau-1}). \quad (49)$$

Here \tilde{v} is the value of adjusting to a new price, given by

$$\tilde{v}(p, z, \mathbf{b}_\tau, \eta) = \max_{\tilde{\pi} \in \Delta(\tilde{\mathbf{p}})} \sum_{\tilde{p} \in \tilde{\mathbf{p}}} \tilde{\pi}(\tilde{p} | p, z, \mathbf{b}_\tau, \eta) \hat{v}(\tilde{p}, z, \mathbf{b}'_\tau, \eta) - \kappa w(\mathbf{b}_\tau) \mathcal{D}(\tilde{\pi} | \tilde{\eta}(\cdot | \mathbf{b}_\tau)) \quad (50)$$

$$\text{where } \tilde{\mathbf{p}} \equiv \mathbf{p} \setminus \{p\}, \quad \tilde{\eta}(\tilde{p} | \mathbf{b}_\tau) = \frac{\eta(\tilde{p} | \mathbf{b}_\tau)}{1 - \eta(p | \mathbf{b}_\tau)}, \quad \text{and } \mathbf{b}'_\tau \equiv (\tilde{p}, \mathbf{b}_{\tau-1}). \quad (51)$$

In (48) and (50), \hat{v} represents the real value of selling at a given log real price p :

$$\hat{v}(p, z, \mathbf{b}'_\tau, \eta) \equiv u(p, z) + E(\delta(z') v(p - i', z', \mathbf{b}'_\tau - i', \eta) | z). \quad (52)$$

Given this two-step representation of the problem, we can describe the decision as a two-step nested logit. In the first step, the adjustment probability is a weighted binary logit:

$$\lambda(p, z, \mathbf{b}_\tau, \eta) \equiv \frac{\bar{\lambda}(\mathbf{b}_\tau) \exp\left(\frac{\tilde{v}(p, z, \mathbf{b}_\tau, \eta)}{\kappa w(\mathbf{b}_\tau)}\right)}{\bar{\lambda}(\mathbf{b}_\tau) \exp\left(\frac{\tilde{v}(p, z, \mathbf{b}_\tau, \eta)}{\kappa w(\mathbf{b}_\tau)}\right) + (1 - \bar{\lambda}(\mathbf{b}_\tau)) \exp\left(\frac{\hat{v}(p, z, \mathbf{b}_\tau, \eta)}{\kappa w(\mathbf{b}_\tau)}\right)} \in [0, 1]. \quad (53)$$

Which price is chosen, conditional on adjustment, is then given by a weighted multinomial logit:

$$\tilde{\pi}(\tilde{p} | p, z, \mathbf{b}_\tau, \eta) \equiv \frac{\tilde{\eta}(\tilde{p} | \mathbf{b}_\tau) \exp\left(\frac{\hat{v}(\tilde{p}, z, \mathbf{b}_\tau, \eta)}{\kappa w(\mathbf{b}_\tau)}\right)}{\sum_{p' \in \tilde{\mathbf{p}}} \tilde{\eta}(p' | \mathbf{b}_\tau) \exp\left(\frac{\hat{v}(p', z, \mathbf{b}_\tau, \eta)}{\kappa w(\mathbf{b}_\tau)}\right)} \quad (54)$$

The following proposition summarizes the equivalence relations between the three recursive CC problems (38), (47), and (48)-(52), and also states conditions under which they are equivalent to a sequential RI problem, (55) below.

Proposition 6

(i). Let $\pi(P | z, \mathcal{B}_\tau, N)$ be the solution of the CC problem (38) when $N(P | \mathcal{B}_\tau)$ equals the marginal distribution of P , that is, $N(P | \mathcal{B}_\tau) = E_z \pi(P | z, \mathcal{B}_\tau, N)$. Then $\pi(P | z, \mathcal{B}_\tau, N)$ also

solves the following STMRI problem:

$$U(\mathcal{B}_\tau^0) = \max_{\pi(P_t|z_t, \mathcal{B}_\tau^{t-1}) \in \Delta(\mathcal{P})} E \left\{ \sum_{t=1}^{\infty} \delta^t (U(P_t, z_t) - \kappa W(\mathcal{B}_\tau^{t-1}) \mathcal{I}(P_t, z_t | \mathcal{B}_\tau^{t-1})) \middle| \mathcal{B}_\tau^0 \right\}. \quad (55)$$

(ii). Let $U(P, z)$ be HD1 in P , and let \mathbf{p}_t , w , η_t , u , v , and \hat{v} be defined by (40)-(46). Then the real CC problem (47) is equivalent to the nominal CC problem (38).

(iii). Let the benchmark probabilities $\bar{\lambda}$ and $\tilde{\eta}$ be given by (49) and (51), respectively. Then the multi-step CC problem (48)-(52) is equivalent to the one-step CC problem (47).

Part (i) is a direct application of Props. 3 and/or 4, which show that an RI or STMRI problem can be represented as a CC problem with an optimal benchmark distribution. Part (ii) follows from rescaling all value functions of the firm by \bar{P} when profits are HD1. Part (iii) is an application of Prop. 1, the fact that control cost models can be broken down in a step-by-step fashion, after an appropriate transformation of the benchmark probabilities. Two examples will now help clarify the type of stickiness implied by our model.

Example 1 Suppose $\mathcal{A}_t = \mathcal{P} = \mathcal{R}_+$ and the price process solves the STMRI problem (55). Suppose U is quadratic, and z_t and \bar{P}_t are i.i.d. random variables. Then for each possible information set \mathcal{B} , prices are chosen from the set of nominal prices $\mathcal{Q}(\mathcal{B})$ that satisfy (33). $\mathcal{Q}(\mathcal{B})$ is a strict subset of \mathcal{R}_+ , and generically, it is a discrete, finite set.

Example 2 Suppose $\mathcal{A}_t = \mathcal{P} = \mathcal{R}_+$ and the price process solves the STMRI problem (55). Suppose U is quadratic, and \bar{P}_t is a nonstationary random variable with continuous support, but z_t and $i_t \equiv \ln(\bar{P}_t/\bar{P}_{t-1})$ are i.i.d. Then for each possible real information set \mathbf{b} , prices are chosen from the set of real prices $\mathbf{q}(\mathbf{b})$ that satisfy (33). $\mathbf{q}(\mathbf{b})$ is a strict subset of \mathcal{R} , and generically, it is a discrete, finite set.

The multi-step decision procedure in Prop. 6(iii) motivates us to consider Examples 1 and 2, which illustrate the importance of carefully defining the choice set in a rational inattention model. Example 1 considers the case of STMRI (or RI) when the action set is the nonnegative real line: the firm may set any nonnegative nominal price. In this case, if shocks to the model are stationary *in nominal terms*, then the firm faces exactly the same problem (regardless of t) whenever it faces the same information set. Therefore, by Prop. 5, the firm randomizes over the same set of nominal prices $\mathcal{Q}(\mathcal{B})$, and with the exception of knife-edge cases, this is a discrete, finite set of numbers in spite of the fact that the underlying choice set is a continuum. Therefore, there is a strictly positive probability that the firm chooses exactly the same nominal price at time $t + 1$ that it chose at t . This is the form of price stickiness, based on RI, that was explored by Matějka (2016).

However, Example 2 points out that RI, by itself, does not imply *nominal* price stickiness. It assumes the *same* choice set as Example 1: the firm may set any nonnegative *nominal* price. But while Example 1 assumed the nominal aggregate price index was stationary, Example 2 instead considers the more realistic case in which the inflation rate is stationary, making the log aggregate price level $I(1)$. If z is also stationary, then the model becomes stationary in *real* terms, so the firm faces exactly the same real problem (regardless of t) whenever it faces the same real information set. Therefore, by Prop. 5, the firm randomizes over the same set of real prices $\mathbf{q}(\mathbf{b})$, and with the exception of knife-edge cases, this is a discrete, finite set of numbers in spite of the fact that the underlying choice set is equivalent to $\mathbf{p} = \mathcal{R}$. Therefore, there is a nonzero probability that the firm chooses exactly the same *real* price at time $t + 1$ that it chose at t . But of course, this means there cannot be nominal stickiness; given the continuous support of the aggregate price level \bar{P}_t , the probability of setting exactly the same nominal price in periods t and $t + 1$ is zero.

At the risk of stating the obvious, let us pursue these points further. The aggregate price level is not stationary in the data.¹⁹ Therefore the setting considered in Example 1 is empirically rejected. But we have seen in the microdata that setting exactly the same nominal price level from one week to the next is a common event. Therefore, applying Example 2 to the data rejects the joint hypothesis that prices are constrained by rational inattention, and that the firm chooses from the set of all nonnegative prices, $P \in \mathcal{R}_+$. The fact that nominal prices are in fact sticky from one week to the next could of course be explained either by assuming a menu cost, or by assuming that nominal prices are chosen from a discrete subset of \mathcal{R}_+ , for example, all positive integer multiples of 0.01. But neither of these mechanisms, alone, would explain stickiness of multiple price points, as we documented in Sec. 3.

Instead, we maintain that RI and CC models can explain nominal price stickiness as long as we define the action set in a more plausible way. RI and CC models impose no restrictions whatsoever on the form of the choice set, so we are free to be realistic when we set up the problem. In practice, nominal retail prices are set by actions such as marking a number on a price tag, or typing a number into a database, or programming an algorithm that determines future nominal prices as a function of a number of future inputs. In the first two cases, the nominal price remains unchanged if the firm “does nothing”: it only changes when the firm makes a new, deliberate decision.²⁰ Therefore, we propose that the firm’s action set is better described as having two qualitatively different subsets: it may either “do nothing”, in which case its nominal price remains unchanged, or it may choose a new price. The problem is then

19. Cite here any ECB study showing that the price level is not $I(0)$.

20. Third case obviously increasingly important; implies nominal prices may change at each transaction. In this case, the firm’s deliberate decisions involve changing the program (a costly choice of a new *function* linking input data with nominal prices). Modeling the intermittent updating of a function is beyond the scope of this paper.

conveniently written in the multi-step formulation (48)-(52), though it may still be written in the single-step forms (38) or (47) by an appropriate rescaling of benchmark probabilities.²¹

We now discuss how to take the model (48)-(52) to the data. This formulation is convenient because it allows us to observe the benchmark probability of non-adjustment – an easily observable event in the data – directly by revealed preference, instead of computing it within the model.²² Subsequently, we will also discuss how to extend this model to allow for multiple sticky price points.

4.1 Zero-memory simulations: sticky regular prices

Our theoretical results suggest a tractable methodology for modelling retail price microdata in a control cost framework. They also suggest an iterative procedure for fitting a rational inattention model to the data.²³ The zero-memory rational inattention model ($\tau = 0$) is a control cost model in which the benchmark action distribution is the *unconditional* distribution of actions observed in the data. Taking as given the observed benchmark distribution, the problem reduces to computing the value functions and policy functions in the space (p, z) of prices and shocks, which means they can be computed to high accuracy by grid-based methods, both in partial equilibrium and in general equilibrium.²⁴

Allowing the firm to use some memory will improve its decision-making, but makes its state space too large to be conveniently solved by grid-based methods. Therefore, we will instead approximate the value function using a neural network, since these can handle a higher-dimensional state space.²⁵ As a first step, we will fit a neural network to the grid-based, zero-memory value function, which amounts to smoothly interpolating and extrapolating the solution found at the grid points. We report results based on the zero-memory neural network solution in this section. Next, since using memory will slightly improve the value of the firm’s decision, we use a neural network to solve the Bellman equation in the case of memory, initializing from the solution of the zero-memory problem. The results of the extended model, with memory, are reported in Sec. 5.1.

In the present section, we compare two specifications of the zero-memory model. First, we assume that each firm faces a constant elasticity demand function, with elasticity $\epsilon = 7$. Second, we consider an environment with an incentive for stochastic price discrimination,

21. If the choice set includes “do nothing”, but we rewrite the problem in a single-step form, then the benchmark probability assigned to a given nominal price $P \in \mathcal{P}$ is likely to vary greatly depending on whether P coincides with the previous nominal price or not.

22. NEED DISCUSSION of the options to infer benchmarks by observation, then TEST by solving the model, vs. solving the model directly.

23. See Appendix B for more details on our computational approach.

24. Details of a grid-based solution are discussed in the appendix of Costain and Nakov (2019).

25. Our computations are loosely based on Azinović et al. (2019), but we approximate value functions, while they approximate policy functions.

in which any given product is demanded by a subset of loyal consumers with low elasticity $\epsilon_L = 3$, and a subset of bargain hunters with high elasticity $\epsilon_H = 11$. We impose the demand function derived by [Guimaraes and Sheedy \(2011\)](#), who show that with a mixture of sufficiently heterogeneous consumers, the firm has an incentive to play a mixed strategy, randomly setting high or low prices. This allows us to compare the incentives for randomization derived from costly decision-making with those derived from stochastic price discrimination.

Table 2: Common parameters

<i>Symbol</i>	<i>Value</i>	<i>Description</i>
<i>Central Bank</i>		
μ	1	Gross money growth
<i>Households</i>		
β	0.9967	Discount factor
γ	0.5	Intert. elast. of subst.
χ	6	Utility weight of labor
<i>Retail Firms</i>		
κ	0.01767	Noise
p^B	1	Product-specific price index
ρ	0.9	Persistence of productivity
std	0.12	Std dev of productivity


Parameter values that are common across simulations are presented in Table 2. The parameters are based mostly on [Costain and Nakov \(2019\)](#), which estimated the noise parameter $\kappa = 0.01767$, conditional on preference parameters taken from [Goloso and Lucas Jr \(2007\)](#). We compute empirical counterparts to the benchmark action distribution by selecting products from store A of the IRI dataset which are characterized primarily by type 3 price introductions— in other words, we select products that appear to exhibit sales behavior.²⁶ In this sample, the unconditional adjustment probability is 0.58 per month, so this is our calibration for the monthly value of the parameter $\bar{\lambda}$. Likewise, for this sample, we compute the distribution of newly-set log prices, expressed as deviations from the product-specific mean. This distribution is plotted as a histogram in Figure B, and we impose this distribution as the benchmark distribution $\tilde{\eta}(\tilde{p})$ over newly-set prices \tilde{p} .

TO DISCUSS:

- Fig. 5 displays unconditional distribution $\tilde{\eta}$ of new prices, from data.
- We first solve the Bellman equations using the observed distribution $\tilde{\eta}$ as the benchmark distribution. This is step 2 of the algorithm described in Appendix B. The value function and adjustment hazard function are shown in the upper panels of Figs. 7 and 8.


26. In this exercise, we use data characterized by frequent type 3 adjustments both in this section (the zero memory model) and the next (the finite memory model). Thus our focus here is on comparing how changes in model specification affect behavior, instead of fitting different versions of the model to different datasets.

Figure 5: Empirical histogram of price introductions, $\tilde{\eta}(\tilde{p})$.



Notes: Histogram of newly-introduced log prices, store A, after removing product-specific mean.

Figure 6: Benchmark adjustment frequencies: Blahut-Arimoto fixed point, $\tilde{\eta}(\tilde{p})$.



Notes: Benchmark frequencies newly-introduced log prices, calculated from Blahut-Arimoto algorithm, for four model specifications.

- However, the observed $\tilde{\eta}$ represents a mix of price distributions across heterogeneous products.
- Therefore, in the next step, we find the fixed point $\eta^*(p)$ of the benchmark distribution, using equations (27)-(28). This is step 2A of the algorithm described in Appendix B.
- As expected, the fixed point $\eta^*(p)$ is a discrete distribution (over a finite number of points), unlike the continuous distribution $\tilde{\eta}$ taken from the data.
- The price sequences shown in the lower panels of Figs. 7 and 8, and the price statistics in the “no memory” columns of Table 3, are derived using the value function from step 2 and the benchmark distribution from step 2A.

Figures 7 and 8 show numerical results from the zero-memory model, assuming homogeneous and heterogeneous substitution elasticities, respectively.²⁷ All parameters other than these elasticities are identical in the two simulations. The upper panels of these figures show the value function $v(p, z)$ and the adjustment probability function $\lambda(p, z)$. While the value function appears very flat over much of its range, this is largely a matter of scale. The subtle variations in the value function suffice to cause large changes in the adjustment hazard. Along a line of points near the 45° line, the adjustment probability falls almost to zero; away from this line the adjustment probability climbs rapidly but smoothly to one. The band of nonadjustment (resembling the S,s bands of a menu cost model) is narrow when costs are low, due to the high stakes involved in selling large quantities at a low price. At higher costs, the nonadjustment band widens; the losses from setting an excessively high price are relatively small in this range, especially in the heterogeneous elasticity case (but the losses from setting a low price when costs are high are huge, so the adjustment hazard is almost exactly one when costs are high and prices are low).

The figures also illustrate the ergodic distribution, by overlaying blue dots representing a population of 2000 products on top of the value function and hazard function. With constant demand elasticity (Fig. 7), the price distribution shifts smoothly towards higher values as costs increase. With heterogeneous elasticities (Fig. 8), we instead see a break in the distribution, because firms avoid intermediate prices, either setting a low price directed mainly towards high-elasticity consumers, or a high price targeted to low-elasticity types.

The implications for price dynamics can be seen in the bottom panels of the figures, which show a sample of twelve simulated series of costs (red) and prices (blue) over 500 weeks. In the constant elasticity case (Fig. 7), we see that although prices are sticky, they coarsely track the ups and downs of costs, with a roughly constant markup of $\frac{1}{\epsilon-1} = 1/6$. In the heterogeneous

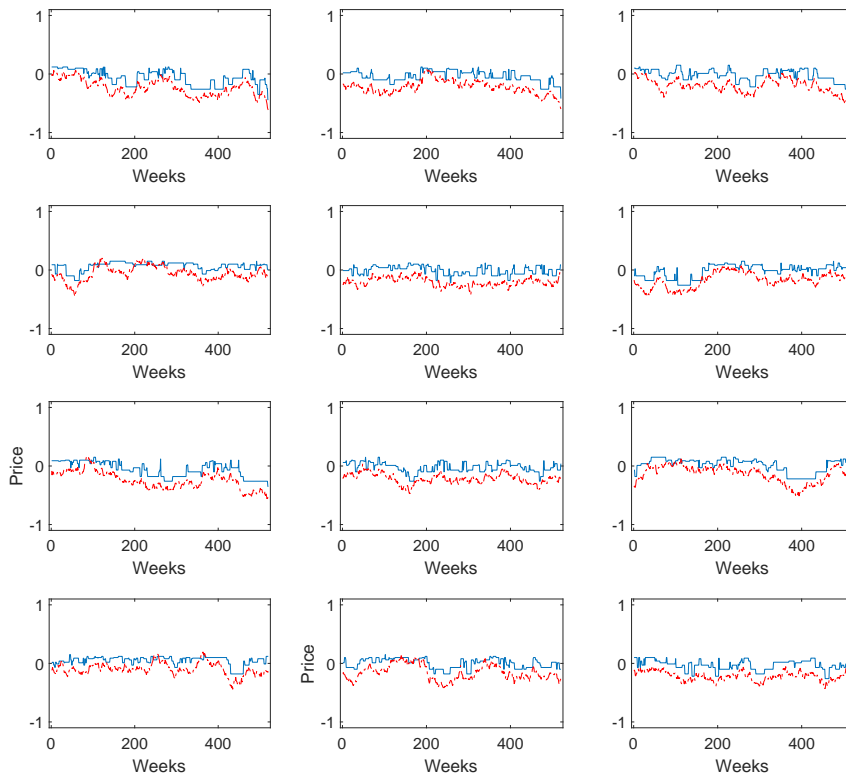

27. GRAPHS FOR THIS VERSION: graphs/corrected-191206.

elasticity case (Fig. 8), we instead see variable markups: when costs rise, firms tend to raise their product prices even more, catering more to the low-elasticity consumers.

Nonetheless, the variable markups of the heterogeneous elasticity case do little to create excess short-run volatility. Table ?? below (first and third columns) reports excess volatility ratios from our simulations, showing that these ratios are all far too low in the zero-memory model, compared to the data reported in Table 1. For example, the excess volatility ratio $VR^{avg}(\tau)$ is 0.98 in the simulation with zero memory and $\epsilon = 7$, while it is 2.7 in the overall data sample from Store A. Likewise, comparing the same simulation and dataset, the excess volatility ratio for regular prices is $VR^{reg}(\tau) = 1.52$ in the simulation, but equals 4.1 in the data. Interestingly, variable markups do not change this conclusion— the volatility ratios actually fall slightly as we move from $\epsilon = 7$ to the heterogeneous demand specification with $\epsilon_L = 3$ and $\epsilon_H = 11$. Intuitively, prices in the zero-memory model track slow-moving costs over time, so even though heterogeneous demand raises volatility by making markups variable, it does not raise short-run volatility more than it raises long-run volatility.

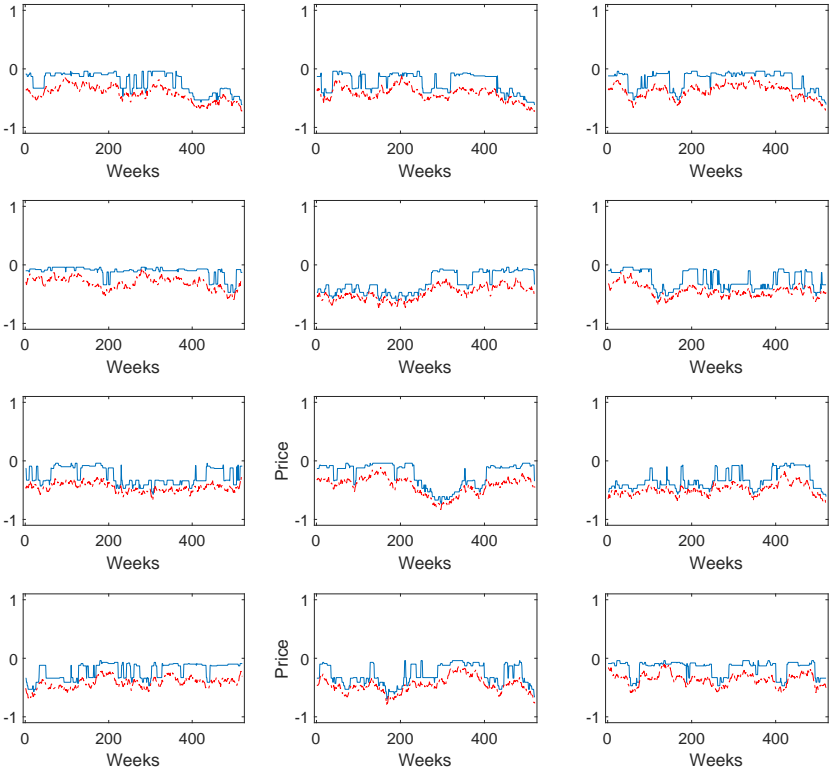

Finally, Table ?? also shows that the zero-memory model does little to produce the patterns of jumps across multiple sticky price points that we see in the data. At store A, roughly 2/3 of price adjustments are recurrences, and almost 90% of introductions are type 3. In the simulations of the model without memory, only 33%-40% of price adjustments are recurrences, and roughly three-quarters of introductions are type 3. We next turn to the model with memory, which better matches the stickiness and jumpiness of the data.

Figure 7: Price adjustment: no memory, substitution elasticity 7



Notes: Simulations of the zero-memory price adjustment model, with elasticity of substitution $\epsilon = 7$.
 Top panels: Value function (left) and adjustment hazard (right) as functions of log cost and current log real price. Simulated distribution shown by blue dots on top of the value function and hazard function.
 Bottom panels: Twelve simulated time series, showing log costs (red) and log real price (blue).

Figure 8: Price adjustment: no memory, heterogeneous substitution elasticities (3 and 11)



Notes: Simulations of the zero-memory price adjustment model, with mix of substitution elasticities in the population, $\epsilon_L = 3$ and $\epsilon_H = 11$.
 Top panels: Value function (left) and adjustment hazard (right) as functions of log cost and current log real price. Simulated distribution shown by blue dots on top of the value function and hazard function.
 Bottom panels: Twelve simulated time series, showing log costs (red) and log real price (blue).

Table 3: Parameter specifications and simulation results (medians)*

<i>Parameter specifications</i>					
	7 NoM	7 Mem	311 NoM	311 Mem	Description
ϵ^{LO}	7	7	3	3	Low elasticity of subst.
ϵ^{HI}	7	7	11	11	High elasticity of subst.
ϕ	1	1	0.9	0.9	Fraction of loyal customers
W	0	12	0	12	Size of memory window
λ	0.58	0.58	0.58	0.58	Benchmark price change hazard (monthly)
λ_Q	0.2587	0.2587	0.2587	0.2587	Benchmark prob. of choosing outside Q
<i>Simulation accuracy</i> [†]					
	2.686	1.866	0.2726	0.2572	Euler residual error x100
<i>Profitability</i> ^{††}					
	0.8667	0.8673	0.7075	0.7183	Mean firm profits as share
	0.9197	0.9497	0.3109	0.3052	Mean firm revenues as share
	0.01904	0.01896	0.01448	0.01381	Mean profit loss as share
<i>Price changes</i> *					
	0.1547	0.1382	0.1572	0.1524	Frequency of adjustment (weekly)
<i>Classifying price changes</i> *					
	58.9	29.5	54.5	30.2	Frequency of introductions
	33.3	67.2	40.0	66.2	Frequency of recurrences
	7.5	3.1	5.2	3.2	Freq. of transitory changes
	8.7	0.0	5.4	0.0	Freq. of type 1 introductions
	15.4	5.9	17.4	5.3	Freq. of type 2 introductions
	75.0	90.9	75.7	93.6	Freq. of type 3 introductions
<i>Short-run volatility</i> *					
	0.98	0.95	0.85	1.33	Ratio $VR^{avg}(\tau)$
	16.5	17.5	14.1	23.3	Ratio $VR^{diff}(\tau)$
	1.52	1.68	1.36	2.36	Ratio $VR^{reg}(\tau)$
	6.8	7.2	5.6	10.4	Ratio $VR^{abs}(\tau)$
	0.41	0.44	0.65	1.78	Ratio $CR^{avg}(\tau)$
	1.28	1.39	1.97	6.2	Ratio $CR^{diff}(\tau)$
	0.52	0.49	0.72	2.5	Ratio $CR^{abs}(\tau)$

Notes: Calibrated parameters and equilibrium statistics from four simulations at weekly frequency: homogeneous (7) vs. heterogeneous (3 and 11) elasticity, and zero memory case vs. memory of the previous 12 months' prices.

*All statistics in sections with asterisks are reported as medians across simulated product histories.

[†]Median Euler residual, expressed as a fraction of consumption, times 100.

^{††}Mean loss in profits of decision-constrained firm, as a fraction of the median revenues of a flexible-price firm.

5 An empirical model of sticky price points

Thus far, we have only worked with the zero-memory limit of the STMRI model. Next, we allow for a finite, nonzero memory. This is interesting for several reasons. First, for a given signal processing capacity, greater memory improves payoffs (CITE). Second, as memory increases, STMRI converges to the infinite-memory RI model that has been the main focus of previous literature, including [Steiner et al. \(2017\)](#). Third, infinite memory is unrealistic, so some model with limited memory and limited signal processing capacity is likely to fit microdata better than the alternatives. Fourth, limited memory immediately suggests an explanation for sticky price points: perhaps firms can economize on information processing by sometimes repeating prices that they already chose in the recent past. [Wilson \(2014\)](#) also argues that limited memory helps explain sticky behavior.

Hence, we generalize our model in a simple way that may help generate multiple sticky price points, taking advantage of Prop. 1 to further decompose the firm's decision into three steps. In the first step, the firm chooses whether to adjust or not, as in the previous section:

$$v(p, z, \mathbf{b}_\tau, \eta) = \max_{\lambda \in [0,1]} (1 - \lambda) \hat{v}(p, z, \mathbf{b}'_\tau, \eta) + \lambda \tilde{v}(p, z, \mathbf{b}_\tau, \eta) - \kappa w(\mathbf{b}_\tau) \mathcal{D}(\lambda || \bar{\lambda}(\mathbf{b}_\tau)) \quad (56)$$

$$\text{where } \bar{\lambda}(\mathbf{b}_\tau) = 1 - \eta(p | \mathbf{b}_\tau), \quad \text{and } \mathbf{b}'_\tau \equiv (p, \mathbf{b}_{\tau-1}). \quad (57)$$

In a second step, we break down the value of choosing a new price, \tilde{v} , into two components. The firm considers whether to return to one of the prices it remembers setting in the past (choosing from the set $\mathbf{b}_\tau \setminus \{p\}$), or to choose a different point not in memory (choosing from the set $\mathbf{p} \setminus \mathbf{b}_\tau$):

$$\tilde{v}(p, z, \mathbf{b}_\tau, \eta) = \max_{\mu \in [0,1]} (1 - \mu) x(p, z, \mathbf{b}_\tau \setminus \{p\}, q_m; \mathbf{b}_\tau, \eta) + \mu x(p, z, \mathbf{p} \setminus \mathbf{b}_\tau, q; \mathbf{b}_\tau, \eta) - \kappa w(\mathbf{b}_\tau) \mathcal{D}(\mu || \bar{\mu}(\mathbf{b}_\tau)), \quad (58)$$

$$\text{where } q_m(\tilde{p}) = \frac{\eta(\tilde{p} | \mathbf{b}_\tau)}{\sum_{p' \in \mathbf{b}_\tau \setminus \{p\}} \eta(p' | \mathbf{b}_\tau)} \text{ for } \tilde{p} \in \mathbf{b}_\tau \setminus \{p\}, \quad \text{and } q(\tilde{p}) = \frac{\eta(\tilde{p} | \mathbf{b}_\tau)}{\sum_{p' \in \mathbf{p} \setminus \mathbf{b}_\tau} \eta(p' | \mathbf{b}_\tau)} \text{ for } \tilde{p} \in \mathbf{p} \setminus \mathbf{b}_\tau. \quad (59)$$

Here, q_m is a benchmark distribution over prices in the memory set \mathbf{b}_τ , q is a benchmark distribution over prices outside the memory set, and x is the value of choosing a price from within a subset $\mathcal{Q} \subseteq \mathbf{p}$:

$$x(p, z, \mathcal{Q}, q; \mathbf{b}_\tau, \eta) = \max_{\tilde{\pi} \in \Delta(\mathcal{Q})} \sum_{\tilde{p} \in \mathcal{Q}} \tilde{\pi}(\tilde{p} | p, z, \mathcal{Q}, q; \mathbf{b}_\tau, \eta) \hat{v}(\tilde{p}, z, \mathbf{b}'_\tau, \eta) - \kappa w(\mathbf{b}_\tau) \mathcal{D}(\tilde{\pi} || q) \quad (60)$$

$$\text{where } \mathbf{b}'_\tau \equiv (\tilde{p}, \mathbf{b}_{\tau-1}).$$

Finally, as before, \hat{v} represents the continuation value at a given log real price p :

$$\hat{v}(p, z, \mathbf{b}'_{\tau}, \eta) \equiv u(p, z) + E(\delta(z')v(p - i', z', \mathbf{b}'_{\tau} - i', \eta)|z). \quad (61)$$

Proposition 7

Let the benchmark probabilities $\bar{\lambda}$, q_b , and q_p be given by (57) and (59). Then the multi-step CC problem (56)-(61) is equivalent to the one-step CC problem (47).

Prop. 7 is simply a direct application of Prop. 1, which showed that CC problems can be broken down in multistep form. Then, following Prop. 6, the multistep model (56)-(61) can also be mapped into the other CC and RI formulations of the price adjustment problem that we considered in the last section.

5.1 Finite memory simulations: sticky price points

TO DISCUSS:

- For quicker calculation, we solve the Bellman equations at monthly frequency.
- For comparability to the data, we then rescale the hazard function in order to simulate the model at weekly frequency.
- As we did for the zero-memory economy, we solve the Bellman equations using the unconditional distribution of new prices as the benchmark distribution (step 3 of the algorithm in Appendix B).
- As we did for the zero-memory economy, we then solve for the benchmark distribution as the fixed point $q^*(p)$ of the Blahut-Arimoto equations (27)-(28). This is step 3A of the algorithm described in Appendix B.
- As expected, the fixed point $q^*(p)$ is a discrete distribution (over a finite number of points), unlike the continuous distribution taken from the data. This distribution is shown in the lower panels of Fig. 6.
- The price sequences shown in the lower panels of Figs. 9 and 10, and the price statistics in the “memory” columns of Table 3, are derived using the value function from step 3 and the benchmark distribution from step 3A.

Figures (9)-(10) show the value functions and hazard functions for the model with memory (top panels), and simulated time series (bottom panels). Again, we overlay the functions with blue dots representing a sample from the ergodic distribution over costs, prices, and memory

states.²⁸ In the model with memory, we see that prices often revisit values that have been set before. This decreases the standard deviation of the price series in the $\epsilon = 7$ case, but in the heterogeneous demand case (Fig. 10) we see an increase in variability, with patterns of “stickiness” and “jumpiness” resembling those in the retail data.

Quantitatively, Table ?? (second and fourth columns) shows that the introduction of memory greatly improves the model’s fit to a number of price change statistics. The frequency of recurrences rises from 40% or less to 66%-67%, close to the value in the data (compare Table 1). Type 3 introductions are now overwhelmingly the most common type, at 91%-94% for the median simulation, slightly above the value of 89%-91% calculated for the median product in the data sample. The median frequency of Type 1 introductions is zero, as in the data, and the median frequency of Type 2 introductions is 5%-6%, matching the figure calculated for Store B.

Adding memory, in the constant elasticity case, leaves measures of excess short-run volatility largely unchanged. But memory has a big effect on excess volatility when coupled with heterogeneous demand elasticity, because when the firm has an incentive to randomize, any prices it has set recently are relatively likely to be valuable prices to choose again. Therefore, the firm has a tendency to jump back and forth across prices in the backwards window of memory. Hence, in the specification with heterogeneous elasticity and memory, each of the excess short-run volatility ratios rises by 30%-45% beyond the highest value seen in the other specifications. Even so, the model displays less excess short-run volatility than the data: $VR^{avg}(\tau) = 1.33$ in the heterogeneous demand model with memory, compared with 2.7 (2.1) at Store A (B). Likewise, $VR^{reg}(\tau) = 2.36$ in the model but equals 4.1 (3.2) in the data from Stores A (B). Nonetheless, the model suggests that heterogeneity of demand is the decisive factor, rather than frictions in decision-making, for driving the short-run volatility of regular vs. sales prices.

Summarizing, both the STMRI model and heterogeneous demand elasticity appear important for matching the dynamics of retail prices. While costly information acquisition tends to shrink the action distribution down to a set of discrete points, by itself this is more likely to produce *real* than *nominal* price stickiness, as Example 2 showed. Nominal stickiness instead arises when we allow for finite memory of past nominal prices. STMRI then generates stickiness of multiple nominal price points, with frequent recurrences and frequent Type 3 but rare Type 2 price introductions. In particular, limited memory generates intermittent introduction and elimination of nominal price points, instead of causing intermittent renewal

28. In Figs. (7)-(8) the blue dots coincided with the surfaces, because the value and the hazard were computed as functions of p and z only. In Figs. (9)-(10), the value and hazard functions depend on the current state of memory too; the surfaces are depicted conditional on a single memory state, while the blue dots are drawn from the ergodic distribution of memory states. Therefore the dots no longer coincide exactly with the surfaces, but instead illustrate how the value and the hazard vary with the state of memory.

of all price points, as does the fixed cost of “replanning” in the model of Stevens (2019). Still, STMRI is not the only element needed to model the rapid “jumpiness” of sales; under constant demand elasticity, STMRI prices just track costs, in a coarse way, through discrete jumps. To reproduce the high short-run volatility of sales, STMRI must be applied to a context of heterogeneous demand elasticity, generating a motive for randomization. In other words, while the stickiness of multiple nominal price points is explained well by limits on information and memory, the main explanation for the “jumpiness” of sales appears to be consumer heterogeneity, as Varian (1980) originally proposed.

6 Monetary policy shocks (incomplete)

Next, we study how monetary policy affects prices in the environment considered above. To do so, we must fill out some additional details about the general equilibrium structure of the economy. We follow [Guimaraes and Sheedy \(2011\)](#) to introduce heterogeneous consumption elasticities in a way that permits easy aggregation. We assume there exists a continuum of households with the period utility function $\frac{1}{1-\gamma}C_t^{1-\gamma} - \chi N_t + \nu(M_t/P_t)$, where C_t is consumption, N_t is labor supply, and M_t/P_t is real money balances. The household maximizes expected, discounted utility (with discount factor β) subject to its budget constraint.

Consumption C_t is a CES aggregate across differentiated products $C_{k,t}$:

$$C_t = \left\{ \int_0^1 (C_{k,t})^{\frac{\epsilon_L-1}{\epsilon_L}} dk \right\}^{\frac{\epsilon_L}{\epsilon_L-1}}, \quad (62)$$

where ϵ_L is the elasticity of substitution across products. Each product k is produced under many different brands j , consumed in quantities $c_{j,k,t}$. Any given household is *loyal* to a specific brand $j(k)$ for a fraction α of the products k it consumes. For these products, $C_{k,t} = c_{j(k),k,t}$. Each household is a *bargain hunter* for the remaining products, meaning that it consumes a CES aggregate of different brands:

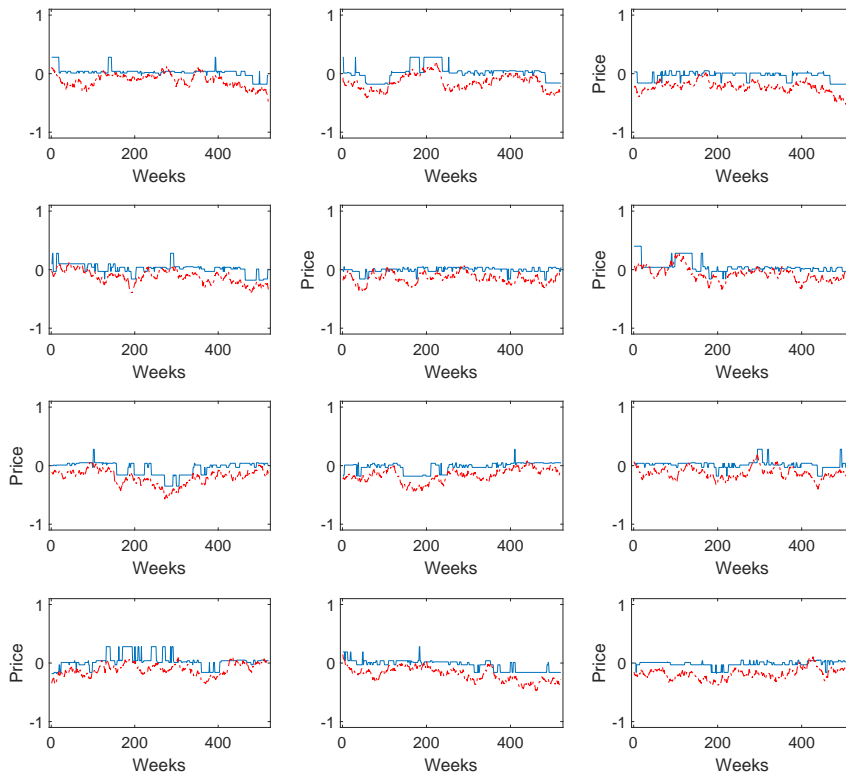

$$C_{k,t} = \left\{ \int_0^1 (c_{j,k,t})^{\frac{\epsilon_H-1}{\epsilon_H}} dj \right\}^{\frac{\epsilon_H}{\epsilon_H-1}}, \quad (63)$$

where $\epsilon_H > \epsilon_L$. As [Guimaraes and Sheedy \(2011\)](#) have shown, by imposing symmetric behavior across all households, products, and brands, we can write total demand for a specific brand j as

$$C_{j,t} = \left[\alpha \left(\frac{P_{j,t}}{P_t} \right)^{-\epsilon_L} + (1-\alpha) \left(\frac{P_t^B}{P_t} \right)^{-\epsilon_L} \left(\frac{P_{j,t}}{P_t^B} \right)^{-\epsilon_H} \right] C_t, \quad (64)$$

where P_t^B is the representative household’s price index across brands, when it hunts for

Figure 9: Price adjustment: memory = 12 months, substitution elasticity 7

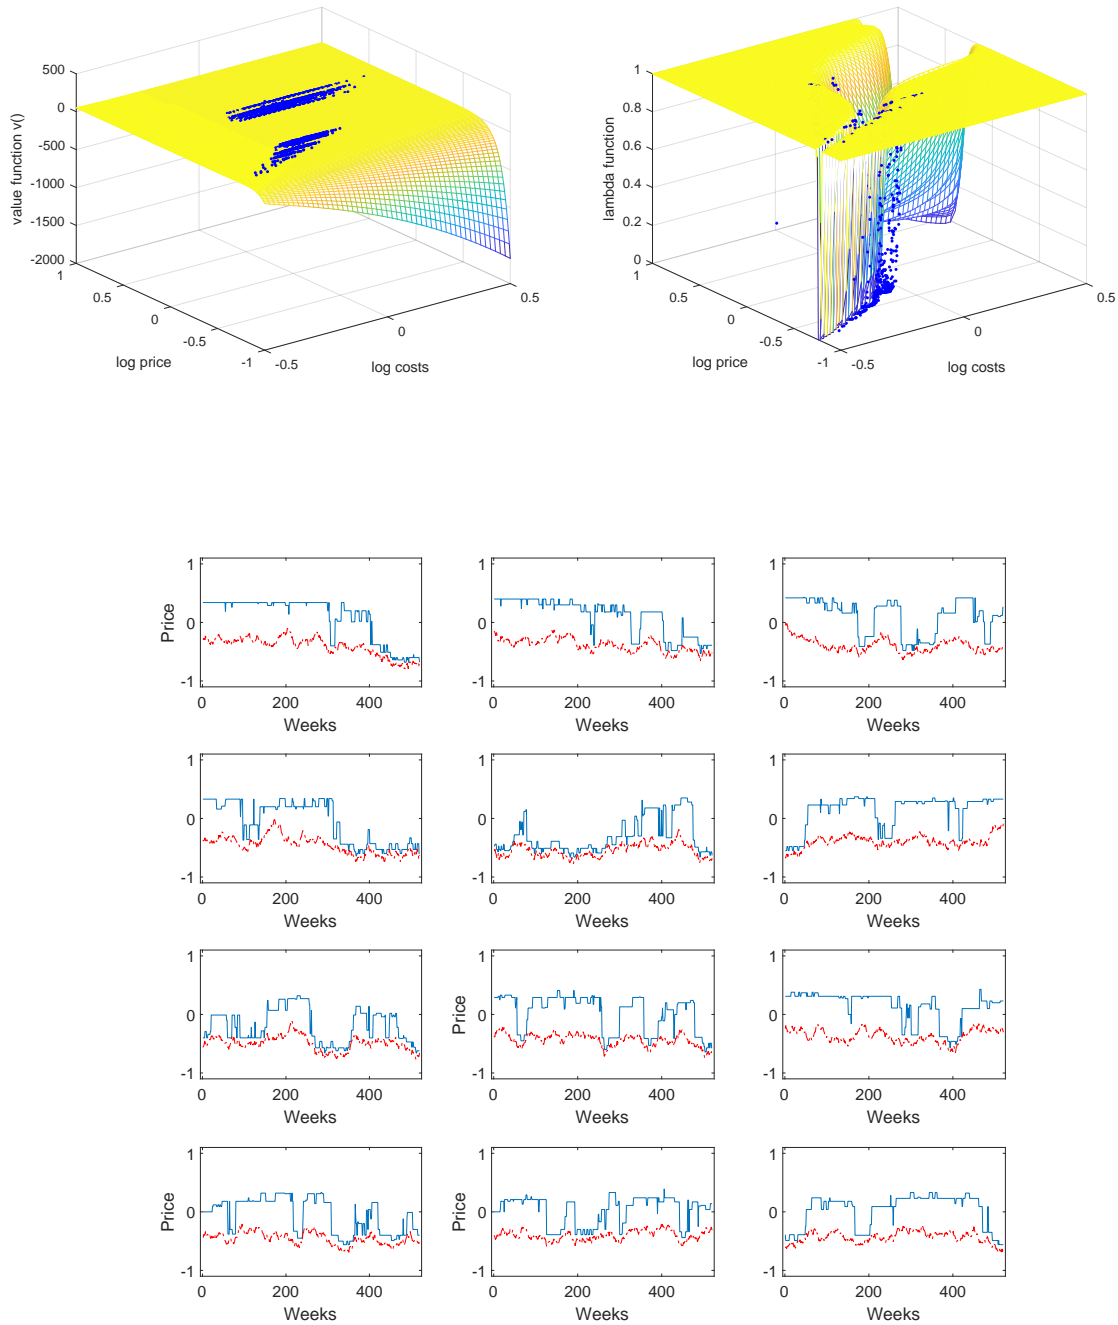


Notes: Simulations of the price adjustment model with twelve months' memory, assuming homogeneous elasticity of substitution $\epsilon = 7$.

Top panels: Value function (left) and adjustment hazard (right) as functions of log cost and current log real price. Simulated distribution shown by blue dots on top of the value function and hazard function.

Bottom panels: Twelve simulated time series, showing log costs (red) and log real price (blue).

Figure 10: Price adjustment: memory = 12 months, heterogeneous substitution elasticities (3 and 11)



Notes: Simulations of the price adjustment model with twelve months' memory, assuming heterogeneous substitution elasticities in the population, $\epsilon_L = 3$ and $\epsilon_H = 11$.

Top panels: Value function (left) and adjustment hazard (right) as functions of log cost and current log real price. Simulated distribution shown by blue dots on top of the value function and hazard function.

Bottom panels: Twelve simulated time series, showing log costs (red) and log real price (blue).

((NOTE: COMPUTED ON BIGGER DATASET. Must choose which one to use for each specification.))

bargains:

$$P_t^B = \left(\int_0^1 (P_{j,t})^{1-\epsilon_H} dj \right)^{\frac{1}{1-\epsilon_H}}, \quad (65)$$

and P_t is the household's price index across products:

$$P_t = \left(\alpha \int_0^1 (P_{j,t})^{1-\epsilon_L} dj + (1-\alpha)(P_t^B)^{1-\epsilon_L} \right)^{\frac{1}{1-\epsilon_L}}, \quad (66)$$

Equation (64) shows that the demand curve for a firm's branded product j exhibits a mixture of two elasticities, ϵ_L and ϵ_H . At low prices, it will primarily face the bargain hunters' elasticity ϵ_H ; at high prices it will sell primarily to loyal customers, with elasticity ϵ_L . When ϵ_L and ϵ_H are sufficiently different, and α lies in a middle range between 0 and 1, the profit function is not concave, and the firm benefits from randomization.²⁹

The household's remaining first-order conditions are:

$$\chi = C_t^{-\gamma} W_t / P_t, \quad (67)$$

$$R_t^{-1} = 1 - \frac{v'(M_t/P_t)}{u'(C_t)} = \frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)}. \quad (68)$$

We close the model by assuming that the nominal money supply is affected by an AR(1) shock process z :³⁰

$$z_t = \phi_z z_{t-1} + \epsilon_t^z, \quad (69)$$

where $0 \leq \phi_z < 1$ and $\epsilon_t^z \sim i.i.d. N(0, \sigma_z^2)$. Here, z_t represents the rate of growth of the money stock:

$$M_t/M_{t-1} \equiv \mu_t = \mu \exp(z_t). \quad (70)$$

((DESCRIBE DISTRIBUTIONAL DYNAMICS HERE? OR IN APPENDIX?))

((TO BE COMPLETED))

7 Conclusions

((TO BE COMPLETED))

29. See [Guimaraes and Sheedy \(2011\)](#), Theorem 1.

30. Alternatively, do a Taylor rule. Both are feasible.

References

- Akerlof, G. and J. Yellen (1985). A near-rational model of the business cycle with wage and price inertia. *Quarterly Journal of Economics* 100(Supplement), 823–838.
- Anderson, S. P., A. de Palma, and J.-F. Thisse (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press.
- Ash, R. B. (1990). *Information Theory*. Dover.
- Azinović, M., L. Gaegouf, and S. Scheidegger (2019). Deep equilibrium nets. Manuscript, ETH Zurich.
- Blahut, R. E. (1972). Computation of channel capacity and rate-distortion functions. *IEEE Transactions on Information Theory* IT-18(4), 1118–46.
- Bronnenberg, B., M. Kruger, and C. Mela (2008). Database paper – the iri marketing data set. *Marketing Science* 27(4), 745–748.
- Carvalho, C. and O. Kryvtsov (2018). Price selection. Bank of Canada Staff WP 2018-44.
- Cavallo, A. (2018). Scraped data and sticky prices. *The Review of Economics and Statistics* 100(1), 105–119.
- Coibion, O., Y. Gorodnichenko, and G. Hong (2015). The cyclicalities of sales, regular, and effective prices: Business cycle and policy implications. *American Economic Review* 105(3), 993–1029.
- Costain, J. and A. Nakov (2015). Precautionary price stickiness. *Journal of Economic Dynamics and Control* 58, 218–234.
- Costain, J. and A. Nakov (2019). Logit price dynamics. *Journal of Money, Credit and Banking* 51(1), 43–78.
- Cover, T. and J. Thomas (2006). *Elements of Information Theory, Second edition*. Wiley Inter-science.
- Csiszár, I. and G. Tusnády (1984). Information geometry and alternating minimization procedures. *Statistics and Decisions* (Supplement 1), 205–237.
- Dixon, H. (2019). Almost-maximization as a behavioral theory of the firm: Static, dynamic, and evolutionary perspectives. *Review of Industrial Organization* Forthcoming.
- Eden, B., M. Eden, and J. Yuen (2019). Temporary sales in response to aggregate shocks. Vanderbilt University Economics WP 19-00003.
- Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011). Reference prices, costs, and nominal rigidities. *American Economic Review* 101(1), 234–62.
- Ellison, S. F., C. Snyder, and H. Zhang (2018). Costs of managerial attention as a source of sticky prices: Structural estimates from an online market. NBER Working Paper 24680.

- Fix, S. L. (1978). Rate distortion functions for squared error distortion measures. *Proceedings of the Sixteenth Annual Allerton Conference on Communication, Control, and Computing*, 704–711.
- Golosov, M. and R. E. Lucas Jr (2007). Menu costs and Phillips curves. *Journal of Political Economy* 115(2), 171–199.
- Gorodnichenko, Y. (2010). Endogenous information, menu costs, and inflation persistence. Manuscript, UC Berkeley.
- Guimaraes, B. and K. D. Sheedy (2011). Sales and monetary policy. *American Economic Review* 101(2), 844–876.
- Ilut, C., R. Valchev, and N. Vincent (2019). Paralyzed by fear: Rigid and discrete pricing under demand uncertainty. Manuscript, Duke Univ.
- Jung, J., J.-H. Kim, F. Matějka, and C. A. Sims (2015). Discrete actions in information-constrained decision problems. Manuscript, CERGE-EI.
- Kehoe, P. and V. Midrigan (2015). Prices are sticky after all. *Journal of Monetary Economics* 75, 35–53.
- Knotek, E. (2011). Convenient prices and price rigidity: cross-sectional evidence. *Review of Economics and Statistics* 93(3), 1076–1086.
- Kryvtsov, O. and N. Vincent (2014). On the importance of sales for aggregate price flexibility. Bank of Canada WP 45/2014.
- Levy, D., S. Dutta, and M. Bergen (2002). Heterogeneity in price rigidity: Evidence from a case study using microlevel data. *Journal of Money, Credit and Banking* 34(1), 197–220.
- Machina, M. (1985). Stochastic choice functions generated from deterministic preferences over lotteries. *The Economic Journal* 95(379), 575–594.
- Mackowiak, B., F. Matějka, and M. Wiederholt (2018). Dynamic rational inattention: Analytical results. *Journal of Economic Theory* 176(C), 650–692.
- Matějka, F. (2016). Rationally inattentive seller: Sales and discrete pricing. *The Review of Economic Studies* 83(3), 1125–1155.
- Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review* 105(1), 272–98.
- Mattsson, L.-G. and J. W. Weibull (2002). Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior* 41(1), 61–78.
- Pastén, E. and R. Schoenle (2016). Rational inattention, multi-product firms, and the neutrality of money. *Journal of Monetary Economics* 80, 1–16.
- Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control* 33(3), 649–665.

- Shannon, C. (1948). A mathematical theory of communication. *Bell Systems Technical Journal* 27, 379–423 and 623–656.
- Shannon, C. (1959). Coding theorems for a discrete source with a fidelity criterion. *Institute of Radio Engineers, National Convention Record* 4, 142–163.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50(3), 665–690.
- Stahl, D. O. (1990). Entropy control costs and entropic equilibria. *International Journal of Game Theory* 19(2), 129–138.
- Steiner, J., C. Stewart, and F. Matejka (2017). Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica* 85(2), 521–553.
- Stevens, L. (2019). Coarse pricing policies. *Review of Economic Studies*, forthcoming.
- Van Damme, E. (2002). *Stability and perfection of Nash equilibria, Second edition*. Springer.
- Varian, H. (1980). A model of sales. *American Economic Review* 70(4), 651–659.
- Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica* 82(6), 2257–2294.
- Zbaracki, M., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004). Managerial and customer costs of price adjustment: direct evidence from industrial markets. *Review of Economics and Statistics* 86, 514–533.

A Proofs of Props. 3 and 4: Equivalence of dynamic RI and CC problems

Showing the equivalence between rational inattention (RI) and control costs (CC) requires more complex notation in a dynamic model than in the static context of [Matějka and McKay \(2015\)](#). Nonetheless, [Steiner et al. \(2017\)](#) showed that the same basic arguments are applicable. We repeat and extend their main result here (Prop. 3), stating it in somewhat less generality for the sake of transparency.³¹ But we also extend their result by considering limits on memory (Prop. 4). Concretely, we assume that the DM costlessly recalls the last $\tau \geq 0$ signals received (i.e. the last τ actions), $\mathcal{B}_\tau^{t-1} \equiv (a_{t-1}, a_{t-2}, \dots, a_{t-\tau})$. The limit $\tau \rightarrow \infty$ corresponds to the standard rational inattention problem in which the history of all past signals, $a^{t-1} \equiv \mathcal{B}_\infty^{t-1}$, is costlessly remembered; this is the case studied by [Steiner et al. \(2017\)](#).

The value of a dynamic RI problem is a function of the DM's prior about the state of the world θ^t , conditional on signals \mathcal{B}_τ^{t-1} recalled at time $t - 1$. So we consider RI problems of the following form:³²

$$U(\mathcal{B}_\tau^0) = \max_{\pi \in \Delta(\mathcal{A})} E \left\{ \sum_{t=1}^{\infty} \delta^t [u(a_t, \theta_t) - \kappa \mathcal{I}(a_t; \theta^t | \mathcal{B}_\tau^{t-1})] \middle| \mathcal{B}_\tau^0 \right\}. \quad (71)$$

$$= \max_{\pi \in \Delta(\mathcal{A})} \sum_{\theta^1} \pi_p(\theta^1 | \mathcal{B}_\tau^0) E \left\{ \sum_{t=1}^{\infty} \delta^t [u(a_t, \theta_t) - \kappa \mathcal{I}(a_t; \theta^t | \mathcal{B}_\tau^{t-1})] \middle| \theta^1, \mathcal{B}_\tau^0 \right\}. \quad (72)$$

The second line partially expands the expectation in terms of the prior $\pi_p(\theta^1 | \mathcal{B}_\tau^0)$ over possible histories θ^1 consistent with the initial information set \mathcal{B}_τ^0 . The choice variable here is the system of distributions of actions $\pi \equiv \{\pi_t(a_t | \theta^t, \mathcal{B}_\tau^{t-1})\}_{t=1}^{\infty}$, which must be chosen from $\Delta(\mathcal{A})$ for each t , each θ^t , and each \mathcal{B}_τ^{t-1} . The expectation is defined by the dynamics of the underlying state, $\pi(\theta_t | \theta^{t-1})$, and by the chosen distribution of actions $\pi_t(a_t | \theta^t, \mathcal{B}_\tau^{t-1})$.³³ Thus, the chosen distribution allows the action a_t to be correlated with the underlying state θ^t , subject to information costs. Since the signals \mathcal{B}_τ^{t-1} are costlessly remembered, the chosen distribution may condition on these signals too, without additional cost. We assume that current profits depend on a_t and θ_t only. More generally, the following arguments will still go through as long as the utility function can be written in the form $u_t(\mathcal{B}_\tau^t, \theta^t)$, which allows for some cases with intertemporal dependence.

31. As a first step, Lemma 1 of [Steiner et al. \(2017\)](#) shows that a problem on an arbitrary signal space can be reduced to a problem in which each signal indicates precisely one action. We take as given that the simplified signal space applies. Also, their model allows for two classes of signals: costly and costless. For brevity, we restrict ourselves to the case where there is no costless signal.

32. The value could be written as a function of the prior $\pi(\theta^t | \mathcal{B}_\tau^{t-1})$, but we can also write it as a function of the information on which the prior depends, namely the history of observed signals a^{t-1} .

33. We economize on notation by using the letter π to represent all the probabilities related to the stochastic processes θ_t , a_t , and their histories. We drop the time subscript on π when it is clear from context.

The information costs at time t refer to the mutual information between a_t and θ^t (or equivalently, between \mathcal{B}_τ^t and θ^t), conditional on time t memories \mathcal{B}_τ^{t-1} :

$$\mathcal{I}(a_t; \theta^t | \mathcal{B}_\tau^{t-1}) \equiv \sum_{\theta^t} \sum_{a_t} \pi(a_t, \theta^t | \mathcal{B}_\tau^{t-1}) \ln \left(\frac{\pi(a_t, \theta^t | \mathcal{B}_\tau^{t-1})}{\pi(a_t | \mathcal{B}_\tau^{t-1}) \pi(\theta^t | \mathcal{B}_\tau^{t-1})} \right) \quad (73)$$

$$= \sum_{\theta^t} \pi(\theta^t | \mathcal{B}_\tau^{t-1}) \sum_{a_t} \pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1}) \ln \left(\frac{\pi(a_t, | \theta^t, \mathcal{B}_\tau^{t-1})}{\pi(a_t | \mathcal{B}_\tau^{t-1})} \right) \quad (74)$$

$$= \min_{q(a_t | \mathcal{B}_\tau^{t-1}) \in \Delta(\mathcal{A})} \sum_{\theta^t} \pi(\theta^t | \mathcal{B}_\tau^{t-1}) \sum_{a_t} \pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1}) \ln \left(\frac{\pi(a_t, | \theta^t, \mathcal{B}_\tau^{t-1})}{q(a_t | \mathcal{B}_\tau^{t-1})} \right). \quad (75)$$

The second expression applies the relation between joint and conditional probabilities. The third line recalls the fact that mutual information represents a minimum of Kullback-Leibler divergence across all distributions $q(a)$ over \mathcal{A} (Cover and Thomas 2006, Lemma 10.8.1).³⁴

Plugging in (75), and expanding the expectations, problem (71) becomes:

$$U(\mathcal{B}_\tau^0) = \max_{\pi, q \in \Delta(\mathcal{A})} \sum_{\theta^1} \pi_p(\theta^1 | \mathcal{B}_\tau^0) E \left\{ \sum_{t=1}^{\infty} \delta^t \left[u(a_t, \theta_t) - \kappa \ln \left(\frac{\pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1})}{q(a_t | \mathcal{B}_\tau^{t-1})} \right) \right] \middle| \theta^1, \mathcal{B}_\tau^0 \right\} \quad (76)$$

$$= \max_{\pi, q \in \Delta(\mathcal{A})} \sum_{\theta^1} \pi_p(\theta^1 | \mathcal{B}_\tau^0) \sum_{t=1}^{\infty} \delta^t \sum_{\theta^t} \sum_{\mathcal{B}_\tau^{t-1}} \pi(\theta^t, \mathcal{B}_\tau^{t-1} | \theta^1, \mathcal{B}_\tau^0) \sum_{a_t} \pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1}) \left[u(a_t, \theta_t) - \kappa \ln \left(\frac{\pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1})}{q(a_t | \mathcal{B}_\tau^{t-1})} \right) \right]. \quad (77)$$

This problem selects a benchmark distribution $q \equiv \{q_t(a_t | \mathcal{B}_\tau^{t-1})\}_{t=1}^{\infty}$ for each time t , and for each information set \mathcal{B}_τ^{t-1} ; the distribution $\pi \equiv \{\pi_t(a_t | \theta^t, \mathcal{B}_\tau^{t-1})\}_{t=1}^{\infty}$ conditions on the underlying state θ^t too. This proves part (i) of Prop. 3, which is analogous to Lemma 2 of Steiner et al. (2017).

Next, we can separate out the first period:

$$U(\mathcal{B}_\tau^0) = \max_{\pi, q \in \Delta(\mathcal{A})} \sum_{\theta^1} \pi_p(\theta^1 | \mathcal{B}_\tau^0) \left\{ \delta \sum_{a_1} \pi(a_1 | \theta^1, \mathcal{B}_\tau^0) \left[u(a_1, \theta_1) - \kappa \ln \left(\frac{\pi(a_1 | \theta^1, \mathcal{B}_\tau^0)}{q(a_1 | \mathcal{B}_\tau^0)} \right) \right] \dots \right. \\ \left. + \sum_{t=2}^{\infty} \delta^t \sum_{\theta^t | \theta^1} \sum_{\mathcal{B}_\tau^{t-1} | \mathcal{B}_\tau^0} \pi(\theta^t, \mathcal{B}_\tau^{t-1} | \theta^1, \mathcal{B}_\tau^0) \left[\sum_{a_t} \pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1}) \left[u(a_t, \theta_t) - \kappa \ln \left(\frac{\pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1})}{q(a_t | \mathcal{B}_\tau^{t-1})} \right) \right] \right] \right\}. \quad (78)$$

Now note that the possible information sets \mathcal{B}_τ^1 at $t = 1$ are given by $\mathcal{B}_\tau^1 \equiv (a_1, \mathcal{B}_{\tau-1}^0)$ for each

34. The second summation in (75) is taken over all histories a^t which represent continuations of the history a^{t-1} . Note that conditional probabilities of actions and continuation histories are interchangeable, that is, $\pi(a_t | a^{t-1}) = \pi(a^t | a^{t-1})$ and $\pi(a_t | \theta^t, a^{t-1}) = \pi(a^t | \theta^t, a^{t-1})$ for all histories a^t that are continuations of a^{t-1} . We will use whichever expression is clearest in context.

possible $a_1 \in \mathcal{A}$. Therefore we can factor transition probabilities as follows:

$$\pi(\theta^2, \mathcal{B}_\tau^1 | \theta^1, \mathcal{B}_\tau^0) = \pi(\theta^2 | \theta^1) \pi(a_1 | \theta^1, \mathcal{B}_\tau^0)$$

for $\mathcal{B}_\tau^1 \equiv (a_1, \mathcal{B}_{\tau-1}^0)$, for each $a_1 \in \mathcal{A}$. This allows us to pull the common factor $\pi(a_1 | \theta^1, \mathcal{B}_\tau^0)$ out of the first and second lines of (78), to obtain:

$$\begin{aligned} U(\mathcal{B}_\tau^0) &= \delta \max_{\pi, q \in \Delta(\mathcal{A})} \sum_{\theta^1} \pi_p(\theta^1 | \mathcal{B}_\tau^0) \sum_{a_1} \pi(a_1 | \theta^1, \mathcal{B}_\tau^0) \left\{ \left[u(a_1, \theta_1) - \kappa \ln \left(\frac{\pi(a_1 | \theta^1, \mathcal{B}_\tau^0)}{q(a_1 | \mathcal{B}_\tau^0)} \right) \right] \dots \right. \\ &\quad \left. + \sum_{\theta^2 | \theta^1} \pi(\theta^2 | \theta^1) \sum_{t=2}^{\infty} \delta^{t-1} \sum_{\theta^t | \theta^2} \sum_{\mathcal{B}_\tau^{t-1} | \mathcal{B}_\tau^1} \pi(\theta^t, \mathcal{B}_\tau^{t-1} | \theta^2, \mathcal{B}_\tau^1) \left[\sum_{a_t} \pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1}) \left[u(a_t, \theta_t) - \kappa \ln \left(\frac{\pi(a_t | \theta^t, \mathcal{B}_\tau^{t-1})}{q(a_t | \mathcal{B}_\tau^{t-1})} \right) \right] \right] \right\}. \end{aligned} \quad (79)$$

Here we have effectively broken the dynamic RI value function $U(\mathcal{B}_\tau^0)$ into pieces $V(\theta^1, \mathcal{B}_\tau^0; q)$ which condition on *full information*. Formally, we can write:³⁵

$$U(\mathcal{B}_\tau^0) = \delta \max_{q_t(a_t | \mathcal{B}_\tau^{t-1}) \in \Delta(\mathcal{A})} \sum_{\theta^1} \pi_p(\theta^1 | \mathcal{B}_\tau^0) V(\theta^1, \mathcal{B}_\tau^0; q). \quad (80)$$

where $V(\theta^1, \mathcal{B}_\tau^0; q)$ is the value of the following full-information recursive CC problem:

$$V(\theta^1, \mathcal{B}_\tau^0; q) = \max_{\pi(a_1 | \theta^1, \mathcal{B}_\tau^0) \in \Delta(\mathcal{A})} \sum_{a_1} \pi(a_1 | \theta^1, \mathcal{B}_\tau^0) \left[u(a_1, \theta_1) - \kappa \ln \left(\frac{\pi(a_1 | \theta^1, \mathcal{B}_\tau^0)}{q(a_1 | \mathcal{B}_\tau^0)} \right) + \delta \sum_{\theta^2 | \theta^1} \pi(\theta^2 | \theta^1) V(\theta^2, \mathcal{B}_\tau^1; q) \right]. \quad (81)$$

We have now proved parts (ii)-(iii) of Prop. 3, which are equivalent to Theorem 1 and Prop. 3 of Steiner et al. (2017). Notice that maximization in this problem treats θ^1 as if it were known, and the transition probabilities governing θ^2 are the true, objective probabilities. Therefore, at each time step, (81) is a static control cost problem formally identical to (15); hence it is solved by a multinomial logit, proving part (iv). Part (v.) of Prop. 3 follows from the first-order condition for q in (75) or (77).

Since we have proved the results for the general case of finite or infinite memory length τ , we have already proved Prop. 4, part (a).

Now note that for any τ , the value functions and probability functions have the same functional forms. Comparing the information sets at any time t in the case of some finite

35. Steiner et al. (2017) write the CC value function V from the point of view of the beginning of period t , but they define the RI problem from the point of view of the end of $t - 1$. We follow the same convention here, to keep notation similar.

memory length $\tau = T$ and in the infinite-memory case $\tau = \infty$, the difference between the two is just the tail of distant memories $\mathcal{B}_\infty^{t-T+1}$. As $T \rightarrow \infty$ the value of this tail for predicting the current state θ_t goes to zero. Therefore, the difference in the decision values achievable by knowing \mathcal{B}_∞^{t-1} rather than \mathcal{B}_T^{t-1} goes to zero as $T \rightarrow \infty$. Therefore the finite-memory value functions $V(\theta^t, \mathcal{B}_T^{t-1}, q)$ converge to the infinite memory values $V(\theta^t, \mathcal{B}_\infty^{t-1}, q)$ as $T \rightarrow \infty$, and hence the probabilities (27)-(28) of the finite-memory problem converge to those of the infinite-memory problem, (23)-(25). This proves Prop. 4, part (b).

B Computation

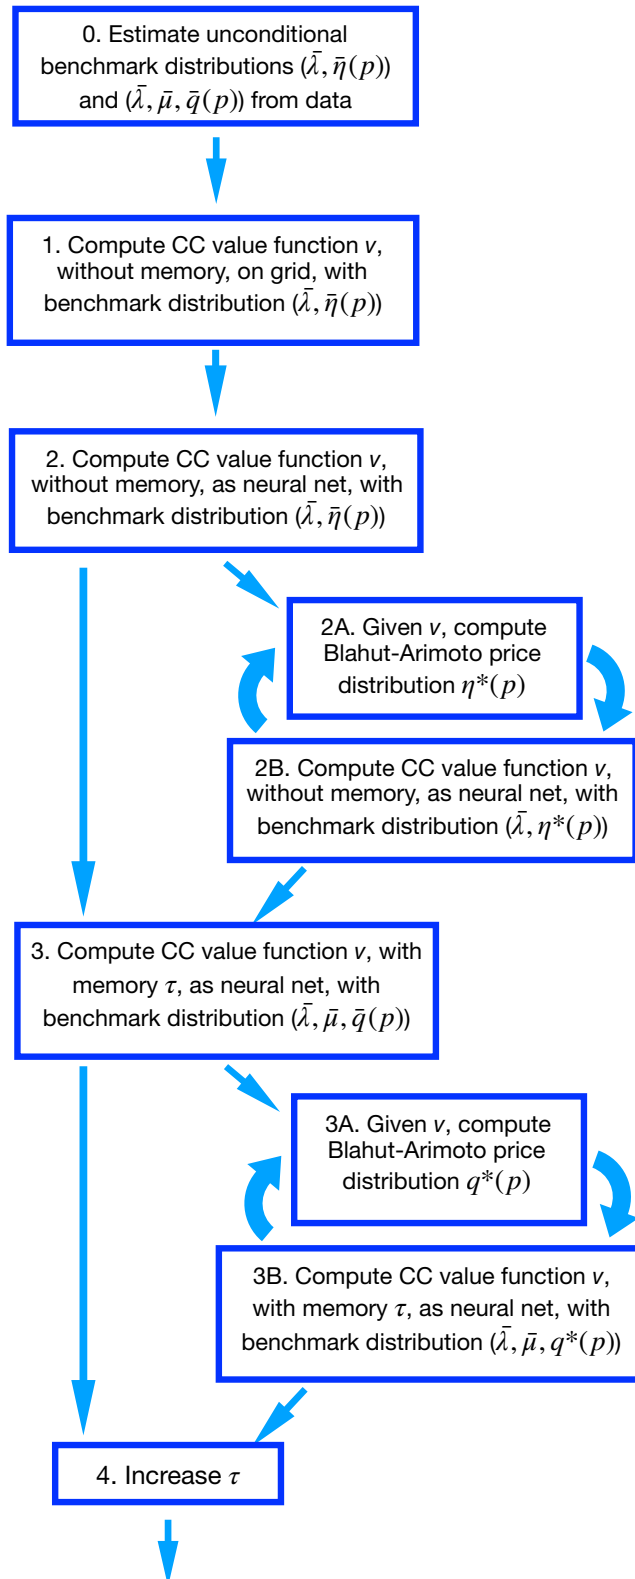
Figure 11 provides an outline of our algorithm, which calculates the steady-state partial equilibrium of a population of retail firms that behave according to the model in Sec. 4 or Sec. 5. The diagram indicates more than one possible direction of flow, representing possible alternative sequences; steps may be skipped or added in order to speed up the calculation or ensure convergence. In principle, following the algorithm for ever larger memory lengths τ , the solution converges to that of the RI model with infinite memory. In practice, one may stop earlier, either because larger τ becomes computationally feasible, or because finite τ provides an adequate model that fits the data well.

As an initial step, we use appropriate retail microdata to estimate the unconditional price adjustment hazard $\bar{\lambda}$, and the unconditional distribution $\bar{\eta}(p)$ of prices that are set if a different price is chosen. In the RI model with limitless memory, these unconditional probabilities are the average benchmark probabilities (taking an expectation over all possible information sets). In the zero-memory RI model, these unconditional probabilities are the benchmark probabilities that are used in all states of the world.

Hence, in step 1, we compute the zero-memory STMRI model, taking the unconditional probabilities $(\bar{\lambda}, \bar{\eta}(p))$ as the benchmark probabilities, using the grid-based method of Reiter (2009). Concretely, we compute the value function $v(p, z)$ over a grid of possible values of prices and productivity, by backwards induction on the Bellman equation. Details of this calculation are described in the computational appendix of Costain and Nakov (2019).

In step 2, we solve the same model, but represent the value function $v(p, z)$ by a neural network $N(p, z, \beta)$, where β is a vector of parameters. To do so, we write a function which calculates the difference between the left and right sides of the Bellman equation, for any state (p, z) . The left side is simply the value function, $N(p, z, \beta)$. The right side contains an expectation over the next period's value function, $N(p', z', \beta)$, which we compute by quadrature conditional on given values of (p, z, β) . Integrating over z' is straightforward, since this is an exogenous shock, so its distribution is given. Crucially, the distribution over p' , conditional on z' , is also known: we can calculate it analytically as a logit, given the value function

Figure 11: Algorithm: schematic outline.



$N(p', z', \beta)$. The fact that our control cost model delivers a known functional form for decision probabilities is key for the tractability of this calculation.

Hence, to find the neural network that solves the Bellman equation for the model in Sec. 4, we simulate a population of products, with current states (p, z) , seeking the parameters β that minimize the difference between the left-hand and right-hand sides of the Bellman equation, while simulating forward given the implied distribution of prices in order to find the steady-state cross-sectional distribution of states (p, z) . To do this, we call a MATLAB minimization routine to find the minimum of the function described in the previous paragraph.

In principle, the previous steps will find the solution of the zero-memory STMRI model. In practice, the unconditional probabilities $(\bar{\lambda}, \bar{\eta}(p))$ extracted from microdata are an average price distribution over many heterogeneous products, not the benchmark probabilities governing the price choice for a single product. To find the benchmark probabilities for a given product, we can next follow the Blahut-Arimoto algorithm. Given the neural network representation of the value function $v(p, z) = N(p, z, \beta)$, we can iterate on (27)-(28) to calculate the zero-memory benchmark distribution $\eta^*(p)$. This is step 2A, which reduces the continuous distribution extracted from the data (see Fig.) to a discrete distribution (see Fig.). Then, in step 2B, we can again solve the Bellman equation by minimizing the difference between the two sides of the Bellman equation with respect to the neural network parameters β . Iterating on steps 2A and 2B, we converge to a solution of the zero-memory STMRI model.

Next, we use the zero-memory solution as a starting point for a calculation with finite memory $\tau > 0$. At this point we return to the data to extract the unconditional probabilities $(\bar{\lambda}, \bar{\mu}, \bar{q}(p))$, where $\bar{\lambda}$ is the probability of adjusting the price, $\bar{\mu}$ is the probability of choosing a price *not* observed in the backwards window \mathbf{b}_τ , and $\bar{q}(p)$ is the distribution over prices chosen when these are *not* among the prices in \mathbf{b}_τ . In step 3, we assume these empirical probabilities are the benchmark probabilities, and we seek neural network parameters β that minimize the Bellman equation for the model described in Sec. 5. The value function is assumed to vary with the memories of the firm, so formally the neural network takes the form $N(p, z, \lfloor_\tau, \beta)$, conditioning on the prices in memory as well as the current p and z . Subsequently, iterating on steps 3A-3B, we can seek a fixed point $\bar{q}^*(p)$ for the benchmark probabilities.

The reader will notice that these steps, as described, do not suffice to solve the STMRI model with memory τ . As described, we have found a single set of benchmark probabilities $\bar{q}^*(p)$; but under rational inattention, the benchmark probabilities should instead vary with the information set \mathbf{b}_τ , taking the form $\bar{q}^*(p|\mathbf{b}_\tau)$. In principle, we could solve the STMRI model by seeking a fixed point of steps 3A-3B separately for each possible information set \mathbf{b}_τ (assuming that prices are drawn from a finite grid, the set of possible information sets is finite). In special cases where the solution can be shown to have a simple Markovian structure, making many possible information sets equivalent, this calculation may be feasible (see the

examples in [Steiner et al. \(2017\)](#)). But this is obviously not feasible in general, as it is subject to a fierce curse of dimensionality.

Therefore, the results reported in the current version of the paper (Jan. 2020) are based on a single set of benchmark probabilities $\bar{q}^*(p)$; the value function varies with memory, but the benchmark probabilities do not. To get closer to the true STMRI solution, there are several alternatives. On one hand, we could seek the fixed point of 3A-3B conditional on a subset of memories; for example we could solve 3A-3B separately conditional on each possible current price p , instead of solving separately for each information set \mathbf{b}_τ . This would take us part of the way to the STMRI solution. Alternatively, we could use a neural network to represent the smooth function C defined in Prop. 5, as a function of p , and memories \mathbf{b}_τ . The points where this function is (approximately) equal to one are the support of q , given the memories. Using the function C to infer q at each set of memories that occurs in simulated data would require computation time proportional to the amount of simulated data, rather than being proportional to the number of possible memory states, so the curse of dimensionality would be less severe with this approach. We hope to explore these extensions in future versions of this paper.