Online Appendix (not for publication)

The Welfare Effects of Supply and Demand Frictions in a Dynamic Pricing Game by Myśliwski, Sanches, Silva Jr., and Srisuma

June 30, 2023

S.1 Unobserved demand heterogeneity

This appendix discusses the effects of time-varying unobserved heterogeneity on our demand estimates. Our analysis is based on three different exercises. First, using aggregate data at brand level we estimated the following regression:

$$s_{jmt} = \alpha + \sum_{j=1}^{J} \beta_j p_{jmt} + \sum_{j=1}^{J} \gamma_j s_{jmt-1} + \delta_j + \delta_t + \varepsilon_{jmt}, \tag{S.1.1}$$

where, s_{jmt} and p_{jmt} are the share and the price of brand j at supermarket m and week t, δ_j is a brand dummy, δ_t is a weekly dummy and ε_{jmt} is an idiosyncratic error term. We see this equation as a reduced form representation of the process governing the evolution of the aggregated market shares as shown by equation (9) in the paper.

We estimated one equation for each supermarket – Morrisons and Tesco – and collected the estimates of the weekly dummies, δ_t. These dummies are interpreted as unobserved demand shocks affecting all butter and margarine brands in each supermarket. Figure S1.1 illustrates the dummies for Morrisons (upper panel) and Tesco (lower panel). The figure on the left hand side shows point estimates of each dummies and 95% confidence intervals. For Morrisons and Tesco virtually all dummy estimates are not significant at 10%. The figures on the right hand side show the same point estimates without the confidence intervals and allows us to better visualise the behaviour of these demand shocks over time. In both figures we do not see any clear pattern (cycle, trend, seasonality) of these demand shocks over time.

As mentioned, the coefficients in Figure S1.1 capture unobserved demand shocks affecting all brands in each supermarket. To see how these shocks differ across brands we estimate a separate regression for each brand separately pooling both supermarkets (and including supermarket dummies). The time dummies in this regression are now interpreted as brand specific unobserved demand shocks. Figure S1.2 shows these time dummies. For each brand, the figure shows two graphs. The first has point estimates with 95% confidence intervals and the second has the point estimates without confidence intervals. A quick inspection of the figure reveals that all the point estimates are not significant at 10%. They also appear to do not show any

Figure S1.1: Weekly Dummies (Demand Shock) for each Supermarket

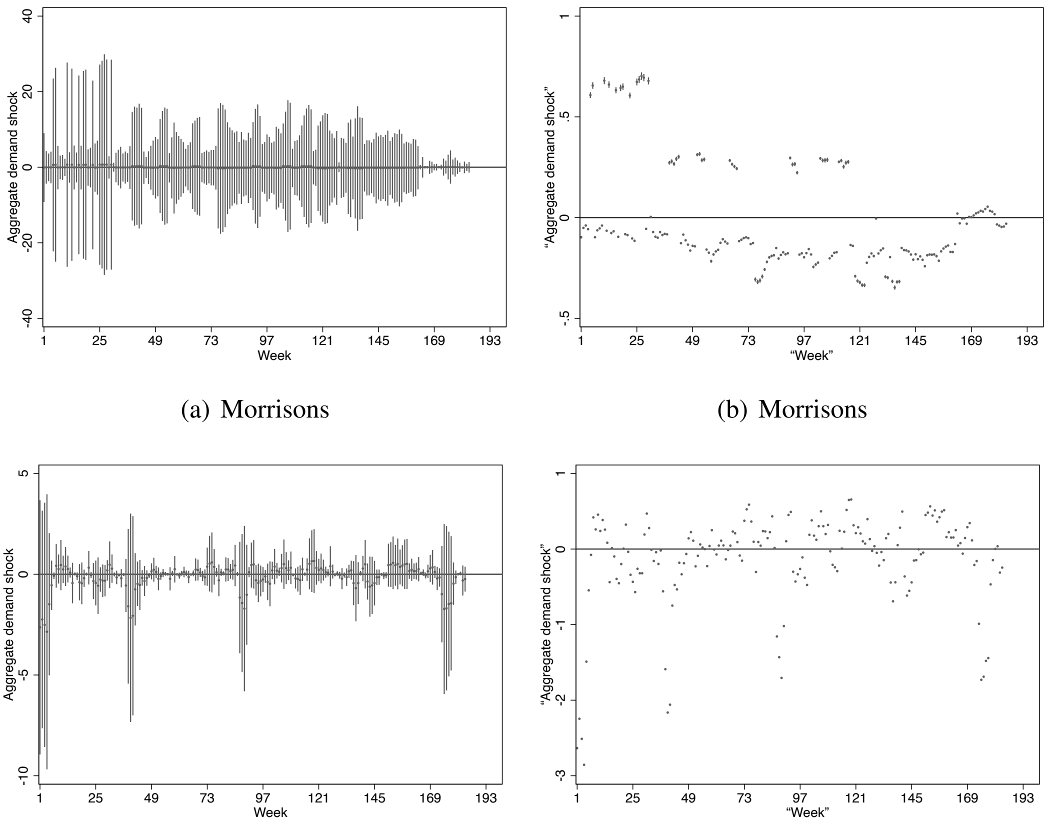

(a) Morrisons (b) Morrisons

(c) Tesco (d) Tesco

clear pattern over time.

For our third exercise, we use the same scanner data that we use to estimate the demand models in our paper. Using a similar idea above, we re-estimate our demand model that now includes weekly dummies in our models. Ideally, we would like to include a time varying brand dummy to mimic BLP's ξ_{jt} – instead of a brand fixed effect as in the previous version of the paper – but estimating a non linear model with more than 1400 dummies showed infeasible. Noting here that we have to perform maximum likelihood estimation to capture the unobserved time effects because we cannot use Berry's inversion. As an alternative, we included 200 weekly dummies shifting the indirect utility from choosing the outside option, so that:

$$u_{0t}^{h} = \delta_{0t} + \xi_{0t}^{h} \tag{S.1.2}$$

The estimated $\hat{\delta}_{0t}$ are interpreted as unobserved demand shocks affecting all brands at the same supermarket and are plotted in Figure S1.3. Similarly to the aggregate-level exercise, we find no clear patterns or cycles. For Morrisons, the majority of shocks are not significantly different from 0. For Tesco, the confidence intervals are narrower, but the magnitudes are

Figure S1.2: Weekly Dummies (Demand Shock) for each Brand

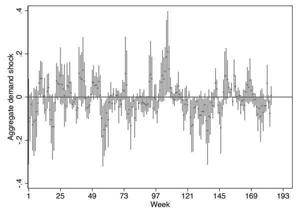

(a) Anchor

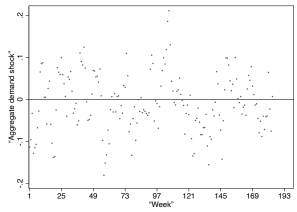

(b) Anchor

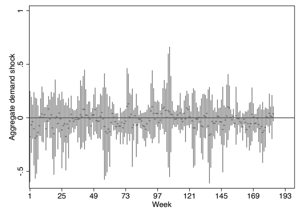

(c) Lurpak

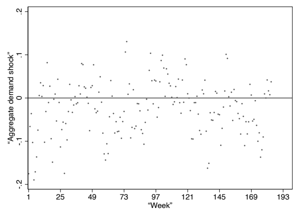

(d) Lurpak

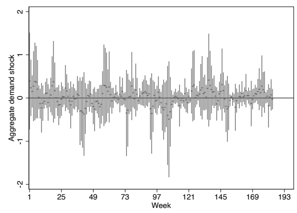

(e) Clover

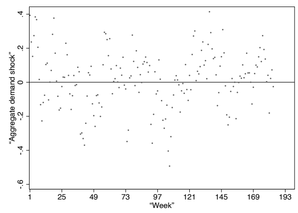

(f) Clover

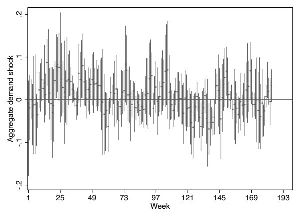

(g) Country Life

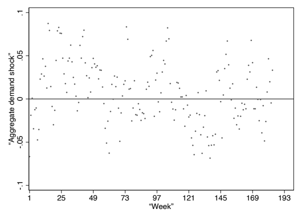

(h) Country Life

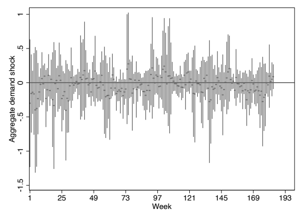

(i) Flora

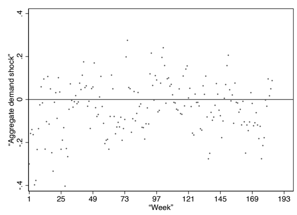

(j) Flora

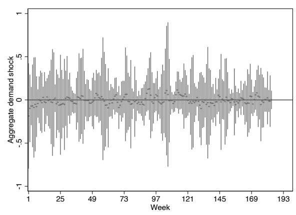

(k) ICBINB

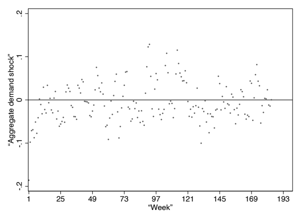

(l) ICBINB

not economically significant either when compared to the estimated brand constants δ_j. The normalisation in the baseline model implies that all those coefficients are 0, whilst here most of the estimates are between 0 and -0.1. We therefore believe that the bias in the baseline model is negligible. Overall, this set of results seems to suggest that time varying unobserved heterogeneity affecting individual demand is not an important issue in this market. A potential explanation for these patterns is that all these brands are well established in the market and the relevant space of characteristics of these products is pretty stable over this period. The same type of arguments were used in Griffith et al. (2017) to model the demand for butter and margarine using the same dataset used in this paper.

Figure S1.3: Weekly Dummies (Demand Shock) for each Supermarket – Scanner Data

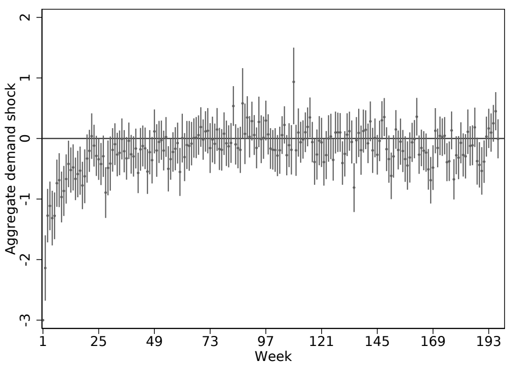

(a) Morrisons

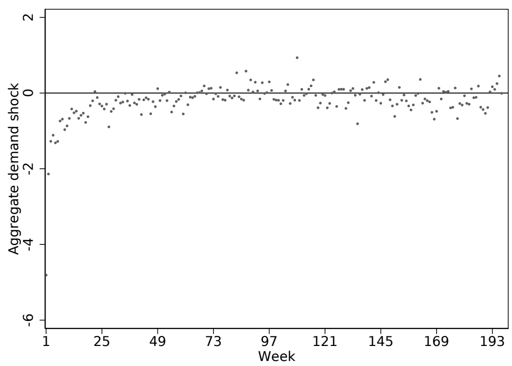

(b) Morrisons

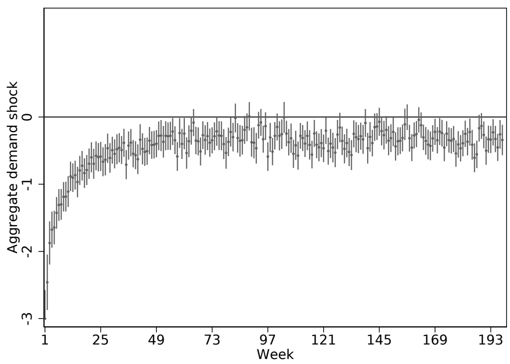

(c) Tesco

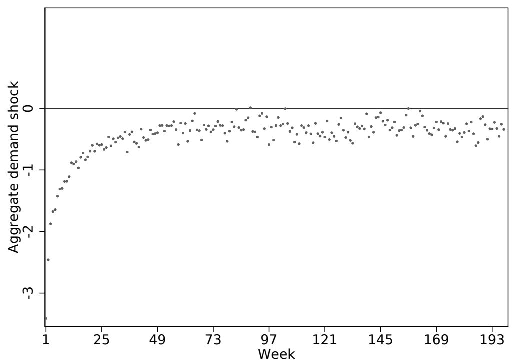

(d) Tesco

S.2 Identification, estimation and model solution

This appendix summarises some of the technical details behind identification, estimation and solution of the model. Appendix S.2.1 shows how relative price adjustment costs can be identified independently of payoff parameters using the arguments in Komarova et al. (2018) and argues how additional, economically meaningful restrictions can be imposed to interpret the parameters as product-specific price adjustment costs (instead of relative differences). Appendix S.2.2 lays out the procedures used to estimate the discount factor, β. Finally, Appendix S.2.3 contains details on the algorithms used to solve for Markov Perfect Equilibrium in the counterfactual scenarios.

S.2.1 Closed-form identification of price adjustment costs

In this appendix we lay out the identification result for the vector of price adjustment costs in our model. To make it self-contained, we will repeat some of the notational assumptions we have been making throughout the main body of the paper. Also, to make the exposition clearer and give the reader an idea of the dimension of the problem, we will be referring to a specific number of firms, actions and cardinality of the set of possible market shares which will be the same as in our empirical application.

Notation recap

There are three firms, producing two products each (four actions per firm). There is also a generic good that can be chosen by consumers, but its price is exogenously given (hence there are 7 lagged market shares to keep track of). The vector of publicly observed state variables is $\mathbf{z}_t = (\mathbf{s}_{t-1}, \mathbf{a}_{t-1})$. We discretise last period's market shares into 3 bins, therefore the dimension of the state space $\boldsymbol{\mathcal{Z}}$ is: $|\boldsymbol{\mathcal{Z}}| = 4^3 \cdot 3^7 = 64 \cdot 2187 = 139,968$. For simplicity we will refer to the action $(p^H_{i_1}, p^H_{i_2})$ as HH. The payoff function of firm i is:

$$\Pi_i(\mathbf{a}_t, \mathbf{z}_t, \boldsymbol{\varepsilon}_{it}) = \pi_i(a_{it}, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) + \sum_{\ell \in \mathcal{A}_i} \zeta \cdot \varepsilon_{it}(\ell) \cdot \mathbf{1}(a_{it} = \ell) - \sum_{\ell \in \mathcal{A}_i} \sum_{\ell' \neq \ell} AC_i^{\ell' \to \ell} \cdot \mathbf{1}(a_{it} = \ell, a_{i,t-1} = \ell') - \sum_{\ell \in \mathcal{A}_i} FC_i^{\ell} \cdot \mathbf{1}(a_{it} = \ell). \tag{S.2.1}$$

To simplify the notation in the derivations that follow, we will assume that $\zeta = 1$, so that $AC = AC'$, as defined in section 4.1. Otherwise one should divide both sides of (S.2.1) by ζ and use the prime notation to denote rescaled primitives.

Without loss of generality, we also slightly abuse the notation and let $\pi_i(\cdot)$ absorb FC_i. This can be done because the component $\sum_{\ell \in \mathcal{A}_i} FC_i^\ell \cdot \mathbf{1}(a_{it} = \ell)$ does not depend on past actions and will be integrated out in the derivation together with the remainder of the deterministic, static payoff.

Derivation of ΔAC

The non-stochastic dynamic payoff from choosing $a_{it} = \ell$ is:

$$\bar{v}_i(\ell, \mathbf{z}_t) = \sum_{\mathbf{a}_{-it} \in \times_{j \neq i} \mathcal{A}_j} \sigma_i(\mathbf{a}_{-it}|\mathbf{z}_t) \Big[\pi_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) + \beta \sum_{\mathbf{z}_{t+1}} G(\mathbf{z}_{t+1}|\mathbf{s}_{t-1}, \ell, \mathbf{a}_{-it}) \tag{S.2.2}$$
$$\underbrace{\cdot \int V_i(\mathbf{z}_{t+1}, \varepsilon_{t+1}) dQ(\varepsilon_{i,t+1})}_{\tilde{V}(\mathbf{z}_{t+1})} \Big] - \sum_{\ell' \neq \ell} AC_i^{\ell' \to \ell} \cdot \mathbf{1}(a_{i,t-1} = \ell')$$

Defining the differences with respect to the reference action HH we have:

$$\begin{aligned}
\Delta \bar{v}_i(\ell, \mathbf{z}_t) &= \bar{v}_i(\ell, \mathbf{z}_t) - \bar{v}_i(HH, \mathbf{z}_t) \\
&= \sum_{\mathbf{a}_{-it} \in \times_{j \neq i} \mathcal{A}_j} \sigma_i(\mathbf{a}_{-it}|\mathbf{z}_t) \Big\{ \underbrace{\pi_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) - \pi_i(HH, \mathbf{a}_{-it}, \mathbf{s}_{t-1})}_{\Delta \pi_i^\ell(\mathbf{a}_{-it}, \mathbf{s}_{t-1})} \Big\} \\
&+ \sum_{\mathbf{a}_{-it} \in \times_{j \neq i} \mathcal{A}_j} \sigma_i(\mathbf{a}_{-it}|\mathbf{z}_t) \Big\{ \beta \sum_{\mathbf{z}_{t+1}} [\underbrace{G(\mathbf{z}_{t+1}|\mathbf{s}_{t-1}, \ell, \mathbf{a}_{-it}) - G(\mathbf{z}_{t+1}|\mathbf{s}_{t-1}, HH, \mathbf{a}_{-it})}_{\Delta G^\ell(\mathbf{z}_{t+1}|\mathbf{a}_{-it}, \mathbf{s}_{t-1})}] \tilde{V}(\mathbf{z}_{t+1}) \Big\} \\
&- \underbrace{\sum_{\ell' \neq \ell} [AC_i^{\ell' \to \ell} \cdot \mathbf{1}(a_{i,t-1} = \ell') - AC_i^{\ell' \to HH} \cdot \mathbf{1}(a_{i,t-1} = \ell')]}_{\Delta AC_i^\ell(a_{i,t-1})}
\end{aligned}$$

Using the newly introduced notation, we have:

$$\Delta \bar{v}_i(\ell, \mathbf{z}_t) = \sum_{\mathbf{a}_{-it} \in \times_{j \neq i} \mathcal{A}_j} \sigma_i(\mathbf{a}_{-it}|\mathbf{z}_t) \underbrace{\Big\{ \Delta \pi_i^\ell(\mathbf{a}_{-it}, \mathbf{s}_{t-1}) + \beta \sum_{\mathbf{z}_{t+1}} \Delta G^\ell(\mathbf{z}_{t+1}|\mathbf{a}_{-it}, \mathbf{s}_{t-1}) \tilde{V}(\mathbf{z}_{t+1}) \Big\}}_{\lambda_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1})}$$
$$- \Delta AC_i^\ell(a_{i,t-1}) \tag{S.2.3}$$

Thinking back about the dimension of the problem, for each of the three remaining (that is, excluding HH) actions of firm i, there are $4^2 \cdot 3^7 = 16 \cdot 2187 = 34992$ $\lambda_i(\ell, *)$ terms. Rewriting (S.2.3) in vector form:

$$\Delta \bar{v}_i(\ell, \mathbf{z}_t) = \boldsymbol{\sigma}_i(\mathbf{z}_t)' \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}) - \Delta AC_i^\ell(a_{i,t-1}), \tag{S.2.4}$$

where $\boldsymbol{\sigma}_i(\mathbf{z}_t) = [\sigma_i(\mathbf{a}_{-it}|\mathbf{z}_t)]_{\mathbf{a}_{-it}}$ and $\boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}) = [\lambda_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1})]_{\mathbf{a}_{-it}}$ are 16×1 column vectors. (S.2.4) holds for all of the 139,968 points in the state space. To make things more explicit, use the fact that $\mathbf{z}_t$ can be partitioned into $(\mathbf{a}_{t-1}, \mathbf{s}_{t-1})$. Furthermore:

$$\begin{aligned}
\mathbf{a}_{t-1} &= \{\mathbf{a}_{t-1}^1, \mathbf{a}_{t-1}^2, \ldots, \mathbf{a}_{t-1}^{64}\} \\
\mathbf{s}_{t-1} &= \{\mathbf{s}_{t-1}^1, \mathbf{s}_{t-1}^2, \ldots, \mathbf{s}_{t-1}^{2187}\}
\end{aligned}$$

For $\mathbf{s}^1_{t-1}$ the system can be written as:

$$\begin{cases} \Delta\bar{v}_i(\ell, \mathbf{a}^1_{t-1}, \mathbf{s}^1_{t-1}) = \boldsymbol{\sigma}_i(\mathbf{a}^1_{t-1}, \mathbf{s}^1_{t-1})'\boldsymbol{\lambda}_i(\ell, \mathbf{s}^1_{t-1}) - \Delta AC^\ell_i(\mathbf{a}^1_{t-1}) \\ \vdots \\ \Delta\bar{v}_i(\ell, \mathbf{a}^{64}_{t-1}, \mathbf{s}^1_{t-1}) = \boldsymbol{\sigma}_i(\mathbf{a}^{64}_{t-1}, \mathbf{s}^1_{t-1})'\boldsymbol{\lambda}_i(\ell, \mathbf{s}^1_{t-1}) - \Delta AC^\ell_i(\mathbf{a}^{64}_{t-1}) \end{cases}$$

Vectorising again:

$$\Delta\bar{\mathbf{v}}_i(\ell, \mathbf{s}^1_{t-1}) = \boldsymbol{\sigma}_i(\mathbf{s}^1_{t-1})\boldsymbol{\lambda}_i(\ell, \mathbf{s}^1_{t-1}) - \Delta\mathbf{AC}^\ell_i, \tag{S.2.5}$$

where $\bar{\mathbf{v}}_i(\ell, \mathbf{s}^1_{t-1}) = [\Delta\bar{v}_i(\ell, \mathbf{a}_{t-1}, \mathbf{s}^1_{t-1})]_{\mathbf{a}_{t-1}}$ is a 64×1 vector, $\boldsymbol{\sigma}_i(\mathbf{s}^1_{t-1}) = [\boldsymbol{\sigma}_i(\mathbf{a}_{t-1}, \mathbf{s}^1_{t-1})']_{\mathbf{a}_{t-1}}$ is a 64×16 matrix and $\Delta\mathbf{AC}^\ell_i = [\Delta AC^\ell_i(\mathbf{a}_{t-1})]_{\mathbf{a}_{t-1}}$ is a 64×1 vector. In matrix notation, for all $\mathbf{s}_{t-1}$, this becomes:

$$\Delta\bar{\mathbf{v}}_i(\ell) = \underbrace{\begin{bmatrix} \boldsymbol{\sigma}_i(\mathbf{s}^1_{t-1}) & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \boldsymbol{\sigma}_i(\mathbf{s}^{2187}_{t-1}) \end{bmatrix}}_{(2187\cdot 64)\times(2187\cdot 16)} \underbrace{\begin{bmatrix} \boldsymbol{\lambda}_i(\ell, \mathbf{s}^1_{t-1}) \\ \vdots \\ \boldsymbol{\lambda}_i(\ell, \mathbf{s}^{2187}_{t-1}) \end{bmatrix}}_{(2187\cdot 16)\times 1} - \Delta\widetilde{\mathbf{AC}}^\ell_i \tag{S.2.6}$$

We will be referring to the block-diagonal matrix containing firm i's beliefs as $\boldsymbol{\sigma}$. It can be written more compactly as a Kronecker product of an identity matrix I and matrix containing beliefs:

$$\begin{aligned} \Delta\bar{\mathbf{v}}_i(\ell) &= \left[I_{2187} \otimes \begin{bmatrix} \boldsymbol{\sigma}_i(\mathbf{s}^1_{t-1}) \\ \vdots \\ \boldsymbol{\sigma}_i(\mathbf{s}^{2187}_{t-1}) \end{bmatrix} \right] \begin{bmatrix} \boldsymbol{\lambda}_i(\ell, \mathbf{s}^1_{t-1}) \\ \vdots \\ \boldsymbol{\lambda}_i(\ell, \mathbf{s}^{2187}_{t-1}) \end{bmatrix} - \Delta\widetilde{\mathbf{AC}}^\ell_i \\ &= \boldsymbol{\sigma}_i\boldsymbol{\lambda}_i(\ell) - \Delta\widetilde{\mathbf{AC}}^\ell_i \end{aligned}$$

Everything we showed so far was for a selected action $\ell \in \mathcal{A}_i \backslash \{HH\}$. We can now define $\Delta\bar{\mathbf{v}}_i = [\bar{\mathbf{v}}_i(HL); \bar{\mathbf{v}}_i(LH); \bar{\mathbf{v}}_i(LL)]'$, so that:

$$\begin{aligned} \Delta\bar{\mathbf{v}}_i &= [I_3 \otimes \boldsymbol{\sigma}_i] \begin{bmatrix} \boldsymbol{\lambda}_i(HL) \\ \boldsymbol{\lambda}_i(LH) \\ \boldsymbol{\lambda}_i(LL) \end{bmatrix} - \begin{bmatrix} \Delta\widetilde{\mathbf{AC}}^{HL}_i \\ \Delta\widetilde{\mathbf{AC}}^{LH}_i \\ \Delta\widetilde{\mathbf{AC}}^{LL}_i \end{bmatrix} \\ &= \mathbf{Z}_i\boldsymbol{\lambda}_i - \Delta\widetilde{\mathbf{AC}}_i \end{aligned} \tag{S.2.7}$$

The dimension of the object on the LHS of (S.2.7) is $(139968 \cdot 3 \times 1) = 419904 \times 1$. Define the following 419904×419904 projection matrix:

$$\mathbf{M}^{\mathbf{Z}}_i = I_{419904} - \mathbf{Z}_i(\mathbf{Z}'_i\mathbf{Z}_i)^{-1}\mathbf{Z}'_i \tag{S.2.8}$$

So far we have not discussed $\Delta\widetilde{\mathbf{AC}}_i$ in detail, but it can be written as: $\Delta\widetilde{\mathbf{AC}}_i = \widetilde{\mathbf{D}}_i\Delta\mathbf{AC}_i$

where $\widetilde{\mathbf{D}}_i$ is a $419904 \times \kappa_i$ matrix of zeros and ones which are a natural consequence of the indicator functions used while defining the profit function. κ_i is the number of dynamic parameters to estimate for firm i and $\Delta\mathbf{AC}_i$ is a $\kappa_i \times 1$ vector of parameters to identify. Multiplying both sides of (S.2.7) by the projection matrix defined in (S.2.8), we have:

$$\begin{aligned}
\mathbf{M}_i^{\mathbf{Z}}\Delta\bar{\mathbf{v}}_i &= -\mathbf{M}_i^{\mathbf{Z}}\widetilde{\mathbf{D}}_i\Delta\mathbf{AC}_i \\
\widetilde{\mathbf{D}}_i'\mathbf{M}_i^{\mathbf{Z}}\Delta\bar{\mathbf{v}}_i &= -\widetilde{\mathbf{D}}_i'\mathbf{M}_i^{\boldsymbol{\lambda}}\widetilde{\mathbf{D}}_i\Delta\mathbf{AC}_i \\
\Delta\mathbf{AC}_i &= -(\widetilde{\mathbf{D}}_i'\mathbf{M}_i^{\mathbf{Z}}\widetilde{\mathbf{D}}_i)^{-1}(\widetilde{\mathbf{D}}_i'\mathbf{M}_i^{\mathbf{Z}}\Delta\bar{\mathbf{v}}_i) \qquad \text{(S.2.9)}
\end{aligned}$$

(S.2.9) defines the identifying correspondence for firm i. We can proceed in an identical fashion to recover the parameters for the remaining firms. There is also a straightforward way to incorporate equality restrictions across firms an estimate $\{\Delta\mathbf{AC}_i\}_{i=1}^N$ for all firms in one step.

Further identifying restrictions

So far we showed how the structure of the model identifies $\{\Delta\mathbf{AC}_i\}_{i=1}^N$, that is the vector of differences in adjustment costs relative to a chosen (baseline) action. For example, with HH being the baseline action, we identify

$$\Delta AC_i^{\ell}(a_{i,t-1} = \ell') = AC_i^{\ell' \to \ell} \cdot \mathbf{1}(a_{i,t-1} = \ell') - AC_i^{\ell' \to HH} \cdot \mathbf{1}(a_{i,t-1} = \ell')$$

Differences are not interpretable as price adjustment costs per se. While it should be straightforward to see that just assuming that switching from any price regime to HH (all products with regular/high price) is always costless would be sufficient to recover $AC_i^{\ell' \to \ell} \cdot \mathbf{1}(a_{i,t-1} = \ell')$ for all ℓ, ℓ', the model remains heavily overparametrised without further restrictions. This is because we are only interested in recovering one parameter per product, which would be interpreted as the cost of putting a particular product on promotion. Following our exposition, there are no reasons to believe that e.g. $AC_i^{HH \to LH}$ should be different from $AC_i^{HL \to LL}$ – where the only difference is that the second product was on promotion in $t-1$ and t, while in the former case i was charging regular (high) price for it. Since the adjustment costs should be invariant to many other combinations of past prices and actions, we spell out three identifying restrictions below as assumptions R1-3. While this is not the only possible set of assumptions allowing for point identification of product-specific adjustment costs, we believe that what we propose has a natural, economically meaningful interpretation.

Assumption (R1). *Adjustment costs are incurred only when switching from high to low price.*

This assumption effectively sets $AC_i^{\ell' \to HH} = 0$ for all ℓ' as well as $AC_i^{\ell' \to HL} = 0$ and $AC_i^{\ell' \to LH} = 0$ if $\ell' = LL$. As discussed above, the first restriction is sufficient to recover absolute, instead of relative, levels of adjustment costs.

Assumption (R2). *Adjustment cost associated with one product is independent of the current and lagged promotional status of other products.*

R2 is a natural assumption, and allows us to impose equality restriction across $\mathbf{a}_{-i,t-1}$ in the switching cost part of (1). Finally, consider the situation in which prices of more than one product of a firm move in the same direction. R3 says that we can express the cost of taking this action as a sum of individual price adjustments of the products involved:

Assumption (R3). *There are no economies of scope associated with price promotions on multiple products of the same firm.*

R1-2 will be sufficient to identify one cost of adjusting prices per product plus the joint cost of putting more than 1 product on promotion at the same time. R3 can then be used to reduce the dimension of the parameter vector to equal to the number of products. The identifying power of our assumptions is summarised by the following proposition:

Proposition 1. *Under assumptions R1-2, the matrix $\widetilde{\mathbf{D}}_i$ satisfies the requirements of theorem 2 in Komarova et al. (2018) and for each firm one can identify $|\mathcal{A}_i| - 1$ parameters in $\mathbf{AC}_i$. Adding assumption R3 reduces the number of parameters to $|\mathcal{J}_i|$.*

We leave the proposition without a proof which amounts to showing that $\widetilde{\mathbf{D}}_i$ has a full column rank when R1-3 are imposed. This can be easily verified numerically. In an earlier version of the paper, we had a simplified duopoly example where we provided an algebraic expression for $\widetilde{\mathbf{D}}_i$ which made it immediately obvious that the matrix had full column rank. For the sake of brevity we suppress this result here, but the derivations can be obtained from the authors upon request.

S.2.2 Discount factor and value function

To estimate the discount factor and subsequently solve the model we have to compute the value functions associated with each element of the state space. Because our state space is large and some state variables are effectively continuous it is computationally infeasible to compute the value function for each state, even with a coarse discretisation of market shares. Likewise we compute the value function for each of the $T = 200$ observed states (for each firm in each supermarket) assuming that value functions can be approximated by a linear function of functions of state variables. The same approach has been used in Sweeting (2013), Barwick and Pathak (2015) and Fowlie et al. (2016). Next we discuss the procedures used to estimate the discount factor.

Using the fact the state transitions in our model are deterministic – see equation (9) – we can write the *ex ante* value function in problem (3) as:

$$V_i\left(\mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right) = \sum_{\mathbf{a_t}\in\mathcal{A}_i} \sigma_i\left(\mathbf{a_t}|\mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right)\left\{\tilde{\Pi}_i\left(\mathbf{a_t}, \mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right) + \beta V_i\left(\mathbf{a_t}, \mathbf{s}\left(\mathbf{a_t}, \mathbf{s_{t-1}}\right)\right)\right\}, \tag{S.2.10}$$

where $V_i(\mathbf{z}_{t+1}) = \int V_i(\mathbf{z}_{t+1}, \varepsilon_{t+1})dQ(\varepsilon_{i,t+1})$ and $\tilde{\Pi}_i\left(\mathbf{a_t}, \mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right)$ is the (conditional) expectation of the payoff function $\Pi_i\left(\mathbf{a_t}, \mathbf{a_{t-1}}, \mathbf{s_{t-1}}, \varepsilon_{it}\left(a_{it}\right)\right)$ with respect to ε_{it} when states are $\left(\mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right)$ and current actions are $\mathbf{a_t}$, and $\mathbf{s}\left(\mathbf{a_t}, \mathbf{s_{t-1}}\right)$ is the vector of current shares – implied by equation (9) – when past shares are $\mathbf{s_{t-1}}$ and current actions are $\mathbf{a_t}$. As in Sweeting (2013) we approximate $V_i(\mathbf{z}_t)$ using the following parametric function:

$$V_i(\mathbf{z_t}) \simeq \sum_{k=1}^{K} \lambda_{ki}\phi_{ki}\left(\mathbf{z_t}\right) \equiv \Phi_i\left(\mathbf{z_t}\right)\lambda_\mathbf{i}, \tag{S.2.11}$$

where λ_{ki} is a coefficient and $\phi_{ki}\left(\cdot\right)$ is a well-defined function mapping the state vector into the set of real numbers. In our case, $\phi_{ki}\left(\cdot\right)$ are flexible functions of shares and prices of the firms. In practice, the variables we use to approximate the value functions include (i) (past) actions of all firms, (ii) second order polynomials of (past) shares of all products, (iii) interactions between (past) actions and shares of the different products and (iv) second order polynomials of the interactions between (past) actions and shares. We experimented with third and fourth order polynomials of shares and interactions between shares and actions but the results did not change significantly.

Notice that under this formulation solving for the value function requires that one computes only K parameters (λ_{ki}'s) for each manufacturer. By substituting this equation into the *ex ante* value function we can solve for $\lambda_\mathbf{i} = [\lambda_{1i}\ \lambda_{2i}\ ...\ \lambda_{Ki}]'$ in closed-form as a function of the primitives of the model, states and beliefs. Substituting (S.2.11) into (S.2.10) we get:

$$\Phi_i\left(\mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right)\lambda_\mathbf{i} = \sum_{\mathbf{a_t}\in\mathcal{A}} \sigma_i\left(\mathbf{a_t}|\mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right)\left\{\tilde{\Pi}_i\left(\mathbf{a_t}, \mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right) + \beta\Phi_i\left(\mathbf{a_t}, \mathbf{s}\left(\mathbf{a_t}, \mathbf{s_{t-1}}\right)\right)\lambda_\mathbf{i}\right\}.$$

To simplify the notation let $\tilde{\Pi}_i^*\left(\mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right)$ and $\Phi_i^*\left(\mathbf{s_{t-1}}\right)$ be the conditional expectations of $\tilde{\Pi}_i\left(\mathbf{a_t}, \mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right)$ and of $\Phi_i\left(\mathbf{a_t}, \mathbf{s}\left(\mathbf{a_t}, \mathbf{s_{t-1}}\right)\right)$ with respect to current actions, respectively. Therefore, we can rewrite equation above as:

$$\left(\Phi_i\left(\mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right) - \beta\Phi_i^*\left(\mathbf{s_{t-1}}\right)\right)\lambda_\mathbf{i} = \tilde{\Pi}_i^*\left(\mathbf{a_{t-1}}, \mathbf{s_{t-1}}\right).$$

Stacking this equation for every possible state in S we have that:

$$\left(\Phi_i - \beta\Phi_i^*\right)\lambda_\mathbf{i} = \tilde{\Pi}_i^*,$$

where Φ_i and Φ_i^* are $N_s \times K$ matrices that depend on states and beliefs and $\tilde{\Pi}_i^*$ is a $N_s \times 1$ vector of expected profits that depends on state, beliefs and parameters, N_s being the number of states observed in the data. Assuming $K < N_s$, this expression can be rewritten as:

$$\lambda_{\mathbf{i}} = \left[\left(\Phi_i - \beta \Phi_i^* \right)' \left(\Phi_i - \beta \Phi_i^* \right) \right]^{-1} \left[\left(\Phi_i - \beta \Phi_i^* \right)' \tilde{\Pi}_i^* \right]. \qquad \text{(S.2.12)}$$

Inserting (S.2.12) into (S.2.11) we obtain the unconditional value functions associated with problem (3); given the logit assumption on ε_{it} we can calculate the probability of each action solving problem (3). Having estimated adjustment costs outside of the dynamic model and having calibrated H and marginal costs, the only parameter to be estimated inside the dynamic model is the discount factor. We do this by choosing the discount factor that minimises the difference between estimated action probabilities and the probabilities implied by the structural model, which are defined based on the approximation explained above (see Komarova et al. (2018)).

S.2.3 Model solution

To solve the model we use an algorithm similar to that described in Sweeting (2013). The algorithm works as follows:

1. In step s we calculate $\lambda(\sigma^s)$ as a function of the vector of beliefs, σ^s, substituting equation (S.2.11) into the *ex-ante* value function and solving for $\lambda = [\lambda_1\, \lambda_2 \dots \lambda_k]$ in closed-form as a function of the primitives of the model, states and beliefs;

2. We use $\lambda(\sigma^s)$ to calculate choice specific value functions for each of the selected states and the multinomial logit formula implied by the model to update the vector of beliefs, $\tilde{\sigma}$;

3. If the value of the euclidian norm $\|\sigma^s - \tilde{\sigma}\|$ is sufficiently small we stop the procedure and save $\tilde{\sigma}$ as the equilibrium vector of probabilities implied by the model, $\tilde{\sigma} = \sigma^*$; if $\|\sigma^s - \tilde{\sigma}\|$ is larger than the tolerance we update $\sigma^{s+1} = \psi\tilde{\sigma} + (1-\psi)\sigma^s$, where ψ is a number between 0 and 1, and restart the procedure.

The tolerance used on $\|\sigma^s - \tilde{\sigma}\|$ was 10^{-3} and the value of ψ used to update σ^s to σ^{s+1} was 0.5. We have made several attempts using lower values for the tolerance on $\|\sigma^s - \tilde{\sigma}\|$ and for ψ. All these attempts generated very similar equilibrium probabilities, but the time to achieve convergence was larger. The initial guess used to start the algorithm, σ^0, is equal to the estimated CCPs evaluated at the corresponding state. To check the robustness of our results to changes in the initial guess we changed arbitrarily the original initial guess multiplying it by several factors between 0 and 1. For all our attempts the resulting equilibrium vector of probabilities was the same.

For the counterfactuals we have to simulate the model for states that are not observed in the data – i.e. we need estimates of σ^* for states that are not in the data. To do this we assumed that the solution of the model, σ^*, for the relevant counterfactual scenario is a logistic function of a linear index of states – i.e. the same function that we used to compute the CCPs. Mathematically, let $\sigma_i^*(a_i = k|\mathbf{z})$ be the probability that firm i plays $a_i = k$ when the state vector is $\mathbf{z}$. We assume that:

$$\sigma_i^*(a_i = k|\mathbf{z}) = \frac{\exp(\mathbf{z}'\gamma_k)}{\sum_{k'}\exp(\mathbf{z}'\gamma_{k'})}. \tag{S.2.13}$$

Dividing it by the probability of an anchor choice, say $a_i = HH$, normalising $\gamma_1 = 0$ and taking logs we have $\ln\{\sigma_i^*(a_i = k|\mathbf{z})\} - \ln\{\sigma_i^*(a_i = HH|\mathbf{z})\} = \mathbf{z}'\gamma_k$. Then the vector of parameters γ_k can be estimated by OLS – one OLS equation is estimated for each $a_i = k$, $k \neq HH$.

The probability function (S.2.13) and the Markovian transitions for actions and shares are used to simulate moments implied by the model. Starting from the initial state vector for each firm in each supermarket we forward simulate 1000 paths of 200 periods of actions and shares and computed profits for each period by averaging period profits for each path.

References

BARWICK, P. J. AND P. A. PATHAK (2015): "The Costs of Free Entry: An Empirical Study of Real Estate Agents in Greater Boston," *RAND Journal of Economics*, 46, 103–145.

FOWLIE, M., M. REGUANT, AND S. P. RYAN (2016): "Market-Based Emissions Regulation and Industry Dynamics," *Journal of Political Economy*, 124, 249–302.

GRIFFITH, R., L. NESHEIM, AND M. O'CONNELL (2017): "Income effects and the welfare consequences of tax in differentiated product oligopoly," *Quantitative Economics*, 9, 305–341.

KOMAROVA, T., F. A. M. SANCHES, D. SILVA JR., AND S. SRISUMA (2018): "Joint Analysis of the Discount Factor and Payoff Parameters in Dynamic Discrete Choice Models," *Quantitative Economics*, 9, 1153–94.

SWEETING, A. (2013): "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry," *Econometrica*, 81, 1763–1803.