Common price dynamics in consumer search estimates

Thiago Cacicedo[*] Tommy Ranyard[†]

Heriot-Watt University University of Edinburgh

Abstract

Non-sequential structural search models infer consumer search intensity from observed price dispersion. We show that one-way demeaning, the standard correction for seller heterogeneity, leaves common time-varying price movements in the pooled distribution. When the model maps the width or variance of this distribution into search intensity, the remaining common component generates spurious changes across estimation windows. We demonstrate that two-way demeaning, by seller and period, removes this additive component exactly. Applied to Spanish diesel prices around a 2015 antitrust sanction, one-way demeaning implies a 27.7 p.p. decline in fuel stations searched. Two-way demeaning eliminates this change, leaving search-relevant price gaps unchanged.

Keywords: consumer search, structural search models, price dispersion, non-stationarity

JEL classification: C14, D83, L40, L71

[*]Email: t.cacicedo@hw.ac.uk

[†]Email: s2854416@ed.ac.uk

1 Introduction

Structural models of consumer non-sequential search provide a way to recover heterogeneity in consumers' search intensity from price data alone. Building on equilibrium models of price dispersion with search frictions (e.g. Burdett and Judd 1983) and empirical approaches that infer search-cost heterogeneity from market price distributions (Hong and Shum 2006; Moraga-González and Wildenbeest 2008; Wildenbeest 2011), this literature exploits the link between pricing equilibrium and consumer search behaviour. In the class of models developed by Wildenbeest (2011), the symmetric mixed-strategy equilibrium implies a direct mapping from the support and shape of the price distribution to the shares of consumers sampling different numbers of sellers. Estimated search shares are therefore interpreted as measures of consumer search intensity, even when individual search behaviour is unobserved.

In applications to small retail markets (e.g. Nishida and Remer 2018*a*,*b*), the relevant price distribution is rarely observed in a single cross-section. A local fuel market, supermarket category, or neighbourhood retail market may contain only a handful of sellers, leading researchers to pool prices over time and estimate the model from the resulting distribution of demeaned prices. This strategy is valid only if the underlying equilibrium price distribution is stationary over the estimating window. If common price dynamics are present, time-series variation is folded into the pooled distribution and may be misinterpreted as equilibrium price dispersion generated by consumer search and firm pricing incentives.

This paper studies the consequences of violating the stationarity assumption. One-way seller demeaning removes persistent seller heterogeneity, but it leaves common time-varying price components in the pooled distribution. Common cost movements, aggregate trends, wholesale shocks, or policy-induced shifts can therefore enter the object from which search is inferred. To isolate the mechanism, we study two closed-form matching exercises: one in which the stationary model is forced to match the support width of the pooled distribution, and one in which it is forced to match its variance. In both cases, the remaining common component is mechanically mapped into higher search, so differences in the common price environment across two windows can appear as differences in search even when behaviour is unchanged. For the full likelihood estimator, the mapping is less transparent because support, scale, and shape are fitted jointly. The empirical question is therefore diagnostic: does an estimated change in search survive once market-period common movements are removed?

This concern is related to an issue already recognised in the literature. Wildenbeest (2011, p. 754) notes that unobserved common shocks may be absorbed into the price variation that the model interprets as the outcome of mixed strategies, and hence as search. He sets the issue aside because, in a static single-window application, firms still face strategic uncertainty whether the dispersion comes from genuine mixed strategies or unobserved shocks. That argument is appropriate for estimating a level of search in one stable window. It does not, however, protect comparisons across windows. If the common price environment changes between two periods, the pooled distribution can change even when consumer behaviour does not. The distinction between levels and differences is therefore central.

The two-way transformation we use (seller and period) is not new. Existing applications already employ period controls as robustness checks and find that they change little. Wildenbeest (2011) re-estimates his model with time fixed effects and reports near-identical estimates. Nishida and Remer (2018*a*) control for wholesale costs directly and, adding time fixed effects, likewise find their estimates unchanged. We argue that this innocuousness is not a general property of the estimator. In those applications, the common component is small, either because the market is stable (a twelve-week grocery panel, Wildenbeest 2011) or because the estimation window is chosen to minimise wholesale cost variation (Nishida and Remer 2018*a*). Our contribution is to show what happens when that condition fails. We characterise how common within-window movements remain in the one-way pooled object, show that two-way demeaning removes the additive component exactly, and use the one-way/two-way gap as a diagnostic for whether estimated changes are carried by market-period common movements. We complement the algebra with a Monte Carlo exercise that checks clean-data recovery and finite-sample demeaning bias in short panels.

The issue becomes especially relevant when the estimator is used for before-and-after comparisons. Such comparisons are motivated by evidence that consumer search frictions respond to changes in market and information environments (Luco 2019; Araujo and Rodrigues 2020). But when estimation windows are fixed by the timing of an event, common price movements are harder to avoid, and period controls that are innocuous in stable panels can become pivotal. Reduced-form studies of search around events typically remove common time variation by design. For example, Allen et al. (2014), studying a merger in the Canadian mortgage market, measure rates net of the contemporaneous bond rate and identify distributional effects against a control group, so the common interest-rate environment does not enter their dispersion estimates. The pitfall we identify is specific to transplanting the pooled, single-market non-sequential search estimator into a before-and-after setting without comparable safeguards. A robustness check that passes in a stable window is therefore not evidence that it would pass across an event.

We illustrate our result in an empirical application focusing on the Spanish diesel retail market. In February 2015, the Spanish competition authority publicly sanctioned major retailer brands for collusive practices, an event that received widespread media coverage.[1] We use station-level price data and Wildenbeest (2011)'s model to recover the distribution of consumer search intensity before and after the announcement in several isolated markets.

Additionally, that period also had severe fluctuations in the international oil market. In the beginning of 2015 the oil price was collapsing due to supply rising and demand growth was weak (International Energy Agency 2015). On 16 January 2015 Brent price rose to \$50.16 before falling back.[2] However, by 17 February it rose up to \$63.53.[3] Therefore, given all the cost fluctuation and possible effects from the antitrust sanction, the price distribution is not stationary during the period, and it

[1] See, for example, https://elpais.com/cat/2015/02/25/economia/1424859682_727919.html, https://www.elmundo.es/economia/2015/02/25/54edb48922601d984a8b457b.html, https://en.kiosko.net/es/2015-02-25/np/eleconomista.html

[2] Source: https://www.reuters.com/article/us-markets-oil-idUSKBN0KP06L20150116/

[3] https://www.reuters.com/article/us-markets-oil-idUSKBN0LL02T20150217/

provides an ideal setting to illustrate our results.

Using one-way demeaning, the estimates suggest that search behaviour decreases following the disclosure of collusion. The expected fraction of stations searched falls by 27.6 p.p. after the announcement, a decline that appears in 95.5% of isolated markets. This is driven by an implausibly high pre-announcement estimate of the exhaustive-search share, i.e. consumers searching all stations in a market. The fraction of consumers estimated to search all stations in their market is 0.349 before the announcement, well above the values reported in the prior literature.

When we apply two-way demeaning, the change disappears. The expected fraction of stations searched falls by only 0.9 p.p., and the decline is negative in roughly half of markets (52.6%, against 95.5% under one-way demeaning). A paired comparison test does not reject equality of the two periods. The implausible pre-announcement level is itself an artefact of the same contamination. Holding the period fixed, the pre-announcement exhaustive-search share falls from 0.349 under one-way demeaning to 0.038 under two-way demeaning, the latter in line with Nishida and Remer (2018*a*), who estimate an exhaustive-search share near 0.06 in three-station isolated gasoline markets in the U.S. The inflated pre-period level and its subsequent collapse both arise from common price movements that one-way demeaning leaves in the data, and both are removed once those movements are absorbed in our two-way demeaned approach.

The rest of the paper is organised as follows. Section 2 presents the Wildenbeest (2011) search model and explains why empirical implementations pool demeaned prices over time. Section 3 introduces the additive common price component and shows how one-way demeaning leaves it in the pooled distribution. Section 4 derives the resulting mapping into search in tractable matching cases, then develops the two-way demeaning diagnostic and its limits. Section 5 illustrates the mechanism numerically and reports the Monte Carlo diagnostic for clean short panels. Section 6 applies the diagnostic to the Spanish diesel market around the 2015 antitrust sanction, and Section 7 concludes. Proofs are in the appendix.

2 Model

2.1 The equilibrium model

We follow Wildenbeest (2011), matching its notation. The model is built in the fixed-sample search framework of Burdett and Judd (1983) and it extends Moraga-González and Wildenbeest (2008) to allow for vertically differentiated products. The key object for our purposes is the distribution of utilities generated by firms' mixed pricing strategies. The empirical estimator recovers the search distribution from this price-distribution object, so any non-stationary component that enters the pooled distribution can be read as a change in search.

A market has N sellers, indexed by j, offering a homogeneous good at unit cost r_j. The good is bundled with store's characteristics, so sellers may be vertically differentiated even though the underlying physical product is the same. There is a unit mass of uninformed consumers, each demanding one unit.

Consumers differ in search costs. Search is non-sequential, i.e. before observing prices, a fraction μ_k of consumers commits to inspecting exactly k sellers, $k = 1, \ldots, N$, with $\sum_k \mu_k = 1$, and then buys from the inspected seller yielding the highest utility.[4]

Consumers value the good of seller j at v_j, so a consumer paying p_j obtains utility $u_j = v_j - p_j$. Under the assumptions on quality production (competitive quality-input markets and constant returns to scale) the valuation-cost markup is constant across sellers,

$$x \equiv v_j - r_j. \tag{1}$$

The equilibrium can therefore be represented as a symmetric mixed strategy over utility levels. Each seller randomises its price so that the implied utility $u_j = v_j - p_j$ is a draw from a common atomless cdf $L(u)$ on a common support $[0, \bar{u}]$. The margin

[4] This is also known as fixed-sample search.

earned when the seller offers utility u is $x - u = p_j - r_j$. Seller specific levels, (v_j, r_j), shift price distributions in levels, but they do not change the common utility distribution L.

A seller that offers utility u is sampled by a type-k consumer with probability k/N. Conditional on being sampled, it wins the sale if its utility exceeds the $k-1$ other sampled draws, an event with probability $L(u)^{k-1}$. Expected profit is therefore

$$\pi(u) = (x-u)\,\frac{1}{N}\sum_{k=1}^{N} k\,\mu_k\,L(u)^{k-1}. \tag{2}$$

In equilibrium $\pi(u)$ is constant on the support of L. This indifference condition links the shape and support of the price distribution to the search shares $\{\mu_k\}$, which is why the model can be estimated from price dispersion.

2.1.1 Why does estimation pool over time?

The econometrician observes prices p_{jt} for sellers $j = 1, \dots, N$ over periods $t = 1, \dots, T$. In local markets where this model is usually taken to the data, a single cross-section contains too few prices to estimate the search distribution. The implementation therefore pools prices over time. Pooling is a sampling device that it is valid only if the observations in the estimation window are draws from the same stationary equilibrium distribution.

Before pooling, empirical applications remove seller-specific price levels. In utility notation, the one-way seller-demeaned object is

$$\hat{u}_{jt} = \bar{p}_j - p_{jt}, \qquad \bar{p}_j = \tfrac{1}{T}\textstyle\sum_t p_{jt}, \tag{3}$$

and the estimator fits the equilibrium distribution implied by (2) to the pooled sample $\{\hat{u}_{jt}\}_{j,t}$.

Seller demeaning removes fixed seller level differences. It does not, by itself, remove a price component that is common across sellers but varies over time. The next section studies exactly that departure from the stationary condition needed for pooling.

2.2 Equilibrium preliminaries

2.2.1 Support width in the general model

LEMMA 1 (Support width). *In any symmetric mixed-strategy equilibrium of the model with search distribution* $\{\mu_k\}$ *such that* $\mu_1 > 0$*, the support of L is* $[0, \bar{u}]$ *with*

$$\bar{u} \;=\; x\left(1 - \frac{\mu_1}{\mathbb{E}[k]}\right), \qquad \mathbb{E}[k] \equiv \sum_{k=1}^{N} k\,\mu_k. \tag{4}$$

$\bar{u}$ *is strictly decreasing in* μ_1 *and strictly increasing in* $\mathbb{E}[k]$ *holding the other fixed.*

Lemma 1 shows why support width is informative about search intensity. Holding x fixed, a wider utility support implies a lower captive-to-search ratio $\mu_1/\mathbb{E}[k]$, and thus more effective search. Conversely, anything outside the stationary equilibrium distribution that mechanically widens the pooled empirical support gives the model a reason to report more search unless another fitted object, such as x, absorbs it.

2.2.2 The two-type case

For the analytical bias results in Section 4, it is useful to focus on the two-type case with captives and exhaustive searchers only. This restriction collapses the search distribution to a single search-intensity parameter, μ_N, and makes the support width and variance of the equilibrium utility distribution explicit functions of that parameter.

The simplification allows us to obtain a one-dimensional benchmark in which the effect of common price dynamics on inferred search can be derived in closed form.

Formally, let $\mu_1 = \alpha$, $\mu_N = 1-\alpha$, and $\alpha \in (0,1)$. Thus, α is the share of captive consumers, i.e. consumers who sample only one seller. Then (2) becomes

$$\pi(u) = (x-u)\left[\frac{\alpha}{N} + (1-\alpha)L(u)^{N-1}\right].$$

Equating this expression to the endpoint profit $\pi(0) = x\alpha/N$ gives

$$L(u)^{N-1} = \frac{\alpha}{N(1-\alpha)} \cdot \frac{u}{x-u} \quad \Longrightarrow \quad L(u) = \left(\frac{A\,u}{x-u}\right)^{\frac{1}{N-1}}, \qquad \text{where } A \equiv \frac{\alpha}{N(1-\alpha)}.$$

Setting $L(\bar{u}) = 1$ gives $\bar{u} = x/(1+A)$, which coincides with (4) since $\mathbb{E}[k] = \alpha + N(1-\alpha)$ implies $1 - \alpha/\mathbb{E}[k] = N(1-\alpha)/[\alpha + N(1-\alpha)] = 1/(1+A)$. For $N = 2$ this is $\bar{u} = x(1-\alpha)/(1-\alpha/2)$.

2.3 Moments of L: quantile representation

The second matching benchmark in Section 4 uses the variance of the pooled demeaned distribution. We first derive the variance of the equilibrium utility distribution in the two-type case, as they provide the closed-form object that the variance-matching exercise later inverts.

In the two-type case these moments are easiest to obtain from the quantile function. Inverting L gives, for $w \in [0,1]$,

$$u(w) \;=\; x\,\frac{w^{N-1}}{A + w^{N-1}} \;=\; x\left(1 - \frac{A}{A + w^{N-1}}\right), \tag{5}$$

and for any integrable g, $\mathbb{E}[g(u)] = \int_0^1 g(u(w))\,dw$.

General N. Using the series $(A + w^s)^{-1} = A^{-1}\sum_{k\geq 0}(-1)^k w^{sk} A^{-k}$ for $A > 1$ (and by analytic continuation otherwise), with $s = N-1$,

$$\begin{aligned}\int_0^1 \frac{dw}{A+w^s} &= \frac{1}{A}\sum_{k\geq 0}\frac{(-1/A)^k}{sk+1}\\ &= \frac{1}{A}\,{}_2F_1\!\left(1, \tfrac{1}{s}; 1+\tfrac{1}{s}; -\tfrac{1}{A}\right),\end{aligned}$$

using $(1)_k = k!$ and $(\frac{1}{s})_k/(1+\frac{1}{s})_k = 1/(1+sk)$. Hence

$$\mathbb{E}[u] = x\left[1 - A\int_0^1 \frac{dw}{A+w^{N-1}}\right] = x\left[1 - {}_2F_1\!\left(1, \tfrac{1}{N-1}; \tfrac{N}{N-1}; -\tfrac{1}{A}\right)\right]. \tag{6}$$

Second moments follow analogously from $u(w)^2 = x^2[1 - A/(A+w^{N-1})]^2$. The cases $N = 2$ and $N = 3$ are elementary.

Case $N = 2$**.** Write $\tilde{L} \equiv \ln\left(1+\frac{1}{A}\right)$, $A = \alpha/(2(1-\alpha))$. Then

$$\mathbb{E}[u] = x\int_0^1 \left(1 - \frac{A}{A+w}\right)dw = x\left[1 - A\,\tilde{L}\right], \tag{7}$$

$$\mathbb{E}[u^2] = x^2\int_0^1 \left(1 - \frac{A}{A+w}\right)^2 dw = x^2\left[1 - 2A\tilde{L} + A^2\int_0^1 \frac{dw}{(A+w)^2}\right] = x^2\left[1 - 2A\tilde{L} + \frac{A}{1+A}\right], \tag{8}$$

using $\int_0^1 (A+w)^{-2}dw = \frac{1}{A} - \frac{1}{1+A} = \frac{1}{A(1+A)}$. Subtracting the square of (7) from (8),

$$\sigma_L^2(\alpha; 2) = x^2 A\left[\frac{1}{1+A} - A\,\tilde{L}^2\right], \qquad \text{where } A = \frac{\alpha}{2(1-\alpha)}. \tag{9}$$

Case $N = 3$**.** Write $\Phi \equiv \sqrt{A}\,\arctan\left(1/\sqrt{A}\right)$, $A = \alpha/(3(1-\alpha))$. Using $\int_0^1 (A+w^2)^{-1}dw = A^{-1/2}\arctan(A^{-1/2})$ and the standard reduction

$$\int_0^1 \frac{dw}{(A+w^2)^2} = \left[\frac{w}{2A(A+w^2)}\right]_0^1 + \frac{1}{2A}\int_0^1 \frac{dw}{A+w^2}$$
$$= \frac{1}{2A(1+A)} + \frac{\arctan(1/\sqrt{A})}{2A^{3/2}},$$

we obtain

$$\mathbb{E}[u] = x\,[\,1 - \Phi\,], \tag{10}$$
$$\mathbb{E}[u^2] = x^2\left[1 - 2\Phi + \frac{A}{2(1+A)} + \frac{\Phi}{2}\right], \tag{11}$$

and therefore

$$\sigma_L^2(\alpha;3) = x^2\left[\frac{\Phi}{2} - \Phi^2 + \frac{A}{2(1+A)}\right], \qquad \Phi = \sqrt{A}\arctan\frac{1}{\sqrt{A}}, \quad \text{where } A = \frac{\alpha}{3(1-\alpha)}. \tag{12}$$

Remark 1 (Non-monotonicity of the variance). *σ_L^2 is hump-shaped in $\mu_N = 1-\alpha$. It vanishes at both $\mu_N \to 0$ (monopoly corner, L degenerate at 0) and $\mu_N \to 1$ (Bertrand corner, L degenerate at x), and peaking at $\mu_N^* \approx 0.698$ for $N = 2$ and $\mu_N^* \approx 0.804$ for $N = 3$. On the increasing branch $\mu_N < \mu_N^*$, which contains the estimates in our application (see Section 6) and in the published literature, σ_L^2 is strictly increasing in μ_N, and the variance-based mapping discussed in the in Section 4 is invertible. The support width $\bar{u}$, by contrast, is globally strictly increasing in μ_N (Lemma 1), so the width-based results require no branch restriction.*

3 Common price dynamics under one-way demeaning

Assumption 1 (Additive common component). *Observed prices satisfy*

$$p_{jt} = \xi_j + \delta_t^c + p_{jt}^*, \qquad p_{jt}^* = \rho - u_{jt}^*, \tag{13}$$

where ξ_j is a seller fixed effect capturing seller specific price levels, δ_t^c is a common price component identical across sellers within the market-period, ρ is a constant, and $u_{jt}^* \overset{iid}{\sim} L(\cdot\,;\mu)$ across j and t, independent of $\{\delta_t^c\}$.

The common component δ_t^c should be interpreted as a price shifter that is common to all sellers in a market-period, for instance input-cost pass-through, wholesale price movements, exchange-rate or tax changes, or seasonal demand. This component is not meant to describe a full re-solution of the static search equilibrium with a time-varying marginal cost. A literal cost shock would also move the reservation margin x and hence the entire equilibrium distribution.

Assumption 1 isolates the effect of an additive common component that lies outside the stationary distribution fitted by the econometrician. It nests the implementation's maintained assumption as the special case in which δ_t^c is constant within the estimation window. Heterogeneous or lagged responses across sellers are discussed in Section 4.3.

Define the in-window common deviations as the following:

$$\delta_t \equiv -\left(\delta_t^c - \bar{\delta}^c\right), \qquad \bar{\delta}^c = \tfrac{1}{T}\textstyle\sum_t \delta_t^c, \qquad \sigma_\delta^2 \equiv \tfrac{1}{T}\textstyle\sum_t \delta_t^2, \qquad R_\delta \equiv \max_t \delta_t - \min_t \delta_t. \tag{14}$$

In the proposition below, the common path refers to the realised sequence $\{\delta_t\}_{t=1}^T$ within the estimating window in a market. Repeated sampling averages over the idiosyncratic equilibrium draws u_{jt}^* while holding this sequence fixed.

Proposition 1 (Contamination representation). *Under Assumption 1, the one-way demeaned utilities (3) satisfy*

$$\hat{u}_{jt} \;=\; \left(u^*_{jt} - \bar{u}^*_j\right) \;+\; \delta_t, \qquad \bar{u}^*_j = \tfrac{1}{T}\sum_t u^*_{jt}. \tag{15}$$

Consequently, pooling over j and t:

(i) ***(Variance.)*** *Under repeated sampling of the idiosyncratic equilibrium draws, conditional on the realised common path $\{\delta_t\}_{t=1}^{T}$ (or, equivalently, averaging over many markets or replications that share the same N, T, and common path), the pooled variance of $\hat{u}$ converges to*

$$\mathrm{Var}(\hat{u}) \;=\; \frac{T-1}{T}\,\sigma^2_L(\mu) \;+\; \sigma^2_\delta. \tag{16}$$

(ii) ***(Support.)*** *As $T \to \infty$ with δ_t drawn from a distribution with bounded support of range R_δ, the pooled $\hat{u}$ are distributed as the convolution of the equilibrium utility distribution and the distribution of δ, up to the irrelevant centering by the mean of L. The support has width*

$$W \;=\; \bar{u}(\mu) \;+\; R_\delta. \tag{17}$$

In empirical applications, the estimation step fits the stationary family $L(\cdot\,;\mu)$ to the pooled demeaned observations. Proposition 1 shows what goes wrong before estimation. One-way demeaning leaves the common component in the pooled object. It adds dispersion and support width that are not part of the stationary equilibrium distribution. How the full likelihood allocates that extra variation between the markup/support parameter and the search shares is a likelihood-specific question. In the following sections, we will show the matching results that isolate the mechanism in transparent cases.

4 How are common dynamics mapped into search?

We characterise the mapping for two transparent matching estimators: width matching and variance matching. They show why common dynamics can be interpreted as search when the model is forced to fit a stationary price distribution to a non-stationary pooled object.

The objects in this section are diagnostic exercises. We do not derive an analogous closed-form result for the full likelihood because the likelihood does not assign the common component to search through a single moment. The full estimator jointly chooses support, scale, and search shares to fit the entire pooled distribution. Each matching exercise, instead, isolates one channel through which the common component left by one-way demeaning can enter the search estimates.

Width matching asks what search distribution would rationalise the observed pooled support width if that width were attributed entirely to the stationary equilibrium distribution. Variance matching asks the analogous question for the pooled variance. The variance result is stated for the two-type restriction because a scalar variance moment cannot identify a full search distribution without additional restrictions. Thus, we use it to show the sign of the mechanism in a one-dimensional case where the model implied variance can be inverted. Thus, both exercises answer the counterfactual question: if the econometrician forced the stationary search model to match one contaminated moment of the one-way demeaned pooled data, what search parameters would be inferred?

Throughout this section, x is treated as known. Remark 2 discusses why the full maximum-likelihood estimator, which also estimates the support and markup, need not reproduce the same magnitudes.

PROPOSITION 2 (Spurious search: level bias). *Let μ^0 denote the true search distribution.*

(i) Width matching, general μ. *Suppose $0 \leq W \leq x$. Any estimator that equates the model implied support width $\bar{u}(\hat{\mu})$ to the observed pooled width W recovers*

$$\widehat{\left(\frac{\mu_1}{\mathbb{E}[k]}\right)} = 1 - \frac{W}{x} = \frac{\mu_1^0}{\mathbb{E}^0[k]} - \frac{R_\delta}{x}. \tag{18}$$

(ii) Variance matching, two-type. *On the increasing branch* $\mu_N < \mu_N^*$ *(Remark 1), and provided the argument below lies in the range of* g_N*, the estimator defined by equating the model-implied variance to the pooled variance,* $g_N(\hat{\mu}_N) = \mathrm{Var}(\hat{u})$ *with* $\mathrm{Var}(\hat{u})$ *given by (16), recovers*

$$\hat{\mu}_N = g_N^{-1}\left(\frac{T-1}{T} g_N(\mu_N^0) + \sigma_\delta^2\right) > \mu_N^0 \quad \textit{whenever } \sigma_\delta^2 > \tfrac{1}{T} g_N(\mu_N^0), \tag{19}$$

where $g_N(\mu_N) \equiv \sigma_L^2(1-\mu_N; N)$ *is given in closed form by (9)–(12) for* $N \in \{2, 3\}$.

Equation (18) shows that the captive-to-search ratio is understated by exactly R_δ/x whenever $R_\delta > 0$. Equivalently, the support-based search index $1 - \mu_1/\mathbb{E}[k]$ is overstated by R_δ/x. If $W > x$, no interior parameter rationalises the observed support width and the estimator is pushed to the boundary of the admissible parameter space.

For variance matching, the familiar expression $\hat{\mu}_N = g_N^{-1}(g_N(\mu_N^0) + \sigma_\delta^2)$ is obtained as the $T \to \infty$ limiting case. On the increasing branch, $\hat{\mu}_N$ is strictly increasing in σ_δ^2, with $\partial\hat{\mu}_N/\partial\sigma_\delta^2 = 1/g_N'(\hat{\mu}_N) > 0$. If the argument exceeds $\max_\mu g_N$, the variance-matching estimator is pushed to a boundary or onto the decreasing branch.

Remark 2 (Estimated x and the limits of moment matching). *In practice* x *is estimated jointly (from the price level and the profit condition, or from support endpoints), and this matters for how the contamination transmits. The matching estimators above route the bias through a single low-order moment, either support width or pooled variance. The full maximum-likelihood estimator does not. Because* x *is recovered from the support of the demeaned data, an additive common component that widens the support can be absorbed largely into* $\hat{x}$ *rather than into* $\hat{\mu}$*. The likelihood reads search shares from the shape of the pooled distribution relative to its estimated support, so the movement in* $\hat{\mu}$ *is a nonlinear* *functional of the entire contaminated distribution rather than a closed-form function of one moment. Width and variance matching establish the mechanism in tractable cases, but they are not a magnitude formula for the full estimator. This is why the empirical strategy is a diagnostic comparison (Section 4.2) rather than a moment-based correction.*

4.1 Event-window comparisons

We now consider the exercise of estimating the model separately across two windows, such as pre- and post-event periods.

PROPOSITION 3 (Spurious change across windows). *Let the search distribution be constant across two windows $w \in \{pre, post\}$, with common-component dispersion $\sigma^2_{\delta,\mathrm{pre}} \neq \sigma^2_{\delta,\mathrm{post}}$ (respectively, ranges $R_{\delta,w}$), each window satisfying the admissibility conditions of Proposition 2. Then the pooled estimator of Proposition 2(ii) recovers*

$$\hat{\mu}_{N,\mathrm{post}} - \hat{\mu}_{N,\mathrm{pre}} = g_N^{-1}\Big(\tfrac{T-1}{T} g_N(\mu_N^0) + \sigma^2_{\delta,\mathrm{post}}\Big) - g_N^{-1}\Big(\tfrac{T-1}{T} g_N(\mu_N^0) + \sigma^2_{\delta,\mathrm{pre}}\Big), \quad (20)$$

which is strictly negative if and only if $\sigma^2_{\delta,\mathrm{post}} < \sigma^2_{\delta,\mathrm{pre}}$. A window with calmer common dynamics is read as a collapse in search, with no change in behaviour. The analogous statement holds for width matching with $R_{\delta,w}$.

COROLLARY 1 (Trends and turning points). *Suppose within window w the common component follows a linear trend, $\delta_t^c = b_w t$, $t = 1, \ldots, T$. Then*

$$\sigma^2_{\delta,w} = b_w^2 \,\frac{T^2 - 1}{12}, \qquad R_{\delta,w} = |b_w|\,(T-1). \quad (21)$$

Hence, (i) levels are biased towards more search in any window containing a common trend; (ii) a spurious decline in measured search arises whenever the post window has a flatter common path, $|b_{\mathrm{post}}| < |b_{\mathrm{pre}}|$, as when the event window straddles a turning point in input costs or other common drivers.

One economic source of asymmetric window steepness in retail settings is asymmetric cost pass-through ("rockets and feathers"), i.e. prices responding sluggishly to falling input costs and rapidly to rising ones (see Bacon 1991; Borenstein et al. 1997; Tappata 2009; Lewis 2011).

In empirical before-and-after analyses it is common to test for similar effects on random dates as a placebo exercise. Proposition 3 clarifies what such a placebo can and cannot rule out. A placebo window pair with symmetric common dynamics ($\sigma^2_{\delta,\text{pre}} \approx \sigma^2_{\delta,\text{post}}$) yields $\Delta\hat{\mu}_N \approx 0$ even though both levels are contaminated. A date-shift placebo is therefore informative about whether the result is specific to the chosen calendar date, but it is much less informative about whether the within-window price distribution is stationary. A placebo window without comparable common price dynamics may pass, or produce a much smaller effect, even when the before-and-after estimate is driven by nonstationary common movements.

4.2 Removing common components: a two-way demeaning diagnostic

Define the two-way (seller and market-period) residualised utilities

$$\tilde{u}_{jt} = -\left(p_{jt} - \bar{p}_{j\cdot} - \bar{p}_{\cdot t} + \bar{p}\right), \tag{22}$$

with $\bar{p}_{j\cdot}$, $\bar{p}_{\cdot t}$, and $\bar{p}$ denoting the seller, period, and grand means within the market-window. The two-way transformation has a diagnostic role. It asks whether the estimated change survives after removing market-period movements that are common to all sellers.

Proposition 4 (Exact removal and scale). *Under Assumption 1:*

(i) $\tilde{u}_{jt}$ does not depend on $\{\delta^c_t\}$: the common component is removed exactly, for any common path and any window.

(ii) For a balanced $N \times T$ panel,

$$\mathbb{E}\big[\tilde{u}_{jt}^2\big] = \sigma_L^2(\mu)\left(1 - \frac{1}{N}\right)\left(1 - \frac{1}{T}\right), \tag{23}$$

so the degrees-of-freedom corrected variance statistic $\hat{\sigma}_L^2 = \widehat{\mathrm{Var}}(\tilde{u}) / \big[(1-1/N)(1-1/T)\big]$ *is consistent for* $\sigma_L^2(\mu)$ *under the stated repeated sampling. In the two-type case, and on the increasing branch of* g_N*, inverting* g_N *at this corrected variance recovers* μ_N *under the additive common-component DGP.*

COROLLARY 2 (The diagnostic localises the source of the change). *Suppose the one-way estimator yields a change in effective search across two windows,* $\Delta\hat{\mu}^{\text{ow}} \neq 0$*, while the two-way diagnostic yields* $\Delta\hat{\mu}^{\text{tw}} \approx 0$ *on the same markets. Then, under Assumption 1, the one-way change is carried by the market-period common component removed by the two-way transformation. Equivalently, the estimated change is not present in the component of price variation that remains after common market-period movements are removed. This conclusion uses part (i), exact removal of* $\{\delta_t^c\}$*, and does not require knowing which moment or shape feature transmits the difference in the one-way likelihood.*

Part (ii) of Proposition 4 does not imply that applying the original likelihood to two-way residuals is correctly specified. It shows that a variance-based diagnostic can recover the latent dispersion after a degrees-of-freedom correction under the additive common-component DGP. There are, however, three further caveats. The first is related to shape, not just scale. $\tilde{u}_{jt}$ is distributed as a draw of u^* minus an own-inclusive mean of N and T draws; within a period the N residuals are negatively correlated by construction, and their marginal distribution is not in the $L(\cdot\,;\mu)$ family. MLE applied to $\tilde{u}$ as if it were iid from L is therefore misspecified. The moment-based statistic in (ii), or simulated MLE on the transformed family, is the internally consistent implementation. In an empirical application, the two-way likelihood estimates should therefore be read as a diagnostic comparison rather than as a new correctly specified structural estimator.

The second caveat is related to small N. At $N = 3$ the within-period mean absorbs one third of idiosyncratic variance, and degenerate residual variation can make the likelihood flat or undefined, thus generating estimation failures in three-seller markets. Failures are then selected on synchronisation and must be reported, not discarded silently.

The third caveat is what we call conservatism. If true mixed strategies contained a genuinely common market-wide component, the two-way transformation removes it too. The gap between one-way and two-way exercises therefore measures the empirical importance of common market-period dynamics. This is why our result is a diagnostic: when the model is forced to use price variation net of market-period common movements, does the estimated change in model-implied search survive?

4.3 Heterogeneous and lagged responses

Assumption 1 imposes a homogeneous, contemporaneous common component. If instead sellers respond to common shocks heterogeneously (different pass-through rates β_j) or with seller-specific lags, a common shock generates genuine cross-seller dispersion within a period, which temporary price differences a searching consumer could exploit. Two-way demeaning removes the common part and retains residual cross-seller dispersion, which is closer to the price differences consumers can exploit through search, but such dispersion reflects common-shock dynamics interacting with conduct rather than stationary mixed strategies, and search-theoretic models of asymmetric pass-through (Tappata 2009; Lewis 2011) imply it is itself equilibrium-linked to search. The static model cannot represent within-window dynamics of this kind. The one-way and two-way implementations should therefore be interpreted as sensitivity diagnostics rather than as formal bounds on the behavioural object of interest.

5 Numerical illustration

This Section provides a numerical benchmark for the closed-form matching results above. Our exercise asks a simple question: if true search is unchanged, how large a spurious change would be implied by estimators that match the contaminated pooled support or variance?

As our empirical application in the next Section uses data from a retail diesel sector, we use a similar setting in this Section. Specifically, we use the two-type model with $N = 3$ sellers, $T = 6$ periods, margin $x = 10$ cents/litre, and true exhaustive-search share $\mu_N^0 = 0.05$. This is the market size and window length that matter most in the empirical application, and the implied equilibrium dispersion is modest ($\sigma_L = 0.41$ c/L). Search is held fixed across the two windows. The only difference is the common component: the pre window has a common retail trend of 1.0 c/L per period, while the post window has a flatter trend of 0.3 c/L per period.

By Corollary 1, these paths imply $\sigma_{\delta,\text{pre}} = 1.71$ c/L and $\sigma_{\delta,\text{post}} = 0.51$ c/L. Because the pre window has more common within-window movement, the one-way pooled distribution is wider and more dispersed before the event even though the true search distribution is identical in the two periods. Table 1 reports the finite-T variance-matching estimates from (19) and the width-matching estimates based on Proposition 1.

Table 1. Closed-Form Matching Illustration

Period	True μ_N	$\hat{\mu}_N$ (variance matching)	$\hat{\mu}_N$ (width matching)
Pre	0.050	0.282	0.368
Post	0.050	0.080	0.118
Post - pre	0.000	-0.202	-0.250

Note: The table reports the two-type closed-form matching illustration with $N = 3$, $T = 6$, $x = 10$, and true $\mu_N = 0.05$. The pre and post common trends are 1.0 and 0.3 cents per litre per period. The entries are matching-estimator benchmarks, not full-likelihood predictions.

We focus on the sign and the mechanism of these results. A calmer post-window common path is read by these matching estimators as a fall in exhaustive search, even though true search is constant. The variance benchmark moves from 0.282 to 0.080; the width benchmark moves from 0.368 to 0.118. These numbers show that the common component left by one-way demeaning can be large enough, in empirically plausible windows, to generate the qualitative pattern of a pre-period search spike followed by a post-period collapse.

5.1 Monte Carlo diagnostic

The analytical results above describe how common price components enter the pooled demeaned distribution. We complement them with a small Monte Carlo diagnostic. The exercise is designed to check how the estimator behaves in clean simulated equilibrium panels with the short time dimension.[5]

In this exercise, we simulate markets with different number of sellers. For each market size $N = 3, \ldots, 8$, we simulate $T = 6$ periods of equilibrium prices from the Wildenbeest (2011) model. We use a fixed search distribution with $\mu_1 = 0.70$, $\mu_N = 0.10$, and the remaining mass on μ_2. We then apply the same seller-demeaning and likelihood estimator used in the empirical analysis of Section 6. The reported statistic is the Monte Carlo mean of $\widehat{E[k]/N} - E[k]/N$ across 200 replications.[6]

This design differs from the two-type closed-form illustration because here the goal is to evaluate the implemented estimator across the range of market sizes observed in our data. Holding the captive and exhaustive-search shares fixed while placing the remaining mass on two-search consumers gives a stable comparison across N.

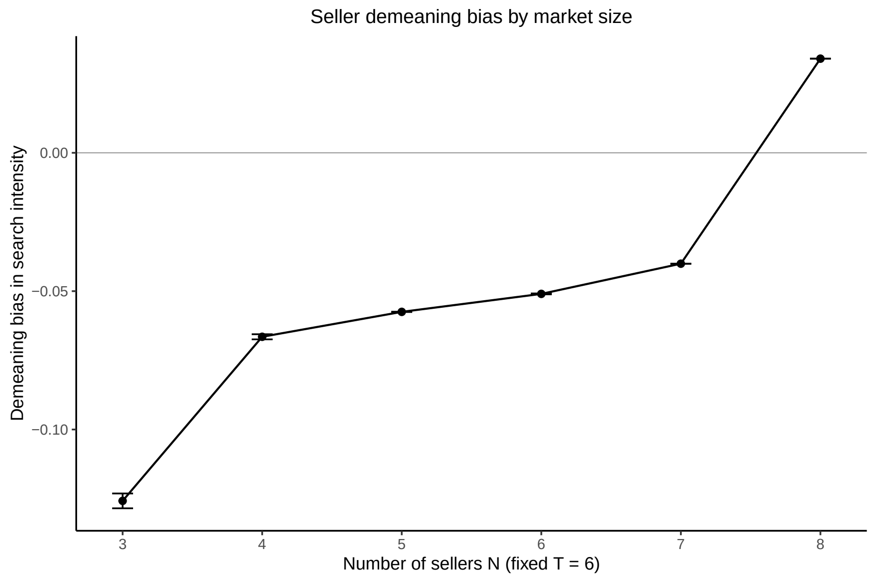

Figure 1. Monte Carlo Demeaning Bias by Market Size

Figure 1 shows that the finite-sample demeaning bias is largest in the smallest markets, which are also the markets that dominate the empirical sample. The bias becomes much smaller as market size increases, although the endpoint $N = 8$ is not monotone in this calibration. We therefore interpret the exercise as evidence that short-panel demeaning can be quantitatively relevant in small isolated markets, not as a monotonic comparative static in N.

[5] Which is the case used in the application in Section 6.

[6] Given the very small standard errors we obtain, we did not increase the number of replications.

The Monte Carlo does not simulate the common price component studied in the main theory. In clean equilibrium simulations, additive common paths are largely absorbed by the estimated support and markup in the full likelihood and do not mechanically reproduce the empirical pre/post collapse. For that reason, the common-dynamics claim rests on Proposition 4 and on the empirical same-sample comparison between one-way and two-way estimates.

6 Empirical Application

6.1 Institutional Setting and Data

In the past decade, the Spanish Competition Authority (Comisión Nacional de los Mercados y Competencia, CNMC) has published several reports highlighting insufficient competition in the fuel distribution sector as well as retail prices that are higher than those in neighbouring countries (González and Moral 2019). Due to worrying evidence laid down by CNC (2012) and CNE (2012), the CNMC launched an investigation into suspected anticompetitive behaviour in the fuel retail sector in May 2013. The investigation found multiple infringements of competition law, and as a result a first round of charges were publicly announced on 20 February 2015. The CNMC resolution charged retailers for coordination relating to prices, customers and commercial conditions, along with exchanging commercially sensitive information. On 25 February, the case attracted extensive media coverage, and it is the date we use in our empirical application (see footnote 1 for examples of media coverage).

We use station-level diesel price data assembled by González and Moral (2019) and publicly available in González and Moral (2020). The dataset includes daily diesel prices for stations across mainland Spain, which we aggregate to weekly level. We focus on a 12-week window around the public announcement of antitrust sanction, i.e. 6 weeks before and 6 weeks after it.

In the empirical model (see Section 2), we need to define markets where retailers compete. Following Nishida and Remer (2018*a*,*b*), we define local markets geographically. We construct isolated markets consisting of stations within a two-mile radius that do not overlap with other markets. This yields a set of small, local markets where competition and consumer search are plausibly contained.

The final constraint in our sample is to include isolated markets with between three and eight stations. This lower bound is necessary to be able to have price dispersion, and then estimate the model. The upper bound is in line with the literature, and puts limits on how we define isolated markets; if we were considering markets with more stations we would be sampling much bigger municipalities. We end up with 187 local markets, and most of them contain between 3 and 5 stations.

6.2 Measuring Search Behaviour

We estimate μ_k separately for the six weeks before and after the announcement date, for each isolated market, and obtain the implied search cost distributions. We then summarise the search behaviour using the expected number of searches, which we interpret as a measure of search intensity, as follows:

$$\mathbb{E}(k) = \sum_k k \cdot \mu_k. \tag{24}$$

6.3 Baseline estimation

Table 2 shows the results of the estimation, for each market size, before and after the sanction using one-way demeaning. The estimates indicate a sharp reduction in consumer search intensity. Across all markets, the estimated fraction of consumers who search all stations, μ_N, falls from 0.349 before the announcement to 0.010 afterwards. The expected fraction of stations searched, $E[k]/N$, falls 27.6 p.p. (from 0.592 to 0.315).

The pre-announcement estimate of μ_N is substantially higher than in the literature (e.g. Nishida and Remer 2018*a*), while the post-announcement estimate is more in line with prior work.

The last two columns of Table 2 present statistics of model fit: the integrated absolute difference (IAD) between the empirical and fitted price CDFs, normalised by the empirical price range in the corresponding market-period. In the pre-announcement period, the average absolute distance between the empirical and fitted CDFs is about 5.3 percentage points over the observed price support, whilst in the post-announcement period it is 5.9 p.p. These values suggest that the model provides a reasonable approximation to the observed price distributions in both periods. Standard fit diagnostics, therefore, do not reveal the source of the identifying problem.

Table 2. Baseline Search Model Estimates by Market Size

Market size	Markets	μ_N pre	μ_N post	$E[k]/N$ pre	$E[k]/N$ post	$\Delta E[k]/N$	Share lower post	IAD/range pre	IAD/range post
All	187	0.349	0.010	0.592	0.315	-0.276	0.936	0.053	0.059
3	117	0.360	0.003	0.603	0.342	-0.260	0.915	0.059	0.064
4	49	0.374	0.011	0.584	0.280	-0.305	0.980	0.045	0.056
5	14	0.238	0.038	0.545	0.259	-0.286	0.929	0.038	0.035
6	3	0.479	0.096	0.679	0.282	-0.397	1.000	0.029	0.042
7	2	0.000	0.056	0.614	0.210	-0.404	1.000	0.029	0.025
8	2	0.023	0.033	0.291	0.169	-0.122	1.000	0.049	0.041

Note: μ_N is the estimated fraction of consumers who search all stations in the market. $E[k]/N$ is the expected fraction of stations searched. IAD/range normalizes the integrated absolute difference between empirical and fitted price CDFs by the empirical price range.

6.4 What moves the baseline estimator

The baseline estimates raise an immediate question: what did change in the price data to make the estimator read so much less search after the announcement? Table 3 shows that the pooled, one-way-demeaned posted-price distribution compresses sharply across the event.

Table 3. Paired Pre/Post Tests for Market-Level Price Moments

Outcome	Pre mean	Post mean	Post - pre	SE	t-stat.	p-value	95 pct. CI	Share lower
Mean price	117.444	126.868	9.425	0.079	119.189	< 0.001	[9.269, 9.581]	0.000
Within-station SD	4.821	1.592	-3.228	0.081	-40.085	< 0.001	[-3.387, -3.070]	0.995
Within-station IQR	7.770	1.433	-6.336	0.128	-49.603	< 0.001	[-6.588, -6.085]	0.995
Within-station p90-p10	10.001	2.969	-7.032	0.159	-44.097	< 0.001	[-7.346, -6.717]	0.995
Abs. weekly price change	2.304	1.367	-0.937	0.038	-24.878	< 0.001	[-1.011, -0.863]	0.964
Price-change frequency	0.944	0.892	-0.052	0.009	-5.902	< 0.001	[-0.070, -0.035]	0.386
Rank reversal rate	0.195	0.193	-0.002	0.010	-0.211	0.833	[-0.022, 0.018]	0.406

Note: The unit of observation is an isolated market. Tests are paired by market and use the six-week windows on either side of the public announcement. The null hypothesis is that the average market-level post-minus-pre change is zero.

The within-station standard deviation falls by about 3.23 cents and the p90-p10 spread by about 7.03 cents. Taken at face value, this is exactly the compression the search model maps into a fall in search intensity. A narrower pooled distribution relative to its support is rationalised by fewer consumers comparing many stations.

This compression, however, is almost entirely a common, market-wide movement rather than a change in the price differences consumers actually face. When we remove week fixed effects in addition to station effects, i.e. the two-way transformation of Section 4.2, Table 4 shows the compression nearly disappears.

Table 4. Pre/Post Changes in Week-FE Residual Price Moments

Outcome	Pre mean	Post mean	Post − pre	SE	t-stat.	p-value	Share lower
Week-FE residual price	-0.021	-0.022	-0.001	0.078	-0.009	0.993	0.492
Within-station SD, week-FE residual	1.227	1.082	-0.145	0.043	-3.371	< 0.001	0.538
Within-station IQR, week-FE residual	1.475	1.120	-0.355	0.068	-5.236	< 0.001	0.614
Within-station p90-p10, week-FE residual	2.483	2.119	-0.364	0.090	-4.048	< 0.001	0.584
Abs. weekly change, week-FE residual	1.003	0.987	-0.016	0.032	-0.505	0.614	0.497

Note: Paired market-level tests. Week-FE residuals are residuals from a regression of posted prices on station fixed effects and event-week fixed effects.

The residual within-station standard deviation falls by only about 0.15 cents (against 3.23 in the raw series) and the residual p90-p10 by about 0.36 cents (against 7.03). In other words, almost all of the post-announcement narrowing in the pooled distribution is accounted for by a common component that shifts every station in a market together week by week. Once that component is removed, the cross-station dispersion is essentially unchanged across the event.

The distinction matters because consumer search acts on the price differences within a market at a point in time, not on the common level. Table 5 reports the search-relevant gaps directly: the dispersion of prices across stations within a market-week, and the saving a consumer obtains by inspecting an additional station.

Table 5. Search-Relevant Price Dispersion Before and After the Announcement

Outcome	Pre	Post	Post − pre	SE	t-stat.	p-value	Share lower
Raw within-week SD	2.366	2.430	0.064	0.057	1.131	0.259	0.487
Raw within-week p90-p10	3.698	3.799	0.102	0.090	1.135	0.258	0.492
Raw max-min spread	4.667	4.810	0.143	0.109	1.314	0.190	0.477
Raw mean saving to cheapest	2.679	2.751	0.072	0.061	1.183	0.238	0.472
Raw cheapest-second gap	2.747	2.794	0.047	0.083	0.563	0.574	0.508
Station/global-week FE residual SD	1.159	1.181	0.021	0.045	0.480	0.632	0.426
Station/global-week FE residual p90-p10	1.847	1.907	0.059	0.070	0.848	0.397	0.431
Station/global-week FE residual range	2.354	2.415	0.061	0.090	0.675	0.501	0.421
Station/market-week FE residual SD	1.159	1.181	0.021	0.045	0.480	0.632	0.426
Station/market-week FE residual p90-p10	1.847	1.907	0.059	0.070	0.848	0.397	0.431
Station/market-week FE residual range	2.354	2.415	0.061	0.090	0.675	0.501	0.421

Note: The unit of observation is an isolated market. Cross-sectional moments are first computed within market-week and then averaged within market-period. Residual moments use fixed-effect residuals as indicated in the row labels.

These quantities do not decline after the announcement. The price differences a searching consumer could have exploited are the same before and after the event, even though the one-way estimator reports a collapse in search. The collapse is thus located in the common component that one-way demeaning leaves in the pooled data, exactly as Proposition 4 and Corollary 2 predict, and not in any change in the price variation that consumer search responds to.

6.5 Search estimates under two-way demeaning

Having shown that the post-announcement compression is a common, market-wide movement, we re-estimate the search model on the two-way demeaned data. To make the comparison clean, we restrict attention to the 156 isolated markets for which both the one-way and two-way estimators converge, so that any difference reflects the transformation rather than a change in sample.[7]

Table 6. Baseline and Market-Week-FE Diagnostic Search Estimates

Specification	Markets	Pre	Post	Change	Median change	Share lower	p-value
Baseline posted prices: $E[k]/N$	156	0.592	0.308	-0.283	-0.295	0.955	< 0.001
Market-week-FE diagnostic: $E[k]/N$	156	0.347	0.338	-0.009	-0.000	0.526	0.282
Baseline posted prices: μ_N	156	0.347	0.011	-0.336	-0.375	0.840	< 0.001
Market-week-FE diagnostic: μ_N	156	0.038	0.020	-0.018	-0.000	0.545	0.041

Note: The sample is restricted to markets successfully estimated in both specifications. $E[k]/N$ is the expected fraction of stations searched. μ_N is the estimated share of consumers who sample all stations. p-values are from paired market-level tests of equality between post- and pre-period estimates.

Table 6 reports the two implementations side by side on this common sample. Under one-way demeaning the expected fraction of stations searched falls by 28.3 p.p. across the announcement (from 0.592 to 0.308), with the decline appearing in 95.5% of markets. This movement is driven by an implausibly high pre-period exhaustive-search share, $\mu_N = 0.347$, which collapses to 0.011 afterwards. Under two-way demeaning the same comparison yields a fall of only 0.9 p.p. (from 0.347 to 0.338), present in just 52.6% of markets, and a paired test does not reject equality of the two periods. The pre-period exhaustive-search share is $\mu_N = 0.038$ under two-way

[7] Of the 187 markets estimated under one-way demeaning, 156 also converge under the two-way transformation. The 31 that fail are all three-station markets, where removing both station and week means leaves too little residual variation for the likelihood to fit. This restriction does not select on the outcome. The dropped markets have one-way profiles close to those of the retained markets with a mean pre-period μ_N of 0.359 against 0.347, an identical mean pre-period $E[k]/N$ of 0.592, and a mean one-way decline in $E[k]/N$ of 0.241 against 0.283. If anything, their one-way decline is smaller, so restricting to the common sample works against the divergence we report rather than towards it.

demeaning, an order of magnitude below the one-way figure and in line with the prior literature (Nishida and Remer 2018*a*).

The divergence between the two columns is the empirical counterpart of the diagnostic. A researcher running the standard one-way implementation would report a 28-point collapse in search around the antitrust announcement and an implausibly high pre-period exhaustive-search share. The two-way comparison reveals that neither survives once the common market-week component is removed. The gap between the implementations is large in exactly the setting our analysis identifies as dangerous, i.e. a short window fixed by an event, in a market with volatile common costs, and it is the size of that gap, not either estimate on its own, that signals the contamination.

7 Conclusion

Structural non-sequential search models recover consumer search intensity from the dispersion of prices within a market. When the data are pooled over time, as they must be in the small markets to which these models are typically applied, the standard one-way demeaning removes seller heterogeneity but leaves any common, time-varying price component in the pooled distribution. In this paper, we have shown analytically how this component enters the object fitted by the econometrician, and in tractable matching cases how its support width or variance can be mapped into higher search. Two-way demeaning removes the additive common component exactly. The gap between one-way and two-way estimates is therefore a diagnostic: it tells us whether the estimated change survives after market-period common movements are removed.

This distinction has been treated as settled in the literature, where the two-way transformation is run as a robustness check and found to make little difference. We have argued that its innocuousness is conditional on a small common component, a condition met in the stable or carefully chosen windows of existing applications but generically violated when the estimation window is fixed by the timing of an event. In the Spanish diesel market around the 2015 antitrust sanction, and a period with great volatility in the international oil market, the one-way estimator implies a sharp collapse in search. The two-way comparison shows that the collapse, and the implausibly high pre-period search level that accompanies it, are not present in the component of price variation left after common market-week movements are removed. The search-relevant price gaps consumers can exploit are essentially unchanged across the antitrust event.

The implication for applied work is practical. Before reading an estimated change in search across an event as behaviour, the one-way and two-way estimates should be compared, and their divergence interpreted as evidence that the result is carried by common price dynamics rather than by search-relevant relative prices. The transformation is no harder to implement than the estimator itself, and in before-and-after applications around events it can separate a robust behavioural conclusion from a result that disappears once market-period common movements are removed.

References

Allen, J., Clark, R. and Houde, J.-F. (2014), 'The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry', *American Economic Review* **104**(10), 3365–3396.

Araujo, J. P. and Rodrigues, M. (2020), 'Evidence on search costs under hyperinflaton in Brazil: The effect of Plano Real', *Brazilian Review of Econometrics* **40**(1), 75–110.

Bacon, R. W. (1991), 'Rockets and feathers: The asymmetric speed of adjustment of UK retail gasoline prices to cost changes', *Energy Economics* **13**(3), 211–218.

Borenstein, S., Cameron, A. C. and Gilbert, R. (1997), 'Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?', *The Quarterly Journal of Economics* **112**(1), 305–339.

Burdett, K. and Judd, K. L. (1983), 'Equilibrium Price Dispersion', *Econometrica* **51**(4), 955–969.

CNC (2012), 'Informe de seguimiento del mercado de distribución de carburantes de automoción en españa'. Comisión Nacional de la Competencia.

CNE (2012), 'El mercado español de la distribución de la gasolina y gasoleo a través del canal de estaciones de servicio'. Comisión Nacional de la Energía.

González, X. and Moral, M. J. (2019), 'Effects of antitrust prosecution on retail fuel prices', *International Journal of Industrial Organization* **67**, 102537.

González, X. and Moral, MaríaJ. (2020), 'Panel data to investigate pricing behavior in the Spanish retail fuel market', *Data in Brief* **28**, 104880.

Hong, H. and Shum, M. (2006), 'Using price distributions to estimate search costs', *The RAND Journal of Economics* **37**(2), 257–275.

International Energy Agency (2015), 'Oil market report', https://iea.blob.core.windows.net/assets/c8c659da-ec0f-43ae-b4ef-a9105ad238d6/OMR_2015.pdf.

Lewis, M. S. (2011), 'Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market', *Journal of Economics & Management Strategy* **20**(2), 409–449.

Luco, F. (2019), 'Who Benefits from Information Disclosure? The Case of Retail Gasoline', *American Economic Journal: Microeconomics* **11**(2), 277–305.

Moraga-González, J. L. and Wildenbeest, M. R. (2008), 'Maximum likelihood estimation of search costs', *European Economic Review* **52**(5), 820–848.

Nishida, M. and Remer, M. (2018*a*), 'The Determinants and Consequences of Search Cost Heterogeneity: Evidence from Local Gasoline Markets', *Journal of Marketing Research* **55**(3), 305–320.

Nishida, M. and Remer, M. (2018*b*), 'Lowering consumer search costs can lead to higher prices', *Economics Letters* **162**, 1–4.

Tappata, M. (2009), 'Rockets and Feathers: Understanding Asymmetric Pricing', *The RAND Journal of Economics* **40**(4), 673–687.

Wildenbeest, M. R. (2011), 'An empirical model of search with vertically differentiated products', *The RAND Journal of Economics* **42**(4), 729–757.

*

Appendix A. Proofs

Proof of Lemma 1. The standard equilibrium has an atomless utility distribution on an interval $[0, \bar{u}]$. Because sellers mix on this support, expected profit must be constant at every utility in the support.

At the lower bound $u = 0$, $L(0) = 0$. A seller offering the lowest utility can sell only to consumers who sample one seller, so every term in (2) with $k \geq 2$ drops out and

$$\pi(0) = x\frac{\mu_1}{N}.$$

At the upper bound $u = \bar{u}$, $L(\bar{u}) = 1$. The seller wins whenever it is sampled, so

$$\pi(\bar{u}) = (x - \bar{u})\frac{1}{N}\sum_{k=1}^{N} k\mu_k = (x - \bar{u})\frac{\mathbb{E}[k]}{N}.$$

Equating endpoint profits gives

$$(x - \bar{u})\,\frac{\mathbb{E}[k]}{N} = x\,\frac{\mu_1}{N} \quad\Longrightarrow\quad \bar{u} = x\left(1 - \frac{\mu_1}{\mathbb{E}[k]}\right).$$

The monotonicity statement follows directly from this expression. The interval support and atomlessness are standard properties of this class of equilibria (see Burdett and Judd 1983; Wildenbeest 2011). □

Proof of Proposition 1. Substituting (13) into (3):

$$\bar{p}_j = \xi_j + \bar{\delta}^c + \rho - \bar{u}_j^*,$$

so

$$\hat{u}_{jt} = \bar{p}_j - p_{jt} = (u^*_{jt} - \bar{u}^*_j) - (\delta^c_t - \bar{\delta}^c) = (u^*_{jt} - \bar{u}^*_j) + \delta_t,$$

which is (15).

(i) For each seller, $u^*_{jt} - \bar{u}^*_j$ is a deviation from an own-inclusive sample mean of T iid draws, so its variance is $\sigma^2_L(1 - 1/T)$. Conditional on the common path, δ_t has in-window mean zero and second moment σ^2_δ. The idiosyncratic equilibrium draws are independent of the common path, so the cross-product has mean zero and vanishes in the pooled average under the repeated-sampling sequence stated in the proposition.

The factor $(T-1)/T$ therefore applies only to the idiosyncratic term. The common deviations δ_t are already expressed relative to their in-window mean and enter with unit weight.

(ii) As $T \to \infty$, the seller-specific sample mean $\bar{u}^*_j$ converges to the mean of L, so the first term in (15) has the same support width as L. For independent random variables with supports $[a_1, b_1]$ and $[a_2, b_2]$, the support of the sum is $[a_1 + a_2, b_1 + b_2]$ (the Minkowski sum). Hence the support width of the pooled object is the width of L plus the range of the common component, $\bar{u} + R_\delta$. □

Proof of Proposition 2. (i) By (17), $W = \bar{u}(\mu^0) + R_\delta = x(1 - \mu^0_1/\mathbb{E}^0[k]) + R_\delta$ (Lemma 1). Width matching imposes $x(1 - \widehat{\mu_1/\mathbb{E}[k]}) = W$. Solving gives (18), which is admissible precisely when $W \leq x$.

(ii) By (16), $g_N(\hat{\mu}_N) = \mathrm{Var}(\hat{u}) = \frac{T-1}{T} g_N(\mu^0_N) + \sigma^2_\delta$, which is (19). The stated condition $\sigma^2_\delta > \frac{1}{T} g_N(\mu^0_N)$ is exactly what makes the argument exceed $g_N(\mu^0_N)$, hence $\hat{\mu}_N > \mu^0_N$. (An estimator that instead applies the finite-T factor to its own implied variance solves $\frac{T-1}{T} g_N(\hat{\mu}_N) = \frac{T-1}{T} g_N(\mu^0_N) + \sigma^2_\delta$ and is biased upwards for any $\sigma^2_\delta > 0$; (19) is the conservative case.) On the branch the map is a strictly increasing bijection. □

Proof of Proposition 3. Immediate from Proposition 2 and strict monotonicity of g_N^{-1}

on the branch. $\square$

Proof of Corollary 1. With $\delta_t = -b_w(t - \bar{t})$ and $\bar{t} = (T+1)/2$, $\sigma_\delta^2 = b_w^2 \cdot \frac{1}{T}\sum_t (t-\bar{t})^2 = b_w^2(T^2-1)/12$, and the range of $t - \bar{t}$ is $T-1$. $\square$

Proof of Proposition 4. (i) $\delta_t^c + \xi_j$ is additively separable in j and t; the two-way transformation annihilates any function of the form $a_j + b_t$.

(ii) For x_{jt} iid with variance σ^2, the two-way demeaned residual has second moment $\sigma^2(1-1/N-1/T+1/(NT)) = \sigma^2(1-1/N)(1-1/T)$; apply to $x_{jt} = -u^*_{jt}$. Consistency of the corrected variance statistic follows under the same repeated-sampling interpretation used above. $\square$

Proof of Corollary 2. By part (i), $\tilde{u}_{jt}$ is invariant to $\{\delta_t^c\}$ and equals the two-way residual of $-u^*_{jt}$. Under Assumption 1, the only component treated differently by one-way and two-way demeaning is the market-period common term. Therefore, if the one-way transformation produces a pre/post change that disappears after two-way residualisation on the same markets, the difference between the two exercises is carried by $\{\delta_t^c\}$. The statement is diagnostic rather than a claim that the original likelihood is correctly specified on two-way residuals. $\square$