UNIVERSITÉ DU QUÉBEC EN OUTAOUAIS

DÉPARTEMENT DES SCIENCES ADMINISTRATIVES

WORKING PAPER NO. 36

The Hidden Dynamics of Used-Car Markets: Price Discovery, Inventory Liquidity, and Strategic Dealer Behavior

Evidence from a Daily Panel of 273,265 Quebec Listings

Simon-Pierre Boucher

Département des sciences administratives
Université du Québec en Outaouais

simon-pierre.boucher@uqo.ca

Gatineau – Pavillon Alexandre-Taché
283, boulevard Alexandre-Taché
Gatineau, Québec, Canada J9A 1L8

November 2026

Version 1.0

Abstract

We study price discovery and liquidity in the Quebec used-car market with a daily panel of 273,265 dealer listings observed by the aggregator Auto-Ka between 12 August and 1 November 2026, with every asking-price change and listing disappearance recorded about seven times a day. Hedonic models with flexible age and mileage splines reject linear depreciation and improve dealer-held-out fit by three percent; profiles differ by powertrain. One listing in five receives a sustained markdown (median 2.8 percent of the entry price); the markdown hazard doubles over the first three weeks on the lot, and vehicles that entered above their quality-adjusted value or above their local competitors are cut sooner and more deeply. A competitor event study with market-by-week and model-year-by-week fixed effects finds a same-day co-movement of price cuts among close local competitors (0.9 percentage points against a daily base rate of 1.4 percent) but no lagged response, and a within-make–model placebo yields nothing: we interpret this as common shocks and synchronised repricing calendars, not strategic reaction. Quality-adjusted regional premia across Quebec's seventeen regions are economically small (all within four percent of Montréal) and jointly insignificant once trim is held fixed ($p = 0.53$), although coarser quality controls detect small but significant differences. Listing exit is strongly predicted by mileage, price and listing age; cross-source matching by vehicle identification number shows that 94 percent of exits coincide within a week on two independent sources while at least 10 percent are followed by a relisting, so exit overstates sales. A transparent five-component Automobile Market Liquidity Index captures 56 percent of the common variation in regional turnover and, built without the exit component, predicts next-week exit intensity within regions under z-score standardisation, a result that weakens under rank standardisation; deep markdowns load as a symptom of illiquidity. No external forecasting claim is possible with an eleven-week window, and the paper documents what listing data can and cannot say about transactions.

Keywords: used cars, asking prices, hedonic pricing, markdown hazard, listing survival, dealer competition, market liquidity index, Quebec

JEL Classification: D22, D83, L11, L62, R32, C41

1. Introduction

Used cars are the largest durable good that most households trade, and the market that reallocates them is organised almost entirely through posted asking prices. A dealer lists a vehicle at a price, waits, cuts the price if the vehicle does not move, and eventually the listing disappears, either because the car was sold, wholesaled to another dealer, or simply withdrawn. Economists have studied the resulting equilibrium with transaction data from auctions and registration records (Genesove, 1993; Gavazza et al., 2014; Larsen, 2021), but the price-discovery process itself, the sequence of asking prices and the decision to cut them in the face of local competition, has remained largely hidden, because no public source records the daily asking-price paths of a whole market. This paper opens that black box for the Canadian province of Quebec.

We assemble a daily panel of 273,265 used-vehicle listings observed between 12 August and 1 November 2026 by Auto-Ka, an aggregator that re-reads the inventory pages of 131 dealer websites, two full-inventory portals and several dealer groups roughly seven times a day, logs every asking-price change, and records when a listing disappears from its source. The data are asking prices, not transaction prices, and listing disappearance is not a verified sale. We take these limits seriously throughout: every dynamic outcome is defined as an observed event on the dealer's website, and the paper's central methodological contribution is to show what can and cannot be learned about market liquidity from such listing data alone.

The empirical programme was preregistered before any model was estimated (Section 5; the full plan is in the replication package) and has six parts. First, we estimate hedonic asking-price models with flexible age and mileage profiles and ask whether depreciation is nonlinear and powertrain-specific (H1). Second, we model the hazard of a first sustained markdown and the size of the cut as functions of listing age, local competition and the vehicle's relative price (H2). Third, we estimate an event study of a dealer's repricing around the price cuts of close local competitors, with market-by-week and model-year-by-week fixed effects, treating the estimates explicitly as associations (H3). Fourth, we test whether quality-adjusted asking prices differ across Quebec's seventeen administrative regions (H4). Fifth, we estimate the hazard of listing exit and validate exit as a sale proxy using vehicles observed simultaneously on a dealer's own site and on a portal, and by detecting relistings of the same vehicle identification number elsewhere (H5). Sixth, we construct an Automobile Market Liquidity Index (AMLI) from five transparent components and validate it internally with a held-out-component design (H6).

Five findings stand out. (i) Depreciation of asking prices is strongly nonlinear in both age and mileage: the linear specification is rejected at any conventional level and splines improve dealer-grouped out-of-sample fit by about 3 percent, a gain that gradient boosting only modestly extends. Age profiles differ by powertrain, with battery-electric vehicles depreciating faster than comparable

combustion vehicles at young ages. (ii) Markdowns are frequent but shallow: one listing in five receives at least one sustained cut, the median first cut is 2.8 percent of the entry price, and the daily markdown hazard doubles over the first three weeks of a listing before plateauing. Listings that enter above their hedonic value, or above their local competitors, are cut sooner and more deeply; the number of local competitors has no stable effect. (iii) Repricing is synchronised rather than sequential: a close competitor's cut is associated with a 0.5–0.9 percentage-point higher probability that the dealer cuts the same day, against a daily base rate of 1.4 percent, but the association vanishes on the following days and reappears only at seven-day harmonics; it survives dealer-by-weekday fixed effects at $k = 0$ but not as a lagged response. A placebo that permutes competitor series within make–model cells produces no effect. We read this as common local shocks and shared repricing calendars rather than strategic reaction, and we do not claim causality. (iv) Quality-adjusted asking prices are close to integrated across Quebec: the sixteen regional premia relative to Montréal all lie within ± 4 percent, and with make–model–trim fixed effects and dealer-clustered inference they are jointly insignificant ($p = 0.53$); with coarser make–model controls the joint test rejects ($p = 0.004$) but the premia remain within ± 3 percent, so the economically relevant conclusion is that regional differences are small. (v) Listing exit is strongly predicted by mileage, price and listing age, and exit is a reasonably consistent event across sources: for vehicles observed both on a dealer's site and on a portal, 80 percent of exits coincide within one day and 94 percent within a week; 10 percent of exits are followed by a relisting of the same vehicle within 30 days, 6 percent at a different location, which bounds the share of non-sale exits from below. The AMLI's first principal component explains 56 percent of component variance, and an index built without the exit component predicts next-week exit intensity within regions, although this internal validation is sensitive to the standardisation of the components. We are explicit that the 81-day window does not allow any test of the index's forecasting value for external economic variables.

The paper contributes to three literatures. To the economics of durable-goods markets (Akerlof, 1970; Hendel and Lizzeri, 1999; Stolyarov, 2002; Gavazza et al., 2014), it adds listing-level evidence on how dealers search for the price at which a used car clears, the counterpart on the sell side to the buyer-search evidence of Murry and Zhou (2020) and the intermediary margins of Biglaiser et al. (2020). To the literature on asking prices and time on market, developed mostly for housing (Knight, 2002; Han and Strange, 2016; Anenberg, 2016; Guren, 2018), it brings a market with far shorter spells, near-complete price histories and many close substitutes. To the measurement of market liquidity, where industry days'-supply and days-to-turn statistics (Cox Automotive, 2022; Canadian Black Book, 2026) exist but are opaque, it offers a reproducible, component-based index with documented validation limits.

The remainder of the paper is organised as follows. Section 2 describes the Quebec used-car market and the data-generating process of the aggregator. Section 3 reviews the related literature.

Section 4 documents the data, the quality controls and the sample construction. Section 5 sets out the econometric framework and identification. Sections 6–8 report the results on price formation, dealer dynamics and the liquidity index. Section 9 presents robustness checks and limitations, and Section 10 concludes.

2. Institutional background and data-generating process

2.1. The Quebec used-car market

Quebec is Canada's second-largest provincial vehicle market. New motor vehicle sales in the province were 41,657 units in August 2026 (Statistics Canada table 20-10-0085-01), and the used market is conventionally two to three times larger in units. Used vehicles are sold by franchised dealers, independent used-car dealers, dealer groups operating under a common banner, and private sellers. Almost every dealer publishes its inventory on its own website, typically through one of a handful of website platforms (D2C Media, SM360, AMVOQ's Theme Central, EvalAuto, SyncAuto) that expose structured listings, and most dealers simultaneously syndicate their inventory to portals such as Otogo (operated by LC Media), AutoUsagée (operated by the dealer association AMVOQ), AutoTrader.ca, CarGurus and Kijiji Autos. Prices are posted in Canadian dollars, usually before taxes and fees; Quebec's consumer-protection law requires the advertised price to be the all-in price excluding sales taxes, which limits the scope for hidden fees and makes posted prices comparable across dealers.

Transactions are not public. The provincial registry (SAAQ) publishes annual vehicle-stock files but no ownership-transfer series, and the only Canadian transaction-based used-vehicle index (the Statistics Canada used-vehicle CPI component, built from J.D. Power dealer transaction data since June 2022) is national and monthly (Taylor et al., 2022). Consequently, the asking price and the listing's presence on the dealer's website are the only high-frequency, vehicle-level signals available to a researcher, as they are to a buyer.

2.2. Auto-Ka and the listing panel

Auto-Ka (`https://www.auto-ka.com`) is an independent aggregator of Quebec used-vehicle listings built and operated by the author. It maintains a registry of 143 sources: 131 dealer websites, 5 dealer groups, 2 full-inventory portals (Otogo and AutoUsagée), 2 sampled portals (AutoTrader and CarGurus), the private-seller section of Kijiji and two small portals. A dedicated connector per website platform reads each source directly, normalises each listing into a common schema (make, model, trim, model year, asking price, mileage, transmission, fuel, drivetrain, body type, colours, VIN when published, stock number, dealer name and location), and hashes the content. A

synchronisation cycle re-reads every source; over the sample the cycle took 3.3 hours on average, giving 7.3 observations per listing per day, with no gap longer than 13.6 hours for any source. Each cycle classifies every listing as new, changed or unchanged, appends every asking-price change to a price log, and increments a miss counter for listings that are no longer returned by the source; a listing is declared inactive after two consecutive misses, unless a drift alert (a fall in the source's volume below a quarter of its recent median, or a jump in missing prices) suspends removals for that source. Listings carrying the same valid 17-character VIN on several sources are de-duplicated in favour of the dealer's own site.

The research dataset is a read-only copy of the production database taken on 1 November 2026 after the last synchronisation of the morning, from which the user-account table was removed before transfer. No scraping was performed for this paper and no public API was used; the platform's data-governance rules, the field-level audit, and the restricted status of VINs, URLs, free text and exact coordinates are documented in the replication package.

Three features of this data-generating process matter for the analysis. First, the panel starts on 12 August 2026 for dealer websites, on 18 August for Otogo and on 23 August for AutoUsagée; the stock present on each source's first crawl is *left-censored* (its true listing date is unknown) and is excluded from every duration analysis. Second, a listing's disappearance is an observed event on the source website; it may reflect a retail sale, a wholesale transfer, a withdrawal, or a temporary feed failure. We therefore require seven days of continuous absence before counting an exit, and we validate exits across sources and against relistings in Section 7. Third, two portals are deliberately sampled by the platform (AutoTrader returns a sliding window of at most 2,000 of more than 4,000 listings; CarGurus is capped at 500) and cannot distinguish a listing leaving the market from a listing leaving the sample; they are excluded from all dynamic analyses and from the baseline hedonic sample.

3. Related literature and mechanisms

Hedonic prices and depreciation. Quality-adjusted pricing of automobiles is the founding application of hedonic methods (Court, 1939; Griliches, 1961; Ohta and Griliches, 1976), formalised by Lancaster (1966) and Rosen (1974) and extended to flexible functional forms and identification by Pakes (2003), Ekeland et al. (2004) and Bajari and Benkard (2005); Triplett (2006) and Hill (2012) survey measurement practice. For used vehicles, the price of a given model falls with age and with accumulated mileage, and Engers et al. (2009) show that annual miles, not only odometer readings, drive used-car prices. The Quebec-specific evidence on depreciation is thin (Yanes, 2010). Our hedonic models follow this tradition with asking rather than transaction prices and ask how far the flexible-form conclusions carry over.

Durable-goods markets, adverse selection and liquidity. Since Akerlof (1970), used-car markets have been the canonical setting for asymmetric information (Kim, 1985; Bond, 1982; Genesove, 1993; Hendel and Lizzeri, 1999; Emons and Sheldon, 2009; Peterson and Schneider, 2014), with online evidence from eBay Motors (Lewis, 2011; Adams et al., 2011). Stolyarov (2002), Gavazza (2011) and Gavazza et al. (2014) model trading frictions and the resulting pattern of turnover and prices over the vehicle's life; Larsen (2021) measures bargaining efficiency in wholesale auctions and Biglaiser et al. (2020) the premium that dealers earn as intermediaries. Dealers in these models hold inventory precisely because buyers arrive stochastically; the markdown and exit hazards we estimate are the empirical shadows of that search process.

Dealer pricing and inventory. Retail price cuts are the standard tool for clearing inventory (Lazear, 1986; Pashigian and Bowen, 1991; Pesendorfer, 2002; Aguirregabiria, 1999), and the dynamic-pricing literature characterises optimal markdown paths for perishable or seasonal goods (Gallego and van Ryzin, 1994; Elmaghraby and Keskinocak, 2003; Sweeting, 2012). For cars, Zettelmeyer et al. (2006) and Scott Morton et al. (2001) document how internet referrals compressed dealer margins, Busse et al. (2006) how promotions pass through, and Israeli et al. (2022) how new-car dealers vary prices with inventory. Murry and Zhou (2020) show that used-car dealers co-locate to attract searching buyers. Industry sources report that a substantial share of used listings receive price drops (AutoTrader.ca, 2024; iSeeCars.com, 2013), but to our knowledge no peer-reviewed study estimates markdown hazards or cut sizes on listing-level paths.

Search, dispersion and asking prices. Price dispersion for homogeneous goods is the signature of costly search (Stigler, 1961; Diamond, 1971; Varian, 1980; Burdett and Judd, 1983; Sorensen, 2000; Lach, 2002; Baye et al., 2006; Kaplan and Menzio, 2015). The role of the asking price as a signal and a commitment device has been studied mainly for housing: time on market and list-price cuts (Knight, 2002; Anglin et al., 2003; Genesove and Mayer, 2001; Carrillo, 2012; Han and Strange, 2016), momentum and strategic complementarity in list prices (Guren, 2018), and listing-based price indexes (Anenberg and Laufer, 2017; Lyons, 2019). Liquidity itself is measured with price impact in financial markets (Amihud and Mendelson, 1986; Amihud, 2002), with search-and-matching models in over-the-counter and housing markets (Duffie et al., 2005; Lagos and Rocheteau, 2009; Krainer, 2001; Head et al., 2014; Ngai and Tenreyro, 2014), and with days'-supply or days-to-turn statistics in the car trade (Cox Automotive, 2022; Manheim Consulting , Cox Automotive; Canadian Black Book, 2026). Our index formalises the latter as a transparent multi-component measure.

Survival analysis. We rely on discrete-time hazard models (Allison, 1982; Singer and Willett, 1993), the Cox model and its diagnostics (Cox, 1972; Grambsch and Therneau, 1994; Lin and

Wei, 1989; Therneau and Grambsch, 2000), flexible parametric and competing-risk alternatives (Royston and Parmar, 2002; Fine and Gray, 1999; Putter et al., 2007), and the econometric treatment of duration data with censoring and unobserved heterogeneity (Kiefer, 1988; Lancaster, 1990; Heckman and Singer, 1984).

Spatial competition and regional integration. Localised competition follows Hotelling (1929) and Salop (1979); empirical work measures spatial differentiation in retail gasoline (Houde, 2012; Pinkse et al., 2002; Chandra and Tappata, 2011; Byrne and de Roos, 2019) and hotels (Kalnins, 2003), and regional price integration through border effects (Engel and Rogers, 1996; Gorodnichenko and Tesar, 2009; Ceglowski, 2003); Haan and de Boer (2010) document regional differences in Dutch used-car listing prices. Spatial inference follows Conley (1999) and Anselin (1988).

Identification of strategic interactions. Whether a dealer's price cut *causes* a neighbour's cut is a reflection problem (Manski, 1993; Brock and Durlauf, 2001; Bramoullé et al., 2009; Angrist, 2014; Blume et al., 2011). We use event-study and distributed-lag designs with the diagnostics developed in the recent difference-in-differences literature (Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021; Borusyak et al., 2024; Roth, 2022; Rambachan and Roth, 2023; Roth et al., 2023), cluster-robust inference (Bertrand et al., 2004; Cameron et al., 2008, 2011; Abadie et al., 2023), and multiple-testing control (Benjamini and Hochberg, 1995; Romano and Wolf, 2005); predictive comparisons use Diebold and Mariano (1995), Giacomini and White (2006), Brier (1950) and gradient boosting (Friedman, 2001; Ke et al., 2017; Athey and Imbens, 2019; Mullainathan and Spiess, 2017).

What this paper adds. Relative to this literature the paper offers (i) the first listing-level daily panel of a whole regional used-car market with complete asking-price histories, (ii) hazard-based evidence on markdowns and exits with dealer-clustered inference, (iii) a competitor event study that distinguishes same-day synchronisation from lagged response, (iv) a precisely estimated null on regional premia, and (v) a transparent liquidity index. It does not add transaction prices, and we acknowledge that days'-supply measures already capture part of what the index measures.

4. Data, measurement and descriptive evidence

4.1. Sources and coverage

Table 1 summarises the 143 sources. Over the 81-day window (12 August to 1 November 2026, 586 synchronisation cycles) the platform observed 273,265 distinct listings, of which 62,821 were

Table 1. Data sources, coverage and sample membership

Source class	Sources	Dealers	Listings	Active	Found/cycle	S1	S2
Dealer websites	131	131	33 061	12 222	12 523	Yes	Yes (113 of 131)
Full-inventory portals	2	929	107 903	39 125	36 693	Yes	Yes
Dealer groups	5	5	17 162	5 579	5 131	Yes	Yes (4 of 5)
Sampled portals	2	1 313	103 027	2 654	2 359	No	No
Private sellers (Kijiji)	1	1	11 874	3 158	3 041	No	No
Small portals	2	2	238	83	194	Yes	No
All sources	143	2 381	273 265	62 821	59 941		

Notes: Dealers = distinct dealer identifiers (dealer websites and groups are one dealer each; portals identify the selling dealer by name). Active = active on 2026-11-01. Found/cycle = mean number of listings returned per synchronisation cycle. S1 = hedonic sample; S2 = dynamic sample (Table 2). Sampled portals are capped sliding windows and cannot reveal exits.

active on the snapshot date. Dealer websites contribute 33,061 listings from 131 dealers; the two full-inventory portals contribute 107,903 listings from 929 named dealers, with exact dealer coordinates and VINs for 82 percent of Otogo listings; dealer groups add 17,162. The sampled portals (103,027 listings) and Kijiji private sellers (11,874) are observed but excluded from the analysis samples for the reasons given in Section 2. Figure 1 shows the daily stock and flows of the reliable-source universe, the crawl intensity, and the regional distribution of the active stock. After the two portals come online, the reliable stock is stable at 44,000–49,000 listings, with roughly 1,000 entries and 1,000 confirmed exits per day and a pronounced weekly cycle in both.

4.2. Variables

Asking price. The platform logs the entry price and every subsequent change. Price logs are contaminated by two artefacts that any listing-based study must address: placeholder prices (2,164 listings at $1 or less; 5 above $1 million) and *oscillations*, in which a connector alternately returns two variants of the same listing (3,451 listings exhibit at least one A→B→A reversion, some twice a day for weeks). We therefore work with a *sustained price path*: a price level that lasts less than 48 hours and reverts to the previous level is deleted, listings with three or more such transients are flagged as oscillating and excluded from markdown analyses, and a *markdown* is a sustained decrease of at least 1 percent and $100. Under these rules the price log yields 76,261 cuts and 26,486 rises across all sources; among reliable sources cuts outnumber rises nine to one.

Vehicle age is measured in years from 1 October of the year preceding the model year to the listing's entry date; *mileage* in units of 10,000 km. *Powertrain* (combustion, hybrid, plug-in hybrid, battery-electric, diesel) is read from the fuel field and, when missing, recovered from the title or from the make–model(–trim–year) cell when at least 90 percent of known listings in the cell agree (84,753 recoveries; 11,099 listings remain unknown). *Body type* is normalised to eight classes and recovered from the make–model modal class when missing. *Region* is the administrative region

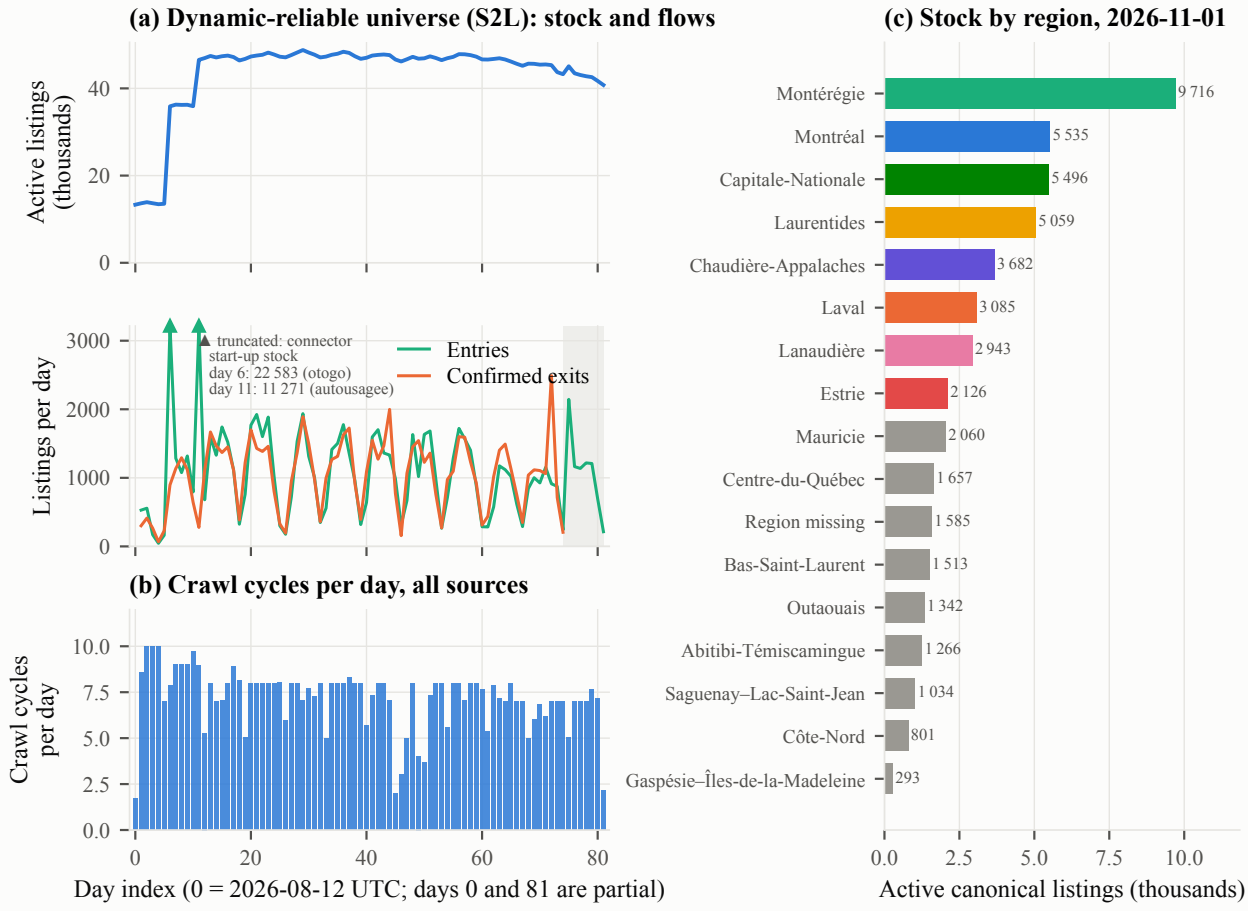


Figure 1. Coverage across time, region and source. (a) Active listings, entries and confirmed exits per day in the reliable-source universe (S2L); the two spikes are the connector start-ups of Otogo (day 6) and AutoUsagée (day 11), whose stock is treated as left-censored. (b) Synchronisation cycles per day. (c) Active canonical listings by administrative region on 1 November 2026.

Table 2. Sample construction and exclusions

Step		Remaining	Dropped
S0 all listings in snapshot		273 265	–
− kind != auto		271 992	1 273
− VIN duplicate (non-canonical)		259 133	12 859
− model year outside [2005,2026] or missing		256 651	2 482
− mileage outside [1000,400000] km or missing		249 958	6 693
− entry price outside [2000,250000] $ or missing/placeholder		241 321	8 637
− make or model missing		241 298	23
− uid identity collision	⇒ S1 (all sources)	241 297	1
− sampled portal (autotrader/cargurus) or private marketplace (kijiji)	⇒ S1 hedonic	136 393	104 904
− flickering/failed connector or small portal		128 144	8 249
− oscillating price path	⇒ S2L	128 091	53
− left-censored stock (day 1 or source start-up day)		80 856	47 235

Notes: Sequential exclusions from the 2026-11-01 snapshot. S1 = hedonic sample (reliable sources); S2L = dynamic sample including the left-censored stock; S2 = dynamic sample used for hazards, competition and the liquidity index.

reported by the source, recovered from the city or from the nearest known city when missing (13,073 recoveries; 12,434 listings, mostly on excluded sources, remain without region). *Dealer* identifiers are the source for dealer websites and the normalised dealer name within portal for portals; coordinates are the listing's, or the dealer's median location when missing.

Listing age is days since entry; *vehicle age* is distinct. *Exit* is defined as the last observation of a listing that is inactive at the snapshot and has been absent for at least seven days; exits within the final seven days of the window are right-censored, as are listings still active. The exit time is placed at the last observation plus half the source's median crawl interval, which treats exit as interval-censored between two crawls.

4.3. Sample construction and quality controls

Table 2 reports the attrition from the raw snapshot to the analysis samples. The hedonic sample S1 (136,393 listings) keeps canonical automobile listings from reliable sources with model years 2005–2026, mileage between 1,000 and 400,000 km, an entry price between $2,000 and $250,000, and non-missing make and model. The dynamic sample S2 (80,856 listings) further excludes the sampled portals and Kijiji, six connectors with documented failures or flicker (up to 94 percent of their exit events were later reversed), 53 oscillating listings, and the 47,235 listings that were already present on the first crawl of their source. The intermediate sample S2L retains the left-censored stock and is used as the competitor universe and for sensitivity analysis. Every excluded listing is written to an auditable exclusion log with the rule that removed it.

Three further quality controls are reported in the replication package: a per-source reactivation share (the fraction of exit events later reversed, which motivated the seven-day confirmation rule and the exclusion of flickering connectors), the stability of listing identifiers within sources (only

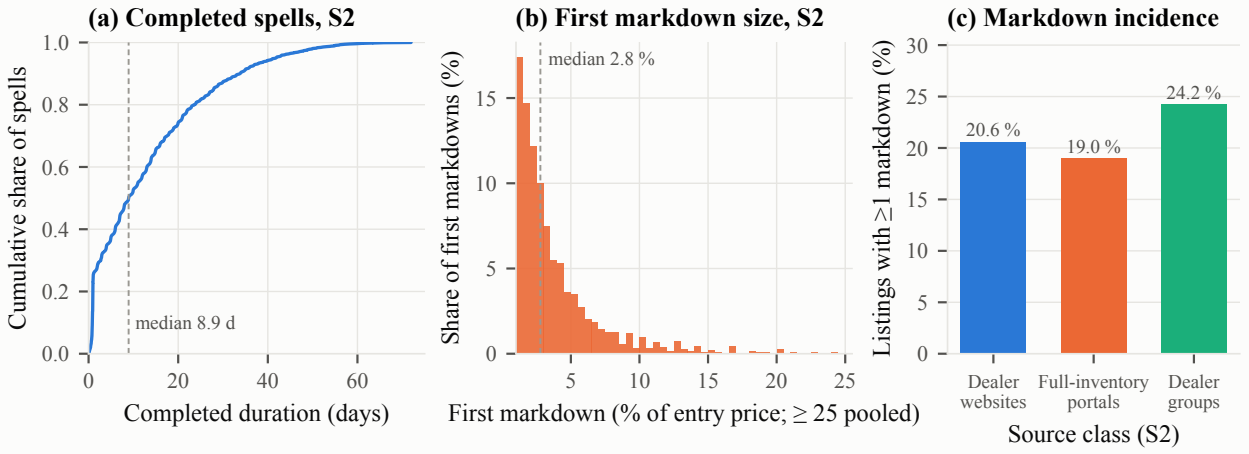


Figure 2. Listing durations and markdowns (S2). (a) Empirical distribution of completed spell lengths. (b) Size of the first sustained markdown as a percentage of the entry price. (c) Share of listings with at least one sustained markdown, by source class.

1.7 percent of Otogo VINs appear under more than one identifier, almost all genuine dealer-to-dealer transfers), and the distribution of crawl gaps (median 3.1 hours, maximum 13.6 hours).

4.4. Descriptive evidence

Table 3 reports descriptive statistics for S1 and S2 and by region. The median listing in S1 asks $25,627 for a 5.7-year-old vehicle with 68,629 km; sport-utility vehicles make up 57 percent of listings, sedans 21 percent and pickups 11 percent; battery-electric vehicles are 6 percent and hybrids 4 percent. Completed listing spells in S2 are short: the median confirmed exit occurs 8.9 days after entry (interquartile range 1.0–20.6 days), 44 percent of completed spells end within a week and 87 percent within 30 days; 46 percent of S2 listings are right-censored. One listing in five (19.7 percent) receives at least one sustained markdown and 6.6 percent receive two or more; the median first cut is 2.8 percent of the entry price (mean 4.6; three quarters are 5 percent or less, 7.6 percent exceed 10 percent). Figure 2 displays these distributions and Figure 3 shows six representative price paths selected by a deterministic rule described in the caption.

5. Econometric framework and identification

The hypotheses, estimands, samples and tests were written down before estimation and are reproduced in the replication package together with every amendment (three amendments were made after the data audit: quality fixed effects at the make–model level rather than make–model–year, because model-year effects absorb vehicle age in an 81-day window; a linear-probability specification for the competition event study; and a per-source definition of left-censoring after the portal start-ups

Table 3. Descriptive statistics by sample and region

	S1 hedonic	S2 dynamic
Listings (N)	136 393	80 856
Dealers	1 046	982
Entry asking price, mean ($)	29 644	29 942
Entry asking price, median ($)	25 627	25 995
Mileage, mean (km)	77 817	74 387
Vehicle age at entry, mean (years)	5.7	5.5
Powertrain (% of listings)		
ICE	81.3	80.4
HEV	3.6	3.7
PHEV	1.7	1.9
BEV	6.1	7.1
Diesel	1.7	1.4
Unknown	5.6	5.5
Body type, top 5 (% of listings)		
SUV	56.7	57.1
Sedan	20.7	20.7
Pickup	10.6	10.3
Hatchback	4.2	4.4
Van	3.1	3.0
Listing dynamics		
Completed spells (confirmed exits)	47 934	43 734
Duration of completed spells, median (days)	8.9	8.9
Listings with ≥ 1 sustained markdown (%)	27.2	19.7
First markdown, median (% of entry price)	2.9	2.8

Region	N (S1)	Median price ($)	Median km	Median age (y)	Exit $\leq$ 30 d (%)
Montréal	12 120	28 995	56 032	4.0	65.1
Laval	8 968	26 495	68 084	4.9	58.1
Montérégie	26 490	25 488	73 319	4.9	61.3
Laurentides	12 654	25 696	71 500	4.9	53.8
Lanaudière	8 549	26 495	44 422	3.9	56.3
Capitale-Nationale	17 015	26 555	55 315	4.0	63.6
Chaudière-Appalaches	9 176	25 322	76 918	5.0	53.2
Estrie	4 627	23 995	78 690	5.0	56.8
Mauricie	7 146	23 995	68 997	4.9	61.1
Centre-du-Québec	3 839	25 688	72 640	4.9	52.0
Outaouais	4 348	24 988	78 546	5.1	55.9
Bas-Saint-Laurent	4 423	23 995	76 619	5.9	60.3
Saguenay–Lac-Saint-Jean	4 009	25 995	72 220	5.9	69.2
Abitibi-Témiscamingue	3 338	26 970	73 252	5.0	53.2
Côte-Nord	2 909	25 490	65 486	4.9	57.6
Gaspésie–Îles-de-la-Madeleine	741	23 995	79 782	5.0	50.8
Region missing	6 041	21 995	85 479	5.9	43.2
All regions	136 393	25 627	68 629	4.9	59.3

Notes: Entry price in Canadian dollars; vehicle age in years at entry; duration = completed spells (confirmed exits). Exit within 30 days is computed on S2 entrants observed for at least 37 days. "Region missing" = listings whose region could not be recovered.

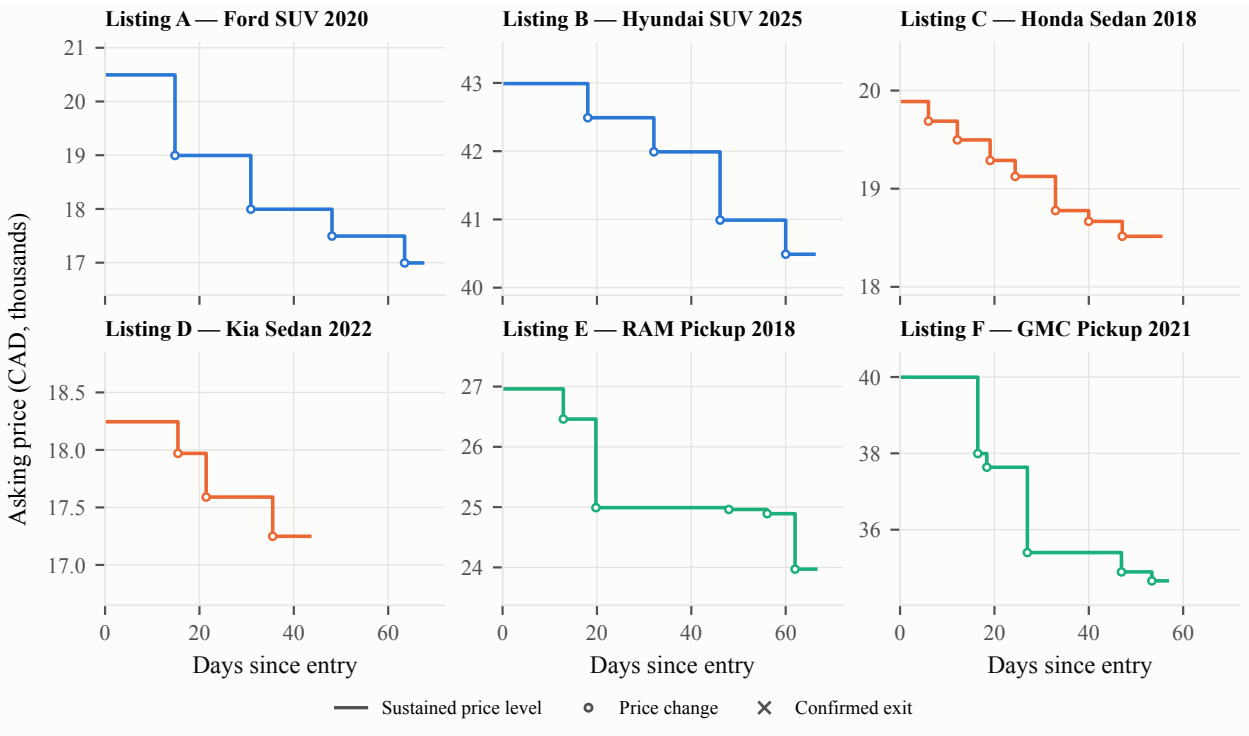


Figure 3. Asking-price paths of six anonymised listings. Selection rule: S2 listings with at least three sustained cuts and a duration between 30 and 70 days, in the three largest body classes; within each class the two listings with the lowest SHA-256 digest of their anonymised identifier. Only make, body and model year are shown.

were discovered). We summarise the six models here.

5.1. Model A: hedonic asking prices (H1, H4)

For listing i with entry price P_i,

$$\log P_i = f(\mathrm{Age}_i, \mathrm{KM}_i) + X_i'\beta + \mu_{m(i)} + \tau_{w(i)} + \theta_{s(i)} + \rho_{r(i)} + \varepsilon_i, \tag{1}$$

where f is a sum of cubic B-splines in age and mileage (five degrees of freedom each, knots at quintiles), optionally augmented by their tensor product (A2t) or interacted with powertrain (A3); X_i contains transmission, drivetrain, powertrain and body class; μ_m are make–model (A1–A3) or make–model–trim (A4) fixed effects; τ_w are entry-week effects; θ_s source-class effects; and ρ_r region effects with Montréal as reference. A5 adds dealer fixed effects, which absorb region. Standard errors are clustered by dealer. The nonlinear components of the spline bases are orthogonalised on the linear term so that H1 is a cluster-robust Wald test that the nonlinear coefficients are zero; H4 is a Wald test that the sixteen region effects are jointly zero in A4. Out-of-sample performance is assessed by five-fold cross-validation grouped by dealer, against a linear benchmark and a gradient-boosting model (A6) with the same inputs.

5.2. Models B and C: markdown and exit hazards (H2, H5)

Listings in S2 are expanded to a listing-day risk set (1.66 million listing-days). For the first markdown, the risk set runs from entry to the day of the first sustained cut (318,323 listing-days, 15,692 events); for exit it runs to the confirmed exit or censoring day (43,311 events). We estimate discrete-time complementary log-log hazards

$$\mathrm{cloglog}\, h_{it} = g(a_{it}) + \gamma \log(1 + N_{it}) + \delta\, r_i + \lambda\, \pi_{it} + Z_i'\beta + \rho_{r(i)} + \tau_t + \theta_{s(i)}, \tag{2}$$

with listing-age bins $g(\cdot)$, the number N_{it} of competing listings (same make and model, model year within one year, different dealer, within 25 km) active on day t, the entry hedonic residual r_i from A2 (the listing's price relative to its quality-adjusted value), the price position π_{it} relative to the mean log price of its competitors (zero with an indicator when there is no competitor), vehicle characteristics Z_i, region, calendar-week and source-class effects; the exit model adds the number of prior cuts. Standard errors are clustered by dealer. Cox proportional-hazards models with dealer-clustered sandwich errors, Schoenfeld tests and a log-logistic accelerated-failure-time model are estimated at the listing level with entry covariates; the conditional size of the first cut is modelled by least squares on log cut size with make–model, region and week fixed effects.

5.3. Model D: competitor price response (H3)

For listing-day (i,t) in S2, let $C_{i,t-k}$ indicate that at least one competitor in the 25-km set cut its price on day $t-k$. We estimate

$$\mathbb{1}[\text{cut}_{it}] = \sum_{k=-7}^{14} \beta_k C_{i,t-k} + \gamma \log(1+N_{it}) + g(a_{it}) + \alpha_{r(i),w(t)} + \alpha_{my(i),w(t)} + u_{it}, \tag{3}$$

with region-by-week and make–model–year-by-week fixed effects, by linear probability (the daily base rate is 1.4 percent, so the linear approximation is local), with errors clustered by dealer and, alternatively, by region. The leads $k<0$ test for pretrends and reverse timing; a placebo permutes competitor series across listings within make–model cells, breaking geography while preserving each series' own dynamics; a robustness specification adds dealer-by-weekday fixed effects to absorb repricing calendars. We interpret all estimates as conditional associations: there is no exogenous source of variation in competitor cuts, and common local demand shocks are the leading alternative explanation.

5.4. Model E: regional segmentation (H4)

Region effects from A4, within-specification dispersion (the interquartile range of A4 residuals by region), and a comparison of the Outaouais region, which borders Ontario, with the Montréal-area regions.

5.5. Model F: prediction

Log entry price, a first markdown within 30 days, and an exit within 30 days are predicted from entry information only, with a regularised linear or logistic baseline and gradient boosting, under chronological, dealer-holdout and region-holdout splits; we report root-mean-squared error, log loss, Brier score, calibration and bootstrap intervals.

5.6. The liquidity index (H6)

For market m (region, or price tier, or body class) and panel week w, we compute exit intensity (confirmed exits per listing-day), markdown frequency (cuts per listing-day), markdown depth (median cut size), quality-adjusted dispersion (interquartile range of hedonic residuals among active listings, sign reversed) and arrival rate (entries per listing-day); higher values of the index mean a *more liquid* market. Components are standardised across market-weeks, markets with fewer than 30 listing-days are set to missing, and three weighting schemes are compared: equal weights, the first principal component (signed so that exit intensity loads positively), and supervised weights

from a ridge regression of next-week exit intensity on the non-exit components fitted on weeks 2–5. Internal validation regresses next-week exit intensity on an index built *without* the exit component, with market fixed effects, on weeks 6–10. Uncertainty bands come from a block bootstrap over dealers (200 replications). Weeks 0–1 are excluded because of the portal start-ups. No external forecasting test is possible with 81 days of data, and none is claimed.

5.7. Inference and multiple testing

The six primary tests (H1 joint nonlinearity, H2 listing-age profile, H3 lag sum, H4 region joint test, H5 relative-price coefficient in the exit hazard, H6 held-out validation slope) are reported with Holm-adjusted p-values in Section 9; secondary tests are unadjusted and labelled exploratory.

6. Price formation and depreciation

6.1. Hedonic asking prices

Table 4 reports six specifications of equation (1) on the 136,393 listings of S1. The linear benchmark (column 1) implies that asking prices fall by 8.5 percent per year of age and 3.2 percent per 10,000 km, conditional on make–model, entry week, source class and region; make–model effects alone explain most of the cross-sectional variance ($R^2 = 0.92$). Adding spline terms (column 2) raises the within-R^2 from 0.836 to 0.846 and lowers the in-sample root-mean-squared error from 0.167 to 0.162 log points; the tensor interaction (column 3) and the powertrain interactions (column 4) add little in sample. Trim fixed effects (column 5) reduce the error to 0.107 and dealer fixed effects (column 6) to 0.091, which shows that a substantial share of residual price variation is dealer-specific pricing rather than vehicle heterogeneity.

6.2. Nonlinear depreciation (H1)

The preregistered test of H1 is decisive (Table 5): the nonlinear age and mileage terms are jointly significant ($\chi^2(12) = 664$), as are the tensor interaction ($\chi^2(49) = 667$) and the interaction of the age profile with powertrain ($\chi^2(35) = 312$). The out-of-sample criterion is met as well: with dealers held out, the spline specification lowers the root-mean-squared error from 0.176 to 0.171 log points (2.8 percent), beyond the 0.5 percent threshold set in the analysis plan; the tensor model does not improve on the additive splines out of sample (0.172), and gradient boosting, which can exploit make, model and trim interactions freely, reaches 0.159. Figure 4 shows the estimated surface. Relative to a one-year-old vehicle with 10,000 km of the same make and model, the quality-adjusted asking price is 14 percent lower at one year and 50,000 km, 27 percent lower at three years, 37 percent lower at five years and 57 percent lower at ten years (all at 50,000 km); at

Table 4. Hedonic asking-price models (S1, dependent variable: log entry price)

	(1) A1 linear	(2) A2 splines	(3) A2t tensor	(4) A3 splines × powertrain	(5) A4 splines, trim FE	(6) A5 + dealer FE
Vehicle age (years)	-0.085***	-0.085***	0.010	-0.085***	-0.083***	-0.080***
	(0.001)	(0.001)	(0.015)	(0.001)	(0.001)	(0.001)
Mileage (10 000 km)	-0.031***	-0.032***	0.076	-0.032***	-0.030***	-0.028***
	(0.000)	(0.000)	(0.442)	(0.000)	(0.001)	(0.000)
Battery electric	0.056***	0.060***	0.059***	0.099***	0.026	0.039***
	(0.011)	(0.011)	(0.011)	(0.023)	(0.019)	(0.014)
Hybrid	0.072***	0.075***	0.074***	0.047***	0.044***	0.031***
	(0.008)	(0.008)	(0.008)	(0.013)	(0.011)	(0.010)
Plug-in hybrid	0.106***	0.118***	0.117***	0.135	0.031***	0.033***
	(0.017)	(0.019)	(0.019)	(0.719)	(0.012)	(0.011)
Diesel	0.079***	0.091***	0.090***	-0.011	0.057***	0.047***
	(0.010)	(0.009)	(0.009)	(0.014)	(0.015)	(0.013)
Manual transmission	-0.034***	-0.023**	-0.023**	-0.023**	-0.060***	-0.055***
	(0.011)	(0.009)	(0.009)	(0.009)	(0.011)	(0.008)
Front-wheel drive	-0.029***	-0.024***	-0.024***	-0.024***	-0.018***	-0.024***
	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)	(0.004)
Rear-wheel drive	-0.043**	-0.039*	-0.043**	-0.039*	-0.051**	-0.060***
	(0.021)	(0.020)	(0.019)	(0.020)	(0.020)	(0.019)
Full-inventory portal	-0.041***	-0.032***	-0.033***	-0.032***	-0.030***	absorbed
	(0.012)	(0.010)	(0.010)	(0.010)	(0.012)	
Dealer group	-0.011	-0.021**	-0.021**	-0.021**	-0.020	absorbed
	(0.011)	(0.010)	(0.010)	(0.010)	(0.013)	
Nonlinear age/km terms	no	yes	yes	yes	yes	yes
Tensor age × km	no	no	yes	no	no	no
Quality FE	make-model	make-model	make-model	make-model	make-model-trim	make-model-trim
Entry-week FE	yes	yes	yes	yes	yes	yes
Region dummies	yes	yes	yes	yes	yes	absorbed
Dealer FE	no	no	no	no	no	yes
Observations	136 393	136 393	136 393	136 393	136 393	136 393
R^2	0.923	0.928	0.929	0.928	0.968	0.977
Within R^2	0.836	0.846	0.848	0.847	0.801	0.777
In-sample RMSE (log)	0.167	0.162	0.161	0.161	0.107	0.091
Grouped-dealer CV RMSE	0.176	0.171	0.172	–	0.161	–

Notes: Standard errors clustered by dealer in parentheses; *, **, ***: 10, 5, 1 percent. Columns (2)–(6) add orthogonalised cubic B-spline terms in age and mileage (not shown); in column (3) the linear age and mileage coefficients are not separately interpretable because of the tensor terms. Reference categories: combustion engine, automatic transmission, all-wheel drive, dealer website, Montréal. Grouped-dealer CV = root-mean-squared error in five-fold cross-validation with dealers held out; a gradient-boosting model with the same inputs attains 0.159.

Table 5. Hypothesis tests on the hedonic models (cluster-robust Wald tests)

Test	χ^2	df	p
H1: nonlinear age terms = 0 (A2)	340.1	6	2.01e-70
H1: nonlinear mileage terms = 0 (A2)	412.8	6	5.06e-86
H1: all nonlinear terms = 0 (A2)	663.6	12	2.73e-134
H1: age × mileage tensor = 0 (A2t)	661.2	49	3.69e-108
H1: age profile × powertrain = 0 (A3)	312.1	35	3.37e-46
H4: region effects = 0 (A4)	14.9	16	0.534

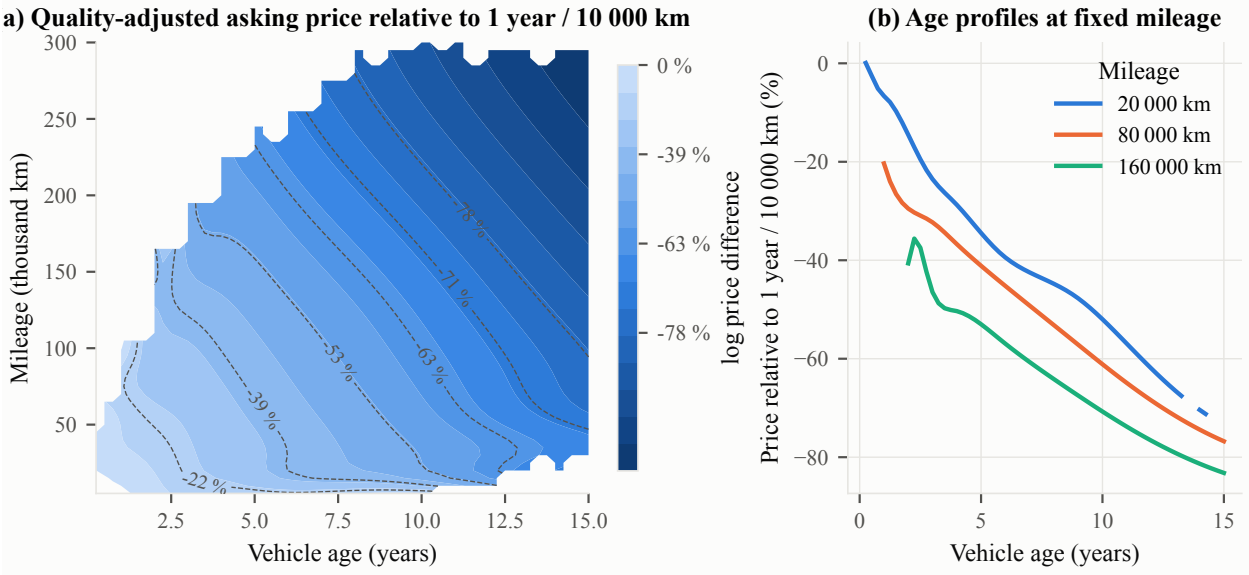


Figure 4. Quality-adjusted depreciation surface (Model A2t). (a) Predicted log asking price relative to a one-year-old vehicle with 10,000 km of the same make and model, over age and mileage; cells without data support are blank. (b) Age profiles at fixed mileage.

five years, moving from 10,000 to 200,000 km lowers the price from 28 to 59 percent below the reference. Age profiles are concave early and close to linear in logs afterwards, and the mileage gradient is steeper for younger vehicles, which is the signature of a market that prices expected remaining service rather than the odometer alone (Engers et al., 2009).

6.3. Regional segmentation (H4)

Table 6 and Figure 5 report the region effects. With make–model–trim fixed effects, mileage, age, powertrain, body, transmission, drivetrain, entry week and source class held fixed, the sixteen regional premia relative to Montréal range from -1.0 percent (Lanaudière) to $+3.7$ percent (Bas-Saint-Laurent); none is individually significant at the 5 percent level with dealer-clustered standard errors, and the joint test does not reject equality ($\chi^2(16) = 14.9$, $p = 0.53$). The point estimates have a mild geographic pattern, with the remote regions (Bas-Saint-Laurent, Abitibi-Témiscamingue, Gaspésie, Saguenay) 2–4 percent above Montréal, consistent with transport costs and thinner local supply, but the data cannot distinguish these from zero. The Outaouais, which borders Ontario, asks 2.1 percent more than Montréal ($p = 0.14$), so there is no evidence of a cross-border discount. Within-specification dispersion, measured by the interquartile range of A2 residuals, is larger in the remote regions, that is, prices for the same vehicle are more dispersed where there are fewer competing listings. The conclusion is, however, sensitive to the quality controls: with make–model fixed effects only (Model A2, Table 6, fourth column) the region effects are jointly significant ($p = 0.004$) even though every premium remains within ± 3 percent, and adding the sampled portals

Table 6. Quality-adjusted regional premia and dispersion

Region	A4 premium (%)	95% CI	A2 premium (%)	Residual IQR
Bas-Saint-Laurent	3.7	[-2.5, 9.9]	2.7	0.168
Abitibi-Témiscamingue	3.0	[-0.1, 6.2]	2.8	0.158
Gaspésie–Îles-de-la-Madeleine	2.7	[-1.1, 6.5]	2.1	0.154
Mauricie	2.4	[-0.6, 5.4]	2.0	0.131
Saguenay–Lac-Saint-Jean	2.3	[-1.5, 6.1]	2.5	0.156
Outaouais	2.1	[-0.7, 4.9]	1.4	0.137
Estrie	1.5	[-1.9, 4.9]	0.7	0.148
Montérégie	1.3	[-1.4, 4.0]	0.4	0.159
Centre-du-Québec	1.2	[-2.4, 4.9]	1.0	0.154
Chaudière-Appalaches	1.1	[-3.3, 5.6]	-0.1	0.166
Laval	1.1	[-3.2, 5.3]	0.7	0.131
Côte-Nord	0.9	[-2.0, 3.9]	0.6	0.129
Capitale-Nationale	0.0	[-3.2, 3.2]	0.9	0.143
Laurentides	-0.5	[-4.0, 2.9]	-1.3	0.149
Lanaudière	-1.0	[-6.2, 4.1]	-0.8	0.147
Montréal (reference)	0	–	0	0.166

Notes: A4 = Model A4 (make–model–trim FE); A2 = Model A2 (make–model FE); 95% confidence intervals from dealer-clustered standard errors. Residual IQR = interquartile range of A2 residuals among S1 listings in the region.

to the sample also produces a rejection with trim fixed effects ($p = 0.014$; Section 9). We therefore read H4 not as a strict null but as a bound: regional differences in quality-adjusted asking prices exist at most at the level of a few percent, smaller than the within-region dispersion for the same vehicle, and they are not statistically detectable once trim-level quality is held fixed.

7. Dealer repricing, competition and inventory dynamics

7.1. Markdown hazards and cut sizes (H2)

Figure 6(a) shows the empirical daily hazard of a first sustained markdown by listing age. It is low in the first days (0.7 percent per day), doubles by the end of the first week (1.1 percent) and reaches a plateau of 1.4–1.5 percent per day between the second and fourth weeks, after which it declines slightly among the surviving listings. In the discrete-time model of Table 7, column (1), estimated on the 1.37 million listing-days at risk of a first cut, the listing-age bins are jointly significant ($\chi^2(8) = 157$) and the coefficients rise monotonically from 0.91 (days 1–3, relative to the entry day) to 1.8 (beyond three weeks), so the first half of H2 is supported: the probability of a markdown increases with the time a vehicle has been on the lot and then plateaus. Relative price matters: a listing that enters 10 percent above its hedonic value has a 7.7 percent higher daily markdown hazard (coefficient 0.74 on the residual, $p < 0.001$), and a 10 percent higher price than its local competitor set adds 3.9 percent ($p < 0.001$); the listing-level Cox model with dealer-clustered sandwich errors (column 2) gives almost identical coefficients (0.72 and 0.36). Higher-mileage vehicles are cut

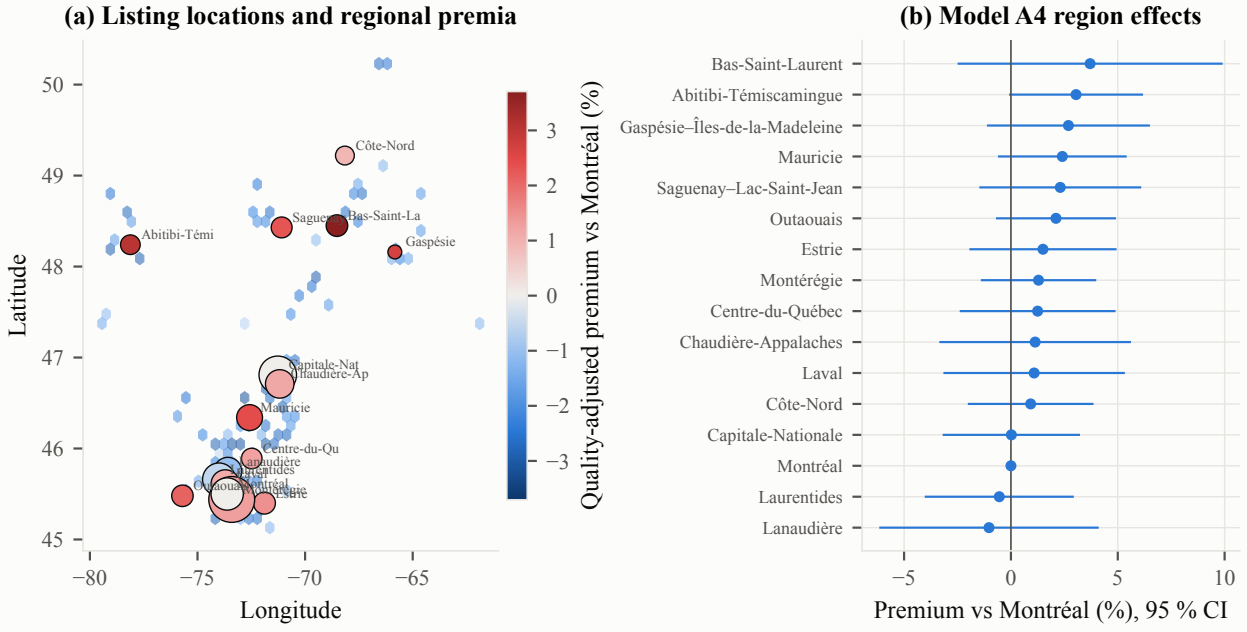


Figure 5. Regional premia. (a) Listing locations (hexagonal bins, log scale) and region effects (marker size proportional to the number of listings, colour = premium versus Montréal). (b) Region effects from Model A4 with 95% confidence intervals.

more often (a 10 percent higher mileage raises the hazard by about 1 percent), while the price level itself and the vehicle's age are not related to the markdown hazard once the relative price is held fixed. The density of local competitors is the one preregistered covariate without a stable effect: its coefficient is positive and significant in the Cox model (0.11) but close to zero and insignificant in the discrete-time model with region and calendar-week effects (0.02, standard error 0.03), and the robustness suite shows that it changes sign across markdown definitions. The cut-size regression (column 5) is unambiguous: the first cut is 2.7 percent of the entry price at the median, it is larger for listings that entered above their hedonic value (a 10 percent overpricing raises the cut by 22 percent) and for expensive vehicles in absolute terms, and it declines with listing age at the time of the cut, that is, late cuts are smaller than early ones.

7.2. Competitor price response (H3)

Figure 7 and Table 8 report the event study of equation (3). The daily probability that a dealer cuts the price of a listing is 1.44 percent on average. On the day a close competitor (same make and model, model year within one year, within 25 km) cuts its price, that probability is 0.89 percentage points higher (standard error 0.21), a 60 percent increase relative to the base rate. The association is, however, entirely contemporaneous: the coefficients for days $+1$ to $+6$ are zero or negative, and the only other positive coefficients appear at $k = +7$, $k = +14$ and $k = -7$, a weekly periodicity.

Table 7. Markdown and listing-exit models (S2)

	First markdown		Listing exit		Cut size
	cloglog	Cox	cloglog	Cox	log(%)
log(1 + competitors within 25 km)	0.019	0.114***	0.033	0.050*	-0.027**
	(0.029)	(0.025)	(0.030)	(0.029)	(0.011)
Entry hedonic residual (relative price)	0.744***	0.716***	0.180	0.129	2.224***
	(0.223)	(0.206)	(0.225)	(0.209)	(0.221)
Price position vs competitors at entry	0.379***	0.356***	–	-0.156	–
	(0.113)	(0.116)		(0.125)	
Current price position vs competitors	–	–	-0.159	–	–
			(0.139)		
Prior sustained cuts (count)	–	–	0.033	–	–
			(0.036)		
log entry price	0.007	0.091	-0.307***	-0.255***	-1.526***
	(0.091)	(0.079)	(0.111)	(0.095)	(0.117)
log mileage	0.116***	0.106***	-0.512***	-0.490***	-0.194***
	(0.037)	(0.039)	(0.070)	(0.070)	(0.018)
Vehicle age at entry (years)	-0.022	-0.010	0.034**	0.038***	-0.069***
	(0.013)	(0.012)	(0.015)	(0.014)	(0.010)
Own-dealer same-model stock	-0.000	–	0.009	–	–
	(0.004)		(0.006)		
Listing-age bins (joint χ^2, df 8)	157.0	–	103.6	–	–
Listing-days / listings	1 371 918	79 698	1 663 763	79 698	15 692
Events	15 692	15 692	43 310	43 310	–
Concordance / within R^2	–	0.575	–	0.648	0.233
PH test: covariates with $p < 0.01$	–	5	–	8	–

Notes: cloglog = discrete-time complementary log-log hazard on listing-days with listing-age bins, powertrain, body, source class, region and calendar-week effects, standard errors clustered by dealer. Cox = proportional-hazards model at the listing level with entry covariates and dealer-clustered sandwich errors. Cut size = OLS on the log of the first cut (percent of entry price) with make–model, region and week fixed effects. The entry hedonic residual is from Model A2. *, **, ***: 10, 5, 1 percent.

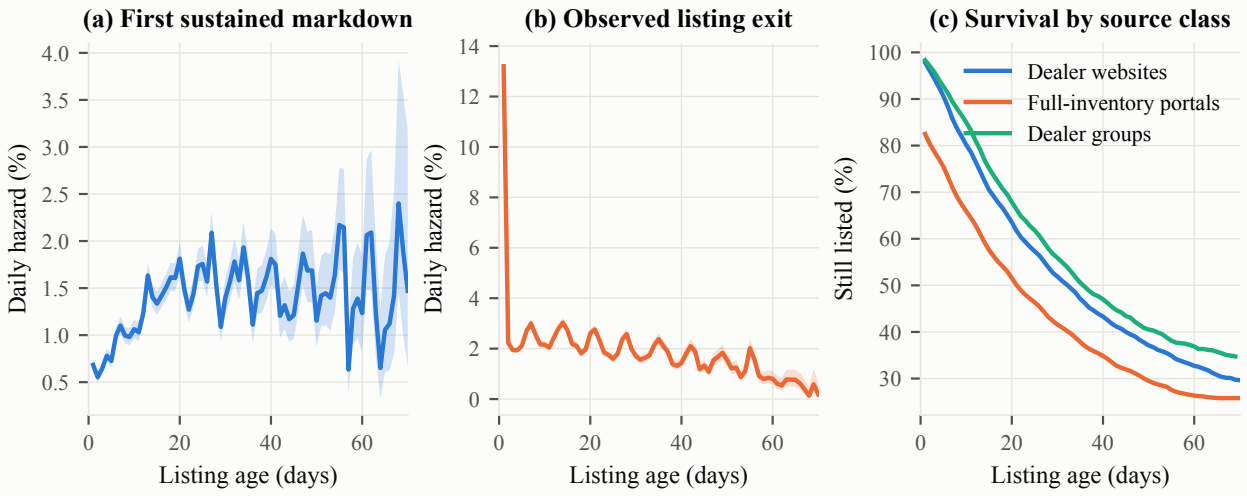


Figure 6. Empirical daily hazards by listing age (S2) with 95% Wilson intervals. (a) First sustained markdown. (b) Observed listing exit. (c) Implied survival by source class.

Table 8. Competitor price-response event study (dependent variable: own sustained cut on day t, percentage points)

	Main, 25 km	Cluster: region	10 km	50 km	Coarse FE	Placebo	+ dealer × weekday FE	First cut, + dealer × weekday
$k = -7$	0.42	0.42	0.59	0.40	0.35	-0.05	0.12	0.12
	(0.15)	(0.11)	(0.15)	(0.12)	(0.14)	(0.04)	(0.14)	(0.11)
$k = -3$	-0.04	-0.04	0.09	-0.07	0.04	0.07	0.00	0.00
	(0.06)	(0.09)	(0.12)	(0.05)	(0.06)	(0.04)	(0.05)	(0.05)
$k = -1$	-0.07	-0.07	-0.00	-0.08	0.10	-0.01	-0.03	0.01
	(0.12)	(0.08)	(0.20)	(0.08)	(0.11)	(0.03)	(0.10)	(0.09)
$k = 0$	0.89	0.89	1.80	0.63	1.10	0.01	0.55	0.49
	(0.21)	(0.21)	(0.55)	(0.13)	(0.20)	(0.04)	(0.20)	(0.17)
$k = +1$	-0.11	-0.11	0.11	-0.20	0.07	0.03	-0.01	0.06
	(0.20)	(0.12)	(0.44)	(0.17)	(0.20)	(0.04)	(0.18)	(0.15)
$k = +2$	-0.24	-0.24	-0.15	-0.28	-0.11	-0.04	-0.16	-0.11
	(0.07)	(0.07)	(0.11)	(0.06)	(0.06)	(0.04)	(0.07)	(0.06)
$k = +3$	-0.08	-0.08	0.12	-0.10	0.04	-0.04	-0.03	0.02
	(0.08)	(0.08)	(0.10)	(0.07)	(0.08)	(0.04)	(0.06)	(0.06)
$k = +7$	0.45	0.45	0.66	0.54	0.42	0.07	0.17	0.16
	(0.08)	(0.08)	(0.17)	(0.06)	(0.08)	(0.06)	(0.06)	(0.07)
$k = +14$	0.59	0.59	0.81	0.55	0.57	0.08	0.33	0.33
	(0.11)	(0.09)	(0.22)	(0.09)	(0.11)	(0.08)	(0.09)	(0.11)
Sum of leads $k = -7..-1$	0.37	0.37	0.90	0.24	0.69	0.04	0.50	0.42
joint p	3.2e-05	1.5e-13	3.7e-06	3.8e-07	0.00087	0.26	2.5e-05	4.4e-06
Sum of lags $k = 0..7$	0.95	0.95	2.92	0.73	1.70	-0.03	0.71	0.85
joint p	4.3e-13	1.2e-147	2.9e-08	3e-21	6.5e-18	0.42	0.00021	5.5e-05
Daily base rate (%)	1.44	1.44	1.44	1.44	1.44	1.44	1.44	1.14
Listing-days	1 663 763	1 663 763	1 663 763	1 663 763	1 663 763	1 663 763	1 663 763	1 371 918

Notes: Linear probability models on S2 listing-days with region-by-week and make–model–year-by-week fixed effects, log competitor density and listing-age effects; coefficients are percentage points, standard errors clustered by dealer (by region in column 2) in parentheses. Placebo: competitor-cut series permuted across listings within make–model cells. "+ dealer × weekday FE" adds dealer-by-weekday effects.

The leads are jointly significant ($p < 10^{-4}$), so the pretrend test fails in the sense that a dealer's cut today also predicts a competitor's cut a week later; lagged and lead sums are of the same order. Adding dealer-by-weekday fixed effects, which absorb each dealer's repricing calendar, reduces the same-day coefficient to 0.55 points and the seven-day harmonics to 0.17 and 0.33 points, but does not change the picture. The placebo, in which competitor series are permuted across listings within make–model cells, yields coefficients indistinguishable from zero at every horizon (joint $p = 0.18$), which confirms that the same-day association is spatially structured and not an artefact of the fixed effects. Results are similar at 10 km (larger coefficients, fewer competitors) and 50 km (smaller), when errors are clustered by region, and with coarser fixed effects. For the first markdown only (panel b), the same-day coefficient is 0.73 points against a base rate of 1.14 percent, again with no lagged response and the same weekly harmonics (0.49 points with dealer-by-weekday effects).

We draw two conclusions. First, there is no evidence of a lagged strategic response: dealers do not follow a competitor's cut within the next one to six days. Second, the strong same-day co-movement, its weekly harmonics and its robustness to dealer-by-weekday effects point to common local shocks and synchronised repricing routines (dealers and dealer groups that update their inventory feeds on fixed days, manufacturer incentive calendars, local demand shocks) rather than to observation and reaction. Because the design has no exogenous source of variation in competitor cuts, the estimates are associations and we make no causal claim; the reflection problem of Manski (1993) applies with full force.

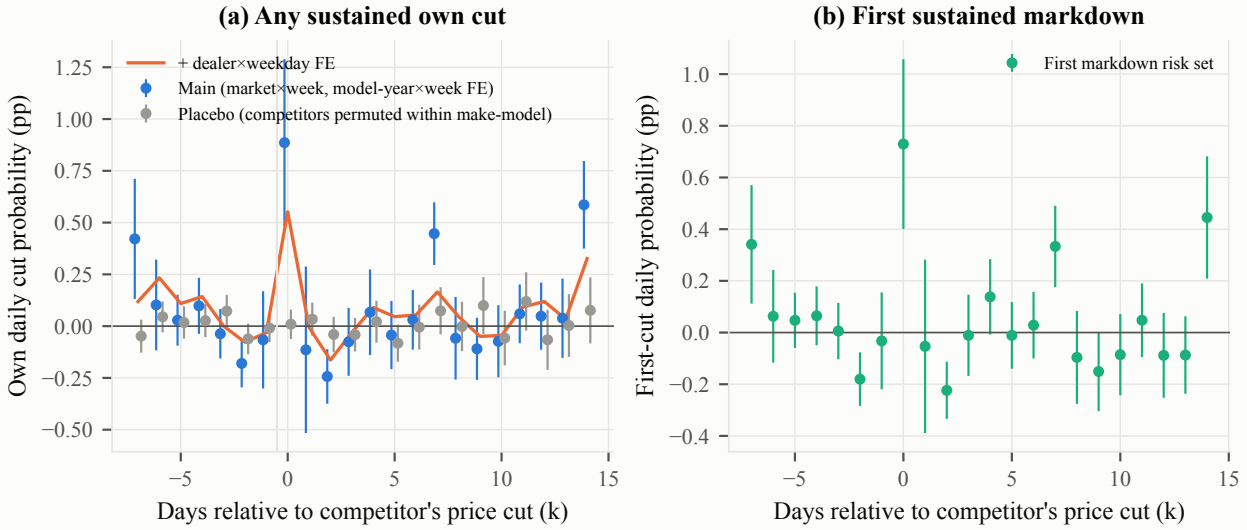


Figure 7. Own repricing around a competitor's price cut (S2). (a) Any sustained own cut: main specification, placebo with permuted competitors, and the specification with dealer-by-weekday fixed effects. (b) First sustained markdown. Bars are 95% confidence intervals (dealer-clustered).

7.3. Listing exit and its validation (H5)

Figure 6(b) shows the exit hazard. It is very high on the first day after entry (13 percent), which reflects listings that are posted and withdrawn within a day (dealer-to-dealer transfers, duplicate postings, feed corrections), then falls to 2 percent, rises to about 3 percent per day in the second and third weeks, and declines thereafter to under 1 percent per day beyond two months; the implied survival curves by source class in panel (c) are similar. The discrete-time exit model (Table 7, column 3) confirms strong negative duration dependence after the first days and shows that cheaper and, above all, higher-mileage vehicles disappear faster (a 10 percent higher mileage lowers the exit hazard by 5 percent; a 10 percent higher price by 3 percent), and that older vehicles exit faster. The preregistered relative-price test of H5 is not supported: the coefficient on the entry hedonic residual is positive (0.18 in the cloglog, 0.13 in the Cox model) but imprecise once errors are clustered by dealer ($p = 0.42$ and 0.54), and the current price position relative to competitors has a negative but insignificant coefficient. Overpriced listings are cut sooner and more deeply (previous subsection), but they do not leave the market faster or more slowly within the horizon of the data. Competitor density raises the exit hazard mildly (Cox coefficient 0.05, $p = 0.09$). The Schoenfeld tests reject proportional hazards for most covariates in both models, which is expected given the shape of the first-day hazard; the discrete-time models with listing-age bins do not impose proportionality and the log-logistic accelerated-failure-time estimates (replication package) agree in sign.

Because exit is not a sale, we validate it in three ways. First, for the 17,898 vehicles observed simultaneously on a dealer's own website and on the Otogo portal (matched by VIN, overlapping

spells), exits coincide closely: among the 8,541 pairs where both listings exited, 80 percent exited within one day of each other and 94 percent within a week (median gap 0.4 days). Second, the exceptions are informative about feed failures: in 11 percent of the dealer-site exits the portal copy was still listed seven days later, against 3 percent in the other direction, so dealer websites generate more spurious exits than the portal. Third, relistings bound the share of non-sale exits from below: of the 47,398 confirmed exits observable for 30 days, 10.2 percent were followed by a new listing of the same VIN within 30 days on a reliable source and 5.8 percent by a listing at a location more than 5 km away; the median gap for these dealer-to-dealer transfers is 1.1 days (upper quartile 8.7 days). Over the whole window, 10,139 of 95,968 VINs (10.6 percent) were listed at two or more locations at least 5 km apart. Listing exit is therefore a consistent, cross-source event that overstates retail sales by at least one tenth, and we label it "observed disappearance" throughout.

7.4. Prediction (Model F)

Table 9 and Figure 8 report the out-of-sample prediction of a first markdown and of an exit within 30 days of entry, from entry information only. On the chronological split (entries before 15 September used to predict entries from 15 to 25 September, the latest cohort whose 30-day outcomes are fully observable), gradient boosting attains a log loss of 0.405 for markdowns (area under the curve 0.785; base rate 0.19) and 0.585 for exits (0.727; base rate 0.59), both well calibrated (expected calibration error below 0.02), and improves on the logistic baseline. The gains disappear when dealers or regions are held out: areas under the curve fall to 0.58–0.61 and the logistic model does as well or better, because the single most informative feature is the dealer's own markdown or exit rate in the training data (a third of the gradient-boosting gain). Thirty-day repricing and exit are, to a first approximation, dealer practices rather than vehicle attributes, and they do not transfer to unseen dealers. A Cox model for time to exit with entry covariates attains a concordance of 0.66 on the later cohort.

8. The Automobile Market Liquidity Index

8.1. Construction

Table 10 summarises the index. Five components are computed for each of the seventeen regions and each panel week from week 2 to week 10 (the two first weeks are contaminated by the portal start-ups): exit intensity, markdown frequency, markdown depth, quality-adjusted dispersion (sign reversed) and arrival rate. Each is standardised across the 153 region-weeks; higher values mean a more liquid market under the preregistered sign convention. Figure 9 plots the province-level components and index and the regional indices with dealer-bootstrap bands.

Table 9. Out-of-sample prediction of 30-day markdown and exit (S2, entry information only)

Target	Split	Model	Log loss [95% CI]	Brier	AUC	n_{test}	Base rate
Markdown $\leq$ 30 d	Chronological	Logistic	0.428 [0.417, 0.438]	0.135	0.748	12 953	0.193
		LightGBM	0.406 [0.395, 0.416]	0.129	0.784	12 953	0.193
Markdown $\leq$ 30 d	Dealer holdout	Logistic	0.476 [0.469, 0.483]	0.150	0.580	44 003	0.184
		LightGBM	0.484 [0.477, 0.492]	0.151	0.578	44 003	0.184
Markdown $\leq$ 30 d	Region LORO	Logistic	0.508 [0.500, 0.516]	0.163	0.606	28 947	0.198
		LightGBM	0.529 [0.518, 0.539]	0.167	0.581	28 947	0.198
Exit $\leq$ 30 d	Chronological	Logistic	0.609 [0.603, 0.616]	0.212	0.705	12 953	0.592
		LightGBM	0.585 [0.580, 0.592]	0.204	0.727	12 953	0.592
Exit $\leq$ 30 d	Dealer holdout	Logistic	0.695 [0.694, 0.697]	0.251	0.566	44 003	0.593
		LightGBM	0.682 [0.680, 0.684]	0.244	0.606	44 003	0.593
Exit $\leq$ 30 d	Region LORO	Logistic	0.690 [0.688, 0.693]	0.248	0.584	28 947	0.601
		LightGBM	0.680 [0.675, 0.684]	0.243	0.613	28 947	0.601

Notes: Chronological = train on entries before 2026-09-15, test on 2026-09-15 to 2026-09-25 (latest entries with fully observable 30-day outcomes); dealer holdout = five-fold grouped cross-validation by dealer (pooled out-of-fold); region LORO = leave-one-region-out over the six largest regions (mean). 95% intervals from 200 bootstrap resamples of test rows. Exit = observed disappearance, not a sale.

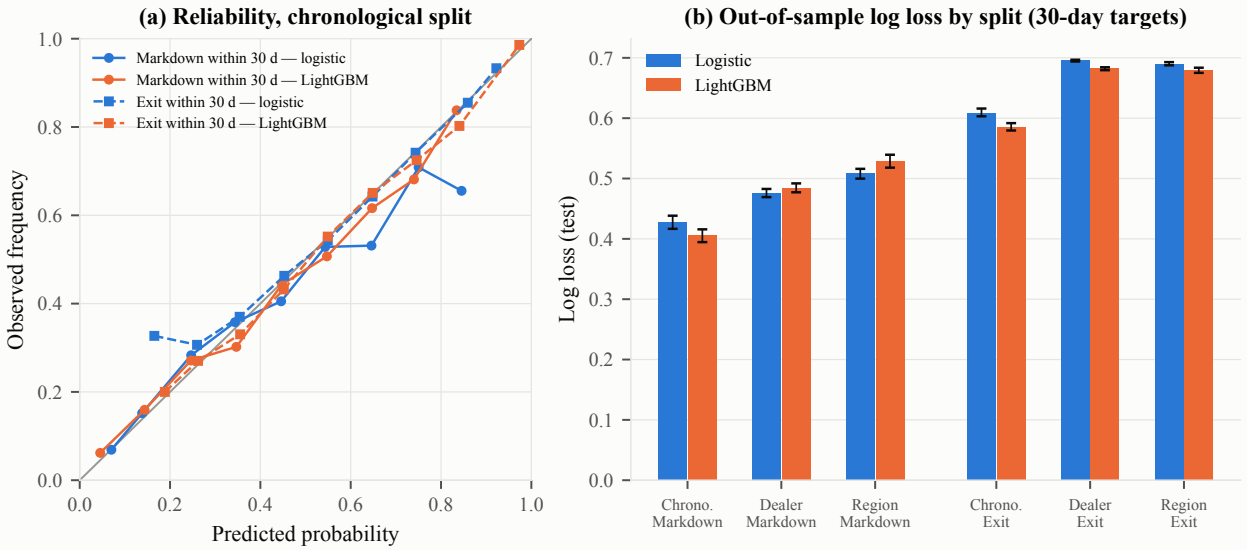


Figure 8. Prediction performance. (a) Reliability diagrams on the chronological split. (b) Log loss by split and model with bootstrap intervals.

Table 10. AMLI construction, weights and internal validation (region × week, weeks 2–10)

Component	Equal	PCA loading	Supervised	Leave-one-out corr.
Exit intensity	0.20	+0.41	–	0.94
Markdown frequency	0.20	+0.35	+0.0011	0.91
Markdown depth	0.20	-0.50	-0.0039	0.95
Residual dispersion (sign reversed)	0.20	+0.51	-0.0015	0.91
Arrival rate	0.20	+0.44	+0.0015	0.94
PCA first-component share		0.56		
Corr(equal, PCA) / Corr(equal, supervised)		0.85 / 0.68		
Leave-exit-out AMLI → next-week exit (region FE)		0.96 pp [0.69, 1.24], $n = 68$		
+ lagged exit		1.02 pp [0.70, 1.35], $n = 68$		
pooled OLS		0.98 pp [0.59, 1.36], $n = 68$		
Price-tier markets: PCA share / held-out slope		0.44 / 0.46 pp [-0.64, 1.55]		
Body-type markets: PCA share / held-out slope		0.46 / 0.44 pp [0.05, 0.82]		

Notes: PCA loading = first principal component of the five standardised components, signed so that exit intensity loads positively. Supervised = ridge coefficients of next-week exit intensity on the four non-exit components, fitted on weeks 2–5. Leave-one-out corr. = correlation between the equal-weight index without the component and the full index. Held-out validation: coefficient of next-week exit intensity on an equal-weight index built without the exit component, region fixed effects, weeks 6–10, standard errors clustered by region; slope in percentage points of daily exit intensity per index unit.

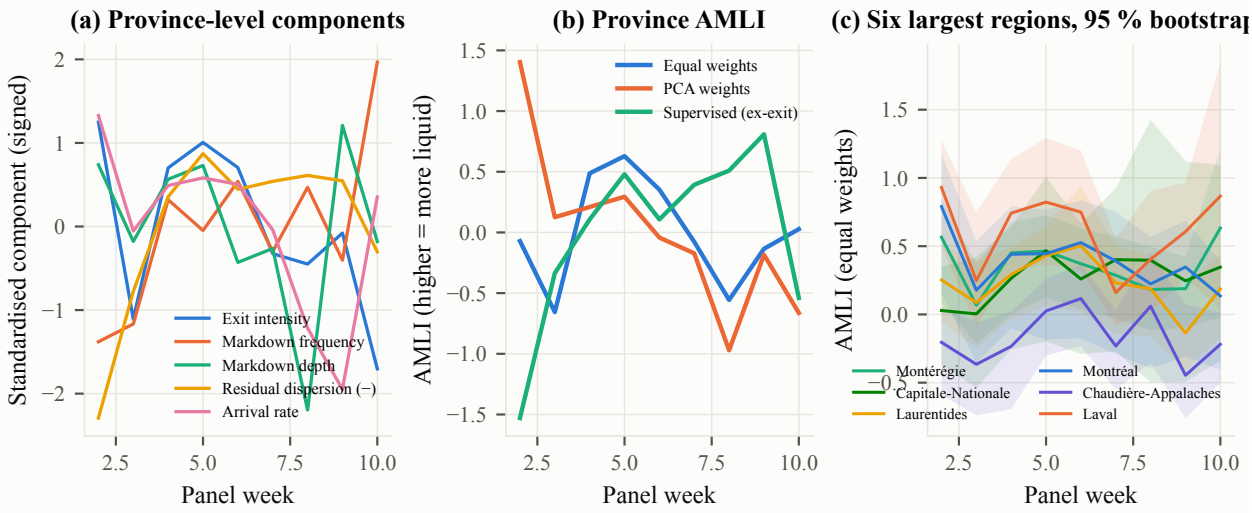


Figure 9. AMLI time series. (a) Province-level standardised components (signed). (b) Province-level index under three weighting schemes. (c) Equal-weight index for the six largest regions with 95% dealer-block-bootstrap bands (200 replications).

8.2. What the components say about liquidity

The component correlations (Figure 10(b)) carry the paper's most useful lesson about liquidity measurement. Exit intensity co-moves positively with the arrival rate ($\rho = 0.62$) and, weakly, with markdown frequency (0.19), but *negatively* with markdown depth (-0.35) and positively with reversed dispersion (0.44): region-weeks in which vehicles leave quickly are region-weeks with many new arrivals, frequent but *shallow* price cuts, and *tight* quality-adjusted price distributions; region-weeks with deep cuts and dispersed prices are slow. Markdown depth and dispersion are themselves strongly correlated (0.85). The first principal component explains 56 percent of the variance of the five components, above the 40 percent threshold set in the analysis plan, and its loadings are positive on exit, markdown frequency, arrival and reversed dispersion, but *negative* on markdown depth. The preregistered sign for depth (deeper cuts = more liquid) is therefore contradicted by the data, which treat deep cuts as a symptom of illiquidity, in the spirit of the list-price-cut literature in housing (Knight, 2002; Anenberg, 2016). We keep the preregistered equal-weight index as the primary measure, report the PCA-weighted index (correlation 0.85 with the equal-weight index) and the supervised index (0.68) alongside it, and record this sign reversal as a substantive finding rather than a tuning choice. At the province level, with only nine weekly observations, the first component explains 44 percent of the variance and its loadings are unstable; the regional panel is the appropriate level for the index in a window this short.

8.3. Internal validation

The analysis plan called for a held-out-component test. An equal-weight index built *without* exit intensity, from the markdown, dispersion and arrival components of week w, predicts the exit intensity of week $w+1$ within regions: the slope is 0.97 percentage points of daily exit intensity per index unit (95% interval 0.69 to 1.24; the mean daily exit intensity is 2.3 percent), with an R^2 of 0.66 including region effects, and it is unchanged when the current exit intensity is added as a control (1.03, interval 0.70 to 1.35). Figure 10(a) shows the within-region relation. The index is robust to leaving any single component out (correlations with the full index between 0.91 and 0.95; the held-out slope stays between 0.58 and 0.82 points), and to the minimum cell size (100 listing-days: unchanged), but the held-out validation is weaker under rank standardisation of the components (0.46 points, standard error 0.29) and uninformative at two-week aggregation, which leaves only four observations per region (Section 9). The same construction on price-tier and body-class markets gives first components explaining 44–46 percent of the variance and held-out slopes of 0.46 points, significant for body classes (interval 0.06 to 0.86) but not for the three price tiers (12 market-weeks only).

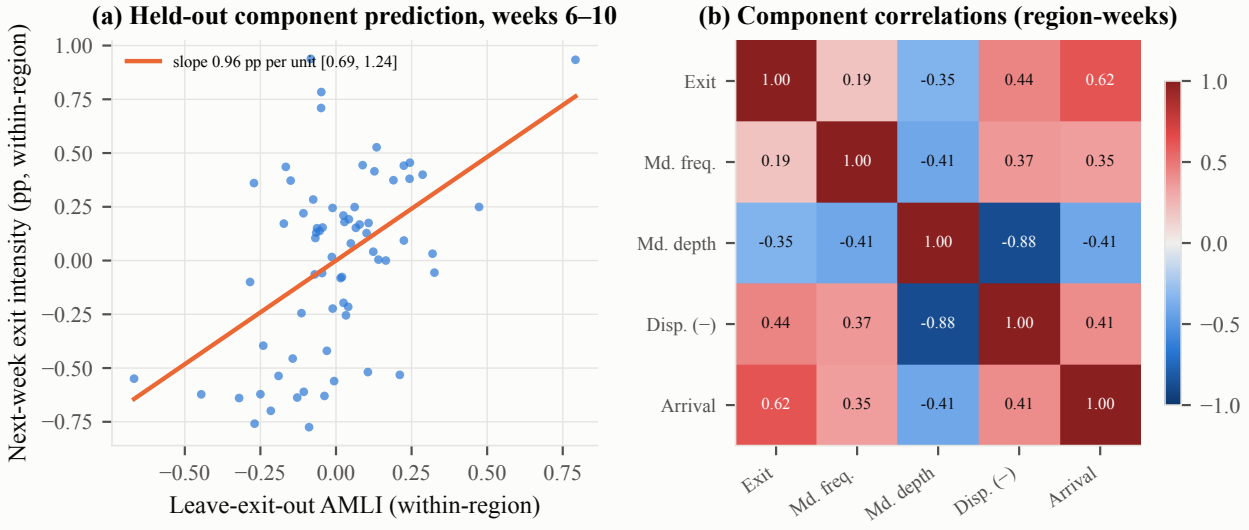


Figure 10. AMLI internal validation. (a) Within-region relation between the leave-exit-out index in week w and exit intensity in week $w+1$, weeks 6–10. (b) Correlations among the five standardised components across region-weeks.

8.4. What the index cannot do yet

The window contains eleven weeks. No weekly or daily independent benchmark of used-vehicle turnover exists for Quebec, and the monthly or quarterly series that do exist (Statistics Canada new-vehicle sales, the Canadian Black Book retention index, the Statistics Canada used-vehicle CPI, AutoTrader's quarterly price index) offer at most three overlapping observations, most of them national. We therefore make no claim that the AMLI forecasts any economic variable, and we regard its validation as internal and partial: it summarises coherent variation in observed turnover across regions and weeks, and its non-exit components anticipate next week's exits. Establishing external validity requires a panel of at least one to two years and, ideally, access to registration transfers, which is the natural next step for the platform.

9. Robustness, identification threats and limitations

9.1. Sensitivity analyses

Table 11 summarises the sensitivity of the headline estimates. Three results are not robust and are reported as such: the joint test of regional premia rejects with make–model controls or when the sampled portals are included, although every premium stays within a few percent; the exit-hazard coefficient on the entry residual changes sign when the left-censored stock is included (that stock has unknown listing age, so this variant is a lower bound on age and not a preferred specification); and the held-out validation of the liquidity index loses significance under rank standardisation and

at two-week aggregation. The depreciation slopes, the listing-age profiles of both hazards, the relative-price effects on the markdown hazard and the cut size, the competitor event study and the leave-one-component-out checks are stable; the competitor-density coefficient in the markdown hazard is not (it changes sign across markdown definitions and is never significant with dealer clusters). An earlier version of the first-markdown model had conditioned on listings that eventually receive a cut because of a null-handling error in the risk set; it was discovered by the robustness audit, corrected, and the whole pipeline was rerun (experiment log E-016). The remainder of the table lists the sensitivity to the main measurement and sample choices, following the list fixed in the analysis plan: alternative markdown thresholds (0.5 and 2 percent), transient rules (24 and 72 hours), exit-confirmation windows (3 and 14 days), inclusion of the sampled portals and of private-seller listings in the hedonic sample, wider price filters, inclusion of the left-censored stock in the hazards, heterogeneity of the depreciation slopes by powertrain, model-year band, body class and source class, composition reweighting of the dealer-website sample to the portal's make–model–year distribution, alternative standardisation, aggregation and minimum-size rules for the liquidity index, and competitor radii of 10 and 50 km (Table 8).

9.2. Multiple testing

The six primary preregistered tests are: H1 (joint nonlinearity, $p = 3 \times 10^{-134}$), H2 (listing-age profile of the markdown hazard, $p = 7 \times 10^{-30}$), H3 (joint test of the competitor-cut lags, $p = 1 \times 10^{-14}$, an association), H4 (joint region test in A4, $p = 0.53$), H5 (entry hedonic residual in the exit hazard, $p = 0.42$) and H6 (held-out validation slope, $p = 2 \times 10^{-13}$). Holm's step-down adjustment (`results/robustness/holm_primary_tests.csv`) leaves the conclusions unchanged: H1, H2, H3 and H6 remain significant at any conventional level (adjusted $p \leq 6 \times 10^{-13}$) and H4 and H5 are not rejected (adjusted $p = 0.85$). The null results are therefore not artefacts of the number of tests, and the positive ones are not artefacts of selective reporting; every estimated specification, including failed and null ones, is retained in the experiment log.

9.3. Identification threats

Asking versus transaction prices. Every price in this paper is a dealer's posted price. The hedonic coefficients measure how dealers *price* quality, not what buyers pay; negotiated discounts, financing and trade-in terms are unobserved. Because Quebec law requires all-in advertised prices, posted prices are comparable across dealers, but the level of the asking–transaction gap and its variation with listing age are unknown here, as in the housing literature on list prices (Han and Strange, 2016; Haurin et al., 2010).

Table 11. Identification threats and sensitivity summary

Check	Quantity	Baseline	Variant	Change
Hedonic (Model A2 splines, make–model FE, dealer-clustered)				
+ sampled portals (autotrader, cargurus)	age slope (log $/yr)	-0.0851 (0.0009)	-0.0865 (0.0006)	-0.0015
+ sampled portals	H4 region joint p (A4 trim FE / A2)	0.534 / 0.004	0.014 / $<$0.001	flips
+ Kijiji private sellers	age slope (log $/yr)	-0.0851 (0.0009)	-0.0879 (0.0013)	-0.0028
+ Kijiji private sellers	H4 region joint p (A4 trim FE / A2)	0.534 / 0.004	0.407 / 0.002	same
Price filter 1 000–500 000 $	age slope (log $/yr)	-0.0851 (0.0009)	-0.0857 (0.0009)	-0.0007
Price filter 1 000–500 000 $	H4 region joint p (A4 trim FE / A2)	0.534 / 0.004	0.539 / 0.007	same
Reweighted to otogo make–model–year mix	age slope (log $/yr)	-0.0851 (0.0009)	-0.0846 (0.0009)	+0.0005
Reweighted to otogo make–model–year mix	H4 region joint p (A4 trim FE / A2)	0.534 / 0.004	0.475 / 0.014	same
Heterogeneity: powertrain	age slope: ICE vs BEV	-0.0852 (0.0009)	-0.0942 (0.0031)	-0.0090
Heterogeneity: model-year band	age slope: $\leq$2015 vs 2021+	-0.0934 (0.0027)	-0.0751 (0.0019)	+0.0183
Heterogeneity: body	km slope: SUV vs Pickup	-0.0306 (0.0004)	-0.0255 (0.0009)	+0.0051
Heterogeneity: source class	age slope: direct_dealer vs full_portal	-0.0854 (0.0015)	-0.0844 (0.0011)	+0.0010
Hazards (cloglog, dealer-clustered; $\leq$800k listing-day subsample)				
Markdown $\geq$2 %/200 $	B: coef. ldens	0.017 (0.030)	-0.022 (0.029)	-0.040
Markdown $\geq$0.5 %/50 $	B: coef. ldens	0.017 (0.030)	0.029 (0.029)	+0.012
Transient rule 24 h	B: coef. ldens	0.017 (0.030)	0.007 (0.029)	-0.010
Transient rule 72 h	B: coef. ldens	0.017 (0.030)	0.022 (0.030)	+0.005
Exit confirmation 3 d	C: coef. resid_A2	0.211 (0.234)	0.151 (0.227)	-0.060
Exit confirmation 14 d	C: coef. resid_A2	0.211 (0.234)	0.202 (0.242)	-0.008
S2L (left-censored stock included)	C: coef. resid_A2	0.211 (0.234)	-0.395 (0.166)	-0.605
S2L (left-censored stock included)	C: coef. ldens	0.033 (0.030)	0.057 (0.023)	+0.023
AMLI (leave-exit-out index $\rightarrow$ next-week exit intensity, region FE)				
Rank standardisation	slope $\hat{\beta}$ (s.e.)	0.0096 (0.0013)	0.0046 (0.0029)	-0.0051
Min. 100 listing-days per cell	slope $\hat{\beta}$ (s.e.)	0.0096 (0.0013)	0.0096 (0.0013)	+0.0000
Two-week aggregation	slope $\hat{\beta}$ (s.e.)	0.0096 (0.0013)	-0.0001 (0.0051)	-0.0097
Leave-one-component-out (min: drop mdfreq)	slope $\hat{\beta}$ (s.e.)	0.0096 (0.0013)	0.0057 (0.0019)	-0.0039
Leave-one-component-out (max: drop arrival)	slope $\hat{\beta}$ (s.e.)	0.0096 (0.0013)	0.0082 (0.0012)	-0.0014
Province vs region aggregation	corr. equal-weight index (weekly)	1.000	0.852	–
Multiple testing (six primary tests)				
Holm correction	tests rejected at 5 % (raw $\rightarrow$ Holm)	4	4	0 lost

Notes: Each row re-estimates the baseline model under the alternative rule and reports the quantity in the second column. Details and all variants are in the replication package (`results/robustness/`).

Exit versus sale. Listing disappearance is an event on a website. Cross-source matching shows that it is consistent (94 percent of matched exits coincide within a week) and that at least a tenth of exits are relistings; the true share of retail sales among exits is unknown and may vary by dealer and vehicle, which would bias the exit hazards if the variation is correlated with covariates. The exit results should be read as describing *inventory turnover on dealer websites*.

Composition and selection. The hedonic sample is the universe of listings on registry sources, not a random sample of Quebec dealers; the registry over-represents the Capitale-Nationale, Bas-Saint-Laurent, Chaudière-Appalaches and Côte-Nord regions relative to population, and under-represents Montréal. Region effects are identified from variation across dealers located in different regions and cannot be separated from dealer effects for single-location dealers; the reweighting exercise addresses composition but not unobserved dealer quality. Left-censored stock is excluded from durations, which removes the vehicles that had already been on the lot longest; the sensitivity analysis with the left-censored stock included (entry set at first observation, a lower bound on listing age) is reported in Table 11.

Reflection and common shocks. The competitor event study cannot distinguish a dealer's response to a neighbour's cut from a common shock that moves both; the same-day concentration, the weekly harmonics, the significant leads and the robustness to dealer-by-weekday effects all point to the latter. An exogenous shifter of competitor repricing (for example, a platform-specific feed outage or a manufacturer programme that affects one brand's dealers only) would be needed for a causal estimate, and we have not found one in this window.

Window length and seasonality. Eleven weeks of late summer and autumn cannot identify seasonal patterns, the response to macroeconomic events (the Bank of Canada raised its policy rate on 28 October 2026, three days before the end of the panel), or the long-run dynamics of vehicle cohorts. Listing durations longer than about 70 days are right-censored by construction, and the plateau of the markdown hazard beyond three weeks is estimated on a shrinking risk set.

Measurement of the index. The liquidity index is validated internally only; its components are mechanically related to one another through the stock of listings (exits and arrivals both enter per listing-day), and the held-out test uses next-week exits, which are themselves observed disappearances. The index should be read as a descriptive summary of turnover on dealer websites until external benchmarks become available.

9.4. Economic and policy implications

Three implications follow with appropriate caution. For dealers, the evidence that cuts are frequent, shallow and increasing in listing age, that overpriced vehicles are cut sooner and more deeply, and that deep cuts characterise slow markets, suggests that entry prices close to hedonic value shorten time on the lot without larger total concessions, a hypothesis that transaction data could test. For consumers and regulators, the regional null means that a Quebec buyer gains little, on average, from searching across regions for the same vehicle once quality is held fixed, although dispersion within remote regions is larger. For statistical agencies, the paper shows that listing panels can deliver a timely, reproducible turnover index at the regional level, provided that connector artefacts (sampled portals, start-up stock, oscillating feeds, flickering connectors) are diagnosed and excluded; the audit trail documented here is a template for that work.

10. Conclusion

Using a daily panel of every listing published by 131 dealer websites, two full-inventory portals and five dealer groups in Quebec over eleven weeks of 2026, we have described how the price of a used car is discovered on the dealer's lot. Asking prices depreciate nonlinearly in age and mileage; one listing in five is cut, by 2.8 percent at the median, with a markdown hazard that doubles over the first three weeks; overpriced vehicles are cut sooner and more deeply; and listings disappear after nine days at the median, faster for cheap, high-mileage and older vehicles. Repricing co-moves across neighbouring dealers on the same day and at weekly harmonics but not with a lag, which we interpret as common shocks and shared routines rather than strategic reaction. Quality-adjusted asking prices differ across the province's regions by at most a few percent, and not detectably once trim is held fixed. A transparent five-component liquidity index summarises more than half of the common variation in regional turnover, and its non-exit components anticipate next week's exits under the preregistered standardisation, less clearly under alternatives; the data also tell us that deep markdowns are a sign of slow, not liquid, markets.

The limits are as important as the findings. Asking prices are not transaction prices; exits overstate retail sales by at least a tenth, as our relisting analysis shows; the competitor event study is descriptive; the regional null applies to one quarter; the thirty-day predictability of repricing is mostly dealer practice and does not transfer to unseen dealers; and the liquidity index has no external validation yet. A longer panel, matched to registration transfers, would allow the markdown and exit hazards to be modelled as competing risks with a verified sale outcome, the regional null to be tested across seasons, and the index to be tested as a leading indicator. The pipeline that produced this paper, with its immutable snapshot, exclusion log, preregistered plan and claims ledger, is designed to make that extension a matter of rerunning it.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge, "When Should You Adjust Standard Errors for Clustering?," *Quarterly Journal of Economics*, 2023, *138* (1), 1–35.

Adams, Christopher P., Laura Hosken, and Peter W. Newberry, "Vettes and lemons on eBay," *Quantitative Marketing and Economics*, 2011, *9* (2), 109–127.

Aguirregabiria, Victor, "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies*, 1999, *66* (2), 275–308.

Akerlof, George A., "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 1970, *84* (3), 488.

Allison, Paul D., "Discrete-Time Methods for the Analysis of Event Histories," *Sociological Methodology*, 1982, *13*, 61.

Amihud, Yakov, "Illiquidity and stock returns: cross-section and time-series effects," *Journal of Financial Markets*, 2002, *5* (1), 31–56.

_ **and Haim Mendelson**, "Asset pricing and the bid-ask spread," *Journal of Financial Economics*, 1986, *17* (2), 223–249.

Anenberg, Elliot, "Information Frictions and Housing Market Dynamics," *International Economic Review*, 2016, *57* (4), 1449–1479.

_ **and Steven Laufer**, "A More Timely House Price Index," *Review of Economics and Statistics*, 2017, *99* (4), 722–734.

Anglin, Paul M., Ronald Rutherford, and Thomas M. Springer, "The Trade-off Between the Selling Price of Residential Properties and Time-on-the-Market: The Impact of Price Setting," *Journal of Real Estate Finance and Economics*, 2003, *26* (1), 95–111.

Angrist, Joshua D., "The perils of peer effects," *Labour Economics*, 2014, *30*, 98–108.

Anselin, Luc, *Spatial Econometrics: Methods and Models*, Springer Netherlands, 1988.

Athey, Susan and Guido W. Imbens, "Machine Learning Methods That Economists Should Know About," *Annual Review of Economics*, 2019, *11* (1), 685–725.

AutoTrader.ca, "AutoTrader Price Index: June 2024," Editorial report, 22 July 2024 2024.

Bajari, Patrick and C. Lanier Benkard, "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach," *Journal of Political Economy*, 2005, *113* (6), 1239–1276.

Baye, Michael R., John Morgan, and Patrick Scholten, "Information, Search, and Price Dispersion," in Terrence Hendershott, ed., *Handbook on Economics and Information Systems*, Vol. 1, Elsevier, 2006, pp. 323–375.

Benjamini, Yoav and Yosef Hochberg, "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 1995, *57* (1), 289–300.

Bertrand, M., E. Duflo, and S. Mullainathan, "How Much Should We Trust Differences-In-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Biglaiser, Gary, Fei Li, Charles Murry, and Yiyi Zhou, "Intermediaries and product quality in used car markets," *RAND Journal of Economics*, 2020, *51* (3), 905–933.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Yannis M. Ioannides, "Identification of Social Interactions," in Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, eds., *Handbook of Social Economics*, Vol. 1, North-Holland, 2011, pp. 853–964.

Bond, Eric W., "A Direct Test of the "Lemons" Model: The Market for Used Pickup Trucks," *American Economic Review*, 1982, *72* (4), 836–840.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin, "Identification of peer effects through social networks," *Journal of Econometrics*, 2009, *150* (1), 41–55.

Brier, Glenn W., "Verification of Forecasts Expressed in Terms of Probability," *Monthly Weather Review*, 1950, *78* (1), 1–3.

Brock, W. A. and S. N. Durlauf, "Discrete Choice with Social Interactions," *Review of Economic Studies*, 2001, *68* (2), 235–260.

Burdett, Kenneth and Kenneth L. Judd, "Equilibrium Price Dispersion," *Econometrica*, 1983, *51* (4), 955.

Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer, "$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions," *American Economic Review*, 2006, *96* (4), 1253–1270.

Byrne, David P. and Nicolas de Roos, "Learning to Coordinate: A Study in Retail Gasoline," *American Economic Review*, 2019, *109* (2), 591–619.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

_ , _ , and _ , "Robust Inference With Multiway Clustering," *Journal of Business & Economic Statistics*, 2011, *29* (2), 238–249.

Canadian Black Book, "Used Vehicle Retention Index," Web page, accessed 1 November 2026 2026.

Carrillo, Paul E., "An Empirical Stationary Equilibrium Search Model of the Housing Market," *International Economic Review*, 2012, *53* (1), 203–234.

Ceglowski, Janet, "The law of one price: intranational evidence for Canada," *Canadian Journal of Economics*, 2003, *36* (2), 373–400.

Chandra, Ambarish and Mariano Tappata, "Consumer search and dynamic price dispersion: an application to gasoline markets," *RAND Journal of Economics*, 2011, *42* (4), 681–704.

Conley, T.G., "GMM estimation with cross sectional dependence," *Journal of Econometrics*, 1999, *92* (1), 1–45.

Court, Andrew T., "Hedonic Price Indexes with Automotive Examples," in "The Dynamics of Automobile Demand," New York: General Motors Corporation, 1939. Scan hosted by A. C. Goodman (Wayne State University); see also Goodman (1998).

Cox Automotive, "Used-Vehicle Days' Supply Breaks 50; Average Listing Price Surpasses $28,000," Market Insights, 20 January 2022 (revised 17 February 2022) 2022.

Cox, D. R., "Regression Models and Life-Tables," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 1972, *34* (2), 187–202.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Diamond, Peter A, "A model of price adjustment," *Journal of Economic Theory*, 1971, *3* (2), 156–168.

Diebold, Francis X. and Roberto S. Mariano, "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, 1995, *13* (3), 253–263.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen, "Over-the-Counter Markets," *Econometrica*, 2005, *73* (6), 1815–1847.

Ekeland, Ivar, James J. Heckman, and Lars Nesheim, "Identification and Estimation of Hedonic Models," *Journal of Political Economy*, 2004, *112* (S1), S60–S109.

Elmaghraby, Wedad and Pınar Keskinocak, "Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions," *Management Science*, 2003, *49* (10), 1287–1309.

Emons, Winand and George Sheldon, "The market for used cars: new evidence of the lemons phenomenon," *Applied Economics*, 2009, *41* (22), 2867–2885.

Engel, Charles and John H. Rogers, "How Wide Is the Border?," *American Economic Review*, 1996, *86* (5), 1112–1125.

Engers, Maxim, Monica Hartmann, and Steven Stern, "Annual miles drive used car prices," *Journal of Applied Econometrics*, 2009, *24* (1), 1–33.

Fine, Jason P. and Robert J. Gray, "A Proportional Hazards Model for the Subdistribution of a Competing Risk," *Journal of the American Statistical Association*, 1999, *94* (446), 496–509.

Friedman, Jerome H., "Greedy function approximation: A gradient boosting machine.," *Annals of Statistics*, 2001, *29* (5), 1189–1232.

Gallego, Guillermo and Garrett van Ryzin, "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons," *Management Science*, 1994, *40* (8), 999–1020.

Gavazza, Alessandro, "The Role of Trading Frictions in Real Asset Markets," *American Economic Review*, 2011, *101* (4), 1106–1143.

_ , Alessandro Lizzeri, and Nikita Roketskiy, "A Quantitative Analysis of the Used-Car Market," *American Economic Review*, 2014, *104* (11), 3668–3700.

Genesove, D. and C. Mayer, "Loss Aversion and Seller Behavior: Evidence from the Housing Market," *Quarterly Journal of Economics*, 2001, *116* (4), 1233–1260.

Genesove, David, "Adverse Selection in the Wholesale Used Car Market," *Journal of Political Economy*, 1993, *101* (4), 644–665.

Giacomini, Raffaella and Halbert White, "Tests of Conditional Predictive Ability," *Econometrica*, 2006, *74* (6), 1545–1578.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Gorodnichenko, Yuriy and Linda L Tesar, "Border Effect or Country Effect? Seattle May Not Be so Far from Vancouver After All," *American Economic Journal: Macroeconomics*, 2009, *1* (1), 219–241.

Grambsch, Patricia M. and Terry M. Therneau, "Proportional hazards tests and diagnostics based on weighted residuals," *Biometrika*, 1994, *81* (3), 515–526.

Griliches, Zvi, "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change," in Price Statistics Review Committee, ed., *The Price Statistics of the Federal Government*, New York: National Bureau of Economic Research, 1961.

Guren, Adam M., "House Price Momentum and Strategic Complementarity," *Journal of Political Economy*, 2018, *126* (3), 1172–1218.

Haan, Marco A. and Henk-Wim de Boer, "Has the Internet Eliminated Regional Price Differences? Evidence from the Used Car Market," *De Economist*, 2010, *158* (4), 373–386.

Han, Lu and William C. Strange, "What is the role of the asking price for a house?," *Journal of Urban Economics*, 2016, *93*, 115–130.

Haurin, Donald R., Jessica L. Haurin, Taylor Nadauld, and Anthony Sanders, "List Prices, Sale Prices and Marketing Time: An Application to U.S. Housing Markets," *Real Estate Economics*, 2010, *38* (4), 659–685.

Head, Allen, Huw Lloyd-Ellis, and Hongfei Sun, "Search, Liquidity, and the Dynamics of House Prices and Construction," *American Economic Review*, 2014, *104* (4), 1172–1210.

Heckman, J. and B. Singer, "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data," *Econometrica*, 1984, *52* (2), 271.

Hendel, Igal and Alessandro Lizzeri, "Adverse Selection in Durable Goods Markets," *American Economic Review*, 1999, *89* (5), 1097–1115.

Hill, Robert J., "Hedonic Price Indexes for Residential Housing: A Survey, Evaluation and Taxonomy," *Journal of Economic Surveys*, 2012, *27* (5), 879–914.

Hotelling, Harold, "Stability in Competition," *Economic Journal*, 1929, *39* (153), 41.

Houde, Jean-François, "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline," *American Economic Review*, 2012, *102* (5), 2147–2182.

iSeeCars.com, "iSeeCars.com Reveals How Used Car Shoppers Can Get a Better Deal by Waiting," Study of 30 million used-car listings; press release 2013.

Israeli, Ayelet, Fiona Scott Morton, Jorge Silva-Risso, and Florian Zettelmeyer, "How Market Power Affects Dynamic Pricing: Evidence from Inventory Fluctuations at Car Dealerships," *Management Science*, 2022, *68* (2), 895–916.

Kalnins, Arturs, "Hamburger Prices and Spatial Econometrics," *Journal of Economics & Management Strategy*, 2003, *12* (4), 591–616.

Kaplan, Greg and Guido Menzio, "The Morphology of Price Dispersion," *International Economic Review*, 2015, *56* (4), 1165–1206.

Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu, "LightGBM: A Highly Efficient Gradient Boosting Decision Tree," in "Advances in Neural Information Processing Systems 30 (NIPS 2017)" 2017.

Kiefer, Nicholas M., "Economic Duration Data and Hazard Functions," *Journal of Economic Literature*, 1988, *26* (2), 646–679.

Kim, Jae-Cheol, "The Market for "Lemons" Reconsidered: A Model of the Used Car Market with Asymmetric Information," *American Economic Review*, 1985, *75* (4), 836–843.

Knight, John R., "Listing Price, Time on Market, and Ultimate Selling Price: Causes and Effects of Listing Price Changes," *Real Estate Economics*, 2002, *30* (2), 213–237.

Krainer, John, "A Theory of Liquidity in Residential Real Estate Markets," *Journal of Urban Economics*, 2001, *49* (1), 32–53.

Lach, Saul, "Existence and Persistence of Price Dispersion: An Empirical Analysis," *Review of Economics and Statistics*, 2002, *84* (3), 433–444.

Lagos, Ricardo and Guillaume Rocheteau, "Liquidity in Asset Markets with Search Frictions," *Econometrica*, 2009, *77* (2), 403–426.

Lancaster, Kelvin J., "A New Approach to Consumer Theory," *Journal of Political Economy*, 1966, *74* (2), 132–157.

Lancaster, Tony, *The Econometric Analysis of Transition Data*, Cambridge University Press, 1990.

Larsen, Bradley J, "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions," *Review of Economic Studies*, 2021, *88* (2), 851–882.

Lazear, Edward P., "Retail Pricing and Clearance Sales," *American Economic Review*, 1986, *76* (1), 14–32.

Lewis, Gregory, "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors," *American Economic Review*, 2011, *101* (4), 1535–1546.

Lin, D. Y. and L. J. Wei, "The Robust Inference for the Cox Proportional Hazards Model," *Journal of the American Statistical Association*, 1989, *84* (408), 1074–1078.

Lyons, Ronan C., "Can list prices accurately capture housing price trends? Insights from extreme markets conditions," *Finance Research Letters*, 2019, *30*, 228–232.

Manheim Consulting (Cox Automotive), "Summary Methodology for Manheim Used Vehicle Value Index," Technical note, January 2023 rebasing 2023.

Manski, Charles F., "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 1993, *60* (3), 531.

Morton, Fiona Scott, Florian Zettelmeyer, and Jorge Silva-Risso, "Internet Car Retailing," *Journal of Industrial Economics*, 2001, *49* (4), 501–519.

Mullainathan, Sendhil and Jann Spiess, "Machine Learning: An Applied Econometric Approach," *Journal of Economic Perspectives*, 2017, *31* (2), 87–106.

Murry, Charles and Yiyi Zhou, "Consumer Search and Automobile Dealer Colocation," *Management Science*, 2020, *66* (5), 1909–1934.

Ngai, L. Rachel and Silvana Tenreyro, "Hot and Cold Seasons in the Housing Market," *American Economic Review*, 2014, *104* (12), 3991–4026.

Ohta, Makoto and Zvi Griliches, "Automobile Prices Revisited: Extensions of the Hedonic Hypothesis," in Nestor E. Terleckyj, ed., *Household Production and Consumption*, National Bureau of Economic Research, 1976.

Pakes, Ariel, "A Reconsideration of Hedonic Price Indexes with an Application to PC's," *American Economic Review*, 2003, *93* (5), 1578–1596.

Pashigian, B. P. and B. Bowen, "Why Are Products Sold on Sale?: Explanations of Pricing Regularities," *Quarterly Journal of Economics*, 1991, *106* (4), 1015–1038.

Pesendorfer, Martin, "Retail Sales: A Study of Pricing Behavior in Supermarkets," *Journal of Business*, 2002, *75* (1), 33–66.

Peterson, Jonathan R. and Henry S. Schneider, "Adverse selection in the used-car market: evidence from purchase and repair patterns in the Consumer Expenditure Survey," *RAND Journal of Economics*, 2014, *45* (1), 140–154.

Pinkse, Joris, Margaret E. Slade, and Craig Brett, "Spatial Price Competition: A Semiparametric Approach," *Econometrica*, 2002, *70* (3), 1111–1153.

Putter, H., M. Fiocco, and R. B. Geskus, "Tutorial in biostatistics: competing risks and multi-state models," *Statistics in Medicine*, 2007, *26* (11), 2389–2430.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Romano, Joseph P. and Michael Wolf, "Stepwise Multiple Testing as Formalized Data Snooping," *Econometrica*, 2005, *73* (4), 1237–1282.

Rosen, Sherwin, "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, 1974, *82* (1), 34–55.

Roth, Jonathan, "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 2022, *4* (3), 305–322.

_ , Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Royston, Patrick and Mahesh K. B. Parmar, "Flexible parametric proportional-hazards and proportional-odds models for censored survival data, with application to prognostic modelling and estimation of treatment effects," *Statistics in Medicine*, 2002, *21* (15), 2175–2197.

Salop, Steven C., "Monopolistic Competition with Outside Goods," *Bell Journal of Economics*, 1979, *10* (1), 141.

Singer, Judith D. and John B. Willett, "It's About Time: Using Discrete-Time Survival Analysis to Study Duration and the Timing of Events," *Journal of Educational Statistics*, 1993, *18* (2), 155–195.

Sorensen, Alan T., "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," *Journal of Political Economy*, 2000, *108* (4), 833–850.

Stigler, George J., "The Economics of Information," *Journal of Political Economy*, 1961, *69* (3), 213–225.

Stolyarov, Dmitriy, "Turnover of Used Durables in a Stationary Equilibrium: Are Older Goods Traded More?," *Journal of Political Economy*, 2002, *110* (6), 1390–1413.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Sweeting, Andrew, "Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets," *Journal of Political Economy*, 2012, *120* (6), 1133–1172.

Taylor, Lance, Taylor Mitchell, and Gerry O'Donnell, "Measuring Price Change for Used Vehicles in the Canadian Consumer Price Index," Prices Analytical Series 62F0014M2022008, Statistics Canada 2022.

Therneau, Terry M. and Patricia M. Grambsch, *Modeling Survival Data: Extending the Cox Model*, Springer New York, 2000.

Triplett, Jack, *Handbook on Hedonic Indexes and Quality Adjustments in Price Indexes*, Paris: OECD Publishing, 2006.

Varian, Hal R., "A Model of Sales," *American Economic Review*, 1980, *70* (4), 651–659.

Yanes, Ana, "Depreciation of Light Duty Road Vehicles: Evidence from the Canadian Red Book," in "Proceedings of the 45th Annual Conference of the Canadian Transportation Research Forum" Toronto 2010. Transport Canada.

Zettelmeyer, Florian, Fiona Scott Morton, and Jorge Silva-Risso, "How the Internet Lowers Prices: Evidence from Matched Survey and Automobile Transaction Data," *Journal of Marketing Research*, 2006, *43* (2), 168–181.

Appendices

A. Data-generating process and quality audit

Synchronisation and exit rule. Each cycle re-reads every source; a listing absent from its source for two consecutive cycles is declared inactive by the platform. The research definition of exit (seven days of absence before the snapshot) is stricter: the per-source "reactivation share" (exit events later reversed, divided by exit events) is 0.44 for Otogo, 0.04 for AutoUsagée, below 0.1 for most dealer websites, and 0.52–0.94 for the four connectors excluded from the dynamic sample. Listings whose last observation falls in the final seven days of the window are right-censored.

Sustained price paths. Of 428,484 price-log rows, 2,257 record a withdrawn price; 121,239 are price changes, of which 11,479 occur less than 24 hours after the previous change. The transient rule (a level shorter than 48 hours that reverts to the preceding level is deleted) removes the alternating-feed artefacts illustrated by a listing that alternated between $37,840 and $40,794 twice a day for four weeks; 70 listings with three or more transients are excluded from markdown analyses. Robustness to a 24-hour and a 72-hour rule is reported in Section 9.

Identifier stability. Within Otogo, 1,059 of 63,821 VINs (1.7 percent) appear under more than one listing identifier; inspection shows these are the same vehicle re-listed by a different dealer after a transfer, not identifier churn. Dealer websites assign stable identifiers.

Governance. The snapshot excludes the platform's user table; VINs, URLs, images, free text and exact coordinates are retained only in the restricted raw and interim layers, never in released tables, which carry hashed listing and dealer identifiers and coordinates rounded to 0.01 degree.

B. Preregistration and amendments

The research questions and the analysis plan (v1.0, 2026-11-01) were written after the data audit and before estimation. Amendments recorded during execution: (D-010) make–model and make–model–trim quality fixed effects instead of make–model–year, because model-year effects absorb vehicle age within an 81-day window; (D-011) linear-probability estimation of the competitor event study with high-dimensional fixed effects; (D-014) per-source left-censoring after the portal start-up stock (33,224 listings) was found to be treated as entries; (D-015) exclusion of panel weeks 0–1 from the liquidity index. The complete decision log, experiment log and claims ledger are distributed with the replication package.

C. Additional estimates

Additional coefficient tables (full spline bases, listing-age bins, powertrain and body effects, the log-logistic accelerated-failure-time model, the 10-km and 50-km competitor sets, the dealer-by-weekday specifications, LightGBM feature importances and calibration tables) are provided as machine-readable files in the replication package (`results/`).