Discussion of "Regulating Inaction: The Case of Price Walking"

Marco Cosconati
IVASS, Bank of Italy

October 22, 2026
Chicago Fed

Disclaimer: The views expressed are solely those of the author and do not represent those of IVASS or the Bank of Italy.

Why this paper matters for policy

Large price walking, large distributional stakes

+5.4%

paid premium per year of tenure

Preferred IV estimate

−33%

average paid premiums, 2013–2021

Same estimation sample

Price walking can coexist with falling average premiums. Policies aimed at containing premiums can therefore have large distributional consequences across customers.

Policy relevance

Price walking and differential pricing are already part of the Italian and European policy debate.

This paper shows how insurers adapt when renewal pricing is constrained and is therefore likely to be highly relevant in policy circles.

Source: Cosconati, *Loyal but Not Captive: Price Walking and Consumer Switching in Auto Insurance* (2026).

My reading of the paper

Main message: a ban on direct price walking does not eliminate the value of customer inertia. It changes the instruments firms use to monetize it.

- Clean theory: linking new and renewal prices makes products with large inherited back books less competitive.
- Rich data: complete quote menus, product rankings, product availability, click-through choices, and repeated searches.
- Main empirical contribution: strong evidence of firms' supply responses inside the comparison-site market.
- My discussion: what exactly is identified, and how far the evidence travels toward market-wide incidence and equilibrium outcomes.

Five predictions: discussion roadmap

Prediction	Empirical result	Discussion question
1. Inertia	Likely passive customers lose introductory discounts.	Is measured inertia stable across policy regimes, and does the platform sample change?
2. Back books	Large-back-book products become less competitive.	How is the back-book response implemented across customer types?
3–4. Dynamics	Product shares oscillate; large-back-book firms introduce products.	Do new-product gains offset legacy-product losses within the same insurer?
5. Search gains	The old product remains far above the cheapest available quote.	When does a quote-menu gap identify a renewal-price penalty?

Focus: search and sample composition; heterogeneous supply responses; and the mapping from web quotes to renewal payments.

Theory ingredients

A clean back-book mechanism

Environment

- Two symmetric Hotelling firms.
- Consumers live for two periods.
- All consumers search when new.
- At renewal, share λ becomes passive.
- Firms face demand shocks Δ_t.

y_t = lagged market-share advantage.

Regulation

- Before: firms set new and renewal prices separately.

$$p_t^i \neq r_t^i$$

- After: same product, same price.

$$p_t^i = r_t^i$$

- A large back book makes discounting costly.

$$p_R^0 = c(m + a\Delta_t + by_t), \qquad b > 0.$$

Useful abstraction: firms are symmetric; there is no heterogeneous risk-rating precision, cost structure, or pricing algorithm.

Theory: what does firm symmetry buy?

The back-book mechanism is robust. The sharper aggregate predictions rely more on the assumption that insurers are symmetric.

Most robust

- Back-book pricing
- Product proliferation

$$BB_{jt} \uparrow \Rightarrow \text{less aggressive pricing}$$

Heterogeneous response

- Teaser discounts
- Search gains

$$b \longrightarrow b_j$$

Most symmetry-dependent

- Oscillating shares
- Firm-level aggregation

$$y_{t+1} = -by_t + \left(\frac{1}{2} - a\right)\Delta_t$$

Why this matters: in my ReStud paper, firms differ substantially in risk-rating precision, cost structures, and pricing strategies. With this heterogeneity, the same regulation may produce different responses across insurers.

Suggested extension: estimate the five predictions by insurer or insurer group. Are the aggregate effects broad-based, or driven by a subset of business models?

A third margin: risk may respond to the ban

Insurance prices combine markups and expected claims

$$\text{quoted price} = \underbrace{\text{expected claims cost}}_{\text{risk, effort, sorting}} + \underbrace{\text{markup}}_{\text{back book, competition}}$$

Benchmark in the paper

- The model focuses on inertia and supplier pricing.
- Risk and expected claims costs are essentially held fixed.
- Price changes are interpreted mainly as changes in competitive pricing and back-book harvesting.

Additional insurance channel

- Price walking and switching can weaken accident penalties.
- If penalties change, driving incentives and claims may change.
- Switching can also change the risk composition of legacy products.

Implication: the observed price response may combine strategic repricing and changes in expected costs. This matters most for price levels and welfare; product proliferation remains direct evidence of supplier adaptation.

Concern 1: Does the score describe search propensity?

Does the same X_i predict the same behavior across policy regimes?

Among customers who initially search and select a product, the paper estimates

$$\widehat{I}_i = \widehat{\Pr}(\text{no platform return around } t+12 \mid X_i)\,.$$

Role of the score

- It ranks customer types by their predicted likelihood of not returning to Compare the Market.
- Prediction 1 compares quote changes across these predicted-inertia groups.
- The interpretation is that a higher score identifies customers more likely to become passive.

Measurement concern

The reform changes prices, product menus, and the gains from search. Hence,

$$X_i \longmapsto \Pr(\text{no platform return} \mid X_i)$$

may differ before and after the policy.

- A pooled score may lose calibration.
- If responses differ across X_i, it may also change the ranking of customer types.

Suggested check: estimate the score using pre-policy cohorts with complete follow-up, freeze the groups, and test their calibration and ranking after the reform.

Concern 2: Is search propensity policy invariant?

A fixed score does not hold consumers' subsequent behavior fixed

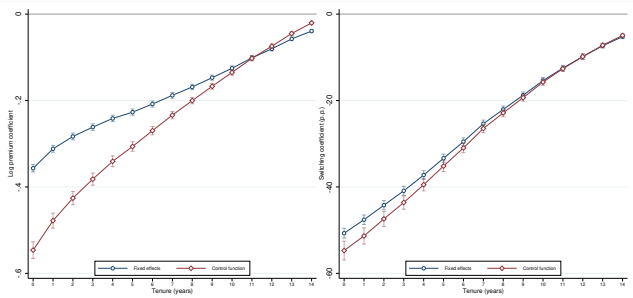

Cosconati: tenure profiles from fixed-effects and control-function specifications.

Economic concern

- The model treats the passive share

$$\lambda$$

  as fixed.
- The reform may change whether consumers search, not only which insurer an active shopper chooses.
- Renewal prices, quote dispersion, and product proliferation all affect the value of searching.

Italian motivation

Premiums and switching rise in the reported conditional tenure profiles.
This illustrates that observed passivity need not be a fixed consumer trait.

Suggested check: using fixed pre-policy score groups, plot return-search rates by renewal cohort around the reform. The Italian figure motivates the channel; it does not identify the UK policy effect.

Concern 3: Does the sample of searchers change?

The policy may alter who appears in the observed quote data

$$\mathcal{S}_t = \{i : \text{consumer } i \text{ submits an enquiry on Compare the Market at } t\}.$$

The paper observes menus and choices for consumers in $\mathcal{S}_t$, not for the full population of policyholders.

Why composition may change

- The reform may change platform participation and the choice of search channel.
- Consumers may substitute between comparison sites, direct insurers, and brokers.
- Pre- and post-policy enquiries may differ even within the same observable X_i groups.

Checks around the reform

- Event-study balance tests for customer, vehicle, coverage, and driving-history characteristics.
- Reweight post-policy enquiries to the pre-policy observable distribution.
- Compare discretionary with less discretionary searches: new car, new licence, or first policy.
- Track composition within fixed inertia-score groups.

Observed quote change = firm response + change in the composition of platform users. Reweighting addresses observable composition, not all selection.

Prediction 2: How is back-book pricing implemented?

Customer groups and the empirical test

$$\theta(i) = \text{birth year} \times \text{car-value quintile} \times \text{coverage level}.$$

$BB_{\theta(i)jt}$ uses earlier platform purchases of product j by customers in shopper i's group, proxying its inherited book.

$$Top_{ijt} = \gamma BB_{\theta(i)jt} + \rho\big(Post_t \times BB_{\theta(i)jt}\big) + FE + \varepsilon_{ijt}.$$

i: enquiry; j: firm-specific product; t: date.
$Top = 1$: cheapest quote; $Post = 1$: after the ban; FE: fixed effects.

Finding: $\rho < 0$. Larger same-group back books are associated with a relative loss of quote competitiveness after the ban.

A possible implementation

A low new-business quote also constrains renewal prices for equivalent customers. Groups with larger inherited books therefore have more renewal revenue to protect.

An insurer may raise the pricing weights on characteristics prevalent in those groups—for example, charging older drivers more when many renewers are older. Equivalent new older customers also face higher quotes.

Extension: does inherited-book composition predict post-ban changes in pricing coefficients?

Appendix: Italian direct-channel estimates show that age effects on tariffs and paid premiums change over time.

Predictions 3–4: From product cycling to proliferation

Strong corroborating evidence for the supply-side mechanism

product wins new customers $\Rightarrow$ larger inherited book $\Rightarrow$ less competitive next year

$\Downarrow$

introduce a new product with little or no inherited book $\Rightarrow$ compete again for new customers.

What the paper shows

- After the ban, product-share persistence turns negative.
- Previously successful products become less competitive as their back books grow.
- Firms with larger back books introduce more products.

A natural within-insurer extension

- Separate each insurer's platform share into legacy and newly introduced products.
- Ask whether new-product gains offset legacy-product losses.
- This reveals whether product proliferation preserves the insurer's overall ability to attract new business.

Question for the authors:

Do new-product gains offset legacy-product losses within the same insurer, or does platform business shift to competitors?

Prediction 5: From repeated menus to a 40% gap

Repeated quote menus; the passive-renewal offer is not observed

The same customer is observed again only if she returns to Compare the Market. Between searches, contracts, payments, and off-platform choices are not observed.

Observed in the two quote menus

Year 1: customer i selects product j_i^0.

Year 2: if i returns and j_i^0 reappears, observe

$$q_{i2}^{old} = \text{current web quote for } j_i^0,$$

$$q_{i2}^{min} = \text{cheapest quote in the Year-2 menu.}$$

The Year-2 purchase is not required.

Quote gap versus renewal penalty

$$G_i^Q = \frac{q_{i2}^{old} - q_{i2}^{min}}{q_{i2}^{min}}.$$

G_i^Q is directly an observed quote-menu gap.

The passive-renewal offer is not observed:

$$r_{i2} = \text{unobserved renewal offer.}$$

Reading G_i^Q as a renewal-price penalty requires

$$q_{i2}^{old} \simeq r_{i2}.$$

Italian benchmark: in direct-channel contracts, how close is the scheduled tariff T to the premium actually paid P, where $T = P + D$?

The full enquiry structure and the mapping from model objects to data are in the appendix.

Scheduled price versus premium paid at renewal

Italian direct-channel evidence

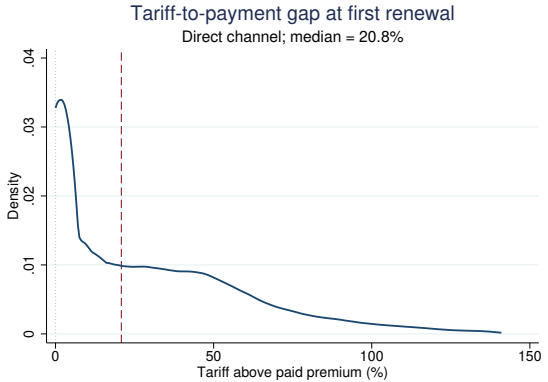

The median tariff is 20.8% above the premium paid. Italian tariffs are not UK web quotes: this is an external benchmark, not a direct test of $q_{i2}^{old} = r_{i2}$.

Concluding remarks

What the paper establishes

- Regulation changes firms' equilibrium quote and menu strategies.
- Products with large back books become less competitive after implementation.
- Back-book pricing, oscillating product shares, and product proliferation form a coherent supply-side response.

What remains to pin down

- Is the inertia ranking—and the pool of platform searchers—stable across policy regimes?
- Are responses broad-based across heterogeneous insurers, and how are back-book and product responses implemented across customer types?
- Do observed quotes map into renewal payments and changes in expected claims costs?

Bottom line: a very good paper with convincing evidence on how suppliers adapt. The remaining empirical task is to map menu responses into who searches, which firms respond, and what consumers actually pay.

Policy lesson: restricting tenure-based pricing within a product may shift segmentation across products, permitted pricing variables, and search channels.

Additional evidence and data details

Material for the authors and for questions

Data structure: repeated searches, not contract histories

Repeated search behavior is the empirical basis for the inertia score

One enquiry e is one form submission by customer i at date t_e, specifying a driver–vehicle–coverage profile X_{ie}.

$$\underbrace{(i, e, X_{ie})}_{\text{search request}} \longrightarrow \underbrace{\text{insurers } f}_{\text{many firms}} \longrightarrow \underbrace{\text{products } k}_{\text{possibly several per firm}} \longrightarrow \underbrace{\mathcal{M}_{ie} = \{(f, k, q^{web}_{iefk})\}}_{\text{quote menu}}.$$

Within one enquiry

- One customer profile X_{ie}
- Many insurer–product quotes
- Median menu: **64 products from 26 insurers**
- If the customer clicks through, observe

$$(f^*_{ie}, k^*_{ie})$$

Across enquiries

- Customer identifiers link repeated searches:

$$e = 1, \ldots, E_i.$$

- 64% appear in multiple years
- 53% search multiple times within a year
- Between searches, contracts, payments, and off-platform choices are not observed

Why this matters for Prediction 1: after a customer purchases at t, whether she returns around $t + 12$ is used to predict which customer types are more likely to be passive.

Data: directly observed vs. constructed

Directly observed

For enquiry e by customer i:

$$\mathcal{M}_{ie} = \left\{ (f, k, q^{web}_{iefk}) \right\},$$

where f is the insurer and k is one of its products.

- Driver, vehicle, and requested coverage
- Complete menu of web quotes
- Product availability and quote ranking
- Insurer–product selected when the customer clicks through
- Further enquiries when the customer returns

A rich snapshot of offers to the same shopper.

Constructed or inferred

- **New-business price paid:** the selected web quote is observed, but there is no separate administrative record of the issued premium.
- **Customer inertia:** predicted from whether a customer who purchased fails to return about 12 months later.
- **Product back book:** earlier platform purchases of the same product by similar customers; not the insurer's full back book.
- **Product market share:** share of observed click-through choices on the platform, not total market-wide purchases.
- **Renewal price:** not observed in the main data; Section 5.5 uses the current web quote for last year's product as a proxy.

Scope: direct evidence on offers, menus, and product responses conditional on searching on Compare the Market. Passivity, full back books, market-wide shares, and renewal incidence are measured more indirectly.

Italian evidence: pricing loadings can change over time

Direct-channel new customers; 2013 normalized to zero

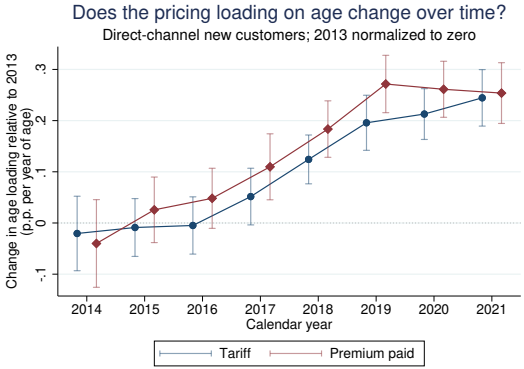

What the figure shows

- The conditional loading on age changes substantially over time.
- By 2019–2021, it is roughly 0.20–0.27 p.p. higher per year of age than in 2013.
- Tariff and paid-premium loadings move together.

Interpretation

This illustrates flexibility in the pricing map. It is not causal evidence that the UK ban induced proxying: age is also a genuine risk factor.

Product proliferation is a visible response; changing pricing coefficients may be a less visible one.

Discounting at entry and first renewal

A descriptive comparison in Italian direct-channel contracts

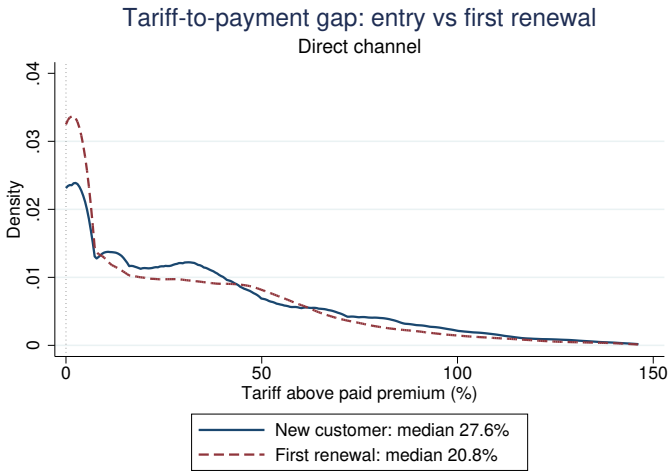

What does the Italian benchmark tell us?

The mapping from scheduled prices to payments is not automatic

- Even in the direct channel, the pre-discount tariff is often different from the premium actually paid.
- At first renewal, the median tariff is **20.8% above** the paid premium.
- At entry, the corresponding raw median gap is **27.6%**.
- Discounts are quantitatively important at both stages; the entry-versus-renewal comparison is descriptive.
- This does **not** directly test $q_{i2}^{old} = r_{i2}$: an Italian pre-discount tariff is not a UK consumer-facing web quote.

Implication for the paper: the 40% statistic is a clean quote-menu gap. Reading it as a price actually paid by passive renewers requires an additional empirical assumption. The core back-book and product-proliferation results are unaffected.