Journal of Revenue and Pricing Management
https://doi.org/10.1057/s41272-026-00604-6

RESEARCH

Willingness-to-pay estimation for airline ancillary seats under censored price exposure

Kevin Duijndam[1,2] · **Ger Koole**[1] · **Tim von Lücken**[1,2] · **Rob Van Der Mei**[1,3]

Received: 11 February 2026 / Accepted: 15 June 2026

Abstract

Ancillary seat pricing requires estimating willingness-to-pay (WTP) from transaction logs in which offered prices are unobserved for non-purchases and operationally distorted for purchases. We develop a two-stage approach. *Synthetic Offer Reconstruction* (SOR) reverses known loyalty and operational pricing rules around the filed base price to recover a price-exposure feature for every passenger. A *Proximal Causal Inference* (PCI) refinement then estimates the price-purchase relationship with a cross-fitted two-stage sieve bridge, using operational variables as upstream proxies and leave-one-out within-flight statistics as downstream proxies for the latent exposure regime. We evaluate on synthetic airline data containing an unobserved flight-level confounder, over 100 random seeds, against an oracle restricted to the same price grid, alongside qualitative evidence from real operational airline data supporting the reconstruction and proxy choices. When training and deployment regimes coincide, the bridge matches a direct gradient-boosted baseline. When the confounder distribution shifts between training and deployment, it recovers 90–93% of oracle revenue versus roughly 86% for the baseline and an inverse-probability-weighting benchmark. We further document a *calibration paradox*: supplying the same proxies to a flexible classifier as ordinary covariates attains the best observed calibration on held-out data, yet collapses to a near-fixed-price policy capturing only 74–84% of oracle revenue in the synthetic evaluation. Structural use of proxy information through the bridge's moment condition - not proxy availability itself - is what yields transportable WTP estimates under missing-not-at-random price exposure.

keywords Dynamic Pricing · Willingness-to-pay Estimation · Airline Ancillary Pricing

Introduction

Airlines increasingly monetize seat attributes (legroom, location) as paid ancillaries. For such products, demand is heterogeneous and inventory perishes at departure, which increases the importance of setting the right price. In current practice, base seat prices are usually defined by destination and travel month, through industry-standard fare tables filed with the Airline Tariff Publishing Company (ATPCO) (see section Background: seat pricing and ATPCO for more detail). After these prices are set, they are adjusted by all kinds of rules: there can be discounts in the last days before departure to try to fill remaining open seats; potentially large last-minute markdowns at check-in when other economy sections are too full; and regular Frequent Flyer (FF) discounts (usually higher-tier levels of FF programs get a 10-50% discount). Besides, airlines frequently partner with other airlines, and cross-sales effects (e.g., joint-venture bookings) further perturb observed prices. Consequently,

✉ Kevin Duijndam
kevin.duijndam@klm.com

Ger Koole
ger.koole@vu.nl

Tim von Lücken
Tim-von.Luecken@klm.com

Rob Van Der Mei
mei@cwi.nl

1 Department of Mathematics, Section Stochastics, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

2 KLM Royal Dutch Airlines, Amstelveen, The Netherlands

3 Department of Probability and Stochastic Networks, Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

Published online: 02 July 2026

price exposure has higher uncertainty than in many more standard retail settings. Usually, logs register accepted transactions, but generally not the prices seen. This is the case both for customer who purchased and by those who did not purchase, while due to the heavy operational effects on ancillary seat pricing, the logged transaction prices can be heavily distorted, creating a censoring problem when estimating Willingness-To-Pay (WTP). Many airline seat price optimization papers focus on optimization given a demand signal or based on stated WTP through surveys. We focus here on creating a reliable dataset to learn WTP from observed data. We then consequently apply robust causal inference algorithms to learn purchase probabilities. Finally, we use these learned purchase probabilities to optimize price setting. Rather than focusing on inventory-constrained dynamic pricing, we consider the related problem of estimating willingness-to-pay when price exposure itself is partially unobserved or distorted. This paper extends and substantially develops earlier work presented in von Lücken et al. (2023).

Problem definition

Let $y \in \{0, 1\}$ indicate purchase of a seat product on a given flight instance, with features *x* and some price *p*. Operational logs store p^{paid} for $y{=}1$, but for $y{=}0$ the offered price p^{seen} is unobserved or Missing Not At Random (MNAR). Even p^{paid} confounds offered price with various potential discount adjustments. Without reconstructing exposure, price sensitivity cannot properly be identified, leading to biased elasticities and suboptimal price setting.

Background: seat pricing and ATPCO

Seat products (e.g., extra legroom seats) have an ATPCO base price by destination and month (full detail of how this works is available in Macrakis et al. (2016)). Around this base, various business rules are normally applied:

- *Load-based markdowns* up to $\approx 10\%$ close to departure if the section is selling less than expected.
- *Check-in markdowns* when economy seating is overfilled, to avoid free upgrades; these can be much larger.
- *Frequent-flyer (FF) discounts* granted at all times (usually range of 10%-50% off the seen price, with top-tier passengers often getting such seats for free).
- *Partner artifacts* (e.g., partner bookings that are not credited as an FF discount since the FF status comes from the partner airline) that cause anomalously high "discounts".

These mechanics make the observed paid prices a mixture of operational and loyalty effects on top of the ATPCO base, distinct from the offered price exposure that non-buyers saw. Therefore, ancillary seat pricing faces more exposure uncertainty than products with stable list prices and transparent promotions.

If we fit $\Pr(y{=}1 \,|\, x, p)$ using only purchases with p^{paid} and treat zeros as $p{=}0$, we bias both slopes and levels: non-buyer exposures are unknown, and buyer prices are adjusted based on business rules. In other words, we are basically then assuming that our pricing business rules have no effect on purchase probability. De-biasing requires reconstructing p^{seen} for all records and an effective price notion for modeling.

Related work

Ancillary pricing

We focus on products with perishable inventory, since airline seats cannot be sold once the flight has departed. In such settings, accurate demand prediction depends on learning both the customer arrival process and the distribution of WTP, as emphasized by Yang et al. (2022).

Gallego and Talebian (2012) introduces a rolling horizon framework that dynamically selects prices while simultaneously updating unknown parameters governing arrival rates and customer valuations for different seat locations over the sales horizon. Their model assumes that all seats share a common but unknown base value, with known premiums or discounts applied relative to this base value. Demand arrivals follow a Poisson process, with Gamma priors used for Bayesian updating of arrival rates. The base value is then estimated via maximum likelihood, and pricing decisions are derived through dynamic programming. Talebian et al. (2014) develop a dynamic programming model that integrates assortment and pricing decisions for perishable and seasonal products. Their approach incorporates Bayesian demand learning, and demonstrates that reducing prices earlier in the time window helps learn this demand faster.

More recently, Yang et al. (2022) propose a revenue management framework combining an LSTM neural network with a Markov Decision Process (MDP) to determine prices. The neural network is trained on historical data to estimate the WTP distributions parameters. These estimates are then embedded within the MDP to optimize pricing decisions, and since the network is already trained, the estimates can already be used early in the selling horizon.

Within the airline domain, Mumbower (2025) provide a broad review of ancillary offer management and dynamic pricing, highlighting WTP estimation as a central unresolved challenge. Wang et al. (2023) develop a Markov chain choice model for dynamic ancillary bundling, assuming that the relationship between price and demand has already been estimated from historical data. In contrast, their work does not focus on recovering this relationship from censored observations. Long and Belobaba (2024) analyze segmented continuous pricing enabled by NDC standards (NDC is a relatively new way of distributing price offers for airlines to their agents which enables more control of the final offer for the airline), concentrating on ticket pricing. They show that finer price granularity can generate revenue gains, though such gains depend critically on accurate demand estimation of the type studied here. Our work complements these contributions by developing a methodology to recover the price sensitivity estimates required for downstream optimization.

Another stream of literature estimates customer WTP using survey-based approaches, leveraging individual-level characteristics to determine personalized prices. Examples include (Warnock-Smith et al. 2017; Shukla et al. 2019), and Zhao et al. (2021).

However, estimating air passenger price elasticity from transactional data presents significant challenges. As discussed by Perera and Tan (2019), two issues are particularly relevant. First, the level of data granularity required for rigorous elasticity estimation is often difficult to obtain, as detailed passenger booking data is proprietary and commercially sensitive. Second, prices and passenger demand are jointly determined, giving rise to endogeneity: pricing decisions depend on anticipated demand, while realized demand is influenced by prices.

Addressing personalization in upgrade pricing, Thirumuruganathan et al. (2023) introduce a price elasticity model (PREM) aimed at identifying customers likely to accept economy-to-premium upgrade offers. Their approach employs denoising autoencoders to construct embedding-based representations and reduce dimensionality. A cost-sensitive classification model is then used to identify customers with higher upgrade propensity while penalizing false negatives. Using three binary autoencoders, customers are assigned to personalized price buckets (low, middle, or high), after which upgrade offers are optimized to maximize expected revenue. Unlike PREM, which classifies customers into discrete price buckets from fully observed offer-response pairs, our setting must first reconstruct the price exposure itself, since offered prices are unobserved for non-buyers and distorted for buyers.

In contrast, Duijndam et al. (2025) examine ancillary seat pricing from a revenue management perspective. Without assuming prior knowledge of WTP, they learn price-demand relationships at the flight level through adaptive pricing strategies and subsequently optimize inventory control for a given seat type. Their framework does not incorporate individual customer characteristics. We differ in two respects: we estimate a contextual purchase probability $\pi(x, p)$ at the passenger level rather than a flight-level demand curve, and we treat the recovery of price exposure under censoring as the central estimation problem rather than assuming it is available.

Censored data

An important element of our approach is that we have censored data, in the sense that from our logging data, we don't know exactly in our offline modeling what price was originally proposed to the customer. In the pricing literature, this is not an unknown scenario. In relatively recent work, Bu et al. (2023) studies a similar problem, however a key difference is that the focus there is on unconstraining of demand. So the focus is learning how much more could have been sold if inventory levels would not be reached. Price elasticity is assumed to be linear, and through varying prices in a clever way and solving a linear regression to find the demand-price relation, optimal prices are found. The problem is similar in Chen et al. (2020), where a combined pricing and inventory decision is studied under censored demand. Here the relation between price and demand is assumed to belong to some (known) distribution family parameterized by an unknown z. The goal is to set price and inventory levels under various conditions to learn this z, and with that set the optimal policy. Another interesting approach is in Sachs and Minner (2014), where the same lost sales inventory problem is studied, but then by looking at more granular data to estimate the true customer demand and with that optimal inventory or pricing – here also studied on real-life data.

Closer to our case, Tang et al. (2025) study feature-based price optimization based on offline learning from historical data with a finite inventory. Also in this paper, the censoring focus is on recovering expected true conditional demand if inventory levels would not be met. Here methods from survival analysis are used to recover this relation. The distinction from this unconstraining literature (Bu et al. 2023; Tang et al. 2025) is central to our contribution: these methods impute missing *quantity* (how much more would have sold if there would be no capacity limits), whereas Synthetic Offer Reconstruction (SOR) pipeline reconstructs the missing *price exposure* that buyers and non-buyers faced, and additionally corrects buyer-side distortion from loyalty and operational rules rather than treating only non-buyer missingness.

Sample selection and missing-not-at-random demand

The censoring we face is a form of sample selection: the offered price is observed only for a self-selected subset (buyers), and even then it is distorted. The canonical econometric treatment is the two-step selection model of Heckman (1979), in which a first-stage equation models the probability that an outcome is observed and a second-stage outcome equation is corrected via the inverse Mills ratio. A nonparametric alternative widely used in causal inference reweights observed records by the inverse of an estimated response or exposure propensity (inverse-probability weighting (IPW); (Horvitz and Thompson 1952), and doubly robust estimators combine an outcome model with a propensity model so that consistency is retained if either is correctly specified (Robins et al. 1994).

These methods are the natural benchmarks for our problem, and we draw on them directly: we include an IPW response-reweighting baseline in our experiments (Sect. Empirical evaluation). We do not adopt the Heckman model as our primary estimator for three reasons specific to the airline setting. First, the censoring here is double-sided: non-buyer offers are entirely missing *and* buyer prices are distorted by loyalty and operational rules, whereas the Heckman correction targets one-sided selection on the outcome. Second, Heckman's identification leans on a bivariate-normality assumption between the selection and outcome errors that is difficult to justify for a binary purchase outcome with heavy operational discounting. Third, a valid exclusion restriction (a selection instrument affecting observation but not purchase) is at least as demanding as the proxy conditions required by the approach we adopt. In the end, most importantly, neither selection-correction method returns a reconstructed price feature. They de-bias a demand or outcome model conditional on the data at hand, whereas our first stage must recover the price exposure required for downstream optimization over a price grid. We therefore treat selection correction as a benchmark for the de-biasing objective rather than as a substitute for exposure reconstruction.

Proximal causal inference

Classical causal inference relies on the assumption that all confounders affecting treatment and outcome are observed. When this is violated, identification typically fails unless it's possible to introduce valid proxy variables. Proximal Causal Inference (PCI), formalized by Tchetgen et al. (2024); Cui et al. (2024), provides a general identification framework for such settings by leveraging pairs of observable proxies for unmeasured confounders.

In the PCI framework, one distinguishes an upstream proxy S associated with the latent exposure mechanism, and a downstream proxy W associated with the outcome conditional on the exposure. Three conditions are required in this theory: (i) *relevance*, that the proxies are associated with the latent factors; (ii) *exclusion*, that each proxy is independent of the potential outcome or exposure given the unmeasured factors; and (iii) *completeness*, ensuring uniqueness of the solution. The causal relation of interest can be identified through a bridge function h_ϕ satisfying a conditional moment equation

$$\mathbb{E}[\,Y - h_\phi(X, A, W) \mid X, S\,] \;=\; 0,$$

where A denotes the treatment or exposure variable, Y the outcome, and X observed covariates. The function h_ϕ can be estimated using flexible nonparametric or machine-learning regressors, such as kernel methods or gradient-boosted trees.

The theoretical foundations of PCI, including the bridge function construction and identification conditions, are developed in Tchetgen et al. (2024) and Cui et al. (2024). We provide an accessible intuition for our application in Sect. Proximal causal inference for WTP estimation.

PCI has been successfully applied in econometrics, epidemiology, and reinforcement-learning contexts where exposure is only partially observed. In our setting, the offered seat price corresponds to the latent exposure, the observed paid price and operational variables serve as upstream and downstream proxies (S, W), and the bridge moment above defines the correction applied after Synthetic Offer Reconstruction. This provides a theoretically grounded way to recover purchase probabilities $\pi(x, p)$ despite unobserved price exposures.

Closest to our work, Shen and Cui (2025) develop a general proxy-aided demand learning framework that applies proximal causal inference bridge estimators to pricing under endogeneity, with regret guarantees for online and offline settings. Our setting differs in that price exposure itself is partially unobserved and operationally distorted rather than only confounded: non-buyer offers are entirely missing, and buyer prices are perturbed by loyalty and operational rules. Relative to Shen and Cui (2025), whose contribution is a general proxy-aided estimator with regret guarantees, our increment is threefold and specific to censored ancillary exposure: (i) we precede the bridge estimator with an operationally grounded reconstruction (SOR) that exploits known ATPCO mechanics to recover a price feature, which their framework assumes available; (ii) we estimate the bridge with a cross-fitted two-stage sieve (2SLS) estimator that projects a polynomial basis of (X, A, W) onto a polynomial basis of the instrument set (X, A, S), with categorical airline fields handled by one-hot encoding, and we compare it against gradient-boosted alternatives that ignore the moment structure; and (iii) we provide practitioner guidance for when the proximal refinement pays for itself versus when the reconstruction alone suffices, grounded in the bias-variance pattern observed across our evaluation regimes and in proxy diagnostics that practitioners can compute on their own logs (Sect. Proxy diagnostics). In our pipeline PCI is a robustness refinement on top of SOR rather than the primary demand estimator.

Relation to censored-demand literature. The censored-demand literature in operations research and revenue management typically addresses unobserved sales through unconstraining or survival-based imputation methods (Bu et al. 2023; Chen et al. 2020; Tang et al. 2025), as discussed earlier in Sect. Censored data. These approaches assume a known functional form for price sensitivity or inventory effects and aim to estimate missing demand under capacity constraints. Proximal causal inference differs in its focus. Instead of assuming a specific parametric censoring process, it identifies the effect of latent exposure through observable proxies under weaker, testable moment conditions.

Our contribution

This paper makes two complementary contributions, that together strengthen the estimation of WTP.

1. **Synthetic Offer Reconstruction (SOR).** We structure the censoring problem in ancillary seat logs, where purchases are recorded with loyalty and operational discounts while non-purchases lack exposure. We then propose a transparent *Synthetic Offer Reconstruction* pipeline that recovers offered prices by reversing loyalty rules, removing operational artifacts, estimating flight-level discount dispersion around the ATPCO base, imputing non-buyer exposures, and normalizing onto a price grid. This yields a feature-based exposure variable suitable for learning willingness-to-pay from revealed behavior (as opposed to stated preferences) and is distinct from unconstraining demand for inventory.
2. **A valid proximal pair and a structural bridge for WTP.** Building on SOR, we estimate the price-purchase relationship under MNAR exposure with a proximal bridge, using PCI theory. Two design choices make the proxy structure relevant. First, the downstream proxy W is built from *leave-one-out, within-flight* statistics - the mean reconstructed normalized price and the purchase rate of the *other* passengers on the same flight - so that W is a downstream consequence of the shared flight-level exposure regime, without coinciding with the passenger's own treatment (price) or outcome. Second, the bridge is estimated by a cross-fitted two-stage sieve estimator in which the basis of $H = (X, A, W)$ is projected onto the basis of $Z = (X, A, S)$, which is exactly the empirical analogue of the conditional moment (2). We show that this structural use of the proxies is essential: supplying the same W (and S) to a flexible classifier as ordinary covariates improves *observed* calibration but produces a degenerate, revenue-destroying pricing policy. The proximal bridge is therefore not a cosmetic addition to SOR, but the mechanism by which proxy information can be exploited safely.

Taken together, the combined domain structure of airline seat pricing, where SOR reconstructs operationally consistent exposure, and the PCI-based WTP estimation framework, provides a principled correction for MNAR bias and yields robust, feature-based elasticities without requiring a fully specified demand model.

Data and the censoring problem

In order to properly estimate passenger WTP, we make use of a synthetic data set. Every data point represents a passenger with a corresponding flight, and related characteristics. We observe passenger instances $i = 1, \ldots, N$ with features x_i (itinerary, timing, connection structure, demographics, booking channel, etc.), a label $y_i \in \{0, 1\}$ indicating paid seat purchase, and, if $y_i{=}1$, an effective paid seat price p_i^{paid}. Each instance belongs to a flight $f(i)$ with ATPCO base price $b_{f(i)}$ (destination, month). Non-buyers ($y_i{=}0$) have $p^{\text{paid}}{=}0$ with no associated offered price recorded. Therefore, the exposure p_i^{seen} is missing-not-at-random, and needs to be corrected. As described in Sect. Background: seat pricing and ATPCO, there are various reasons and mechanics that create differences from the base price. Non-FF anomalies are rare (about a few percent of passengers) but distort flight-level discount estimates and are excluded during cleaning. After removing check-in and non-FF artifacts, remaining discounts are approximately normal around a small mean with moderate dispersion; trimming 3σ tails tightens dispersion markedly.

We emphasize the role of the synthetic environment: the simulator emits a *censored operational log* of the kind an airline data warehouse would contain - non-buyer offered prices are withheld, and buyer paid prices are distorted by loyalty and operational rules - and all estimation (SOR and every purchase model) is applied to this log only. The simulator's ground truth is used exclusively to evaluate policies against the oracle, never as an input to any estimator.

Table 1 Notation used in the Synthetic Offer Reconstruction (SOR) Section

Symbol	Definition / meaning
i	Passenger identifier
f	Flight identifier; passengers indexed by $i \in f$
$y_i \in \{0, 1\}$	Purchase indicator for passenger i (1 if seat purchased)
b_f	ATPCO base price for flight f
p_i^{paid}	Logged paid price (may differ from p_i^{eff} due to artifacts)
p_i^{seen}	Price seen by passenger i (offer price before loyalty effects)
p_i^{eff}	Effective price actually faced by i, after loyalty adjustment
$d_i = 1 - p_i^{\text{seen}}/b_f$	Discount fraction relative to ATPCO base
$\Delta_f = \mathcal{N}(\hat{\mu}_f, \hat{\sigma}_f^2)$	Model used for sampling non-buyer exposure
μ_f, σ_f	Mean and standard deviation of flight-level discount distribution
$\tilde{\delta}_i$	Imputed discount draw from per-flight distribution Δ_f
δ_i^{FF}	Fractional discount determined by FF tier rules
$\hat{\mu}_f$, $\hat{\sigma}_f$	Estimated flight-level discount parameters after cleaning
P	Discretized grid of normalized prices ($\{0.475, 0.500, 0.525, \ldots, 1.075\}$)
$p_i^{\text{norm}} = p_i^{\text{eff}}/b_f$	Normalized effective price (used for learning)
FF tier_i	Loyalty tier of passenger i

Therefore, in order to have a proper relation between WTP and offered prices, we should reconstruct p^{seen} for non-buyers, reverse loyalty effects from buyer prices, and define a consistent price feature for all records. In the next section we describe our main method for this.

Synthetic offer reconstruction (SOR)

As described, we need reliable estimate for the prices offered, to use that to learn passenger WTP. SOR first undoes loyalty adjustments on observed purchases, removes atypical operational markdowns and partner airline artifacts, and fits a per-flight discount distribution around the ATPCO base price. It then imputes plausible seen prices for non-buyers, reapplies FF rules to form an effective price for modeling, and normalizes/bins prices to stabilize estimation and support grid-based optimization. First, we define some terminology in table 1.

Step 1: Discount reversal and outlier filtering

For y=1, so for purchases, starting with p_i^{paid}, we reverse FF discounts to approximate pre-loyalty prices. We remove heavy check-in discounts using operational flags, exclude partner airline artifacts, and trim 3σ d_i discount outliers around the global mean, where this trimming reduces dispersion substantially (see Fig. 1).

Discounts are defined as $d_i = 1 - p_i^{seen}/b_f$. Therefore, $d_i < 0$ corresponds to a price above the ATPCO base, which we use in the simulator to represent moderate route- or instance-level surcharges which might also occur

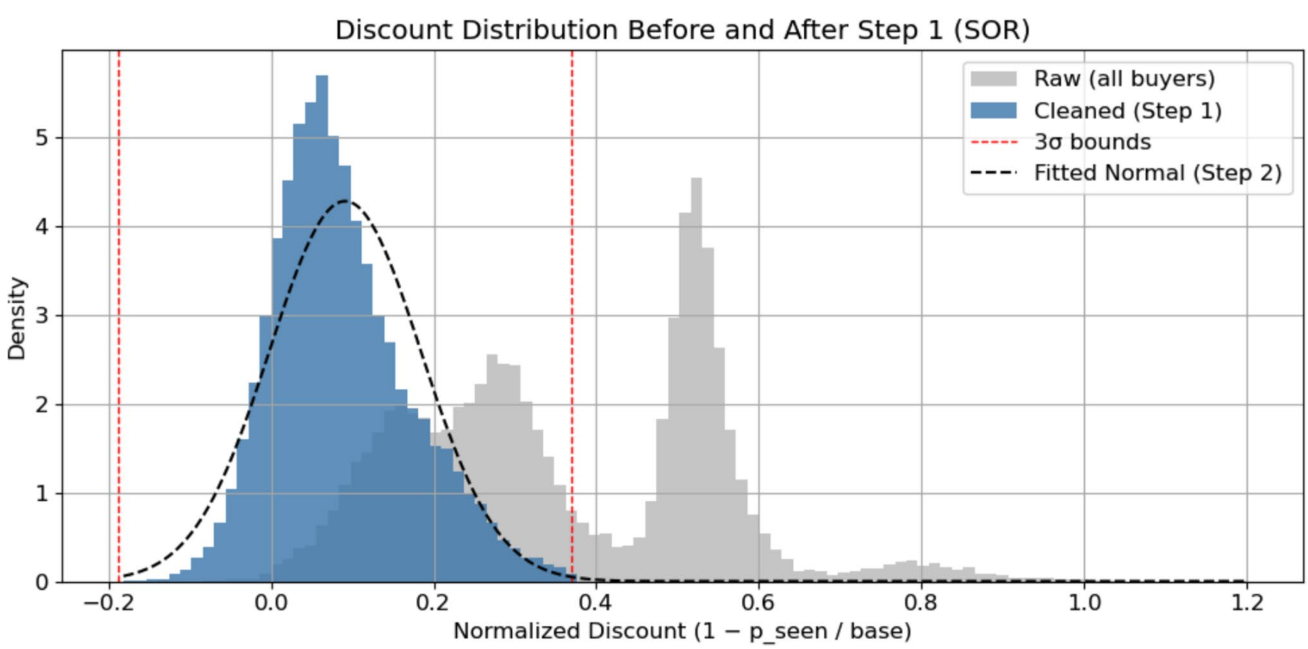

Fig. 1 The dispersion of price discounts with and without trimming on synthetic data. Negative discounts correspond to above-base prices. Observations in the economically impossible tail are excluded before fitting the per-flight Normal discount model

occasionally in real-life data. Values with $d_i \geq 1$ would imply non-positive seen prices and are treated as operational artifacts. Such observations are removed by the Step Step 1: Discount reversal and outlier filtering filtering and the subsequent 3σ trimming before the per-flight Normal discount model is fitted. Figure 1 therefore shows the raw synthetic discount distribution for diagnostic purposes, while the fitted Normal is estimated only on the cleaned and trimmed sample.

Step 2: Flight-level discount estimation

On the cleaned buyer set (from Sect. Step 1: Discount reversal and outlier filtering), we define discount fractions $d_i = 1 - p_i^{\text{seen}}/b_f$. After trimming 3σ tails, we assume $d_i \sim \mathcal{N}(\mu_f, \sigma_f^2)$ per flight. We estimate $(\hat{\mu}_f, \hat{\sigma}_f)$ by sample moments with $\hat{\sigma}_f \geq \sigma_{\min}$. The outlier filtering from Step 1: Discount reversal and outlier filtering serves as effective truncation, so a simple Normal distribution is adequate after this cleaning step. This Normal model is used in next steps both to impute non-buyer offers and to regularize outlier handling.

Step 3: Estimation of seen and effective prices

For any passenger *i* on *f*, draw $\tilde{\delta}_i \sim \Delta_f$ and set

$$p_i^{\text{seen}} \;=\; b_f \cdot (1 - \tilde{\delta}_i), \qquad p_i^{\text{eff}} \;=\; p_i^{\text{seen}}\,(1 - \delta_i^{\text{FF}}),$$

where δ_i^{FF} is defined by fixed business rules for the FF program. For buyers we analogously reconstruct p^{seen} by undoing FF rules and relevant operational rules.

Step 4: Price feature construction

The price feature is critical for optimization. In our approach, we define four variants for input to probe robustness and control correlations:

- `ATPCO`: tests whether a fully reconstructed exposure variable provides a more consistent measure of willingness-to-pay than observed operational prices. So we do a full overwrite of all prices using Steps 1–3, including for buyers. We apply FF discounts in this $\Rightarrow$ effective prices reflect tier benefits.
- `ATPCO_nodisc`: tests whether the apparent effect of price is confounded by frequent-flyer status when discounts are embedded directly in the price feature. So we do a full overwrite of all prices, including for buyers, but remove FF discounts from the feature. We include FF tier as a separate covariate to decorrelate price and tier.
- `ATPCO_partov`: tests whether correcting only extreme discounts and operational artifacts is sufficient, or whether systematic reconstruction of all prices is required. So we overwrite only outliers/heavy discounts for buyers, but keep in-sample purchases otherwise.
- `ATPCO_noov`: tests the consequences of using raw logged prices without reconstruction, so only available for buyers, and therefore serves as an operational baseline.

Together, these variants isolate two key design choices: (i) whether observed prices should still be reconstructed and (ii) whether loyalty discounts should be represented through the price feature itself or through separate covariates. As can be seen in Fig. 2, `ATPCO` yields the most coherent, decreasing elasticity. The `nodisc`-variant avoids price-tier collinearity, but harms low-price elasticity. The `noov` variant simply uses the raw logged prices, and is therefore left out of this evaluation. In the following sections, we will use the `ATPCO` method to input prices in the dataset, which means we overwrite all prices, including for buyers.

Step 5: Normalization and binning

To stabilize estimation across routes and support grid search in optimization, we normalize prices by b_f and discretize into 2.5% bins:

$$P = \{0.475, 0.500, 0.525, \ldots, 1.075\},$$

where the lowest bin is represented by 0.475 to reflect mass just below 0.5. This base-price normalization eases comparison and bound handling downstream.

Proximal causal inference for WTP estimation

The SOR pipeline reconstructs a plausible price exposure variable by removing artifacts, imputing missing offers around ATPCO base levels, and normalizing variation per flight. This already enables effective supervised learning of demand curves. However, SOR alone does not guarantee robustness under missing-not-at-random (MNAR) exposure: whether a passenger sees a given price may depend on operational factors (e.g., heavy check-in load or partner flight logic) that are not fully captured by the reconstructed offer, or the data itself might not be completely correct. To further reduce exposure-driven bias and improve stability, we introduce a refinement step based on proximal causal inference ideas.

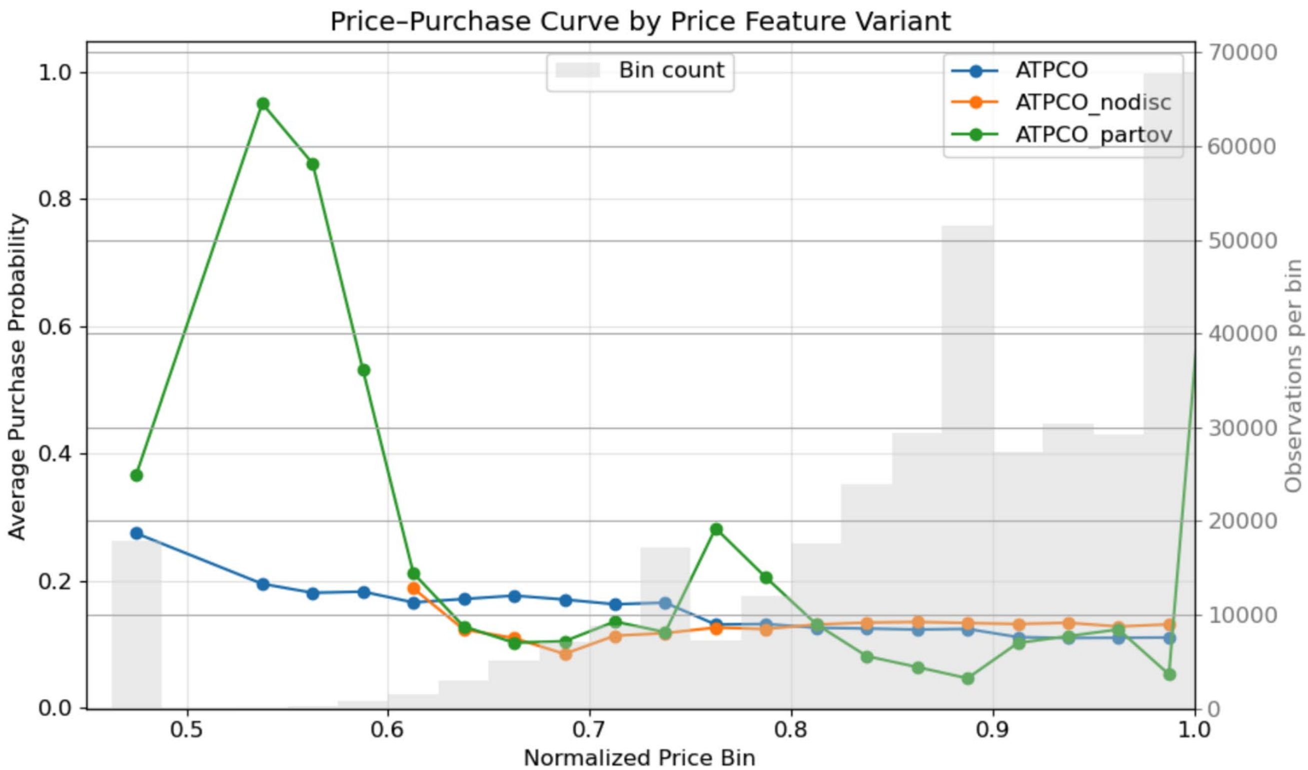

Fig. 2 The price elasticity based on the various price input methods on synthetic data. Since the `ATPCO_noov` is the original price for buyers, there is no average purchase probability available and it is therefore left out of the plot. The grey bars indicate bin volume

Core intuition. The key challenge in estimating willingness-to-pay from observational data is that the offered price is not randomly assigned: passengers who saw lower prices (due to FF discounts, operational markdowns, or partner effects) may differ systematically from those who saw higher prices. If we ignore this selection into price exposure, our estimated price sensitivity will be biased.

Proximal causal inference addresses this by exploiting two types of observable variables that serve as proxies for the unobserved confounding mechanism:

- **Upstream proxies** (S): variables that drive variation in price exposure but do not directly affect purchase except through price - here operational state (capacity, partner-booking flags) and temporal triggers (booking month) of discount rules.
- **Downstream proxies** (W): variables that are downstream consequences of the realized exposure and are observed alongside the outcome. We construct W from leave-one-out, within-flight statistics of the other passengers on the same flight (Sect. Latent-offer model and proxies), so that W reflects the shared flight-level exposure regime without coinciding with the focal passenger's own price or purchase.

The identification logic is an instrumental one. The upstream set S moves the exposure proxy W without acting on purchase directly, so the part of W that S predicts identifies how exposure maps to purchase, while the part of demand correlated with the latent flight-level confounder is purged. The bridge $h_\phi(x, p, W)$ that solves the resulting moment equation therefore recovers the causal price-purchase relationship rather than the confounded association. This is well suited to airline pricing: operational adjustments create the selection, SOR supplies a reconstructed price feature, and routinely logged operational variables supply both the upstream instruments S and the within-flight proxies W.

Latent-offer model and proxies

For passenger i on flight f, let b_f be the ATPCO base and write the (normalized) seen price as

$$\bar{p}_i^{\text{seen}} = \frac{p_i^{\text{seen}}}{b_f}, \qquad \bar{p}_i^{\text{eff}} = \bar{p}_i^{\text{seen}}\,(1 - \delta_i^{\text{FF}}),$$

with δ_i^{FF} the loyalty discount fraction determined by business rules. We model the operational discount $\delta_i := 1 - \bar{p}_i^{\text{seen}}$ at flight level by moments (μ_f, σ_f) estimated from SOR and treat $\bar{p}_i^{\text{seen}}$ as latent for non-buyers. Let x_i denote customer and itinerary features. We define S_i as operational variables that influence exposure but are excludable from the direct purchase decision (capacity, partner artifacts, booking month). The downstream proxy W_i is not the passenger's own reconstructed price. Instead, for each passenger we form leave-one-out, within-flight statistics

$$W_i^{\mathrm{p}} = \frac{1}{|f(i)|-1} \sum_{j \in f(i), j \neq i} \bar{p}_j^{\mathrm{eff}}, \qquad W_i^{\mathrm{y}} = \frac{1}{|f(i)|-1} \sum_{j \in f(i), j \neq i} y_j, \tag{1}$$

i.e. the mean reconstructed normalized price and the purchase rate of the other passengers on the same flight *f*(*i*). Both are downstream consequences of the shared flight-level exposure regime (and of any unobserved flight-level demand pressure), but exclude the current passenger of interest's own treatment $A_i = \bar{p}_i^{\mathrm{eff}}$ and outcome y_i. Together, (S_i, W_i) with $W_i = (W_i^{\mathrm{p}}, W_i^{\mathrm{y}})$ form a proximal pair: S_i captures upstream operational drivers of exposure, while W_i is a downstream proxy of the realized exposure that carries information about the latent flight-level confounder without being a deterministic function of S_i (Sect. Proxy diagnostics).

Note also that using the outcomes of other passengers in W^{y} does not leak the focal passenger's own label. Purchases within a flight are conditionally independent given the flight-level state (U_f, α_f) and individual covariates, so W_i^{y} carries information about the shared latent state - which is exactly what a downstream proxy should do - while having no direct causal channel into y_i. The leave-one-out construction, rather than cross-fitting, is what enforces this.

Proximal bridge for MNAR exposure

Let $\pi(x, p) = \Pr(y{=}1 \mid x, p)$ denote the purchase probability with context x and price p. Following proximal causal inference (e.g. Tchetgen et al. (2020); Cui et al. (2024); Tchetgen et al. (2024)), suppose there exists a bridge function $h_\phi(x, p, W)$ such that

$$\mathbb{E}\big[\, y_i - h_\phi\big(x_i, \bar{p}_i^{\mathrm{eff}}, W_i\big) \;\big|\, x_i, S_i \,\big] \;=\; 0. \tag{2}$$

The conditioning set is (x_i, S_i) and excludes W_i, even though h_ϕ depends on W_i. This is by design: the moment integrates over *W* given (*x*, *S*), and it is this integration that yields identification of the causal price-purchase relationship $\pi(x, p)$. Intuitively, S acts as an instrument that drives the exposure proxy W without affecting purchase directly, so the component of W that is predictable from S identifies how exposure maps to purchase, while the component of demand that is correlated with the latent flight-level pressure is purged.

We estimate (2) with a cross-fitted two-stage sieve estimator. Let Φ_H be a polynomial-sieve basis of $H = (X, A, W)$ and Φ_Z a polynomial-sieve basis of the instrument set $Z = (X, A, S)$, with categorical airline fields (destination, frequent-flyer tier, partner flag) one-hot encoded. The first stage projects each column of Φ_H onto Φ_Z by ridge regression, yielding fitted sieve features $\widehat{\Phi}_H$; the second stage regresses *y* on $\widehat{\Phi}_H$ by ridge regression. This is the standard two-stage least-squares analogue of (2): the instrument set *Z* removes the part of *H* that is correlated with the structural error. We compute out-of-fold predictions with $K = 5$-fold cross-fitting (Chernozhukov et al. 2018) so that the bridge is never evaluated on the data used to fit it, and refit on the full sample for deployment-time prediction. Because the second stage is linear, predictions are clipped to [0, 1] as a guardrail before being interpreted as purchase probabilities. We measure the clipped share on the raw (so pre-clipped) out-of-fold predictions. Across regimes it is at most 0.03% of observations (e.g. 0.026% in the moderate regime, 0.002% under shift), and clipping binds exclusively at the lower boundary - the 99th percentile of raw predictions is below 0.28 and no prediction exceeds one. Lower-boundary clipping occurs when the bridge extrapolates to high prices where the true probability is near zero. A prediction clipped to zero contributes zero expected revenue and therefore cannot attract the revenue argmax, so clipping cannot mask miscalibration in the direction that matters for pricing.

Assumption 6.1 (*Proxy conditions*) The proxy variables *S* and *W* satisfy:

1. **Relevance:** *S* contains information about the latent operational exposure factors that influence the effective price observed by a passenger.
2. **Exclusion:** conditional on the latent exposure factors, *S* does not directly affect purchase outcomes (it acts only through exposure), and *W* does not directly affect purchase outcomes beyond its association with the latent exposure. Equivalently, $W \perp Y \mid (\mathrm{exposure}, X)$ up to the latent factors, so that *W* is a valid downstream proxy rather than a direct demand shifter.
3. **Completeness:** a bridge function h_ϕ satisfying (2) exists and is identifiable from the observed data; or equivalently, *S* induces enough variation in *W* to pin down the exposure-purchase map.

Under these conditions $\pi(x, p)$ is identified up to the approximation error of the chosen function classes.

Proxy diagnostics

A proximal pair is only informative if (i) the downstream proxy W carries variation about the latent exposure that is not already in *S*, and (ii) the upstream set *S* is, at least approximately, excludable from purchase given price and covariates. We assess both here. Table 2 summarizes the diagnostics across the regimes in our synthetic data we test our models on, generated by the process described in Sect. Data-generating process.

For non-degeneracy, we regress W^{p} on S out-of-fold and report the unexplained share $1 - R^2(W^{\mathrm{p}} \mid S)$. About 82% of the variation in W is independent of S, so the pair does not collapse into a regression on observed covariates. The proxy also tracks the latent factor in the expected directions: a higher latent flight pressure U_f lowers the offered price $(\mathrm{corr}(U, A) < 0)$ and raises within-flight purchases $(\mathrm{corr}(U, W^{\mathrm{y}}) > 0)$, while S alone explains under a third of U $(R^2(U \mid S) \approx 0.29)$. This is the configuration the instrument-based identification in (2) requires: W supplies first-order information about exposure beyond S. Finally, the first-stage strength of the instrument set is non-trivial. Beyond covariates and price, S explains an additional 4–6% of the variation in W^{p} and 14–17% of W^{y} (partial R^2, last two rows of the top block). This is the empirical counterpart of the completeness requirement that S moves W: a near-zero value here would signal a weak-instrument problem for the bridge, analogous to weak instruments in two-stage least squares.

For approximate exclusion, we compare the out-of-fold predictive performance of a model on $(x, \bar{p}^{\mathrm{eff}})$ against one that additionally includes S. Adding S improves AUC by only 0.2–0.8 percentage points and Brier score by $\mathcal{O}(10^{-4})$ once price and customer features are controlled. The contribution is small but not exactly zero, and it grows with the strength of the confounder, consistent with S being approximately - not perfectly - excludable. S correlates with U_f, which carries a residual direct demand channel. We therefore position PCI as a robustness refinement under approximate exclusion rather than as exact identification (Sect. Discussion).

For a practitioner applying these diagnostics to their own logs, each block guards against a distinct failure mode. If the independent share of W were near zero, W would be a deterministic function of S, the proximal pair would collapse, and the bridge would degenerate into an ordinary regression on observed covariates - the degenerate configuration in which the pair provides no identification. If the correlations between the latent factor and the proxies were near zero (in an applied setting: if within-flight statistics showed no co-movement with suspected confounding channels such as load or partner mix), W would carry no information about the confounder and the first stage would project onto noise. And if adding S to a demand model produced a large predictive gain, S would be acting as a direct demand shifter, the exclusion restriction would fail, and the "correction" would remove genuine demand signal rather than confounding. The desirable configuration is therefore the one largely observed here: W is informative about the latent factor, largely independent of S, with S adding little direct predictive power yet measurably moving W.

Diagnostic assessment. The proxy conditions cannot be fully verified in applied settings, but the diagnostics of Sect. Proxy diagnostics make their plausibility checkable. The non-degeneracy of the pair and the approximate exclusion of S are both quantified out-of-fold. Completeness is not directly testable and is treated as a maintained assumption, as is standard in the proximal-causal-inference literature. It is plausible here given the rich variation in flight-level discounting, temporal pricing rules, and partner-specific adjustments. We stress that, because observed calibration can reward a model that fits the confounded conditional (Sect. Price optimization results), it is not a reliable criterion for deciding whether to apply the refinement. Robustness to regime change and agreement with the deconfounded target are the relevant checks.

Proxy specification in our experiments. The assignment of variables to x, S, and W requires domain knowledge of the pricing mechanism. Direct demand shifters belong in the covariates x. The upstream set S is reserved for variables

Table 2 Proxy diagnostics, mean ± s.d. over 50 seeds. Top block: relevance and non-degeneracy of the proxy pair. Bottom block: approximate exclusion of S (out-of-fold gain from adding S to a model already containing x and price)

Quantity	Moderate	Strong	Shifted
$\mathrm{corr}(U, A)$	-0.30 ± 0.02	-0.28 ± 0.02	-0.41 ± 0.01
$\mathrm{corr}(U, W^{\mathrm{p}})$	-0.77 ± 0.02	-0.73 ± 0.03	-0.88 ± 0.01
$\mathrm{corr}(U, W^{\mathrm{y}})$	$+0.60 \pm 0.03$	$+0.74 \pm 0.02$	$+0.56 \pm 0.03$
$R^2(W^{\mathrm{p}} \mid S)$	0.18 ± 0.02	0.18 ± 0.02	0.18 ± 0.02
indep. share of W	0.82 ± 0.02	0.82 ± 0.02	0.82 ± 0.02
$R^2(U \mid S)$	0.29 ± 0.02	0.29 ± 0.02	0.29 ± 0.02
partial $R^2(W^{\mathrm{p}} \sim S \mid X, A)$	0.06 ± 0.01	0.06 ± 0.01	0.04 ± 0.00
partial $R^2(W^{\mathrm{y}} \sim S \mid X, A)$	0.14 ± 0.02	0.16 ± 0.02	0.16 ± 0.02
ΔAUC from S	0.0028 ± 0.0009	0.0077 ± 0.0019	0.0031 ± 0.0006
ΔBrier from S	1.2×10^{-4}	2.5×10^{-4}	1.4×10^{-4}

whose primary pathway to purchase runs through exposure, and W is the within-flight leave-one-out proxy. Concretely:

- x_i = is_business, has_connection, channel_online, FF_tier, length_hours, dest;
- S_i = capacity, partner_artefact, month;
- $W_i = (W_i^{\mathrm{p}}, W_i^{\mathrm{y}})$, the leave-one-out within-flight mean reconstructed price and purchase rate of Sect. Latent-offer model and proxies.

This keeps known demand shifters (destination, trip length, connection status, business-travel status, loyalty tier) in x, which ensures that operational variables with a demand pathway should not be treated as upstream proxies. Note that flight identity enters only through the construction of W, not as an S covariate. This avoids using a near-unique flight key as an instrument while still exploiting flight-level exposure variation through the leave-one-out statistics.

Proximal causal inference for WTP algorithm

Starting from SOR moments (μ_f, σ_f), we run:

1: **Input:** SOR outputs (normalized grid P, binned $\bar{p}^{\,\mathrm{eff}}$, flight moments (μ_f, σ_f)); covariates X; treatment $A = \bar{p}^{\,\mathrm{eff}}$; upstream proxies S; downstream proxies $W = (W^{\mathrm{p}}, W^{\mathrm{y}})$.
2: **Build** W**:** for each passenger, compute leave-one-out within-flight mean reconstructed price W_i^{p} and leave-one-out within-flight purchase rate W_i^{y}.
3: **Sieve bases:** form polynomial-sieve features Φ_H of $H = (X, A, W)$ and Φ_Z of $Z = (X, A, S)$, one-hot encoding categorical fields.
4: **Cross-fit (**K **folds):** for each held-out fold, on the remaining folds fit first-stage ridge $\Phi_H \leftarrow \Phi_Z$ to obtain $\widehat{\Phi}_H$, then second-stage ridge $y \leftarrow \widehat{\Phi}_H$; predict $\widehat{\pi}_i = \mathrm{clip}\big(\hat{h}(x_i, \bar{p}_i^{\,\mathrm{eff}}, W_i), [0,1]\big)$ on the held-out fold.
5: **Refit:** refit both stages on the full sample to obtain the deployment-time bridge $\hat{h}(\cdot)$.
6: **Output:** (μ_f, σ_f), $\widehat{\pi}(x, p)$.

Algorithm 1 Cross-fitted 2SLS sieve bridge for WTP refinement on top of SOR

The estimator is purely additive-free: there is no separate $c_\psi(S)$ term to fit and clip. The role formerly played by an explicit residual correction is now performed implicitly by the first-stage projection onto the instrument set $Z = (X, A, S)$, which removes the confounded component of H before the second-stage fit. Cross-fitting follows the double/de-biased machine-learning prescription (Chernozhukov et al. 2018) and removes the own-observation regularization bias that a single-sample fit would inherit.

Positioning and advantages

This step builds on the SOR estimated prices. SOR provides a deterministic operational reconstruction aligned with known rules. On top, the PCI-based WTP estimation treats SOR as initialization and refines the latent exposure-demand relationship. Compared to directly fitting a gradient-boosted model (e.g. CatBoost) on reconstructed prices, the PCI-based estimator is designed to correct for residual exposure bias and shows improved robustness when the data is perturbed.

A subtle but important point is that flexibility is not the same as correctness. A flexible classifier given the proxy variables S and W as ordinary inputs can fit the training outcome - and even the observed calibration on held-out data - at least as well as the bridge, because the proxies are informative. But without the moment restriction, it absorbs the proxy information in a way that does not respect the exposure mechanism, and the resulting price-purchase relation is no longer usable to set new prices. As Sect. Price optimization results shows, this results in a near-degenerate pricing policy. The cross-fitted 2SLS bridge instead uses S as an instrument for W, so the proxy information enters only through the channel that identifies the causal price effect.

Empirical evaluation

We evaluate the end-to-end pipeline in two steps. First, we present qualitative evidence from real operational airline data that the structural mechanisms encoded in the pipeline are present in practice (Sect. Evidence from real operational airline data). Second, we run a controlled synthetic evaluation containing a genuine unobserved confounder (Sect. Data-generating process), comparing the proposed proximal bridge against a direct SOR+CatBoost classifier, three classifiers that add the proxy variables as ordinary covariates, and an IPW benchmark, over 100 seeds with three complementary metrics.

Evidence from real operational airline data

Before turning to the synthetic evaluation, we report qualitative evidence from a real operational dataset for a paid seat ancillary product. For commercial-confidentiality reasons we cannot release the underlying records, nor disclose the specific routes, date ranges, seat product, or the exact variable set used in production. The selection of predictive

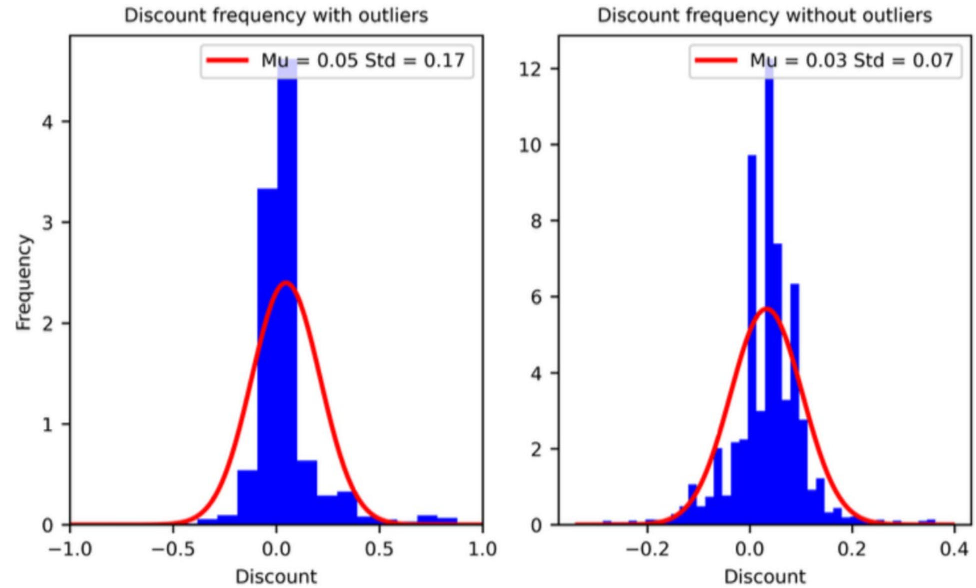

Fig. 3 Real operational data: per-record discount distribution before (left) and after (right) Step 1 cleaning. Trimming tightens dispersion around a small positive mean, but the cleaned distribution remains right-skewed, so the fitted Normal understates the right tail. Axes are shown in normalized discount units; absolute scales are withheld for confidentiality

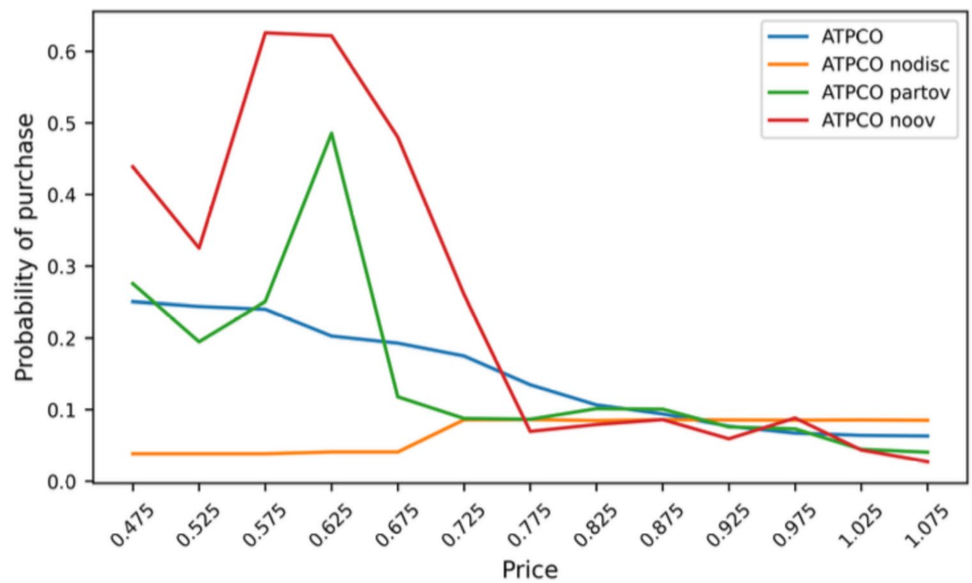

Fig. 4 Real operational data: average purchase probability over the normalized price grid under the four SOR price-feature variants. The qualitative ordering matches the synthetic ablation of Fig. 2: `ATPCO` is coherently monotone-decreasing, `noov` is outlier-inflated, `nodisc` is flat, and `partov` is noisy at low prices

variables is itself product- and carrier-specific and, in our experience, requires dedicated experimentation and domain expertise for each application. We therefore present this evidence not as a measurement of incremental revenue, which, as we argue below, requires a counterfactual oracle and hence a synthetic environment, but as a support that the structural mechanisms encoded in our pipeline are observed in practice.

Measuring incremental value on synthetic data Quantifying the revenue gain of one WTP estimator over another requires evaluating each estimator's pricing policy against the true purchase probability $\pi^*(x, p)$ at counterfactual prices. In live operations this counterfactual is unavailable. Synthetic data is consequently necessary for the oracle-relative comparisons of Sect. Results. It is the only setting in which ground-truth WTP is known. The real-data evidence below instead validates the qualitative behavior of the reconstruction and the plausibility of the feature and proxy choices.

Reconstructed discount distribution.

Figure 3 shows the per-record discount distribution on real data before and after the Step 1 cleaning of Sect. Step 1: Discount reversal and outlier filtering. As in the synthetic case (Fig. 1), trimming heavy check-in markdowns, partner artifacts, and 3σ outliers tightens the dispersion markedly and centers the distribution on a small positive mean. Consistent with the caveat in Sect. Discussion, the cleaned real distribution is right-skewed and not perfectly Normal. A single Gaussian captures the central mass relatively well, but slightly understates the tails. This motivates the imputation sensitivity analysis of Appendix A, which delineates which deviations the pipeline tolerates (mild skew around a single mode, as here) and which it does not (well-separated modes with no identifying flag).

Price-purchase response under the SOR variants.

Figure 4 reproduces, on real data, the price-feature ablation of Sect. Step 4: Price feature construction (cf. Fig. 2). The qualitative ordering matches the synthetic findings. The `ATPCO` reconstruction yields a coherent, monotone-decreasing price–purchase relationship; `noov` inflates apparent price sensitivity through uncorrected operational outliers; `nodisc` flattens the curve by removing loyalty-driven variation; and `partov` is unstable at the low-price end where few clean observations remain. That the same ordering emerges on operational data supports the SOR feature-construction choice independently of the synthetic generator. The `noov` is plotted here using estimated prices for non-buyers to get to an empirical purchase probability.

Demand heterogeneity and proxy specification.

Figure 5 shows observed purchase rates across a subset of customer and itinerary characteristics. Loyalty tier and corporate-traveler status shift purchase probability substantially and monotonically, confirming that these are genuine demand-side covariates. This empirical pattern supports our identification strategy in Sect. Proximal causal inference for WTP estimation. Variables with a direct relation to demand are assigned to the preference covariates x rather than to the upstream operational proxy set S, for which we reserve variables whose association with purchase plausibly runs through price exposure (e.g. capacity, partner flags, temporal discount triggers).

Data-generating process

The Data-Generating Process (DGP) is specified by the following structural equations. Each flight f has an ATPCO base b_f, a discount regime (μ_f, σ_f), a demand random intercept $\alpha_f \sim \mathcal{N}(0, \sigma_\alpha^2)$, and (in the confounded regimes) a latent factor $U_f \sim \mathcal{N}(\bar{u}, \sigma_u^2)$ that is never observed by any estimator. For passenger i on flight $f(i)$:

$$\begin{aligned}&(\text{exposure}) \quad \delta_i^{\text{op}} \sim \text{Trunc}\big(\mathcal{N}(\mu_f, \sigma_f^2), [0,1)\big),\\ &\quad \delta_i^U = \min\big(c_\delta \, \max(U_f, 0),\, d_{\max}\big),\end{aligned} \tag{3}$$

$$\bar{p}_i^{\text{seen}} = 1 - \big(\delta_i^{\text{op}} + \delta_i^U\big), \qquad \bar{p}_i^{\text{eff}} = \bar{p}_i^{\text{seen}}\,(1 - \delta_i^{\text{FF}}), \tag{4}$$

$$(\text{WTP shift}) \quad \theta_{d(i)} = w_{d(i)} + c_{\text{wtp}}\, U_f, \tag{5}$$

$$(\text{utility}) \quad z_i = \beta_0 + \beta_{d(i)}\big(\bar{p}_i^{\text{eff}} - \theta_{d(i)}\big) + \gamma^\top x_i + \alpha_{f(i)} + c_{\text{dem}}\, U_f, \tag{6}$$

$$(\text{purchase}) \quad \pi^*(x_i, p_i) = \sigma(z_i), \qquad y_i \sim \text{Bernoulli}\big(\pi^*(x_i, p_i)\big), \tag{7}$$

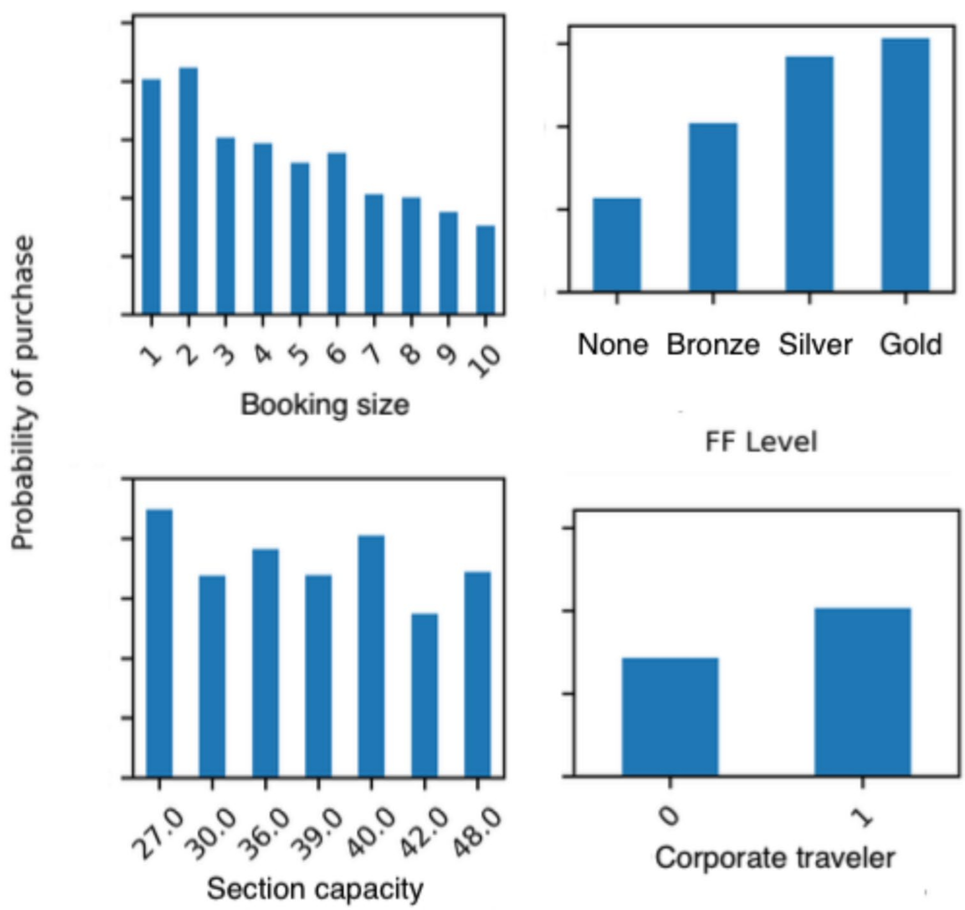

Fig. 5 Real operational data: observed purchase rates across selected customer and itinerary characteristics. Loyalty tier and corporate-traveller status are strong, monotone demand shifters, supporting their assignment to the preference covariates x rather than the operational proxy set S. Bars are marginal purchase rates by bin, not price elasticities. Variable scales and category labels are partially abstracted for confidentiality

where σ is the logistic function, $\beta_d < 0$ is the destination-specific price slope, w_d the destination WTP location, $\gamma^\top x_i$ collects the customer terms (business, connection, online channel, trip length, gold-tier), and δ_i^{FF} the FF-tier discount. The latent factor enters through two channels: it lowers the offered price (via δ_i^U, with coupling c_δ) and raises demand (via c_{dem}, and optionally WTP via c_{wtp}). Non-buyers' $\bar{p}^{\mathrm{seen}}$ is not logged (MNAR). The *deconfounded* (zero-U) oracle is the same model with $c_{\mathrm{dem}} = c_{\mathrm{wtp}} = 0$, isolating the causal price response; the *full* oracle retains all terms. The flight random intercept α_f introduces flight-level demand heterogeneity that is independent of price and is therefore not a confounder; only U_f creates endogeneity, because it moves price and demand jointly. Covariates x_i, calendar and capacity variables, and partner flags are drawn as specified in Table 4. Capacity binds on a relevant share of flights (Table 2 reports the realized binding share), and the partner artifact zeroes the *logged* FF discount for affected records while leaving the decision price unchanged. We then provide this dataset to all subsequent models. The full code for this can be seen in Anonymous Authors (2026).

An unobserved flight-level confounder and the evaluation regimes

We add a genuine *unobserved* confounder. Each flight *f* carries a latent factor $U_f \sim \mathcal{N}(\bar{u}, \sigma_u^2)$, never available to any estimator, that simultaneously (i) increases the operational discount, hence lowers the offered price, and (ii) raises baseline purchase propensity. A naive fit of purchase on reconstructed price is therefore biased. Low prices co-occur with high demand for reasons unrelated to price sensitivity, and this operates through a channel SOR cannot reconstruct from the recorded log.

We then study three regimes, each over 100 seeds. **Moderate**: U_f moderately raises demand. **Strong**: U_f raises both demand and willingness-to-pay more aggressively. **Shifted**: the latent factor has a positive mean in training and a negative mean at test ($\bar{u}_{\mathrm{train}} = +0.5$, $\bar{u}_{\mathrm{test}} = -0.5$). In the case of the shifted regime, an estimator that has merely memorized the in-sample association between low prices and high demand will misprice when the confounder flips sign, whereas an estimator that has recovered the deconfounded price effect will not. The diagnostics of Table 2 confirm the shift is working as designed ($\mathrm{corr}(U, A)$ strengthens from -0.30 to -0.41, and the train/test latent means separate cleanly).

Because the latent factor enters the outcome, we report against two oracles. The *full* oracle uses the true purchase probability including U_f and is the realistic target against which revenue is scored. The *deconfounded* (zero-U) oracle removes the latent demand effect and represents the pure causal price response the bridge targets. Reporting both makes explicit that a method which chases the full conditional - including the non-transportable U_f component - can attain a marginally better in-sample probability fit while pricing worse out of regime if the latent factor would shift.

Exact parameter configurations can be found in Table 3 and Table 4.

Table 3 Latent-confounder couplings and means by regime. The shifted regime uses opposite-signed latent means at train and test

Regime	$\bar{u}_{\mathrm{train}}$	$\bar{u}_{\mathrm{test}}$	c_{dem}	c_{wtp}
Moderate	0.0	0.0	0.50	0.00
Strong	0.0	0.0	0.75	0.10
Shifted	+ 0.5	− 0.5	0.50	0.00

Training objectives

For the two approaches, we define here how the model is trained, and specifically what training objective the model has.

SOR-only (baseline model). Fit a probabilistic classifier $\hat{\pi}_{\mathrm{SOR}}(x, p)$, which we implement through CatBoost (Prokhorenkova 2018), on inputs $(x, \bar{p}^{\mathrm{eff}})$. Since our training objective is to predict a seat purchase or not, we have a binary outcome. So we use Binary Cross-Entropy (BCE, log-loss):

$$\mathcal{L}_{\mathrm{BCE}} = -\frac{1}{n}\sum_{i=1}^{n}\Big(y_i \log \hat{\pi}(x_i, \bar{p}_i^{\mathrm{eff}}) + (1 - y_i)\log\big(1 - \hat{\pi}(x_i, \bar{p}_i^{\mathrm{eff}})\big)\Big).$$

SOR+PCI (sieve bridge). We estimate the bridge $h_\phi(x, \bar{p}^{\mathrm{eff}}, W)$ by the cross-fitted two-stage sieve estimator of Sect. Proximal bridge for MNAR exposure (Algorithm 1), using $Z = (x, \bar{p}^{\mathrm{eff}}, S)$ as instruments.

Estimator settings. The sieve bridge uses degree-2 polynomial bases for both $H = (X, A, W)$ and $Z = (X, A, S)$, with categorical fields one-hot encoded. The first stage is ridge with $\alpha = 10$ and the second stage ridge with $\alpha = 1$, cross-fitted over $K = 5$ folds. The SOR-only baseline and the $+S$ classifier use CatBoost with depth 4, learning rate 0.05, and ℓ_2 leaf regularization 20; the $+W$ and $+S$+W classifiers use depth 3, learning rate 0.03, and ℓ_2 leaf regularization 50 (the stronger regularization shows their failure is structural, not a tuning artefact).

Comparators. We compare six estimators on identical SOR outputs. **SOR + CatBoost** fits a gradient-boosted classifier on $(x, \bar{p}^{\mathrm{eff}})$ and is the primary baseline. **SOR + CatBoost** $+S$, **SOR + CatBoost** $+W$, and **SOR + CatBoost** $+S + W$ add the proxy variables to that classifier as ordinary covariates (with stronger regularization for the *W* variants). These isolate what the moment structure contributes beyond mere proxy availability. **IPW** reweights a CatBoost

Table 4 Fixed DGP constants (shared across regimes unless noted)

Parameter	Value
Destination price slopes β_d	NYC -1.6, SFO -2.4, BKK -2.3, JNB -2.6, NBO -2.4
Destination WTP locations w_d	NYC 0.55, SFO 0.30, BKK 0.20, JNB 0.10, NBO 0.45
Intercept β_0	-1.05
Latent discount coupling c_δ, cap $d_{\max}$	0.10; 0.30 (Moderate/Shifted), 0.40 (Strong)
Latent s.d. σ_u	1.0
Covariate shares	business: 0.35, connection: 0.30, online channel: 0.70
Customer coefficients γ	business: 0.65, connection: -0.25, online: 0.15, length: 0.05
Flight random-intercept s.d. σ_α	0.2
Number of flights / mean pax per flight	500 / 300
ATPCO base distribution (mean, sd)	NYC 65, 10, SFO 70, 10, BKK 85, 15, JNB 90, 10, NBO 80, 12
FF-tier discount schedule δ^{FF}	Bronze: 0.1, Silver: 0.25, Gold: 0.5
Partner-artefact rate	0.02
Outlier trim	3σ
Operational discount moments	$\mu_f \sim \mathcal{N}(0.03,\ 0.02^2)$, $\sigma_f \sim \lvert\mathcal{N}(0.07,\ 0.01^2)\rvert$
FF tier shares	None: 0.60, Bronze: 0.25, Silver: 0.10, Gold: 0.05
Trip length normal distribution (mean, sd)	8, 3, with minimum of 1
Capacity / load	per-flight capacity; binding on $\approx$ 29–45% depending on regime
Price grid $\mathcal{P}$	$\{0.475, 0.500, \ldots, 1.075\}$ (25 points)

purchase model by the inverse estimated exposure-response propensity, the standard selection-aware benchmark from Sect. Sample selection and missing-not-at-random demand. Our IPW benchmark is a two-stage response-reweighting estimator. Because reliable price exposure is observed essentially only for buyers, we estimate a propensity of observed exposure $\hat{e}(x, S) = \Pr(\text{exposure observed} \mid x, S)$ with a gradient-boosted classifier (buyer status as the exposure-observed label), clip it to [0.05, 1], and form stabilized weights $w_i = \min\{1/\hat{e}_i, 20\}$. A CatBoost purchase model on $(x, \bar{p}^{\,\mathrm{eff}})$ is then fit with these weights. This is a *response-reweighting* benchmark rather than a classical exposure-propensity or doubly robust estimator. Like the Heckman correction it de-biases the outcome model conditional on the data at hand but returns no reconstructed price feature, which is why we treat it as a comparator for the de-biasing objective rather than a substitute for SOR. **SOR + PCI (sieve bridge)** is the proposed estimator.

Evaluation metrics

We report the calibration loss over binned prices. More specifically, let $\mathrm{price\_bin}$ be created by discretizing the effective normalized price p_eff into *N* bins. For each bin *b*, let $\widehat{\mathrm{pred}}_b$ be the mean predicted probability in that bin, and act_b the observed purchase rate in bin *b*. The *calibration loss* is the average of squared bin-wise gaps:

$$\mathcal{L}_{\mathrm{cal}} = \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \left(\widehat{\mathrm{pred}}_b - \mathrm{act}_b\right)^2.$$

Observed calibration measures agreement between predicted probabilities and realised purchases on the test log. Because the test log itself reflects the latent confounder, observed calibration can reward a model that fits the confounded conditional. We therefore additionally report the mean squared error between each estimator's predicted probability and the simulator's true probability $\pi^*(x, p)$, under both the full and deconfounded oracles,

$$\mathrm{MSE}_\pi = \frac{1}{N} \sum_{i=1}^{N} \left(\widehat{\pi}(x_i, p_i) - \pi^*(x_i, p_i)\right)^2, \tag{8}$$

and the downstream relative revenue of Sect. Results. These three lenses - observed calibration, probability error against truth, and revenue - do not need to agree, and disagreement between them is itself informative.

Experimental set-up

For each seed we draw an operational log, split it 80/20 into train and test before any reconstruction, estimate SOR moments on the training log only, and reconstruct both splits with the training-estimated discount table. All purchase models are trained on the training split and evaluated on the held-out test split. In the shifted regime, train and test are drawn from configurations with opposite-signed latent means, as described in Sect. An unobserved flight-level confounder and the evaluation regimes. In the other regimes a single configuration is split. We repeat the full procedure over 100 seeds and report means with standard deviations and paired tests across seeds. The comparators are the six estimators of Sect. Training objectives.

Results

We evaluate each estimator on three lenses: observed calibration on the held-out test log, mean squared error against the simulator's true purchase probability under both oracles, and oracle-relative revenue. The revenue evaluation iterates over the price grid $\mathcal{P}$ and selects, for each estimator, the price maximizing estimated expected revenue at three decision granularities: a single Global price, one price per Route, and a Passenger-level price. As Sect. Price optimization results shows, the three lenses do not agree, and that disagreement is one of the interesting empirical findings of this section.

Price optimization performance

Relative revenue definition. For a given estimator $\hat{\pi}$ and decision granularity (global, route, or passenger), let $p_i^{\hat{\pi}}$ denote the price assigned to passenger *i* by maximizing estimated expected revenue over the price grid $\mathcal{P}$:

$$p_i^{\hat{\pi}} \;=\; \arg\max_{p \in \mathcal{P}} \; p \cdot \hat{\pi}(x_i, p),$$

where for global (resp. route) level the same price is chosen for all passengers (resp. all passengers on the same route). The oracle selects prices using the true purchase probability π^*:

$$p_i^* \;=\; \arg\max_{p \in \mathcal{P}} \; p \cdot \pi^*(x_i, p).$$

Expected revenue under any policy is evaluated against the true purchase function:

$$R(\hat{\pi}) \;=\; \frac{1}{N} \sum_{i=1}^{N} p_i^{\hat{\pi}} \cdot \pi^*(x_i,\, p_i^{\hat{\pi}}),$$

$$R^* \;=\; \frac{1}{N} \sum_{i=1}^{N} p_i^* \cdot \pi^*(x_i,\, p_i^*).$$

We report the *relative revenue* $R(\hat{\pi})/R^*$ as the relative share of revenue captured of the total revenue possible.

Price optimization results

Three patterns are visible, and all are stable across 100 seeds (Table 5; paired *t*-tests below). First, in the moderate and strong regimes the proximal bridge is slightly better than the direct SOR+CatBoost baseline, however the effect is limited. The deconfounded price effect and the confounded conditional recommend similar prices when the confounder distribution is the same at train and test. Second, under confounder *distribution shift* the proximal bridge separates more clearly from the alternatives, recovering about 90–93% of oracle revenue against about 86% for SOR+CatBoost and IPW. Third, and most striking, adding the proxy variables to a flexible classifier as ordinary covariates is actively harmful. SOR+CatBoost+W and SOR+CatBoost+S+W collapse to 74–84% of oracle revenue in every regime, choosing a near-fixed price (mean absolute normalized price error ≈ 0.48, almost never within two grid points of the oracle price). Heavy regularization does not rescue these models, which indicates the failure is a structural effect of including these variables rather than a matter of capacity control. The $+W$ and $+S+W$ columns are merged together, since both estimators collapse to the same fixed grid price in essentially every seed and are therefore reported together. Their probability surfaces differ slightly (their observed calibration losses in Table 6 differ), but the degenerate argmax is identical, so the realized revenue under π^* is too.

The collapse is invisible to observed calibration and is the reason we report probability error against truth. Table 6 shows

Table 5 Relative revenue against the (full) oracle, mean ± s.d. over 100 seeds, at three decision granularities. Higher is better; 1.0 is the oracle. The right-most column adds the proxy variables to a gradient-boosted classifier as ordinary covariates, $+W$ and $+S+W$ are reported in a single column because both collapse to the same degenerate grid price in essentially every seed (see text)

Regime	Level	SOR+CatBoost	SOR+PCI	IPW	+S	$+W$ / $+S+W$
Moderate	Global	0.932 ± 0.090	$\mathbf{0.975 \pm 0.033}$	0.915 ± 0.093	0.902 ± 0.101	0.783 ± 0.009
	Route	0.935 ± 0.047	$\mathbf{0.954 \pm 0.029}$	0.911 ± 0.057	0.925 ± 0.057	0.764 ± 0.007
	Passenger	0.921 ± 0.033	$\mathbf{0.935 \pm 0.023}$	0.904 ± 0.035	0.914 ± 0.035	0.756 ± 0.007
Strong	Global	0.995 ± 0.006	0.995 ± 0.009	0.987 ± 0.033	0.993 ± 0.009	0.843 ± 0.011
	Route	0.982 ± 0.012	0.984 ± 0.008	0.973 ± 0.027	0.979 ± 0.016	0.823 ± 0.009
	Passenger	0.954 ± 0.009	$\mathbf{0.959 \pm 0.008}$	0.945 ± 0.017	0.951 ± 0.012	0.801 ± 0.007
Shifted	Global	0.858 ± 0.095	$\mathbf{0.927 \pm 0.022}$	0.860 ± 0.098	0.830 ± 0.094	0.760 ± 0.009
	Route	0.870 ± 0.031	$\mathbf{0.908 \pm 0.019}$	0.866 ± 0.042	0.861 ± 0.038	0.742 ± 0.008
	Passenger	0.866 ± 0.025	$\mathbf{0.896 \pm 0.016}$	0.865 ± 0.028	0.854 ± 0.028	0.737 ± 0.007

Table 6 The calibration paradox (moderate and shifted regimes, mean over 100 seeds). Observed calibration loss disagrees with probability error against the deconfounded oracle and with revenue. Lower is better for the first two columns; higher is better for revenue

Regime	Method	Observed cal. loss	Deconf. prob. MSE	Rel. revenue (Global)
Moderate	SOR+CatBoost	3.7×10^{-5}	3.2×10^{-3}	0.932
	SOR+PCI (bridge)	3.4×10^{-5}	3.2×10^{-3}	$\mathbf{0.975}$
	IPW	4.6×10^{-5}	3.7×10^{-3}	0.915
	SOR+CatBoost+W	4.3×10^{-5}	5.0×10^{-3}	0.783
	SOR+CatBoost+S+W	4.3×10^{-5}	5.1×10^{-3}	0.783
Shifted	SOR+CatBoost	8.2×10^{-4}	3.5×10^{-3}	0.858
	SOR+PCI (bridge)	8.2×10^{-4}	3.6×10^{-3}	$\mathbf{0.927}$
	IPW	6.7×10^{-4}	3.7×10^{-3}	0.860
	SOR+CatBoost+W	$\mathbf{2.9 \times 10^{-4}}$	5.0×10^{-3}	0.760
	SOR+CatBoost+S+W	$\mathbf{2.6 \times 10^{-4}}$	5.1×10^{-3}	0.760

the moderate and shifted regimes. In the moderate regime, observed calibration is essentially tied across all methods (all $\approx 4 \times 10^{-5}$), yet the naive-proxy models already lose 15 percentage points of revenue and carry roughly 60% higher probability error against the deconfounded oracle - calibration is silent about a gap that is already large. Under shift the disagreement becomes an inversion: the naive-proxy models attain the *best* observed calibration loss on held-out data, yet the *worst* error against the true probability and the worst revenue. They are well calibrated to a confounded conditional that does not transport, and the argmax over the price grid amplifies that error into a degenerate policy. The proximal bridge, by contrast, is best (or tied) on probability error against the deconfounded oracle and best on revenue, while its observed calibration is only middling. This supports our claim that the structural use of proxies, not their availability, is what matters.

Significance. Paired t-tests across the 100 seeds, comparing relative revenue against SOR+CatBoost, confirm the ranking. In the moderate regime the proximal bridge gains $+4.3$ percentage points of oracle revenue at the global level (mean paired difference $+0.043$, $p < 0.001$), while SOR+CatBoost+W and $+S$+W lose about 15 percentage points (-0.149, $p < 0.001$). Under distribution shift the bridge gains $+6.9$ percentage points at the global level and $+3.0$ at the passenger level (both $p < 0.001$), and the naive-proxy models again lose about 10 percentage points ($p < 0.001$). IPW is not significantly different from the baseline ($p \approx 0.11$ at the global level). On probability error against the deconfounded oracle the bridge ties or beats the baseline, while the naive-proxy models are significantly worse ($p < 0.001$), consistent with the paradox of Table 6.

Interpretation

The results support the bias-variance reading of the proximal correction. When the exposure mechanism is stable between training and deployment, the confounded conditional and the deconfounded price effect recommend similar prices, so a direct classifier and the proximal bridge perform comparably and the extra machinery buys little to none. When the exposure mechanism shifts - the regime airlines would normally face as discount policy, partner mix, and operational pressure drift over time - the deconfounded price effect is what can be transported to a new regime, and the proximal bridge recovers this. The naive-proxy baselines make the cost of the wrong tool very visible. Proxy information is genuinely useful, but used as ordinary features it is absorbed in a non-causal way that the downstream argmax turns into a revenue-destroying policy.

When to use PCI refinement. The results suggest clear guidance for practitioners deciding between SOR-only and SOR+PCI estimation:

1. **Use SOR-only** when the operational discount mechanism is stable and well documented and the deployment population matches the training population. Here the bridge adds little and the simpler estimator suffices.
2. **Use SOR + PCI (sieve bridge)** when the exposure mechanism may drift between training and deployment, when partner or codeshare effects introduce unmeasured price variation, or whenever the deployment regime is not guaranteed to match the training regime—precisely the conditions under which the proximal bridge's robustness advantage materialises.
3. **Do not add proxy variables to a flexible demand model as ordinary covariates.** Doing so improves observed calibration but yields a non-causal surface and a degenerate pricing policy. If the proxy information is to be used at all, it must be used through the bridge's moment structure.

Table 7 Implied arc elasticity at the revenue-maximizing price of each destination-level demand curve, by destination, mean ± s.d. over 100 seeds (Moderate and Shifted regimes), against the deconfounded oracle. The true elasticity at the revenue-maximizing price is -1 by construction (unconstrained revenue maximization implies unit elasticity), so the check is whether each estimator's chosen price sits at a unit-elastic point of the true demand curve; values collapsing toward 0 are the demand-curve signature of a degenerate fixed-price policy

Regime	Destination	True (deconf.)	SOR+CatBoost	SOR+PCI
Moderate	NYC	-1.00	0.00 ± 0.00	-0.91 ± 0.18
	SFO	-1.02	-0.87 ± 0.45	-0.99 ± 0.05
	BKK	-1.01	-0.91 ± 0.42	-1.00 ± 0.04
	JNB	-1.05	-1.07 ± 0.35	-1.00 ± 0.03
	NBO	-0.99	-0.61 ± 0.57	-0.97 ± 0.13
Shifted	NYC	-1.00	-0.04 ± 0.26	-1.00 ± 0.04
	SFO	-1.01	-0.85 ± 0.50	-1.00 ± 0.02
	BKK	-1.01	-0.98 ± 0.39	-1.00 ± 0.02
	JNB	-1.05	-1.05 ± 0.25	-1.00 ± 0.03
	NBO	-0.99	-0.10 ± 0.34	-1.00 ± 0.03

More generally, PCI is most valuable when the bias introduced by unmeasured confounders exceeds the variance introduced by the additional model complexity.

Implied price elasticities

Since this is a synthetic study, we do not claim to measure real-world ancillary-seat elasticities, which depend heavily on product, route, and customer mix and are best estimated per application. Instead, we ask whether each estimator recovers the planted causal slope. Using the deconfounded oracle, we compute the true arc elasticity at the revenue-maximizing price and compare it with the elasticity each estimator implies at its own chosen price, per destination, averaged over 100 seeds (Table 7). The proximal bridge tracks the true elasticities in sign and approximate magnitude. The direct classifier does so in distribution but is distorted under confounder shift, and the naive-proxy models collapse toward a near-zero implied elasticity, leading to the demand-curve signature of the fixed-price policy of Sect. Price optimization results. Because the optimization is unconstrained revenue maximization over the grid, demand at a correctly chosen price is unit-elastic by construction. The managerially relevant signal is therefore the *deviation* from -1. The direct classifier recovers approximately unit-elastic prices for most destinations in-regime, but already shows a degenerate implied elasticity for NYC in the moderate regime - its chosen price sits on a flat segment of its estimated curve - and the degeneracy spreads to NBO under confounder shift. The naive-proxy models collapse toward a near-zero implied elasticity everywhere, the demand-curve signature of the fixed-price policy of Sect. Price optimization results.

Discussion

SOR translates operational knowledge (ATPCO base, load-based and check-in markdowns, FF tiers, partner artifacts) into a clean, auditable reconstruction of seen prices based on known factors. This resolves the core issue created by censored logs and yields a dataset suitable for learning price sensitivity. By modeling $\pi(x,p)$ rather than an unconditional curve, elasticities vary with itinerary, timing, capacity indicators, and customer attributes. However, SOR only corrects for known distortions in the data. By adding the proximal causal inference step on top of the outcome of SOR, we also correct for unknown perturbations in the data, and obtain a robust estimator for the probability of purchase $\hat{\pi}(x,p)$.

With calibrated $\hat{\pi}(x,p)$ on a discrete grid and business-consistent bounds on prices, optimizing the price becomes a simple argmax operation, or more advanced methods can be built on top of this.

Computational cost. The proximal refinement is light. The bridge is a pair of ridge regressions on a polynomial sieve, fit K times under cross-fitting; each fit is a closed-form linear solve and is cheaper than the gradient-boosted SOR-only baseline. The dominant cost remains the SOR reconstruction itself (per-flight aggregation and discount imputation). In our experiments the full cross-fitted bridge completes in well under a minute on a standard workstation, so the refinement imposes negligible additional cost when deployed alongside SOR.

Distributional assumptions in SOR. The SOR pipeline models per-flight discounts as normally distributed after outlier filtering (Sect. Step 1: Discount reversal and outlier filtering). This assumption is a pragmatic choice justified by the cleaning step: once check-in markdowns, partner artifacts, and 3σ outliers are removed, the remaining operational discounts cluster around a modest mean with moderate spread, and the Normal fit provides an adequate summary. However, in settings where discount mechanics produce genuinely bimodal distributions - for example, routes with both a standard discount tier and a promotional campaign running simultaneously - a mixture model or nonparametric kernel estimate for Δ_f may better capture the exposure distribution. This would affect the imputed prices for non-buyers in Step 3 and could improve calibration on such routes. The bridge provides some robustness to mild deviations such as the moderate right-skew we observe in cleaned real data (Fig. 3). It does not, however, extend to unflagged multi-modality. Appendix A shows that when discounts come from two well-separated modes with no observable flag, both the single-Normal and a correctly specified mixture imputation lose substantial revenue, because the error enters the treatment variable itself. The boundary of applicability is therefore operational: modes must be resolvable from logged metadata, as the check-in and partner modes in Fig. 1 are. Therefore, practitioners should inspect the discount distribution per route family and consider richer imputation models when clear non-normality is present, ensuring the relevant indicators are captured in the model.

Approximate exclusion and proxy validity. The proxy conditions in Assumption 6.1 require that upstream variables S affect purchase only through price exposure. In practice, this exclusion is approximate: variables such as route characteristics may correlate with unobserved demand heterogeneity beyond their role in the pricing mechanism. For instance, flights to leisure destinations may have systematically different purchase rates even at the same effective price. The PCI framework remains useful under approximate exclusion provided that the dominant confounding channel runs through price exposure, which is the case when operational discount rules are the primary source of endogeneity. The empirical results support this interpretation. Practitioners should focus the S specification on variables with a clear operational pathway to price exposure (partner flags, capacity indicators, temporal discount triggers) and avoid including variables whose primary association with purchase runs through demand rather than supply. A second validity question concerns the proximal pair itself: if the downstream proxy were a function of the upstream set, the bridge would degenerate. Our leave-one-out construction guards against this, and the diagnostic of Sect. Proxy diagnostics confirms that roughly four-fifths of the variation in W is not explained by S, so the pair supplies genuinely complementary information about the latent exposure.

Cross-product interactions and capacity constraints. The current framework treats each seat product in isolation: the purchase probability $\pi(x, p)$ models a single product without accounting for substitution between seat types (e.g., extra-legroom versus standard window seats) or complementarities with other ancillaries (e.g., baggage, lounge access). In practice, a passenger deciding whether to purchase a premium seat may trade off against other spending options, and the availability of cheaper alternatives affects the effective price elasticity. Extending the framework to a multi-product setting would require joint modeling of $\pi(x, \mathrm{p})$ where p is a vector of prices across products, and SOR would need to reconstruct exposure for each product simultaneously. Similarly, capacity constraints are omitted to isolate the WTP estimation contribution: the argmax over the price grid treats each passenger independently, without accounting for the finite number of seats in each category. In a full deployment, $\hat{\pi}(x, p)$ would serve as the demand input to a capacity-constrained optimization that jointly allocates seats and sets prices. The WTP estimates produced by our pipeline are directly compatible with such downstream optimization, since they provide the probability of purchase at any price in the grid for each passenger profile.

Deployment considerations. Several practical factors affect the performance of the pipeline in production. First, the flight-level moments $(\hat{\mu}_f, \hat{\sigma}_f)$ require sufficient buyer observations per flight to be estimated reliably. For cold-start routes or new seasonal schedules with few historical transactions, pooling moments across similar routes (by destination region, aircraft type, or haul length) can provide reasonable priors until flight-specific data accumulates. And related to this, the downstream proxy W is built from leave-one-out within-flight statistics of the historical log and is therefore available offline, where the WTP estimation problem lives. When scoring a future flight, W is naturally replaced by the accumulated within-flight statistics observed so far in the booking horizon, or by the historical statistics of comparable flights (route, season, departure time) for flights with no bookings yet; the bridge itself is fit offline either way. Second, the pipeline should be re-estimated periodically to reflect changes in discount policy, loyalty program rules, and market conditions. The frequency depends on the volatility of the operational environment. Third, the SOR reconstruction assumes that the business rules used for discount reversal (Sect. Step 1: Discount reversal and outlier filtering) are correctly specified. Errors in the rule specification - for example, an outdated FF discount table - propagate through the entire pipeline. Regular validation of the rule set against operational documentation is therefore important, and discrepancies between SOR-predicted and observed discount patterns can serve as an early warning of rule drift. Finally, logging offered-price exposures directly in production systems, rather than only transaction prices, would substantially reduce the censoring problem and allow the PCI refinement to be validated against ground-truth exposure data.

Limitations. Beyond the assumptions discussed above, several limitations should be noted. The evaluation is conducted on synthetic data designed to mimic realistic airline pricing mechanics. While the data-generating process incorporates the key features of real-world seat pricing (ATPCO bases, FF discounts, operational markdowns, MNAR censoring), it necessarily simplifies some aspects of the full operational environment, such as multi-leg itineraries with different seat products on each leg, group bookings with coordinated seat selection, and time-varying competitive dynamics. Validation on proprietary airline transaction data is a natural next step. Additionally, the PCI framework assumes that the proxy pair (S, W) is sufficiently rich to identify the latent exposure effect. If unmeasured confounders exist that are orthogonal to both operational proxies and the SOR-reconstructed price - for example, A/B testing artifacts or display effects that vary across booking platforms - the bridge function cannot correct for them. In such cases, the SOR-only estimator may be preferred for its lower variance, and the practitioner guidance in Sect. Interpretation applies.

Bounds on seen prices serve as pragmatic guardrails, but are no substitute for explicit capacity-aware optimization in a full revenue management system.

Conclusion

We addressed ancillary seat pricing under censored exposure by combining an operationally grounded reconstruction with a statistically grounded refinement step on top. The *Synthetic Offer Reconstruction* pipeline maps ATPCO base prices, loyalty rules and operational markdowns to an exposure variable suitable for learning. On top of this, a *PCI-based WTP estimation* estimates latent offers and demand aligning predictions with outcomes conditional on operational proxies. We show that using this combined approach yields stable and reliable estimates for the probability of purchase given input characteristics. Beyond the estimator itself, our evaluation carries a caution for pricing practice: observed calibration on a censored, confounded log is not a sufficient acceptance criterion for a pricing model, since the best-calibrated models in our study produced the worst pricing policies. Decision-quality checks - robustness to regime change and, where available, agreement with a deconfounded target - should accompany calibration in model validation.

A natural next step is to integrate these elasticities into a revenue-management framework that explicitly models capacity, for example by embedding $\hat{\pi}(x, p)$ in a capacity-constrained control, and validating the behavior online. Logging offered-price exposures in production will further strengthen and accelerate learning.

Appendix A

Sensitivity to the SOR discount-imputation model

SOR Step 2 models per-flight discounts as Normal after outlier filtering. To test whether this matters when the truth is non-Normal, we generate a regime with genuinely bimodal flight-level discounts (a standard tier and a simultaneous promotional tier for instance) and compare two imputation choices in Step 2: the baseline single Normal and a two-component Gaussian mixture. Table 8 reports calibration loss and oracle-relative revenue for both, over 50 seeds. The experiment tests whether a richer marginal imputation model restores performance when the unimodal assumption fails outright.

Both imputation choices lose a large share of oracle revenue ($0.56 \sim 0.57$ of oracle, versus $0.93 \sim 0.99$ in the unimodal regimes of Table 5), and the richer mixture imputation does not repair the loss. The reason for this is a relevant insight in itself. Step 3 imputes each passenger's exposure as an *independent draw* from the fitted distribution. When the post-cleaning distribution is unimodal and tight, an independent draw is close to the truth for almost every passenger, so the price feature carries little measurement error. When the distribution has two well-separated modes and no observable variable indicates which mode applied to which passenger, an independent draw assigns a relevant fraction of passengers to the wrong mode regardless of how well the marginal distribution is matched. The failure is one of individual-level assignment, not of distributional shape, so matching the marginal - the Gaussian-mixture imputation - cannot help. The resulting measurement error in the price feature attenuates the estimated price response, and the revenue argmax turns the attenuated curve into a poorly chosen price. The proximal bridge cannot repair this either. The corruption sits in the treatment variable itself rather than in confounding between exposure and demand, which is outside the bridge's moment structure.

Notably, observed calibration remains excellent throughout ($\approx 2 \times 10^{-5}$). When imputed prices are nearly independent of true prices, the observed purchase rate is almost flat across imputed-price bins, and a near-flat predictor matches it perfectly. This is the calibration paradox of Sect. Price optimization results in a different setting: a pricing model can be impeccably calibrated on the log it is given, while carrying no usable price signal.

The practical implication is a boundary-of-applicability statement rather than a tuning recommendation. Multimodal discount regimes must be resolved *before* imputation by observable operational metadata - a promotion flag, a campaign date range, a channel identifier - exactly as Step 1 already resolves the heavy check-in and partner modes visible in Fig. 1 using their flags. When the modes are observable, each component reduces to the unimodal case and the pipeline applies unchanged. When they are not, no marginal imputation model, however rich, can recover individual exposure, and direct logging of offered prices (see Deployment considerations, Sect. Discussion) is the only complete remedy.

Table 8 Bimodal-discount regime (50 seeds): single-Normal vs. two-component Gaussian-mixture imputation in SOR Step 2, for the SOR+CatBoost baseline and the proximal bridge. Calibration loss on the held-out test log and global and route oracle-relative revenue, mean ± s.d

Imputation	Estimator	Calibration loss	Rel. revenue (Global)	Rel. revenue (Route)
Single Normal	SOR+CatBoost	2×10^{-5}	0.57 ± 0.01	0.56 ± 0.01
	SOR+PCI	2×10^{-5}	0.56 ± 0.01	0.56 ± 0.004
Gaussian mixture	SOR+CatBoost	2×10^{-5}	0.57 ± 0.02	0.56 ± 0.01
	SOR+PCI	2×10^{-5}	0.56 ± 0.01	0.56 ± 0.01

Author contributions KD: methodology, model development, data analysis, simulations, writing original draft. TL: methodology, model development, data analysis. GK and RM: supervision, guidance, review & editing. All authors contributed to the interpretation of results and reviewed and approved the final manuscript.

Funding This research was supported by the institutions affiliated with the authors, with no external funding received.

Data availability All data used in this study are synthetically generated. The simulation framework, including code required to reproduce the data and experiments, is described in the paper and publicly available. We further ground the synthetic generation in real-life data, for which we can only produce aggregate statistics.

Declarations

Competing interests The authors declare no competing interest.

References

Kevin Duijndam(2026). *Code for SOR + PCI notebook*. https://github.com/KevinDuijndam/SOR-PCI-code/blob/c2e46155a3e3c8d5bae70b37fe531ef56a194d6e/SOR%20%2B%20PCI%20notebook.ipynb
von Lücken, Tim. 2023. Airplane Seat pricing: Prediction and Optimization based on historical data, MSc thesis. Vrije Universiteit Amsterdam .
Bu, Jinzhi, David Simchi-Levi, and Li. Wang. 2023. Offline Pricing and Demand Learning with Censored data. *Management Science* 69 (2): 885–903.
Chen, Boxiao, Xiuli Chao, and Yining Wang. 2020. Data-based Dynamic Pricing and Inventory Control with Censored Demand and Limited Price Changes. *Operations Research* 68 (5): 1445–1456.
Chernozhukov, Victor et al. 2018. Double/debiased machine learning for treatment and structural parameters.
Cui, Yifan, et al. 2024. Semiparametric Proximal Causal Inference. *Journal of the American Statistical Association* 119 (546): 1348–1359.
Duijndam, Kevin, Koole, Ger, and van der Mei, Rob. 2025. Bayesian Reinforcement Learning to Optimize Paid Ancillary Revenue in the Airline Industry. In: *Journal of Revenue and Pricing Management*, pp. 1–17.
Gallego, Guillermo, and Masoud Talebian. 2012. Demand Learning and Dynamic Pricing for Multi-Version Products. *Journal of Revenue and Pricing Management* 11 (3): 303–318.
Heckman, James J. 1979. Sample selection bias as a specification error. *Econometrica: Journal of the econometric society* 153–161.
Horvitz, Daniel G., and Donovan J. Thompson. 1952. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association* 47 (260): 663–685.
Long, Zheng, and Peter Belobaba. 2024. Airline revenue management with segmented continuous pricing: methods and competitive effects. *Journal of Revenue and Pricing Management* 23 (3): 185–202.
Macrakis, Stavros, Michael Leen, and Daniel Burkhard. 2016. Pricing Airline Optional Services, Technical Disclosure Commons.
Mumbower, Stacey. 2025. Emerging research on airline ancillaries: a review of offer management and dynamic pricing literature. *Journal of Revenue and Pricing Management* 24:1–15.
Perera, Suranga, and David Tan. 2019. In search of the Right Price for Air Travel: First Steps towards Estimating Granular Price-Demand Elasticity. *Transportation Research Part A: Policy and Practice* 130:557–569.
Prokhorenkova, Liudmila. 2018. CatBoost: Unbiased Boosting with Categorical Features. In: *Advances in Neural Information Processing Systems*, 31.
Robins, James M., Andrea Rotnitzky, and Lue Ping Zhao. 1994. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association* 89 (427): 846–866.
Sachs, Anna-Lena., and Stefan Minner. 2014. The Data-Driven Newsvendor with Censored Demand Observations. *International Journal of Production Economics* 149:28–36.
Shen, Tao, and Yifan Cui. 2025. Proxy-Aided Demand Learning with an Application to Various Pricing Problems. Operations Research.
Shukla, Naman et al. 2019. Dynamic Pricing for Airline Ancillaries with Customer Context. In: *Proceedings of the 25th ACM SIGKDD International Conference on knowledge discovery & data mining*, pp. 2174–2182.
Talebian, Masoud, Natashia Boland, and Martin Savelsbergh. 2014. Pricing to Accelerate Demand Learning in Dynamic Assortment Planning for Perishable Products. *European Journal of Operational Research* 237 (2): 555–565.
Tang, Jingwen et al. 2025. Offline Feature-Based Pricing under Censored Demand: A Causal Inference Approach. *Manufacturing & Service Operations Management* 535–553.
Tchetgen, Eric J Tchetgen et al. 2024. An Introduction to Proximal Causal Inference. *Statistical Science* 39 (3): 375–390.
Tchetgen, Eric J Tchetgen et al. 2020. An Introduction to Proximal Causal Learning. In: *arXiv preprint* arXiv:2009.10982.
Thirumuruganathan, Saravanan, et al. 2023. Will they take this offer? A Machine Learning Price Elasticity Model for Predicting Upselling Acceptance of Premium Airline Seating. *Information & Management* 60 (3) : 103759.
Wang, Xinwei, Michael Wittman, and Thomas Fiig. 2023. Dynamic offer creation for airline ancillaries using a Markov chain choice model. *Journal of Revenue and Pricing Management* 22 (4): 276–290.
Warnock-Smith, David, John F. O'Connell, and Mahnaz Maleki. 2017. An Analysis of Ongoing Trends in Airline Ancillary Revenues. *Journal of Air Transport Management* 64:42–54.
Yang, Yang, Wan-Ling. Chu, and Cheng-Hung. Wu. 2022. Learning Customer Preferences and Dynamic Pricing for Perishable Products. *Computers & Industrial Engineering* 171 : 108440.
Zhao, Guihong, Yue Cui, and Shaoyu Cheng. 2021. Dynamic Pricing of Ancillary Services based on Passenger Choice Behavior. *Journal of Air Transport Management* 94 : 102058.