



Fare families pricing optimisation

Closing the efficiency gap in fare family buy-up pricing

Most airlines operate on **Static Fare Ladders**: fixed price increments between (e.g., Basic / Standard / Flex) products based on legacy business rules. While simple to manage, these static gaps do not reflect the actual price sensitivity of the passenger on a specific route, season, or day.

Why fixed fare family gaps leave value on the table

	Static Fare Families (Legacy)	WTP-Driven Fare Families (ADC)
Logic	Fixed increments (e.g., always +\$55)	Data-driven "buy-up" differentials
Foundation	Historical averages & legacy rules	Scientific experimentation & elasticity
Granularity	Uniform across the entire network	Route, season, and segment specific
Objective	Simplification of filing	Capture maximum consumer surplus
Upsell Risk	High risk of "buy-down" or missed upsell conversion	Optimised "Nudge" towards higher tiers

WTP-Driven Pricing replaces these fixed rules with evidence-based increments. By understanding the true willingness to pay for the next tier of service, you maximise both conversion and yield.

The opportunity

Optimising the price gaps between your fare families is one of the highest impact opportunities in Revenue Management, with relatively low operational overhead.

Conversion Optimisation: Large price gaps discourage upselling; gaps that are too small fail to maximise margin. We find the optimal balance for every route.

Revenue uplift: WTP-based optimisation typically delivers ~1%–4% revenue improvement versus traditional business rules pricing.

Our transition framework

We don't just replace one static list with another. We build the capability to find these optimal points continuously.

- 1. Identify the bias:** We analyse your current "buy-up" rates to identify where legacy gaps are currently suppressing demand or diluting yield.
- 2. Measure the elasticity:** Through controlled experiments, we determine the true elasticity of your fare families, moving beyond legacy rules and historical averages.
- 3. Calibrate the differentials:** We apply models that recommend the optimal WTP-driven price increments for your specific commercial context (e.g., business vs. leisure routes).

Built for operational reality

This approach is **WTP-driven, not necessarily high-frequency dynamic**. It integrates seamlessly with:

- **ATPCO Filing:** Recommendations are designed to be compatible with standard fare structures.
- **RMS Integration:** Our WTP forecasting works with your existing bid-price or class-based systems.
- **Regulatory Guardrails:** Ensures transparency and consistency in your public-facing offers.

Own the science, not just the results

At the end of our engagement, ADC transfers the full solution to your team. You own the **elasticity models** and the **decision logic** with full transparency into the reasoning, plus an **enterprise-grade codebase** and documentation. This ensures you maintain optimised fare families as market conditions shift in 2026 and beyond.

Ready to unlock the revenue hidden in your fare families?

Let's scope a proof-of-concept for your highest-impact routes

Frequently asked questions (FAQ)

What is the level of granularity for these price optimisations?

The ADC engine is designed for maximum flexibility. Optimisation can be applied at the **individual flight level** or the **O&D (Origin & Destination) level** to capture hyper-local demand. However, for markets with lower density, the system can automatically aggregate data into specific market groups or regional clusters, ensuring statistical significance while still outperforming "one-size-fits-all" network rules

How does this handle competitive price matching?

Our model treats competitor pricing as a contextual variable. It calculates the optimal "buy-up" increment relative to the market, ensuring you don't lose volume while avoiding unnecessary price wars.

Will experimentation hurt our current revenue?

No. Unlike traditional A/B testing, our adaptive testing techniques (such as Multi-Armed Bandit) minimize revenue risk. It quickly abandons price points that don't convert, protecting your baseline while it learns.

How does this integrate with our existing RMS?

The engine is RMS-agnostic. It takes your base fare or bid price as an input and provides the optimal "add-on" values for your fare families, filing them back into your system or injecting them into your NDC offer engine.

Can we use this for both Business and Leisure routes?

Yes. The model identifies different elasticity curves for different personas, allowing for wider gaps on high-yield business routes and more aggressive, conversion-focused increments on leisure segments.

Let's talk

Leading airlines are shifting from gut instinct to evidence-led decisions. Curious what this looks like for your network? Let's start the conversation.



Vladimir Antsibor
Commercial Lead



Joël Gastelaars
Practice Lead