AIRFRANCEKLM GROUP

PATAT 2026

TOWARDS RESILIENT AIRLINE PLANNING

FROM FLEET PLANNING TO OPERATIONS CONTROL

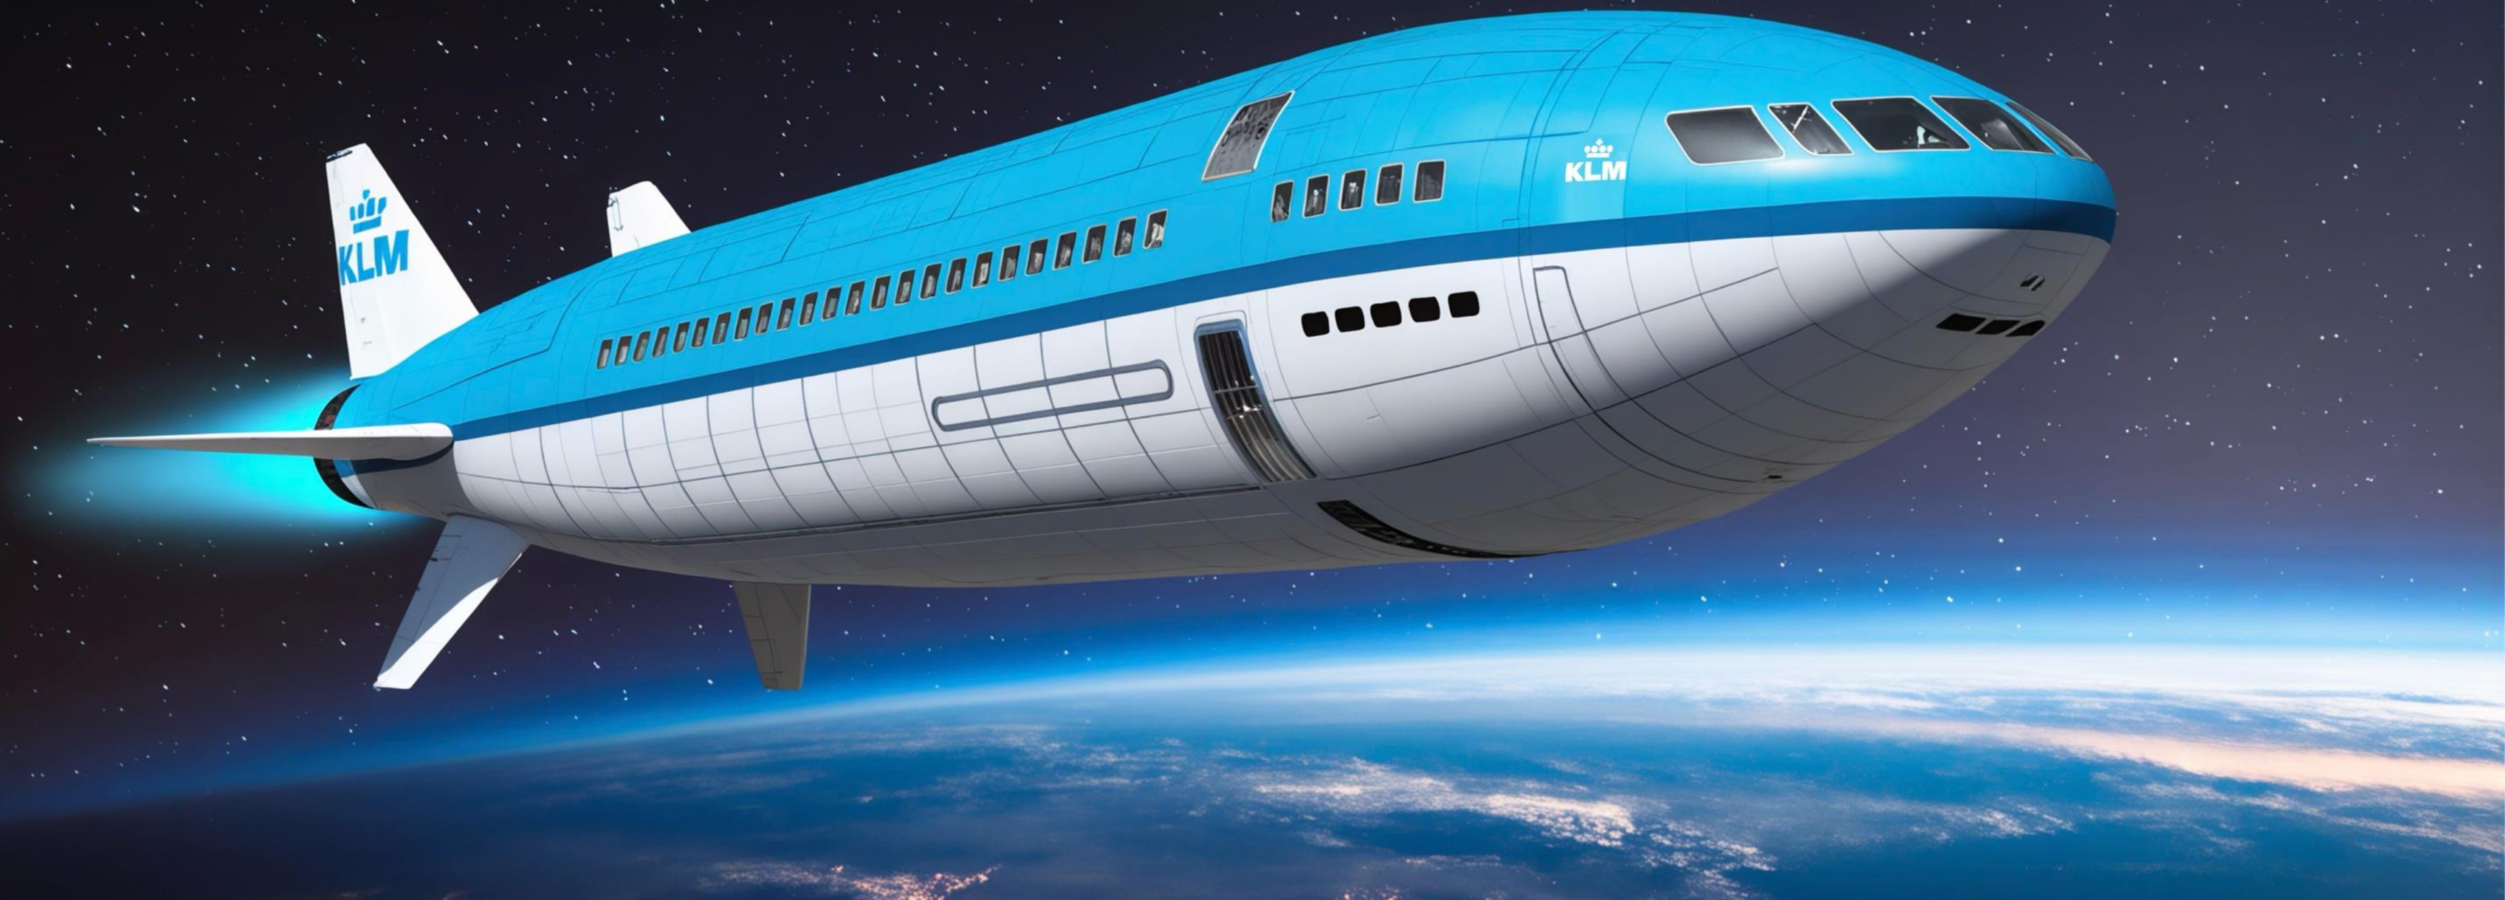

Dr. Jeroen Mulder | Innovation Lead @ KLM Bluelabs

August 2026 | Nottingham

TOWARDS RESILIENT AIRLINE PLANNING

INSIDE AIR FRANCE-KLM
WHO ARE WE?

✓ **AIR FRANCE-KLM is a leader in international air transport departing from Europe**

- We are a major global player with 3 areas of expertise: Air passenger transport, Cargo transport, and Aircraft maintenance.
- We have a fleet of 564 aircraft, divided between Air France, KLM Royal Dutch Airlines, and Transavia.
- We provide services to up to 320 destinations in 100 countries,
- Mainly from our central hubs in Paris-Charles de Gaulle and Amsterdam-Schiphol.

✓ **Our objective (-function) is**

- To be the most **customer-centric**, innovative and **efficient** European network carrier
- Design the **most attractive** Air France-KLM network including our partners, for **all stakeholders**, …
- … based on the **optimal** match between **commercial wishes** & **operational capabilities**.

AIRFRANCEKLM GROUP

TOWARDS RESILIENT AIRLINE PLANNING

INSIDE AIR FRANCE-KLM
AIRLINE PLANNING PROCESS *)

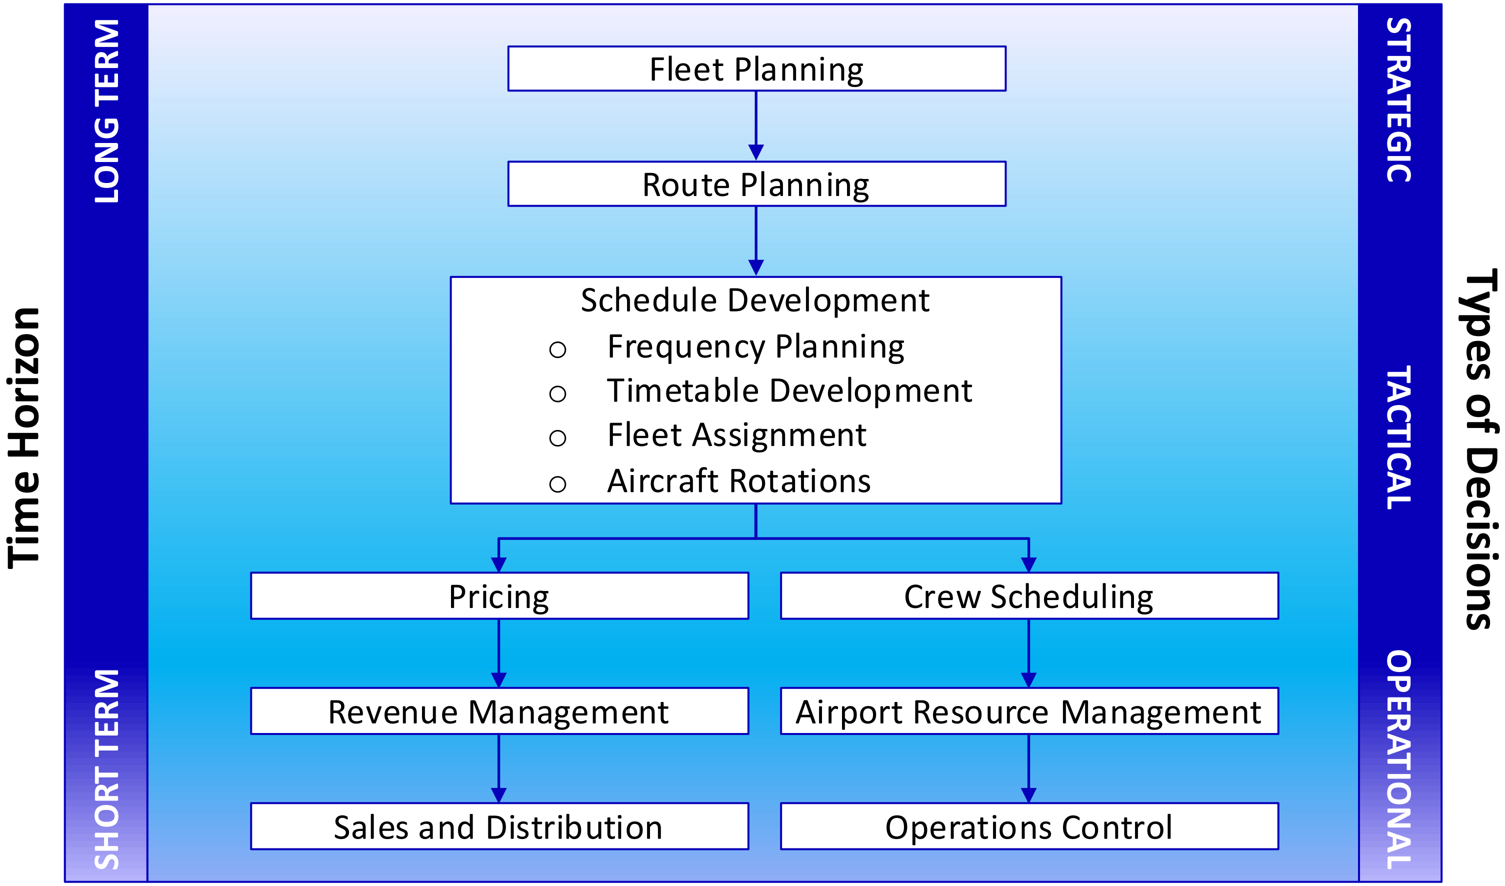

Source: Cynthia Barnhart

- ✓ **Challenges**
  - Cross-silo optimization
  - Organizational structures
  - Huge amounts of data
  - Data quality
  - Distinction noise vs patterns
  - Scalable AI/ML/OR

*) From **The Global Airline Industry** by Peter Belobaba, Amedeo Odoni and Cynthia Barnhart

General

AIRFRANCEKLM GROUP

WHY THIS IS HARD

SEVEN MINUTES AT ONE STAND
HOW A SMALL EVENT BECOMES A NETWORK PROBLEM

Every box in the planning chain has its own model and its own KPI — but the day of operations does not respect those boundaries. A disturbance **enters at the bottom and surfaces as cost at the top**.

07:12	**A baggage truck arrives 7 minutes late** One task, one stand — invisible on any dashboard	1 task
07:26	**The turnaround has no slack — late pushback** Tight turnarounds are a margin decision, not an accident	1 flight
09:40	**Connecting passengers at risk — hold the wave, or lose them** A hub carrier’s dilemma: protect connections and delay others	**4 flights**
13:05	**Rotation broken, crew hits its duty limit** An unplanned tail and crew swap — two domains, one decision	**9 flights**
next day	**Tomorrow starts degraded** Aircraft and crews end up in the wrong places	**the network**

WHY IT AMPLIFIES

No slack. Capital cost and margins push turnarounds to the minimum.

Shared resources. Teams, equipment and gates are scarce and contested.

Synchronization. Tasks must coincide; one absence blocks a whole chain.

Hub waves. Connections couple flights that share nothing else.

THE POINT The event is local; the cost is not. **No single silo can see it, and no deterministic plan can absorb it.**

AIRFRANCEKLM GROUP

WHY THIS IS HARD

WHAT DELAYS A DEPARTURE — AND WHEN

AVERAGE DELAY PER DEPARTING FLIGHT, IN MINUTES

Knock-on (reactionary) delay **snowballs through the day**: from ~5 minutes per departing flight at dawn to ~28 by 22:00 — delays inherited from earlier flights dominate every other cause.

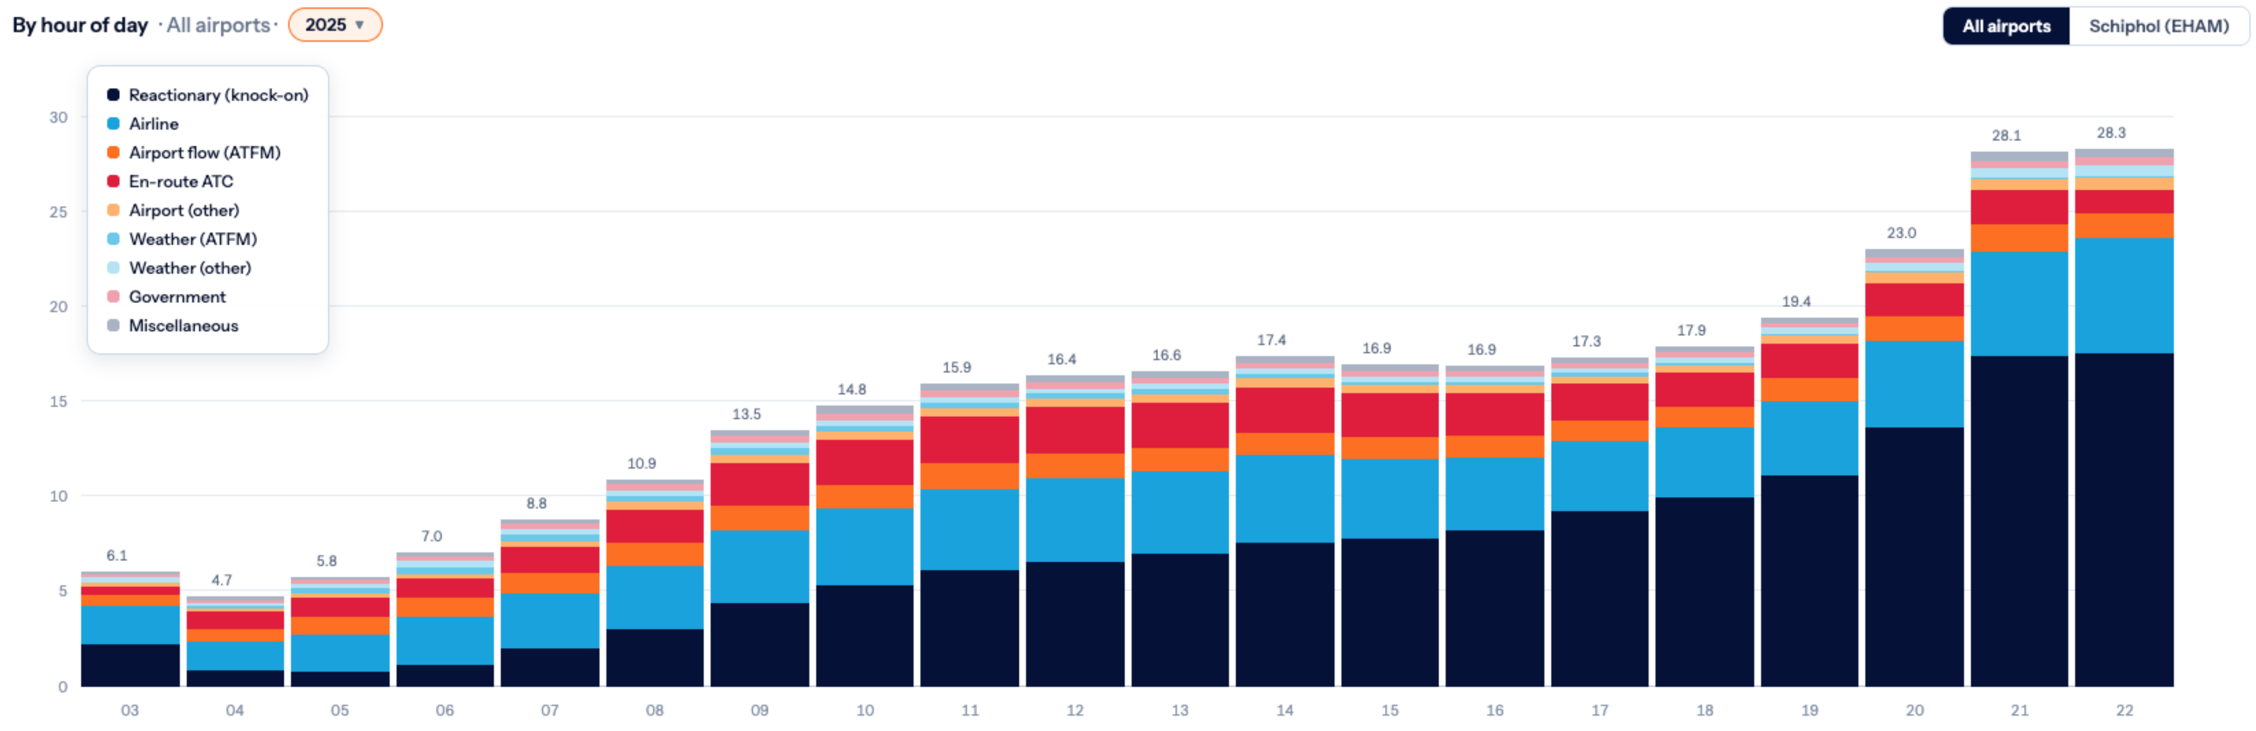

AIRFRANCEKLM GROUP

WHY THIS IS HARD

WHAT DELAYS A DEPARTURE — AND WHEN

AVERAGE DELAY PER DEPARTING FLIGHT, IN MINUTES

At our hub the evening is extreme: **106 minutes of average delay at 22:00 — almost 4× the European average** , nearly all of it knock-on. Transfer passengers must not strand overnight, so the last wave absorbs the whole day.

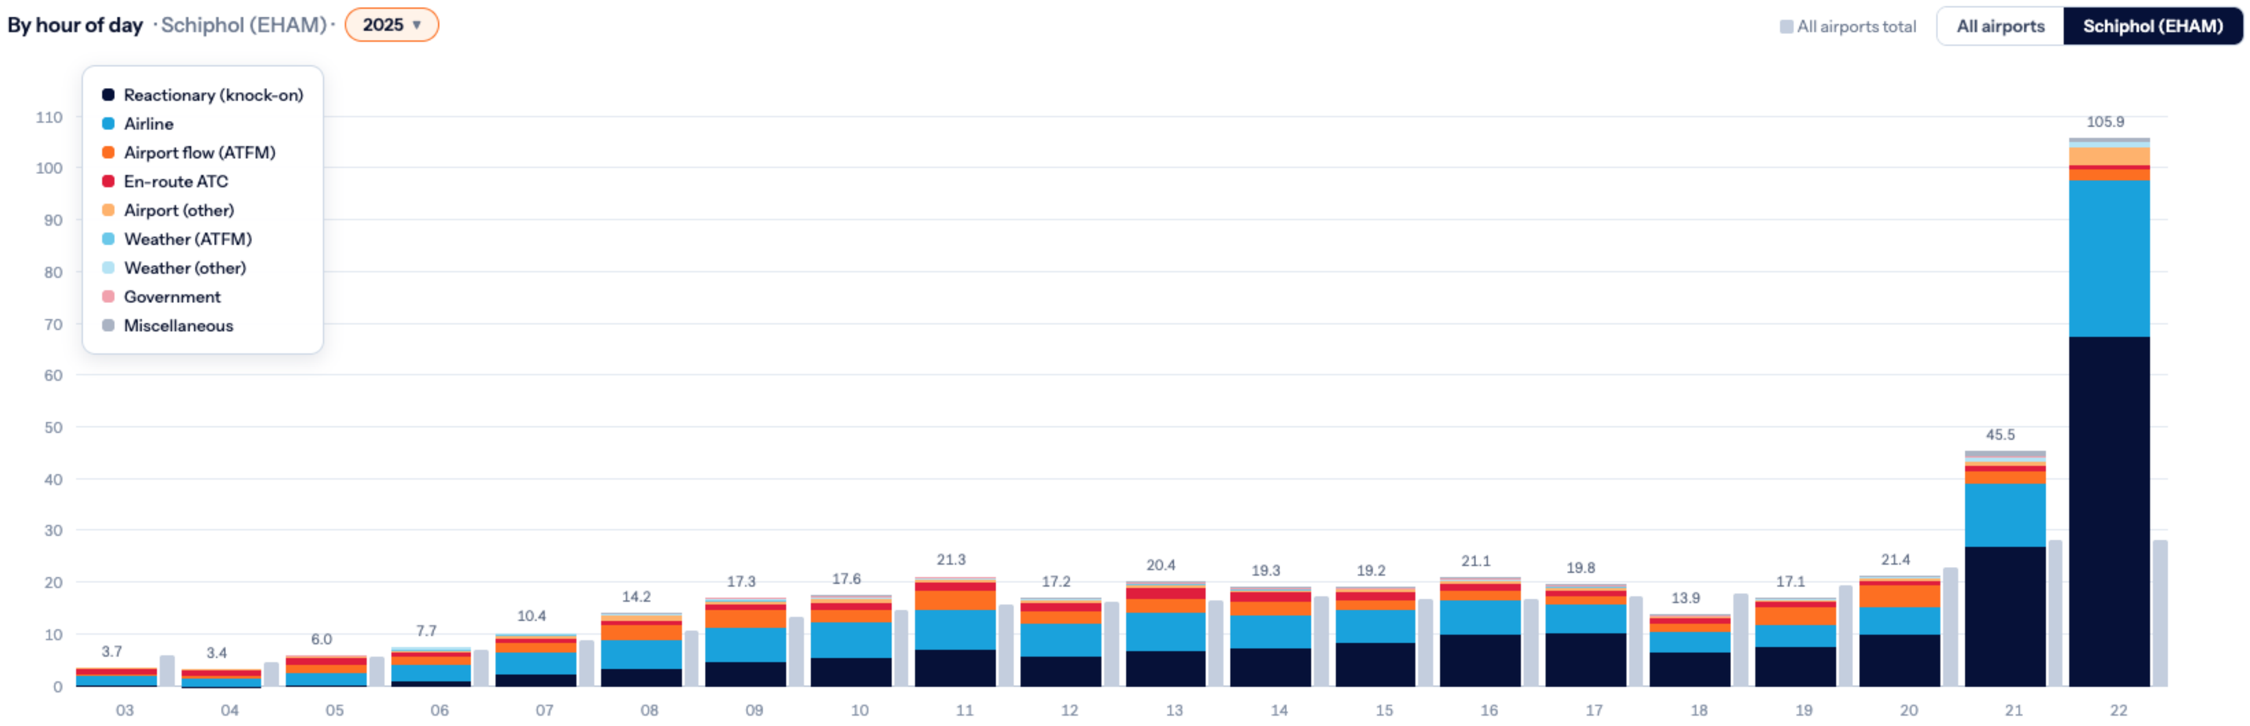

AIRFRANCEKLM GROUP

WHY THIS IS HARD

DEPARTURES: EARLY, ON TIME, OR LATE?

SHARE OF DEPARTURES (%) IN EACH PUNCTUALITY BAND

About **a third of departures leave on time (±5 min)**; since 2021 the late bands, 5-15 and 16-30 minutes, have clearly grown.

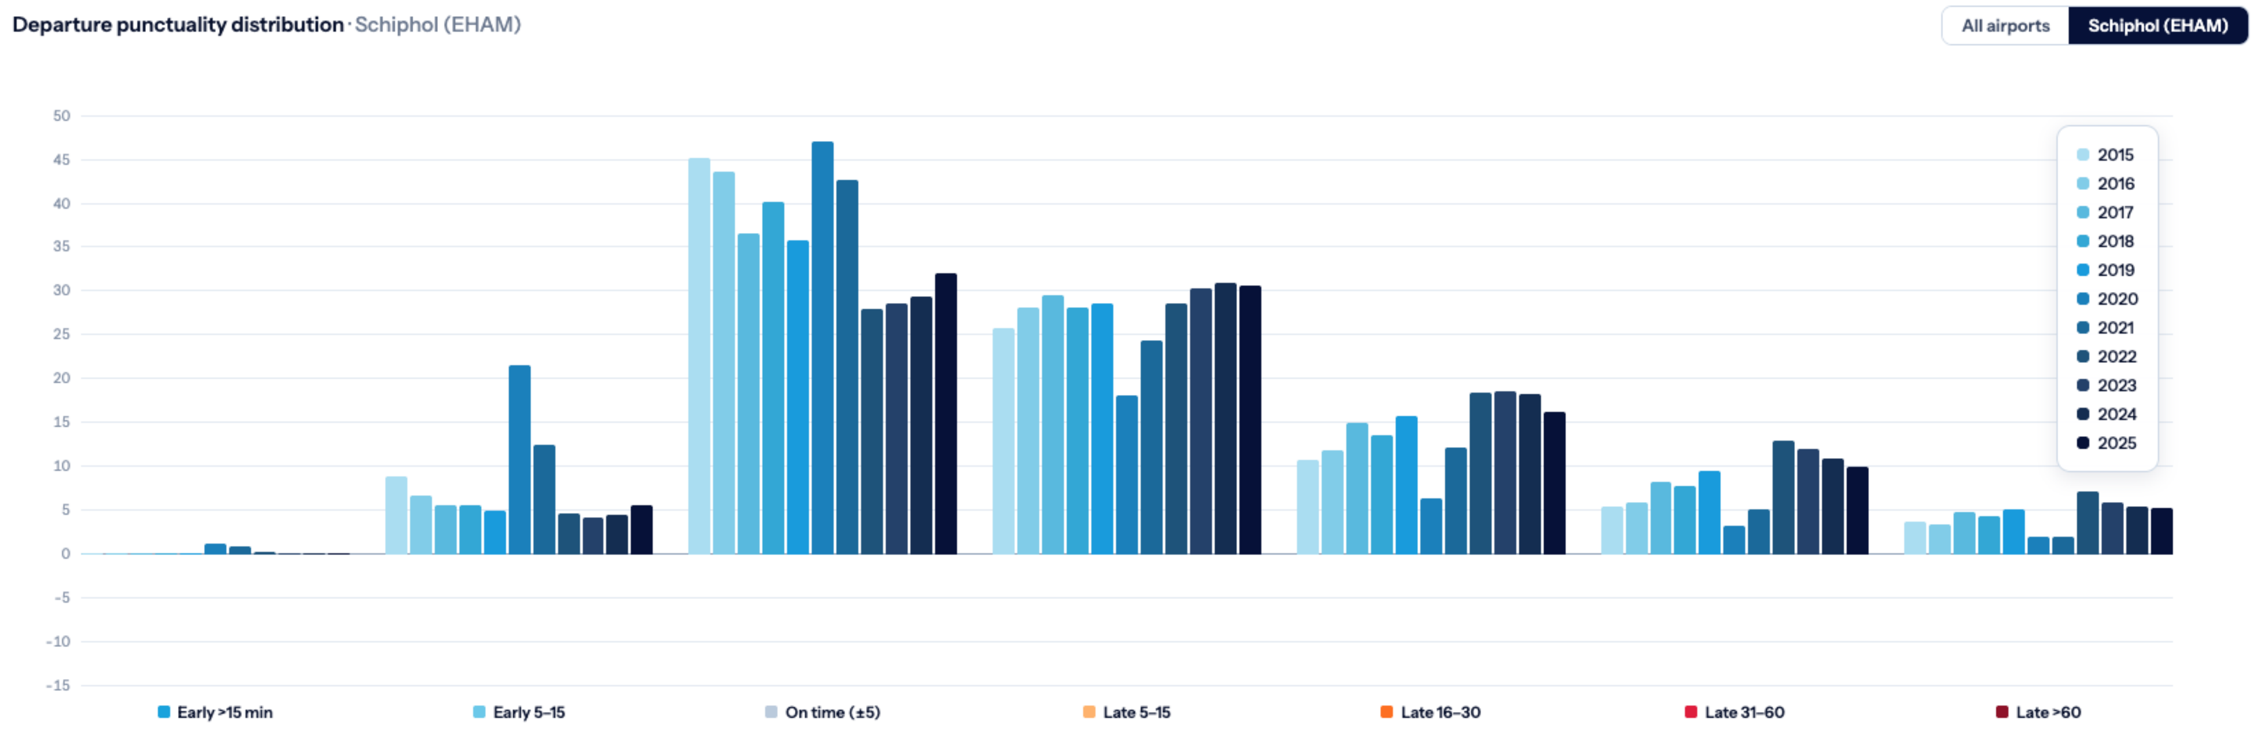

General

AIRFRANCEKLM GROUP

WHY THIS IS HARD

DEPARTURES: EARLY, ON TIME, OR LATE?

SHARE OF DEPARTURES (%) IN EACH PUNCTUALITY BAND

Against all airports we **depart early less often** and sit clearly above average in **Late 5–15 and 16–30,** consistent with the delay minutes piling up at the end of the day.

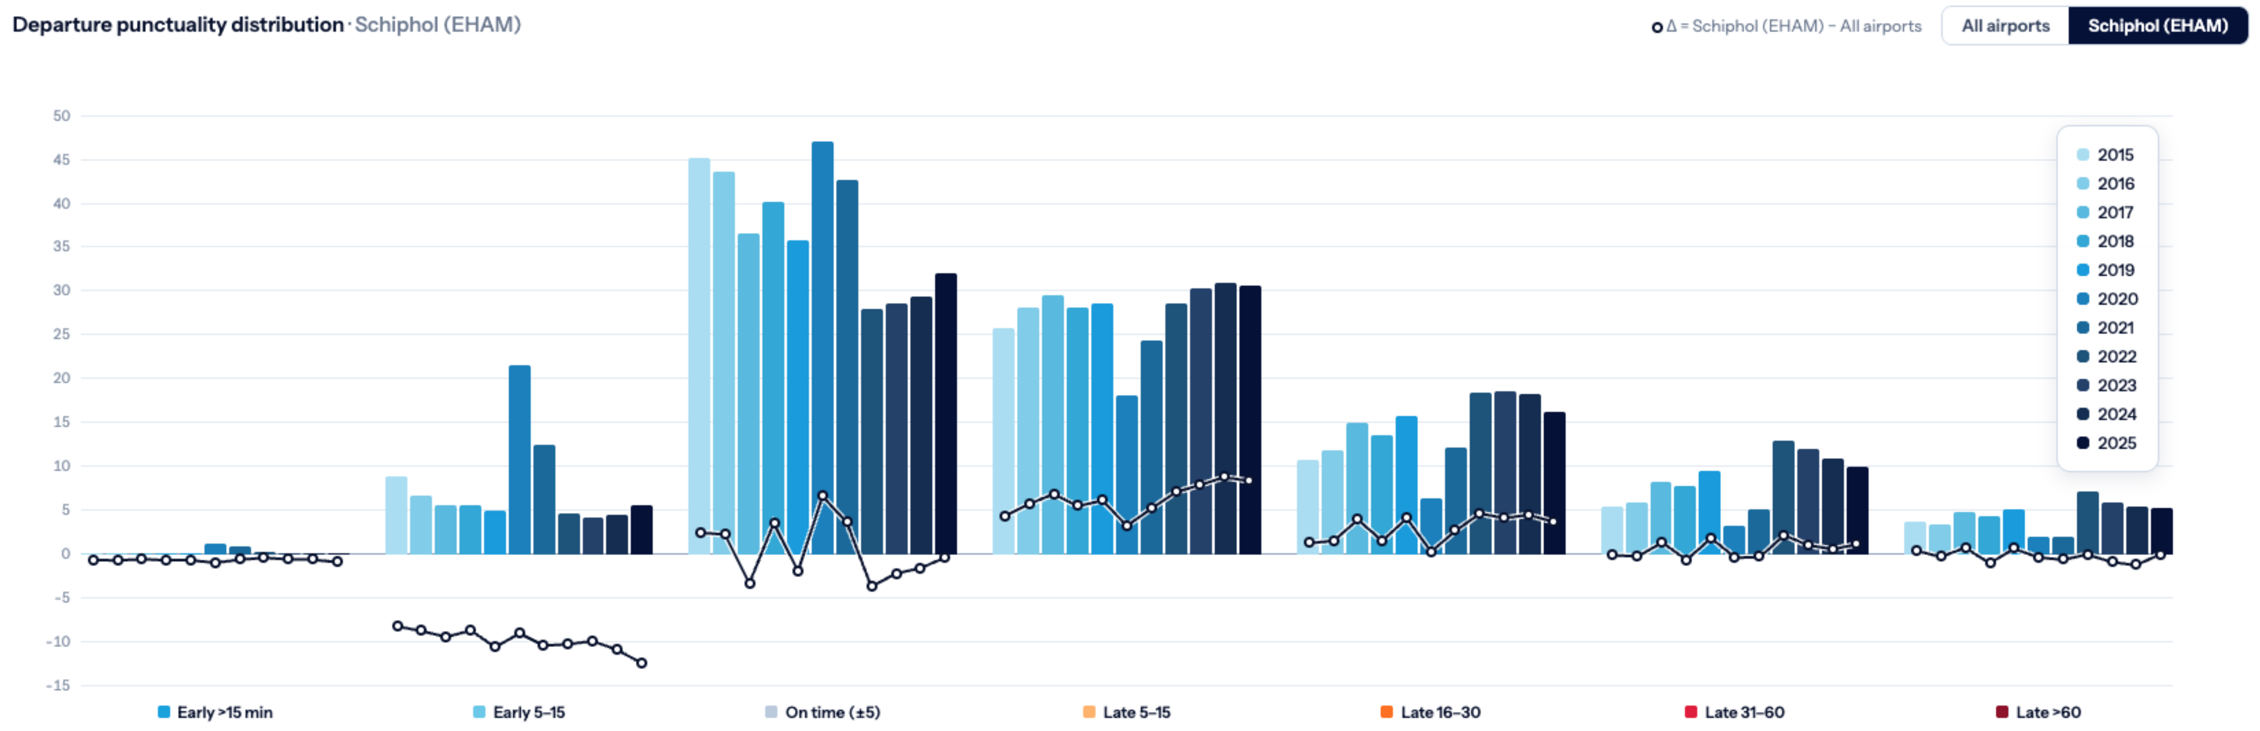

AIRFRANCEKLM GROUP

WHY THIS IS HARD

ARRIVALS: EARLY, ON TIME, OR LATE?

SHARE OF ARRIVALS (%) IN EACH PUNCTUALITY BAND

On the arrival side, **early and on-time arrivals look like the European picture**: roughly half of all flights arrive ahead of schedule.

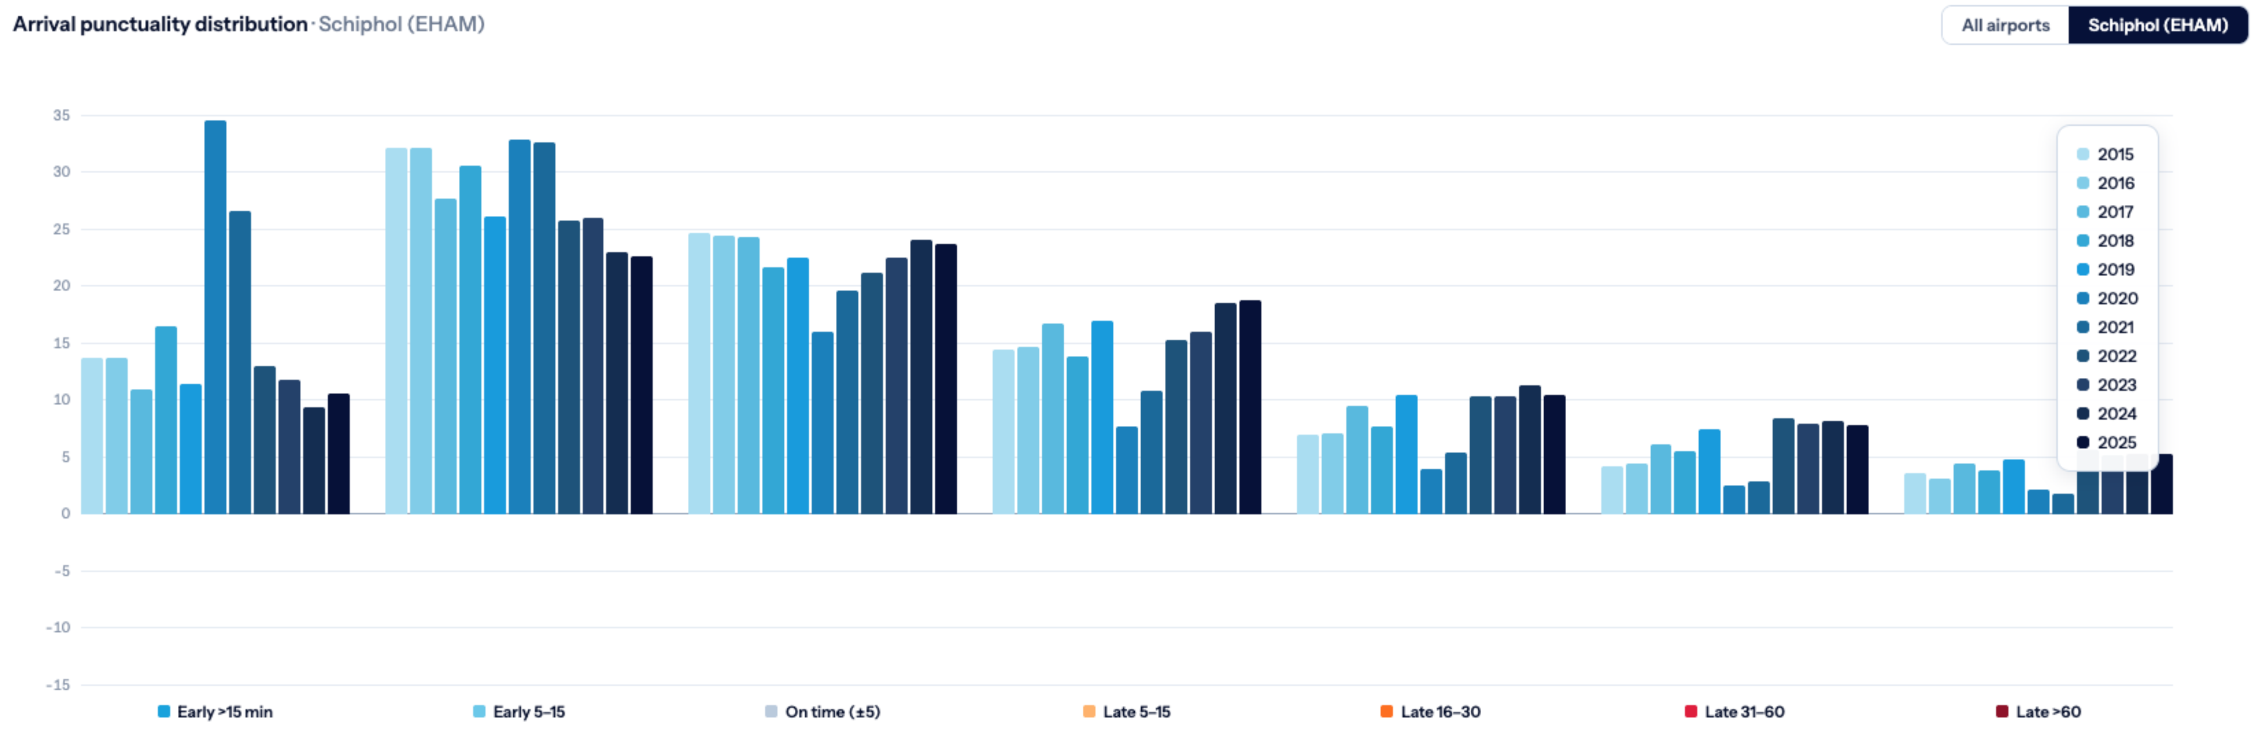

General

AIRFRANCEKLM GROUP

WHY THIS IS HARD

ARRIVALS: EARLY, ON TIME, OR LATE?

SHARE OF ARRIVALS (%) IN EACH PUNCTUALITY BAND

> The surprise: despite departing later, we **arrive late less often than the average airport** — we make it up in the air. That points to deliberate buffer in our flight times.

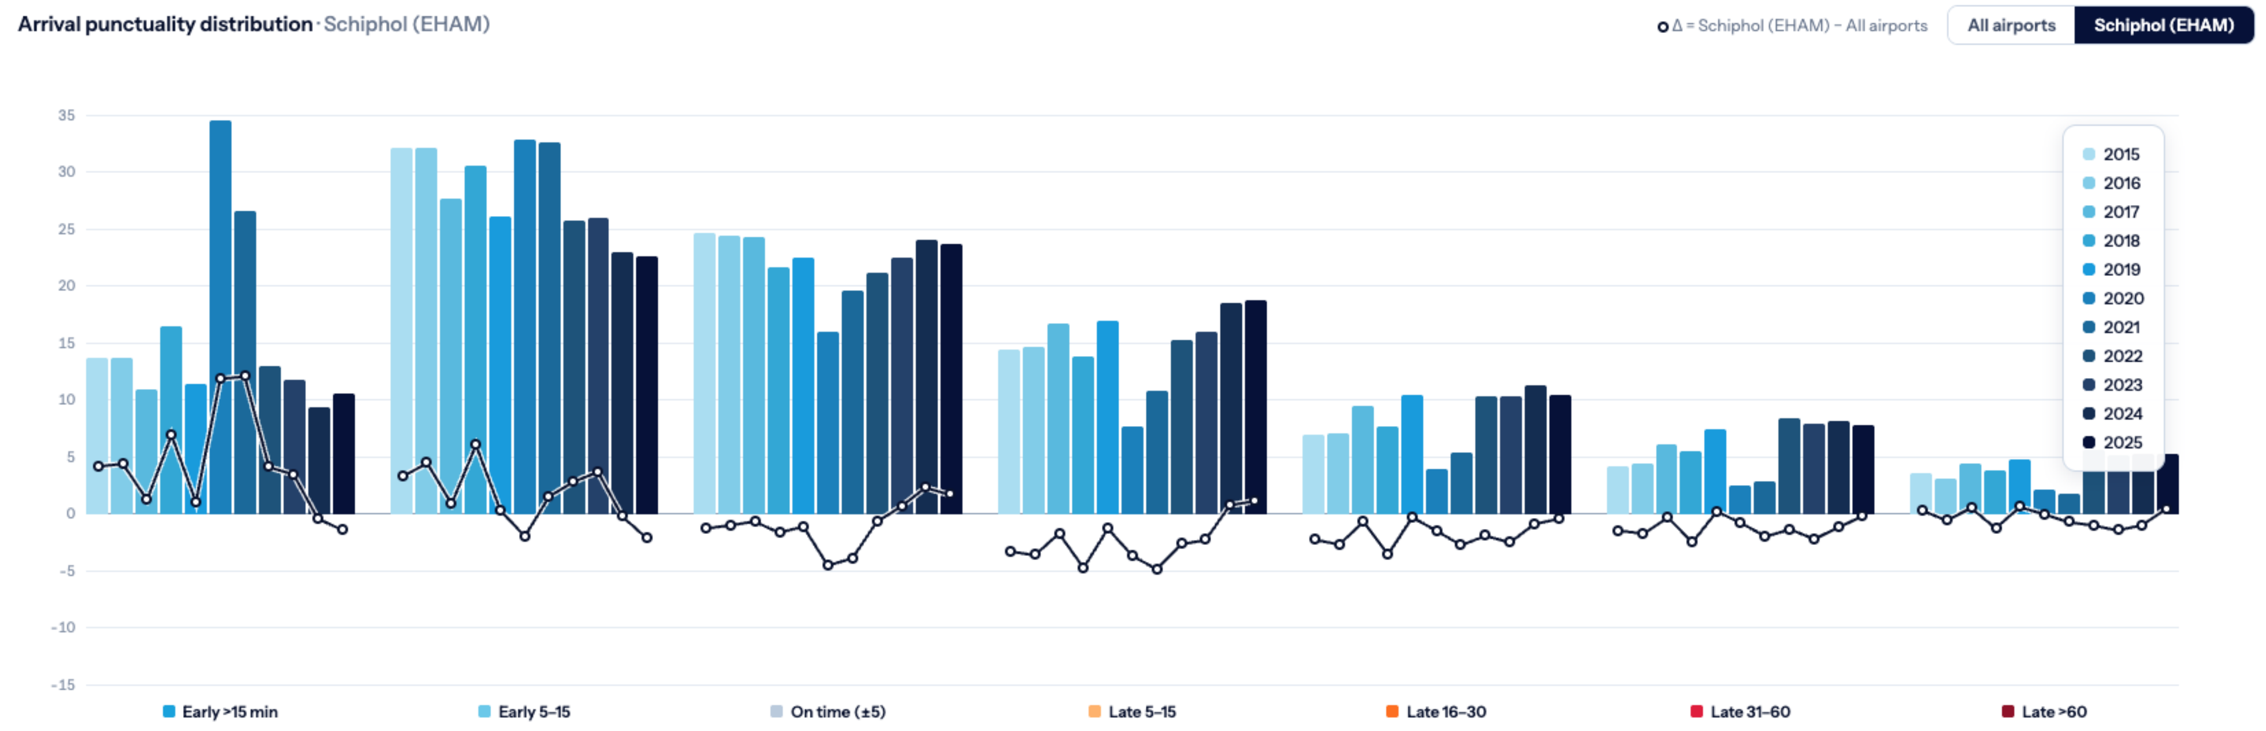

AIRFRANCEKLM GROUP

WHY THIS IS HARD

BLOCK TIME COVERAGE & DELAY RECOVERY

TWO BUFFER MEASURES, THE HIGHER THE MORE BUFFER

Around **80% of flights complete within their scheduled block time** , gaining 5–8 minutes on average, buffer built into the schedule, year after year.

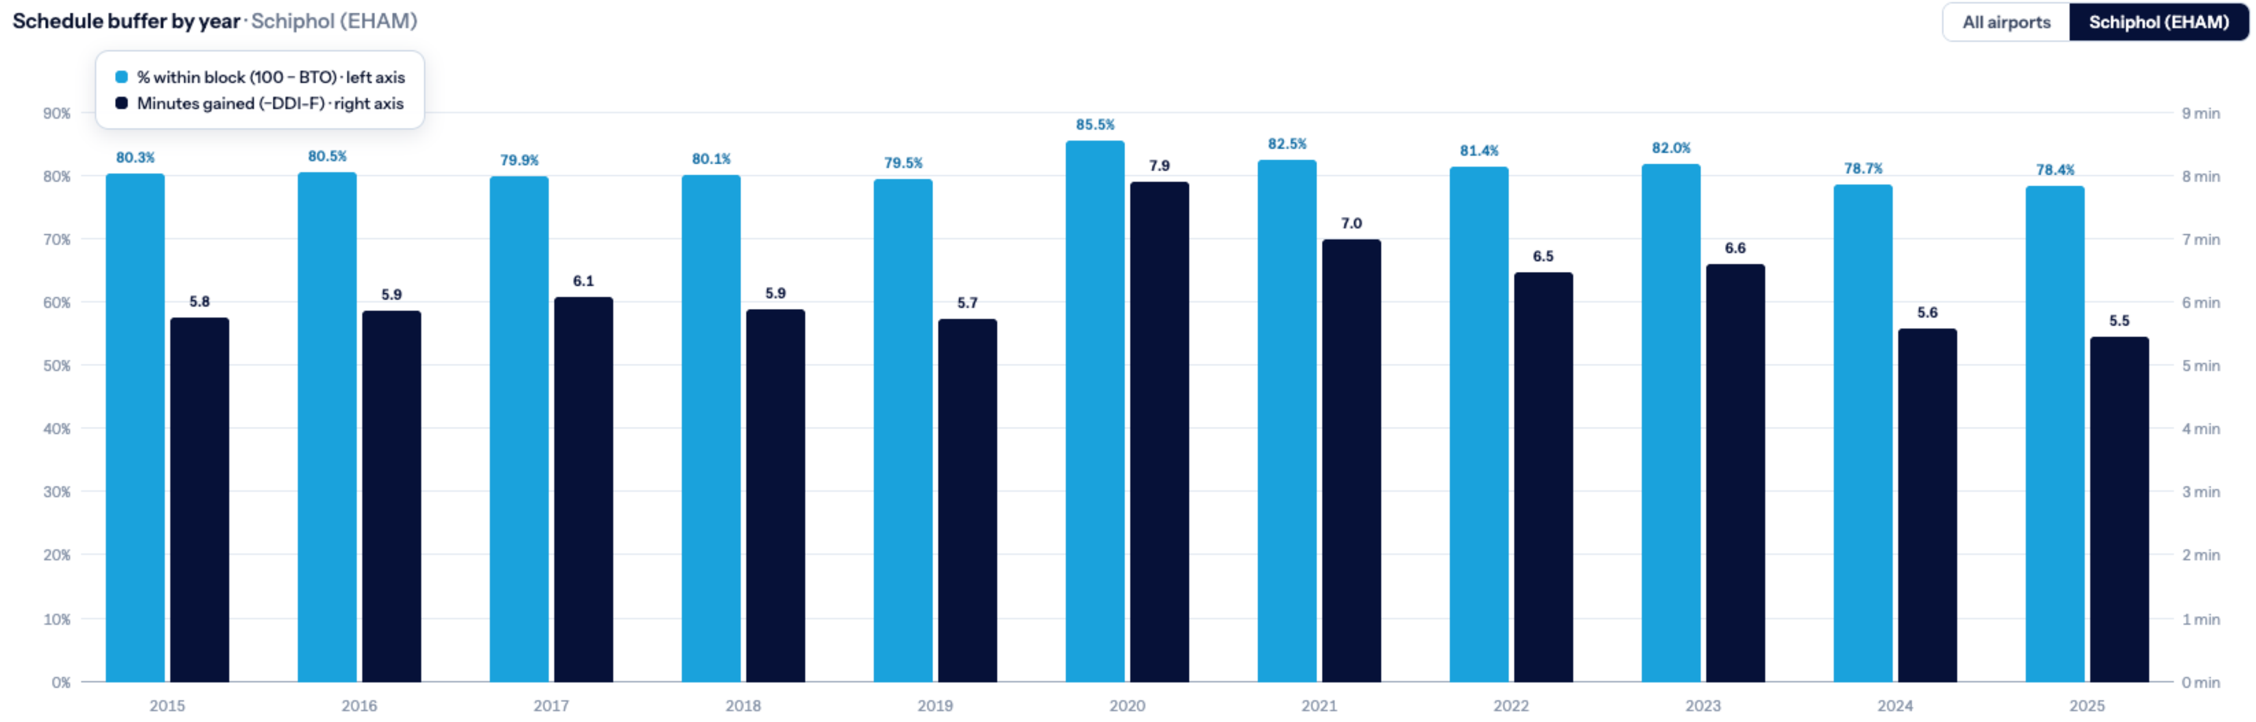

AIRFRANCEKLM GROUP

WHY THIS IS HARD

BLOCK TIME COVERAGE & DELAY RECOVERY

TWO BUFFER MEASURES, THE HIGHER THE MORE BUFFER

And **more of it than elsewhere** : Schiphol sits consistently above the all-airport average on both buffer measures, we buy robustness with schedule slack.

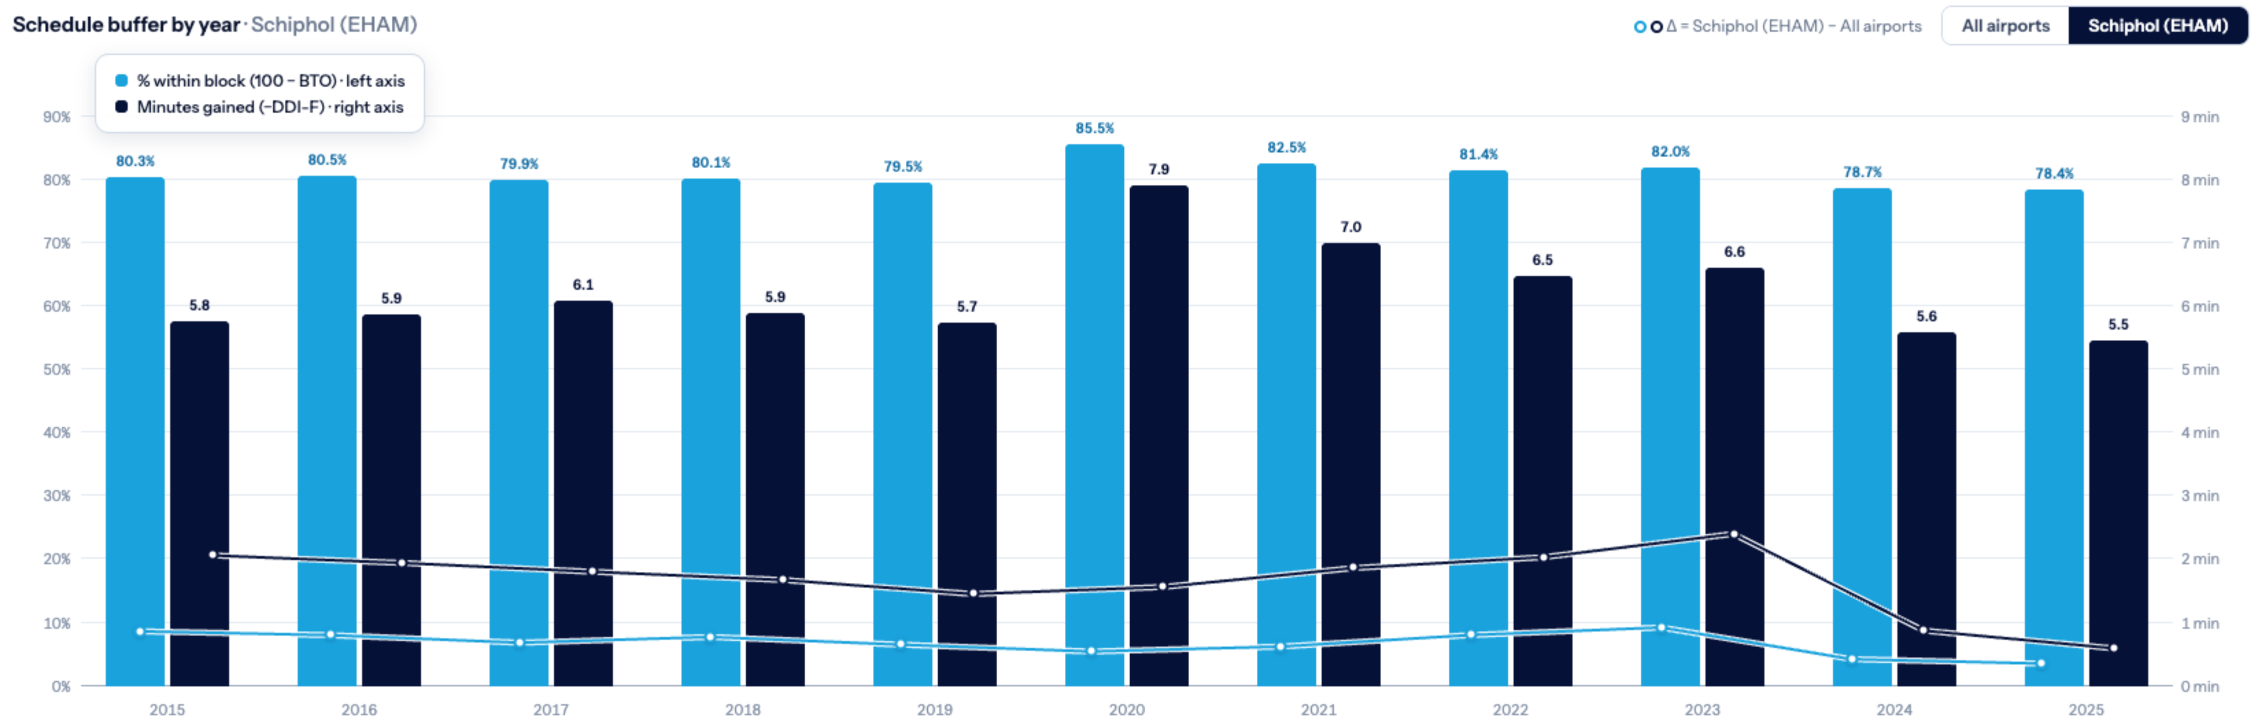

AIRFRANCEKLM GROUP

WHY THIS IS HARD

SLACK IS THE PRICE OF A HUB
WHAT THE DATA TOLD US

FINDING 01
Delay snowballs

Knock-on delay builds all day; at our hub the evening explodes, 106 min per departure at 22:00, ~4× the European average.

FINDING 02
We push back late

Fewer early departures than all airports, more in Late 5–15 and 16–30, the price of holding the evening bank together.

FINDING 03
…yet we land on time

Late arrivals are rarer than average despite the late departures, we make it up in the air.

FINDING 04
Because we buy buffer

~80% of flights beat their scheduled block time, gaining 5–8 minutes, consistently more slack than other airports.

THE WAY OUT — THREE LEVERS, ONE PER USE CASE AHEAD

Stochasticity
Plan against distributions, not averages, schedules that hedge instead of hope.

Integrality
Optimize across silos: fleet, crew, ground, revenue; one decision, not four.

Fast custom algorithms
Tailor-made methods quick enough to re-decide the moment reality deviates.

THE BRIDGE
Buffers buy robustness with idle aircraft, idle crew and unsold capacity. **The use cases that follow buy it with algorithms instead.**

AIRFRANCEKLM GROUP

TOWARDS RESILIENT AIRLINE PLANNING

BRIDGING THE SILOS
RESEARCH TOGETHER WITH TWENTE & UTRECHT

KLM partially funds the research positions of **2 professors and 4 PhD students**, supervised by **3 Data Science & Analytics leads**

UNIVERSITY OF TWENTE.

Dennis Prak
Assistant Professor

Mahekha Dahanayaka
PhD candidate

Antonio Montaruli
PhD candidate

Utrecht University

Marjan van den Akker
Associate Professor ·
President of the Dutch OR Society

Philip de Bruin
PhD candidate

Lisanne Heuseveldt
PhD candidate

KLM SUPERVISORS

Matthijs Kieskamp
Data Science & Analytics Leader

Jelle Dusseldorp
Data Science & Analytics Leader

Rohit Gupta
Data Science & Analytics Leader

WHY IT WORKS Jointly appointed academics give KLM **direct access to the latest science and emerging technologies** staying at the forefront of AI while solving complex operational challenges.

AIRFRANCEKLM GROUP

LITERATURE REVIEW

FROM GATE TO RUNWAY
A SYSTEMATIC REVIEW OF AIRPORT GROUND OPERATIONS

199 peer-reviewed studies (2004–2024) on everything between arrival at the gate and departure from the runway, classified into seven functional areas.

FUNCTIONAL AREA	SHARE OF REVIEWED PAPERS
Airport gate assignment	34%
Boarding processes & strategies	17%
Turnaround process	15%
Aircraft taxiing & ground movements	13%
Ground support equipment handling	11%
Baggage handling	6%
Ground workforce planning & scheduling	4%

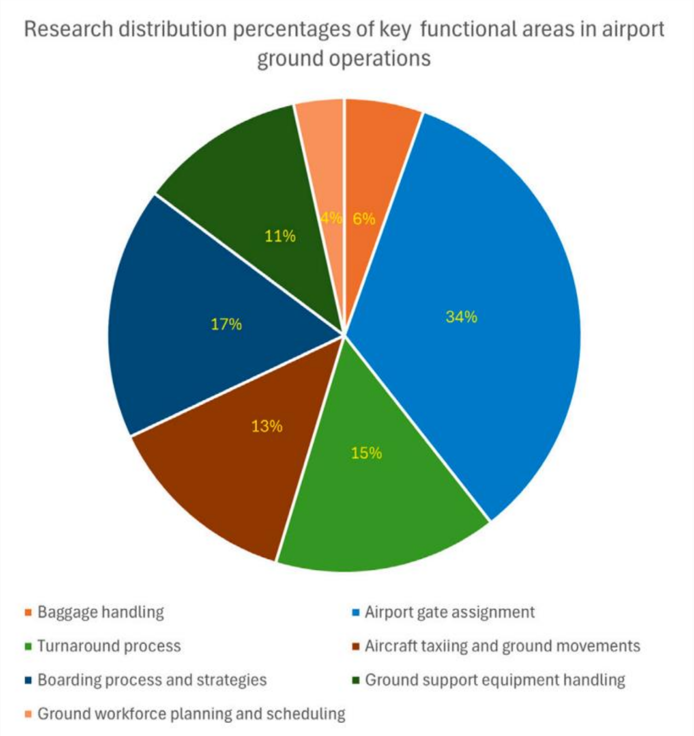

Dahanayaka, Prak & Mes (2026), "From gate to runway: a systematic review of airport ground operations optimization", Journal of Air Transport Management 135, 103013

AIRFRANCEKLM GROUP

LITERATURE REVIEW

STRONG PROGRESS ON SUB-PROBLEMS
BUT REAL OPERATIONS LOOK DIFFERENT

Ground processes are tightly coupled: a late inbound or one delayed task **cascades through gates, turnaround, crews and equipment into network-wide delay** . High capital costs and low margins force tight turnarounds — yet most published models assume the plan will execute perfectly.

GAP 01

A deterministic world

Deterministic service times, static schedules, limited disruption handling — assumptions that inflate performance on paper and break under real variability in arrivals, tasks, staffing and equipment.

GAP 02

Silos optimized in isolation

Gates, boarding, turnaround, baggage, taxiing, GSE and workforce are each optimized locally — while it is exactly the interdependencies between them that drive delay propagation.

GAP 03

Too slow for the day of operations

Exact methods on large combinatorial problems do not answer in operational time — disruptions demand decisions in minutes, not hours.

THE FINDING

Individual sub-problems are well studied — what is missing are **integrated, uncertainty-aware solutions** that work in practice.

AIRFRANCEKLM GROUP

THE FRAME FOR THIS TALK

TOWARDS RESILIENT AIRLINE PLANNING
THREE INSIGHTS THAT SHAPE WHAT FOLLOWS

01
Plan with uncertainty, not around it

Include **stochasticity** explicitly — stochastic and robust optimization that anticipates variability and disruptions, instead of deterministic plans that break on contact with the day of operations.

STOCHASTIC & ROBUST OPTIMIZATION · ROLLING HORIZON · REAL-TIME ADAPTATION

02
Optimize integrally, not per silo

Integral planning and optimization that captures the (inter-)dependencies across processes and departments — joint models consistently outperform sequential, siloed planning.

GATES + TURNAROUND + GSE + WORKFORCE · END-TO-END COORDINATION

03
Custom algorithms for near real-time

Handmade, **customized algorithms** — like simulated annealing and other tailored (meta-)heuristics — that cut computation from hours to near real-time, so optimization can support decisions as they happen.

METAHEURISTICS · DECOMPOSITION · HYBRID OPTIMIZATION + SIMULATION

The remainder of this talk shows these insights at work in a series of use cases across the planning chain — from airline network planning, to staff scheduling and gate agent planning, to disruption management to recover from disruptions.

AIRFRANCEKLM GROUP

AIRFRANCEKLM GROUP

MULTI-HUB AIRLINE NETWORK PLANNING

Antonio Montaruli, Matthijs Kieskamp, Sebastian Birolini, Dennis Prak, Martijn Mes

AIRFRANCEKLM GROUP

THE PROBLEM

MULTI-HUB AIRLINE NETWORK PLANNING
GROWTH AT ONE HUB HAS RUN OUT OF ROOM

Amsterdam Schiphol is bounded by slots, noise and runway capacity. A second hub is the obvious lever — but **no-one can currently say whether, and when, operating more than one hub actually creates value.**

TODAY - SINGLE HUB

THE QUESTION - MULTI-HUB

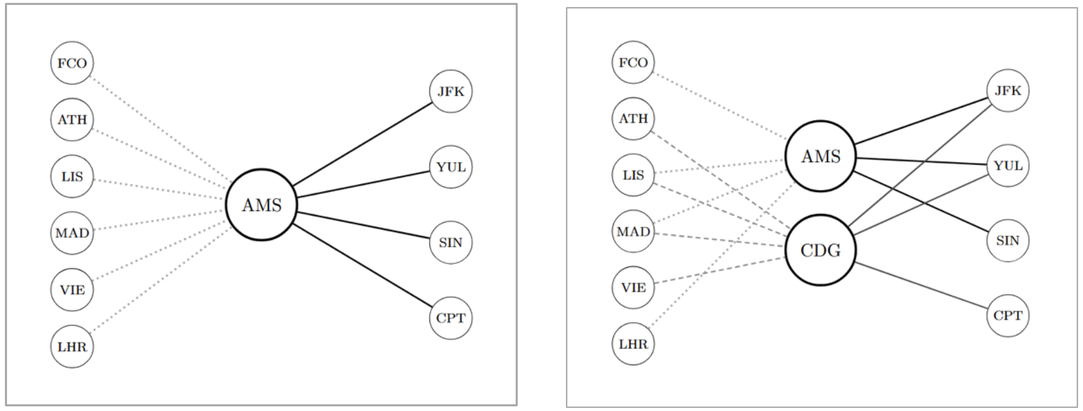

WHY IT IS HARD TODAY

- Routes, frequencies and fleet are decided **sequentially**, never jointly
- Competitors and passenger choice are **left out** of strategic models
- Upside and risk of a second hub are **never quantified in the same currency**

POTENTIAL UPSIDE

Broader coverage, access to thin markets

Relief of slot and capacity constraints

Better aircraft utilization

REAL RISK

Demand cannibalization between hubs

Loss of economies of density

Higher complexity and cost

UNIVERSITY OF TWENTE.

AIRFRANCEKLM GROUP

A STRATEGIC OPTIMIZATION FRAMEWORK
ONE MODEL INSTEAD OF A SEQUENTIAL CASCADE

max Revenue – Operating cost – Aircraft ownership cost mixed-integer nonlinear program

Four decisions taken jointly, in one competitive market model:

Route selection

Which airport pairs and hub itineraries to operate

Flight frequency

How often each itinerary is served

Fleet composition

How many aircraft of each type, and where

Demand allocation

Share of each market won against competitors

Tested at real scale: AMS and CDG, 152 candidate destinations, ≈30k markets and ≈60k itineraries, ≈200 competing airlines, 10 aircraft types.

UNIVERSITY OF TWENTE.

AIRFRANCEKLM GROUP

MULTI-HUB AIRLINE NETWORK PLANNING · RESULTS

AMSTERDAM SCHIPHOL & PARIS CHARLES DE GAULLE
WHAT A SECOND HUB ACTUALLY DELIVERS

+37.0%
Average revenue
versus the single-hub network

67→82%
Transfer passengers
share of all carried passengers

+66
Destinations added
served from one hub only — 86 from both

+19 / +30%
Aircraft required
long-haul / short-haul — the price of the upside

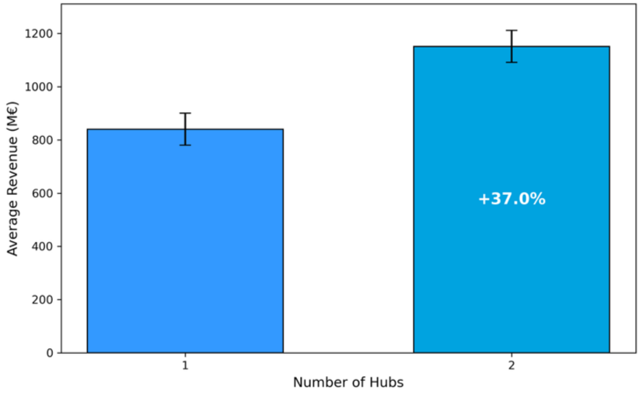

Average revenue (M€) by number of hubs

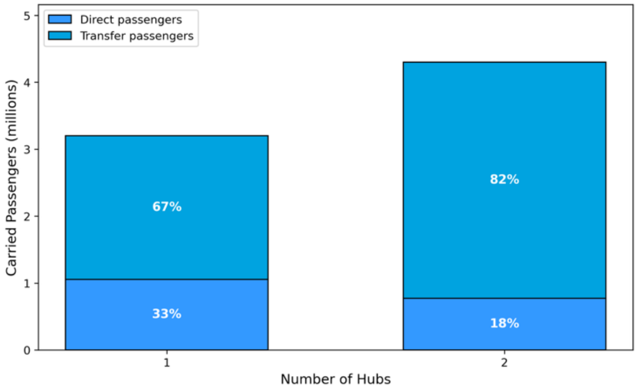

Carried passengers: direct vs. transfer

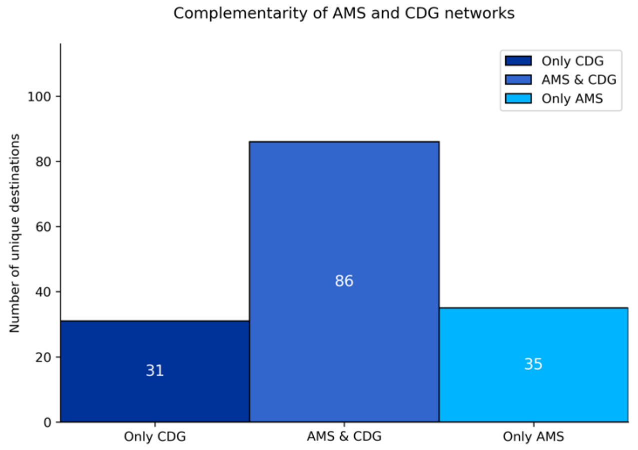

Destinations: only CDG · both · only AMS

General

UNIVERSITY OF TWENTE.

AIRFRANCEKLM GROUP

MULTI-HUB AIRLINE NETWORK PLANNING · RESULTS · ROBUSTNESS

DOES THE GAIN HOLD UP?
UNDER COMPETITION — AND UNDER SCRUTINY

Multi-hub remains the most profitable configuration at **every level of competitive pressure tested (0–30%)** , with higher load factors than either single-hub network. The advantage is real — but it is conditional.

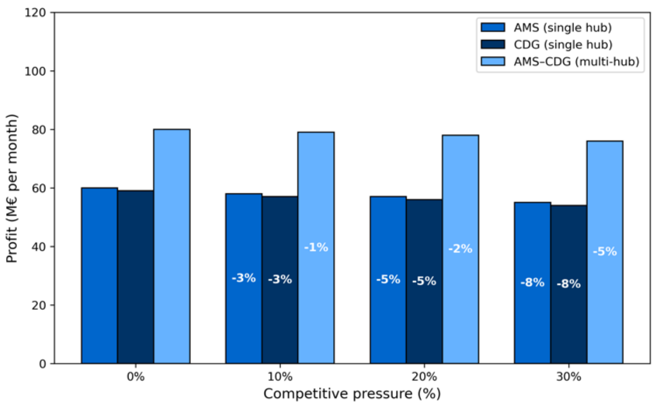

Profit (M€/month) vs. competitive pressure

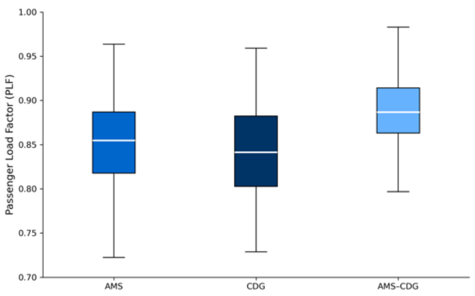

Passenger load factor distribution

ONLY IF THE DESIGN IS RIGHT

The hubs must complement, not duplicate

Overlapping catchments cannibalize demand — the second hub then buys traffic from the first.

Added complexity must earn its cost

Thinner flows cost economies of density; utilization and fleet mix are what pay for them.

Market share is endogenous

A case that looks profitable in isolation must survive competitors’ frequencies.

UNIVERSITY OF TWENTE.

AIRFRANCEKLM GROUP

WHAT MADE IT WORK

MULTI-HUB AIRLINE NETWORK PLANNING
THE THREE INSIGHTS, AT WORK

INSIGHT 01 · STOCHASTICITY

Test the strategy, not the forecast

A hub commitment outlives any demand forecast, so the network is judged across a **sweep of competitive pressure (0–30%)** rather than one point estimate — the multi-hub answer has to stay ahead in every scenario, not just the expected one.

INSIGHT 02 · INTEGRALITY

Hub choice only makes sense jointly

Routes, frequencies, fleet assignment and demand allocation are solved in **one program** instead of a sequential cascade — with competition against ≈200 carriers and aircraft ownership cost inside the same objective, so cannibalization and utilization are visible where the hub decision is made.

INSIGHT 03 · CUSTOM ALGORITHMS

Scope chosen to keep it solvable

Frequency-driven market share is nonlinear, and the network spans **≈60k itineraries** . Restricting to two-leg itineraries and leaving out rotations and scheduling is the deliberate simplification that makes a strategic question answerable at all.

NEXT STEPS

- Couple with flight scheduling and aircraft rotations
- Disruption scenarios — test hub strategies on resilience, not only profit
- Learning-based demand and pricing instead of fixed fares

TAKEAWAY A structured, data-driven tool to assess multi-hub strategies quantitatively — at the interface of **commercial planning and operations** .

UNIVERSITY OF TWENTE.

AIRFRANCEKLM GROUP

DRL FOR AIRLINE REVENUE MANAGEMENT

Antonio Montaruli, Florentina Hager, Dennis Prak, Matthijs Kieskamp, Willem van Jaarsveld, Martijn Mes

AIRFRANCEKLM GROUP

THE PROBLEM

DEEP REINFORCEMENT LEARNING

PRICING A NETWORK OF SEATS THAT EXPIRE AT DEPARTURE

> Offer the right product, to the right customer, at the right time, at the right price — while **demand arrives stochastically, willingness to pay is unobserved, and every leg's price affects connecting itineraries on other legs.**

THE NETWORK EFFECT

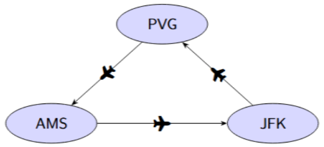

Possible Flight Legs	
Origin Airport	Destination Airport
PVG	AMS
AMS	JFK
JFK	PVG

Possible Combinations	
Origin Airport	Destination Airport
PVG	AMS
PVG	JFK
AMS	PVG
AMS	JFK
JFK	AMS
JFK	PVG

Price one leg too high and you lose a connecting passenger worth revenue on *two* legs; too low and you give away high-yield direct demand.

WHY TRADITIONAL RM SYSTEMS STRUGGLE

- **Curse of dimensionality** — the state space grows exponentially with network size
- **Reliance on external demand models** — historical data may not represent future conditions
- **Sensitive to forecast accuracy** — performance degrades when estimates are off
- **No exploration, no autonomous learning** — systems only adapt when manually recalibrated

UNIVERSITY OF TWENTE.

AIRFRANCEKLM GROUP

TAMING THE ACTION SPACE
PRICE ONE LEG AT A TIME, KEEP THE NETWORK COUPLING

JOINT PRICING — EXPONENTIAL

$$|P|^{|L|}$$

Choosing prices for all legs simultaneously explodes the action space: with 10 price points, a 6-leg network already means a million joint actions per decision — the curse of dimensionality that breaks classical dynamic programming.

LEG-WISE DECOMPOSITION — LINEAR

$$|L| \times |P|$$

The agent prices legs **one at a time** , cycling through them with intermediate state transitions. Shared seat capacities and itinerary consumption preserve every cross-leg effect — at linear cost.

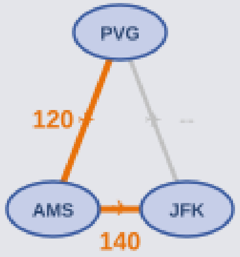

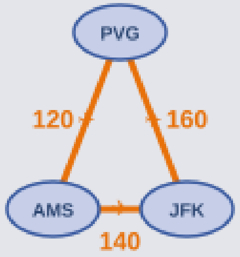

A **handmade, customized formulation** — not a bigger solver — is what makes network-scale DRL computationally feasible (insight 03).

UNIVERSITY OF TWENTE.

AIRFRANCEKLM GROUP

DRL FOR AIRLINE REVENUE MANAGEMENT · THE APPROACH

LEARN THE PRICING POLICY
INSTEAD OF FORECASTING DEMAND FOR IT

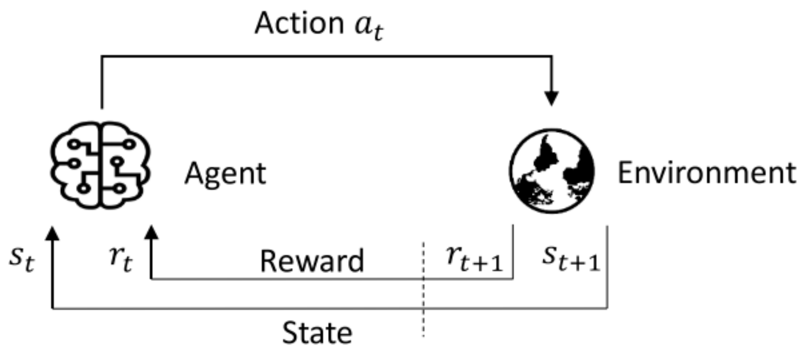

The pricing problem becomes a **finite-horizon Markov Decision Process**: the agent posts prices, a simulated booking environment answers with stochastic requests, and the policy improves from the realized revenue — **no predefined demand model**.

State	Remaining seats per class and leg + time to departure
Action	The posted price per leg, from a discrete fare ladder
Reward	Realized ticket revenue from accepted bookings
Transition	Shared capacities carry every cross-leg network effect

TWO DRL ALGORITHMS

DCL (Deep Controlled Learning) — built for highly stochastic settings: Monte Carlo rollouts per action, policy improvement as supervised classification.

PPO (Proximal Policy Optimization) — the standard on-policy actor-critic reference with clipped updates.

UNIVERSITY OF TWENTE.

BENCHMARKED ON NINE NETWORKS, 3 TO 50 LEGS
DRL WINS ON EVERY ONE OF THEM

Against bid-price control and a ladder of heuristic pricing rules, the learned policies deliver the highest revenue on **all nine network configurations — up to +11.5%** — and DCL does it with markedly lower variability across runs.

$\lvert\mathcal{L}\rvert$	$\lvert\mathcal{I}\rvert$	Best DRL	Revenue	Best benchmark	Revenue	Δ (%)
3	6	DCL	45 984	IRCR	44 325	+3.7
6	12	DCL	89 700	IRDT	80 945	+10.8
10	30	DCL	151 900	IRCT	137 523	+10.5
15	50	DCL	233 810	IRCT	209 787	+11.5
20	60	DCL	310 552	FP	280 124	+10.9
25	80	DCL	369 269	FP	341 689	+8.1
30	100	DCL	455 950	FP	423 937	+7.6
40	120	PPO	536 643	IRCT	521 213	+3.0
50	150	PPO	678 334	IRDT	669 510	+1.3

Best DRL policy vs. best benchmark per network (revenue, Δ%)

up to +11.5%

revenue over the best benchmark, growing with network complexity (DCL up to 30 legs, PPO beyond)

σ 72 vs. 598

revenue variability across runs, DCL vs. the best heuristic — the stability that makes automation trustworthy

Seconds, not minutes

decisions come from a trained network at posting time — the MILP-based re-optimization benchmark is orders slower

DRL FOR AIRLINE REVENUE MANAGEMENT · WHAT MADE IT WORK

DEEP RL FOR AIRLINE REVENUE MANAGEMENT
THE THREE INSIGHTS, AT WORK

INSIGHT 01 · STOCHASTICITY

Learn from interaction, not from a forecast

Uncertainty in arrivals and willingness to pay is handled by **learning** — the policy absorbs stochasticity from millions of simulated bookings instead of depending on a brittle external demand model.

INSIGHT 02 · INTEGRALITY

Price the network, not each leg in isolation

Shared capacities couple every decision, so the agent balances a connecting passenger's two-leg revenue against high-yield direct demand — the interdependency siloed pricing misses.

INSIGHT 03 · CUSTOM ALGORITHMS

Tailored machinery buys tractability

Leg-wise decomposition, rollouts with common random numbers, sequential halving in DCL — handmade choices that turn an intractable problem into one solved in operational time.

NEXT STEPS

- Scale to real-world hub-and-spoke networks at reasonable computational cost
- Integrate overbooking and no-show dynamics
- Model competition and demand substitution across carriers
- Price each route directly, not as a sum of legs

TAKEAWAY Where traditional RM relies on external models of demand, the DRL framework **learns to balance immediate revenue against future capacity directly from interaction** — a scalable, adaptive complement to today's systems.

INTEGRATED AIRLINE FLEET AND CREW RECOVERY

Philip de Bruin, Kunal Kumar, Marc Paelinck, Marjan van den Akker

WHEN THE PLAN BREAKS, MINUTES COUNT
REDECIDE FLEET, CREW AND PASSENGERS TOGETHER

Weather, maintenance, staff shortages — a disruption makes the planned tail and crew assignment **infeasible** , and delays cascade through an interconnected network. Controllers need a recovered schedule they can communicate **now**.

≈500
flights per day — KLM's European network, one hub, 6 aircraft types

72 h
recovery window — long enough to catch knock-on effects, then rejoin the plan

15–30 min
runtime of (basic) integrated models in the literature — too slow for the control room

30 s
the budget this work sets itself — good solutions fast beat optimal too late

WHY SEQUENTIAL RECOVERY FALLS SHORT

- Practice solves **tails first, then crew** — but decisions in one problem constrain the other
- A tail fix can **break crew legality** : a propagated delay becomes duty-time overtime
- Passenger costs like **missed connections** are decided implicitly, never traded off like costs due to delays and cancellations

The question: can **integrated** fleet + crew + passenger recovery run at control-room speed?

AIRFRANCEKLM GROUP

INTEGRATED AIRLINE FLEET AND CREW RECOVERY · THE KEY TRICK

SEARCH THE ISSUES, NOT THE WHOLE SCHEDULE

DIRECTED LOCAL SEARCH TARGETS WHAT IS BROKEN

TRICK 01

Every move fixes a live issue

Instead of random neighbors everywhere, pick one current **issue** — an overlap, a duty violation, a missed connection, a delay or cancellation — and generate a neighbor that addresses **it** . Effort lands where the schedule is infeasible or costly.

TRICK 02

Soften the hard constraints — temporarily

Dense schedules leave no slack, so **overlaps due to delays and duty/rest violations are allowed at a penalty** that rises over runtime. The search space stays connected early; the final solution is pushed to feasibility by the end.

TRICK 03

Disruption = a blocking task

An **aircraft or crew suddenly unavailable** for 6–12 h becomes an **unmovable task** on that resource. The resulting overlaps are ordinary issues the search removes by swap, delay or cancellation — one mechanism for every disruption type.

SIMULATED ANNEALING ON TOP

Worse moves are accepted with a probability that cools over time — the search escapes poor local minima early, then converges. And it can be **stopped at any moment** with the best solution so far: exactly what a control room needs.

Utrecht University

MOVES THAT RESPECT THE PHYSICS OF OPERATIONS

ROTATIONS, NOT FLIGHTS — SO NOTHING GETS STRANDED

SwapTail

exchange rotations between two aircraft of the same type

SwapCrew

exchange rotations between qualified cockpit crews

SwapBoth ★

keep tail + crew together — qualification travels with the pair, so swaps work **across aircraft types**

MoveTail

move a rotation from one aircraft to another

Delay · Cancel · Undo

delays can *save* connections; cancellation applies to full rotations; undo removes earlier delays or cancellations when the search finds better

CHOOSING THE NEXT MOVE

RandomNeighbour

one move per issue — fastest, near-instant answers

Candidates · greedy or SoftMin

generate many candidate fixes, pick the best or probability-weighted — slightly better cost, more compute

WHY ROTATIONS

Single-flight swaps strand aircraft and crew at wrong airports. Hub-to-hub rotations preserve location continuity — the schedule can rejoin the plan at the window’s end.

INTEGRATED BEATS SEQUENTIAL — AT 30 SECONDS

≈40% LOWER PASSENGER DISRUPTION COST

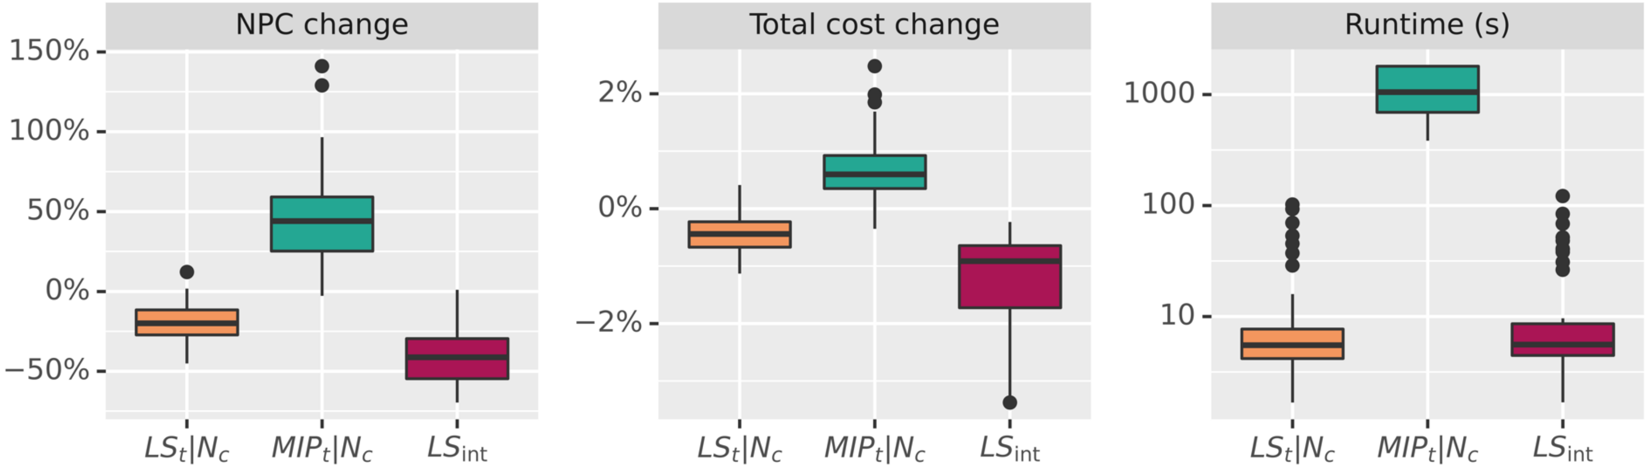

Change vs. naive delay propagation (N_{int}), KLM European network. Left to right: passenger non-performance cost, total cost, runtime — sequential local search ($LS_t|N_c$), sequential MIP ($MIP_t|N_c$, 30-min budget, reduced 36-h horizon), integrated local search (LS_{int}).

−40%

passenger non-performance cost vs. naive — the only approach that consistently improves it

≈30 s

to a high-quality integrated solution — vs. ~1,000 s for the sequential MIP

100%

of test instances solved — naive propagation manages only ~64%

2 families

of disruptions tested: widespread delays and 6–12 h resource unavailability

INTEGRATED RECOVERY
THE THREE INSIGHTS, AT WORK

INSIGHT 01 · STOCHASTICITY

Uncertainty, on the day it strikes

Recovery is the realized tail of uncertainty. An **anytime** algorithm — stoppable at 30 seconds with the best solution so far — answers it: good solutions fast beat optimal solutions too late.

INSIGHT 02 · INTEGRALITY

The strongest lever only exists jointly

SwapBoth — moving tail and crew as a pair across aircraft types — is **impossible in sequential recovery**, where a type change reverberates into crew legality.
Fleet + crew + passenger costs, one problem.

INSIGHT 03 · CUSTOM ALGORITHMS

Directed search fits the time budget

Issue-driven neighbors, adaptive penalties, rotation-level moves — a search built around the problem's physics does in 30 seconds what general-purpose MIPs need 30 minutes for.

NEXT STEPS

- More levers: deadheading crew as passengers
- From single-hub to multi-hub networks
- Possibly ground resources
- Robustness: recover today, dampen tomorrow's disruption

TAKEAWAY

Integrated recovery, long "too slow for practice", runs in 30 seconds — the shift from **resource-by-resource firefighting to network-wide, passenger-aware control** .

AIRLINE GROUND STAFF SCHEDULING

Mahekha Dahanayaka, Jeroen Assink, Rohit Gupta, Tomas Pippia, Dennis Prak

AIRFRANCEKLM GROUP

MANAGING UNCERTAINTY IN AIRLINE GROUND OPERATIONS

STAFF SCHEDULING UNDER UNCERTAINTY
FOR THE AVERAGE DAY - NOT THE DAY OF OPERATION

Luggage-handling shifts at Schiphol are fixed months ahead with deterministic models and planner experience. Then flights run late, tasks overrun, workload migrates through the day — and the plan misaligns, forcing **outsourcing, borrowed staff or absorbed flight delays.**

≈620

flights handled daily at Schiphol

≈15,000

employees in ground operations — a critical cost driver

<1%

of the scheduling literature addresses stochastic shift scheduling in aviation ground ops

Robustness

the missing third pillar, next to effectiveness and efficiency

WHY DELAYS BREAK THE ROSTER

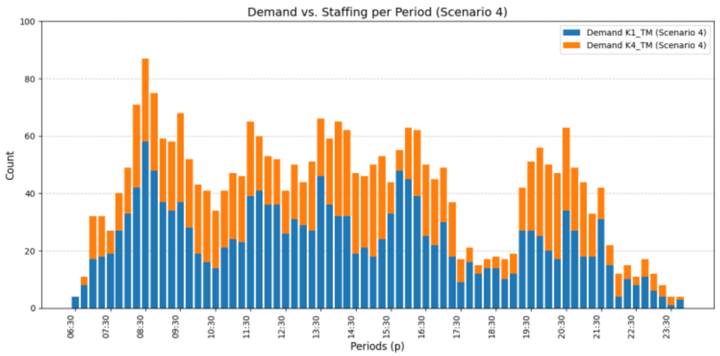

Planned staffing vs. realized workload per period: delays shift demand into hours the deterministic roster never covered.

UNIVERSITY OF TWENTE.

AIRFRANCEKLM GROUP

MANAGING UNCERTAINTY IN AIRLINE GROUND OPERATIONS · THE KEY TRICK

FROM TASKS TO PERIODS

COVER DEMAND PER TIME SLOT, NOT THOUSANDS OF TASKS

TASK ASSIGNMENT MODEL — TODAY

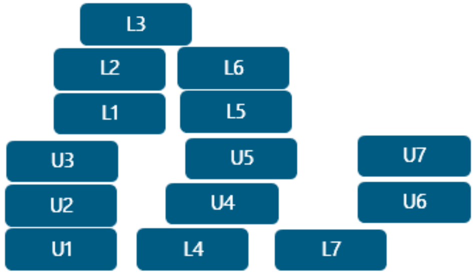

Every task modelled and assigned to a shift, with overlap constraints — a mixed-integer program with **millions of variables** that takes ~12 minutes per run. Too heavy to solve once per scenario.

DEMAND ACCUMULATION MODEL — PROPOSED

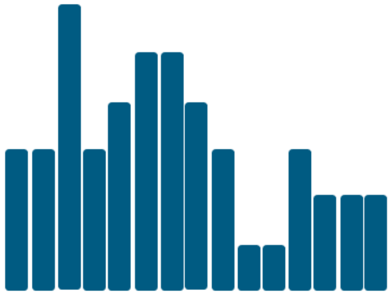

Aggregate workload into **time-based demand periods** and schedule staff to cover each period — the same shifts question, at a fraction of the size.

−99.7%
variables

−99.9%
constraints

12 min → <1 s
runtime, near-constant as task volume grows

99–100%
of original tasks still assignable under the period plan

General

UNIVERSITY OF TWENTE.

AIRFRANCEKLM GROUP

STOCHASTIC DEMAND ACCUMULATION
OPTIMIZE THE ROSTER OVER DISRUPTION SCENARIOS

1 **Fit delay distributions from real data**

Arrival and departure delays fitted separately (noncentral t-distribution) from KLM operations.

2 **Sample disruption scenarios**

Each scenario shifts task timings, reproducing how workload migrates through the day — **≈40 scenarios** suffice for stable estimates.

3 **Solve one two-stage stochastic program**

Shift types and staffing fixed *before* the day of operation; period-by-period staff exchange between pools acts as recourse per scenario. Sample Average Approximation, **optimal in ≈20 s**.

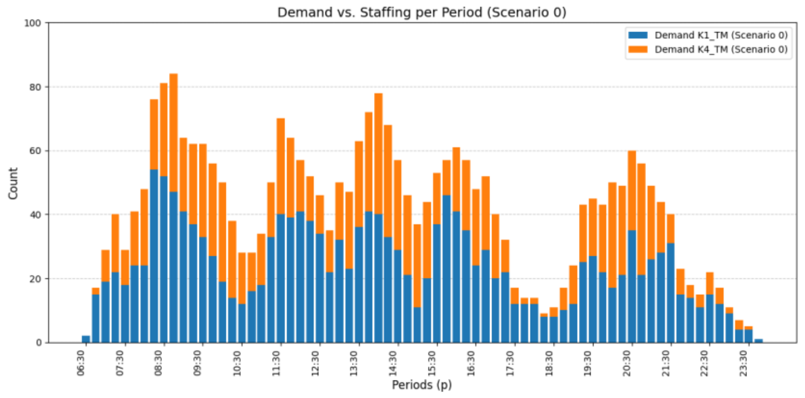

One sampled scenario: per-period demand vs. staffing — the roster must hold across all of them

WHY IT ONLY WORKS NOW

A 12-minute model × 40 scenarios was never an option. A sub-second model makes scenario-based robustness **computationally free**.

UNIVERSITY OF TWENTE.

VALIDATED ON A WEEK OF REAL OPERATIONS

PLANNING FOR UNCERTAINTY PAYS — TWICE

4.6–36.5%

cost savings vs. deterministic

Value of the Stochastic Solution, across a week of task sets

≈7%

left for perfect foresight

EVPI: 40 scenarios already capture most of the attainable benefit

+10.2%

from staff exchangeability

letting teams cover each other's work, across departments and skills

THREE MAIN CONCLUSIONS

01 **Modeling uncertainty creates substantial value** — especially when demand is volatile, understaffing is costly, and workload peaks are narrow or multiple; deterministic schedules break on temporal mismatch.

02 **Perfect information adds little** — a well-designed stochastic model already captures most of the achievable benefit from improved foresight.

03 **Break flexibility helps, but modestly** — its contribution is small relative to the gains from stochastic planning itself and from employee exchangeability.

THE BYCATCH: EXCHANGEABILITY

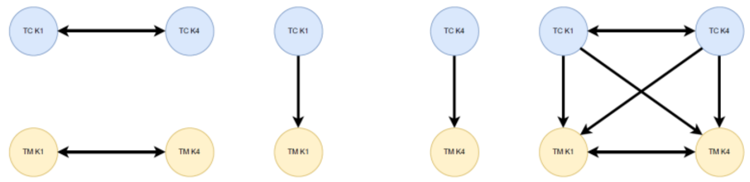

(a) Horizontal (b) Vertical (c) Cross

Exchangeability is the key driver of robustness — covering tasks across pools and skill levels consistently improves cost efficiency, especially under correlated disruptions. Prioritize structural flexibility before sophisticated break-adjustment policies.

For practitioners: when penalty parameters are hard to calibrate, service-level constraints directly manage the cost–reliability trade-off.

MANAGING UNCERTAINTY IN AIRLINE GROUND OPERATIONS · WHAT MADE IT WORK

STOCHASTIC STAFF SCHEDULING
THE THREE INSIGHTS, AT WORK

INSIGHT 01 · STOCHASTICITY

Plan against scenarios, not the average day

Empirically fitted delay distributions drive a representative scenario set; the roster is chosen to perform **across** them — cutting understaffing, overstaffing and last-minute interventions.

INSIGHT 02 · INTEGRALITY

Pool the workforce across boundaries

Departments and skill levels optimized jointly, with staff exchange as recourse — the interdependency a per-department roster can never exploit, worth ≈10% on its own.

INSIGHT 03 · CUSTOM ALGORITHMS

Reformulate before you optimize

The task-to-period reformulation — not a bigger solver — shrank the model by 99.7% and made scenario-based optimization run in seconds, at operational scale.

NEXT STEPS

- Adopt DAM/SDAM to streamline today's two-step shift-type selection
- Richer telemetry: task-duration variability, not only flight delays
- Correlated disruptions: weather, IT, time-of-day structure
- Generalize beyond luggage handling to other turnaround activities

TAKEAWAY

A reformulation that makes the model **small enough to plan against uncertainty itself** — robust rosters at real scale, in real planning runtimes.

PLANNING GATE AGENTS AT SCHIPHOL AIRPORT

Bas Nijmeijer, Ron Kaslasi, Ali Poursaeidesfahani, Marjan van den Akker, Han Hoogeveen

PLANNING GATE AGENTS AT SCHIPHOL AIRPORT · THE PROBLEM

BOARDING RUNS ON PEOPLE, IN SYNC
A STATIC DECISION TREE PLANS A DYNAMIC OPERATION

Today each flight gets its agents based on a **static decision tree**; but agent need varies with aircraft type, destination, passenger mix and regulations. Under-staff a gate and boarding delays cascade; over-staff it and scarce agents idle.

≈150 *)
KLM departures per day at Schiphol to plan agents for

≈ 300
gate and service agents working in shifts with strict start and end times

NP-hard
related to multi-depot vehicle scheduling, with synchro-nized non-preemptive tasks

Agents ↔ time
more agents shorten boarding, a discrete time/resource trade-off inside the schedule

WHAT MAKES IT HARD

- **Synchronization:** several agents must be at the gate at the same moment
- **Walking times** between gates couple every assignment to the next one
- **Tight coupling:** a small disruption propagates through a schedule without slack, it’s the trade-off between efficiency and robustness

The aim: **data-driven agent planning** — fewer delay-induced costs, robust schedules with real buffers.

*) Provided instances contained about 150 departures, in general around 500 departures

SPLIT THE PROBLEM WHERE IT BENDS
ANNEALING DECIDES, A FLOW NETWORK SCHEDULES

1 **Master: explore the hard decisions**

Simulated annealing searches over **team compositions per flight** and deliberate delays — the NP-hard time/resource trade-off — minimizing delay-related costs.

2 **Subproblem: build the schedule exactly**

For each candidate, a **min-cost max-flow** model routes agents through tasks — respecting shift bounds and walking times, rewarding slack, penalizing distance. A repair heuristic links the correct start- and endtimes

3 **Iterate**

The flow model is polynomial — fast enough to evaluate **every** SA candidate with an exact, feasibility-checked schedule instead of an estimate.

THE SUBPROBLEM AS A FLOW NETWORK

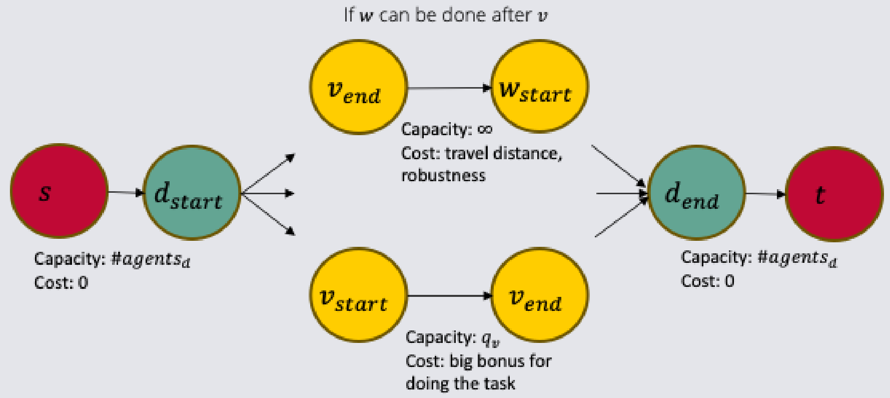

Nodes	every shift and every task (needing qv agents at once)
Task → task	only if reachable in time — cost carries walking distance and a **reward for slack**
Coverage	a large bonus for doing each task pulls flow through every flight

Agents flow from source to sink; the cheapest max flow *is* the agent schedule — robustness priced in, not patched on.

AIRFRANCEKLM GROUP

PLANNING GATE AGENTS AT SCHIPHOL AIRPORT · RESULTS

NEARLY EVERY FLIGHT BOARDED ON TIME
ROBUSTNESS DOMINATES SCHEDULE QUALITY

≈100%

of flights planned without introducing delays

on real KLM flight data from Schiphol *)

5–7 min

to a full feasible day plan

well within a planning cycle — and flexible where the decision tree is rigid

Slack wins

robustness dominates schedule quality

buffer time between tasks outweighs every other objective term

WHAT THE PLANS LOOK LIKE

Agents are used **more efficiently** than under the rule tree: teams carry over between nearby flights, walking is kept short, and slack is placed where disruptions would otherwise cascade.

WHERE ELSE IT APPLIES

Any operation with **synchronized teams, shift bounds and disruption risk** : ground handling, security staffing, healthcare logistics — the framework is domain-agnostic.

The striking finding: shaving walking minutes matters far less than placing buffers; in a tightly coupled airport, slack is the cheapest insurance.

*) Not all duty desk tasks were included

Utrecht University

GATE AGENT PLANNING
THE THREE INSIGHTS, AT WORK

INSIGHT 01 · STOCHASTICITY

Buy buffers before the disruption

Uncertainty is answered **proactively** : slack between tasks is rewarded inside the objective, and it proves to be the term that dominates schedule quality — robustness by construction, not repair.

INSIGHT 02 · INTEGRALITY

Staffing, timing & routing: one problem

Team size trades against boarding duration, deliberate delays are a planning lever, and walking times couple every gate — decided jointly where the rule tree decided flight by flight.

INSIGHT 03 · CUSTOM ALGORITHMS

A heuristic with an exact engine

The decomposition puts simulated annealing on the NP-hard decisions and a **polynomial min-cost max-flow** on schedule construction — each part doing what it is best at.

NEXT STEPS

- Real-time disruptions and partial schedule execution
- Tighter integration with existing airline planning systems
- Transfer: ground handling, security staffing, healthcare logistics

TAKEAWAY

Advanced optimization, **tailored to operational reality** , turns rule-based staffing into data-driven planning — reliable boarding and efficient agents at airport scale.

LUGGAGE HANDLING RECOVERY

Lisanne Heuseveldt, Ali Poursaeidesfahani, Marjan van den Akker, Egon van den Broek

AIRFRANCEKLM GROUP

LUGGAGE HANDLING RECOVERY · THE PROBLEM

SYNCHRONIZED TEAMS, MOVING PLATFORM, NO TIME
ONE LATE ARRIVAL IMPACTS THE WHOLE TASK ASSIGNMENT

Loading and unloading is **routing and scheduling at once** — teams and equipment cross a large airport area — and tasks are synchronized on several levels. Aircraft arrivals shift, handlers call in sick, and a feasible plan becomes infeasible in minutes.

LONG TURNAROUND — DIFFERENT TEAMS, MINIMUM GAP

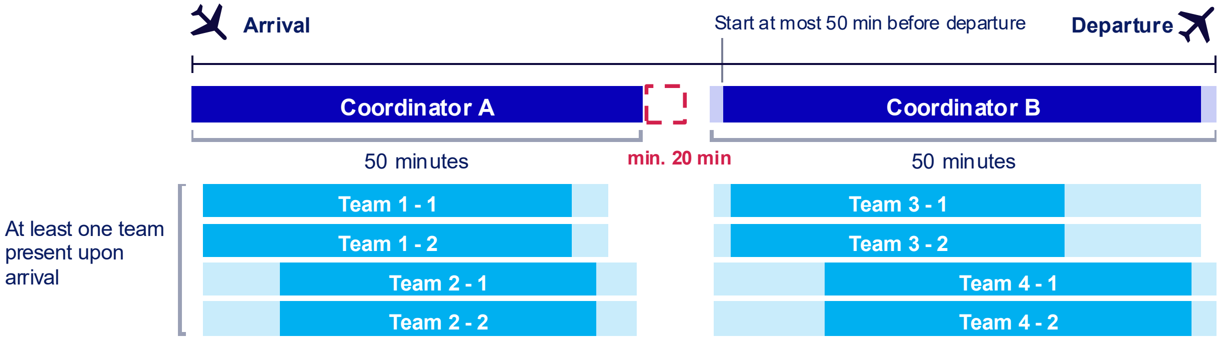

SHORT TURNAROUND — SAME COORDINATOR, SAME TEAMS

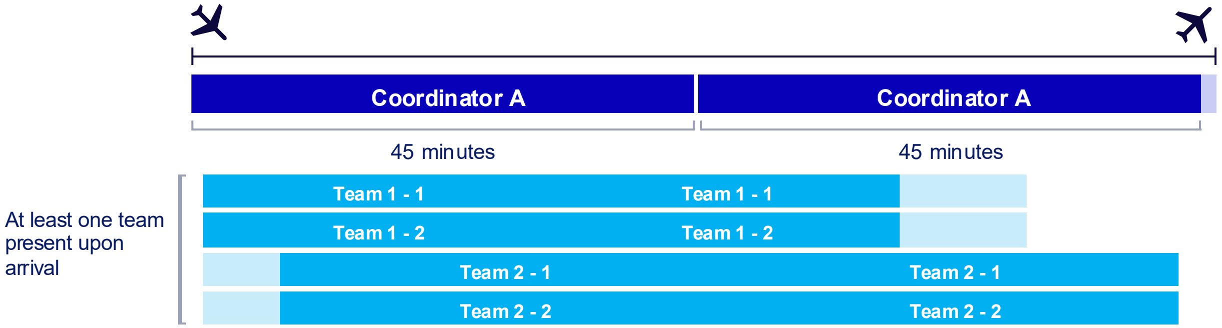

- A **coordinator** must be present during the work of every (un)loading team — and at least one team upon arrival
- Departure work starts **at most 50 min before departure** ; arrival and departure need different teams — except on short turnarounds, where they must be the *same*

THE SPEED WALL

A mathematical optimizer needs **far too long** for a live platform decision. Heuristics can find good solutions **within seconds** — and seconds is the budget.

Utrecht University

AIRFRANCEKLM GROUP

LUGGAGE HANDLING RECOVERY · THE KEY TRICK

LET THE GROUND PROPOSE THE DELAY
IF A DELAY IS UNAVOIDABLE, CHOOSE ONE THAT HELPS

TRADITIONALLY
Ground handling follows the OCC's timing

The flight schedule is an input. Handling absorbs whatever it gets — and when scarce resources make it infeasible, delays happen anyway, **unchosen and uncoordinated**.

OCC → GROUND HANDLING

IN THIS WORK
Delay becomes a decision variable

A **delay neighbor** in the search lets recovery propose the **least bad delay** — priced against unassigned tasks and changes, so the trade-off is explicit rather than emergent.

OCC ↔ GROUND HANDLING

AND: CHANGE ITSELF IS A COST

Repair, don't replan from scratch
The search focuses on the **disrupted parts** of the schedule — the rest of the day stays as the crews already know it

Penalty for plan changes
A scenario prices **every change** to the original assignment — because reshuffles are hard to execute on the ramp

Unassign, then rebuild
Overlapping tasks are never allowed: conflicting tasks become **unassigned at a penalty**, the search must place them again

Utrecht University

AIRFRANCEKLM GROUP

LUGGAGE HANDLING RECOVERY · APPROACH & RESULTS

CONFLICT-DIRECTED SEARCH, SEVEN SECONDS
RECOVERY AT DECISION SPEED

THE EXPERIMENTAL PIPELINE

Initial schedule *)
real operational plan - optimizer output mixed with manual decisions

→

Disrupted schedule
delays and absences injected, many tasks left unassigned

→

Resolved schedule
conflict-directed local search + simulated annealing

Conflict-directed	moves target the tasks actually in conflict — effective changes, not random ones
Simulated annealing	specific local search that escapes the local optima that a pure repair heuristic settles into
Three scenarios	base · penalty for changes · delay suggestions
Next step	min-cost max-flow in the subproblem — the same engine as the gate agent case

20-25%
average score improvement
from heavily disrupted schedules with many unassigned tasks

7-15 s
average runtime
real-time viable on the platform — real KLM data

Ongoing work. The delay neighbor and the change penalty are being weighed against each other: how much churn, and how much proposed delay, an operation is actually willing to accept.

*) Heuristic repair to make the plan feasible

Utrecht University

AIRFRANCEKLM GROUP

LUGGAGE HANDLING RECOVERY
THE THREE INSIGHTS, AT WORK

INSIGHT 01 · STOCHASTICITY

Repair, don't replan from scratch

Disruption is met where it lands: only the **disrupted parts** are reopened, so the recovered plan stays recognizable to the crews who have to execute it within seconds.

INSIGHT 02 · INTEGRALITY

Integration pointing upward

Ground handling stops being a pure taker of the flight schedule: by **proposing delays back to the OCC** , scarce platform resources enter the timing decision itself.

INSIGHT 03 · CUSTOM ALGORITHMS

Seconds is a hard constraint

Multi-level synchronization and routing put an exact optimizer out of reach in real time; conflict-directed local search with annealing recovers **20-25% in 7-15 seconds**.

WHERE THE PATTERN TRAVELS

- Turnaround coordination across all ground processes
- Gate and stand support with shared scarce equipment
- Any crew-task problem where synchronization is the binding constraint

TAKEAWAY

Fast, minimal, feasible — and allowed to **talk back to the flight schedule** : recovery as a negotiation between the ramp and operations control.

Utrecht University

AIRFRANCEKLM GROUP

CONCLUSIONS

SIX USE CASES, THREE LESSONS AND ONE INVITATION

STOCHASTICITY

Planning for the average is planning to fail

Every use case gained by planning against distributions instead of averages — schedules that hedge, not hope.

INTEGRALITY

The best local decision is rarely the best decision

Fleet with crew, gates with rosters, delays with passengers — optimizing across silos beat optimizing within them. Every time.

SPEED

A perfect answer that arrives late is a wrong answer

Tailor-made algorithms delivered near-optimal decisions in seconds — fast enough for the operation, not just the paper.

THE COLLABORATION

Working with **Utrecht & Twente** paid off: **clear potential identified** across use cases, new methods, and a knowledge base we keep — now it must earn its way into the operation.

Bridge your silos. Embrace the uncertainty. Out-pace it with tailor-made algorithms.
Today we buy robustness in minutes of slack — tomorrow, let the buffer live **in the algorithm, not on the clock**.


20 YEARS
AIRFRANCEKLM
GROUP
ON BEHALF OF THE ENTIRE CREW
I'D LIKE TO THANK YOU FOR “FLYING” WITH ME TODAY

APPENDIX

ACADEMIC COLLABORATION PORTFOLIO

PUBLICATIONS, THESES AND RESEARCH CONTRIBUTIONS

Utrecht University × University of Twente × KLM · consolidated overview, August 2026

APPENDIX · ACADEMIC OUTPUT

17 FORMAL ACADEMIC OUTPUTS
DELIVERED AND IN THE PIPELINE

Completed theses, publication-stage research and conference contributions; **delivered outcomes and a continuing research pipeline**.

9

COMPLETED THESES

7 Master's · 2 Bachelor's

4

ARTICLES & PREPRINTS

2 peer-reviewed · 2 arXiv

4

CONFERENCE CONTRIBUTIONS

AGIFORS and JOPT · 2024–26

Excluded for now: Luggage Handling Recovery — ongoing/unpublished, no citable output yet.

AIRFRANCEKLM GROUP

APPENDIX · ACADEMIC OUTPUT

FOUR PUBLICATION-STAGE OUTPUTS
THE STRONGEST FORMAL EVIDENCE

PEER-REVIEWED

Assink, J., Dahanayaka, M., Prak, D., Gupta, R., & Mes, M. (2026). *Integrated shift and break scheduling for airline ground staff under flight delay uncertainty.* European Journal of Operational Research. Advance online publication.

doi.org/10.1016/j.ejor.2026.07.008

PEER-REVIEWED

Dahanayaka, M., Prak, D., & Mes, M. (2026). *From gate to runway: A systematic review of airport ground operations optimization.* Journal of Air Transport Management, 135, 103013.

doi.org/10.1016/j.jairtraman.2026.103013

PREPRINT

de Bruin, P., van den Akker, M., Kumar, K., Heuseveldt, L., & Paelinck, M. (2025). *Integrated airline fleet and crew recovery through local search* [Preprint]. arXiv.

doi.org/10.48550/arXiv.2505.04274

PREPRINT

de Bruin, P., Elderhorst, B., van den Akker, M., & Hoogeveen, H. (2026). *Simulation strategies for an efficient local search to solve stochastic scheduling problems* [Preprint]. arXiv.

doi.org/10.48550/arXiv.2605.23675

FOUR MASTER'S THESES AT UTRECHT
SCHEDULING, ML, ROUTING, PREDICTION

MASTER'S THESIS

Boot, S. (2024). *Predicting the amount of excess hand luggage on an aircraft using machine learning* [Master's thesis, Utrecht University].

MASTER'S THESIS

de Heij, M. (2025). *Machine learning augmented (re-)optimization of a mixed integer linear programming scheduling problem* [Master's thesis, Utrecht University].

MASTER'S THESIS

Diele, M. (2025). *Solving a technician routing and scheduling problem using column generation with a genetic algorithm* [Master's thesis, Utrecht University].

MASTER'S THESIS

Nijmeijer, B. (2025). *Planning gate agents at Schiphol Airport* [Master's thesis, Utrecht University].

FIVE COMPLETED THESES AT TWENTE
OPERATIONS AND PREDICTION

MASTER'S THESES

Assink, J. G. (2025). *Managing uncertainty in airline ground operations: A stochastic staff scheduling approach* [Master's thesis, University of Twente].

Bergsma, Y. (2025). *Estimating misconnecting passengers to enhance decision-making on the minimal connecting time for KLM* [Master's thesis, University of Twente].
purl.utwente.nl/essays/107742

van Oost, L. (2025). *Maintenance slot capacity planning under uncertainty: Integrating workload prediction and stochastic optimisation* [Master's thesis, University of Twente].
purl.utwente.nl/essays/107912

BACHELOR'S THESES

Heinrich, C. V. (2023). *Hand luggage overflow prediction at KLM* [Bachelor's thesis, University of Twente].
purl.utwente.nl/essays/96988

Ochoa Barnuevo, M. L. (2023). *Enhancing baggage handling duration predictions for KLM: A data-driven and machine learning approach using camera and sensor data* [Bachelor's thesis, University of Twente].
purl.utwente.nl/essays/96204

APPENDIX · CONFERENCE CONTRIBUTIONS

FOUR CONFERENCE CONTRIBUTIONS
EXTENDING THE PIPELINE, 2024–2026

CONFERENCE ABSTRACT

de Bruin, P., van den Akker, M., Kumar, K., Heuseveldt, L., & Paelinck, M. (2024). *Simulated annealing to solve the integrated airline fleet and crew recovery problem.* Abstract from 2024 AGIFORS Airline Operations Conference, Hamburg, Germany.

Utrecht University research record

CONFERENCE PRESENTATION

Montaruli, A., Kieskamp, M., Prak, D., Mes, M., & Birolini, S. (2025). *Multi-hub airline network planning model* [Conference presentation]. AGIFORS Symposium 2025.

AGIFORS Symposium programme

CONFERENCE PRESENTATION

Dahanayaka, M. (2026). *A simulation-optimization approach for airline ground staff scheduling* [Conference presentation]. Joint Optimization Conference, Montreal, Canada.

JOPT 2026 programme

CONFERENCE PRESENTATION

Montaruli, A. (2026). *Deep reinforcement learning for airline revenue management: A pricing network-based approach* [Conference presentation]. Joint Optimization Conference, Montreal, Canada.

JOPT 2026 programme