The Price of a Promise:

Faster Scheduling for Automated Kitchens

Matteo Bonsignore · Nelson Lu

Faculty Advisor Prof. Alexandre Jacquillat, MIT Sloan

Company Advisor Evan Fox, Senior Data Scientist · Wonder Group, Inc.

THE PROBLEM

1. **Wonder's Infinite Kitchens** cook many restaurant concepts at once; the **Kitchen Display System (KDS)** decides when every cooking task is released to staff.
2. That one decision — *sequencing* — drives delivery speed, food quality, labor, and compute cost.
3. No single algorithm fits: customer ETAs need **milliseconds**; the kitchen's schedule gets **tens of seconds**.

320,000 orders per year suffer from "ETA shock"

Quotes ignore **live kitchen capacity**, so the ETA drifts >10 min. Affected customers return at **76.9% vs 83.3%** — a 6.4 pp retention gap on above-average baskets.

THE DATA

3.3M completed orders analyzed (trailing 180 days)	**19M+** historical order items
100+ Infinite Kitchens network-wide	**3TB+** production data · MySQL, MongoDB, BigQuery

Worked **embedded in Wonder's production environment** — same data, code, and review process as the sequencing team.

Benchmark: **1,000 historical kitchen states**, re-solved as a **deterministic paired A/B**.

OUR APPROACH — ONE JOB, TWO TIME SCALES, TWO ENGINES

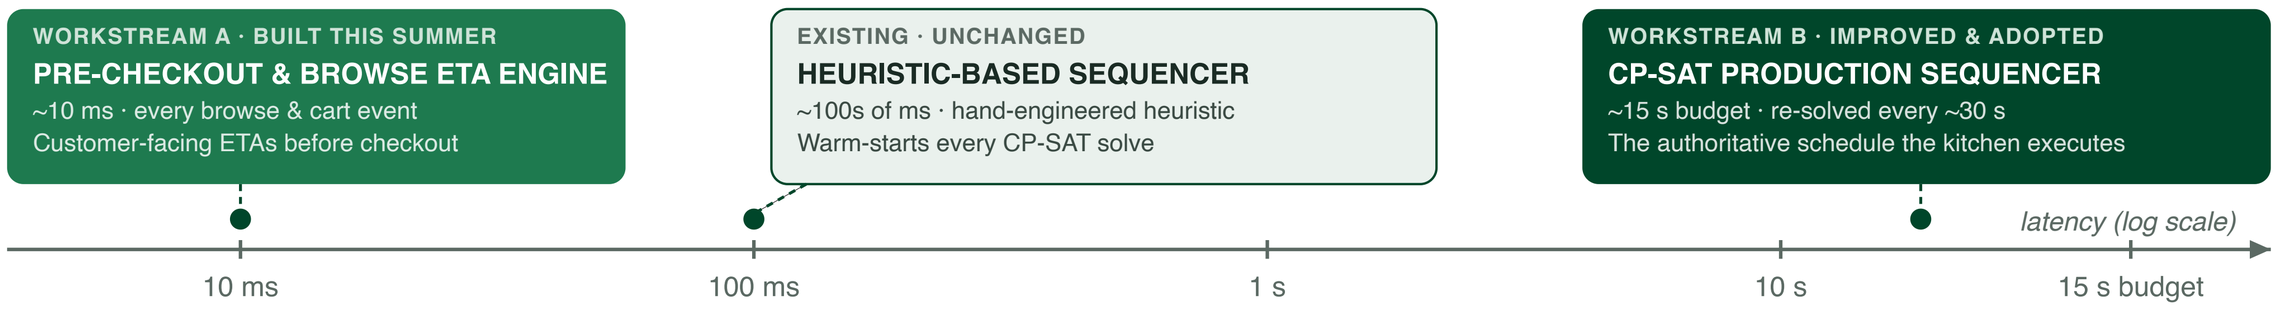

WORKSTREAM A — AN ETA ENGINE FAST ENOUGH FOR EVERY BROWSE

The solver can never answer a browsing customer — so we built a **kitchen-assessment engine**: a simulation that overlays the basket onto the live kitchen state and returns a capacity-aware ETA in **under 10 ms**.

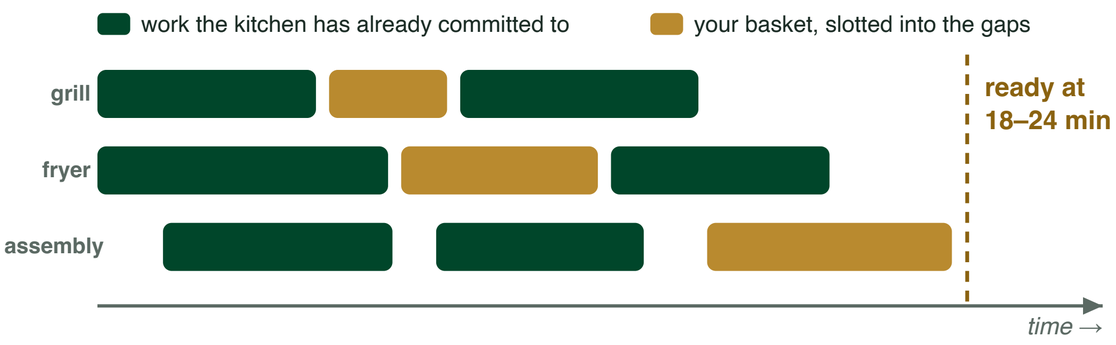

1. **RECONSTRUCT** — Rebuild the kitchen's committed schedule from the optimizer's published per-step timings
2. **OVERLAY BASKET** — Insert the customer's hypothetical basket items onto live pod queues
3. **ETA FORECAST** — Nominal point ETA prediction → "Ready in 18 min"

1.7 ms per answer, vs a 10 ms budget	**4.3 min** held-out MAE vs 5.3 incumbent ML (–19%)	**95.3%** baskets successfully placed

The same engine unlocks adjacent wins at negligible cost: **appliance auto-pausing**, add-to-order ETAs, upsell gates, courier-dispatch timing.

WORKSTREAM B — BREAKING SYMMETRY IN THE PRODUCTION SOLVER

The insight. The solver's search space is **dominated by symmetry**: interchangeable kitchen steps form **k!-sized orbits** of equal-objective twin solutions that CP-SAT must refute one by one.

1. **Define the automorphisms exactly** — steps are interchangeable only if identical on every attribute the model *actually reads*.
2. **Make the model label-invariant first** — the audit also fixed a wrong-pod bug (**–7% median solve**).
3. **Pin one representative per orbit** — $start(a) \le start(b)$: duplicates pruned, no distinct schedule deleted.
4. **Prove safety, not just speed** — paired A/B on 1,000 states: **100% of proven optima preserved**.

TWO ORDERINGS, ONE PLAN

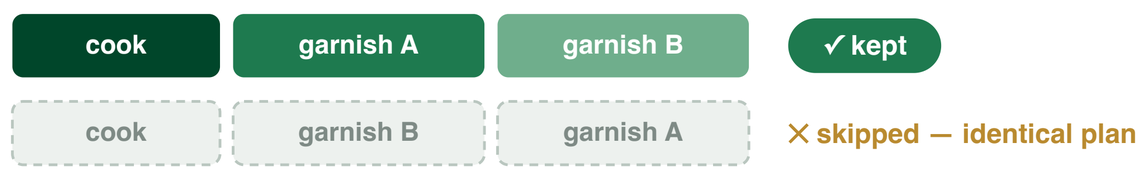

Steps identical in every way → the solver keeps exactly one ordering: $start(a) \le start(b)$

–7.3% total solve time, n=1,000	**–20%** median solve time	**72.6→75.4%** proven-optimal rate	**100%** proven optima preserved

✔ MERGED & RUNNING NETWORK-WIDE ACROSS 100+ KITCHENS

BUSINESS IMPACT

Example context: ~18,000 orders/day · $35 AOV · ~$235M annual gross order value

SIZED OPPORTUNITY — Pre-checkout ETA engine

$3M /yr recoverable

- ✓ **~320,000 ETA-shock events/yr** addressed at the source
- ✓ **6 pp retention gap**, half of it closed on ~20,000 at-risk customers/yr
- ✓ **Validated by design**: accuracy-gated rollout, A/B-confirmed impact

EST. IMPACT — Production sequencer

$1M /yr est.

- ✓ **–7% solver compute (–20% typical solve)** — direct infra savings
- ✓ **Better peak-hour schedules**: 4–7% more plans proven optimal at mid-load
- ✓ **Scalability headroom**: growth before quality erodes at the fixed 15 s budget

TECH STACK

Python · OR-Tools CP-SAT · Java · MySQL · MongoDB · BigQuery

With gratitude to the Wonder Sequencing & KDS teams and Prof. Alexandre Jacquillat · MIT Master of Business Analytics · Capstone Showcase · August 2026