Reserving Capacity, Restoring Fairness: Scheduling Elective Caesarean Sections in the NHS

London Business School

Yuhang Du [1] Yueyang Zhong [1] Catherine Aiken [2] Pedra Rabiee [2]

[1]Management Science and Operations, London Business School [2]Department of Obstetrics and Gynaecology, University of Cambridge

Motivation

NHS maternity services face persistent capacity constraints. Demand for elective caesarean sections generally exceeds available operating theatre slots, yet no national guidance exists on managing the resulting system congestion.

- **Early engagers (E):** Requests placed well in advance of the delivery deadline (e.g., previous caesarean sections).
- **Late engagers (L):** Requests placed closer to the delivery deadline (e.g., breech, fetal cause, unstable lie).

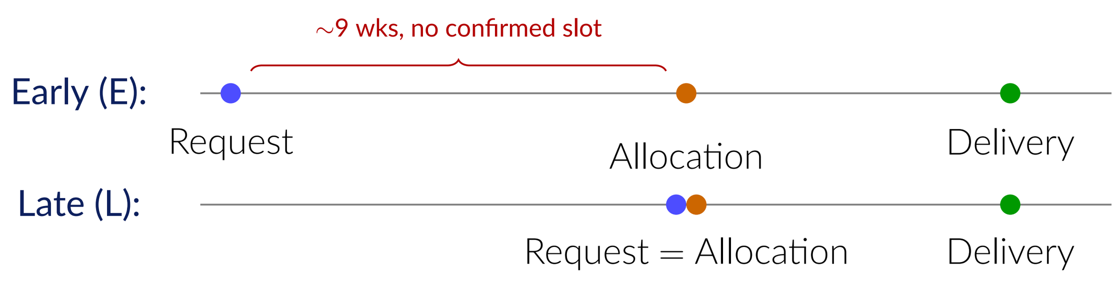

Scheduling Policy: Earliest-Deadline-First (EDF)

When an operating theatre slot opens, the patient whose delivery deadline is nearest is scheduled next.

Under EDF: Nobody Is Happy

- **Early engagers (unsatisfied):** Wait ∼9 weeks with no confirmed slot. Repeated follow-up calls or even in-person visits to scheduling midwives.
- **Late engagers (unprotected):** Despite complex indications, receive no scheduling priority over early engagers.

T: Lead time (request to deadline) S: Service time ($S \le T$)

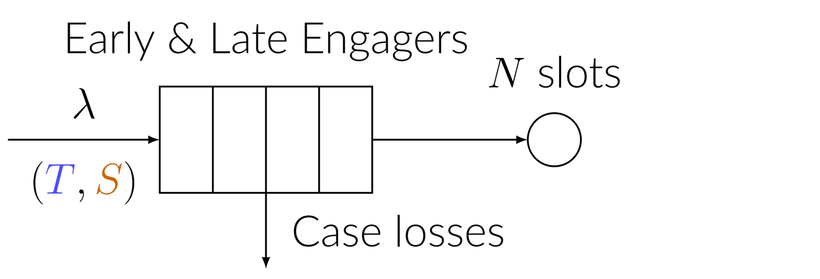

Figure 1. Pooled queue: both types compete for N slots under EDF.

Model & Fluid Analysis

- **Arrivals.**
  - Poisson process: rate λ.
  - Early engagers' arrival rate $\lambda^E = \alpha\lambda$;
  - Late engagers' arrival rate $\lambda^L = (1-\alpha)\lambda$.
- **Requests.**
  - Service time $S_i \sim f_S$;
  - Lead time $T_i \sim f_T$ (request to deadline).
- **System.**
  - N slots & overloaded ($\lambda\mathbb{E}[S] > N$);
  - Non-preemptive service.

Theorem 1 (steady-state threshold)

Let $h(\tau) := \int_0^\tau \sigma f_S(\sigma) d\sigma$. Under EDF in overload,

$$\tau^* = h^{-1}(N/\lambda)$$

Case loss fractions: $\Pr(S^E > \tau^*)$ and $\Pr(S^L > \tau^*)$. Case loss rates: $A^E = \lambda^E \Pr(S^E > \tau^*)$ and $A^L = \lambda^L \Pr(S^L > \tau^*)$.

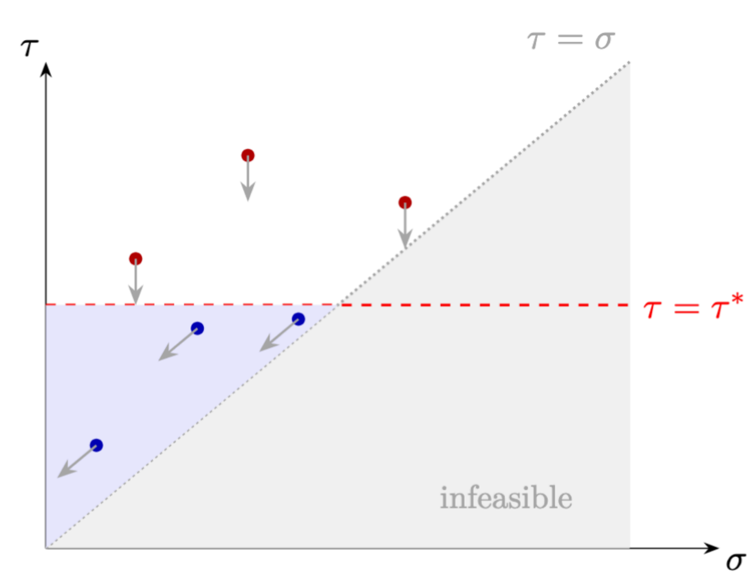

Figure 2. Phase diagram of the EDF fluid model (σ represents remaining service time and τ represents remaining time to deadline). τ^* separates service (blue) and waiting (white) regions.

Key Insight

All non-idling policies (EDF, FCFS, LCFS, LAS) produce the same case loss fractions under overload. **Only capacity reallocation can help.**

Capacity Reservation Policy

Reserve fraction $\gamma \in [0,1]$ of N slots for early engagers. The system splits into two dedicated queues, each under EDF.

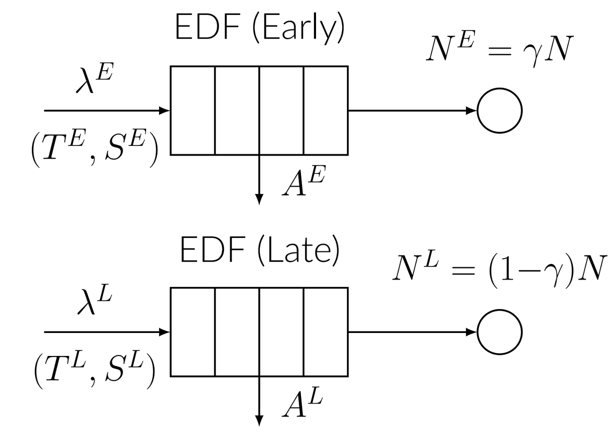

Figure 3. Dedicated queues under the reservation policy.

Define $h^k(\tau) := \int_0^\tau \sigma f_{S^k}(\sigma)\, d\sigma$ for each type $k \in \{E, L\}$.

Proposition 1 (type-specific thresholds and case loss rates)

$$\tau^{E*}(\gamma) = (h^E)^{-1}(\gamma N/\lambda^E) \qquad A^E(\gamma) = \lambda^E \Pr(S^E > \tau^{E*}(\gamma))$$

$$\tau^{L*}(\gamma) = (h^L)^{-1}((1-\gamma)N/\lambda^L) \qquad A^L(\gamma) = \lambda^L \Pr(S^L > \tau^{L*}(\gamma))$$

Theorem 2 (optimal γ^*)

Minimise the total case-loss cost $C(\gamma) = c^E A^E(\gamma) + c^L A^L(\gamma)$ subject to $|A^E(\gamma)/\lambda^E - A^L(\gamma)/\lambda^L| \le \Delta$.

$$\gamma^* = \begin{cases} \underline{\gamma}, & \text{if } \gamma^u < \underline{\gamma} \text{ (early-engager protection binds)} \\ \gamma^u, & \text{if } \underline{\gamma} \le \gamma^u \le \overline{\gamma} \text{ (unconstrained optimum)} \\ \overline{\gamma}, & \text{if } \gamma^u > \overline{\gamma} \text{ (late-engager protection binds)} \end{cases}$$

Corollary: γ^u equalises priority indices: $c^L/\tau^{L*}(\gamma^u) = c^E/\tau^{E*}(\gamma^u)$.

Proposition 2 (early allocation — "free lunch")

Optimal γ^* and case losses are **invariant** to allocation time. Confirm slot **upon request** with no performance loss.

Early-Engager Satisfaction

Each reserved slot allocated at request time generates a **satisfaction value** $v \ge 0$, reducing patient anxiety and eliminating follow-up contacts.

Theorem 3 (extended cost)

Minimise the extended cost $\tilde{C}(\gamma) = c^E A^E(\gamma) + c^L A^L(\gamma) - v\gamma N$ subject to $|A^E(\gamma)/\lambda^E - A^L(\gamma)/\lambda^L| \le \Delta$.

Optimal $\tilde{\gamma}^*$ has same piecewise structure as Theorem 2. As v increases, $\tilde{\gamma}^*$ shifts toward greater early-engager protection.

Case Study: Cambridge University Hospitals

Data (2024) & Parameters

1,202	18	73/27%	1.32×
C-section patients	slots/week	early/late split	demand/supply

Dedicated-EDF vs. Pooled-EDF

Policy	γ^*	N^E	N^L	p^E(%)	p^L(%)	Case-loss cost	Cost savings
Pooled-EDF	–	–	–	22.2	22.0	337.6	–
Ours Δ=20%	.66	11.9	6.1	29.6	9.7	322.3	**+15.3**
Ours Δ=15%	.68	12.2	5.8	27.1	13.9	327.7	+9.9
Ours Δ=10%	.70	12.6	5.4	25.3	17.5	335.7	+1.9
Ours Δ=5%	.71	12.8	5.2	24.2	19.7	340.5	−2.9

c^E=1, c^L=2, v=0. N^E=γN, N^L=$(1-\gamma)N$. Mean across 50 replications.

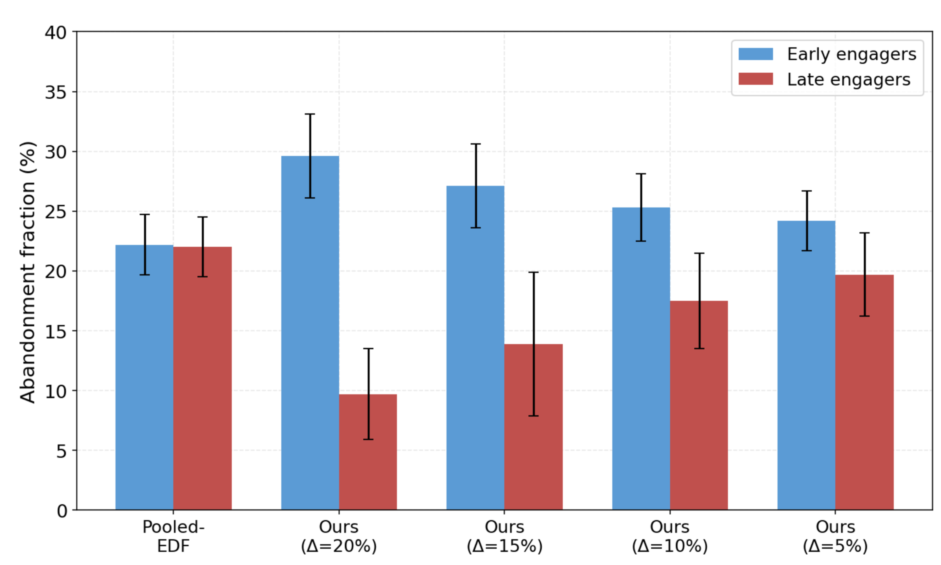

Figure 4. Case loss fractions under Pooled-EDF and Dedicated-EDF at different equity tolerances.

Key Contributions

1. Problem. Early (low-risk) vs. late (high-risk) elective caesarean section cases create a tradeoff between rewarding early action and clinical priority.

2. Methodology. Rigorous queueing model with fluid analysis that fully characterizes the optimal policy balancing efficiency and fairness.

3. Impact. Easy-to-implement slot reservation by type that boosts early-engager satisfaction and guarantees late-engager priority.

Effect of Satisfaction Level v

v	$\tilde{\gamma}^*$	N^E	N^L	p^E(%)	p^L(%)	Case-loss cost	Ext. cost
0	0.66	11.9	6.1	29.6	9.7	322.3	322.3
0.1	0.67	12.1	5.9	28.3	11.5	322.8	260.1
0.3	0.67	12.1	5.9	28.3	11.5	322.8	134.7
0.5	0.69	12.4	5.6	26.4	15.3	330.5	7.6
0.7	0.72	13.0	5.0	23.2	21.7	344.5	−127.3
1.0	0.77	13.9	4.1	18.1	34.3	381.4	−339.3

Dedicated-EDF, Δ=20%, c^E=1, c^L=2. N^E=γN, N^L=$(1-\gamma)N$. Mean across 50 replications. Negative ext. cost = satisfaction benefit exceeds case-loss cost.

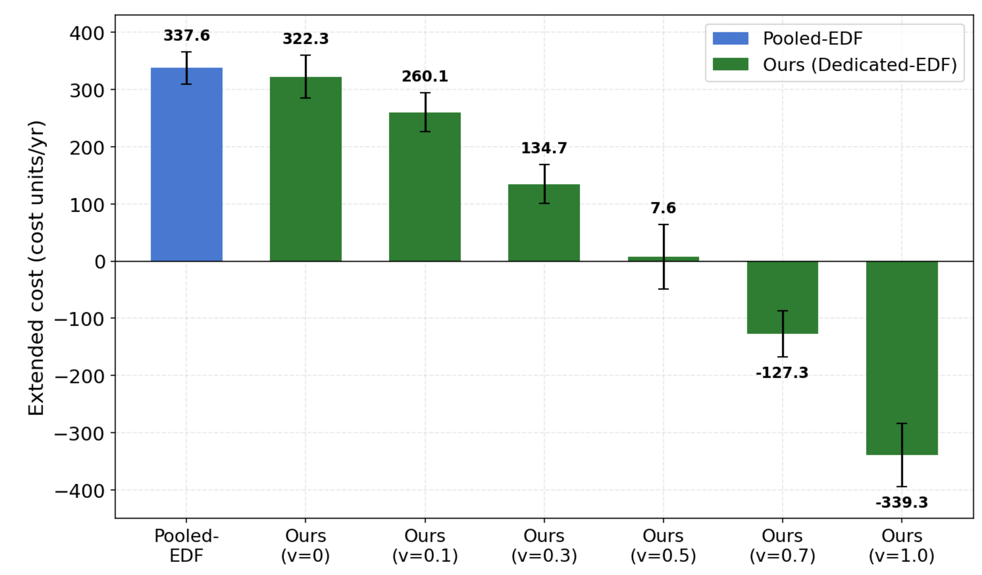

Figure 5. Extended cost under Pooled-EDF and Dedicated-EDF at different satisfaction levels v.

Sensitivity & Robustness

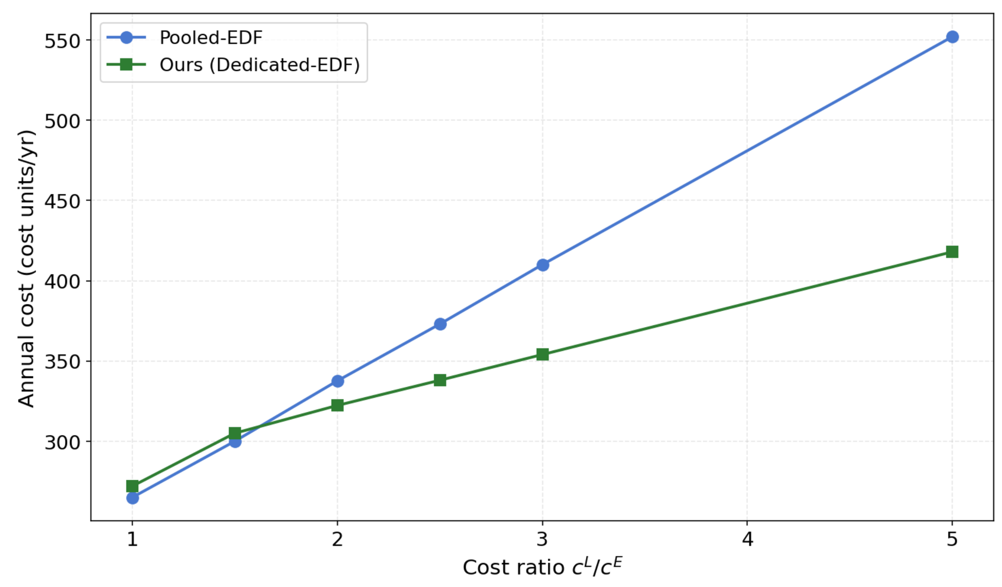

Figure 6. Case-loss cost vs. cost ratio: Dedicated-EDF (Δ=20%, v=0) vs. Pooled-EDF.

Parameter	Range	γ^*	Cost savings
No-show rate	0–5%	.66–.67	+15.3 to +16.3
Spillover (N)	16–18	.66	+15.3 to +30.5
Demand (α)	.50–.80	.44–.73	+0.1 to +36.5

Dedicated-EDF, Δ=20%, c^E=1, c^L=2, v=0. Mean across 50 replications.

Key Findings & Clinical Impact

Status quo (Pooled-EDF): 22% case losses for both groups; early engagers get no advance confirmation, and high-risk late engagers receive no priority.

Our policy (Dedicated-EDF with Δ=20%, γ^*=0.66):

- **Early engagers gain certainty and satisfaction:** > 70% (∼12/week; 618/year) receive a confirmed slot upon request.
- **Late engagers gain clinical priority:** Case loss fraction drops from 22% to 9.7% (∼0.8 fewer losses per week; 40 fewer per year).
- **System gains efficiency:** Case-loss cost decreases by 4.5%; with satisfaction valued (v=0.5), extended cost drops by 97.6%.