

Precisely, if $\pi_b^{(w)}$ is the execution order of block b under the window w , we can define

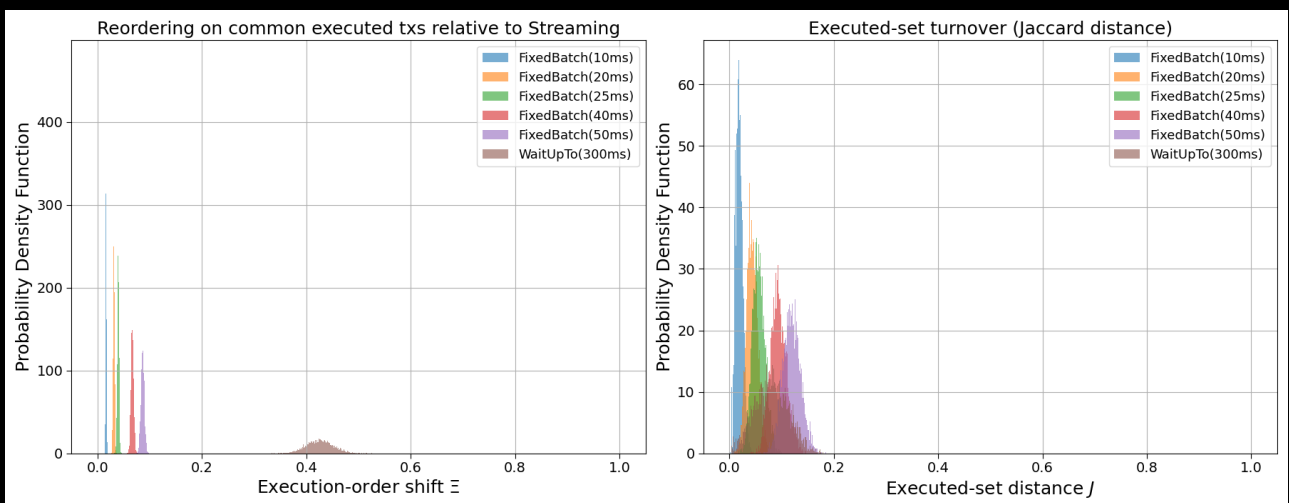
$$\bar{d}_b(w_1, w_2) = d\left(\pi_b^{(w_1)}, \pi_b^{(w_2)}\right),$$

where d is a normalized footrule distance, computed over the intersection of executed txs in both cases. This then helps us to answer the question

How much does the order of realization change when we increase w ?

The Jaccard distance is computed on the sets of executed txs. This helps us to answer the question

How does the set of txs executed change when I change w ?



For FixedBatch, both quantities increase smoothly with the batch window, indicating a gradual displacement away from Streaming in both ordering and selection.

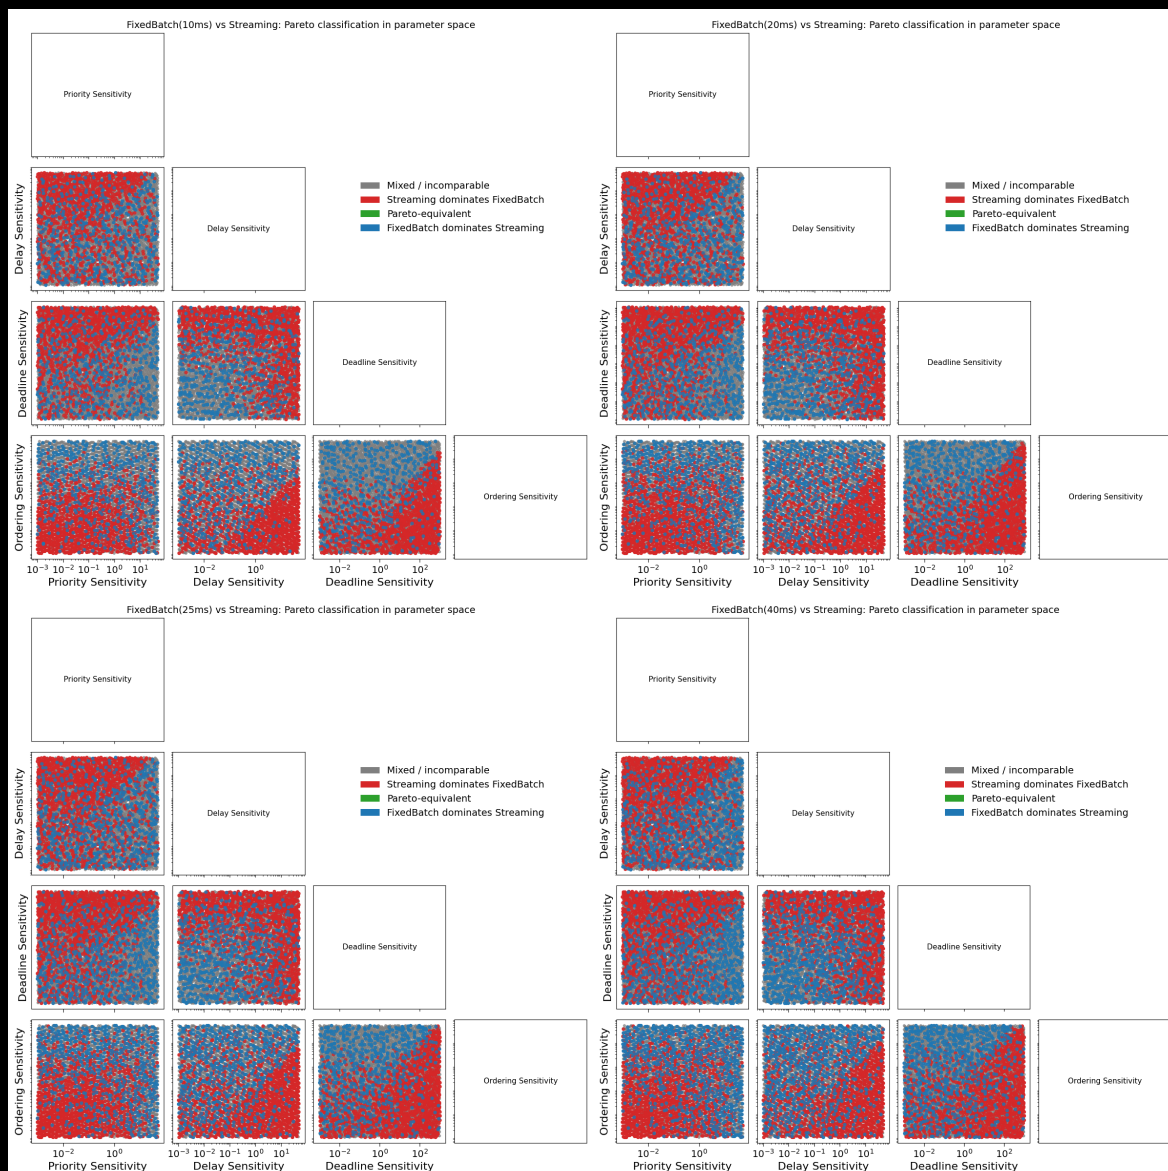
It's clear then that FixedBatch departs from Streaming along two margins: it reorders common executions and it changes the executed set itself. Both effects increase smoothly with the batch window.

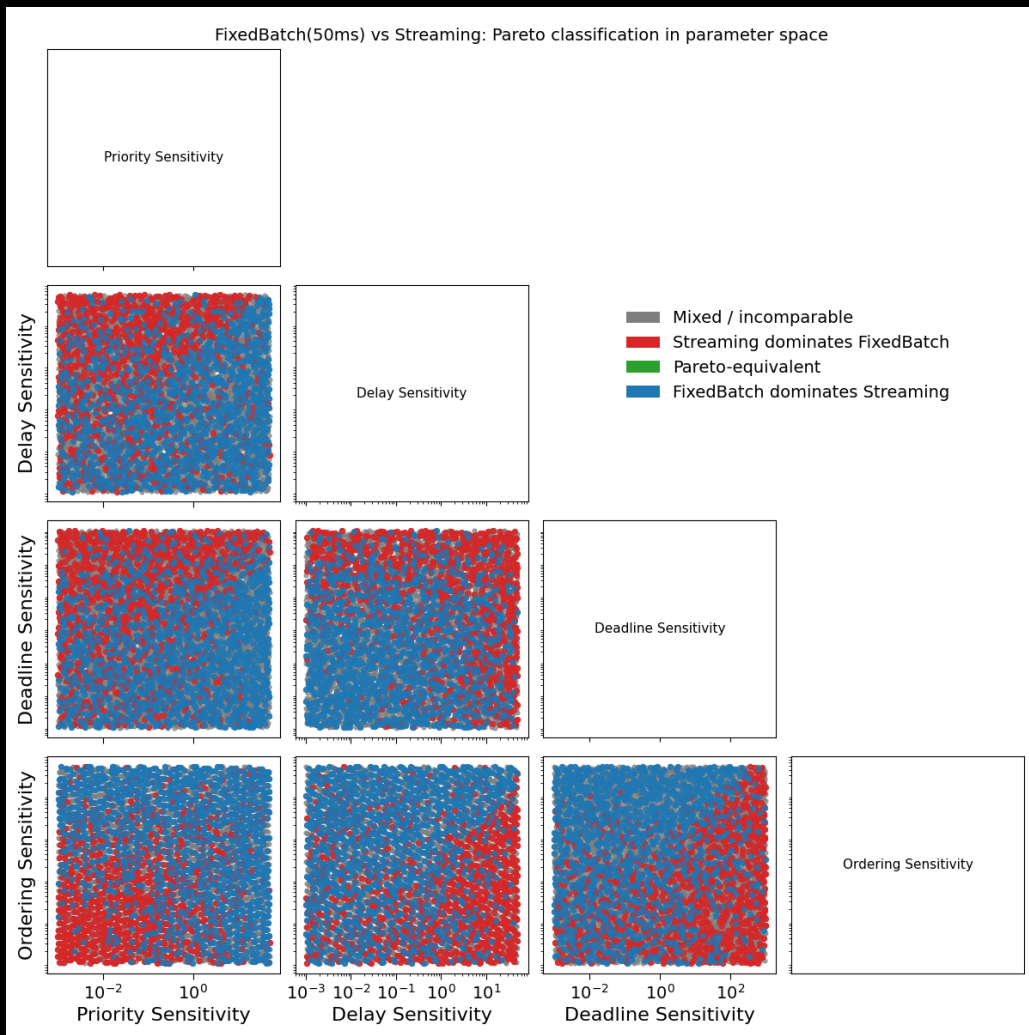
For FixedBatch, the increase in executed-set turnover J should be interpreted with care. In our implementation, transactions arriving in the final batch window are released only at the slot boundary T , and therefore cannot complete execution within the same slot. This means that part of the increase in J is mechanically induced by the treatment of the last batch, rather than by an intrinsic change in scheduler selection throughout the slot.

Pairwise Pareto Dominance

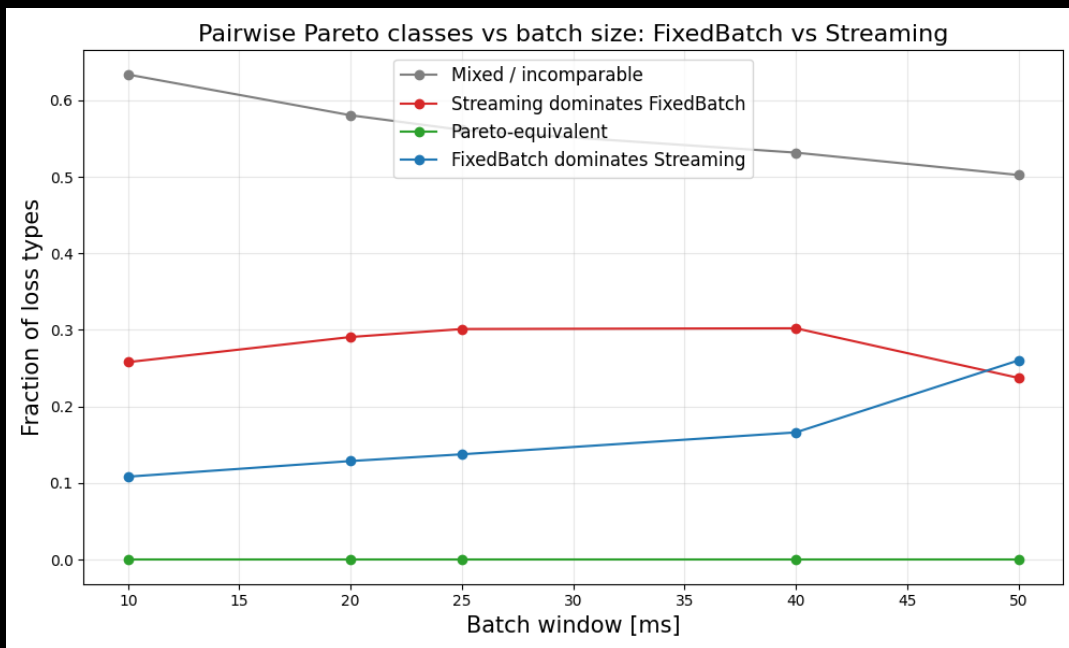
Here we compare each sampled profile for FB(w) against Streaming using four classes for comparison:

- FixedBatch dominates Streaming
- Streaming dominates FixedBatch
- Pareto-equivalent
- Mixed / incomparable





The idea is to understand for each agent type when the we have that FB largely Pareto dominates.



Across the tested grid, Mixed / incomparable remains the largest class. At 10 ms, it is roughly 63% of sampled profiles; by 50 ms, it falls to about 50%. So the main effect of increasing the batch size is not an abrupt switch from one winner to another. It is a gradual conversion of ambiguous cases into more decisive ones.

At the same time, the balance between the two decisive classes moves steadily. For small windows, the share of profiles in which Streaming dominates FixedBatch is larger than the share in which FixedBatch dominates Streaming. That remains true at 10 ms, 20 ms, 25 ms, and 40 ms. Only at 50 ms does the ranking flip: the FixedBatch-dominant share rises to about 26%, while the Streaming-dominant share falls to about 24%.

That is the central result.

It does not mean that FixedBatch wins universally at 50 ms. Half the profiles remain mixed. But it does mean that

Among the tested values, 50 ms is the first batch size at which the pairwise balance turns favourable to FixedBatch

This is the precise sense in which sub-50 ms windows look weak.

Critical-Region Estimator

A natural question is if a phase transition occurs when moving the batch size. To study it we can define

$$\rho = \log_{10} \left(\frac{\kappa}{\beta + \gamma_{\text{deadline}}} \right).$$

Here

- κ is the penalty for losing the ordering race
- β is the marginal cost of delay
- γ_{deadline} is the marginal penalty for delay beyond the deadline.

The ratio

$$\frac{\kappa}{\beta + \gamma_{\text{deadline}}}$$

has units of time. It is the delay-equivalent value of avoiding one ordering loss. In plain words: how many seconds of effective delay a user would trade away to avoid losing the race.

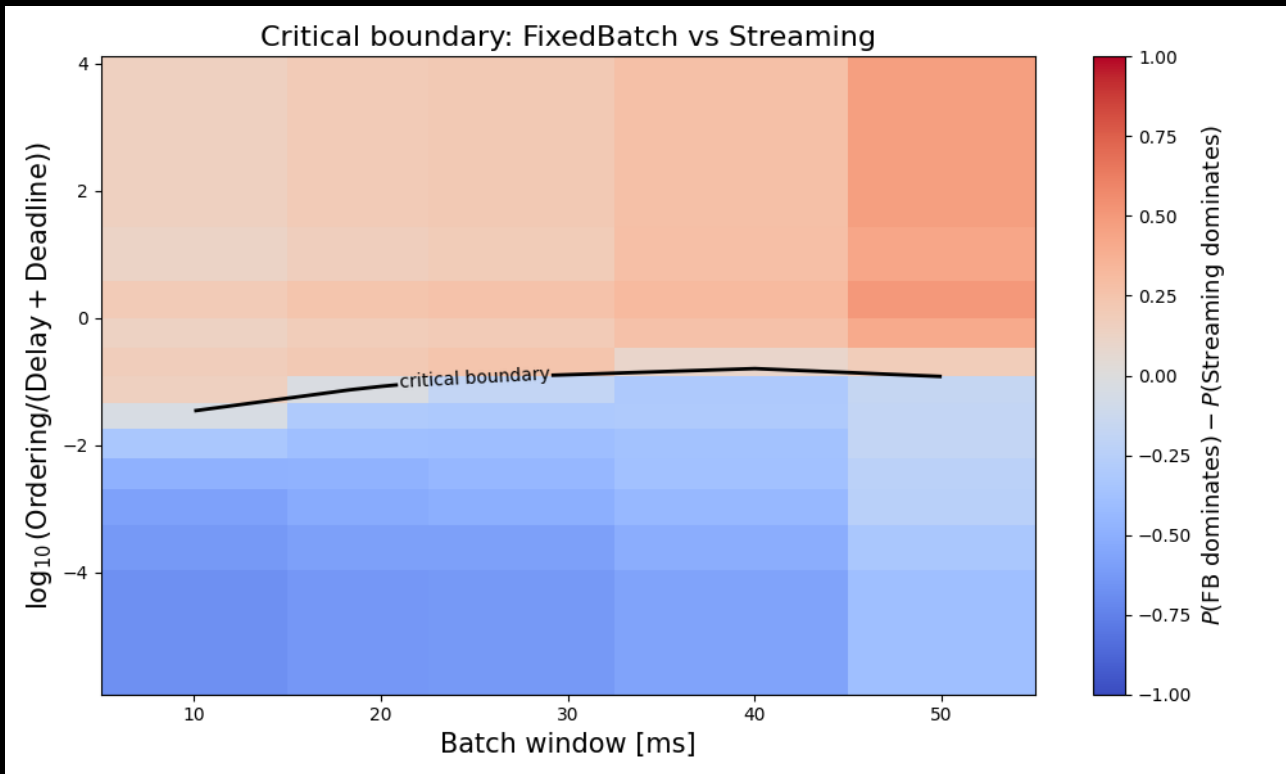
For example, $\rho \sim -1$ means

one avoided ordering loss is worth about 100 ms of effective delay.

At this point we can define a critical-region estimator built from the signed pairwise advantage

$$S(w, \rho) = \Pr(\text{FB dominates Streaming} \mid w, \rho) - \Pr(\text{Streaming dominates FB} \mid w, \rho),$$

so that $S = 0$ defines a critical boundary: above it, the pairwise balance favours FixedBatch; below it, it favours Streaming.



The plots do not show a genuine phase transition in the strong sense. The critical-boundary plot shows that the decisive control variable is not the batch size alone. It is the interaction between w and ρ .

For low ρ , delay and deadline costs dominate, so Streaming remains preferable. For high ρ , ordering quality matters enough that FixedBatch becomes attractive. Increasing w shifts the balance toward FixedBatch, but the shift is gradual rather than discontinuous.

Conclusions

The results are clear.

As the batch window grows, FixedBatch moves progressively away from Streaming, both by reordering common executions and by changing the executed set. This shift is gradual rather than abrupt: we do not observe a true phase transition, but a critical boundary.

In the tested grid, 50 ms is the first batch size for which FixedBatch pairwise dominates Streaming on a larger share of loss profiles than the reverse.

