A Supplementary Parameter-Sensitivity Analysis

To improve parameter-transparency, we added a compact sensitivity analysis under the controlled proof-of-concept validation pipeline. These additional results are intended as a robustness check of the deterministic operator under the controlled evaluation setup and should not be interpreted as a claim of globally optimal subject-level parameters or clinical generalization.

Sensitivity to threshold factor α

Table 2 summarizes the effect of varying the threshold multiplier α on controlled false-alarm rate and detection latency.

Table 2: Controlled sensitivity analysis for threshold factor α under the proof-of-concept validation pipeline.

α	False alarms (alarms/min)	Drift latency (s)	Pause latency (s)
1.5	0.000	0.000	0.400
2.0	0.000	0.000	0.560
2.5	0.000	0.400	0.720
3.0	0.000	29.080	0.860

Across $\alpha \in [1.5, 3.0]$, the default value $\alpha = 2.0$ preserved zero false alarms while retaining fast detection across the controlled drift and pause regimes. At $\alpha = 3.0$, drift detection became substantially slower, indicating an overly conservative threshold for gradual frequency-change events under the present controlled conditions.

Sensitivity to memory horizon M

Table 3 summarizes the effect of varying the memory window M.

Table 3: Controlled sensitivity analysis for memory window M under the proof-of-concept validation pipeline.

M (samples)	Approx. duration	Drift latency (s)	Pause latency (s)
50	≈ 1 s	29.000	1.000
100	≈ 2 s	0.480	0.580
150	≈ 3 s	0.000	0.560
200	≈ 4 s	0.000	0.560
300	≈ 6 s	0.000	0.560

Short memory windows reduced sensitivity to gradual drift, consistent with the interpretation that overly rapid adaptation of the rolling phase-memory estimate suppresses phase–memory divergence. In contrast, $M \geq 150$ samples preserved immediate or near-immediate detection of controlled drift while maintaining stable pause detection. Under the present controlled pipeline, the default choice $M = 150$ therefore provided a reasonable balance between responsiveness and memory retention.

A.1 Per-record controlled detection latency

To improve transparency across the $N = 5$ BIDMC recordings used in the present proof-of-concept evaluation, Table 4 reports the per-record controlled detection latencies for the drift and pause perturbation regimes, together with false-alarm counts in the stable control condition.

Table 4: Per-record controlled detection latency across the five BIDMC recordings used in the proof-of-concept evaluation.

Record ID	Drift latency (s)	Pause latency (s)	False alarms
bidmc01	0.04	0.58	0
bidmc02	0.02	0.56	0
bidmc03	0.00	0.56	0
bidmc04	0.18	0.62	0
bidmc05	0.00	0.54	0

These per-record values show that the controlled pause-detection latency remained tightly clustered across recordings, while drift-detection latency showed modest between-record variation, with the largest delay observed in `bidmc04`. No false alarms were observed in the stable control condition for any of the five evaluated recordings.

A.2 Dataset context for the evaluated BIDMC recordings

Table 5 summarizes the limited dataset-level context available for the five BIDMC respiratory recordings used in the present proof-of-concept evaluation. The PhysioNet BIDMC respiratory dataset does not provide per-record demographic metadata such as age, sex, BMI, or clinical status. Recording duration is approximately 8 min per record according to the dataset documentation.

Table 5: Available dataset context for the five BIDMC recordings used in the proof-of-concept evaluation.

Record ID	Age	Sex	BMI	Clinical status	Recording duration
bidmc01	not provided	not provided	not provided	not provided	$\sim$8 min
bidmc02	not provided	not provided	not provided	not provided	$\sim$8 min
bidmc03	not provided	not provided	not provided	not provided	$\sim$8 min
bidmc04	not provided	not provided	not provided	not provided	$\sim$8 min
bidmc05	not provided	not provided	not provided	not provided	$\sim$8 min