Scientific Research Publishing

Communications and Network, 2026, 18(4), 59-76
https://www.scirp.org/journal/cn
ISSN Online: 1947-3826
ISSN Print: 1949-2421

# Candidate-Specific Uncertainty-Aware Predictive Handover for Dense Mobile Networks

**Utham Kumar Anugula Sethupathy*, Vijayanand Ananthanarayanan**

Independent Researcher, Atlanta, USA
Email: *mailuthamkumar@gmail.com, vijayanand31@gmail.com

**How to cite this paper:** Sethupathy, U.K.A. and Ananthanarayanan, V. (2026) Candidate-Specific Uncertainty-Aware Predictive Handover for Dense Mobile Networks. *Communications and Network*, **18**, 59-76.
https://doi.org/10.4236/cn.2026.184005

**Received:** September 10, 2026
**Accepted:** October 31, 2026
**Published:** November 3, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc.
This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
http://creativecommons.org/licenses/by/4.0/

Open Access

## Abstract

Predictive handover can react early to mobility but may become unstable when forecast errors differ across candidate cells. This study presents uncertainty-aware predictive handover (UAPHO), which estimates candidate-specific uncertainty from matured 0.5-s one-step forecast residuals while using a separate 1.5-s forecast for handover decisions. Twenty paired seeds were evaluated for three user densities, three speed ranges, eight policies, and five uncertainty coefficients, producing 1440 main and 100 sensitivity trials. At 80 users and 70 - 100 km/h, full UAPHO had a 1.08% handover-failure rate and 0.00% ping-pong rate, compared with 2.03% failure for the no-uncertainty predictive variant and 3.48% ping-pong without the return guard. Conventional A3 and load-aware A3 achieved lower failure rates of 0.04% and 0.00%, respectively, and tenth-percentile throughputs of 1.98 and 2.08 Mbit/s, versus 1.51 Mbit/s for full UAPHO. UAPHO therefore controls predictive instability but trades lower-tail throughput for stability in this synthetic corridor. Its value depends on uncertainty calibration and the selected guardrails.

## Keywords

5G Mobility, Predictive Handover, Forecast Uncertainty, Quality of Service, Sensitivity Analysis

## 1. Introduction

NR mobility management combines measurement configuration, handover control, and execution procedures [1]. Radio-resource control specifies the reporting and control messages that carry those decisions [2]. A predictive policy must

therefore distinguish the information available at decision time from radio outcomes observed only after execution.

A scalar uncertainty penalty shared by every candidate cannot influence target ranking or a target-serving comparison. UAPHO instead estimates a distinct uncertainty state for each user-cell pair, evaluates cell-specific margins, recomputes occupancy after executed handovers, and scores failure over a short post-handover interval.

The study asks whether cell-specific forecast residual uncertainty can stabilize short-horizon mobility decisions and what portion of the result comes from the return-cell, load, and viability guards. It does not assume future channel values. Figure 1 shows the measurement, forecast, decision, and execution sequence.

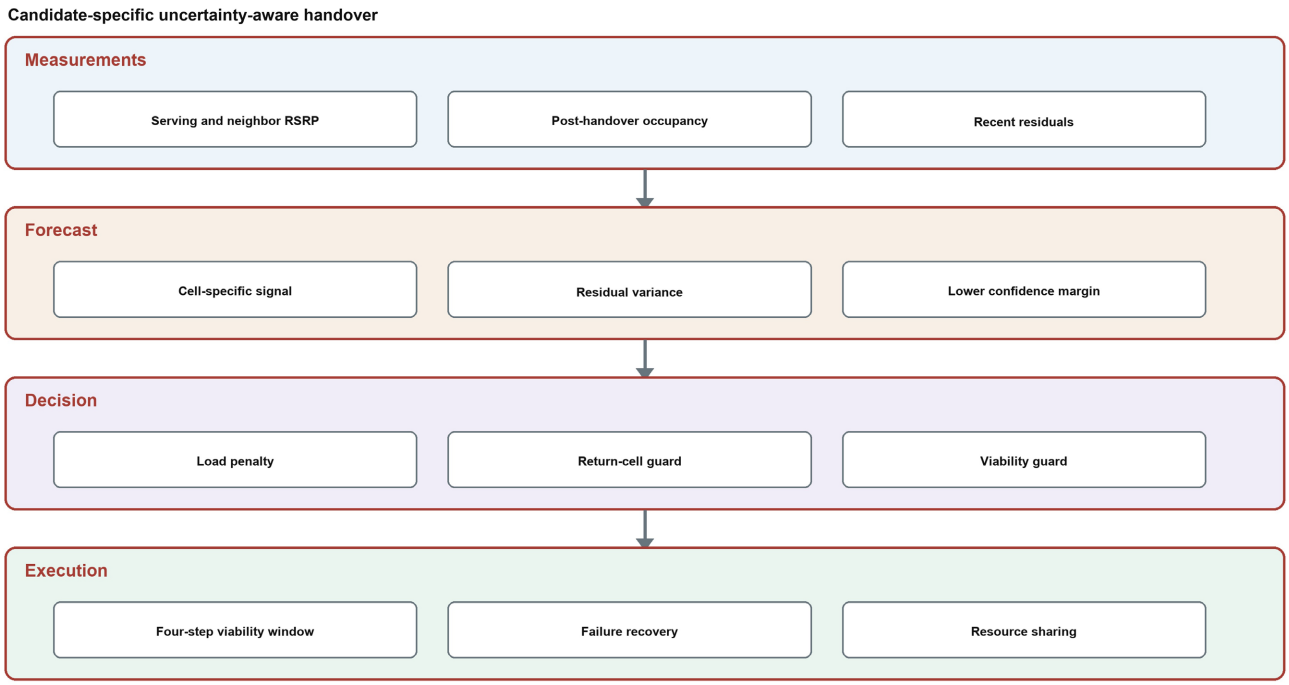



**Figure 1.** UAPHO pipeline with candidate-specific residual uncertainty, guarded decisions, and post-handover execution assessment.

## 2. Related Work

The channel model follows the distance-loss and correlated-shadowing concepts documented for 0.5 - 100 GHz evaluation [3]. Dense 5G systems motivate small cells, directional links, and frequent mobility events [4], while higher-frequency propagation increases sensitivity to blockage and geometry [5].

Recursive prediction methods formalize estimation from noisy observations [6]. UAPHO uses a transparent local trend rather than a trained model so that residual uncertainty and every decision term remain inspectable. Packet-level LTE simulation [7] and learned radio coordination [8] offer richer alternatives but would obscure the specific guardrail question addressed here.

Velocity and mobility-state inference can use handover histories [9]. IMT-2020

requirements identify mobility and user-experienced data rate as linked evaluation dimensions [10]. Heterogeneous-network studies [11] and media-independent handover standards [12] further motivate ping-pong and failure analysis. The NR system text [13] supplies broader architectural context. Simulation limits are reported explicitly in line with established cautions about synthetic network models [14].

Recent studies predict mobility-management demand from operational signaling [15], optimize handover jointly with load distribution [16], and reduce association churn through soft handover in cell-free systems [17]. Load-sensitive parameterization remains important in heterogeneous 5G networks [18], while multi-agent learning [19] and proportionally fair mobility control [20] illustrate broader stability-throughput objectives. Self-optimization of handover parameters [21] and predictive mechanisms for 5G and beyond [22] further motivate transparent forecast uncertainty and guardrail analysis.

## 3. Mobility and Radio Model

Figure 2 depicts the 6 km wraparound corridor with nine sites spaced 750 m apart. Users move at one of three speed ranges through a correlated shadowing field. Each trace is generated once per seed-density-speed scenario and replayed for every policy.

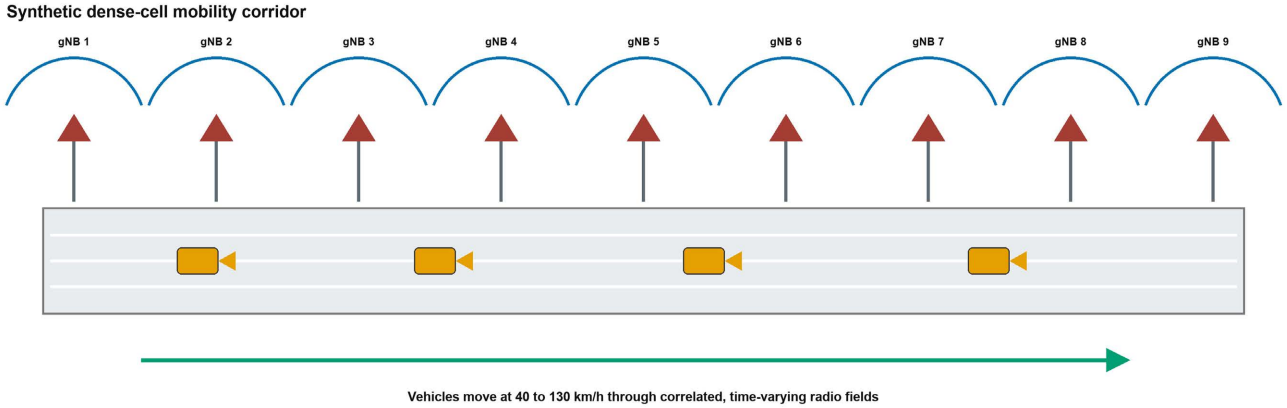



Figure 2. Synthetic nine-site mobility corridor used for all handover-policy comparisons.

Table 1 lists the radio and mobility assumptions. The model is intentionally controlled: it is designed to test decision logic, not to reproduce a specific operator deployment.

Table 1. Radio and mobility assumptions.

| Parameter | Value |
|---|---|
| Sites | 9, spaced 750 m |
| Corridor | 6 km with wraparound |
| Trace duration | 50 s |

**Continued**

| | |
|---|---|
| Decision interval | 0.5 s |
| Users | 40, 80, 120 |
| Speed ranges | 40 - 70, 70 - 100, 100 - 130 km/h |
| Shadowing | 4 dB correlated Gaussian |
| Bandwidth | 20 MHz per cell |
| RSRP model | $-43 - 30.5 \log_{10}(\max(10, d)/10) + X$ dBm |
| Shadowing SD | 4 dB |
| Shadowing temporal correlation | 0.94 per 0.5-s step |
| Candidate-cell set | All 9 cells |
| Interference abstraction | Strongest non-serving cell |
| SINR-like offset | 1.5 dB |
| Explicit carrier frequency | Not modeled |
| Explicit Tx power/antenna gains | Not modeled |
| Thermal noise/noise figure | Not modeled |

### 3.1. Radio Propagation and Measurement Model

The study uses a controlled received-power abstraction rather than a complete 3GPP link-budget model. Carrier frequency, absolute transmit power, antenna gains and heights, thermal-noise power, and receiver noise figure are not represented as separate parameters. Their aggregate effect is absorbed into the synthetic RSRP-distance relationship so that the study can isolate handover-decision behavior.

For user $i$ and cell $j$, the modeled RSRP is

$$P_{i,j}(t) = -43 - 30.5 \log_{10}\left(\frac{\max\left(10, d_{i,j}(t)\right)}{10}\right) + X_{i,j}(t) \text{ dBm}, \quad (1)$$

where $d_{\{i,j\}}(t)$ is the one-dimensional user-site distance in meters and $X_{\{i,j\}}(t)$ is the shadowing term. Distances below 10 m are clipped to 10 m.

Initial shadowing is sampled independently as

$$X_{i,j}(0) \sim \mathcal{N}\left(0, 4^2\right) \text{dB}.$$

Thereafter, shadowing follows the temporal first-order process

$$X_{i,j}(k) = 0.94 X_{i,j}(k-1) + \sqrt{1-0.94^2}\, Z_{i,j}(k),$$

with

$$Z_{i,j}(k) \sim \mathcal{N}\left(0, 4^2\right).$$

The model therefore uses temporal correlation of 0.94 per 0.5-s measurement interval rather than a fixed spatial correlation distance. Because user speeds vary, the implied traveled distance over one correlation step also varies.

All nine cells are measurable candidates for every user at each decision epoch. Interference is represented by the strongest non-serving-cell RSRP rather than by a full sum of received interference powers. A synthetic SINR-like margin is computed as

$$\Gamma_i(k) = P_{i,s_i}(k) - \max_{j \neq s_i} P_{i,j}(k) - 1.5\,\text{dB}.$$

Thus, this simulation is intended to compare handover control policies under a common radio trace; it is not a calibrated carrier-frequency-specific radio-access-network model.

### 3.2. Mobility Trace Generation

For each seed-density-speed scenario, the mobility trace is generated once and replayed unchanged across all eight handover policies. User initial positions are independently sampled from a uniform distribution over the 6-km wraparound corridor:

$$x_i(0) \sim U(0, 6000)\,\text{m}.$$

All users move in the positive corridor direction; travel direction is not independently randomized. Initial speed is sampled uniformly within the configured scenario range:

$$v_i(0) \sim U(v_{\min}, v_{\max}),$$

after conversion from km/h to m/s.

At every 0.5-s step, speed evolves according to

$$v_i(k) = \text{clip}\left(v_i(k-1) + \epsilon_i(k), v_{\min}, v_{\max}\right),$$

where

$$\epsilon_i(k) \sim \mathcal{N}\left(0, 0.32^2\right)\text{m/s}.$$

Position is then updated as

$$x_i(k) = \left[x_i(k-1) + v_i(k)\Delta t\right] \bmod 6000.$$

The initial serving cell is the cell with the highest modeled RSRP at the first recorded measurement:

$$s_i(0) = \arg\max_j P_{i,j}(0).$$

This pairing ensures that every policy evaluated for a common seed-density-speed scenario sees the same user positions, speeds, shadowing process, and RSRP trace.

## 4. Candidate Specific Forecast Uncertainty

### 4.1. Observable Forecast and Residual Variance

At each decision epoch $k$, separated by $\Delta t = 0.5\,\text{s}$, the controller computes the local RSRP trend for user $i$ and cell $j$ as

$$\dot{P}_{i,j}(k) = \frac{P_{i,j}(k) - P_{i,j}(k-1)}{\Delta t}.$$

Two forecasts are derived from the same observable trend. The handover decision uses a 1.5-s-ahead forecast:

$$\hat{P}_{i,j}^{(1.5)}(k) = P_{i,j}(k) + 1.5\dot{P}_{i,j}(k).$$

Uncertainty estimation uses a separate one-step forecast for the next 0.5-s measurement:

$$\hat{P}_{i,j}^{(0.5)}(k+1|k) = P_{i,j}(k) + \Delta t\,\dot{P}_{i,j}(k).$$

At epoch $k$, the matured one-step residual is therefore

$$e_{i,j}(k) = P_{i,j}(k) - \hat{P}_{i,j}^{(0.5)}(k|k-1).$$

The candidate-specific residual variance is updated as

$$\sigma_{i,j}^2(k) = 0.86\sigma_{i,j}^2(k-1) + 0.14e_{i,j}^2(k). \quad (2)$$

The initial residual variance is 9 $dB^2$. Thus, no 1.5-s prediction is evaluated prematurely: the 1.5-s forecast drives the handover decision, while uncertainty is estimated from matured 0.5-s forecast errors.

### 4.2. Lower Confidence Margin

Occupancy/load term:

$$R_{i,j}^{L}(k) = \ln\left(\frac{\tilde{n}_j(k)+1}{\tilde{n}_{s_i}(k)}\right)$$

where $\tilde{n}_j = n_j + 1$.

For serving cell s and candidate *j*, Equation (3) is evaluated with candidate-specific and serving-cell variances. The uncertainty term can no longer cancel because $\sigma_{i,j}$ varies by candidate and the combined margin changes the handover condition.

$$M_{i,j}^{\mathrm{LCB}}(k) = \left[\hat{P}_{i,j}(k+h) - \hat{P}_{i,s}(k+h)\right] - \kappa\sqrt{\sigma_{i,j}^2(k) + \sigma_{i,s}^2(k)} - \lambda\rho_j(k). \quad (3)$$

Full UAPHO uses $\kappa = 1$ and $\lambda = 2.8$. A candidate becomes eligible when $M_{i,j}^{\mathrm{LCB}} \geq 3$ dB for two consecutive decision epochs.

### 4.3. Guardrails and Execution

Handover persistence is evaluated at the 0.5-s decision interval. When a candidate first satisfies the handover margin, the persistence timer begins. The same candidate must continue satisfying the margin for two consecutive decision epochs, corresponding to 1.0 s, before handover execution is initiated.

Once a handover is executed at epoch $k$, the post-handover viability window begins with the first subsequent measurement and spans the next four 0.5-s measurement intervals, *i.e.*, epochs $k+1$ through $k+4$. Failure is recorded if the tar-

get violates either the absolute-power or relative-margin criterion at any point within this 2.0-s execution-assessment window.

$$P_{i,s_i}(k) < -112\,\text{dBm}$$

or

$$P_{i,s_i}(k) - \max_{j \neq s_i} P_{i,j}(k) < -11\,\text{dB}.$$

Cell occupancy is recomputed after all handovers and failures in the interval. Equation (4) then shares cell bandwidth among the updated occupants.

$$\begin{aligned} \Gamma_i(k) &= P_{i,s_i}(k) - \max_{j \neq s_i} P_{i,j}(k) - 1.5 \\ \eta_i(k) &= \text{clip}\left[\log_2\left(1 + 10^{\Gamma_i(k)/10}\right), 0, 6\right] \\ R_i(k) &= \frac{B}{N_{s_i}(k)} \log_2\left(1 + \text{SINR}_i(k)\right) \end{aligned} \tag{4}$$

where $B = 20$ MHz is the per-cell bandwidth and $N_{\{s_i\}}(k)$ is the post-handover number of users attached to user $i$'s serving cell.

A user that executes a handover during the interval receives 90% of the calculated rate, representing a 50-ms interruption during the 0.5-s service interval.

Table 2 defines the baselines and ablations. Forecast-only intentionally omits every guard so that prediction instability is visible.

Table 2. Policies and ablations.

| Policy | Uncertainty | Return guard | Load | Viability |
|---|---|---|---|---|
| A3 baseline | No | No | No | Execution only |
| Load-aware A3 | No | No | Yes | Execution only |
| Full UAPHO | Cell specific | Yes | Yes | Yes |
| No uncertainty term | $\kappa = 0$ | Yes | Yes | Yes |
| No return guard | Cell specific | No | Yes | Yes |
| No load penalty | Cell specific | Yes | No | Yes |
| No viability guard | Cell specific | Yes | Yes | No |
| Forecast only | No | No | No | No |

These values are held fixed across all reported seeds, densities, and speed ranges. The uncertainty coefficient $\kappa$ is additionally evaluated from 0 to 2 in the predefined sensitivity experiment.

| Parameter | Value | Rationale |
|---|---|---|
| Decision interval | 0.5 s | Common measurement/control interval |

**Continued**

| | | |
|---|---|---|
| Decision forecast | 1.5 s | Anticipates conditions three control intervals ahead |
| Residual forecast | 0.5 s | Allows residual to mature at the next measurement |
| Residual EWMA $\beta$ | 0.86 | Smooths short-term forecast-error variation |
| Initial variance | 9 dB$^2$ | Initial 3-dB residual standard deviation |
| $\kappa$ | 1.0 | Central value of predefined 0–2 sweep |
| $\lambda$ | 2.8 | Moderate load penalty |
| UAPHO decision threshold | 2.8 dB | Requires material lower-confidence advantage |
| A3 threshold | 3.0 dB | Fixed baseline comparison threshold |
| TTT | 1.0 s | Two qualifying 0.5-s intervals |
| Return guard | 8 s | Longer than 5-s ping-pong window |
| Viability threshold | −107 dBm | Rejects predicted weak target |
| Failure power threshold | −112 dBm | Severe post-HO degradation |
| Failure relative threshold | −11 dB | Target collapse vs strongest competitor |
| HO interruption | 50 ms | Successful-handover service loss |
| Failure interruption | 180 ms | Failed-HO recovery penalty |

### 4.4. Baseline Decision Rules

The A3 baseline uses current measured RSRP. At each decision epoch, the target is the cell with the highest instantaneous RSRP:

$$j^* = \arg\max_j P_{i,j}(k).$$

The A3 advantage is

$$A_i^{A3}(k) = P_{i,j^*}(k) - P_{i,s_i}(k).$$

A handover is initiated when the target differs from the serving cell and $A_i^{A3}(k) > 3.0$ dB for two consecutive 0.5-s decision intervals, corresponding to a 1.0-s time-to-trigger.

Load-aware A3 uses the same 3.0-dB threshold and 1.0-s time-to-trigger but adjusts target ranking using cell occupancy. Let $n_j(k)$ denote the number of users currently attached to cell $j$, and define the pseudocount-adjusted occupancy $\tilde{n}_j(k) = n_j(k) + 1$. The load term for candidate $j$ relative to serving cell $s_i$ is

$$R_{i,j}^L(k) = \ln\left(\frac{\tilde{n}_j(k)+1}{\tilde{n}_{s_i}(k)}\right).$$

The load-adjusted score is

$$S_{i,j}^L(k) = P_{i,j}(k) - 3.0R_{i,j}^L(k),$$

and the target is $j^* = \arg\max_j S_{i,j}^L(k)$. Handover occurs when the target's load-adjusted advantage over the serving cell exceeds 3.0 dB for 1.0 s.

## 5. Experimental Design and Statistics

Table 3 records the executed factorial design. Each policy sees the same RSRP trace within a seed-density-speed scenario. The sensitivity experiment reruns the principal scenario at $\kappa$ values from 0 to 2.

Table 3. Executed design.

| Factor | Values |
|---|---|
| Seeds | 20, numbered 2001-2020 |
| Densities | 40, 80, 120 users |
| Speed ranges | 40 - 70, 70 - 100, 100 - 130 km/h |
| Policies | 8 |
| Main records | 1440 |
| Kappa sensitivity | 5 values × 20 seeds = 100 records |
| Primary outcomes | Failure, ping-pong, outage, P10 throughput |

Equation (5) gives paired 95% confidence intervals. All means and standard deviations are calculated across fixed seeds.

$$CI_{95\%} = \bar{\Delta} \pm t_{0.975,19} \frac{s_\Delta}{\sqrt{20}}. \tag{5}$$

Differences are formed seed by seed before the mean and standard error are calculated.

## 6. Results

### 6.1. Principal Scenario

Figure 3 reports post-handover failures at 80 users and 70 - 100 km/h. Full UAPHO failed 1.08% of handovers (SD 0.82); no uncertainty failed 2.03% (SD 1.42). The paired full-minus-no-uncertainty difference was −0.95 percentage points (95% CI −1.60 to −0.29).

Figure 4 isolates ping-pong behavior. Full UAPHO recorded 0.00% and the no-return-guard ablation recorded 3.48%. The paired difference was −3.48 points (95% CI −4.30 to −2.67). This comparison, rather than the full-policy zero alone, quantifies the return guard's effect.

Figure 5 shows lower-tail service. Full UAPHO produced 1.51 Mbit/s (SD 0.17); A3 produced 1.98 Mbit/s (SD 0.15). The paired difference was −0.46 Mbit/s (95% CI −0.56 to −0.36).

Handover-failure rate is calculated as

$$R_{\text{fail}} = 100 \frac{N_{\text{failed handovers}}}{N_{\text{executed handovers}}}.$$

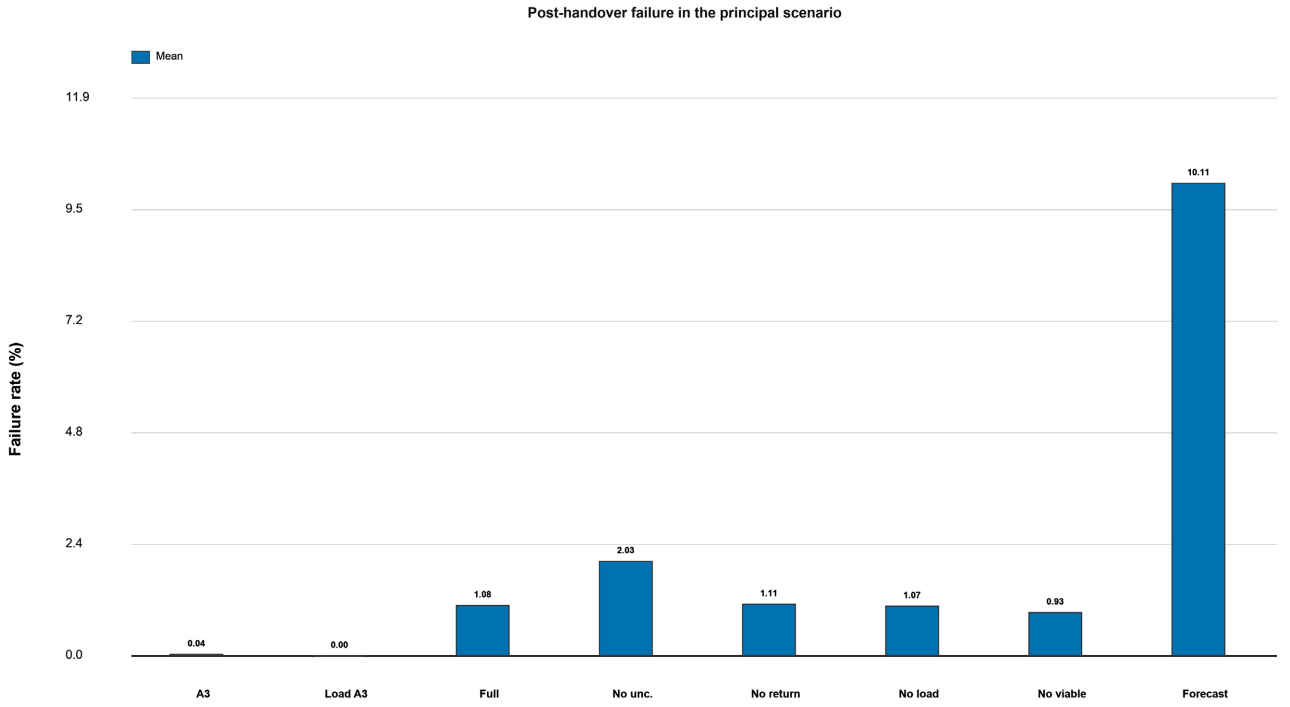



**Figure 3.** Mean post-handover failure rate in the principal scenario.

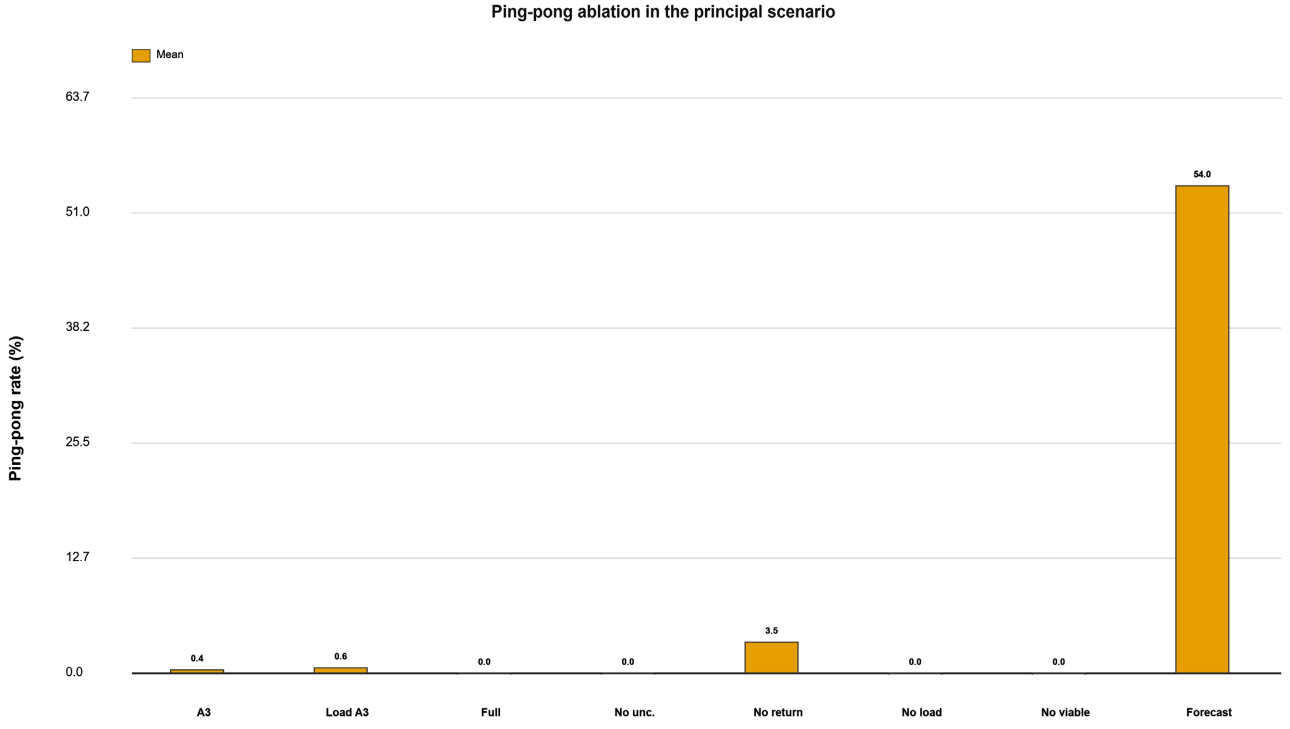



**Figure 4.** Ping-pong rate with and without the return-cell and other guardrails.

Ping-pong rate is

$$R_{\text{pp}} = 100 \frac{N_{\text{return handovers within 5s}}}{N_{\text{executed handovers}}}.$$

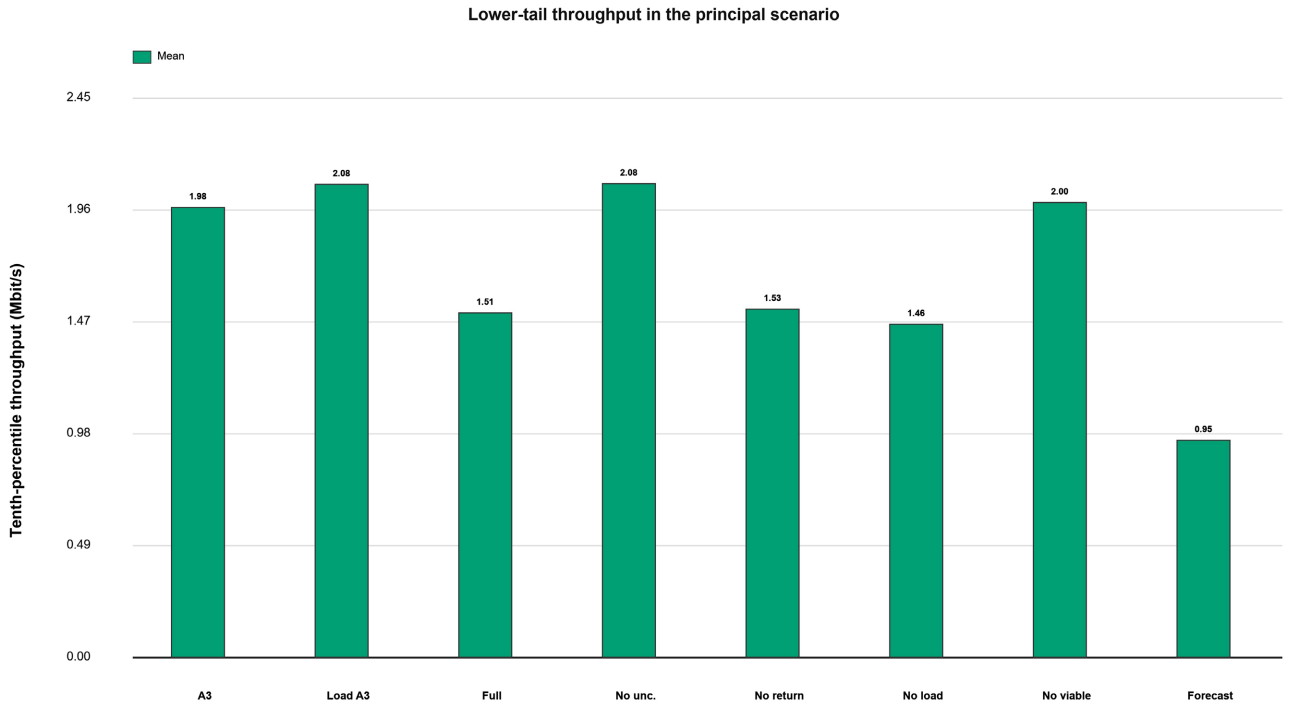



**Figure 5.** Tenth-percentile throughput after recomputing post-handover occupancy.

For each seed, handover-failure rate is

$$HFR = 100\frac{N_{\text{failed}}}{N_{\text{handover}}}.$$

Ping-pong rate is

$$PPR = 100\frac{N_{\text{return within 5s}}}{N_{\text{handover}}}.$$

Outage rate is computed over all user-decision intervals. An interval is classified as outage when the serving-cell RSRP is below −110 dBm or when the SINR-like margin $\Gamma_i(k)$ is below −7 dB:

$$O = 100\frac{N_{\text{outage user intervals}}}{N_{\text{users}}\left(N_{\text{steps}}-1\right)}.$$

The reported tenth-percentile throughput is the 10th percentile of all user-interval throughput samples accumulated over the trace.

Table 4 consolidates failure rate, ping-pong rate, outage rate, and tenth-percentile throughput.

**Table 4.** Principal scenario results.

| Policy | Mean HOs ± SD | Total HOs | Failed HOs | Ping-pong events | Failure mean ± SD (%) | Ping-pong (%) | Outage (%) | P10 Mbit/s |
|---|---|---|---|---|---|---|---|---|
| A3 baseline | 136.4 ± 5.1 | 2728 | 1 | 12 | 0.04 ± 0.17 | 0.43 | 0.49 | 1.98 |
| Load-aware A3 | 137.7 ± 4.6 | 2754 | 0 | 18 | 0.00 ± 0.00 | 0.64 | 0.56 | 2.08 |

**Continued**

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Full UAPHO | 133.9 ± 3.8 | 2677 | 29 | 0 | 1.08 ± 0.82 | 0.00 | 4.85 | 1.51 |
| No uncertainty term | 153.1 ± 5.7 | 3061 | 62 | 0 | 2.03 ± 1.42 | 0.00 | 1.39 | 2.08 |
| No return guard | 139.2 ± 4.8 | 2783 | 31 | 98 | 1.11 ± 0.82 | 3.48 | 4.76 | 1.53 |
| No load penalty | 132.1 ± 3.5 | 2642 | 28 | 0 | 1.07 ± 0.83 | 0.00 | 4.47 | 1.46 |
| No viability guard | 140.2 ± 4.6 | 2804 | 26 | 0 | 0.93 ± 0.70 | 0.00 | 1.72 | 2.00 |
| Forecast only | 479.0 ± 23.9 | 9580 | 969 | 5178 | 10.11 ± 0.87 | 53.98 | 5.57 | 0.95 |

Failure and ping-pong percentages are computed within each seed and then averaged across the 20 seeds; therefore, the displayed mean percentages need not exactly equal the ratios calculated from pooled event totals.

### 6.2. Stability-Throughput Trade-Off

Full UAPHO eliminated measured ping-pong in the principal scenario and reduced failure relative to the no-uncertainty predictive variant (1.08% versus 2.03%). That stability did not translate into across-the-board service gains. A3 and load-aware A3 had lower failure rates (0.04% and 0.00%) and higher tenth-percentile throughput (1.98 and 2.08 Mbit/s) than full UAPHO (1.51 Mbit/s). The principal result is therefore a stability-throughput trade-off: uncertainty-aware guardrails suppress unstable predictive handovers but can retain a weaker serving cell long enough to reduce lower-tail service.

The no-return-guard result attributes the zero ping-pong rate to the explicit dwell constraint, while the full-versus-no-uncertainty comparison isolates the effect of the candidate-specific confidence term. Together, the ablations show that forecast uncertainty and dwell protection address predictive instability, but neither should be interpreted as a general replacement for conventional A3 in this model.

### 6.3. Speed and Density Sensitivity

Figure 6 compares failure rate across speed ranges. Figure 7 reports full-UAPHO lower-tail throughput for every density-speed cell. These plots use all 20 seeds per cell.

### 6.4. Uncertainty Coefficient Sensitivity

Figure 8 shows that $\kappa$ changes failure, ping-pong, and lower-tail throughput. A larger uncertainty penalty suppresses marginal handovers but can retain a degrading serving cell too long. Table 5 provides the executed means.

Table 5. Kappa sensitivity.

| Kappa | Failure (%) | Ping-pong (%) | P10 Mbit/s | Outage (%) |
|---|---|---|---|---|
| 0.0 | 2.03 | 0.00 | 2.08 | 1.39 |

**Continued**

| | | | | |
|---|---|---|---|---|
| 0.5 | 1.58 | 0.00 | 1.80 | 3.47 |
| 1.0 | 1.08 | 0.00 | 1.51 | 4.85 |
| 1.5 | 0.58 | 0.00 | 1.25 | 5.68 |
| 2.0 | 0.52 | 0.00 | 0.97 | 6.60 |

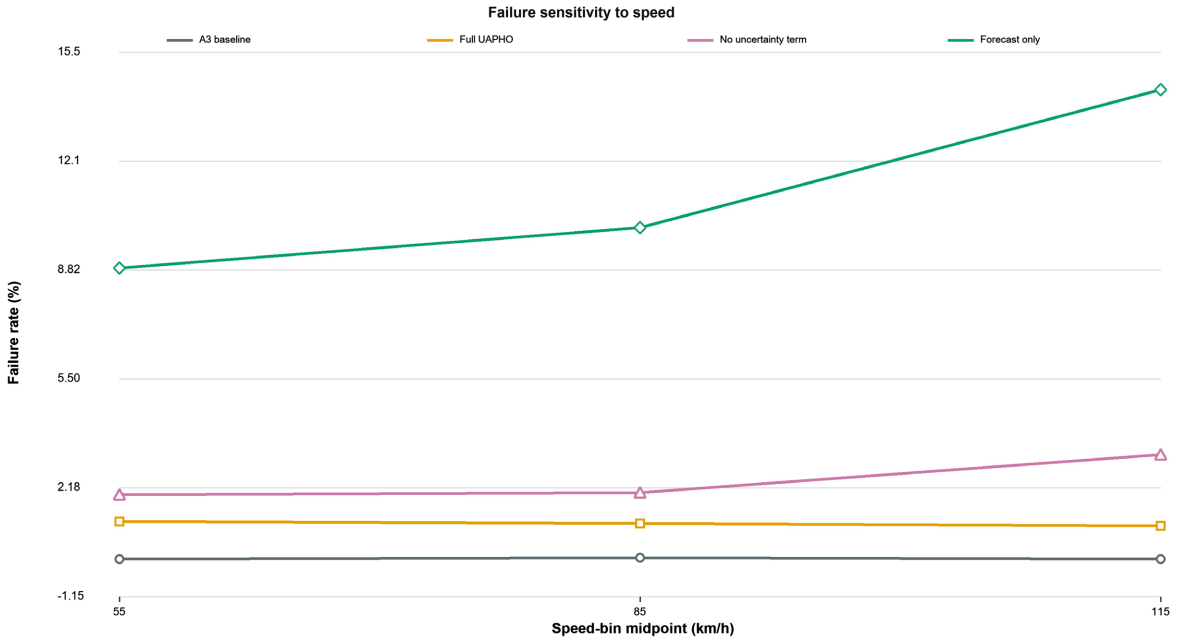



**Figure 6.** Handover-failure sensitivity across the three vehicle-speed ranges at 80 users.

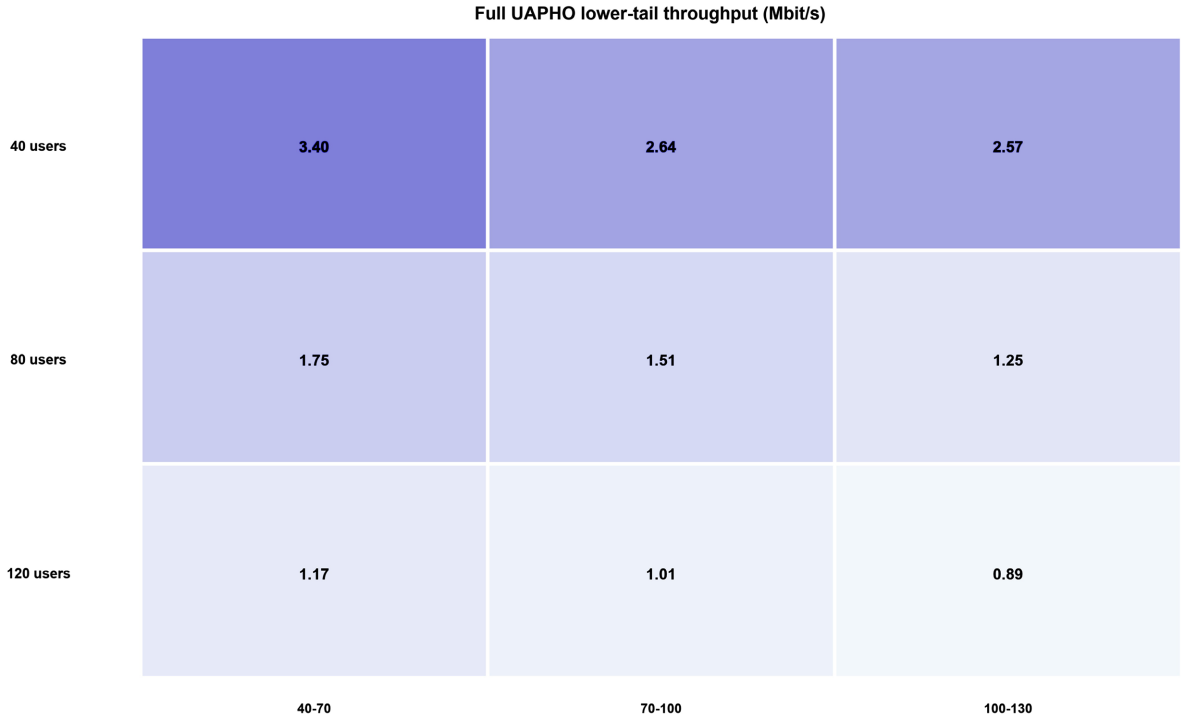



**Figure 7.** Full-UAPHO tenth-percentile throughput across all density and speed combinations.

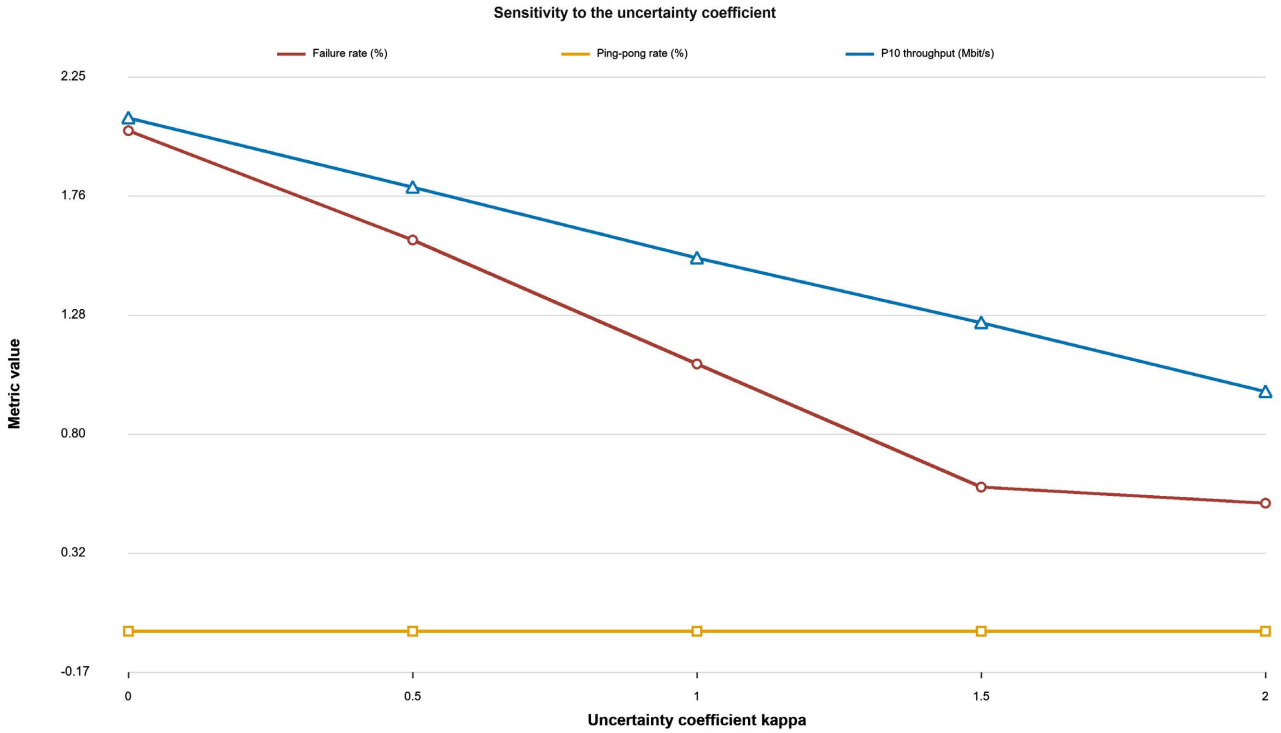



**Figure 8.** Sensitivity of the principal scenario to the candidate-specific uncertainty coefficient.

## 6.5. Resource Sharing after Handover

Figure 9 reports mean throughput calculated from occupancy recomputed after handover decisions, so each interval's resource sharing reflects the executed serving-cell assignments.

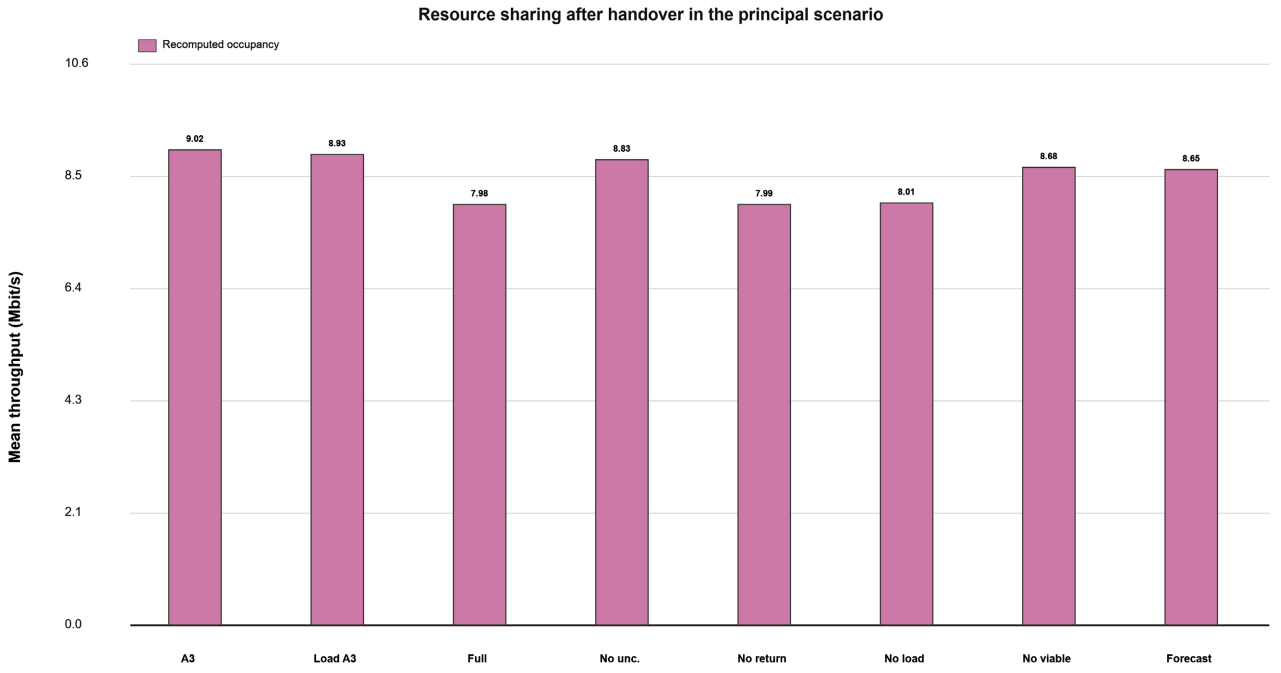



**Figure 9.** Mean throughput using post-handover cell occupancy in the principal scenario.

## 7. Discussion

Candidate-specific uncertainty affects both the margin and the executed decision set. The results do not justify a universal superiority claim. In this controlled corridor, $\kappa = 0$ can outperform $\kappa = 1$ on some throughput or failure metrics, indicating that uncertainty calibration matters as much as including the term.

The return-cell guard is the direct cause of the full policy's ping-pong suppression. The no-return ablation restores ping-pong while leaving the uncertainty estimator intact. The method's novelty therefore lies in the combined, auditable decision rule and the separation of each guard's contribution, not in attributing a mechanically enforced zero to prediction.

The viability guard is evaluated prospectively from present information, while failure is scored from a later execution interval. This temporal separation removes future leakage. Recomputed occupancy also exposes the throughput cost when many users select the same target.

## 8. Threats to Validity and Limitations

The corridor is one-dimensional, the channel equation is simplified, sites share equal bandwidth, and the model omits beam management, scheduler details, core-network delay, and signaling contention. Results are simulation outcomes rather than field measurements. The model intentionally abstracts explicit carrier frequency, transmit power, antenna characteristics, thermal noise, and summed inter-cell interference; therefore, absolute throughput values should be interpreted as internally consistent simulation metrics for policy comparison rather than operator-specific 5G performance predictions.

Residual variance is an EWMA of matured 0.5-s one-step forecast errors and is used as an engineering uncertainty proxy for the separate 1.5-s decision forecast. It is not a calibrated probabilistic confidence interval. The lower-confidence expression uses that variance as an engineering penalty. Trace-driven or packet-level evaluation is required before interpreting $\kappa$ as a confidence level.

Twenty seeds quantify within-model variation but do not cover all road layouts or fading processes. Failure thresholds and execution-window length are explicit sensitivity candidates for future work.

## 9. Conclusion

UAPHO uses cell-specific residual uncertainty, recomputes load after handovers, and evaluates target viability over a four-step execution window. In the principal scenario, full UAPHO recorded 1.08% failures, 0.00% ping-pong, and 1.51 Mbit/s tenth-percentile throughput. It reduced failure relative to the no-uncertainty predictive variant, and the no-return ablation demonstrates that the return guard accounts for the ping-pong result. Conventional A3 policies nevertheless achieved lower failure and better lower-tail throughput. The contribution is therefore control of prediction-induced instability under uncertain forecasts, with an explicit stability-throughput trade-off rather than across-the-board superiority.

## Data and Code Availability

The study uses no proprietary or personal data. The complete simulation source, fixed random-seed ranges, seed-level result files, aggregated tables, and publication figures are included in the accompanying reproducibility package. The code regenerates all numerical results and figures reported in this manuscript.

## Author Contributions

**Utham Kumar Anugula Sethupathy:** Conceptualization, methodology, software, validation, formal analysis, investigation, data curation, writing—original draft, writing—review and editing, and visualization. **Vijayanand Ananthanarayanan:** Validation, formal analysis, writing—review and editing.

## Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

## References

[1] 3rd Generation Partnership Project (2025) NR and NG-RAN Overall Description; Stage 2. 3GPP TS 38.300, Release 18. https://www.3gpp.org/dynareport/38300.htm

[2] 3rd Generation Partnership Project (2025) NR; Radio Resource Control Protocol Specification. 3GPP TS 38.331, Release 18. https://www.3gpp.org/dynareport/38331.htm

[3] 3rd Generation Partnership Project (2022) Study on Channel Model for Frequencies from 0.5 to 100 GHz. 3GPP TR 38.901, Version 17.0.0. https://www.3gpp.org/dynareport/38901.htm

[4] Andrews, J.G., Buzzi, S., Choi, W., Hanly, S.V., Lozano, A., Soong, A.C.K., *et al.* (2014) What Will 5G Be? *IEEE Journal on Selected Areas in Communications*, **32**, 1065-1082. https://doi.org/10.1109/jsac.2014.2328098

[5] Rappaport, T.S., Sun, S., Mayzus, R., Zhao, H., Azar, Y., Wang, K., *et al.* (2013) Millimeter Wave Mobile Communications for 5G Cellular: It Will Work! *IEEE Access*, **1**, 335-349. https://doi.org/10.1109/access.2013.2260813

[6] Kalman, R.E. (1960) A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, **82**, 35-45. https://doi.org/10.1115/1.3662552

[7] Baldo, N., Miozzo, M., Requena-Esteso, M. and Nin-Guerrero, J. (2011) An Open Source Product-Oriented LTE Network Simulator Based on NS-3. *Proceedings of the 14th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, Miami, 31 October-4 November 2011, 293-298. https://doi.org/10.1145/2068897.2068948

[8] Mismar, F.B. and Evans, B.L. (2019) Deep Learning in Downlink Coordinated Multipoint in New Radio Heterogeneous Networks. *IEEE Wireless Communications Letters*, **8**, 1040-1043. https://doi.org/10.1109/lwc.2019.2904686

[9] Merwaday, A. and Guvenc, I. (2016) Handover Count Based Velocity Estimation and Mobility State Detection in Dense HetNets. *IEEE Transactions on Wireless Communications*, **15**, 4673-4688. https://doi.org/10.1109/TWC.2016.2544302

[10] ITU-R (2017) Minimum Requirements Related to Technical Performance for IMT-2020 Radio Interface(s). Report ITU-R M.2410-0. https://www.itu.int/pub/R-REP-M.2410-2017

[11] 3rd Generation Partnership Project (2012) Mobility Enhancements in Heterogeneous Networks. 3GPP TR 36.839, Version 11.1.0. https://www.3gpp.org/dynareport/36839.htm

[12] IEEE (2017) IEEE Standard for Local and Metropolitan Area Networks Part 21: Media Independent Services Framework. IEEE 802.21-2017. https://standards.ieee.org/ieee/802.21/5608/

[13] Dahlman, E., Parkvall, S. and Skold, J. (2020) 5G NR: The Next Generation Wireless Access Technology. 2nd Edition, Academic Press.

[14] Paxson, V. and Floyd, S. (1997) Why We Don't Know How to Simulate the Internet. *Proceedings of the* 29*th Conference on Winter Simulation—WSC*'97, Atlanta, 7-10 December 1997, 1037-1044. https://doi.org/10.1145/268437.268737

[15] Makai, L.B. and Varga, P. (2023) Predicting Mobility Management Demands of Cellular Networks Based on User Behavior. *NOMS* 2023-2023 *IEEE/IFIP Network Operations and Management Symposium*, Miami, 8-12 May 2023, 1-6. https://doi.org/10.1109/noms56928.2023.10154215

[16] Kim, K., Lee, E., Bae, C. and Pack, S. (2024) Load-Aware Handover Optimization in Heterogeneous Networks: A Multi-Objective Learning Approach. 2024 *IEEE* 100*th Vehicular Technology Conference* (*VTC*2024-*Fall*), Washington DC, 7-10 October 2024, 1-5. https://doi.org/10.1109/vtc2024-fall63153.2024.10758043

[17] Zaher, M., Björnson, E. and Petrova, M. (2024) Soft Handover Procedures in mmWave Cell-Free Massive MIMO Networks. *IEEE Transactions on Wireless Communications*, **23**, 6124-6138. https://doi.org/10.1109/twc.2023.3330199

[18] Gures, E., Shayea, I., Saad, S.A., Ergen, M., El-Saleh, A.A., Ahmed, N.M.O.S., *et al.* (2023) Load Balancing in 5G Heterogeneous Networks Based on Automatic Weight Function. *ICT Express*, **9**, 1019-1025. https://doi.org/10.1016/j.icte.2023.03.008

[19] Alizadeh, A., Lim, B. and Vu, M. (2024) Multi-Agent Q-Learning for Real-Time Load Balancing User Association and Handover in Mobile Networks. *IEEE Transactions on Wireless Communications*, **23**, 9001-9015. https://doi.org/10.1109/twc.2024.3357702

[20] Prado, A., Stöckeler, F., Mehmeti, F., Krämer, P. and Kellerer, W. (2023) Enabling Proportionally-Fair Mobility Management with Reinforcement Learning in 5G Networks. *IEEE Journal on Selected Areas in Communications*, **41**, 1845-1858. https://doi.org/10.1109/jsac.2023.3273705

[21] Mbulwa, A.I., Yew, H.T., Chekima, A. and Dargham, J.A. (2024) Self-Optimization of Handover Control Parameters for 5G Wireless Networks and beyond. *IEEE Access*, **12**, 6117-6135. https://doi.org/10.1109/access.2023.3346039

[22] Sulaiman, T.H. and Al-Raweshidy, H.S. (2025) Predictive Handover Mechanism for Seamless Mobility in 5G and beyond Networks. *IET Communications*, **19**, e12878. https://doi.org/10.1049/cmu2.12878

## Appendix. Reproduction Procedure

Run the supplied Python file using the following command: python run_experiments.py --study cn. The script generates the nine scenario cells, eight-policy comparisons, five-level $\kappa$ sweep, figures, seed-level CSVs, and grouped summaries; Table A1 lists the expected checks.

Table A1. Reproduction checks.

| Check | Expected value |
|---|---|
| Seed range | 2001-2020 |
| Main records | 1440 |
| Sensitivity records | 100 |
| Occupancy timing | After handover execution |
| Future state used by decision | No |
| Decision interval | 0.5 s |
| Decision forecast horizon | 1.5 s |
| Residual forecast horizon | 0.5 s |
| Residual EWMA | $\beta = 0.86$ |
| A3 threshold | 3.0 dB |
| UAPHO threshold | 2.8 dB |
| TTT | 1.0 s |
| Return guard | 8 s |
| Ping-pong window | 5 s |
| Viability threshold | −107 dBm |
| Failure power threshold | −112 dBm |
| Failure relative margin | −11 dB |
| Post-HO assessment | 4 × 0.5 s |
| Successful HO interruption | 50 ms |
| Failed HO interruption | 180 ms |
| Candidate cells | All 9 |