Carrying the Sky: Causal Evidence for Umbrella-Mediated Precipitation Suppression Across 14 Metropolitan Centers

Authors: Elise M. Vanderhoof[1], Kenta Fujimori[2], Anja R. Petterssen[3], David Okonkwo-Lyle[1]

[1] Centre for Preparedness Meteorology, University of Edinburgh

[2] Department of Atmospheric Behavioral Science, Kyoto Institute of Pluvial Studies

[3] Nordic Laboratory for Commuter Climatology, University of Bergen

Corresponding author: e.vanderhoof@ed.ac.uk

Abstract

The relationship between personal rain-preparedness behavior and localized precipitation outcomes has remained largely unexamined since the foundational work of Hargreaves and Lunt (2011), despite anecdotal reports spanning centuries and across cultures. Recent advances in ambulatory meteorological sensing now permit rigorous investigation of what practitioners term the Umbrella Suppression Hypothesis (USH): namely, that the act of carrying an umbrella exerts a statistically significant inhibitory effect on rainfall within the carrier's immediate mesoscale environment. The present study asks whether umbrella carriage state (UCS) — operationalized as a binary preparedness variable — causally modulates precipitation probability at the hyper-local level, or whether the observed correlation reflects mere survivorship bias in commuter self-report data. Employing the Precipitation-Readiness Inversion Sampling Method (PRISM) across 14 metropolitan centers (N = 2,437 commuter-days), we observed a mean rainfall suppression index (δ_ρ) of 0.38 (95% CI [0.29, 0.47], $p < .001$), consistent with a moderate-to-strong preparedness-precipitation coupling. These findings held after controlling for seasonal confounds via the modified Galosh-Petterssen barometric adjustment. The implications for municipal drought-mitigation strategy — specifically, subsidized umbrella distribution programs — are discussed with reference to Kahneman and Tversky's (1979) prospect theory framework.

1 Introduction

It has long been observed, across cultures and centuries, that the deliberate act of carrying an umbrella appears to coincide with an absence of rain. Early references to this phenomenon surface in the meteorological diaries of 19th-century European naturalists (see Lindqvist, 1883, as cited in Fjeldstad & Gruber, 2004), though systematic inquiry remained dormant until well into the modern era. The formalization of the Umbrella Suppression Hypothesis (USH) is generally attributed to Hargreaves and Lunt (2011), whose three-year observational study of London commuters first documented a statistically reliable negative association between umbrella carriage and same-day rainfall events. Their work, published in *Atmospheric Behavioral Research*, catalyzed a small but increasingly well-funded research program dedicated to elucidating what they termed the "preparedness paradox" — the counterintuitive finding that readiness for precipitation appears to render precipitation less probable.

The theoretical landscape has since bifurcated into two competing accounts. The strong form of the USH posits a genuinely causal relationship: that the physical act of carrying an umbrella exerts a suppressive influence on hyper-local precipitation probability, possibly mediated by mechanisms at the boundary of atmospheric physics and behavioral intention (Kowalczyk & Sørensen, 2016). The weak form, by contrast, attributes the observed preparedness-precipitation coupling to selection artifacts, retrospective self-report bias, and what Nakamura (2019) has characterized as "dry-day mnemonic salience" — the tendency for commuters to more readily recall days on which their preparedness was vindicated by clear skies. Distinguishing between these accounts has proven methodologically intractable, owing largely to the field's historical reliance on unvalidated survey instruments. As Björklund, Taniguchi, and Weiß (2021, *Nordic Journal of Pluvial Cognition*, 12, 44–61) noted in their critical review, the umbrella-carriage literature prior to 2020 suffered from a near-total absence of objective meteorological measurement at the individual-carrier level, rendering causal inference "premature at best and epistemologically hazardous at worst."

The present study addresses this gap directly. By deploying the Precipitation-Readiness Inversion Sampling Method (PRISM) — a multi-site, randomized crossover protocol pairing participant-level umbrella assignment with body-mounted ambulatory hygrometry — we provide the first large-scale test of the USH under conditions sufficient to support causal claims. The study's design permits estimation of the rainfall suppression index (δ_ρ) with individual-level precision, while controlling for seasonal and geographic confounds via the Galosh-Petterssen barometric adjustment, a correction procedure adapted from standard synoptic climatology (Holton & Hakim, 2013) and validated for commuter-scale atmospheric sampling by Petterssen and Ng (2019).

The theoretical contribution of this work extends beyond pluviometric phenomenology. Following Ajzen's (1991) theory of planned behavior, we conceptualize umbrella carriage state (UCS) as a behavioral intention variable — one whose consequences, if the strong USH obtains, propagate from the cognitive-motivational domain into the physical atmosphere. Such a framing would position the USH within a broader class of observer-environment coupling phenomena, with implications not only for meteorological prediction but for the philosophy of dispositional causation more generally.

The primary research question may therefore be stated as follows: does umbrella carriage state (UCS), operationalized as a binary preparedness variable, causally modulate hyper-local precipitation probability, as indexed by the rainfall suppression index (δ_ρ), after adjustment for seasonal, geographic, and barometric confounds? In this paper, we report the results of a 14-city, 2,437-commuter-day field trial designed to resolve this question, present moderator and dose-response analyses bearing on the mechanism of the preparedness-precipitation coupling, and discuss implications for municipal drought-mitigation policy.

2 Methods

Participants and Design

A total of 2,437 commuter-days were sampled across 14 metropolitan centers (Amsterdam, Berlin, Buenos Aires, Dublin, Hong Kong, London, Melbourne, Mumbai, New York, Oslo, Phoenix, São Paulo, Seoul, and Vancouver) between March 2024 and February 2025. Participants ($N = 347$; 51.3% female, $M_{age} = 38.4$ years, $SD = 11.2$) were recruited via transit-authority partnerships and screened for eligibility using the Whitford-Tanaka Canopy Integrity Scale (Whitford & Tanaka, 1998, *Proceedings of the International Conference on Personal Meteorological Apparatus*, 3, 112–119), which rates umbrella structural condition on a 1–5 ordinal metric. Only participants possessing an umbrella rated ≥ 3 were admitted to the study; those carrying umbrellas with visible rib deformation or canopy micro-perforations exceeding 4 mm^2 in aggregate were excluded, as prior work has demonstrated that compromised canopy integrity introduces unacceptable noise into the suppression signal (Falk & Lindström, 2003). Each participant was enrolled in a randomized crossover design with three conditions administered across seven-day blocks: UCS-carried (umbrella physically carried throughout the commute), UCS-not-carried (umbrella left at home, confirmed via RFID tag verification at domicile), and UCS-null (umbrella owned but consciously and deliberately left behind, with participants completing a brief preparedness-intention affirmation prior to departure).

Instrumentation

Hyper-local precipitation was measured using body-mounted micro-hygrometers (Model PL-440, Vaisala Corp., Helsinki) affixed to the participant's left shoulder strap or lapel, recording relative humidity (φ), barometric pressure, and droplet-contact events at 30-second intervals. Following the protocol established by Ekström and Murakami (1993, *Advances in Ambulatory Atmospheric Sensing*, 7, 200–214), a personal precipitation envelope (PPE) was defined as the cylindrical volume extending 2 meters radially and 3 meters vertically from the participant's center of mass. Precipitation events were registered when the droplet-contact sensor recorded ≥ 5 discrete impacts within any 120-second window, a threshold validated against municipal tipping-bucket rain-gauge data (Pearson's $r = .91$, $p < .001$) during a six-week calibration phase.

Analytic Pipeline

The rainfall suppression index (δ_ρ) was computed for each commuter-day as the proportional reduction in precipitation probability relative to the site-matched meteorological baseline:

$$\delta_\rho = 1 - (\pi_{obs} / \pi_{exp}) \times \kappa_{GP}$$

where π_{obs} denotes the observed precipitation event rate within the PPE, π_{exp} is the expected rate derived from municipal gauge data for the same temporal window, and κ_{GP} is the Galosh-Petterssen barometric adjustment coefficient, a correction factor accounting for synoptic-scale pressure variation that would otherwise confound individual-level suppression estimates (Petterssen & Ng, 2019). The coefficient κ_{GP} was recalculated daily for each site using the procedure described in the original Galosh-Petterssen technical report and ranges from 0.88 to 1.14 in the present dataset.

Primary analyses employed mixed-effects logistic regression with participant as a random intercept and UCS condition, site, season, and day-of-week as fixed effects. The parameter α_u captured the main effect of umbrella carriage. Sensitivity analyses were conducted via bootstrapped permutation tests ($k = 10{,}000$ iterations) to confirm robustness to distributional assumptions, following the resampling framework advocated by Efron and Tibshirani (1993). Moderator analyses examined canopy diameter (continuous, in centimeters), umbrella color (categorical, five levels), and participant self-reported confidence in rain likelihood (λ_{conf}, measured on a 0–100 visual analog scale administered each morning).

The study protocol was reviewed and approved by the University of Edinburgh Joint Ethics Board for Atmospheric-Behavioral Research (Protocol No. 2024-USH-03), with a supplementary amendment mandating the provision of emergency ponchos to participants assigned to the UCS-not-carried condition on days when the national weather service issued precipitation warnings exceeding 70% probability.

3 Results

Primary Analysis

The omnibus mixed-effects model revealed a significant main effect of umbrella carriage state (UCS) on hyper-local precipitation probability ($\chi^2(2) = 41.73$, $p < .001$). Participants in the UCS-carried condition experienced a mean rainfall suppression index (δ_ρ) of 0.38 (95% CI [0.29, 0.47], $p < .001$), representing a 38% reduction in the probability of a precipitation event within the personal precipitation envelope relative to the site-matched meteorological baseline. This effect held after application of the Galosh-Petterssen barometric adjustment (κ_{GP} range: 0.91–1.12 across sites and seasons). The magnitude of δ_ρ is consistent with the range reported by Chakraborty, Strand, and Fleisch (2022) in their eight-city pilot but exceeds the more conservative estimate of $\delta_\rho = 0.21$ obtained by Kimura and Voss (2018, *Archives of Applied Pluviology*, 9, 155–168) using a non-randomized observational design. Table 1 summarizes the primary effect estimates by condition, including adjusted odds ratios and within-participant variance components.

The UCS-Null Condition

The exploratory UCS-null condition — in which participants owned an umbrella but consciously elected to leave it at home — yielded an intermediate δ_ρ of 0.14 (95% CI [0.03, 0.25], $p = .016$). A planned contrast confirmed that UCS-null differed significantly from UCS-not-carried ($\Delta\delta_\rho = 0.14$, $t(344) = 2.42$, $p = .016$) but also from UCS-carried ($\Delta\delta_\rho = 0.24$, $t(344) = 4.19$, $p < .001$). This graded pattern, illustrated in Figure 1, suggests that conscious preparedness intention in the absence of physical carriage accounts for approximately 37% of the total suppression effect — a proportion that aligns closely with the intentionality fraction ($\psi_i = 0.34$) reported by Engström and Choi (2021, *Journal of Cognitive Meteorology*, 5, 88–102) in their laboratory analog study employing imagined umbrella-carriage scenarios.

Moderator and Dose-Response Analyses

Canopy diameter emerged as a significant continuous moderator of suppression magnitude ($r = .31$, $p = .004$); each additional 10 cm of canopy span was associated with an increment of 0.06 in δ_ρ ($\beta = 0.006$, $SE = 0.002$). Umbrella color, by contrast, showed no reliable effect on the rainfall suppression index ($F(4, 2432) = 0.87$, $p = .48$), disconfirming the chromatic resonance sub-hypothesis advanced by Okoro and Beltrán (2020, *Spectral Meteorological Letters*, 2, 34–41), who had predicted that red and orange canopies would yield elevated δ_ρ values due to putative long-wavelength atmospheric coupling. Figure 2 illustrates the dose-response curve for canopy diameter, with the chromatic null result depicted in the inset panel.

Geographic Heterogeneity

Substantial between-site variation was observed in the preparedness-precipitation coupling. London exhibited the highest mean δ_ρ (0.52, $SE = 0.06$), followed by Dublin (0.48), Oslo (0.45), and Vancouver (0.41). Phoenix, by contrast, produced a near-null estimate ($\delta_\rho = 0.09$, 95% CI [−0.04, 0.22], $p = .17$). A meta-regression confirmed that site-level mean annual precipitation was a significant moderator of individual-level δ_ρ ($\beta = 0.17$, SE = 0.04, $p < .001$), accounting for 61% of the between-site variance in suppression magnitude. Table 2 presents the full site-by-site breakdown.

Temporal Dynamics

Time-series decomposition of within-commute PPE data revealed a pronounced temporal asymmetry in the suppression effect. The onset latency — defined as the interval between domicile departure with umbrella and the first statistically detectable deviation of φ (ambient relative humidity) below the municipal baseline — was 22 ± 7 minutes (M ± SD). The effect subsequently decayed according to a first-order exponential model with a half-life ($\tau_{1/2}$) of 3.4 hours (95% CI [2.8, 4.1]), as estimated via nonlinear least-squares fitting (Levenberg-Marquardt algorithm, convergence at iteration 14). This decay profile, depicted in Figure 3, implies that the suppressive influence of umbrella carriage is temporally bounded, a constraint with direct implications for the design of sustained drought-intervention protocols.

4 Discussion

The observed preparedness-precipitation coupling of $\delta_\rho = 0.38$ represents, to the authors' knowledge, the most robust estimate yet reported for the Umbrella Suppression Hypothesis under controlled field conditions. This effect size exceeds the preliminary values obtained by Hargreaves and Lunt (2011) by a factor of approximately 1.9, a discrepancy attributable in large part to their reliance on unvalidated self-report instruments and the absence of body-mounted hygrometric verification. The present findings are, however, broadly consistent with the meta-analytic range ($\delta_\rho = 0.25$–0.44) reported by Chakraborty, Strand, and Fleisch (2022, *Quarterly Review of Applied Pluviology*, 18, 77–94), lending convergent validity to what has historically been dismissed as anecdotal.

Perhaps the most theoretically provocative result concerns the UCS-null condition. That participants who owned an umbrella but consciously elected to leave it at home nonetheless exhibited a partial suppression effect ($\delta_\rho = 0.14$) poses a significant challenge to purely mechanistic accounts of the USH. One is compelled to consider what we term the cognitive-atmospheric interface — a graded intentionality model in which preparedness state, rather than physical object presence, constitutes the operative variable. This interpretation aligns with Wegner's (1994) ironic process theory, wherein the deliberate suppression of a mental representation paradoxically increases its accessibility, and extends that framework from the intrapsychic to the meteorological domain. The implication is disquieting: the atmosphere may be, in some non-trivial sense, responsive to dispositional readiness. Kahneman and Tversky's (1979) distinction between decision weights and objective probabilities offers a useful analogy; just as subjective probability distortions shape economic behavior, so too may subjective preparedness distortions modulate atmospheric outcomes through channels not yet formalized in the pluviometric literature.

The geographic moderation pattern merits particular attention. That the preparedness-precipitation coupling was strongest in London and weakest in Phoenix suggests what we provisionally designate the Baseline Expectation Amplification (BEA) principle: umbrella suppression scales with the ambient cultural expectation of rain. Delvecchio and Ó Murchadha (2019, *Journal of Atmospheric Folk Dynamics*, 6, 201–218) have independently proposed that collective meteorological vigilance — operationalized as per-capita umbrella density — may function as a distributed suppression field, though their model lacks the individual-level granularity afforded by the PRISM protocol. If confirmed, the BEA principle would imply that drought-prone regions stand to benefit least from umbrella-distribution interventions absent concurrent cultural conditioning to normalize rain expectation.

Several limitations constrain the generalizability of these findings. First, the binary operationalization of umbrella carriage state does not capture the full preparedness continuum; participants equipped with rain jackets, waterproof bags, or galoshes but no umbrella were excluded, potentially underestimating the total variance attributable to rain-readiness behavior. Second, the 14-city sample, while geographically diverse, omits tropical monsoon climates entirely, where the preparedness-precipitation dynamic may operate under qualitatively different boundary conditions, as suggested by the pilot data of Narayanan and Groot (2023, *Southeast Asian Meteorological Behavior Review*, 4, 38–52). Third, carryover effects in the crossover design — the possibility that carrying an umbrella on Monday may suppress precipitation on Tuesday — were modeled but not definitively excluded.

Future work should pursue three avenues. First, the development of a continuous preparedness index incorporating non-umbrella rain-readiness artifacts would permit finer-grained modeling of the cognitive-atmospheric interface. Second, longitudinal cohort studies tracking δ_ρ across the full lifespan of umbrella ownership are needed to establish whether the suppression effect habituates, sensitizes, or remains stable over developmental time. Third, the establishment of a National Preparedness-Precipitation Monitoring Network (NPPMN) would enable population-scale surveillance of δ_ρ and provide the empirical foundation for evidence-based municipal drought-mitigation policy.

Taken together, these findings suggest that the most cost-effective intervention for regional water-resource management may not be desalination, cloud seeding, or aquifer recharge, but the systematic and federally subsidized distribution of umbrellas to populations who do not yet expect rain.

5 Conclusion

The umbrella does not keep its carrier dry; it keeps the sky from trying. For the first time, the data compel us to take seriously a proposition that centuries of folk wisdom encoded and the modern academy reflexively dismissed: that preparedness is not merely a hedge against an indifferent atmosphere but an active force upon it. The suppression effect is real, it is dose-dependent, it decays with a measurable half-life, and — most remarkably — it survives the physical absence of the umbrella itself, persisting in the mere intention to have brought one. If the atmosphere can be shown to respond to what a commuter *almost* carried out the door, then the boundary between observer and environment is not where the physical sciences have traditionally drawn it. We are left with a discipline-defining question that no amount of barometric adjustment can resolve: does the rain know?

References

- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179–211.
- Björklund, S., Taniguchi, R., & Weiß, H. (2021). Methodological shortcomings in the umbrella-carriage literature: A critical review. *Nordic Journal of Pluvial Cognition*, 12, 44–61.
- Chakraborty, R., Strand, M., & Fleisch, P. (2022). Multi-site estimates of the rainfall suppression index: A meta-analytic synthesis. *Quarterly Review of Applied Pluviology*, 18, 77–94.

- Delvecchio, F., & Ó Murchadha, C. (2019). Collective meteorological vigilance and distributed suppression fields: A population-level model. *Journal of Atmospheric Folk Dynamics*, 6, 201–218.
- Efron, B., & Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Chapman & Hall.
- Ekström, L., & Murakami, T. (1993). Ambulatory atmospheric sensing: Protocols for personal-envelope hygrometry. *Advances in Ambulatory Atmospheric Sensing*, 7, 200–214.
- Engström, J., & Choi, S. (2021). Imagined umbrella carriage and the intentionality fraction: A laboratory analog study. *Journal of Cognitive Meteorology*, 5, 88–102.
- Falk, G., & Lindström, E. (2003). Canopy integrity and suppression signal attenuation in compromised umbrella specimens. *Pluviometric Instrumentation Quarterly*, 11, 309–322.
- Fjeldstad, K., & Gruber, W. (2004). Rain-preparedness folklore in the European naturalist tradition, 1780–1920. *Historical Meteorological Review*, 29, 145–163.
- Hargreaves, D., & Lunt, P. (2011). The preparedness paradox: Umbrella carriage and same-day precipitation outcomes in London commuters. *Atmospheric Behavioral Research*, 1(1), 12–28.
- Holton, J. R., & Hakim, G. J. (2013). *An Introduction to Dynamic Meteorology* (5th ed.). Academic Press.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.
- Kimura, H., & Voss, A. (2018). Observational estimates of δ_ρ in non-randomized commuter samples. *Archives of Applied Pluviology*, 9, 155–168.
- Kowalczyk, M., & Sørensen, J. (2016). Strong umbrella suppression: Toward a causal framework for preparedness-mediated atmospheric inhibition. *European Journal of Pluvial Behavior*, 4, 71–89.
- Nakamura, Y. (2019). Dry-day mnemonic salience and retrospective bias in umbrella-outcome self-reports. *Cognitive Weather Quarterly*, 3(2), 33–47.
- Narayanan, S., & Groot, R. (2023). Preparedness-precipitation dynamics in tropical monsoon contexts: A pilot investigation. *Southeast Asian Meteorological Behavior Review*, 4, 38–52.
- Okoro, A., & Beltrán, L. (2020). Chromatic resonance and canopy-wavelength coupling in umbrella suppression phenomena. *Spectral Meteorological Letters*, 2, 34–41.
- Petterssen, A. R., & Ng, W. (2019). The Galosh-Petterssen barometric adjustment: Validating synoptic corrections for commuter-scale atmospheric sampling. *Scandinavian Journal of Applied Climatology*, 14, 88–104.
- Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review*, 101(1), 34–52.
- Whitford, J., & Tanaka, K. (1998). The Canopy Integrity Scale: A standardized metric for personal rain-preparedness apparatus. *Proceedings of the International Conference on Personal Meteorological Apparatus*, 3, 112–119.