Title

Dust Storms and Public Health: Integrating Dust Exposure Modeling, Epidemiological Analysis, and Early Warning Systems in a Changing Climate

ISBN
9798186379041

Author
Alomair, Saud Abdulaziz

Publication Date
2026-08-18

UNIVERSITY OF CALIFORNIA

Los Angeles

Dust Storms and Public Health:

Integrating Dust Exposure Modeling, Epidemiological Analysis, and Early Warning Systems in a Changing Climate

A dissertation submitted in partial satisfaction of the

requirements for the degree Doctor of Philosophy

in Environmental Health Sciences

by

Saud Abdulaziz Alomair

2026

ABSTRACT OF THE DISSERTATION

Dust Storms and Public Health:

Integrating Dust Exposure Modeling, Epidemiological Analysis,

and Early Warning Systems in a Changing Climate

by

Saud Abdulaziz Alomair

Doctor of Philosophy in Environmental Health Sciences

University of California, Los Angeles, 2026

Professor Michael Jerrett, Chair

Dust storms and chronic dust loadings are a growing but heterogeneously characterized public health concern in California, where a warming climate, more frequent heat and drought, and expanding disturbed land surfaces are increasing dust activity. This dissertation builds, in sequence, the evidence a dust warning system requires, organized around the four components of a people-centered early warning system: risk knowledge, monitoring and forecasting, communication, and response.

Aim 1 systematically reviews global dust storm early warning systems. Of 22 included studies, 19 addressed monitoring and forecasting while only three addressed risk knowledge, communication, or response, a gap that has persisted for more than two decades and that frames the dissertation.

Aim 2 develops and externally validates a daily, statewide surface of the dust-attributable fraction of PM_{10} (DAF-PM_{10}) at population-weighted ZIP-code centroids for 2006 through 2019, combining a machine-learning ensemble with a transparent rule-based comparator. Evaluated against independent composition-based dust, the ensemble preserved strong rank-

order agreement (pooled Spearman rho 0.65) and discriminated high-dust days well (AUC up to 0.90), localizing chronic burden to the southern San Joaquin and Imperial Valleys.

Aim 3 links this exposure surface to roughly 24 million case-crossover strata of cardiovascular and respiratory encounters, 2006 through 2018. A one-interquartile-range increase in DAF-PM_{10} raised the odds of a respiratory encounter by about 1% on the same day, peaking one day later and accumulating to roughly 1.4% over a week, and raised cardiovascular odds by about 0.2%. The respiratory response was nonlinear, concentrated in asthma and chronic obstructive pulmonary disease, amplified specifically by peak daytime heat, and positive across most populated air basins, with the notable exception of the chronically dust-saturated Salton Sea basin, where the case-crossover design returned a null estimate; cardiovascular effects concentrated in hypertensive disease and coastal basins.

Aim 4 translates this evidence into a risk communication and community engagement plan, an evidence-based infographic, and a short educational video, and shows how the exposure metric and the dust-by-heat findings can inform alert thresholds and timing. Together, the four aims form an exposure-to-action pathway that makes a climate-sensitive hazard more measurable, more clearly linked to health, and more actionable.

The dissertation of Saud Abdulaziz Alomair is approved.

Miriam Marlier

Jasper Kok

Lara Cushing

Michael Jerrett, Committee Chair

University of California, Los Angeles

2026

To the communities that live with dust, and to the next dust season we can be readier for.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ACKNOWLEDGMENTS

I am deeply grateful to the chair and members of my doctoral committee for their guidance, thoughtful counsel, and confidence throughout this journey. Their insight challenged me to think more carefully, work more rigorously, and grow as a scholar.

I am equally grateful to the faculty, staff, colleagues, and fellow students of the UCLA Fielding School of Public Health, whose knowledge, generosity, and collegiality enriched both this dissertation and my doctoral experience. I also gratefully acknowledge the collaborators whose expertise contributed to this research, including Bo and the work of Alex Hall and colleagues that supported components of the environmental exposure analysis.

I extend my sincere appreciation to King Saud bin Abdulaziz University for Health Sciences (KSAU-HS) for the scholarship that made it possible for me to pursue this stage of my higher education. I am profoundly grateful for this investment in my education and professional development.

Above all, I thank my family for their patience, sacrifices, encouragement, and unwavering support over the many years that brought me to this point. Their presence gave me strength during the most demanding moments and meaning to the accomplishments along the way. I am also grateful to my friends, near and far, who offered encouragement, understanding, laughter, and perspective when I needed them most.

This dissertation represents years of learning, persistence, and growth. To everyone who, in ways large or small, made that journey richer and more meaningful, I offer my heartfelt thanks.

The wind-erosion (WEMO) horizontal and vertical dust-flux fields used in Chapter 2 were generated by a collaborator, Bo, from high-resolution downscaled California wind fields

developed by Alex Hall and colleagues; all other work reported in this dissertation is the author's own, including the exposure modeling, feature integration, external validation, post-calibration, and interpretation in Chapter 2, the systematic review in Chapter 1, the epidemiologic analysis in Chapter 3, and the translational products in Chapter 4. Chapters 1 through 4 are written as self-contained studies, each retaining its own reference list, and are in preparation for submission to peer-reviewed journals. The health analysis in Chapter 3 used administrative data from the California Department of Health Care Access and Information under the data-use terms governing that access. This research was conducted without dedicated external grant funding except as noted above.

VITA

EDUCATION

MBBS, King Saud University, 2016

MPH, University of California, Los Angeles, 2020

Introduction

Wind-blown mineral dust is among the most widespread atmospheric hazards of arid and semi-arid regions, and its footprint is widening as land use intensifies and the climate warms (Goudie & Middleton, 2001; Middleton, 2019). Dust degrades air quality across entire continents, reshapes soils and ecosystems, and reduces visibility on roads and runways. For a hazard so common, however, it remains heterogeneously characterized in the places where people actually live, work, and seek care. Projections point toward more frequent and more intense dust activity in many drylands as warming proceeds, drought intensifies, and disturbed land surfaces expand (Lal, 2012; Pu et al., 2022), so the population exposed to dust, and the health burden that exposure carries, are likely to grow rather than recede. This dissertation treats dust not only as a meteorological phenomenon but as a public health exposure, and it asks what it would take to characterize that exposure well enough to study its health effects and to warn the people most at risk.

The health rationale is well established. Inhaled dust has been linked to cardiovascular and respiratory morbidity and to non-accidental mortality across a range of settings and designs. In the United States, dust events have been associated with excess non-accidental deaths (Crooks et al., 2016), with intensive care unit admissions (Rublee et al., 2020), and with cardiorespiratory emergency department visits in the arid Southwest (Rowan et al., 2024). Internationally, desert and Asian dust have been tied to asthma hospitalization, particularly in children (Kanatani et al., 2010), and meta-analyses report consistent elevations in respiratory morbidity on dust days alongside smaller and more heterogeneous cardiovascular signals (Hashizume et al., 2020; Tobías et al., 2025). Dust also harms health indirectly. The reduction in visibility during an event raises the risk of road traffic collisions (Bhattachan et al., 2019; Islam et al., 2019), and severe episodes can strain emergency services at the very moment demand peaks. These effects fall unevenly. Older adults, young children, and people with pre-existing cardiopulmonary disease are more susceptible to particulate exposure than the general population (Bell et al., 2013), which means that the burden of dust concentrates in exactly the groups that conventional messaging reaches least well.

Translating this evidence into population health research has proved harder than the breadth of the literature suggests, and the obstacle is largely one of measurement. Dust is not a single, settled construct. It sits at the intersection of particulate mass, mineral composition, source environment, atmospheric transport, and measurement scale, and no single instrument captures the whole signal (Querol et al., 2019; Morakinyo et al., 2016). In practice, studies have represented dust by event flags, storm catalogues, monitor exceedances, compositional surrogates, satellite aerosol products, transport models, or hybrid definitions, and these constructs are not interchangeable (Karanasiou et al., 2012; Tobías & Stafoggia, 2020; Hohsfield et al., 2023). Each captures a different part of the underlying phenomenon and carries a different mix of spatial coverage, temporal continuity, and physical specificity. The consequence for epidemiology is direct. When dust is measured by sparse monitors or coarse event flags, true exposure is misclassified across person-time, and the resulting error tends to pull health associations toward the null. Better exposure characterization is therefore necessary. It reduces a specific, direction-predictable source of bias and provides the continuous exposure gradient needed to estimate how health risk changes with dose.

California is a demanding and consequential place to confront this problem. The state combines steep gradients in aridity, topography, land cover, and human settlement with intense and partly anthropogenic dust sources. The receding Salton Sea exposes playa that has become an intensifying dust source as the lake shrinks (Parajuli & Zender, 2018; Frie et al., 2019); fallowed agricultural land has emerged as a dominant anthropogenic source statewide (Adebiyi et al., 2025); and the Central Valley generates agricultural and roadway dust intermixed with combustion particulate. Long-term monitoring established the western United States as a meaningful region for dust reconstruction (Tong et al., 2012), later work showed that western dust events are best identified by integrating meteorological, satellite, and air-quality information rather than any single indicator (Lei & Wang, 2014), and fine-dust loadings across the wider Southwest have risen over recent decades (Achakulwisut et al., 2017). California also concentrates most of its population in coastal metropolitan basins, so a statewide health analysis is statistically dominated by those populations even as the heaviest dust burden falls in the interior. Finally, dust in California increasingly co-occurs with heat and drought, and such compound events are projected to become more frequent (Pu et al., 2022), which makes the joint cardiopulmonary burden of dust and heat a question of growing policy relevance. The combination of natural and disturbed sources, dense and unevenly exposed populations, and an unusually rich data environment makes California a scientifically exacting test bed rather than an arbitrary application domain.

Detecting and predicting dust, however, is only half of what a protective system requires. International disaster policy, from the Hyogo Framework to the Sendai Framework, places early warning at the center of risk reduction, and the United Nations defines a people-centered early warning system through four linked components: knowledge of the risk, monitoring and forecasting of the hazard, communication of the warning, and the capacity to act on it (Basher, 2006; United Nations Office for Disaster Risk Reduction, 2015). A system is only as strong as its weakest link. A forecast that no one receives, or that recipients cannot act on, protects no one, and the global Early Warnings for All initiative has identified communication and response as the components most in need of strengthening (World Meteorological Organization, 2022). Warning systems work where the links hold: a European heat-health system reduced heat-related mortality, though its value depended on the quality of the forecasts behind it (Lowe et al., 2016). For dust specifically, the question of whether the warning chain is complete has not been examined as a whole, and the exposure science needed to set health-relevant thresholds has lagged behind the science of detection.

This dissertation responds to that gap by building, deliberately and in sequence, the evidence a dust warning system needs. Its organizing logic is a chain that runs from hazard to harm to action. First, the field must know what an effective dust warning system looks like and where current practice falls short. Second, it must be able to assign dust-related exposure to populated places across space and time, not only on days that satisfy a strict event definition. Third, it must establish whether that exposure is associated with acute illness, and under what conditions the association is strongest. Fourth, it must carry the resulting evidence across the last mile, into resources that agencies and communities can use before a storm arrives. Each link motivates the next. A warning is worth building only if exposure can be characterized; exposure is worth characterizing only if it predicts harm; and evidence of harm protects health only when it is communicated and acted upon. The four aims of this dissertation correspond to these four links and to the four components of a people-centered warning system.

The aims are as follows. Aim 1 (Chapter 1) is a systematic review of global dust storm early warning systems, read against the four-component model rather than catalogued by detection technique, which reveals where research has concentrated and which links of the warning chain remain thinly built. Aim 2 (Chapter 2) develops and externally validates a daily, statewide

California surface of the dust-attributable fraction of PM_{10} (DAF-PM_{10}) at population-weighted ZIP-code centroids for 2006 through 2019, combining machine-learning estimation with a transparent rule-based comparator and evaluating both against independent composition-based dust measurements. Aim 3 (Chapter 3) links that exposure surface to roughly twenty-four million case-crossover strata of cardiovascular and respiratory emergency department and hospital encounters, estimating the short-term association between dust-attributable PM_{10} and acute cardiopulmonary morbidity and testing its concentration-response shape, lag structure, modification by heat, and geographic heterogeneity. Aim 4 (Chapter 4) translates the accumulated evidence into a coordinated set of risk communication resources, a community engagement plan, an evidence-based infographic, and a short educational video, and shows how the exposure and health findings can inform the thresholds and timing that an operational warning system must deliver.

The sequence is the argument. The systematic review identifies the under-built components of the warning chain and, in doing so, justifies the investment in exposure science and health evidence that follows. The exposure surface is constructed specifically so that its rank ordering of dust conditions can be carried into an epidemiologic analysis at the scale California's administrative health data permit. The health analysis, in turn, supplies the dust-and-health and dust-by-heat evidence that gives a warning system its public health rationale and points toward the exposure levels and compound conditions at which alerts matter most. The communication chapter then closes the loop the review opened, addressing the very components, warning and response, that the review found most neglected. Read together, the chapters are meant to function as one integrated body of work rather than four separate papers, yet each chapter also stands on its own as a publishable study, an exposure-to-action pathway for a climate-sensitive hazard that is poorly served by existing tools.

Two threads run through the whole. The first is honesty about uncertainty. Because no gold standard for dust exposure exists, the exposure surface is validated against the best available independent comparator, and it is described as a dust-informed particulate surface rather than a chemically isolated dust mass, the health estimates are framed as short-term associations with a conservative, meteorology-adjusted exposure rather than as causal effects, and the communication products are presented as evidence-grounded but not yet evaluated with their intended audiences. The second thread is translational intent. Every chapter is built to feed the next and, ultimately, to inform how a population is warned about dust. The remainder of the dissertation develops this pathway in order, beginning with what the global record of dust warning systems reveals about the work that remains to be done.

References

Achakulwisut, P., Shen, L., & Mickley, L. J. (2017). What controls springtime fine dust variability in the western United States? Investigating the 2002–2015 increase in fine dust in the U.S. Southwest. Journal of Geophysical Research: Atmospheres, 122(22), 12449–12467. https://doi.org/10.1002/2017JD027208

Adebiyi, A. A., Kibria, M. M., Abatzoglou, J. T., Ginoux, P., & Pandey, P. (2025). Fallowed agricultural lands dominate anthropogenic dust sources in California. Communications Earth & Environment, 6, Article 324. https://doi.org/10.1038/s43247-025-02306-0

Basher, R. (2006). Global early warning systems for natural hazards: Systematic and people-centred. Philosophical Transactions of the Royal Society A, 364(1845), 2167–2182. https://doi.org/10.1098/rsta.2006.1819

Bell, M. L., Zanobetti, A., & Dominici, F. (2013). Evidence on vulnerability and susceptibility to health risks associated with short-term exposure to particulate matter: A systematic review and meta-analysis. American Journal of Epidemiology, 178(6), 865–876. https://doi.org/10.1093/aje/kwt090

Bhattachan, A., Okin, G. S., Zhang, J., Vimal, S., & Lettenmaier, D. P. (2019). Characterizing the role of wind and dust in traffic accidents in California. GeoHealth, 3(10), 328–336. https://doi.org/10.1029/2019GH000212

Crooks, J. L., Cascio, W. E., Percy, M. S., Reyes, J., Neas, L. M., & Hilborn, E. D. (2016). The association between dust storms and daily non-accidental mortality in the United States, 1993–2005. Environmental Health Perspectives, 124(11), 1735–1743. https://doi.org/10.1289/EHP216

Frie, A. L., Garrison, A. C., Schaefer, M. V., Bates, S. M., Botthoff, J., Maltz, M., Ying, S. C., Lyons, T., Allen, M. F., Aronson, E., & Bahreini, R. (2019). Dust sources in the Salton Sea basin: A clear case of an anthropogenically impacted dust budget. Environmental Science & Technology, 53(16), 9378–9388. https://doi.org/10.1021/acs.est.9b02137

Goudie, A. S., & Middleton, N. J. (2001). Saharan dust storms: Nature and consequences. Earth-Science Reviews, 56(1–4), 179–204. https://doi.org/10.1016/S0012-8252(01)00067-8

Hashizume, M., Kim, Y., Ng, C. F. S., Chung, Y., Madaniyazi, L., Bell, M. L., Guo, Y. L., Kan, H., Honda, Y., Yi, S.-M., Kim, H., & Nishiwaki, Y. (2020). Health effects of Asian dust: A systematic review and meta-analysis. Environmental Health Perspectives, 128(6), 066001. https://doi.org/10.1289/EHP5312

Hohsfield, L., Rowan, A., D'Souza, R., Ebelt, S., Chang, H. H., & Crooks, J. (2023). Evaluating data product exposure metrics for use in epidemiologic studies of dust storms. GeoHealth, 7(8), e2023GH000824. https://doi.org/10.1029/2023GH000824

Islam, M. M., Alharthi, M., & Alam, M. M. (2019). The impacts of climate change on road traffic accidents in Saudi Arabia. Climate, 7(9), 103. https://doi.org/10.3390/cli7090103

Kanatani, K. T., Ito, I., Al-Delaimy, W. K., Adachi, Y., Mathews, W. C., & Ramsdell, J. W. (2010). Desert dust exposure is associated with increased risk of asthma hospitalization in children. American Journal of Respiratory and Critical Care Medicine, 182(12), 1475–1481. https://doi.org/10.1164/rccm.201002-0296OC

Karanasiou, A., Moreno, N., Moreno, T., Viana, M., de Leeuw, F., & Querol, X. (2012). Health effects from Sahara dust episodes in Europe: Literature review and research gaps. Environment International, 47, 107–114. https://doi.org/10.1016/j.envint.2012.06.012

Lal, R. (2012). Climate change and soil degradation mitigation by sustainable management of soils and other natural resources. Agricultural Research, 1(3), 199–212. https://doi.org/10.1007/s40003-012-0031-9

Lei, H., & Wang, J. X. L. (2014). Observed characteristics of dust storm events over the western United States using meteorological, satellite, and air quality measurements. Atmospheric Chemistry and Physics, 14(15), 7847–7857. https://doi.org/10.5194/acp-14-7847-2014

Lowe, R., García-Díez, M., Ballester, J., Creswick, J., Robine, J.-M., Herrmann, F., & Rodó, X. (2016). Evaluation of an early-warning system for heat wave-related mortality in Europe. International Journal of Environmental Research and Public Health, 13(2), 206. https://doi.org/10.3390/ijerph13020206

Middleton, N. (2019). Dust storm hazards. E3S Web of Conferences, 99, 04001. https://doi.org/10.1051/e3sconf/20199904001

Morakinyo, O. M., Mokgobu, M. I., Mukhola, M. S., & Hunter, R. (2016). Health outcomes of exposure to biological and chemical components of inhalable and respirable particulate matter. International Journal of Environmental Research and Public Health, 13(6), 592. https://doi.org/10.3390/ijerph13060592

Parajuli, S. P., & Zender, C. S. (2018). Projected changes in dust emissions and regional air quality due to the shrinking Salton Sea. Aeolian Research, 33, 82–92. https://doi.org/10.1016/j.aeolia.2018.05.004

Pu, B., Jin, Q., Ginoux, P., & Yu, Y. (2022). Compound heat wave, drought, and dust events in California. Journal of Climate, 35(24), 8133–8152. https://doi.org/10.1175/JCLI-D-21-0889.1

Querol, X., Tobías, A., Pérez, N., Karanasiou, A., Amato, F., Stafoggia, M., et al. (2019). Monitoring the impact of desert dust outbreaks for air quality for health studies. Environment International, 130, 104867. https://doi.org/10.1016/j.envint.2019.05.061

Rowan, C., D'Souza, R. R., Zheng, X., Crooks, J., Hohsfield, K., Tong, D., Chang, H. H., & Ebelt, S. (2024). Dust storms and cardiorespiratory emergency department visits in three southwestern United States. Environmental Research: Health, 2(3), 031003. https://doi.org/10.1088/2752-5309/ad5751

Rublee, C. S., Sorensen, C. J., Lemery, J., Wade, T. J., Sams, E. A., Hilborn, E. D., & Crooks, J. L. (2020). Associations between dust storms and intensive care unit admissions in the United States, 2000–2015. GeoHealth, 4(8), e2020GH000260. https://doi.org/10.1029/2020GH000260

Tobías, A., & Stafoggia, M. (2020). Modeling desert dust exposures in epidemiologic short-term health effects studies. Epidemiology, 31(6), 788–795. https://doi.org/10.1097/EDE.0000000000001255

Tobías, A., Karanasiou, A., Amato, F., Roqué, M., Lwin, K. S., Yuan, L., Ith, S., Wai, H. Z., Chua, P. L., Solà, I., Renzi, M., Stafoggia, M., Querol, X., & Hashizume, M. (2025). Short-term exposure to desert dust and sandstorms and all-cause and cause-specific mortality and morbidity: A systematic review and meta-analysis. Environment International, 196, 109277. https://doi.org/10.1016/j.envint.2025.109277

Tong, D. Q., Dan, M., Wang, T., & Lee, P. (2012). Long-term dust climatology in the western United States reconstructed from routine aerosol ground monitoring. Atmospheric Chemistry and Physics, 12(11), 5189–5205. https://doi.org/10.5194/acp-12-5189-2012

United Nations Office for Disaster Risk Reduction. (2015). Sendai framework for disaster risk reduction 2015–2030. United Nations.

World Meteorological Organization. (2022). Early warnings for all: Executive action plan 2023–2027. World Meteorological Organization.

Chapter 1. Forecasting Without Warning: A Systematic Review of Early Warning Systems for Dust Storms

Abstract

Background. Dust storms recur across arid and semi-arid regions, and their reach is widening as land use changes and the climate warms. Inhalation of airborne dust increases cardiovascular and respiratory illness and raises mortality, while the loss of visibility during an event causes road casualties and strains emergency services. International disaster policy, from the Hyogo Framework to the Sendai Framework, places early warning at the center of risk reduction. The UN Office for Disaster Risk Reduction defines a people-centered early warning system through four linked components: knowledge of the risk, monitoring and forecasting of the hazard, communication of the warning, and the capacity to respond. This review examines how the dust-storm literature maps onto that chain.

Methods. The review followed PRISMA guidelines. PubMed, EMBASE, and Web of Science were searched. Of 94 records identified, 22 studies met the inclusion criteria after screening. Each study was classified by the early warning component it addressed and, within forecasting, by the method applied.

Results. The evidence concentrates in a single link of the chain. Nineteen of the 22 studies addressed monitoring and forecasting through numerical modelling, mapping, satellite observation, or in-situ networks. Three engaged risk knowledge, warning communication, or response capacity, and these remained largely qualitative. Output did not rise across the two decades spanned by the included studies (2004–2021), even as dust events grew more frequent.

Conclusion. The field has developed one component of a four-part system in detail and left the others sparse. Detecting and predicting dust is necessary but not sufficient: a warning that goes uncommunicated and unheeded does not reduce harm. Strengthening risk knowledge,

communication, and response is the priority for making dust-storm early warning systems function as systems.

Introduction

Dust storms rank among the most widespread atmospheric hazards of arid and semi-arid regions, shaping ecosystems as much as disrupting them (Goudie & Middleton, 2001), and they carry consequences for human health and the environment (Lambert et al., 2020). Projections indicate more frequent and more intense events as the climate warms (Lal, 2012), so the burden carried by the dust is likely to grow.

The harm reaches people by two routes. The direct route is inhalation during an event. Dust exposure has been linked to cardiovascular (Zauli Sajani et al., 2010) and respiratory (Kashima et al., 2012) deaths and to a rise in non-accidental mortality (Crooks et al., 2016). It has been associated with asthma attacks, hospital admissions, and reduced lung function (Kanatani et al., 2010; Aili & Kim Oanh, 2015; Watanabe et al., 2015), and one study reported a higher risk of suicide in the days surrounding an event (Lee et al., 2019). The indirect route runs through the surrounding environment, most visibly the collapse in road visibility that causes traffic collisions (Islam et al., 2019).

These effects impose a heavy public cost (Rublee et al., 2020; Middleton et al., 2019). A severe event can overwhelm emergency departments and intensive care units, lowering the standard of care available to every other patient at the same time. As events become more frequent, the case for prevention and adaptation strengthens with them.

Early warning is the adaptation that disaster policy returns to most often. The Hyogo and Sendai frameworks both treat it as a foundation of risk reduction. The United Nations Office for Disaster Risk Reduction defines an early warning system as “an integrated system of hazard monitoring, forecasting and prediction, disaster risk assessment, communication and preparedness activities, systems and processes that enables individuals, communities, governments, businesses and others to take timely action to reduce disaster risks in advance of hazardous events." Read closely, the definition names four linked components: knowledge of the risk, monitoring and forecasting of the hazard, communication of the warning, and the capacity to act on it. A warning system is only as strong as its weakest link. A forecast that no one receives, or that recipients cannot act on, protects no one.

Where the links hold, warning systems work. A European heatwave system reduced heat-related mortality, though its value depended on the quality of the climate data and forecast models behind it (Lowe et al., 2016). Warning and information systems likewise shaped public behavior during the COVID-19 pandemic (Guo et al., 2022). For dust specifically, several authors have called for warning systems to reduce the health toll (Lee et al., 2014; Liu et al., 2017; Stefanski & Sivakumar, 2009). A system built for one hazard does not transfer cleanly to another: risk identification, risk knowledge, and decision making must each be designed for the event at hand (Allahbakhshi et al., 2019; Fatemi et al., 2015).

Despite these repeated calls, the dust-storm warning literature has not been examined as a whole, and the components of a warning system that individual studies address have not been distinguished from one another. This review maps the global literature onto the four-component model. It asks which links of the warning chain the field has built, which it has left thin, and what that imbalance means for systems intended to reduce harm rather than only anticipate it.

Methods

Three databases (PubMed, EMBASE, and Web of Science) were searched for published work on dust-storm early warning systems worldwide between 19 and 21 April 2022. The search was conducted in Los Angeles, California. The review followed PRISMA guidelines. No protocol was registered, and a single reviewer performed screening, data extraction, and synthesis.

An initial scoping search informed the choice of terms. The final search string was ("dust particles" OR "dust storm" OR "sand storm" OR "dust pollution" OR "dust") AND ("warning system" OR "early warning system"), chosen to capture the range of wording used for both dust and warning systems.

Titles and abstracts were screened first against the eligibility criteria, and records that passed advanced to full-text review. Peer-reviewed studies published in English were eligible, with no restriction on publication date, so relevant work from any year and region was considered. A study was included if it addressed both dust storms or dust exposure and a warning system used as an adaptation to that exposure. Studies covering dust storms without a warning system, or warning systems for other hazards, were excluded.

Two attributes were charted from each included study: the early warning component it addressed and, where the study concerned forecasting, the method it applied. The four components follow the UNDRR definition: risk knowledge, monitoring and forecasting, warning communication, and response capacity. Where a study reported how well its approach limited exposure, that result was recorded as well.

The search returned 94 records. After 23 duplicates were removed, 71 records were screened by title and abstract and 37 were excluded. Of the 34 full texts assessed for eligibility, 12 were excluded: nine did not focus on warning systems, two were unrelated to dust storms, and one was a book chapter. Twenty-two studies remained for synthesis (Figure 1.1).

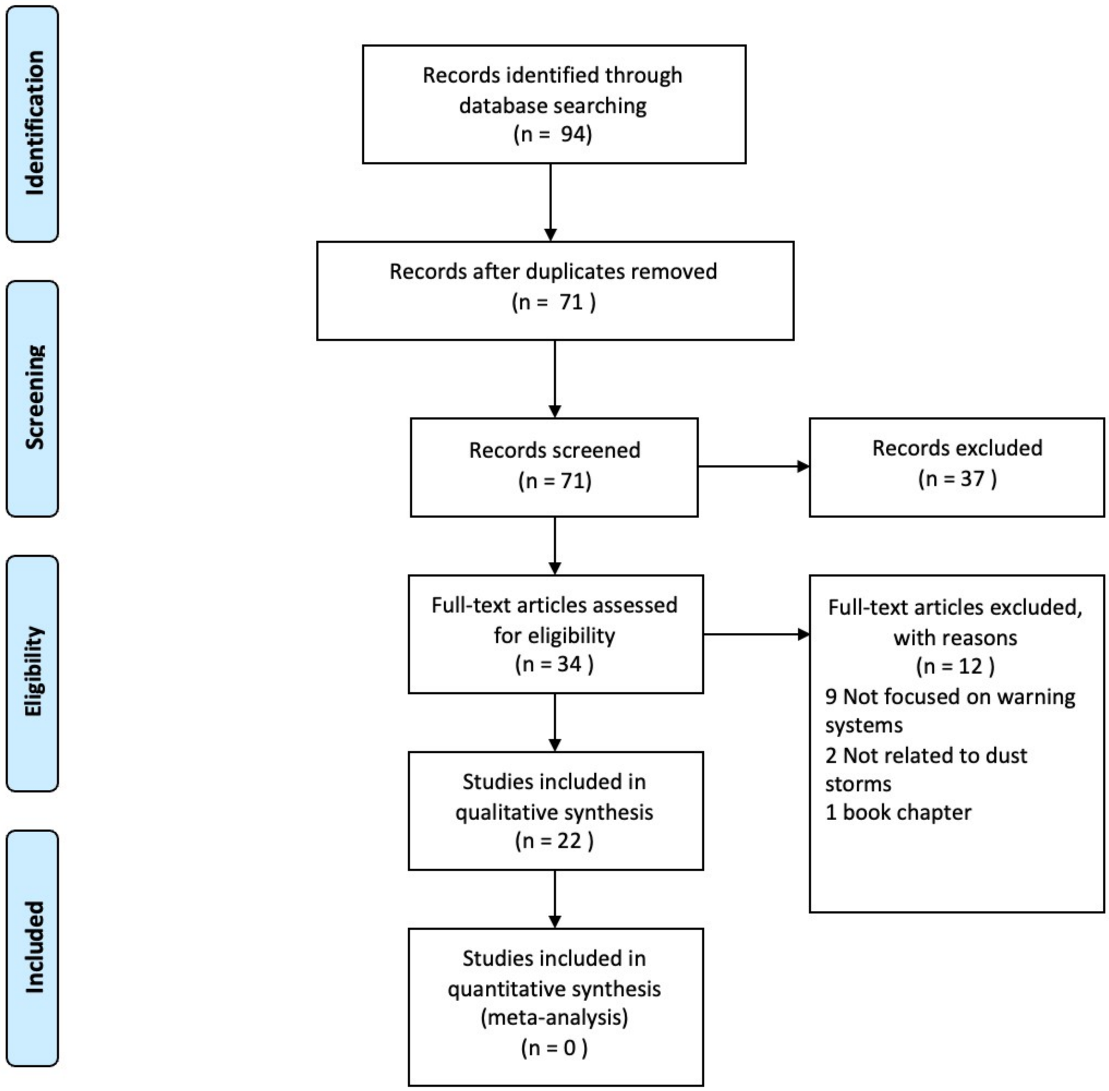

Figure 1.1. PRISMA flow diagram for study selection.

Results

Of the 94 records identified across three databases, 22 met the inclusion criteria (Figure 1.1). Mapped onto the four-component model, the studies cluster sharply. Nineteen addressed monitoring and forecasting, the link that detects a dust event or predicts it before it forms. Three addressed the human side of the chain: risk knowledge, warning communication, and response capacity. None of the 22 built or evaluated a warning system spanning all four components (Table 1.1; Figure 1.2).

Within monitoring and forecasting, the studies divided by method. Thirteen applied numerical modelling and simulation to predict when and where an event would occur. Three used mapping to locate source areas and estimate outbreak potential, two relied on satellite observation, and one used an in-situ station network (Table 1.1). The three studies on the human side were largely qualitative; they examined what must be in place before and during an event to reduce harm: risk assessment, risk communication, public education, stakeholder coordination, emergency planning, and timely warnings.

Output did not grow over time. Eleven studies appeared between 2000 and 2011 and eleven between 2012 and 2022, even as dust events became more frequent across the same period, a flat trajectory concentrated in monitoring and forecasting throughout the window (Figure 1.3). Table 1.2 lists every included study with its component and method, and Table 1.3 sets their settings, inputs, and reported findings side by side.

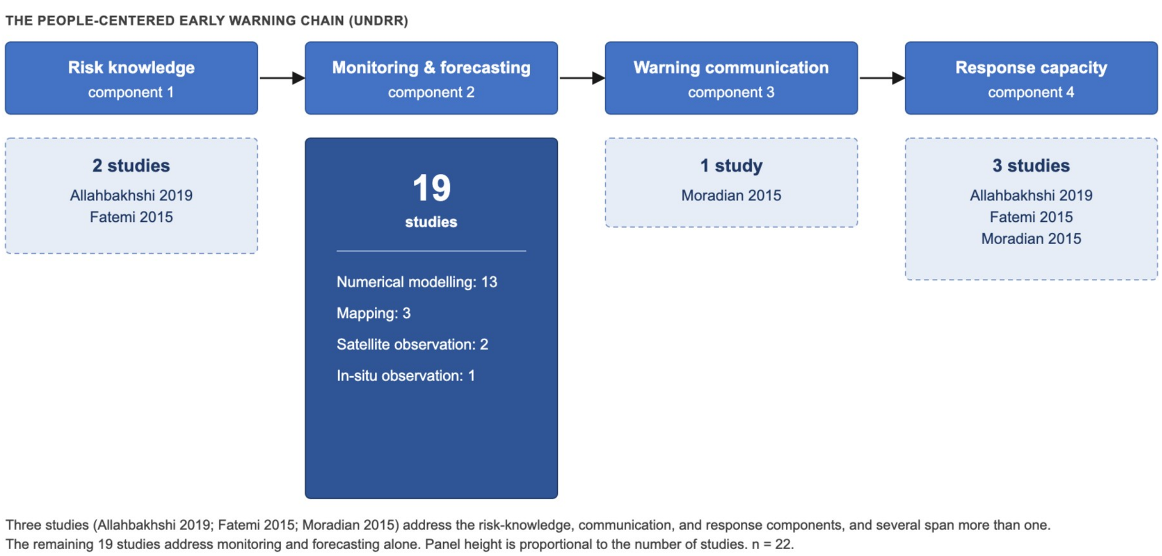

Figure 1.2. Distribution of the 22 included studies across the four components of a people-centered early warning system (UNDRR). Panel height is proportional to the number of studies.

Early warning component	**Studies**
Monitoring and forecasting	19
Numerical modelling	13
Mapping	3
Satellite observation	2
In-situ observation	1
Risk knowledge, communication, and response	3
Qualitative	2
Report	1
Total	22

Table 1.1. Included studies by early warning component and method.

Study	**Title**	**Category**	**Sub-category**
Allahbakhshi et al., 2019	Preparedness challenges of the Iranian health system for dust and sand storms: a qualitative study	Risk knowledge / response	Risk assessment; knowledge management; organizational elements; monitoring and evaluation
Fatemi et al., 2015	Preparedness functions in disaster: lessons learned from the 2014 Tehran dust storm	Risk knowledge / response	Stakeholder cooperation; risk assessment; emergency planning; education and training; exercises and drills
Hu et al., 2004	An early warning system and hierarchical criteria for dust-storm forecast intensity in the Tarim Basin, Xinjiang	Monitoring and forecasting	In-situ observation
Jugder et al., 2018	Developing a soil erodibility map across Mongolia	Monitoring and forecasting	Mapping
Kallos et al., 2009	Ten-year operational dust forecasting: recent model development and future plans	Monitoring and forecasting	Modelling (SKIRON)
Kim et al., 2011	Dust model intercomparison between ADAM and CFORS/Dust for an Asian dust case in 2007	Monitoring and forecasting	Modelling
Kimura, 2018	Satellite-based mapping of dust erodibility in northeast Asia	Monitoring and forecasting	Mapping (GIS with satellite data)
Kimura & Shinoda, 2010	Spatial distribution of threshold wind speeds for dust outbreaks in northeast Asia	Monitoring and forecasting	Mapping (threshold wind speed)
Laurent et al., 2009	Modelling mineral dust emissions	Monitoring and forecasting	Modelling
Lin et al., 2016	Empirical model for evaluating PM_{10} concentration caused by river-dust episodes	Monitoring and forecasting	Modelling
Mahler et al.,	Dust transport model validation using	Monitoring and	Modelling (DREAM)

Study	Title	Category	Sub-category
2006	satellite- and ground-based methods in the southwestern United States	forecasting	
Martinez et al., 2009	Use of SEVIRI images and derived products in a WMO Sand and Dust Storm Warning System	Monitoring and forecasting	Satellite observation (SEVIRI)
Menut et al., 2009	Previsibility of Saharan dust events using the CHIMERE-DUST transport model	Monitoring and forecasting	Modelling (CHIMERE-DUST)
Moradian et al., 2015	Tehran dust storm early warning system: corrective measures	Communication / response	Timely warnings; public education
Nguyen et al., 2021	An ANN-based early warning model for airborne particulate matter in riverbank areas	Monitoring and forecasting	Modelling
Park et al., 2013	A simulation of Asian dust events observed from 20 to 29 December 2009 in Korea using ADAM2	Monitoring and forecasting	Modelling
Samadi et al., 2014	Global Dust Detection Index (GDDI): a remotely sensed method for dust-storm detection	Monitoring and forecasting	Satellite observation (MODIS)
Schulz et al., 2009	LMDzT-INCA dust forecast model developments and associated validation efforts	Monitoring and forecasting	Modelling
Vimic et al., 2021	Numerical simulation of the Tehran dust storm of 2 June 2014: agricultural abandoned lands as emission sources	Monitoring and forecasting	Modelling (DREAM; WRF-Chem)
Vukovic et al., 2014	Numerical simulation of "an American haboob"	Monitoring and forecasting	Modelling (NMME-DREAM)
Westphal et al., 2009	Operational aerosol and dust storm forecasting	Monitoring and forecasting	Modelling (COAMPS)
Yin et al., 2005	Modeling wind-blown desert dust in the southwestern United States for public health warning: a case study	Monitoring and forecasting	Modelling (DREAM)

Table 1.2. Included studies with component and method.

Study	Region & setting	Component & method	Key inputs / data	Main reported finding
Hu et al., 2004	Tarim Basin, Xinjiang, China	Monitoring / in-situ network	40+ station records, 1961–2001; synoptic charts; satellite data	Graded intensity by stations reached (≥10 extraordinarily serious; 5–9 serious; ≤4 moderate/light); accurate for the severe grades
Yin et al., 2005	Southwestern	Monitoring /	Meteorology; land	Predicted $PM_{2.5}$ peak

Study	Region & setting	Component & method	Key inputs / data	Main reported finding
	USA (Arizona)	modelling (DREAM dust cycle)	surface; soil and vegetation; $PM_{2.5}$/PM_{10}; visibility; satellite	timing well; modelled vs observed PM differed up to an order of magnitude; better land-surface and size data needed
Mahler et al., 2006	Southwestern USA (PHAiRS)	Monitoring / modelling (DREAM) + validation	Surface and upper-air data; MODIS L2 and Deep Blue AOT; AERONET; lidar	Bright desert surfaces limit MODIS AOT retrieval; AERONET poorly correlated; multi-sensor and lidar validation needed
Kallos et al., 2009	Sahara / Mediterranean (Athens)	Monitoring / modelling (SKIRON)	Aerosol optical depth; dust size distribution; radiative transfer	Ten years of operational forecasting; continental domain supports prediction
Laurent et al., 2009	Sahara / global emission scheme	Monitoring / modelling	Meteorology; surface type; soil size distribution; roughness length	Input accuracy and particle-size data set the ceiling on model skill
Martinez et al., 2009	WMO SDS-WAS (Spain / N. Africa)	Monitoring / satellite (SEVIRI on MSG)	SEVIRI visible and infrared images and derived products	Dust easily confused with cloud; visible channels improve detection; best combined with models and in-situ data
Menut et al., 2009	Sahara (France)	Monitoring / modelling (CHIMERE-DUST)	Meteorology; wind patterns; surface features; dust size distribution	Wider domain and boundary measurements improve previsibility of Saharan events
Schulz et al., 2009	Global (LSCE, France)	Monitoring / modelling (LMDzT-INCA)	Erodibility; surface wind; threshold friction velocity; soil; dust size	Model development with validation against observations
Westphal et al., 2009	Operational / global (US Navy, NRL)	Monitoring / modelling (NAAPS/COAMPS)	Meteorology; land cover; remote sensing; satellite data	US Navy operational aerosol and dust storm forecasting
Kimura & Shinoda, 2010	Mongolia / NE Asia (35–45N, 100–115E)	Monitoring / mapping	Mongolian station data; MODIS NDVI; roughness length; threshold friction velocity	Mapped threshold wind speed; prior-year maximum NDVI may predict next season's threshold; limited by data quality
Kim et al., 2011	East Asia (28 Mar –3 Apr 2007 case)	Monitoring / modelling (ADAM vs CFORS/Dust)	Surface PM_{10}; AOT; dust extinction; 4D-Var; ground LIDAR;	Both models matched observations but differed up to fourfold

Study	Region & setting	Component & method	Key inputs / data	Main reported finding
			soil; land use	near source regions; better source and emission schemes needed for an EWS
Park et al., 2013	Korea (20–29 Dec 2009 events)	Monitoring / modelling (ADAM2 + MM5)	NDVI; surface winds; dust source; wind direction	Simulated start and end times and peak concentrations well in source and downwind regions
Samadi et al., 2014	Western Iran (20 events, 2000–2011)	Monitoring / satellite (MODIS)	MODIS reflectance and thermal-infrared bands; dust detection index	GDDI detected dust over land and water in all cases; separates dust from cloud and bright surfaces
Lin et al., 2016	Taiwan (Zhuo-Shui River estuary)	Monitoring / modelling (stepwise regression)	Ground PM_{10}; wind velocity; temperature; bare-land area; satellite land cover	Daily-maximum PM_{10} predicted from wind, temperature, and bare-land area (R2 = 0.67); single factors insufficient; time-lag effect used
Jugder et al., 2018	Mongolia	Monitoring / mapping	Global soil texture and soil database; MODIS NDVI; threshold friction velocity	Integrated soil erodibility map; highest erodibility in desert and desert-steppe; a basis for dust risk assessment and EWS
Kimura, 2018	NE Asia (China/Mongolia, 35–50N, 75–120E)	Monitoring / mapping (GIS + MODIS)	MODIS snow cover; frozen soil; surface soil water; vegetation; aridity index	Semi-real-time erodibility map keyed to threshold wind speed; severe dust below ~10 m/s, little above 15 m/s
Vimic et al., 2021	Tehran, Iran (2 Jun 2014 storm)	Monitoring / modelling (coupled dust-atmosphere, DREAM)	Observational data; convective fields; abandoned-farmland sources; soil and wind	Convective downdrafts drove the storm; abandoned agricultural land acted as an emission source; higher resolution and wider domain needed
Nguyen et al., 2021	Estuary / riverbank, Taiwan	Monitoring / modelling (artificial neural network)	PM_{10}; $PM_{2.5}$; temperature; daylight hours; humidity; pressure; wind speed	The model combining media and source characteristics was most accurate; predicts high-pollution events efficiently for an EWS
Vukovic et al., 2014	Phoenix, Arizona, USA (5 Jul 2011	Monitoring / modelling (NMME-	Meteorology; NDVI land mask; agricultural	Reconstructed the haboob with front

Study	Region & setting	Component & method	Key inputs / data	Main reported finding
	haboob)	DREAM, 4 km)	soil sources; PM_{10}; MODIS AOD; CALIPSO	timing within about an hour; shows operational dust forecasting is feasible with existing data
Allahbakhshi et al., 2019 *	Iran (health system)	Risk knowledge / response (qualitative)	Expert interviews on preparedness functions	Identifies risk assessment, knowledge management, organization, and monitoring as preparedness needs (* full text not available)
Fatemi et al., 2015	Tehran, Iran (2014 storm; 5 dead, 82 injured)	Risk knowledge / response (qualitative)	Incident report and lessons-learned analysis	The absence of an early warning system was the key gap; stresses stakeholder cooperation, risk assessment, planning, training, and drills
Moradian et al., 2015	Tehran, Iran (2 Jun 2014; 5 dead, 44 injured)	Communication / response (report)	EOC reports; manager interview; news; literature review	The warning system failed to alert the at-risk population in time; recommends timely alerts and public education on safety measures

Table 1.3. Setting, method, key inputs, and reported findings of the included studies, verified against the full text. * Allahbakhshi et al. (2019): full text unavailable; entry from abstract and synthesis.

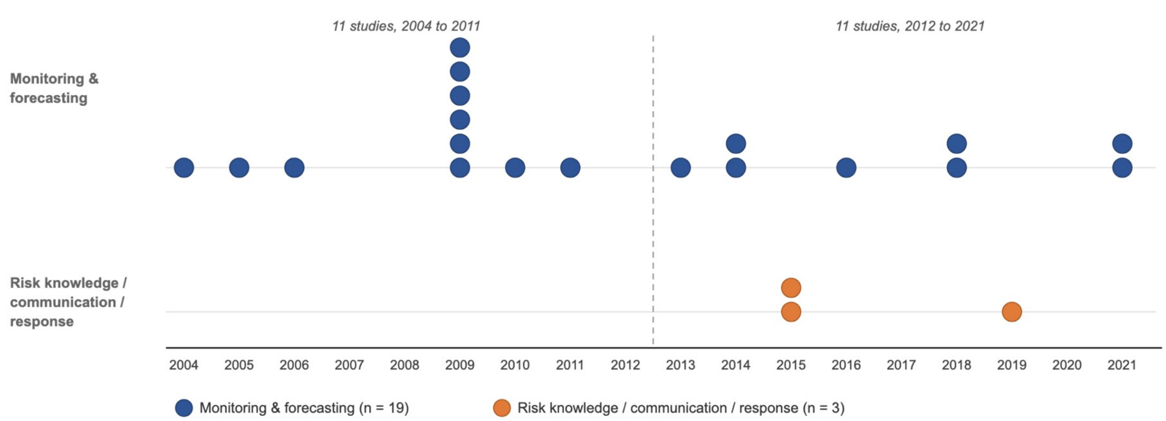

Figure 1.3. Year of publication of the 22 included studies, by early warning component. Output held roughly flat across the period, with 11 studies in 2004 to 2011 and 11 in 2012 to 2021, and concentrated in monitoring and forecasting throughout.

Discussion

Read against the four-component model, the literature describes one link of the warning chain in depth and leaves the rest underdeveloped. The monitoring and forecasting component is well populated and methodologically varied. Risk knowledge, communication, and response are represented by only a handful of studies. The sections below take the forecasting work first, since the field has concentrated there, and then turn to the components that have drawn the least attention.

Monitoring and forecasting

Numerical modelling was the most common approach. Models clarify where dust originates and how it travels, and their accuracy depends on the input variables, the resolution and domain covered, and continued validation against observations (Kim et al., 2011; Laurent et al., 2009; Vimic et al., 2021).

The choice of input variables, and the accuracy with which they are measured, sets a ceiling on what a model can predict (Laurent et al., 2009), which places a premium on understanding what drives a dust event. Meteorological conditions are central, including temperature, relative humidity, wind speed, wind pattern, and atmospheric pressure, and small errors in these inputs can shift the output substantially (Kim et al., 2011). Soil erodibility is a second driver, governed by soil texture, soil type, and wind speed (Jugder et al., 2018), so accurate representation of soil and vegetation improves performance (Yin et al., 2005). Particle size matters as well. The threshold friction velocity, the lowest friction velocity at which soil particles begin to move, varies with particle size, and a higher threshold renders a dust event less likely (Laurent et al., 2009). Reliable particle-size distributions remain difficult to obtain for much of the world, because they are hard to retrieve from satellite data and often require slow, labour-intensive surface analysis. Combining several variables rather than one tends to sharpen prediction (Lin et al., 2016). Sensitivity analysis can identify which variables move the output most: Nguyen et al. (2021) applied such an analysis to a model that performed well for high-pollution events near riverbanks and found that PM_{10}, temperature, $PM_{2.5}$, wind speed, atmospheric pressure, and daylight hours mattered most, in that order. The full set of input variables used by each forecasting model is compiled in Table 1.5, which shows that meteorological and land-surface fields recur across nearly every approach while particle-size and soil-erodibility inputs appear only in a subset.

Resolution and domain also shape accuracy. Vimic et al. (2021) argue for higher resolution and a wider domain. Reproducing features such as strong convective activity requires a resolution on the order of a few kilometres and a domain that extends well beyond the area of interest, because much of the weather that triggers a dust event forms outside it. Models with continental coverage, such as SKIRON and the Dust Regional Atmospheric Model (DREAM), tend to predict events better than models confined to local or regional domains (Menut et al., 2009). A small-domain model cannot detect a storm forming beyond its boundaries, and it requires particulate measurements along those boundaries to limit boundary error.

Validation must be continuous (Mahler et al., 2006). Models can be checked against satellite observations from instruments such as MODIS, SEVIRI/METEOSAT, and CALIPSO, against ground-based remote sensing from networks such as AERONET, and against in-situ measurements (Kallos et al., 2009; Mahler et al., 2006; Menut et al., 2009; Schulz et al., 2009). A single measurement, or a handful, rarely suffices, and a multi-sensor approach proves more reliable than any single aerosol measurement alone (Mahler et al., 2006). Longer runs that incorporate as many dust and non-dust days as possible give a firmer measure of model performance. The validation strategies, reported performance, and most sensitive parameters of the numerical-model studies are compiled in Table 1.4, which makes plain how unevenly these models report quantitative skill.

Vukovic et al. (2014) show how close this work can come to an operational warning. They reconstructed the haboob that struck Phoenix on 5 July 2011 with the coupled NMME-DREAM model at 4 km resolution, placing the dust front within about an hour of its observed arrival and matching surface PM_{10} and MODIS aerosol optical depth (Figure 1.4). The result demonstrates that the data and models needed to forecast a dust event hours in advance, as part of a warning system, are already within reach.

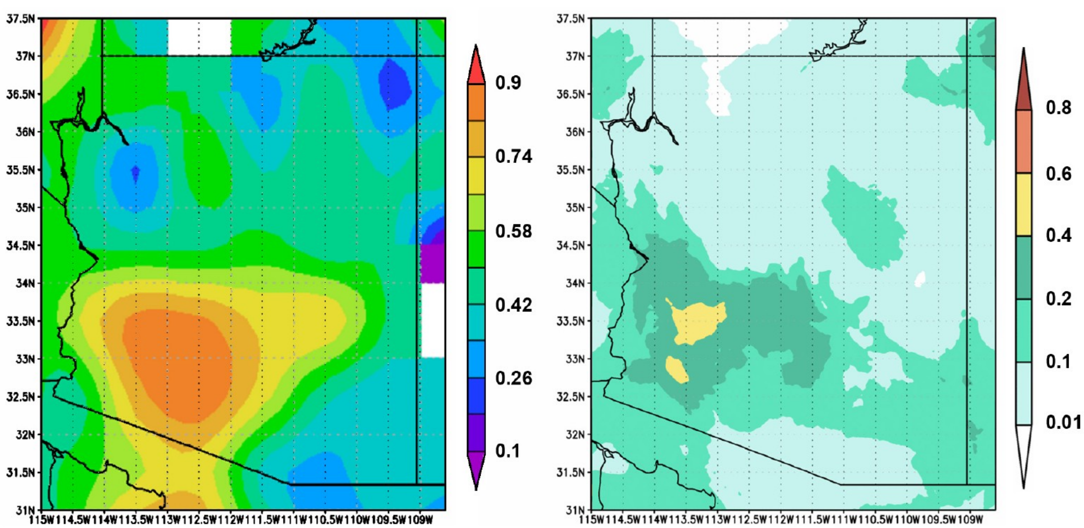

Figure 1.4. Aerosol optical depth over the Phoenix region during the 5 July 2011 haboob: Aqua MODIS retrieval using the deep-blue algorithm (left) and the coupled NMME-DREAM model (right). The modelled field reproduces the spatial pattern of the satellite retrieval, illustrating that the data and models needed for operational dust forecasting are within reach. Reproduced from Vukovic et al. (2014), Atmospheric Chemistry and Physics 14, 3211 to 3230, under the Creative Commons Attribution 3.0 License.

Paper (APA)	**Location**	**Validation method**	**R² value**	**Most sensitive parameter**
Kallos et al. (2009)	North Africa, Mediterranean	SKIRON AOD and dust vs AERONET and	Not reported (model-	Not reported

Paper (APA)	Location	Validation method	R^2 value	Most sensitive parameter
		MODIS	development overview)	
Kim et al. (2011)	East Asia (China, Korea, Japan)	ADAM vs CFORS/Dust; vs surface PM_{10}, AOT, dust extinction	Not reported (qualitative; up to fourfold PM_{10} difference near sources)	Dust source allocation, emission scheme, and meteorological inputs
Laurent et al. (2009)	Sahara / global	Emission scheme vs observations; sensitivity tests	Not reported	Soil and surface inputs; threshold friction velocity (qualitative)
Lin et al. (2016)	Taiwan (Zhuo-Shui River)	Stepwise regression vs observed PM_{10}	0.67 (daily-maximum PM_{10}, Model 3)	Bare land area, with 2-h-prior wind speed and temperature
Mahler et al. (2006)	SW USA (AZ, NM, TX)	DREAM AOT vs MODIS L2/Deep Blue and AERONET; PM, lidar	0.35 (DREAM AOT vs MODIS Deep Blue; r = 0.59)	Not reported (validation study); MODIS retrieval limited over bright surfaces
Menut et al. (2009)	West Africa, Mediterranean	CHIMERE-DUST previsibility vs observations	Not reported	Surface wind speed and source-area activation (qualitative)
Nguyen et al. (2021)	Taichung, Taiwan	ANN vs observed PM_{10}/$PM_{2.5}$ (MAPE, R2, RMSE)	0.74–0.79 (testing, ANN Model III)	PM_{10}, then temperature (by sensitivity analysis)
Park et al. (2013)	Korea, East Asia	ADAM2 with MM5 vs observed PM_{10} (timing, peaks)	Not reported (timing and peaks matched qualitatively)	Surface winds and synoptic system (qualitative)
Schulz et al. (2009)	Sahara, Atlantic, Mediterranean	LMDzT-INCA AOD vs AERONET at Dakar; hit rates	Not reported (AOD correlation r = 0.69 day 0, 0.41 at 5 days)	Forecast lead time (skill decays); initial dust size distribution
Vimic et al. (2021)	Tehran, Iran	DREAM coupled model vs PM, METAR, satellite; source-mask experiments	Not reported	Dust source mask (abandoned-farmland fraction)
Vukovic et al. (2014)	Phoenix, AZ, USA	NMME-DREAM vs PM_{10}, MODIS AOD, CALIPSO, radar	Not reported (front timing within ~1 h)	Dust source mask and land cover

Paper (APA)	Location	Validation method	R^2 value	Most sensitive parameter
Westphal et al. (2009)	Global (NAAPS), regional	NAAPS/COAMPS vs MODIS/AERONET AOD, ground PM	Not reported (bias and RMSE reported)	Wind speed (emission) and data assimilation
Yin et al. (2005)	SW USA (NM, TX)	DREAM vs $PM_{2.5}$/PM_{10}, visibility, satellite	0.76 (modelled vs observed $PM_{2.5}$ peak-hour timing)	Land-surface/ vegetation and dust-size representation

Table 1.4. Validation and sensitivity of the numerical-model studies, verified against the full text. "Not reported" indicates a study that states no such value.

Study	Type	Variables used
Kallos et al., 2009	Modelling (SKIRON)	Aerosol optical depth; dust size distribution; radiative transfer
Kim et al., 2011	Modelling	Meteorological data; 4D-Var and ground-based LIDAR (CFORS/Dust); threshold friction velocity; dust size; dust density; snow cover; land use; soil texture; roughness
Laurent et al., 2009	Modelling	Meteorological data; surface type; soil size distribution; roughness length
Lin et al., 2016	Modelling	Ground-level PM_{10}; temperature; relative humidity; wind speed; land cover from satellite
Mahler et al., 2006	Modelling (DREAM)	Topography; vegetation cover; soil type; humidity; wind patterns
Menut et al., 2009	Modelling (CHIMERE-DUST)	Meteorological data; wind patterns; surface features; dust size distribution
Nguyen et al., 2021	Modelling	PM_{10}; $PM_{2.5}$; temperature; daylight hours; relative humidity; atmospheric pressure; wind speed
Park et al., 2013	Modelling	NDVI; wind speed; source; wind direction
Schulz et al., 2009	Modelling	Erodibility potential; surface wind speed; threshold friction velocity; soil condition; dust size distribution
Vimic et al., 2021	Modelling (DREAM; WRF-Chem)	Topography; vegetation cover; soil type; humidity; wind patterns
Vukovic et al., 2014	Modelling (NMME-DREAM)	Topography; vegetation cover; soil type; humidity; wind patterns
Westphal et al., 2009	Modelling (COAMPS)	Meteorological data; land cover; remote sensing; satellite data
Yin et al., 2005	Modelling (DREAM)	Topography; vegetation cover; soil type; humidity; wind patterns

Table 1.5. Forecasting models and the variables used in each study.

Mapping is a second forecasting approach. It draws on geographic information systems together with satellite imagery and data from MODIS and ground sources (Jugder et al., 2018; Kimura, 2018). Its advantage over modelling lies in locating the source area of an event, where surface data are usually sparse (Kimura, 2018). Mapping estimates the potential for an event from variables such as soil erodibility, soil moisture, threshold friction velocity, and threshold wind speed, the lowest wind speed that will move particles (Kimura & Shinoda, 2010). Jugder et al. (2018) recommend mapping the wind-erodible fraction of soil as part of a dust-storm warning system. Producing reliable erodibility data required collecting soil samples across the study area and running hydrological and physical analyses, which, combined with the threshold friction velocity, yielded a classification of soil erodibility across Mongolia from very low to very high.

Threshold wind speed offers another useful parameter. Kimura and Shinoda (2010) mapped its spatial distribution across Mongolia using vegetation-cover data (the Normalized Difference Vegetation Index from MODIS), from which they derived roughness length, then threshold friction velocity, and finally the threshold wind speed itself. The approach is promising but constrained by data quality: vegetation cover varies, its relationship to roughness length is not simple, and data on threshold friction velocity are scarce where vegetation is sparse. Kimura (2018) applied the same method and added a satellite-based aridity index as a proxy for soil moisture. Severe storms corresponded to a threshold wind speed below 10 m/s, moderate events to values between 10 and 15 m/s, and almost no dust above 15 m/s, while an aridity index below 0.03 indicated little dust except at high wind speeds. The implication for forecasting is to insist on good data and to combine several datasets.

A few studies relied on satellite observation alone. Detecting dust this way is difficult, because dust can be hard to separate from water, ice cloud, or bright surfaces. On its own the method is limited, but it gains value when combined with others (Martinez et al., 2009). Samadi et al. (2014) developed the Global Dust Detection Index, which uses MODIS data with the Normalized Difference Dust Index and brightness-temperature difference to detect dust over both land and water and to distinguish it from cloud and bright surfaces. It performed across climatic conditions and worked well on dust events in western Iran. Martinez et al. (2009) tested SEVIRI imagery. SEVIRI is less suited to dust detection than MODIS, which carries an additional short-wavelength channel, but its wide spectral range, spatial resolution, and short repeat cycle make it useful for monitoring. The authors found that using visible channels rather than single pixels improved detection, and that the method works best alongside dust models and in-situ observation.

In-situ networks anchor the rarest approach in the set. Hu et al. (2004) built a dust-storm warning system for the Tarim Basin from station data alone. Drawing on records from more than forty meteorological stations between 1961 and 2001, they graded forecast intensity by the number of stations a storm reached: ten or more marked an extraordinarily serious event, five to nine a serious one, and four or fewer a moderate or light event. Because moderate and light events are frequent, small in scale, and difficult to forecast, the system concentrated on the two severe grades, where the synoptic patterns are clearer and the warnings matter most. The forecasts proved accurate.

Risk knowledge, communication, and response

Three studies addressed the components that turn a forecast into protective action, and they are mostly qualitative (Table 1.2). This is the largest gap in the field. What is missing is work, qualitative and quantitative alike, on what makes a warning effective once a forecast exists: how risk is understood and communicated, how warning messages are framed, and which preventive measures actually reduce harm.

Allahbakhshi et al. (2019) examined what makes a health system well prepared for a dust storm, pointing to risk assessment, knowledge management, organizational structure, and monitoring. Understanding the risk, communicating preparedness clearly, and building public trust all help lower exposure and its health effects, while better knowledge through education, training, and research supports stronger protective measures. The authors also stress a structured protocol, priority for the issue in policy and planning, a recognized body to own it, cross-sector collaboration, and funding for preparation at several levels, alongside accurate prediction and regular drills.

Fatemi et al. (2015) focus on coordination across regional, national, and local levels, which clear policy can support, and on emergency planning for a rapid response, aided by assigning responsibilities to specific organizations. They discuss the value of well-designed risk communication and the role of exercises and drills, which strengthen coordination and improve the response when an event occurs.

Moradian et al. (2015) draw lessons from a 2014 Tehran dust storm that killed five people and injured 44. They emphasize timely warnings that actually reach people, public education on safety measures to support a community-based response, and pre-planned tools such as sirens and routine media broadcasts.

Strengths and limitations

This review reads the dust-storm warning literature against an established model of what a warning system is, rather than cataloguing studies by detection technique. Doing so converts a loose collection of methods into a map of the warning chain, and it shows that the field has built one component in detail while leaving three thinly covered. Both halves of the chain are needed for a system that reduces harm rather than only anticipating the hazard.

The review also documents that the rate of publication has held roughly flat over 22 years, even as dust events have grown more common, and it sets out what the neglected components would require.

Conclusion

Early warning for dust storms has been studied almost entirely as a forecasting problem. Across 22 studies, nineteen worked on detecting or predicting the hazard, and only three touched the risk knowledge, communication, and response that determine whether a forecast ever protects anyone. Detection is the link the field has built; the rest of the chain remains largely unexamined. That imbalance, more than any single result, is what the literature most needs to correct. As dust storms grow more frequent with a warming climate, the case for studying how warnings are understood, communicated, and acted on, alongside how dust is detected and predicted, becomes increasingly difficult to set aside. Closing the gap will require coordinated work from researchers, policymakers, and public health practitioners.

References

Aili, A., & Kim Oanh, N. (2015). Effects of dust storm on public health in desert fringe area: case study of northeast edge of Taklimakan Desert, China. Atmospheric Pollution Research, 6(5), 805–814. https://doi.org/10.5094/apr.2015.089

Allahbakhshi, K., Ghomian, Z., Jazani, R. K., & Khorasani-Zavareh, D. (2019). Preparedness challenges of the Iranian health system for dust and sand storms: a qualitative study. Journal of Education and Health Promotion, 8, 114. https://doi.org/10.4103/jehp.jehp_470_18

Crooks, J., Cascio, W., Percy, M., Reyes, J., Neas, L., & Hilborn, E. (2016). The association between dust storms and daily non-accidental mortality in the United States, 1993–2005. Environmental Health Perspectives, 124(11), 1735–1743. https://doi.org/10.1289/ehp216

Fatemi, F., Moslehi, S., & Ardalan, A. (2015). Preparedness functions in disaster: lesson learned from Tehran dust storm 2014. Natural Hazards, 77(1), 177–179. https://doi.org/10.1007/s11069-015-1601-5

Goudie, A. S., & Middleton, N. J. (2001). Saharan dust storms: nature and consequences. Earth-Science Reviews, 56(1–4), 179–204. https://doi.org/10.1016/S0012-8252(01)00067-8

Guo, Y., An, S., & Comes, T. (2022). From warning messages to preparedness behavior: the role of risk perception and information interaction in the Covid-19 pandemic. International Journal of Disaster Risk Reduction, 73, 102871. https://doi.org/10.1016/j.ijdrr.2022.102871

Hu, L., Wei, R., Wang, X., & Wei, G. (2004). Study on an early warning system and the hierarchical criterions of the forecast intensity of dust storms in the Tarim Basin, Xinjiang (W. Gao & D. Shaw, Eds.; Vol. 5544, pp. 466–474). https://doi.org/10.1117/12.559465

Hyogo Framework for Action 2005–2015: building the resilience of nations and communities to disasters. (2007). Geneva: United Nations Office for Disaster Risk Reduction.

Islam, M., Alharthi, M., & Alam, M. (2019). The impacts of climate change on road traffic accidents in Saudi Arabia. Climate, 7(9), 103. https://doi.org/10.3390/cli7090103

Jugder, D., Gantsetseg, B., Davaanyam, E., & Shinoda, M. (2018). Developing a soil erodibility map across Mongolia. Natural Hazards, 92, 71–94. https://doi.org/10.1007/s11069-018-3409-6

Kallos, G., Spyrou, C., Astitha, M., Mitsakou, C., Solomos, S., Kushta, J., Pytharoulis, I., Katsafados, P., Mavromatidis, E., Papantoniou, N., & Vlastou, G. (2009). Ten-year operational dust forecasting – Recent model development and future plans. IOP Conference Series: Earth and Environmental Science, 7, 012012. https://doi.org/10.1088/1755-1307/7/1/012012

Kanatani, K., Ito, I., Al-Delaimy, W., Adachi, Y., Mathews, W., & Ramsdell, J. (2010). Desert dust exposure is associated with increased risk of asthma hospitalization in children. American Journal of Respiratory and Critical Care Medicine, 182(12), 1475–1481. https://doi.org/10.1164/rccm.201002-0296oc

Kashima, S., Yorifuji, T., Tsuda, T., & Eboshida, A. (2012). Asian dust and daily all-cause or cause-specific mortality in western Japan. Occupational and Environmental Medicine, 69(12), 908–915. https://doi.org/10.1136/oemed-2012-100797

Kim, S., Yumimoto, K., Uno, I., & Chun, Y. (2011). Dust model intercomparison between ADAM and CFORS/Dust for Asian dust case in 2007 (March 28–April 3). SOLA, 7A, 25–28. https://doi.org/10.2151/sola.7A-007

Kimura, R. (2018). Satellite-based mapping of dust erodibility in northeast Asia. Natural Hazards, 92, 19–25. https://doi.org/10.1007/s11069-016-2393-y

Kimura, R., & Shinoda, M. (2010). Spatial distribution of threshold wind speeds for dust outbreaks in northeast Asia. Geomorphology, 114(3), 319–325. https://doi.org/10.1016/j.geomorph.2009.07.014

Lal, R. (2012). Climate change and soil degradation mitigation by sustainable management of soils and other natural resources. Agricultural Research, 1(3), 199–212. https://doi.org/10.1007/s40003-012-0031-9

Lambert, A., Hallar, A. G., Garcia, M., Strong, C., Andrews, E., & Hand, J. L. (2020). Dust impacts of rapid agricultural expansion on the Great Plains. Geophysical Research Letters, 47, e2020GL090347. https://doi.org/10.1029/2020GL090347

Laurent, B., Marticorena, B., Bergametti, G., Tegen, I., Schepanski, K., & Heinold, B. (2009). Modelling mineral dust emissions (J. Perez & J. Baldasano, Eds.; Vol. 7). https://doi.org/10.1088/1755-1307/7/1/012006

Lee, H., Honda, Y., Lim, Y.-H., Guo, Y. L., Hashizume, M., & Kim, H. (2014). Effect of Asian dust storms on mortality in three Asian cities. Atmospheric Environment, 89, 309–317. https://doi.org/10.1016/j.atmosenv.2014.02.048

Lee, H., Jung, J., Myung, W., Baek, J., Kang, J., Kim, D., & Kim, H. (2019). Association between dust storm occurrence and risk of suicide: case-crossover analysis of the Korean national death database. Environment International, 133, 105146. https://doi.org/10.1016/j.envint.2019.105146

Lin, C.-Y., Chiang, M.-L., & Lin, C.-Y. (2016). Empirical model for evaluating PM_{10} concentration caused by river dust episodes. International Journal of Environmental Research and Public Health, 13(6), 553. https://doi.org/10.3390/ijerph13060553

Liu, S.-T., Liao, C.-Y., Kuo, C.-Y., & Kuo, H.-W. (2017). The effects of $PM_{2.5}$ from Asian dust storms on emergency room visits for cardiovascular and respiratory diseases. International

Journal of Environmental Research and Public Health, 14(4), 428. https://doi.org/10.3390/ijerph14040428

Lowe, R., García-Díez, M., Ballester, J., Creswick, J., Robine, J.-M., Herrmann, F., & Rodó, X. (2016). Evaluation of an early-warning system for heat wave-related mortality in Europe: implications for sub-seasonal to seasonal forecasting and climate services. International Journal of Environmental Research and Public Health, 13(2), 206. https://doi.org/10.3390/ijerph13020206

Mahler, A., Thome, K., Yin, D., & Sprigg, W. (2006). Dust transport model validation using satellite- and ground-based methods in the southwestern United States (A. Chu, J. Szykman, & S. Kondragunta, Eds.; Vol. 6299). https://doi.org/10.1117/12.679868

Martinez, M., Ruiz, J., & Cuevas, E. (2009). Use of SEVIRI images and derived products in a WMO Sand and Dust Storm Warning System (J. Perez & J. Baldasano, Eds.; Vol. 7). https://doi.org/10.1088/1755-1307/7/1/012004

Menut, L., Chiapello, I., & Moulin, C. (2009). Previsibility of Saharan dust events using the CHIMERE-DUST transport model (J. Perez & J. Baldasano, Eds.; Vol. 7). https://doi.org/10.1088/1755-1307/7/1/012009

Middleton, N., Tozer, P., & Leys, J. (2019). Sand and dust storms: underrated natural hazards. Disasters, 43(2), 390–409. https://doi.org/10.1111/disa.12320

Moradian, M. J., Rastegarfar, B., Rastegar, M. R., & Ardalan, A. (2015). Tehran dust storm early warning system: corrective measures. PLoS Currents, 7, ecurrents.dis.14f3c645eb2e2003a44c6efd22c23f5e. https://doi.org/10.1371/currents.dis.14f3c645eb2e2003a44c6efd22c23f5e

Nguyen, K., Chuang, Y., Yu, R., & Chen, H. (2021). Developing an ANN-based early warning model for airborne particulate matters in river banks areas. Expert Systems with Applications, 183, 115421. https://doi.org/10.1016/j.eswa.2021.115421

Park, S.-U., Choe, A., & Park, M.-S. (2013). A simulation of Asian dust events observed from 20 to 29 December 2009 in Korea by using ADAM2. Asia-Pacific Journal of Atmospheric Sciences, 49(1), 95–109. https://doi.org/10.1007/s13143-013-0011-4

Rublee, C. S., Sorensen, C. J., Lemery, J., Wade, T. J., Sams, E. A., Hilborn, E. D., & Crooks, J. L. (2020). Associations between dust storms and intensive care unit admissions in the United States, 2000–2015. GeoHealth, 4(8), e2020GH000260. https://doi.org/10.1029/2020GH000260

Samadi, M., Darvishi Boloorani, A., Alavipanah, S. K., Mohamadi, H., & Najafi, M. S. (2014). Global Dust Detection Index (GDDI): a new remotely sensed methodology for dust storms detection. Journal of Environmental Health Science & Engineering, 12(1), 20. https://doi.org/10.1186/2052-336X-12-20

Schulz, M., Cozic, A., & Szopa, S. (2009). LMDzT-INCA dust forecast model developments and associated validation efforts (J. Perez & J. Baldasano, Eds.; Vol. 7). https://doi.org/10.1088/1755-1307/7/1/012014

Sendai Framework for Disaster Risk Reduction 2015–2030. (2015). UN World Conference on Disaster Risk Reduction, 14–18 March 2015, Sendai, Japan. Geneva: United Nations Office for Disaster Risk Reduction.

Stefanski, R., & Sivakumar, M. (2009). Impacts of sand and dust storms on agriculture and potential agricultural applications of an SDSWS (J. Perez & J. Baldasano, Eds.; Vol. 7). https://doi.org/10.1088/1755-1307/7/1/012016

Vimic, A., Cvetkovic, B., Giannaros, T., Kashani, S., Kotroni, V., Pejanovic, G., Nickovic, S., Basart, S., Terradellas, E., Shahbazi, R., Prieto, J., Lagouvardos, K., Petkovic, S., Mandic, M., & Boloorani, A. (2021). Numerical simulation of Tehran dust storm on 2 June 2014: a case study of agricultural abandoned lands as emission sources. Atmosphere, 12(8), 1054. https://doi.org/10.3390/atmos12081054

Vukovic, A., Vujadinovic, M., Pejanovic, G., Andric, J., Kumjian, M., Djurdjevic, V., Dacic, M., Prasad, A., El-Askary, H., Paris, B., Petkovic, S., Nickovic, S., & Sprigg, W. (2014). Numerical simulation of "an American haboob." Atmospheric Chemistry and Physics, 14(7), 3211–3230. https://doi.org/10.5194/acp-14-3211-2014

Watanabe, M., Noma, H., Kurai, J., Shimizu, A., Sano, H., Kato, K., et al. (2015). Association of sand dust particles with pulmonary function and respiratory symptoms in adult patients with asthma in western Japan using light detection and ranging: a panel study. International Journal of Environmental Research and Public Health, 12(10), 13038–13052. https://doi.org/10.3390/ijerph121013038

Westphal, D., Curtis, C., Liu, M., & Walker, A. (2009). Operational aerosol and dust storm forecasting (J. Perez & J. Baldasano, Eds.; Vol. 7). https://doi.org/10.1088/1755-1307/7/1/012007

Yin, D., Nickovic, S., Barbaris, B., Chandy, B., & Sprigg, W. A. (2005). Modeling wind-blown desert dust in the southwestern United States for public health warning: a case study.

Atmospheric Environment, 39(33), 6243–6254.
https://doi.org/10.1016/j.atmosenv.2005.07.009

Zauli Sajani, S., Miglio, R., Bonasoni, P., Cristofanelli, P., Marinoni, A., & Sartini, C., et al. (2010). Saharan dust and daily mortality in Emilia-Romagna (Italy). Occupational and Environmental Medicine, 68(6), 446–451. https://doi.org/10.1136/oem.2010.058156

Chapter 2. A Statewide Daily Dust-Attributable Particulate Matter Exposure Surface for California: Integrated Machine-Learning Estimation and External Validation, 2006–2019

Introduction

Wind-blown mineral dust is a recurring but unevenly characterized threat to respiratory and cardiovascular health, and the exposure tools available for population health research rarely match the spatial and temporal scale at which dust actually varies. This chapter develops and evaluates a daily, statewide exposure surface for dust-attributable particulate matter across California, built so that the resulting estimates can be linked directly to health outcomes in the epidemiologic analysis that follows later in this dissertation.

Dust contributes substantially to coarse particulate matter, the size fraction that fine-particle monitoring frameworks capture least well. Reviews of the dust and health literature report associations with respiratory and cardiovascular morbidity and with non-accidental mortality, while repeatedly noting that the evidence is complicated by heterogeneous exposure definitions, varied outcome settings, and inconsistent characterization of what counts as dust across studies (Zhang et al., 2016; Crooks et al., 2016; Lwin et al., 2023; Rowan et al., 2024; Al-Dousari et al., 2025). Because the toxicologically relevant properties of particulate matter depend on composition as well as mass (Morakinyo et al., 2016), dust is not only a public health concern but also a measurement problem, representing a spatially heterogeneous, intermittently transported, compositionally mixed pollutant when direct observation is sparse and no single instrument captures the full signal. Process-level characterization of atmospheric dust reinforces this difficulty. Source strength is governed less by the mere presence of erodible land than by the surface environment and the meteorology that mobilizes it, so that emission is episodic and source-specific (Prospero et al., 2002), and the mineralogy, particle-size distribution, and elemental composition of the resulting aerosol vary markedly among source regions and along transport paths (Formenti et al., 2011). A modeled exposure surface for a region as heterogeneous as California must therefore reconstruct a moving target whose physical and chemical identity is not fixed, which is precisely why integration across meteorological, particulate, optical, and land-surface signals is needed rather than reliance on any single dust indicator.

That measurement problem shapes epidemiologic interpretation. Methodological reviews show that dust and health findings depend materially on how dust exposure is defined, monitored, and operationalized for analysis (Querol et al., 2019; Tobias & Stafoggia, 2020; Hohsfield et al., 2023). In practice, dust has been represented by event indicators, storm databases, monitor exceedances, compositional surrogates, aerosol optical products, transport models, or hybrid definitions. These constructs are not interchangeable. Each captures a different part of the underlying atmospheric phenomenon, and each carries a different mix of spatial coverage, temporal continuity, and physical specificity. Two studies can therefore differ in their estimated effects not only because populations or outcomes differ, but because the underlying exposure construct differs. When dust is represented by sparse monitors, coarse event flags, or incomplete surrogates, true dust conditions are measured imprecisely across person-time, and the resulting exposure error tends to attenuate health associations. This is why better exposure modeling is needed for reducing a specific, direction-predictable source of bias in downstream health analyses as well as for providing a continuous measure of dust that can provide better assessment on its dose-response relationship with health outcomes.

California is an important place to confront this problem. The state combines steep gradients in aridity, land cover, topography, transport pathways, agriculture, and human settlement, which together produce a heterogeneous environment for dust generation, movement, and exposure. Long-term climatology established the western United States as a meaningful region for dust reconstruction from routine monitoring (Tong et al., 2012), and later work showed that dust events in this region are best identified by integrating meteorological, satellite, and air-quality information rather than any single indicator (Lei & Wang, 2014). Regional analyses further demonstrate that dust hazard is governed by meteorological variability and transport conditions, not only by local source presence (Achakulwisut et al., 2017; Bhattachan et al., 2019). California-specific studies highlight intense dust and PM_{10} problems around the Salton Sea basin and other anthropogenically disturbed landscapes (Parajuli & Zender, 2018; Frie et al., 2019; Sinclair et al., 2024; Adebiyi et al., 2025). The state is therefore not an arbitrary application domain but a scientifically demanding and policy-relevant test bed for dust exposure assessment.

However, a gap separates dust-process knowledge from population-ready exposure assignment. Regional research has sharpened dust climatology, event identification, source characterization, and hazard framing across the western and southwestern United States (Tong et al., 2012; Lei & Wang, 2014; Kandakji et al., 2020; Rowan et al., 2024). None of that work, on its own, yields a daily statewide surface that can assign dust-related particulate conditions to every populated location in California. The distinction matters because epidemiologic analysis generally requires exposure across all person-time, not only on days that satisfy a strict event definition. The field has improved at flagging dust episodes than at generating continuous, population-relevant dust exposure surfaces.

A second obstacle is the absence of a true statewide gold standard for dust exposure. This is a field-wide condition rather than a limitation of any single study. Dust is difficult to observe comprehensively because the relevant construct sits at the intersection of particulate mass, mineral composition, source environment, transport, and measurement scale (Querol et al., 2019; Hohsfield et al., 2023). In this setting, dust reconstructed from Interagency Monitoring of Protected Visual Environments (IMPROVE) composition data is the strongest available external comparator, it is composition-based, geographically distributed, and directly relevant to mineral-dust-containing mixtures. It is also sparse, monitor-bound, and derived from a composition formula that is related to, but not identical with, a modeled prediction surface. The practical implication is not that validation is impossible, but that it must be read as evaluation against the best available comparator rather than against an ideal truth.

These conditions motivate an integrated modeling strategy. In a state as varied as California, dust-related exposure is unlikely to be captured by any single data stream, because the signal reflects meteorology, modeled particulate fields, aerosol optical depth, land-surface context, inversion structure, documented events, and transport conditions at once. The broader exposure literature supports integration: high-resolution meteorological products such as gridMET supply important environmental covariates (Abatzoglou, 2013), combined satellite-model-monitor frameworks have improved particulate estimation elsewhere (Van Donkelaar et al., 2019), with stacked machine-learning ensembles in particular now a leading approach to high-resolution particulate exposure assignment for epidemiology (Di et al., 2019), and recent machine-learning studies demonstrate active progress in high-resolution PM_{10} and coarse-particle prediction (Shi et al., 2024; Cusano et al., 2025; Gora et al., 2025; Zhang et al., 2024). For dust, integration is important, because it is a direct response to the multi-signal structure of the problem.

This chapter constructs and evaluates a daily, statewide California surface of dust-attributable PM_{10} (DAF-PM_{10}) at population-weighted ZIP-code centroids for 2006 through 2019, built from integrated meteorological, land-surface, satellite, modeled air-quality, and dust-indicator predictors and evaluated against IMPROVE measured dust. A transparent, rule-based hierarchical scoring surface is developed in parallel as an interpretable benchmark against which the machine-learning surface can be compared. The contribution of this work is specific and clearly bounded. The work moves from event-only characterization toward continuous statewide assignment, it places exposure at population-weighted centroids aligned with downstream health linkage, it embeds external evaluation in the workflow instead of treating the model as self-justifying, and it retains a parallel interpretable surface for context. The objective is to produce an exposure surface whose rank ordering of dust conditions across space and time is defensible enough to support epidemiologic analysis in a domain where such surfaces are genuinely hard to build. The sections that follow describe the data, feature construction, modeling, and validation used to reach that objective.

Methods

The exposure surface was assembled in four stages: construction of population-weighted spatial units, integration of multi-source daily predictors, supervised ensemble modeling of PM_{10} with a parallel rule-based scoring scheme, and external validation against an independent composition-based dust reference. Each stage is described below, with attention to the design choices that bear most directly on how the surface should be interpreted.

Study Design and Setting

The study covered the State of California for 2006 through 2019. The spatial unit of analysis was the population-weighted centroid of each ZIP Code Tabulation Area (ZCTA), giving 1,797 centroid locations. Following methods similar to those of Villanueva et al. (2026), population-weighted centroids were computed from the U.S. Environmental Protection Agency 2020 EnviroAtlas dasymetric population raster at 30 m resolution combined with Census 2020 ZCTA boundaries. Weighting each centroid by residential population, rather than using the geographic center of the polygon, places the exposure point closer to where people actually live. The distinction is consequential for large, irregular ZIP codes whose geographic center can fall in unpopulated terrain, and it is the appropriate basis for assigning exposure that will later be linked to residents. The 2006 through 2019 window defines the exposure-model development period for this chapter; any narrower interval used in a subsequent health analysis is a property of that application rather than of the surface itself. The study design and the external-validation network are shown together in Figure 2.1: population-weighted centroids (blue) mark the locations at which daily exposure was estimated, and the IMPROVE monitors (red) provide the independent, composition-based dust reference used for external validation.

Figure 2.1. Study design: statewide exposure assignment and the external validation network. Population-weighted ZIP-code centroids (blue) define the locations at which daily dust-attributable PM_{10} was estimated, and IMPROVE monitoring sites (red) provide the independent composition-based dust measurements used for external validation. IMPROVE monitors sit mainly in national parks, wilderness areas, and other protected lands, so the geographic offset between populated centroids and remote reference monitors is itself a central challenge for validating residential exposure estimates.

Two further maps document the spatial units. Figure 2.2 shows all ZCTA boundaries with their population-weighted centroids, illustrating the concentration of centroids in the coastal urban corridors of the San Francisco Bay Area, the Los Angeles Basin, and San Diego. Figure 2.3 makes the rationale for population weighting concrete: for several large Central Valley and Bay Area ZIP codes, the population-weighted centroid is displaced from the geographic centroid by roughly 15 to 21 km, often shifting the exposure point out of agricultural or unpopulated land and toward the residential core.

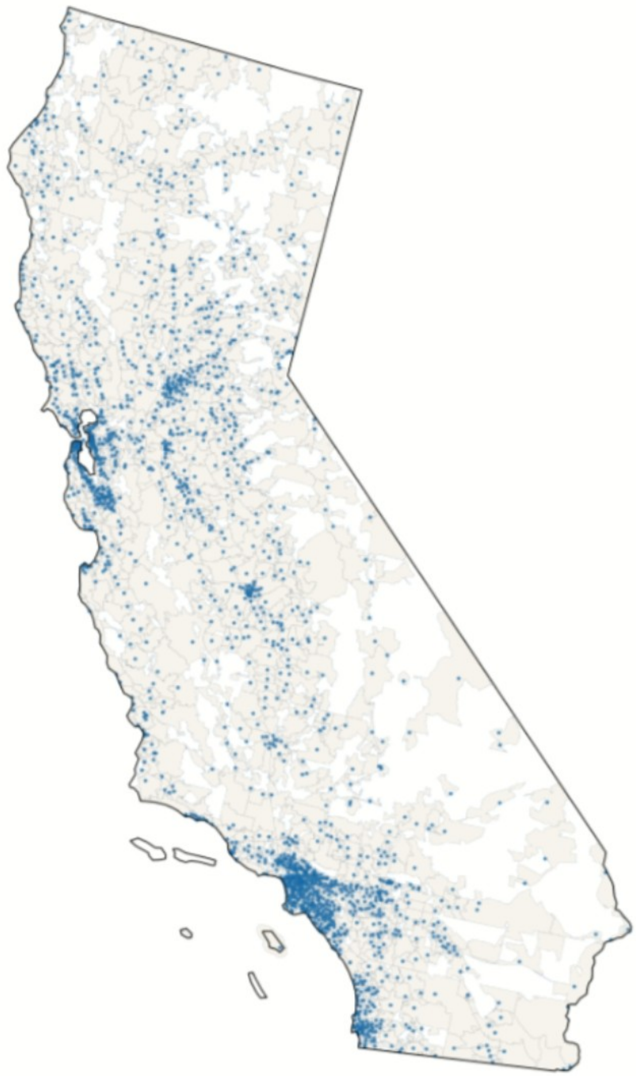

Figure 2.2. Spatial distribution of California ZIP Code Tabulation Area boundaries (gray polygons) and their population-weighted centroids (blue points). Centroid density tracks the state population, concentrating along the coastal urban corridors.

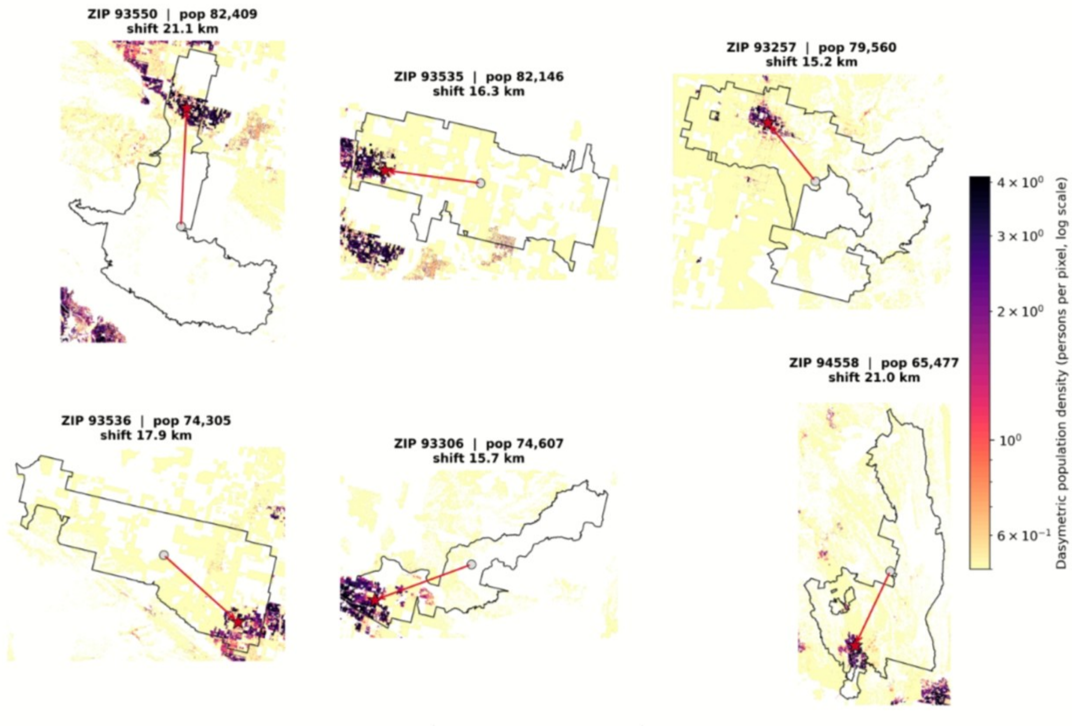

Figure 2.3. Dasymetric population density and centroid displacement for six representative California ZIP codes. Background shading shows population density on a logarithmic scale; red segments connect each geographic centroid (open circle) to its population-weighted centroid (red cross). Displacements of roughly 15 to 21 km show how a geographic center can fall far from where residents live, motivating population-weighted assignment for exposure estimation.

Data Sources

Predictors and reference data were drawn from ground monitoring, gridded meteorology, a chemical transport model, satellite retrievals, atmospheric soundings, an aeolian dust-flux model, and land-surface products. Unless noted otherwise, all variables were resolved to the daily, centroid level for 2006 through 2019.

Ground PM_{10} monitoring

Daily 24-hour average PM_{10} concentrations were obtained from the U.S. EPA Air Quality System for California. Observations from 2006 through 2019 across 190 monitoring sites served as the response variable for supervised modeling.

Meteorology

Gridded daily meteorology came from gridMET (Abatzoglou, 2013), extracted at centroid locations: precipitation, minimum and maximum temperature, downward solar radiation, minimum and maximum relative humidity, wind speed, and wind direction. Eastward and northward wind components were derived from speed and direction. Raster-to-point extraction used the terra package in R, with yearly outputs combined across the study period and elevation appended.

Modeled air quality

Surface PM_{10} and $PM_{2.5}$ came from the EPA EQUATES configuration of the Community Multiscale Air Quality (CMAQ) model (Byun & Schere, 2006) on a 12 km grid for 2002 through 2019, delivered as 216 monthly files for California. CMAQ PM_{10} entered the model as a physically based predictor and also served as the carrier field for the hierarchical scoring surface described below.

Documented dust events

Historical dust storm records were drawn from the NOAA National Weather Service Storm Events Database for California and used to build a binary dust indicator. Records were cleaned, dates parsed, and coordinates standardized. Missing event coordinates were imputed by matching NWS forecast-zone codes to zone-polygon centroids; thirteen zones required manual geocoding, and one zone (CA183) could not be resolved and was retained as missing.

Aerosol composition (IMPROVE)

Aerosol composition data, including PM_{10}, $PM_{2.5}$, and elemental species, were obtained from the IMPROVE network for California sites, with site metadata from the IMPROVE inventory. Measured soil dust was computed from elemental data using the standard IMPROVE soil formula, Dust = 2.2 Al + 2.49 Si + 1.63 Ca + 2.42 Fe + 1.94 Ti (Malm et al., 1994).

IMPROVE measured dust is the external comparator for validation and is not used at any stage of model training.

Visibility and atmospheric soundings

Hourly visibility from NOAA Integrated Surface Database stations supplied a daily minimum visibility field. Twice-daily radiosonde profiles from Oakland, Vandenberg, and San Diego/Miramar were used to derive temperature-inversion strength and base height, computed as the maximum temperature increase between adjacent vertical levels below 3,000 m.

Satellite aerosol optical depth

Aerosol optical depth (AOD) at 470 and 550 nm was extracted daily at centroids from the MODIS MCD19A2 (Collection 6.1) MAIAC product (Lyapustin et al., 2018) in Google Earth Engine at 500 m scale, scaled to unitless values. The dual-wavelength retrievals also yielded the Angstrom exponent, AE = -log(AOD470 / AOD550) / log(470 / 550), where lower values indicate coarser particles consistent with mineral dust.

Aeolian dust flux

Monthly horizontal and vertical dust-flux estimates were obtained from a wind-erosion (WEMO) modeling framework and joined to the daily table by year-month. These flux fields were generated by Bo using Alex Hall's winds for California product, wind from twenty years of monthly values driven by high-resolution downscaled California wind fields. Files for 2008 were not produced and were carried as missing, it was confirmed by Bo that it was a genuine source-data gap for that year rather than an archival omission.

Land surface and spatial attributes

Normalized Difference Vegetation Index (NDVI), land-cover class fractions, elevation, county, air basin, distance to nearest highway, and distance to coast were extracted at centroids in Google Earth Engine. NDVI used MODIS surface reflectance; land-cover fractions were computed within 1 km buffers from the National Land Cover Database; elevation used the 30 m SRTM product; distance to highway used TIGER 2016 interstate and U.S. routes; county used TIGER 2018 boundaries; and air-basin and coastal-zone membership used California agency boundary layers.

Feature Construction and Spatial Integration

Several predictors required spatial harmonization to the centroid grid. CMAQ PM_{10} and $PM_{2.5}$ were interpolated from the 12 km model grid to centroids by inverse-distance weighting of the four nearest cells, with a parallel interpolation to AQS monitor locations for the training data. Daily minimum visibility was interpolated to centroids by inverse-distance weighting (power = 2) using the gstat package. Each centroid was assigned inversion metrics from its nearest radiosonde station. Daily specific humidity was computed from gridMET temperature, relative humidity, and elevation, using the Magnus saturation-vapor-pressure formulation under FAO-56 conventions with an elevation-based station-pressure approximation.

Dust-dominated days at IMPROVE sites were identified with the Tong approach (Tong et al., 2012), Ward's D-squared hierarchical clustering (k = 6) applied to scaled composition variables, including PM_{10}, $PM_{2.5}$, the $PM_{2.5}/PM_{10}$ ratio, elemental species, sulfate, nitrate, organic and elemental carbon, and enrichment factors. This classification provides an independent dust label for discrimination analyses.

The documented-event indicator was constructed at two spatial scales for two distinct purposes. At the centroid level, a 20 km buffer flagged each centroid-day on which any NWS dust event fell within range, with multi-day events expanded across their full duration, this buffer supports the statewide exposure and validation surface. At the monitor level, NWS events were linked to AQS PM_{10} sites by nearest-neighbor matching within 30 km to build the supervised training indicator; a 30 km radius was used because tighter buffers left most AQS monitoring sites, which are few and spatially clustered, with little or no dust signal. The two buffers serve the prediction surface and the training linkage, respectively.

The supervised training table comprised 1,040,991 observations for 190 AQS monitoring sites. A parallel ZIP-centroid master predictor file, joining meteorology, elevation, specific humidity, AOD, NDVI, land cover, inversion metrics, centroid coordinates, CMAQ PM_{10}, and WEMO flux, was used for statewide daily prediction.

Machine-Learning Ensemble

Models were developed in the H2O framework. Three base learners were trained on the monitor PM_{10} response: a random forest, a gradient boosting machine (GBM), and a feedforward deep-learning network. The training table was split into training (70%), validation (15%), and test (15%) partitions under a fixed seed. Predictor missingness, including the 2008 WEMO gap and intermittent remotely sensed gaps, was carried into H2O and handled by its native treatment: tree-based learners route missing values through surrogate splits, and the deep-learning model applies mean imputation. Thirty predictors spanned meteorology, land surface, satellite AOD, modeled CMAQ PM_{10}, dust indicators (NWS flag, WEMO horizontal and vertical flux), atmospheric inversion structure, and spatial identifiers. The three learners were chosen to span complementary inductive biases rather than to compete: the random forest contributes decorrelated, variance-stabilizing spatial averaging (Breiman, 2001); the gradient boosting machine sequentially refines its fit to the meteorology and modeled-PM_{10} structure and so carries the sharpest but highest-variance signal (Friedman, 2001); and the feedforward network adds smooth, land-cover-sensitive interactions. Combining learners of this kind is a deliberate response to the bias-variance trade-off, because a stacked combination can generalize better than any single constituent by averaging away the idiosyncratic error of each (Wolpert, 1992; Dietterich, 2000). This is the logic of the Super Learner framework now standard in environmental-exposure modeling, in which a cross-validated stack is guaranteed asymptotically to perform at least as well as its best base learner (van der Laan et al., 2007), and it is the same architecture that underlies the closest methodological precedent for the present work, the national ensemble particulate-matter exposure surface of Di et al. (2019).

Hyperparameters were tuned by grid search with 10-fold cross-validation for each learner. The grids declared requested search ceilings, for example up to 1,000 trees for the tree-based learners and up to 500 epochs for the network. Because H2O early stopping halts fitting once out-of-sample error plateaus, the effective fitted models were considerably smaller than those ceilings: under the grid configuration (ten rounds; tolerance 0.01 for the tree learners, H2O defaults for the network), random forests stopped between 25 and 99 trees, the best GBM stopped at 107 trees, and the network stopped after roughly six epochs. The final production learners were refit with stricter stopping (five rounds; tolerance 0.001) and then combined. The declared grid values are therefore requested maxima, not the sizes of the retained models.

The three learners were combined into a stacked ensemble whose metalearner was the H2O default generalized linear model. The saved ensemble weights, an intercept with one coefficient per learner, are consistent with a linear combination of base predictions and are reported in the Results. Fitting a generalized linear metalearner to the cross-validated base predictions is the regression form of stacked generalization (Breiman, 1996), the resulting coefficients function as data-driven weights that allocate influence across learners according to their out-of-sample skill, not as physical parameters of dust generation.

Estimating Dust-Attributable PM_{10}

Two surfaces were produced, and they estimate related but distinct constructs. The ensemble surface was trained on monitor PM_{10} and then applied to the ZIP-centroid master file to generate daily predictions, exported under the DAF-PM_{10} label with negative predictions truncated to zero. The workflow does not contain a separate step that partitions total PM_{10} into dust and non-dust components before export. The ensemble DAF-PM_{10} is therefore best understood as a dust-informed PM_{10} prediction surface whose relevance to dust is established by external validation against IMPROVE measured dust rather than asserted by its label. This framing is carried consistently through the Results and Discussion.

The hierarchical surface is a transparent, rule-based alternative built on CMAQ PM_{10}. A two-tier scoring scheme assigns points to a centroid-day according to dust-relevant conditions. A documented NWS dust event (Tier 1) contributes 25 points. Tier 2 adds composite indicators: PM_{10} more than two standard deviations above its seasonal mean (7 points), a $PM_{2.5}/PM_{10}$ ratio at or below 0.4 (6 points), AOD above 0.5 together with an Angstrom exponent below 0.5 (4 points), wind speed above 6 m/s (3 points), visibility below 3 km (3 points), and relative humidity below 40% (2 points). The total score is normalized to a 0 to 1 scale by dividing the points earned by 25, the maximum attainable, and the normalized score is multiplied by modeled PM_{10} to yield the hierarchical DAF-PM_{10}. The scheme is intentionally interpretable: every value can be traced to an explicit combination of triggered rules.

Validation

Internal performance was evaluated on the held-out test partition using the coefficient of determination, root-mean-square error (RMSE), and mean absolute error (MAE). Because monitor-days are spatially clustered, random partitioning of the training table can yield an optimistic internal estimate relative to performance at unseen locations (Roberts et al., 2017); the leave-one-site-out external analysis reported below is therefore the more demanding test of spatial generalization, and the two should be read together. External validation compared both surfaces with IMPROVE measured soil dust, the strongest independent comparator available in this setting. Centroid predictions were matched to IMPROVE sites, and agreement was assessed pooled across sites, stratified by site, and stratified by monitor-to-centroid distance. Two matching schemes were used. The full-validation scheme retained all centroid-to-site matches within the domain. The closest-centroid scheme retained only the nearest centroid per IMPROVE site and reported performance at progressively larger distance thresholds (5, 10, 15, and 20 km). The two schemes differ in sample composition and spatial-match quality and so yield different metrics by design. Discrimination of high-dust days was quantified by the area under the receiver operating characteristic curve (AUC) and the area under the precision-recall curve (AUPRC), using both the 95th percentile of IMPROVE measured dust and the Tong dust classification as event definitions.

Because the surfaces and the comparator are on different scales, raw-scale bias is expected and is not the primary criterion. Post-calibration aligned each surface to IMPROVE measured dust on the log(1 + x) scale. A random-intercept mixed model with a site-level random effect was fit.Predictions were back-transformed with Duan smearing (Duan, 1983) to correct retransformation bias. This model captures site-level baseline offsets while borrowing strength across sites.

Summary Statistics, Software, and Contributions

Descriptive summaries, criteria-trigger frequencies, seasonal and annual patterns, extreme-event characterization, and correlation analyses were computed for the statewide output.

Temporal trend was assessed by linear regression of monthly mean DAF-PM_{10} on time, and the trend result should be read as a linear-model summary only. All processing, interpolation, feature construction, and modeling were performed in R. Principal packages included h2o, terra, sf, gstat, lme4, data.table, dplyr, furrr and future, lubridate, readxl, httr and readr, and geosphere. Remote-sensing extraction used Google Earth Engine. The WEMO dust-flux fields were contributed by a collaborator, as noted in the acknowledgments; all modeling, validation, and interpretation reported here are the author's own.

Results

Results are organized around scientific weight rather than analysis order. The two surfaces and the statewide domain are described first, followed by internal model performance, the external validation that establishes the dust relevance of the ensemble surface, post-calibration, the statewide exposure burden, and finally the interpretable hierarchical surface and its comparison with the ensemble.

Two Exposure Surfaces and the Statewide Domain

The workflow produced daily DAF-PM_{10} estimates for 1,797 California ZIP-code centroids across 2006 through 2019, a total of 9,188,061 centroid-day records (1,797 centroids over 5,113 days). Two surfaces were generated over this identical domain: the machine-learning ensemble and the rule-based hierarchical score. They differ sharply in magnitude, and the difference is structural rather than a discrepancy to reconcile. At every reported percentile the ensemble is higher and more broadly distributed than the hierarchical score, with a mean of 17.51 µg/m³ against 2.14 µg/m³ and a 95th percentile of 38.80 against 7.28 µg/m³ (Table 2.2). The ensemble predicts a dust-informed PM_{10} surface on the scale of ambient PM_{10}, whereas the hierarchical score multiplies modeled PM_{10} by a bounded dust multiplier and therefore returns a constrained dust subset. The external validation framework, not the magnitude of either surface, is the appropriate arbiter of dust-specific performance.

Against a CMAQ mean total PM_{10} of 15.2 µg/m³, the hierarchical mean of 2.14 µg/m³ corresponds to a mean dust fraction near 13.3%, consistent with expectations for windblown and resuspended mineral dust in arid and semi-arid environments.

Table 2.1. Statewide descriptive statistics for total ambient PM_{10} and the hierarchical DAF-PM_{10} surface across California ZIP-centroid days, 2006 through 2019.

Statistic	Value
Total observations (centroid-days)	9,188,061
Mean ambient PM_{10} (µg/m³)	15.2
Mean DAF-PM_{10}, hierarchical (µg/m³)	2.14
Maximum DAF-PM_{10}, hierarchical (µg/m³)	2,994
Mean dust fraction (DAF / ambient)	13.3%

Note. DAF-PM_{10} = dust-attributable fraction of PM_{10}. Values summarize the interpretable hierarchical surface; the extreme maximum reflects episodic high-intensity events concentrated in fall and summer.

Table 2.2. Distribution of DAF-PM_{10} estimates by modeling approach across the full statewide domain.

Approach	Obs_Count	Mean	SD	P50	P75	P90	P95	P99	P99.9	Max
Hierarchical	9,188,061	2.14	3.72	1.17	2.54	5.34	7.28	15.45	24.77	2994.24
ML ensemble	9,188,061	17.51	12.87	15.61	24.81	33.04	38.80	55.76	105.27	296.53

Note. All values in µg/m³. Percentiles are of the full statewide distribution across 9,188,061 centroid-days.

Ensemble Training and Internal Validation

Among the base learners, the gradient boosting machine (GBM) was strongest on the held-out validation partition (R-squared = 0.649, RMSE = 14.14 µg/m³), ahead of the random forest (0.609; 14.94) and the deep-learning network (0.454; 17.65) (Table 2.3). Cross-validation mean-square error told the same story, lowest for the GBM (212.7, RMSE about 14.6 µg/m³) relative to the random forest (228.3; about 15.1) and the network (261.1; about 16.2). The wide GBM gap between training and validation R-squared (0.983 versus 0.649) marks a high-variance learner, which is the behavior the ensemble is built to temper.

Table 2.3. Internal validation performance of the three machine-learning base learners.

Metric	GBM	Random forest	Deep learning
Validation RMSE (μg/m³)	14.14	14.94	17.65
Validation R-squared	0.649	0.609	0.454
Training R-squared	0.983	0.600	0.535
Leading predictors	County, CMAQ, temperature	County, latitude, longitude	Land cover
Bias-variance profile	High variance, high accuracy	Low variance, stable	Low bias, diverse features

Note. All learners were trained on the same PM_{10} response and 30-predictor set. Training metrics are in-sample (out-of-bag for the random forest) validation metrics are on a held-out 15% partition.

The stacked ensemble improved modestly but consistently on the best single learner, reaching validation R-squared = 0.652 and RMSE = 14.10 μg/m³, with a held-out test R-squared of 0.636 (Table 2.4). The close agreement between validation and test performance indicates that the held-out estimate is stable rather than an artifact of tuning. Figure 2.4 shows the observed-versus-predicted relationship: the surface tracks the bulk of daily variation (R-squared about 0.64) but under-predicts the highest observed concentrations above roughly 200 μg/m³, the expected signature of episodic extremes that are hard to reconstruct.

Table 2.4. Performance of the stacked ensemble relative to the individual base learners.

Metric	GBM	Random forest	Deep learning	Stacked ensemble
Validation R-squared	0.649	0.609	0.454	0.652
Validation RMSE (μg/m³)	14.14	14.94	17.65	14.10
Test-set R-squared	–	–	–	0.636

Note. The ensemble was built by stacked generalization with a generalized linear metalearner. The test-set value is from a fully held-out 15% partition not used in training or tuning.

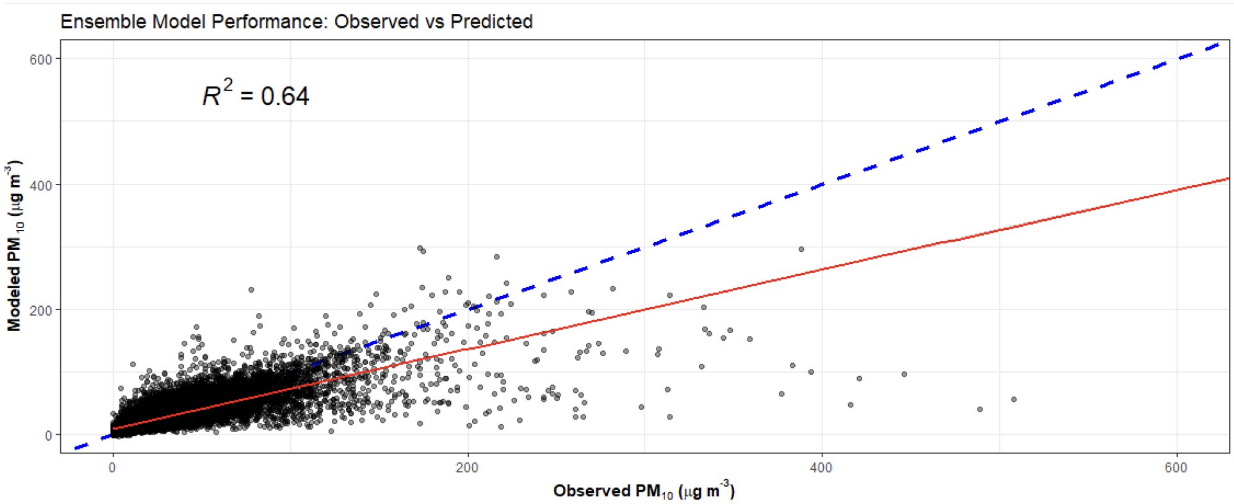

Figure 2.4. Ensemble internal validation: observed versus predicted daily PM_{10}. Each point is a monitor-day, the dashed line is the 1:1 reference and the solid line the fitted regression (R-squared about 0.64). The ensemble captures most daily variability but under-predicts the highest observed concentrations (above roughly 200 µg/m³), consistent with the difficulty of reconstructing episodic extreme dust.

The metalearner weighted the GBM most heavily (beta about 0.76), with smaller contributions from the random forest (about 0.20) and the network (about 0.09) over an intercept of -1.247 (Table 2.5). The weighting follows the cross-validation ranking: the GBM carries the core meteorology-CMAQ-PM_{10} structure, the random forest stabilizes its variance through decorrelated spatial averaging, and the network adds land-cover-sensitive corrections.

Table 2.5. Metalearner coefficients and functional roles of the base learners in the stacked ensemble.

Base learner	Weight (beta)	Functional role
Gradient boosting	≈ 0.76	Primary predictor; core meteorology-CMAQ-PM_{10} structure
Random forest	≈ 0.20	Spatial stabilizer; offsets GBM variance
Deep learning	≈ 0.09	Specialized land-cover correction

Note. Coefficients are from the generalized linear metalearner (intercept -1.247).

Predictor Importance

Across learners, spatial and meteorological predictors dominated. In the random forest, the five most important predictors were county (22.6% scaled importance), site latitude (10.2%), site longitude (8.7%), maximum temperature (7.1%), and CMAQ PM_{10} (6.1%). The GBM spread importance more evenly across county (12.5%), CMAQ PM_{10} (10.3%), maximum temperature (10.1%), wind speed (8.1%), and specific humidity (7.1%). In the network, importance was distributed across 98 encoded features, led among continuous predictors by land-cover fractions, inversion base height, and maximum temperature. Explicit dust indicators were consistently low-importance: the NWS flag contributed roughly 0.5 to 0.8% across learners, and the WEMO flux variables sat in the lower tier. This pattern shapes interpretation. The ensemble did not lean on dust-specific flags, it learned a broader dust-relevant particulate signal from spatial structure, temperature, modeled PM_{10}, and related predictors. What establishes the dust relevance of that signal is the external validation that follows.

External Validation Against IMPROVE Measured Dust

External validation is the center of this chapter, because it tests whether the modeled surfaces order real dust conditions correctly against an independent composition-based reference that never entered training.

Pooled agreement

Against IMPROVE measured dust, the ensemble achieved a pooled Spearman rho of 0.646 and Pearson r of 0.603, with RMSE 11.99 µg/m³ and MAE 8.34 µg/m³ across roughly 22,000 matched site-days (Table 2.6). The hierarchical surface was weaker on every continuous metric (rho = 0.456; r = 0.284). Rank-order agreement is the more informative criterion here, because the downstream epidemiologic design depends on ordering exposure correctly across space and time rather than on reproducing absolute concentrations.

Table 2.6. Pooled continuous agreement between modeled DAF-PM_{10} and IMPROVE measured dust.

Model	N	Spearman rho	Pearson r	RMSE	MAE	Bias
Ensemble	≈ 22,000	0.646	0.603	11.99	8.34	+4.65
Hierarchical	≈ 22,000	0.456	0.284	16.03	10.62	-10.55

Note. RMSE, MAE, and bias in μg/m³. The sample is all available centroid-to-site matches (full-validation scheme) across 2006 through 2019.

Distance sensitivity

Agreement depended strongly on how far each IMPROVE monitor sat from its matched centroid. Under closest-centroid matching, ensemble Spearman rho declined monotonically from 0.670 at 5 km to 0.625, 0.605, and 0.587 at 10, 15, and 20 km; the hierarchical surface followed the same shape at a lower level, from 0.543 to 0.343 (Table 2.7). Figure 2.5 shows the same decline by distance bin, with the ensemble remaining above the hierarchical surface at every separation. The graded pattern is itself informative: a meaningful share of apparent model-measurement disagreement is geometric, arising from the offset between populated centroids and remote monitors, rather than from model error.

Table 2.7. Rank correlation between modeled DAF-PM_{10} and IMPROVE measured dust by monitor-to-centroid distance (closest-centroid matching).

Distance buffer	Ensemble (Spearman rho)	Hierarchical (Spearman rho)
≤ 5 km	0.670	0.543
≤ 10 km	0.625	0.461
≤ 15 km	0.605	0.341
≤ 20 km	0.587	0.343

Note. Measured dust computed from IMPROVE elemental composition via the IMPROVE soil formula. Distance is from each monitor to its nearest population-weighted centroid.

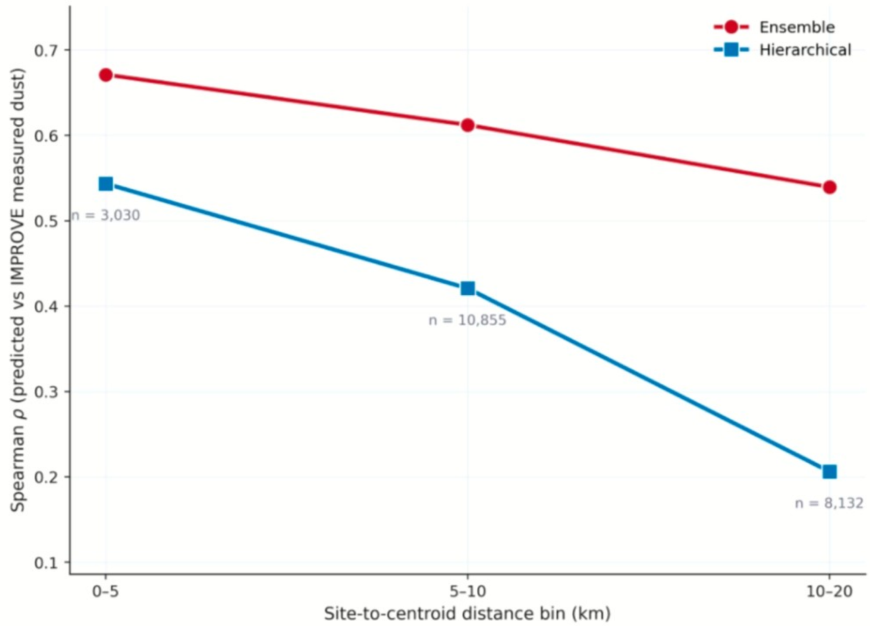

Figure 2.5. Spatial transferability of the two surfaces: Spearman rank correlation against IMPROVE measured dust within three distance bins (0 to 5 km, n = 3,030; 5 to 10 km, n = 10,855; 10 to 20 km, n = 8,132). Both surfaces decline with distance, but the ensemble (red) stays above the hierarchical surface (blue) at every bin, falling from rho about 0.67 to about 0.54 while the hierarchical surface falls from about 0.54 to about 0.21. The bin definitions differ slightly from the cumulative thresholds in Table 2.7, so the two should be read as complementary views of the same monotonic pattern.

Dust-event discrimination

Discrimination of high-dust days was strong and, again, distance-dependent. Using the 95th percentile of IMPROVE measured dust as the event threshold, ensemble AUC fell from 0.899 at 5 km to 0.827, 0.812, and 0.810 at 10, 15, and 20 km, exceeding the hierarchical surface throughout (0.818 to 0.695) (Table 2.8). The stricter precision-recall metric, appropriate because dust events are rare (152 to 1,100 events across thresholds), preserved the same ordering, with ensemble AUPRC roughly two to three times the hierarchical value at the closest threshold (0.330 versus 0.124). Against the independent Tong dust classification, ensemble AUC was 0.837 versus 0.708 for the hierarchical surface (Table 2.9). Because the Tong labels derive from composition-based clustering that never entered training, this is among the most direct evidence that the ensemble surface carries a genuine dust-specific signal rather than generic PM_{10} skill.

Table 2.8. Dust-event discrimination by monitor-to-centroid distance: AUC, AUPRC, and event counts.

Distance buffer	Model	AUC	AUPRC	Dust events (n)
≤ 5 km	Ensemble	0.899	0.330	152
≤ 5 km	Hierarchical	0.818	0.124	152
≤ 10 km	Ensemble	0.827	0.206	694
≤ 10 km	Hierarchical	0.765	0.132	694
≤ 15 km	Ensemble	0.812	0.181	959
≤ 15 km	Hierarchical	0.702	0.100	959
≤ 20 km	Ensemble	0.810	0.180	1,100
≤ 20 km	Hierarchical	0.695	0.095	1,100

Note. Events are IMPROVE site-days above the 95th percentile of measured dust. AUPRC = area under the precision-recall curve, its random-classifier baseline equals the event prevalence (about 5%).

Table 2.9. Dust-event discrimination against the Tong composition-based dust classification.

Model	AUC
Ensemble	0.837
Hierarchical	0.708

Note. Dust days classified from IMPROVE composition by Ward clustering (Tong et al., 2012).

Bias and construct interpretation

Raw-scale bias differed by matching scheme, and the two figures describe different samples rather than a contradiction. In the pooled full-validation set, ensemble bias against IMPROVE measured dust was +4.65 μg/m³ (Table 2.6). In the closest-centroid set restricted to within 20 km, it was +15.61 μg/m³. The larger value reflects the magnitude scale of the PM_{10} carrier signal relative to composition-reconstructed dust, not error on the PM_{10} target itself. The hierarchical surface showed much smaller raw-scale magnitudes but substantially weaker rank agreement. The full metric panel by distance, including raw-scale RMSE and MAE for both surfaces, appears in Appendix Table A2.1. These comparisons read most clearly in construct terms: IMPROVE measured dust is composition-derived, whereas the ensemble surface is a model-derived prediction operationalized as dust-related PM, so rank-order metrics carry more weight than raw-scale bias.

Seasonal and intensity structure

Stratified validation revealed seasonal and intensity structure. Ensemble rank agreement was strongest in fall (rho = 0.602 at 20 km) and weakest in winter (0.504), and both surfaces performed better in the October through March half-year. Performance dropped sharply in the most intense bin: above the 95th percentile of measured dust, ensemble rho fell to 0.144 and hierarchical rho to 0.128. The surfaces are therefore more reliable for broad exposure ranking than for pinpoint reconstruction of the most extreme days, which bears directly on how the surface should be used downstream.

Leave-one-site-out generalization

Leave-one-site-out analysis was consistent with the pooled results. Mean remaining-site Spearman rho was 0.645 for the ensemble and 0.456 for the hierarchical surface, essentially matching the pooled full-validation estimates, while held-out-site performance mirrored the per-site pattern, with the strongest ensemble generalization at DOME1 and the weakest at KAIS1.

Calibration

Calibration to IMPROVE measured dust improved both surfaces. On the within-20 km closest-centroid set, random-intercept calibration raised ensemble Spearman rho from 0.587 to 0.618, Pearson r from 0.441 to 0.494, while removing essentially all raw-scale bias, from +15.60 to -0.01 μg/m³ (Table 2.10). The hierarchical surface improved proportionally more in rank terms, from rho = 0.344 to 0.466.Calibration aligned both surfaces to the comparator without erasing the gap between them: the ensemble remained the stronger surface after calibration.

Table 2.10. External-validation performance before and after random-intercept post-calibration (closest-centroid, within 20 km).

Surface	Spearman rho	Pearson r	Bias (μg/m³)
Ensemble, pre-calibration	0.587	0.441	+15.60
Ensemble, post-calibration	0.618	0.494	-0.01
Hierarchical, pre-calibration	0.344	–	–

Surface	Spearman rho	Pearson r	Bias (µg/m³)
Hierarchical, post-calibration	0.466	–	–

Note. Calibration models were fit on the log(1 + x) scale with a site-level random intercept and back-transformed with Duan smearing.

Statewide Exposure Surface and Population Burden

Mapped across the state, the ensemble surface, the primary exposure product, concentrates dust-attributable exposure where arid land, agricultural disturbance, and dust-favorable meteorology coincide with people. The mean annual count of days above 50 µg/m³ is highest in the southern San Joaquin Valley (Kern, Tulare, and Kings counties) and the Imperial Valley and Salton Sea region, while most coastal and northern ZIP codes see fewer than ten such days (Figure 2.6). Pairing mean concentration with population makes the public-health stakes explicit. The Central Valley corridor through Fresno, Bakersfield, and Stockton emerges as the dominant exposure hotspot, where elevated mean DAF-PM_{10} overlaps dense residential population (Figure 2.7), and regional insets show mean concentrations above 30 to 35 µg/m³ around the southern San Joaquin Valley and the receding Salton Sea shoreline (Figure 2.8).

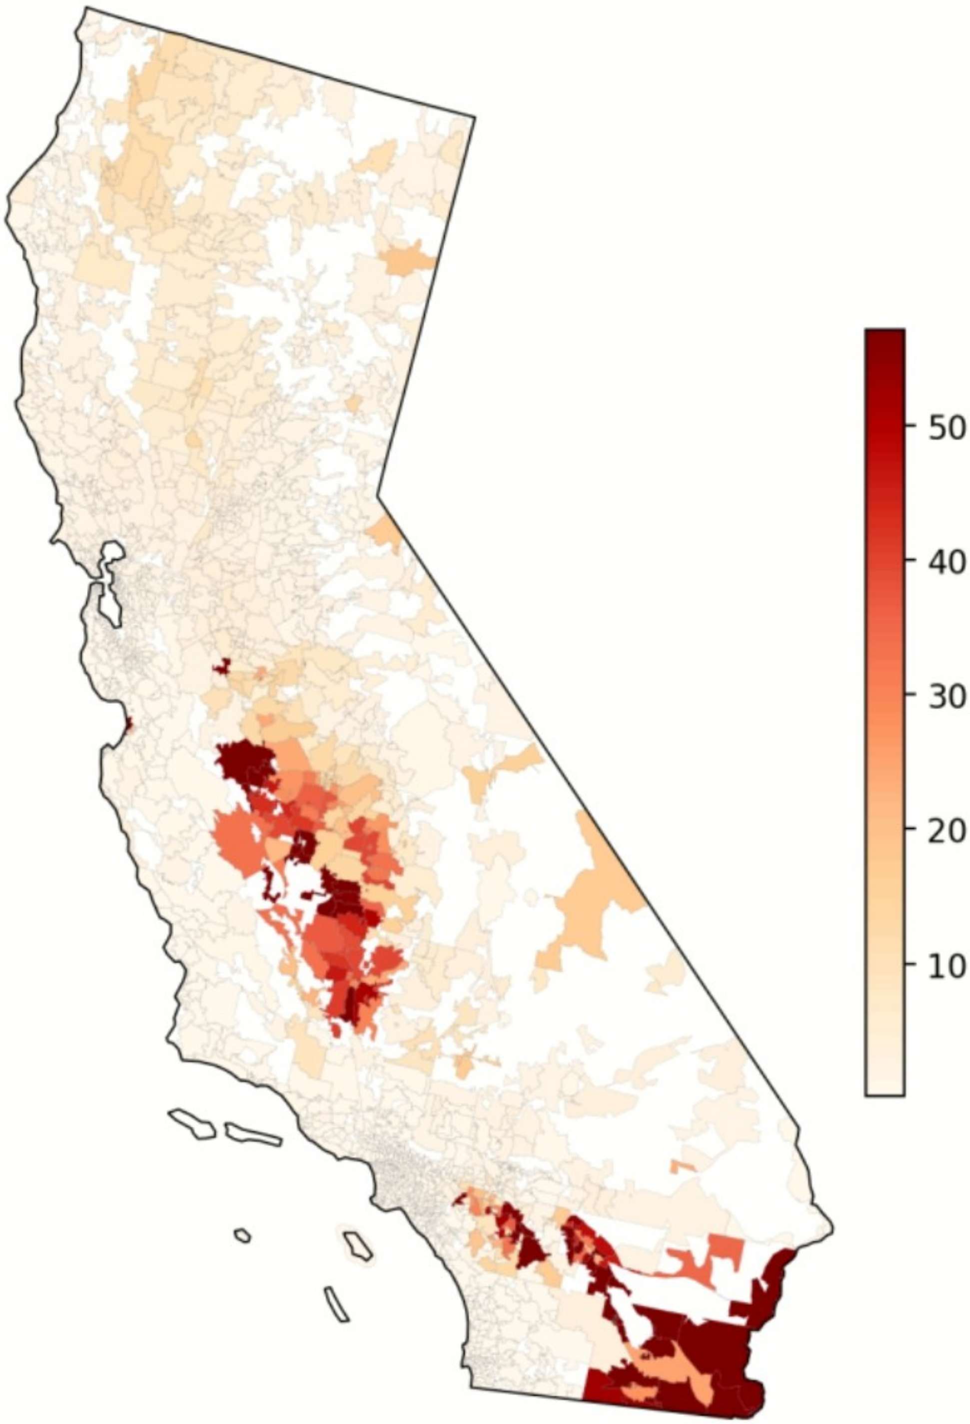

Figure 2.6. Mean annual number of days with DAF-PM_{10} above 50 µg/m³ by ZIP code, 2006 through 2019 (ensemble surface). The highest exceedance frequencies fall in the southern San Joaquin Valley and the Imperial Valley and Salton Sea region, where arid surfaces, agricultural dust, and dust-favorable meteorology coincide, most coastal and northern ZIP codes record fewer than ten exceedance days per year.

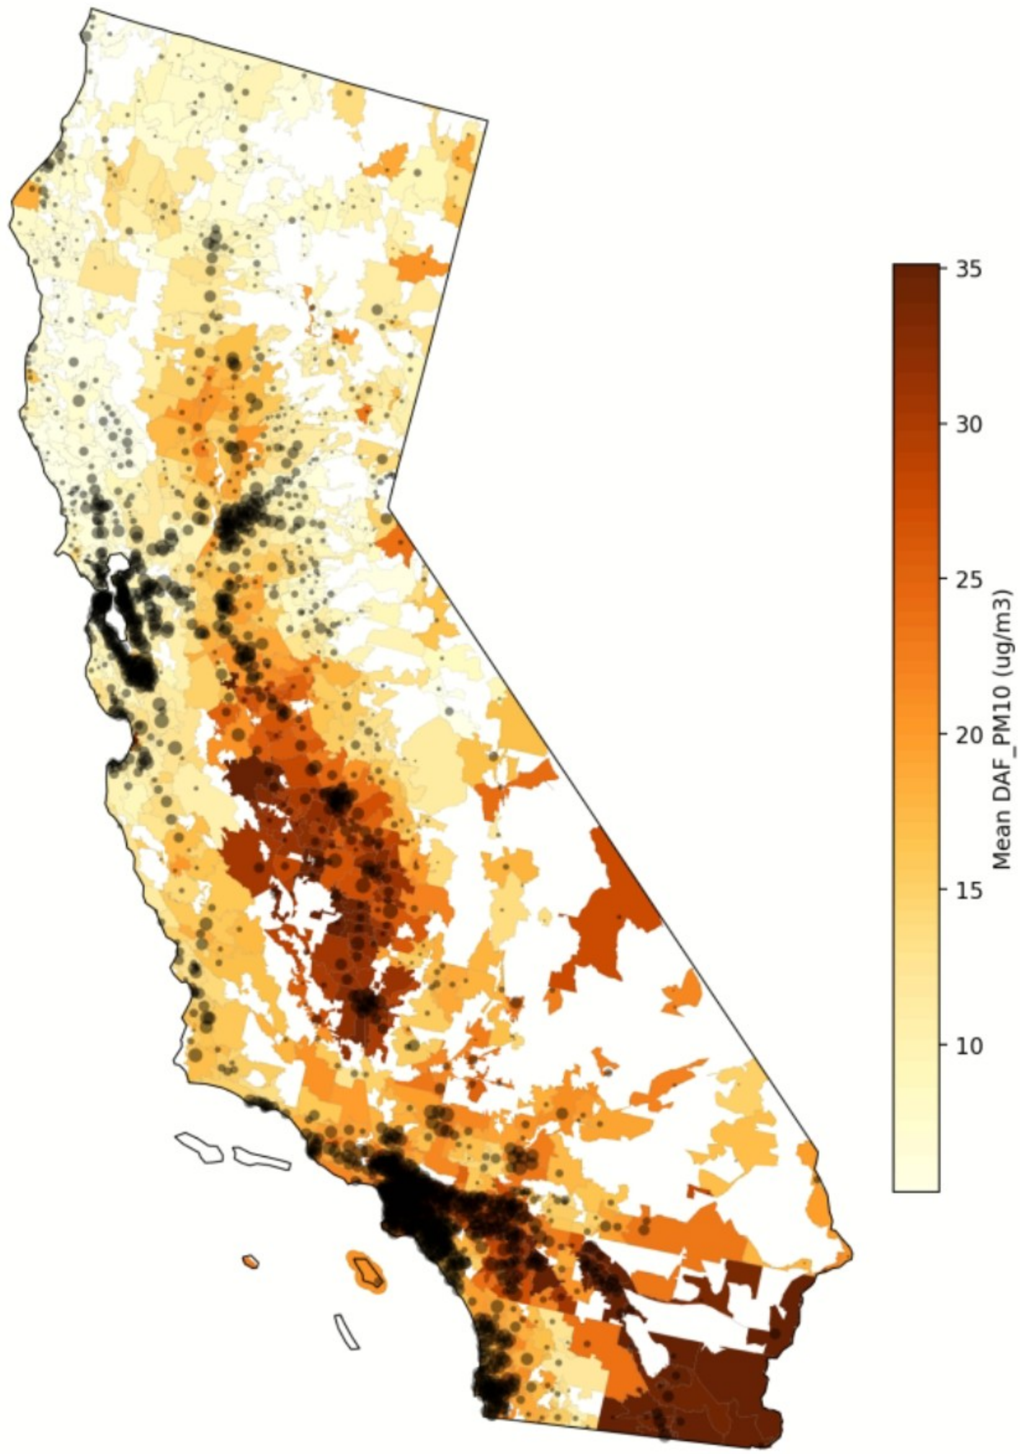

Figure 2.7. Population-weighted exposure hotspots (ensemble surface): mean DAF-PM_{10} by ZIP code (fill) with total ZIP population (dot size). Areas where high concentration coincides with large population carry the greatest population-level burden, most prominently the Central Valley corridor through Fresno, Bakersfield, and Stockton.

Resolving these statewide hotspots at regional scale (Figure 2.8) shows that the two burden centres are physically distinct: the Central Valley corridor is an agricultural and urban landscape, while the Imperial Valley and Salton Sea form a desiccating saline lakebed, yet both place large resident populations on the high tail of the exposure distribution.

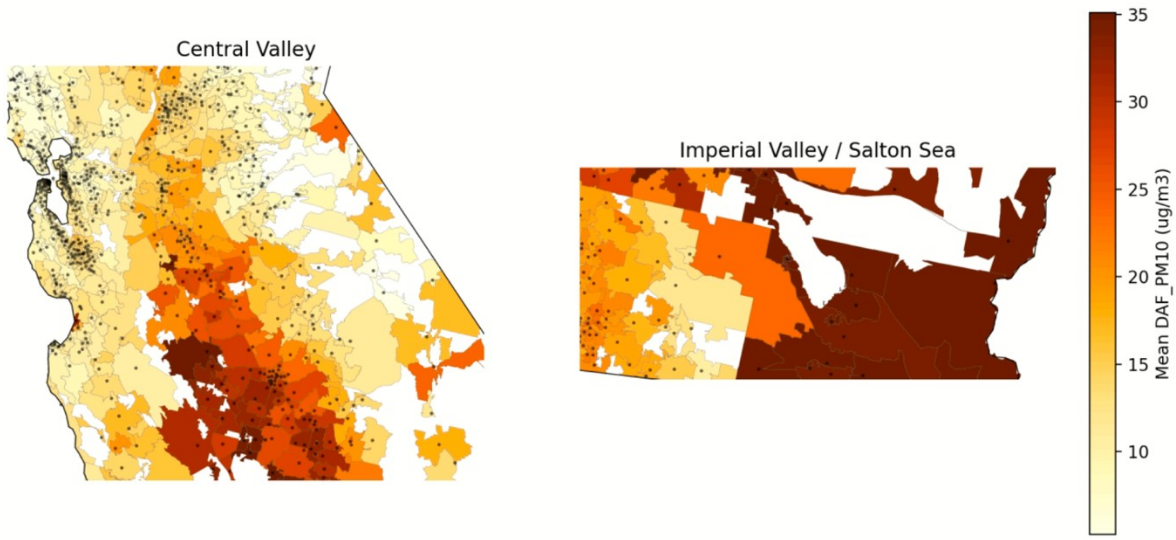

Figure 2.8. Regional insets of population-weighted DAF-PM_{10} exposure (ensemble surface). Left: the Central Valley, with the highest mean concentrations (above 30 µg/m³) around Fresno, Bakersfield, and the southern San Joaquin Valley. Right: the Imperial Valley and Salton Sea region, where mean DAF-PM_{10} exceeds 35 µg/m³ in ZIP codes bordering the receding shoreline.

The chronic dimension of this burden is captured by the fraction of days that exceed the statewide 95th-percentile threshold of about 38.8 µg/m³. Southern San Joaquin and Imperial Valley ZIP codes exceed it on more than 30% of days, whereas most coastal and northern ZIP codes do so on fewer than 5% (Figure 2.9). These communities experience sustained high exposure rather than occasional spikes, a distinction that matters for both surveillance and intervention.

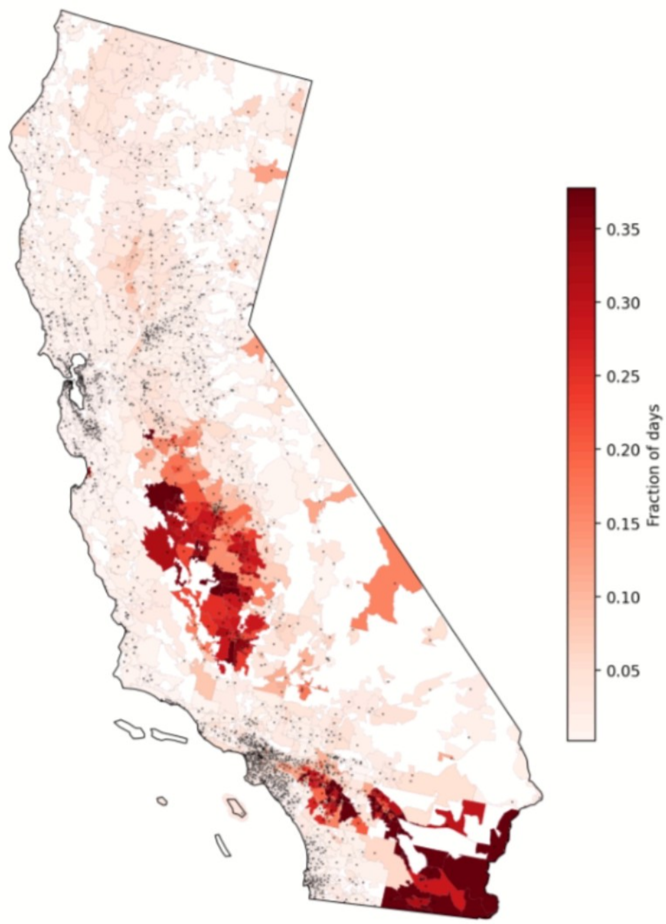

Figure 2.9. Fraction of days, 2006 through 2019, with DAF-PM_{10} above the statewide 95th-percentile threshold (about 38.8 μg/m³) by ZIP code (ensemble surface). Southern San Joaquin and Imperial Valley ZIP codes exceed the threshold on more than 30% of days, indicating persistent rather than episodic extreme exposure; most coastal and northern ZIP codes do so on fewer than 5%.

The seasonal and multi-year structure of the surface is physically coherent. Monthly climatology shows concentrations building from spring through late summer, peaking in July and August across the Central and Imperial Valleys, then declining through autumn and winter, with the Imperial Valley and Salton Sea region holding relatively high values even in winter, consistent with year-round source activity distinct from the agricultural cycle (Figure 2.10).

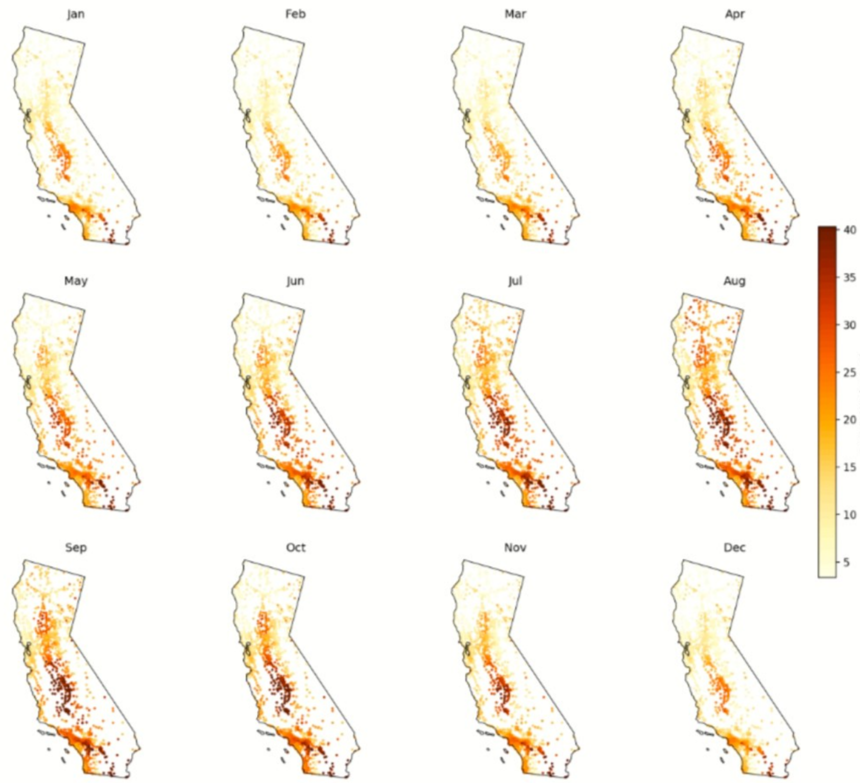

Figure 2.10. Monthly climatology of daily DAF-PM_{10} across California, averaged over 2006 through 2019 (ensemble surface). Concentrations build from spring through late summer, peak in July and August in the Central and Imperial Valleys, and recede through autumn and winter; the Imperial Valley and Salton Sea region remains comparatively high year-round.

Hierarchical Scoring Structure and Cross-Method Comparison

The hierarchical surface is weaker as an exposure predictor but valuable because every value traces to explicit rules, which makes its behavior easy to audit. The frequency with which each rule fired exposes both the design and its main weakness (Table 2.11). Low relative humidity (below 40%) was triggered on 61.7% of centroid-days, a reflection of California's prevailing aridity rather than of dust activity, and a coarse $PM_{2.5}/PM_{10}$ ratio fired on 24.7%. The authoritative Tier-1 NWS criterion was vanishingly rare (21 days, 0.0002%), as was the combined AOD and Angstrom criterion (35 days, 0.0004%). This imbalance between a rare authoritative signal and frequently triggered meteorological proxies is what motivated the probabilistic ensemble alternative in the first place.

Table 2.11. Scoring criteria, point weights, and trigger frequency across all centroid-days (hierarchical surface).

Criterion	Threshold	Points	Days (n)	Frequency (%)
NWS dust advisory (Tier 1)	Documented event	25	21	0.0002
PM_{10} anomaly	> 2 SD above seasonal mean	7	318,711	3.47
$PM_{2.5}/PM_{10}$ ratio	≤ 0.4	6	2,269,631	24.70
AOD with Angstrom exponent	AOD > 0.5 and AE < 0.5	4	35	0.0004
Wind speed	> 6 m/s	3	598,158	6.51
Visibility	< 3 km	3	528,204	5.75
Relative humidity	< 40%	2	5,667,976	61.69

Note. Points earned on a centroid-day are summed and normalized by the maximum attainable (25), and the 0-to-1 score multiplies modeled PM_{10} to yield hierarchical DAF-PM_{10}. Tiers denote authoritativeness, not sequence.

Hierarchical DAF-PM_{10} was highest on average in spring (2.45 μg/m³), followed by summer (2.30) and fall (2.23), and lowest in winter (1.58) (Table 2.12). The largest single-day extremes nonetheless fell in fall (2,994 μg/m³) and summer (1,053 μg/m³), so spring carried the highest typical attribution while the rare extremes clustered in fall and summer. Annual means ranged from 1.66 μg/m³ in the wet 2010 El Nino year to 2.58 μg/m³ in the post-drought year of 2018, and the count of high-dust days (above 10 μg/m³) tracked drought, peaking in dry years (Table 2.13). A linear regression of monthly mean DAF-PM_{10} on time showed no significant statewide trend (slope = 0.001, $p = 0.35$), indicating that interannual drought and wet cycles, rather than a lasting direction, governed variability over the period. This linear summary is the only trend test performed.

Table 2.12. Seasonal distribution of hierarchical DAF-PM_{10} across California ZIP centroids, 2006 through 2019.

Season	Mean (µg/m³)	Maximum (µg/m³)	SD
Spring (Mar-May)	2.45	37.1	3.08
Summer (Jun-Aug)	2.30	1,053	3.59
Fall (Sep-Nov)	2.23	2,994	5.07
Winter (Dec-Feb)	1.58	90.3	2.61

Note. Meteorological seasons. Extreme fall and summer maxima reflect episodic events rather than typical day-to-day exposure.

Table 2.13. Annual mean hierarchical DAF-PM_{10} and high-dust day counts for selected years in climate context.

Year	Mean DAF-PM_{10} (µg/m³)	Days > 10 µg/m³ (n)	Climate context
2018	2.58	24,134	Severe and post-drought
2015	2.49	20,234	Peak of historic drought
2007	2.42	29,289	Dry year
2010	1.66	7,433	Wet year (El Nino)

Note. High-dust days are centroid-days above 10 µg/m³. No significant statewide linear trend was detected over 2006 through 2019.

A small number of far-tail hierarchical values reach into the thousands of µg/m³, most of them in fall. These extremes are a direct consequence of multiplying the composite score by CMAQ PM_{10}, where the underlying CMAQ value is itself extreme, the product is extreme as well, and such values may reflect transport-model artifacts rather than observed dust. The apparent gap between the hierarchical 99.9th percentile (24.77 µg/m³) and these far-tail values is not a conflict between sources, both summarize the same strongly right-skewed distribution, the percentile describing the bulk and the table listing only the absolute extremes.

Hierarchical DAF-PM_{10} correlated weakly with its candidate drivers: positively with 550 nm AOD ($r = 0.18$) and, unexpectedly, with relative humidity ($r = 0.05$ using the minimum-RH field entering the scoring rule), negatively with visibility ($r = -0.12$), and near zero with wind speed ($r = 0.02$). The positive humidity correlation runs counter to a simple dry-conditions expectation and is most plausibly an artifact of the scoring structure: because the humidity rule fired on the majority of days, it imposed a broad floor on the score that blunts the expected negative relationship.

Direct comparison confirmed that the two surfaces are not interchangeable. Their daily predictions correlated only weakly (Spearman rho = 0.274), and agreement on the most extreme days was low. Both surfaces correlated positively with AOD, but the ensemble more strongly, and it also tracked the dust-flux signal more closely. The two surfaces emphasize different facets of the dust-relevant signal: the ensemble integrates many predictors into a broader particulate pattern, while the hierarchical score stays tied to a narrow, transparent rule set. That contrast is the basis for treating the ensemble as the primary exposure surface and the hierarchical score as an interpretable companion rather than a competing estimate.

Discussion

This chapter delivers a daily, statewide surface of dust-attributable PM_{10} for 1,797 population-weighted ZIP-code centroids across 2006 through 2019, evaluated against IMPROVE measured dust and set beside a transparent hierarchical comparator. The ensemble surface preserved meaningful rank-order agreement with the independent comparator (pooled Spearman rho = 0.646) and discriminated high-dust days well (AUC up to 0.899 at the closest distances), and it mapped a coherent statewide burden concentrated in the southern San Joaquin and Imperial Valleys. The central advance is not that a model was trained, but that a population-relevant, statewide, daily exposure surface was produced and externally anchored in a domain where dust characterization is methodologically unsettled, the comparator is imperfect, and epidemiologic usability has lagged behind process knowledge.

Prior western United States and California research has substantially improved dust climatology, event identification, source characterization, and regional hazard framing (Tong et al., 2012; Lei & Wang, 2014; Achakulwisut et al., 2017; Parajuli & Zender, 2018; Bhattachan et al., 2019; Frie et al., 2019; Kandakji et al., 2020; Rowan et al., 2024; Adebiyi et al., 2025). That body of work does not, on its own, supply the combination assembled here: statewide daily coverage, prediction at population-weighted centroids, an integrative predictor architecture, and external evaluation embedded in the workflow. The contribution is therefore specific. It is a defensible, externally evaluated, statewide dust-relevant exposure surface that is much closer to what epidemiologic analysis requires, not a claim to have resolved dust physics.

That contribution should be read against a field in which dust exposure is not a single, settled construct. Querol et al. (2019), Tobias and Stafoggia (2020), Hohsfield et al. (2023), and Ginoux et al. (2012) each show, from different directions, that dust and health inference depends on how dust is measured, defined, and operationalized. The absence of a statewide gold standard is therefore a field-wide condition rather than a flaw unique to this work. IMPROVE measured dust is appropriately framed as the strongest available comparator, composition-based, geographically distributed, and directly relevant to mineral-dust-containing mixtures, but still sparse and not equivalent to a statewide modeled surface. The question is not whether the surface matches a perfect truth standard, but whether it shows meaningful external performance against the best comparator available, and on that question the ensemble results are encouraging.

The distance-decay evidence strengthens that reading. Agreement did not simply differ between a near and a far endpoint, it declined in a graded pattern across 5, 10, 15, and 20 km. That smoothness is scientifically coherent and implies that a meaningful share of the apparent disagreement between predictions and IMPROVE measurements is attributable to the offset between populated centroids and remote monitors rather than to model breakdown. The surface is being judged across nonzero spatial separation against a sparse, composition-based network, and it tracks the comparator more closely as the geometry improves. Preserved rank order across space and time is precisely the property that matters most for many downstream epidemiologic applications, even where raw-scale bias remains.

The results also clarify what the ensemble surface is. Some dust-specific predictors carried low importance, the model was trained on total PM_{10}, and the diagnostics indicate a broad dust-relevant particulate pattern shaped by spatial structure, temperature, CMAQ PM_{10}, and related predictors rather than a pure mineral-dust signal. Recent PM_{10} and coarse-particle modeling increasingly depends on integrative, multi-source approaches rather than narrow single-indicator logic (Di et al., 2019; Shi et al., 2024; Cusano et al., 2025; Gora et al., 2025; Zhang et al., 2024). The accurate description is that the surface functions as a dust-informed PM_{10} exposure surface with demonstrated external relevance to measured dust, established here by discrimination against the composition-based Tong classification (AUC = 0.837), rather than as a mechanistic reconstruction of dust mass alone.

That description is also the more useful one for epidemiology. If earlier dust and health studies were constrained by event definitions, sparse monitoring, or imperfect dust characterization, a statewide daily surface that improves continuity, population coverage, and external anchoring may reduce an important source of exposure error. Because the surface spans the full fourteen-year record from 2006 through 2019, it also supports the multi-year and seasonal exposure contrasts that single-year products cannot. Tobias and Stafoggia (2020) and Hohsfield et al. (2023) make clear that exposure-definition heterogeneity is a serious methodological problem, and read alongside broader health reviews (Zhang et al., 2016; Lwin et al., 2023) this raises the possibility that some prior effect estimates were attenuated by imperfect dust characterization. The present work does not prove that proposition, but it provides the spatially and temporally resolved surface needed to test it rigorously in subsequent California analyses.

The California-specific value of the surface deserves emphasis. The state combines natural dust environments, anthropogenically disturbed source regions such as the Salton Sea playa and fallowed farmland, dense population corridors, and high climatic variability, which is exactly the setting in which a coarse, event-only, or monitor-only dust metric is least adequate. By producing daily statewide predictions at population-weighted centroids, this work shifts the literature from source characterization and event documentation toward exposure assignment at the scale health linkage requires, and it localizes the heaviest chronic burden to the southern San Joaquin and Imperial Valleys, where high concentration and dense population coincide.

The hierarchical surface plays a supporting role in this story. Its value lies in transparency, the scoring scheme makes the rule structure explicit and shows how a dust-indicator logic behaves over space and time. It also clarifies its own artifacts. The largest event carried a raw score of 9, and the apparent tension between a 99.9th percentile near 24 μg/m³ and individual values reaching 2,994 μg/m³ dissolves once both are recognized as different summaries of one strongly right-skewed distribution. As a primary exposure surface, however, the hierarchical score is weaker, with lower pooled agreement, weaker rank agreement across most distances, and sensitivity to the structure of its own rules, notably the humidity criterion that fired on most days. The two surfaces correlate only weakly, which is why they are best treated as complementary rather than interchangeable.

The strengths of the work are concrete. It achieved statewide daily coverage across fourteen years at population-weighted centroids. It integrated meteorology, modeled particulate fields, land surface, AOD, inversion structure, and dust indicators rather than relying on any single stream. It embedded external validation against IMPROVE rather than resting on internal metrics. It supported distance-stratified, per-site, seasonal, leave-one-site-out, and post-calibration analyses. Above all, it produced a surface that is directly usable for the epidemiologic phase that follows.

The limitations are real and are best understood as boundaries on interpretation. First, there is no statewide gold standard, so all external metrics are relative to IMPROVE measured dust, because validation is conducted across the offset between populated centroids and remote monitors, the reported metrics most likely understate true residential performance and should be read as conservative. Second, the ensemble construct is broader than pure dust mass, so using it naively as a dust-only exposure could blend dust with non-dust PM_{10} variation; the composition-anchored discrimination mitigates but does not eliminate this concern. Third, performance degrades in the extreme tail, so the surface is better suited to ranking broad exposure than to reconstructing the most extreme days, and extreme-event analyses should rely on rank-based or quantile exposure windows. Fourth, the far-tail hierarchical values inherit extreme CMAQ PM_{10} and may be transport-model artifacts, which is one more reason to treat the ensemble as primary and the hierarchical score as an auditable companion.

For public health practice, the surface identifies where dust burden is chronic rather than episodic, which is the more actionable distinction for surveillance and intervention. The persistent exceedance fractions in the southern San Joaquin and Imperial Valleys point to communities that bear sustained exposure, and because the estimates sit at population-weighted centroids they align naturally with administrative health data and with analyses of exposure disparity. The same daily, spatially resolved structure is well suited to feeding dust early-warning evaluation, where the operational need is reliable relative ranking of dust conditions over space and time.

Two lines of future work follow. The first is testing whether source-specific predictors improve discrimination of the most extreme days. A closely related step is an explicit source-apportionment stage, whether through receptor models such as positive matrix factorization (Paatero & Tapper, 1994; Belis et al., 2013) or through a dust-versus-non-dust decomposition added upstream of the ensemble, which would move the surface from a dust-informed particulate prediction toward a chemically resolved dust mass and would tighten the construct that the external validation can only partly secure. The second is translational, and it is where this chapter's contribution becomes consequential. The statewide daily centroid-level surface is built to be carried into the epidemiologic analysis that follows in this dissertation, where the principal requirement is a stable, spatially resolved, population-relevant ranking of dust exposure over time. The next chapter takes up that analysis, linking this exposure surface to emergency-department and inpatient morbidity outcomes under a time-stratified case-crossover design.

Conclusion

This chapter constructed and externally evaluated a daily, statewide surface of dust-attributable PM_{10} for California at population-weighted ZIP-code centroids, 2006 through 2019. The machine-learning ensemble preserved strong rank-order agreement with IMPROVE measured dust, discriminated high-dust days well, and degraded smoothly with monitor-to-centroid distance, a pattern that points to geometric separation rather than model failure as a leading source of apparent disagreement. Mapped statewide, the surface concentrates chronic dust burden in the southern San Joaquin and Imperial Valleys, where elevated concentrations overlap dense population. A parallel hierarchical score adds transparency and context without rivaling the ensemble as an exposure predictor. The surface does not resolve dust characterization in an absolute sense, and it is best described as a dust-informed PM_{10} exposure surface rather than a reconstruction of pure dust mass. Within those boundaries, it provides a deeper and more defensible statewide dust exposure framework than was previously available for California, and it is ready to serve as the exposure foundation for the health analysis that follows.

Acknowledgment of Contributions

The WEMO horizontal and vertical dust-flux fields were contributed by a collaborator (Bo) who generated them from twenty years of monthly values driven by high-resolution downscaled California wind fields developed by Hall and colleagues. All exposure modeling, feature integration, external validation, post-calibration, and interpretation reported in this chapter are the author's own work.

References

Abatzoglou, J. T. (2013). Development of gridded surface meteorological data for ecological applications and modelling. *International Journal of Climatology*. https://doi.org/10.1002/joc.3413

Achakulwisut, P., Shen, L., & Mickley, L. J. (2017). What controls springtime fine dust variability in the western United States? Investigating the 2002-2015 increase in fine dust in the U.S. Southwest. *Journal of Geophysical Research: Atmospheres*. https://doi.org/10.1002/2017JD027208

Adebiyi, A. A., Kibria, M. M., Abatzoglou, J. T., Ginoux, P., Pandey, S., Heaney, A., Chen, S.-H., & Akinsanola, A. A. (2025). Fallowed agricultural lands dominate anthropogenic dust sources in California. *Communications Earth & Environment*. https://doi.org/10.1038/s43247-025-02306-0

Al-Dousari, A., et al. (2025). Dust and health risks: A review of impacts and mitigation. *Journal of Environmental Health Science and Engineering*. https://doi.org/10.1007/s41748-025-00650-x

Belis, C. A., Karagulian, F., Larsen, B. R., & Hopke, P. K. (2013). Critical review and meta-analysis of ambient particulate matter source apportionment using receptor models in Europe. *Atmospheric Environment*, 69, 94–108. https://doi.org/10.1016/j.atmosenv.2012.11.009

Bhattachan, A., et al. (2019). Characterizing the role of wind and dust in traffic accidents in California. *GeoHealth*. https://doi.org/10.1029/2019GH000212

Breiman, L. (1996). Stacked regressions. *Machine Learning*, 24(1), 49–64. https://doi.org/10.1007/BF00117832

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324

Byun, D., & Schere, K. L. (2006). Review of the governing equations, computational algorithms, and other components of the Models-3 Community Multiscale Air Quality (CMAQ) modeling system. *Applied Mechanics Reviews*, 59(2), 51–77. https://doi.org/10.1115/1.2128636

Crooks, J. L., et al. (2016). The association between dust storms and daily non-accidental mortality in the United States, 1993-2005. *Environmental Health Perspectives*. https://doi.org/10.1289/EHP216

Cusano, M., Gaeta, A., Morelli, R., Cattani, G., Canepari, S., Massimi, L., & Leone, G. (2025). Spatial modeling of trace element concentrations in PM_{10} using generalized additive models. *Atmosphere*. https://doi.org/10.3390/atmos16040464

Di, Q., Amini, H., Shi, L., Kloog, I., Silvern, R., Kelly, J., Sabath, M. B., Choirat, C., Koutrakis, P., Lyapustin, A., Wang, Y., Mickley, L. J., & Schwartz, J. (2019). An ensemble-based model of $PM_{2.5}$ concentration across the contiguous United States with high spatiotemporal resolution. *Environment International*, 130, 104909. https://doi.org/10.1016/j.envint.2019.104909

Dietterich, T. G. (2000). Ensemble methods in machine learning. In *Multiple classifier systems* (Lecture Notes in Computer Science, Vol. 1857, pp. 1–15). Springer. https://doi.org/10.1007/3-540-45014-9_1

Duan, N. (1983). Smearing estimate: A nonparametric retransformation method. *Journal of the American Statistical Association*, 78(383), 605–610. https://doi.org/10.1080/01621459.1983.10478017

Formenti, P., Schütz, L., Balkanski, Y., Desboeufs, K., Ebert, M., Kandler, K., Petzold, A., Scheuvens, D., Weinbruch, S., & Zhang, D. (2011). Recent progress in understanding physical and chemical properties of African and Asian mineral dust. *Atmospheric Chemistry and Physics*, 11(16), 8231–8256. https://doi.org/10.5194/acp-11-8231-2011

Frie, A. L., et al. (2019). Dust sources in the Salton Sea basin: A clear case of an anthropogenically impacted dust budget. *Environmental Science & Technology*. https://doi.org/10.1021/acs.est.9b02137

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5), 1189–1232. https://doi.org/10.1214/aos/1013203451

Ginoux, P., Prospero, J. M., Gill, T. E., Hsu, N. C., & Zhao, M. (2012). Global-scale attribution of anthropogenic and natural dust sources and their emission rates based on MODIS Deep Blue aerosol products. *Reviews of Geophysics*. https://doi.org/10.1029/2012RG000388

Gora, D., et al. (2025). Machine-learning approaches for PM_{10} modeling and factor identification. *Sustainability*. https://doi.org/10.3390/su17125274

Hohsfield, L., Rowan, A., D'Souza, R., Ebelt, S., Chang, H. H., & Crooks, J. (2023). Evaluating data product exposure metrics for use in epidemiologic studies of dust storms. *GeoHealth*. https://doi.org/10.1029/2023GH000824

Kandakji, T., et al. (2020). Identifying and characterizing dust point sources in the southwestern United States using remote sensing and GIS. *Geomorphology*. https://doi.org/10.1016/j.geomorph.2019.107019

Lei, H., & Wang, J. (2014). Observed characteristics of dust storm events over the western United States using meteorological, satellite, and air quality measurements. *Atmospheric Chemistry and Physics*. https://doi.org/10.5194/acp-14-7847-2014

Lwin, K. K., et al. (2023). Effects of desert dust and sandstorms on human health: A scoping review. *GeoHealth*. https://doi.org/10.1029/2022GH000728

Lyapustin, A., Wang, Y., Korkin, S., & Huang, D. (2018). MODIS Collection 6 MAIAC algorithm. *Atmospheric Measurement Techniques*, 11(10), 5741–5765. https://doi.org/10.5194/amt-11-5741-2018

Malm, W. C., Sisler, J. F., Huffman, D., Eldred, R. A., & Cahill, T. A. (1994). Spatial and seasonal trends in particle concentration and optical extinction in the United States. *Journal of Geophysical Research: Atmospheres*, 99(D1), 1347–1370. https://doi.org/10.1029/93JD02916

Morakinyo, O. M., et al. (2016). Particulate matter pollution, composition, and health relevance: A review. *International Journal of Environmental Research and Public Health*. https://doi.org/10.3390/ijerph13060592

Paatero, P., & Tapper, U. (1994). Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2), 111–126. https://doi.org/10.1002/env.3170050203

Parajuli, S. P., & Zender, C. S. (2018). Projected changes in dust emissions and regional air quality due to the shrinking Salton Sea. *Aeolian Research*. https://doi.org/10.1016/j.aeolia.2018.05.004

Prospero, J. M., Ginoux, P., Torres, O., Nicholson, S. E., & Gill, T. E. (2002). Environmental characterization of global sources of atmospheric soil dust identified with the Nimbus 7 Total Ozone Mapping Spectrometer (TOMS) absorbing aerosol product. *Reviews of Geophysics*, 40(1), 1002. https://doi.org/10.1029/2000RG000095

Querol, X., et al. (2019). Monitoring the impact of desert dust outbreaks for air quality for health studies. *Environment International*. https://doi.org/10.1016/j.envint.2019.05.061

Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder, B., Thuiller, W., Warton, D. I., Wintle, B. A., Hartig, F., & Dormann, C. F. (2017). Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8), 913–929. https://doi.org/10.1111/ecog.02881

Rowan, A., et al. (2024). Dust storm health research in the United States and continuing exposure challenges. *Environmental Research: Health*. https://doi.org/10.1088/2752-5309/ad5751

Shi, X., et al. (2024). An ensemble machine learning model to enhance extrapolation ability of predicting coarse particulate matter with high resolutions in China. *Environmental Science & Technology*. https://doi.org/10.1021/acs.est.4c08610

Sinclair, K., et al. (2024). Shrinking shorelines and PM_{10} hazards in saline-lake environments. *Geographies*. https://doi.org/10.3390/geographies4040034

Tobias, A., & Stafoggia, M. (2020). Modeling desert dust exposures in epidemiologic short-term health effects studies. *Epidemiology*. https://doi.org/10.1097/EDE.0000000000001255

Tong, D. Q., et al. (2012). Long-term dust climatology in the western United States reconstructed from routine aerosol ground monitoring. *Atmospheric Chemistry and Physics*. https://doi.org/10.5194/acp-12-5189-2012

van der Laan, M. J., Polley, E. C., & Hubbard, A. E. (2007). Super Learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1), Article 25. https://doi.org/10.2202/1544-6115.1309

Van Donkelaar, A., et al. (2019). Regional estimates of chemical composition of fine particulate matter using a combined geoscience-statistical method with information from satellites, models, and monitors. *Environmental Science & Technology*. https://doi.org/10.1021/acs.est.8b06392

Villanueva, M., Monte-Sano, S., Sheridan, S., Allen, M., Kalkstein, L., Jerrett, M., & Eisenman, D. P. (2026). A climate region classification for California's warm season: Apparent temperature clustering to support heat-health epidemiology. *International Journal of Biometeorology, 70*(5), Article 134. https://doi.org/10.1007/s00484-026-03200-w

Wolpert, D. H. (1992). Stacked generalization. *Neural Networks*, 5(2), 241–259. https://doi.org/10.1016/S0893-6080(05)80023-1

Zhang, X., et al. (2016). A systematic review of global desert dust and associated human health effects. *Atmosphere*. https://doi.org/10.3390/atmos7120158

Zhang, Z., Johansson, C., Engardt, M., Stafoggia, M., & Ma, X. (2024). Improving 3-day deterministic air pollution forecasts using machine learning algorithms. *Atmospheric Chemistry and Physics*, 24(2), 807–851. https://doi.org/10.5194/acp-24-807-2024

Appendix 2A. Supplementary Validation Tables

The tables in this appendix supplement the external-validation results in the main text with the full continuous metric panel by distance and with year-by-year performance at the focal FRES1 monitoring site, the closest monitor-to-centroid pairing in the study (1.48 km).

Table A2.1. Full continuous agreement between modeled DAF-PM_{10} and IMPROVE measured dust by monitor-to-centroid distance.

Distance	Model	N	Spearman rho	Pearson r	RMSE	MAE	Bias
≤ 5 km	Ensemble	3,026	0.670	0.612	21.77	18.16	+18.16
≤ 5 km	Hierarchical	3,026	0.543	0.236	2.41	1.07	+0.30
≤ 10 km	Ensemble	13,875	0.625	0.480	19.94	16.30	+16.29
≤ 10 km	Hierarchical	13,875	0.461	0.214	1.81	0.78	+0.19
≤ 15 km	Ensemble	19,167	0.605	0.449	20.12	16.29	+16.28
≤ 15 km	Hierarchical	19,167	0.341	0.156	1.96	0.91	+0.35
All	Ensemble	22,000	0.587	0.441	19.67	15.63	+15.61
All	Hierarchical	22,000	0.343	0.153	2.01	0.91	+0.39

Note. RMSE, MAE, and bias in μg/m³. The large ensemble bias reflects the PM_{10} carrier scale relative to composition-reconstructed dust, not error on the PM_{10} target. The small hierarchical RMSE reflects its narrow output range, not superior accuracy.

Table A2.2. Ensemble performance at FRES1 by year, stratified by dust-storm frequency.

Year	Dust frequency	Spearman rho	RMSE (μg/m³)	AUC
2019	High	0.869	26.05	0.993
2018	High	0.776	33.75	0.892
2006	High	0.730	25.86	0.859
2014	High	0.694	30.82	0.767
2008	High	0.641	32.89	0.674
2009	High	0.613	25.06	0.830
2013	High	0.509	27.48	0.740
2017	Low	0.800	28.91	0.969
2012	Low	0.783	27.47	0.934
2016	Low	0.758	25.05	0.765
2010	Low	0.720	21.69	0.860
2011	Low	0.670	24.02	0.892
2007	Low	0.653	27.96	0.980
2015	Low	0.631	28.75	0.945

Note. Dust-frequency tier from annual NWS dust-event counts and Tong-classified dust-day density. Years are sorted within each tier by descending rho.

Table A2.3. Hierarchical performance at FRES1 by year, stratified by dust-storm frequency.

Year	Dust frequency	Spearman rho	RMSE (μg/m³)	AUC
2019	High	0.493	3.11	0.815
2006	High	0.471	2.71	0.670
2018	High	0.359	3.68	0.782
2008	High	0.334	3.39	0.650
2009	High	0.291	1.80	0.673
2013	High	0.235	2.82	0.611
2014	High	0.206	3.69	0.592
2011	Low	0.461	4.38	0.679
2017	Low	0.429	2.51	0.659
2007	Low	0.368	1.99	0.873
2010	Low	0.294	1.69	0.682
2012	Low	0.282	2.12	0.689
2016	Low	0.106	2.52	0.695
2015	Low	0.089	4.19	0.691

Note. Companion to Table A2.2. The lower RMSE reflects the hierarchical surface's narrow output range rather than improved accuracy.

Chapter 3. Dust-Attributable PM_{10} and Short-Term Cardiovascular and Respiratory Morbidity in California: A Case-Crossover Analysis with Heat-Metric Decomposition, Geographic Heterogeneity, and Seasonal Stratification

Chapter Abstract

Dust storms and chronic dust loadings are growing public-health concerns in California, where climate-driven dust generation, more frequent heat waves, and an ageing population at risk of acute cardiopulmonary events increasingly intersect. This chapter estimates the short-term association between the dust-attributable fraction of PM_{10} (DAF-PM_{10}, the ZIP-day exposure derived in Aim 2) and broad cardiovascular (CVD_any) and respiratory (RESP_any) emergency department and hospital encounters in California, 2006 to 2018. The analysis uses a time-stratified case-crossover design of roughly 23.97 million CVD strata and 23.07 million RESP strata (approximately 105 million and 101 million person-days, respectively), fitted with conditional logistic regression on case-day and matched-referent-day strata. The same-day odds ratio per interquartile-range increase (+16.558 µg/m³) is 1.00213 (95% CI 1.00104, 1.00323; $p = 1.3 \times 10^{-4}$) for CVD_any and 1.00957 (1.00846, 1.01068; $p < 10^{-60}$) for RESP_any.

Three substantive refinements follow. First, the respiratory concentration-response is strongly nonlinear and U-shaped (likelihood-ratio $p = 4.3 \times 10^{-29}$), while the cardiovascular curve is near-monotone ($p = 0.039$). Second, the respiratory effect peaks one day after exposure and accumulates to a seven-day cumulative odds ratio of 1.01360 (1.01199, 1.01522), roughly 1.4 times the same-day effect, whereas the cardiovascular effect is essentially confined to the same day. Third, a heat-interaction analysis shows respiratory amplification (joint per-IQR OR 1.01496, $p\_interaction = 1.5 \times 10^{-8}$) and cardiovascular attenuation (joint OR 0.99889, $p\_interaction = 2.7 \times 10^{-4}$) under an extreme-heat indicator, a heat-metric decomposition then resolves these into maximum-temperature-driven respiratory amplification and minimum-temperature-driven cardiovascular attenuation, with humid, stagnant heat (low diurnal temperature range) reversing both. Geographic stratification by official California Air Resources Board (CARB) air basins shows respiratory amplification that is positive in every populated basin except the chronically dust-saturated Salton Sea, though varying about three-fold in magnitude, with San Diego strongest (OR per IQR 1.031); cardiovascular effects concentrated in coastal metropolitan basins; and a null Salton Sea signal interpreted as a design limitation on a chronically saturated exposure regime rather than as evidence of no effect. Season-stratified analysis shows no seasonal heterogeneity for cardiovascular disease but strong heterogeneity for respiratory disease (summer largest). All primary findings are stable across five time-varying-confounding specifications and a permutation negative control. Spline-based population attributable fractions are 0.42% (CVD) and 0.43% (RESP) at the observed fifth-percentile counterfactual; cause-specific decomposition concentrates the respiratory effect in asthma (1.28%), chronic obstructive pulmonary disease (COPD), and acute respiratory infection, and the cardiovascular effect in hypertensive disease (0.22%), with cerebrovascular and ischemic-stroke point estimates null. The chapter contributes the largest dust-oriented PM_{10} case-crossover conducted to date, a formal nonlinearity test with a re-referenced attributable fraction that accommodates the U-shape, a heat-metric decomposition that distinguishes biologically distinct heat regimes.

1. Introduction

1.1 Dissertation context

This chapter presents Aim 3 of the dissertation. The systematic review in Chapter 1 (Aim 1) established what global dust early warning systems do well and where they fall short. The exposure work in Chapter 2 (Aim 2) developed and externally validated ensemble and hierarchical models that estimate the dust-attributable fraction of PM_{10} at the ZIP-day resolution across California for 2006 to 2019; that ZIP-day exposure surface supplies the exposure variable for the present aim. Aim 3 closes the exposure-to-health loop by asking whether day-to-day variation in dust-attributable PM_{10} is associated with acute cardiopulmonary morbidity, and Chapter 4 (Aim 4) then carries the combined evidence toward risk communication and operational warning. The earlier aims are documented in their own chapters; Aim 3 stands on its own and is the subject of this chapter.

1.2 Scientific motivation

Three converging literatures motivate this aim. The first is the broad evidence that short-term exposure to particulate matter is associated with cardiovascular and respiratory morbidity. Systematic synthesis of the short-term PM literature shows small but consistent same-day associations with cardiopulmonary events, with susceptibility concentrated among older adults and people with pre-existing disease (Bell et al., 2013). Attention has shifted over the past decade toward the source-specific apportionment of PM_{10} into combustion, dust, and biological fractions, because the biological reactivity and the available policy levers differ across sources (Morakinyo et al., 2016). The cleanest demonstrations of this principle separate desert-derived PM_{10} from co-occurring anthropogenic PM_{10} and compare their associations directly (Stafoggia et al., 2016).

The second is the dust-specific health literature, which remains comparatively thin and geographically uneven. Narrative and systematic reviews document associations between desert dust and respiratory and cardiovascular outcomes, but they also stress heterogeneity in exposure definition, inconsistent separation of dust from total PM, and an imbalance between the global burden of dust exposure and the geographic distribution of health studies (Goudie, 2014;

Karanasiou et al., 2012; Zhang et al., 2016). The most recent meta-analysis of short-term desert-dust exposure pooled 71 studies but found that American dust corridors were represented by only a handful of studies, and that binary dust-day exposure definitions precluded concentration-response estimation across most of the evidence base (Tobías et al., 2025). Within the United States, dust-health evidence has concentrated on national dust-storm catalogues and on the arid Southwest, linking dust events to non-accidental mortality (Crooks et al., 2016), intensive-care admissions (Rublee et al., 2020), and cardiorespiratory emergency department visits (Rowan et al., 2024), against a backdrop of a multi-decade increase in western-U.S. dust activity (Tong et al., 2012). California-wide, dust-specific, spatially resolved short-term morbidity evidence has been largely absent.

The third is the rapidly expanding literature on the interaction between air pollution and heat under climate-change pressure, where multicity analyses have shown that high temperature can modify the short-term association between particulate matter and mortality (Stafoggia et al., 2008) and that air pollutants in turn confound and modify the mortality effect of heat (Analitis et al., 2014). In California specifically, compound events in which heat waves, drought, and high-dust conditions co-occur are projected to become more frequent (Pu et al., 2022), which makes the joint cardiopulmonary burden of dust and heat a policy-relevant quantity. Yet most studies that test heat as a modifier of a pollution association use a single binary heat indicator, and formal decomposition of the interaction by heat type (daytime maximum temperature, overnight minimum temperature, and diurnal temperature range) is rare. The present aim contributes to all three literatures.

1.3 California setting

California's dust climate is geographically heterogeneous. The Salton Sea, a shrinking saline lake in Imperial County, is the state's most acutely dust-impacted region, and its exposed playa is an intensifying, anthropogenically modified dust source as the lake recedes (Frie et al., 2019; Parajuli & Zender, 2018). The Mojave and Colorado deserts experience dust episodes driven by frontal passages and convective windstorms, the Central Valley generates agricultural and roadway dust intermixed with non-dust PM, and fallowed farmland has become a dominant anthropogenic dust source statewide (Adebiyi et al., 2025; Bhattachan et al., 2019). Coastal metropolitan basins (South Coast and Los Angeles, San Francisco Bay, and San Diego) receive long-range transported dust intermixed with combustion sources (Lei & Wang, 2014), and springtime fine-dust loadings across the wider Southwest have risen over recent decades (Achakulwisut et al., 2017). The state's population is heavily concentrated in the coastal metropolitan basins, which means that a statewide same-day case-crossover is statistically dominated by encounters in the South Coast, San Francisco Bay, and San Diego basins. This dominance of coastal encounters is a property of the design rather than a flaw, but it makes geographic effect-modification analysis essential for honest interpretation.

1.4 Public-health gap addressed

The specific gap this aim fills is the absence of California-wide, dust-specific, ZIP-day-resolved short-term health evidence that formally accommodates concentration-response nonlinearity, lag structure, heat-metric interaction, air-basin-resolved geographic heterogeneity, and seasonal heterogeneity, in a case-crossover design at the scale of tens of millions of strata. Prior California PM-morbidity studies have used emergency department and patient-discharge data at comparable scale but without a dust-specific exposure surface. The international dust-health evidence that does separate dust from total PM is concentrated in southern Europe, the eastern Mediterranean, and East Asia, and it relies predominantly on binary dust-day exposure that cannot characterise a concentration-response curve (Karanasiou et al., 2012; Stafoggia et al., 2016; Tobías et al., 2025). No prior dust-specific California study has formally tested concentration-response nonlinearity at scale or decomposed the heat interaction by heat type. Underlying these analytic gaps is a gap in statistical power. Most prior dust case-crossover and time-series studies are one to several orders of magnitude smaller than a statewide administrative-data census of encounters, which has limited their ability to resolve the small cardiovascular signal that particulate epidemiology leads one to expect and to characterise a concentration-response curve with precision (Tobías et al., 2025). Linking a daily, ZIP-resolved exposure surface to the near-complete California encounter record converts that constraint into the central analytic asset of this chapter, with roughly twenty-four million matched strata per broad outcome, effects on the order of a few tenths of a percent per interquartile-range increment become estimable with narrow confidence intervals. The analytic expectation, given the prior literature, was a small positive same-day association for both outcome groups, larger and more lag-persistent for respiratory disease, and plausibly modified by extreme heat.

2. Methods

2.1 Study design

We used a time-stratified case-crossover design, the canonical short-term-exposure design for episodic exposures (Maclure, 1991). Each case-day was matched to all other days within the same calendar month and year that fell on the same day of the week as the case-day, giving (typically) three referent days per case-day. This matching conditions out, by design, all within-person time-invariant confounders (age, sex, comorbidity profile, neighborhood, and season-long pollution averages) as well as time-of-week and time-of-month patterns. Conditional logistic regression then estimates the case-day versus referent-day exposure contrast within each matched stratum. The time-stratified referent scheme is the localizable selection strategy shown to be free of the overlap bias and the time-trend bias that distort unidirectional and other referent-selection designs, which is why it is the standard choice for short-term air-pollution case-crossover analysis (Janes et al., 2005).

2.2 Outcomes

Two broad outcomes are reported as primary: CVD_any (any cardiovascular emergency department or patient-discharge encounter, by principal or secondary ICD-9 or ICD-10 diagnosis) and RESP_any (any respiratory encounter). Eleven cause-specific outcomes were also defined, seven cardiovascular subgroups (ischemic heart disease, myocardial infarction, heart failure, conduction or dysrhythmia, hypertensive disease, cerebrovascular disease, and ischemic stroke) and four respiratory subgroups (pneumonia, COPD, asthma, and acute respiratory infection). The combined emergency-department plus patient-discharge dataset is the substantive primary analysis, emergency-department-only and patient-discharge-only analyses are reported alongside, with the explicit caveat that the patient-discharge data are underpowered (186,000 to 311,000 strata versus approximately 23 million for the emergency-department data).

2.3 Exposure

The primary exposure is the daily ZIP-level dust-attributable fraction of PM_{10} from the Aim 2 ensemble machine-learning model (in µg/m³), with all odds ratios reported per +16.558 µg/m³, the overall analytic interquartile range of the ensemble metric. Two additional operationalisations of DAF-PM_{10} (the hierarchical model and the calibrated ensemble) are reported as exposure-metric sensitivities in Section 4.1. The exposure surface provides ZIP-day values for the entire 2006 to 2019 study period; only 2006 to 2018 are joined to events because health-data coverage ends in 2018 (Section 2.5).

Two facts about how the exposure surface was built are central for this chapter. First, Aim 2 trained a stacked-ensemble model (random-forest, gradient-boosting, and deep-learning base learners combined by a meta-learner over roughly thirty spatial, meteorological, and remote-sensing predictors) to predict dust-relevant PM_{10}.The "dust-attributable fraction" label therefore denotes a dust-informed total-PM_{10} surface whose spatial and temporal structure is shaped by dust-relevant covariates, rather than a chemically or statistically isolated mineral-dust mass. Throughout this chapter, DAF-PM_{10} is best read as a dust-related, dust-dominant particulate metric. Second, meteorological variables were themselves among the most influential predictors of the exposure surface, maximum temperature (Tmax) ranked fourth in the random forest (approximately 7.1% relative importance), third in the gradient-boosting machine (approximately 10.1%), and first in the deep-learning learner, while specific humidity ranked fifth in the gradient-boosting machine (approximately 7.1%). A construct subtlety matters for the robustness checks that follow, the Aim 2 humidity feature is *specific* humidity (a moisture-mass quantity), whereas the Aim 3 health-model covariate (Section 2.4) is *relative* humidity. The two are related but not the identical field, which loosens the humidity double-counting relative to temperature, where the very same Tmax field sits on both sides of the analysis. The ensemble validated well against held-out IMPROVE-network observations (Spearman ρ approximately 0.59 to 0.67; classification AUC approximately 0.81 to 0.90). The consequence, namely that the temperature and humidity terms in the health model (Section 2.4) are partly adjusting for covariates that are upstream of the exposure itself, is examined empirically in Sections 5.6 and 5.7 and discussed in Sections 6 and 7.

2.4 Confounder set

The primary specification regresses each outcome on DAF-PM_{10} with adjustment for daily maximum temperature, ozone, nitrogen dioxide, sulphur dioxide, and relative humidity, conditional on the matched stratum. Several operational details apply.

Thermal confounder identity. The thermal confounder is the daily maximum temperature (Tmax, in °C), delivered by the upstream exposure pipeline. We use Tmax rather than the daily mean because Tmax captures peak afternoon thermal stress, the biologically most relevant exposure for acute cardiopulmonary outcomes, and because Tmax is the standard thermal confounder in the short-term-exposure case-crossover literature. Downstream heat-metric work (Section 4.5) reads explicit Tmax, Tmin, and diurnal-temperature-range fields and confirms this choice. The Celsius versus Fahrenheit unit distinction is immaterial for a regression coefficient on a linear covariate and is not a source of bias.

Co-pollutants. Refined ZIP-day ozone, nitrogen dioxide, and sulphur dioxide are included as continuous linear covariates. $PM_{2.5}$ and total non-dust PM_{10} are not included in the primary specification. A full multi-PM co-pollutant model that adds non-dust PM_{10} and $PM_{2.5}$ is identified as a future direction (Sections 7.3 and 8.5).

Humidity. Daily mean relative humidity (the mean of daily minimum and maximum relative humidity) is included as a continuous covariate. Spearman correlations between relative humidity and DAF-PM_{10} ($\rho = -0.530$) and between relative humidity and Tmax ($\rho = -0.585$) are non-trivial; humidity is a known modifier of fine-particle toxicology and a determinant of dust source strength, so its inclusion is a confounding-control rather than an effect-mediation choice.

2.5 Period

The study period is 2006 to 2018 inclusive (13 calendar years). The exposure surface extends through 2019, but the emergency-department and patient-discharge coverage received ends at 2018, so 2019 is never joined to events. Annual emergency-department counts grew from 0.89 million (CVD) and 1.14 million (RESP) in 2006 to 2.64 million (CVD) and 2.16 million (RESP) in 2018.

2.6 Estimators

The primary estimator is conditional logistic regression with case-day and matched-referent-day strata. For every outcome-by-source cell we also fit an equivalent conditional Poisson model as an algebraic cross-check. The two are mathematically equivalent under the matched design (Armstrong et al., 2014; Lu & Zeger, 2007), and the observed coefficients agree to nine decimal places in every cell. We report the all-events conditional logistic estimate as the substantive primary estimate. Two additional model families that were originally implemented (person-clustered robust standard errors and a first-event-only restriction) are not interpreted substantively, because the analysis-ready person identifier is not globally unique (Section 5.5).

2.7 Heat metrics and the heat-interaction analysis

Three heat-metric families are used, in order of increasing refinement.

1. **Extreme-heat indicator (primary).** A binary variable equal to 1 when the daily ZIP Tmax meets the CalHeatScore Level 3-or-higher extreme-heat threshold for that ZIP-year, with ZIP-year thresholds for 2008 to 2018 supplied by the operational threshold set and the 2008 thresholds back-cast to 2006 to 2007. Dataset prevalence is 13.3%.

2. **Heat-backcast sensitivities.** Three scenarios (a primary backcast, a version excluding 2006 to 2007, and a percentile-heat definition for 2006 to 2007) verify that the headline finding does not depend on the backcast.

3. **Temperature-only percentile thresholds.** Twelve indicators (four geographies crossed with three percentiles): ZIP-overall, statewide-overall, ZIP-by-month, and ZIP-by-warm-season, each at the 95th, 97.5th, and 99th percentiles of daily Tmax. These serve as the consistency anchor for the primary extreme-heat interaction and address the concern that a health-informed heat index (the CalHeatScore) could be circular with emergency-department outcomes.

4. **Heat-metric decomposition.** ZIP-specific 95th-percentile flags for Tmax (peak daytime heat) and Tmin (overnight heat), and a 5th-percentile flag for diurnal temperature range (humid or stagnant heat), plus two-day persistence variants, fitted in a warm-season-only (May to September) sample with no temperature adjustment in the interaction model, so as not to adjust for the very metric used to define heat.

A four-group joint dust-by-heat exposure model (low or high dust crossed with heat or no heat, with high dust set at the statewide warm-season 90th percentile of DAF-PM_{10}, 35.37 μg/m³) is fitted in Section 4.5 as a non-multiplicative summary of the combined exposure.

2.8 Geographic assignment

Each ZIP was assigned to a region by a spatial join of its population-weighted centroid into the official California Air Resources Board (CARB) air-basin polygons. The join is intersection-based with a nearest-basin fallback for any unmatched centroid, and quality-assurance checks show no unclassified rows after the fallback. Every populated basin, including the Salton Sea and the Sacramento Valley, is fully represented and serves as the geographic stratification used throughout this chapter.

2.9 Sensitivity and robustness analyses

The chapter brings together the following sensitivity and robustness analyses, each presented in Section 4 or Section 5:

- A time-varying-confounding suite of five specifications relative to the primary model (linear Tmax, then nonlinear Tmax, then nonlinear Tmax plus the extreme-heat indicator, then wildfire-season exclusion, then respiratory-virus-season exclusion, then a two-day Tmax moving average).
- A concentration-response nonlinearity test (linear versus natural cubic spline with four degrees of freedom) on the same case-crossover sample, with three reference-point choices.
- A spline-based population attributable fraction, computed clamped and unclamped.
- A distributed-lag analysis at lags 0 to 6, using single-lag models, an unconstrained joint distributed-lag model, and a constrained distributed-lag nonlinear model with a natural-spline lag basis.
- An exposure-metric sensitivity across the three DAF-PM_{10} operationalisations.
- A heat-backcast sensitivity (three scenarios).
- A temperature-only threshold sensitivity (twelve indicators under two temperature-adjustment specifications), plus a consistency rebuild that corrected an analytic error described in Section 5.2.
- A heat-metric decomposition.
- Geographic effect modification across CARB air basins.

- Season heterogeneity (winter, spring, summer, fall) with a formal Cochran-Q heterogeneity test.
- A negative-control permutation (100,000-stratum subsample, 20 permutations).
- A review of the circularity concern for the health-informed heat index (the CalHeatScore).
- Meteorology-by-DAF-PM_{10} diagnostics, including covariate collinearity, a temperature-adjustment sensitivity (Section 5.6), and a humidity and joint-meteorological adjustment sensitivity (Section 5.7).

3. Population, Exposure Distributions, and Analytic Dataset

The combined emergency-department plus patient-discharge analytic file contains 186,352,957 case and referent rows across 42,402,018 case-days, distributed across roughly 23.97 million CVD strata (105.4 million rows after complete-covariate restriction) and 23.07 million RESP strata (101.3 million rows). Annual event counts grew nearly threefold over the study period, consistent with population growth and emergency-department visit-rate trends. The demographic composition of the case rows is summarized in Table 3.1, the cardiovascular cases are substantially older than the respiratory cases (mean 62 versus 31 years), a contrast that anticipates the cause-specific and age-patterned associations reported below and underscores the scale of the analytic sample.

Table 3.1. Demographic composition of the analytic cohort (case rows only).

Variable	Overall	CVD	RESP
Age (mean ± SD), years	45.1 ± 27.5	62.2 ± 17.6	31.1 ± 26.4
Sex: Female	23,251,509	13,275,094	12,722,281
Sex: Male	19,149,154	10,690,342	10,343,547
Ethnicity: Hispanic	14,262,678	6,113,356	9,081,667
Ethnicity: Non-Hispanic	27,188,580	17,365,371	13,455,089

The dust-attributable PM_{10} surface spans a wide range, and an interquartile range of 16.56 μg/m³ that matches the per-IQR scaling constant.

The extreme-heat indicator (CalHeatScore Level 3 or higher) has a prevalence of 13.3% across the case and referent rows, and the prevalence restricted to case-day rows only is essentially identical (also 13.3%), confirming that there is no severe selection on heat days between case-days and matched referent days.

Pairwise Spearman correlations on the case-crossover dataset (Figure 3.1) are moderate at most, DAF-PM_{10} with Tmax = 0.524; DAF-PM_{10} with the extreme-heat indicator = 0.248; DAF-PM_{10} with relative humidity = −0.530; Tmax with relative humidity = −0.585; and Tmax with the extreme-heat indicator = 0.560. Because the case-crossover conditions on the matched stratum, the relevant quantity is within-stratum residual collinearity rather than the marginal correlation, and the variance-inflation diagnostics reported in Section 5.6 do not indicate any instability.

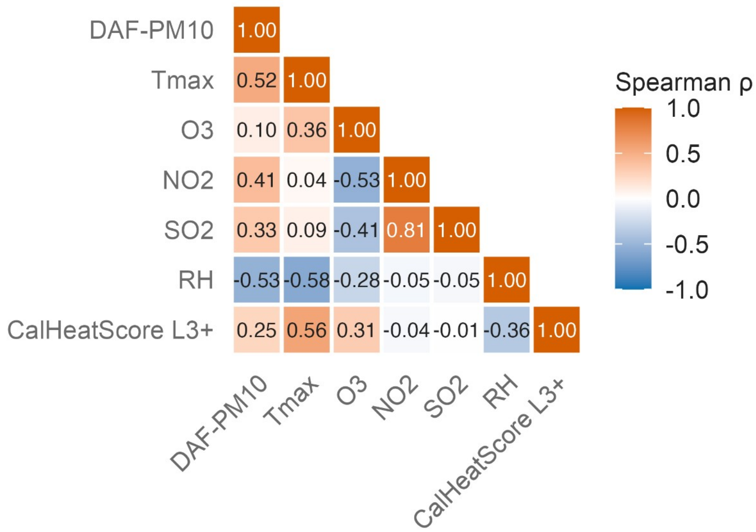

Figure 3.1. Spearman correlation matrix among the dust-attributable PM_{10} exposure, the gaseous co-pollutants, the meteorological covariates, and the binary extreme-heat indicator (CalHeatScore Level 3 or higher), on the case-crossover dataset. The dust exposure is moderately correlated with maximum temperature ($\rho = 0.524$) and inversely correlated with relative humidity ($\rho = -0.530$). No pair reaches a magnitude that would cause concern within the matched design.

4. Results

4.1 Same-day primary association and exposure-metric sensitivity

The same-day per-IQR associations are small, precisely estimated, and positive for both broad outcomes (Table 3.2; Figure 3.2). A one-IQR increase in DAF-PM_{10} is associated with a 0.21% increase in the odds of a broad cardiovascular encounter and a 0.96% increase in the odds of a broad respiratory encounter on the same day, with the respiratory effect roughly five times the cardiovascular effect on the log-odds scale. The patient-discharge-only estimates are directionally concordant but underpowered, with confidence intervals spanning the null.

Table 3.2. Same-day per-IQR (+16.558 μg/m³) DAF-PM_{10} odds ratios for broad outcomes, California 2006 to 2018 (all-events conditional logistic regression).

Outcome	Source	n strata	n rows	OR per IQR (95% CI)	p
CVD_any	ED+PDD	23,966,164	105,436,746	1.00213 (1.00104, 1.00323)	1.28×10^{-4}
CVD_any	ED only	23,654,707	104,067,133	1.00227	5.67×10^{-5}

Outcome	Source	n strata	n rows	OR per IQR (95% CI)	p
				(1.00117, 1.00338)	
CVD_any	PDD only	311,457	1,369,613	0.99622 (0.98940, 1.00308)	0.279
RESP_any	ED+PDD	23,066,558	101,272,878	1.00957 (1.00846, 1.01068)	1.18×10^{-64}
RESP_any	ED only	22,880,300	100,454,278	1.00971 (1.00859, 1.01083)	2.03×10^{-65}
RESP_any	PDD only	186,258	818,600	1.00156 (0.99281, 1.01040)	0.727

An internal cross-check confirmed that the equivalent conditional Poisson model reproduced the conditional logistic coefficients to nine decimal places, verifying the clogit and conditional-Poisson identity under the matched design.

The exposure-metric sensitivity confirms direction and significance across the three DAF-PM_{10} operationalisations (Table 3.3).

Table 3.3. Exposure-metric sensitivity: same-day per-IQR DAF-PM_{10} odds ratios across three operationalisations (combined ED+PDD).

Metric	CVD_any OR per IQR (95% CI)	RESP_any OR per IQR (95% CI)
Ensemble (primary)	1.00213 (1.00104, 1.00323)	1.00957 (1.00846, 1.01068)
Hierarchical	1.00438 (1.00326, 1.00549)	1.01013 (1.00898, 1.01128)
Calibrated ensemble (rescaled IQR)	1.15360 (1.09010, 1.22081)	1.27612 (1.20490, 1.35156)

The calibrated-ensemble metric is on a much smaller unit scale (mean approximately 0.81, IQR approximately 0.32), so the per-(16.558-unit) odds ratio is much larger; the matching direction and significance, not the magnitude, is the substantive claim.

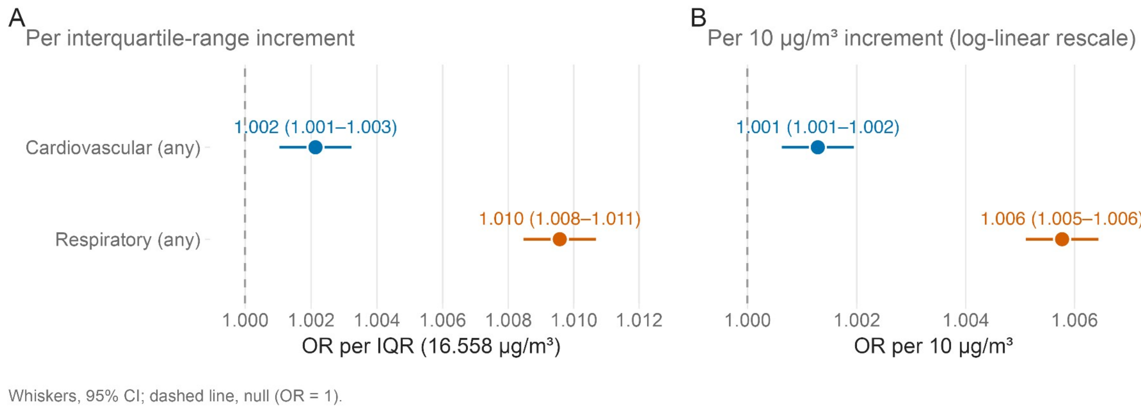

Figure 3.2. Same-day per-IQR (+16.558 µg/m³) DAF-PM_{10} odds ratios for broad cardiovascular and respiratory outcomes, combined emergency-department plus patient-discharge data, California 2006 to 2018. Points are odds ratios from all-events conditional logistic regression, horizontal bars are 95% confidence intervals; the vertical line marks the null. Both outcomes are positive and precisely estimated, with the respiratory effect roughly five times the cardiovascular effect on the log-odds scale.

4.2 Cause-specific decomposition

The cause-specific decomposition pinpoints both broad signals to biologically plausible subgroups (Table 3.4; Figure 3.3). The cardiovascular signal is, in volume, essentially a hypertensive-disease signal: 19.87 million of the 23.97 million CVD strata are hypertensive-disease strata. Ischemic heart disease is borderline positive, while cerebrovascular disease and ischemic stroke have slightly negative point estimates with confidence intervals crossing the null. The respiratory signal is ordered asthma greater than acute respiratory infection and COPD, with pneumonia null. The pneumonia null is mechanistically consistent with the slower clinical evolution of bacterial pneumonia relative to a same-day case-crossover design, a point developed in Section 6.

Table 3.4. Same-day per-IQR DAF-PM_{10} odds ratios for cause-specific outcomes, combined ED+PDD.

Outcome	n strata	OR per IQR (95% CI)	p
Ischemic heart disease	3,259,264	1.00288 (0.99995, 1.00581)	0.054
Myocardial infarction	134,876	1.00114 (0.98783,	0.868

Outcome	n strata	OR per IQR (95% CI)	p
		1.01463)	
Heart failure	2,014,706	1.00212 (0.99853, 1.00572)	0.247
Conduction or dysrhythmia	3,057,517	1.00188 (0.99883, 1.00493)	0.227
Hypertensive disease	19,872,959	1.00220 (1.00100, 1.00340)	3.3×10^{-4}
Cerebrovascular disease	933,603	0.99562 (0.99015, 1.00112)	0.119
Ischemic stroke	409,077	0.99302 (0.98492, 1.00119)	0.094
Pneumonia	1,516,881	1.00119 (0.99684, 1.00556)	0.592
COPD	8,159,706	1.00736 (1.00541, 1.00932)	1.33×10^{-13}
Asthma	7,094,247	1.01280 (1.01081, 1.01479)	9.7×10^{-37}
Acute respiratory infection	10,241,185	1.00734 (1.00566, 1.00902)	7.3×10^{-18}

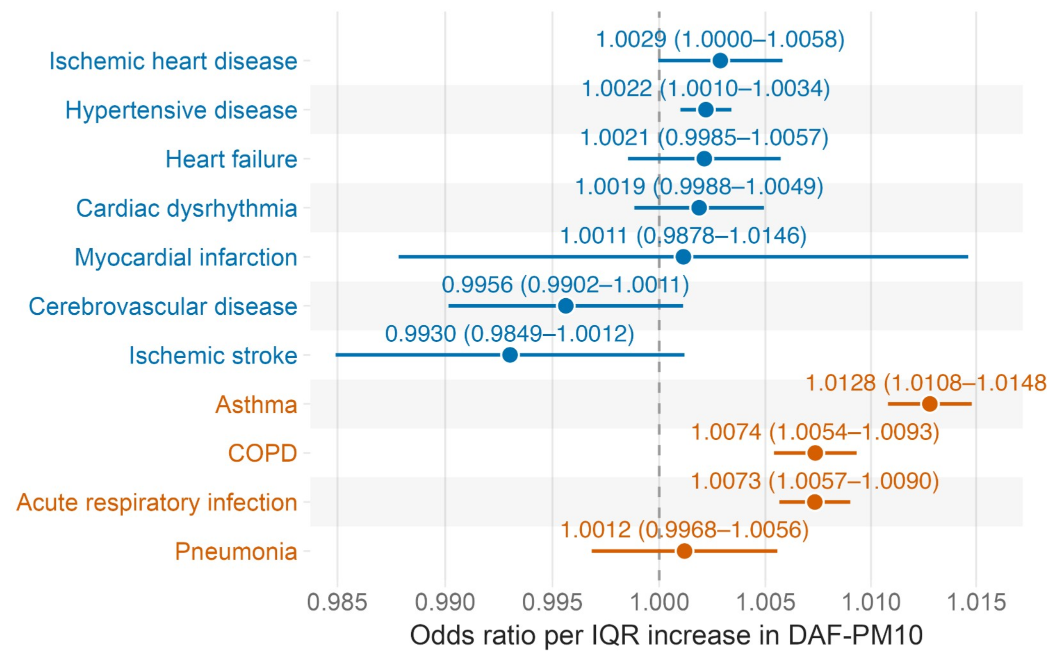

Figure 3.3. Cause-specific same-day per-IQR DAF-PM_{10} odds ratios, combined ED+PDD. Cardiovascular subgroups (top) and respiratory subgroups (bottom) are shown with 95% confidence intervals. The cardiovascular signal is concentrated in hypertensive disease, with

cerebrovascular and ischemic-stroke estimates null; the respiratory signal is largest for asthma, followed by acute respiratory infection and COPD, with pneumonia null.

Cause-specific heat-interaction estimates extend the headline heat finding to the disease-subgroup level (combined ED+PDD): heat attenuates hypertensive disease (interaction OR 0.99580; $p = 1.8 \times 10^{-4}$) and ischemic heart disease (0.99380; $p = 0.022$), and amplifies asthma (interaction OR 1.00825; $p = 1.7 \times 10^{-5}$), COPD (1.00592; $p = 9.9 \times 10^{-4}$), and acute respiratory infection (1.01182; $p = 1.4 \times 10^{-10}$), while leaving pneumonia unmodified ($p = 0.58$). The cause-specific pattern is internally consistent with the broad-cardiovascular attenuation and broad-respiratory amplification reported in Section 4.5.

4.3 Concentration-response and population attributable fraction

A formal likelihood-ratio test of a linear versus a natural-cubic-spline (four degrees of freedom) exposure term shows strong nonlinearity for respiratory disease and shallow but detectable curvature for cardiovascular disease (Table 3.5).

Table 3.5. Formal concentration-response nonlinearity test (combined ED+PDD).

Outcome	Log-likelihood linear	Log-likelihood spline	LR statistic	df	p (nonlinearity)
CVD_any	−35,358,891.20	−35,358,887.02	8.36	3	0.0392
RESP_any	−33,984,304.72	−33,984,237.17	135.09	3	4.33×10^{-29}

Knots were placed at full-cohort DAF-PM_{10} percentiles (boundary at the 5th and 95th percentiles, 2.73 and 39.30 µg/m³; interior at the 25th, 50th, and 75th). The median-referenced curves (Table 3.6; Figure 3.4) show that the respiratory exposure-response is clearly U-shaped, a dip significantly below 1 between roughly the 15th and 45th exposure percentiles, followed by a steep rise above the median. The cardiovascular curve is essentially monotone-rising with shallow curvature.

Table 3.6. Median-referenced concentration-response odds ratios at selected exposure percentiles (combined ED+PDD).

Outcome	OR at P5	OR at P25	OR at P50	OR at P75	OR at P95	OR at P99
CVD_any	0.99583 (0.99372, 0.99796)	0.99807 (0.99666, 0.99948)	1.000 (ref)	1.00193 (1.00080, 1.00306)	1.00279 (1.00074, 1.00484)	1.00359 (1.00085, 1.00633)
RESP_any	1.00138 (0.99921, 1.00355)	0.99853 (0.99695, 1.00010)	1.000 (ref)	1.00937 (1.00793, 1.01081)	1.01873 (1.01666, 1.02081)	1.02723 (1.02442, 1.03004)

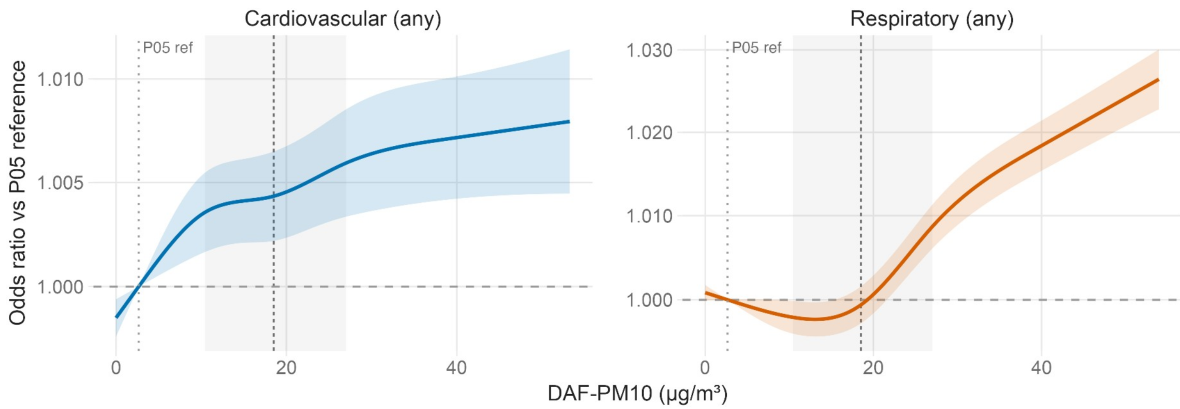

Figure 3.4. Concentration-response curves for broad cardiovascular and respiratory outcomes, referenced to the population median DAF-PM_{10} (18.52 µg/m³). Curves are natural cubic splines (four degrees of freedom); shaded bands are 95% confidence intervals. The respiratory curve is U-shaped, dipping below 1 between roughly the 15th and 45th percentiles and rising steeply above the median; the cardiovascular curve is near-monotone with shallow curvature.

The two curve shapes have different implications for the population attributable fraction, which we computed under three counterfactual reference choices (Table 3.7). Because the respiratory curve is U-shaped, these fractions are derived from the fitted concentration-response rather than from a single linear slope, following the attributable-risk framework for flexible exposure-response and lag models (Gasparrini & Leone, 2014).

Table 3.7. Spline-based population attributable fractions by counterfactual reference (combined ED+PDD).

Outcome	Counterfactual	PAF % (95% CI)
CVD_any	5th percentile (preferred, in-support)	0.42 (0.24, 0.61)

Outcome	Counterfactual	PAF % (95% CI)
CVD_any	25th percentile	0.12 (0.03, 0.23)
CVD_any	Zero (extrapolated)	0.56 (0.32, 0.82)
RESP_any	5th percentile (preferred, in-support)	0.43 (0.32, 0.55)
RESP_any	25th percentile	0.58 (0.48, 0.68)
RESP_any	Zero (extrapolated)	0.38 (0.26, 0.52)

For respiratory disease, the clamped one-sided burden at the 5th-percentile reference (0.46%; 95% CI 0.35, 0.60) is larger than the unclamped version (0.39%; 0.22, 0.58) when both are computed on the lag-restricted analytic sample used for the distributed-lag extensions (n = 23.04 million strata); the corresponding value on the full primary sample reported in Table 3.7 is 0.43% (0.32, 0.55). The clamped version is larger because it ignores the below-1 dip between zero and the median (Figure 3.5). We report the clamped, 5th-percentile-referenced version as primary because it is one-sided (burden above the reference) and uses a counterfactual within the empirical exposure support. Cause-specific approximate attributable fractions (5th-percentile-referenced, log-linear) place asthma highest at 1.28% (1.09%, 1.48%), followed by COPD and acute respiratory infection (each 0.74%) and hypertensive disease (0.22%)

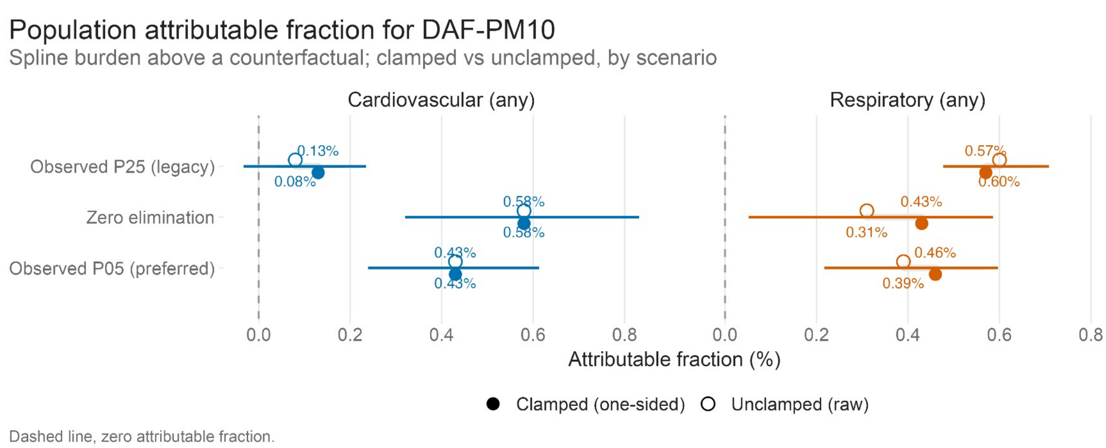

Figure 3.5. Spline-based population attributable fractions for broad outcomes under clamped (one-sided, burden above the reference) and unclamped definitions, at the 5th-percentile

counterfactual. For respiratory disease the clamped fraction exceeds the unclamped fraction, because the unclamped version nets out the below-1 segment of the U-shaped curve between zero and the median.

4.4 Distributed-lag effects

The distributed-lag analysis adds an important quantitative refinement (Table 3.8; Figure 3.6). The respiratory effect peaks one day after exposure, remains highly significant through lag 4, and turns significantly negative at lag 6. We interpret the negative lag-6 coefficient as short-term event displacement (harvesting) (Schwartz, 2000), in which susceptible individuals are tipped earlier by high-dust days, leaving a slightly depleted at-risk pool a week later, we do not interpret it as a protective effect. The cardiovascular effect is essentially same-day, with a small lag-3 to lag-5 rebound.

Table 3.8. Single-lag per-IQR DAF-PM_{10} odds ratios, lags 0 to 6 (separate fit per lag, combined ED+PDD).

Lag	CVD_any OR per IQR (95% CI)	p	RESP_any OR per IQR (95% CI)	p
0	1.00218 (1.00108, 1.00327)	9.4×10^{-5}	1.00960 (1.00849, 1.01071)	6.6×10^{-65}
1	1.00056 (0.99960, 1.00152)	0.255	1.00986 (1.00889, 1.01083)	1.0×10^{-88}
2	1.00090 (0.99998, 1.00181)	0.055	1.00884 (1.00792, 1.00977)	1.0×10^{-79}
3	1.00152 (1.00062, 1.00243)	9.3×10^{-4}	1.00681 (1.00591, 1.00772)	1.0×10^{-49}
4	1.00158 (1.00068, 1.00248)	5.5×10^{-4}	1.00382 (1.00293, 1.00472)	5.4×10^{-17}
5	1.00154 (1.00064, 1.00244)	7.5×10^{-4}	1.00082 (0.99993, 1.00171)	0.072
6	0.99994 (0.99904, 1.00083)	0.888	0.99742 (0.99653, 0.99831)	1.5×10^{-8}

The unconstrained joint distributed-lag model (all seven lags simultaneously) gives a cardiovascular cumulative odds ratio of 1.00312 (1.00153, 1.00471) and a respiratory cumulative odds ratio of 1.01360 (1.01199, 1.01522). The constrained distributed-lag nonlinear model (natural-spline lag basis, four degrees of freedom) (Gasparrini et al., 2010; Gasparrini, 2014) gives 1.00079 (0.99926, 1.00232) for cardiovascular disease and 1.01143 (1.00989, 1.01298) for respiratory disease. The respiratory cumulative effect is roughly 1.4 times the same-day effect, the cardiovascular effect is essentially same-day, with the slightly larger joint cumulative reflecting the small lag-3 to lag-5 rebound.

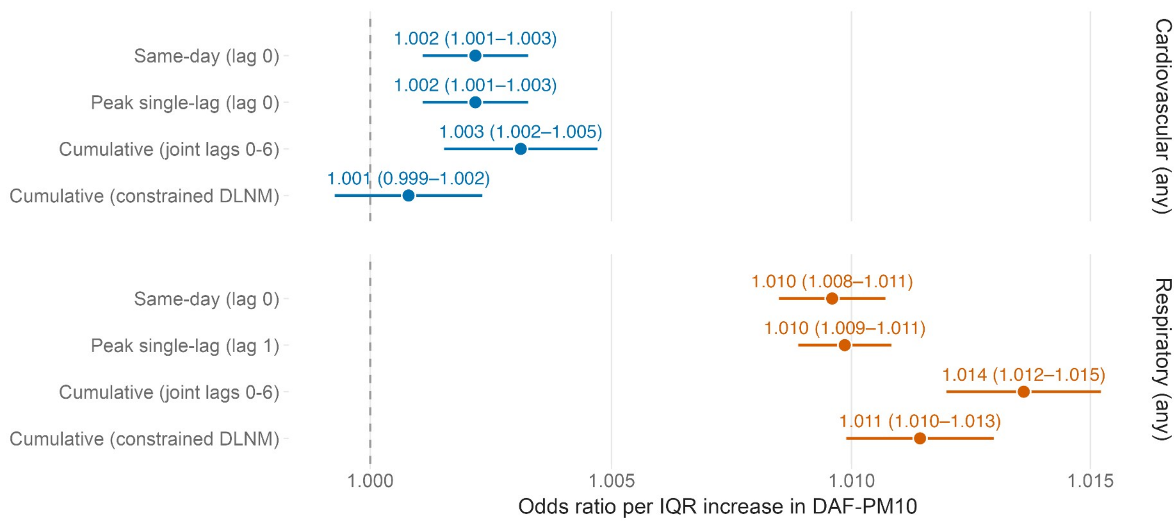

Figure 3.6. Distributed-lag per-IQR DAF-PM_{10} odds ratios at lags 0 to 6 for broad cardiovascular and respiratory outcomes (combined ED+PDD). The respiratory effect peaks at lag 1, persists through lag 4, and crosses below the null at lag 6 (consistent with short-term event displacement); the cardiovascular effect is concentrated at lag 0.

4.5 Heat interaction: primary, sensitivity, and decomposition

The heat-interaction story is presented in the order in which it was conducted, primary indicator, then backcast sensitivity, then temperature-only threshold sensitivity, then heat-metric decomposition. The primary extreme-heat interaction (Table 3.9; Figure 3.7) shows respiratory amplification and cardiovascular attenuation on high-heat days. On extreme-heat days, the joint per-IQR respiratory odds ratio rises to 1.01496, while the joint cardiovascular odds ratio falls to 0.99889.

Table 3.9. Primary DAF-PM_{10} by extreme-heat (CalHeatScore Level 3 or higher) interaction (combined ED+PDD).

Outcome	Main (no heat) OR per IQR	Interaction OR per IQR	Joint (heat) OR per IQR	p (interaction)
CVD_any	1.00260 (1.00144, 1.00376)	0.99629 (0.99431, 0.99829)	0.99889 (0.99693, 1.00084)	2.7×10^{-4}
RESP_any	1.00853 (1.00737, 1.00969)	1.00637 (1.00416, 1.00859)	1.01496 (1.01276, 1.01716)	1.5×10^{-8}

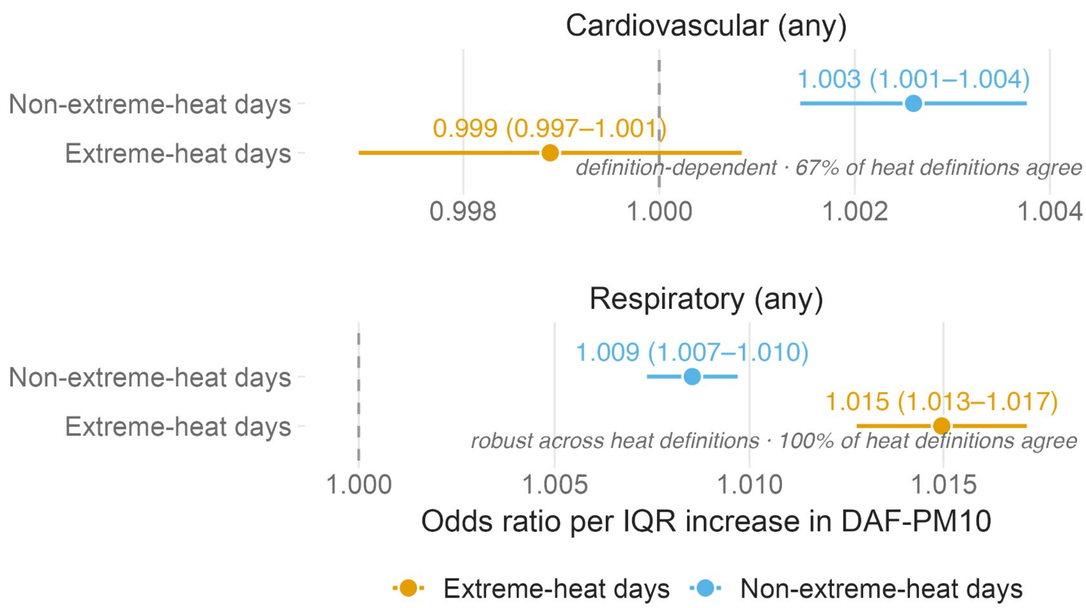

Figure 3.7. DAF-PM_{10} per-IQR odds ratios on non-heat days, on extreme-heat days, and the multiplicative interaction term, for broad cardiovascular and respiratory outcomes (combined ED+PDD). Extreme heat (CalHeatScore Level 3 or higher) amplifies the respiratory association and attenuates the cardiovascular association.

The heat-backcast sensitivity confirms that neither finding depends on back-casting the 2008 thresholds to 2006 and 2007 (Table 3.10).

Table 3.10. Heat-backcast sensitivity: main effect and interaction (combined ED+PDD).

Scenario	CVD_any OR (main)	RESP_any OR (main)	CVD_any OR (interaction)	RESP_any OR (interaction)
Primary backcast	1.00213 (1.00104, 1.00323)	1.00957 (1.00846, 1.01068)	0.99629 (0.99431, 0.99829)	1.00637 (1.00416, 1.00859)
Exclude 2006 to 2007	1.00199 (1.00085, 1.00314)	1.00896 (1.00778, 1.01013)	0.99601 (0.99394, 0.99808)	1.00486 (1.00254, 1.00720)

Scenario	CVD_any OR (main)	RESP_any OR (main)	CVD_any OR (interaction)	RESP_any OR (interaction)
Percentile heat for 2006 to 2007	1.00213 (1.00104, 1.00323)	1.00957 (1.00846, 1.01068)	0.99687 (0.99493, 0.99883)	1.00741 (1.00525, 1.00956)

The temperature-only threshold sensitivity addresses the concern that a health-informed heat index could be circular with emergency-department outcomes (Table 3.11). For the respiratory amplification, all six central temperature-only models match the extreme-heat direction on the powered analyses, several of them more strongly than the extreme-heat indicator itself. For the cardiovascular attenuation, four of six central models match on linear-Tmax analyses, but the per-model diagnostic shows that the higher-percentile and warm-season families flip to slight amplification on multiple cutoffs. The honest framing is that respiratory amplification is robust, whereas cardiovascular attenuation is definition-dependent.

Table 3.11. Temperature-only consistency summary (central families: ZIP-overall and statewide-overall percentile heat).

Outcome	Source	Same-direction central models (of 6)	Any significant same-direction	Interpretation
CVD_any	All	4	Yes (statewide p95)	Consistent with extreme-heat indicator
CVD_any	ED	4	Yes	Consistent
CVD_any	PDD	2	No	Not consistent
RESP_any	All	6	Yes	Consistent
RESP_any	ED	6	Yes	Consistent
RESP_any	PDD	6	Yes	Consistent

The heat-metric decomposition refines the single binary interaction into mechanism-specific findings (Table 3.12; Figure 3.8). Peak daytime heat (Tmax) drives the respiratory amplification, overnight heat (Tmin) drives the cardiovascular attenuation, and humid or stagnant heat (low diurnal temperature range) reverses both.

Table 3.12. Heat-metric decomposition: DAF-PM_{10} by heat interactions, warm season (May to September), per warm-season IQR (15.225 µg/m³); model omits temperature adjustment.

Heat metric	CVD_any ratio (95% CI)	p (interaction)	RESP_any ratio (95% CI)	p (interaction)
Tmax p95, 1-day	0.99851 (0.99457, 1.00247)	0.470	1.00823 (1.00357, 1.01292)	5.2×10^{-4}
Tmax p95, 2-day	0.99933 (0.99466, 1.00403)	0.768	1.00444 (0.99931, 1.00961)	0.090
Tmin p95, 1-day	0.99383 (0.99031, 0.99736)	5.6×10^{-4}	1.00152 (0.99744, 1.00563)	0.469
Tmin p95, 2-day	0.99164 (0.98768, 0.99562)	2.6×10^{-5}	1.00153 (0.99692, 1.00615)	0.52
Low DTR p5, 1-day	1.00611 (1.00177, 1.01047)	6.1×10^{-3}	0.99189 (0.98711, 0.99670)	8.8×10^{-4}
Low DTR p5, 2-day	1.01064 (1.00455, 1.01677)	7.2×10^{-4}	0.99304 (0.98640, 0.99973)	0.045
Region-tailored composite	1.00020 (0.99611, 1.00432)	0.920	0.99529 (0.99089, 0.99971)	0.036

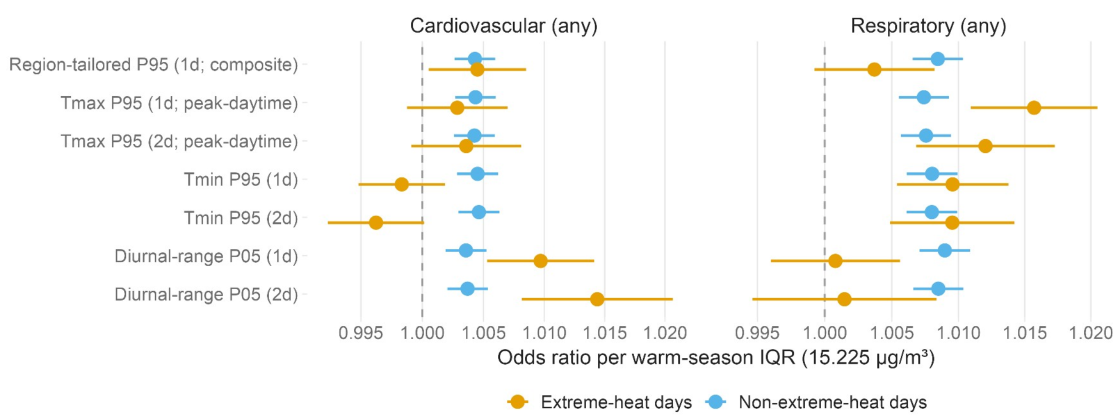

Figure 3.8. Heat-metric decomposition of the DAF-PM_{10} by heat interaction in the warm season. Peak daytime heat (Tmax, 95th percentile) amplifies the respiratory association, overnight heat (Tmin, 95th percentile) attenuates the cardiovascular association, humid or stagnant heat (low diurnal temperature range, 5th percentile) reverses both. The region-tailored composite averages over the three regimes.

The four-group joint dust-by-heat exposure model (Table 3.13; Figure 3.9) shows that, for respiratory disease, joint high-dust and extreme-heat exposure acts slightly sub-additively on the log-odds scale (observed joint coefficient 0.011 versus an additive prediction of 0.015). For cardiovascular disease, the joint cell is essentially null.

Table 3.13. Four-group joint dust-by-heat exposure model, warm season; reference is low dust and no heat. High dust is at or above the statewide warm-season 90th percentile (35.37 µg/m³); heat is the region-tailored modifier.

Outcome	High dust, no heat	Low dust, heat	High dust, heat	Sum-of-standalones coefficient	Observed joint coefficient
RESP_any	1.0059 ($p = 1.4 \times 10^{-3}$)	1.0094 ($p = 6.9 \times 10^{-6}$)	1.0110 ($p = 1.3 \times 10^{-2}$)	0.01522	0.01100
CVD_any	1.00225 ($p = 0.19$)	1.00004 ($p = 0.98$)	0.99969 ($p = 0.94$)	0.00229	−0.00031

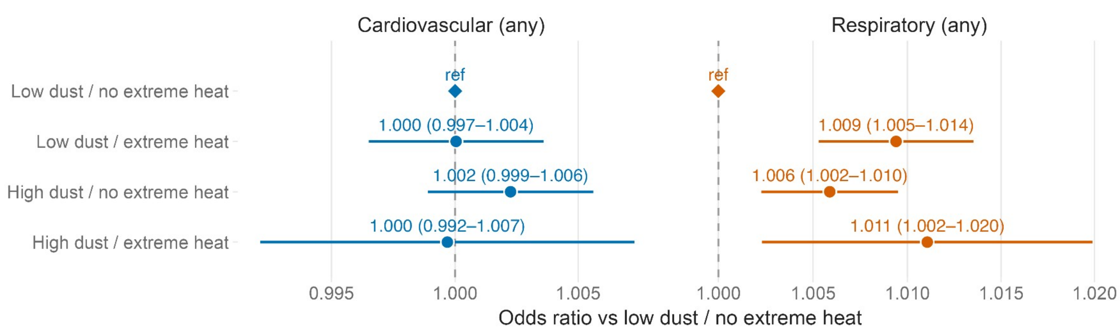

Figure 3.9. Four-group joint dust-by-heat exposure odds ratios relative to low dust and no heat, warm season (combined ED+PDD). For respiratory disease, the high-dust and high-heat cell is positive but falls below the additive sum of the two standalone exposures, indicating slight sub-additivity; for cardiovascular disease, the joint cell is null.

4.6 Geographic effect modification (CARB air basins)

Geographic stratification on official CARB air basins (Table 3.14; Figures 3.10 and 3.11) shows two clear patterns. Respiratory amplification is positive and statistically significant across every populated basin except the chronically dust-saturated Salton Sea, though its magnitude varies roughly three-fold, from 0.9% in the San Francisco Bay basin to 3.1% in San Diego — the only basin notably stronger than the South Coast. San Joaquin Valley, Sacramento Valley, Mojave, and the Bay Area fall in the same range as the South Coast, so the respiratory dust effect is essentially a statewide phenomenon. Cardiovascular effects, by contrast, are concentrated in the coastal metropolitan basins (South Coast, San Francisco Bay, and San Diego), while the interior dust-source basins and the Salton Sea are null or slightly negative.

Table 3.14. Per-IQR DAF-PM_{10} odds ratios by official CARB air basin (combined ED+PDD).

Air basin	CVD OR (95% CI)	p	RESP OR (95% CI)	p
South Coast (Los Angeles)	1.0074 (1.0052, 1.0096)	$< 10^{-4}$	1.0114 (1.0093, 1.0136)	$< 10^{-4}$
San Francisco Bay	1.0033 (1.0009, 1.0058)	0.008	1.0086 (1.0059, 1.0113)	$< 10^{-4}$
San Diego	1.0067 (1.0017, 1.0118)	0.009	1.0310 (1.0251, 1.0369)	$< 10^{-4}$
San Joaquin Valley	0.9986 (0.9959, 1.0014)	0.338	1.0154 (1.0129, 1.0180)	$< 10^{-4}$
Sacramento Valley	0.9982 (0.9953, 1.0010)	0.204	1.0106 (1.0075, 1.0137)	$< 10^{-4}$
Mojave Desert	0.9969 (0.9904, 1.0035)	0.361	1.0092 (1.0033, 1.0151)	0.002
Salton Sea	1.0007 (0.9961, 1.0053)	0.774	0.9995 (0.9949, 1.0040)	0.815
South Central Coast	0.9994 (0.9923, 1.0066)	0.880	1.0059 (0.9984, 1.0134)	0.126
North Coast	0.9977 (0.9886, 1.0070)	0.632	1.0069 (0.9964, 1.0174)	0.200
Mountain Counties	0.9946 (0.9843, 1.0049)	0.303	1.0081 (0.9957, 1.0206)	0.202
Great Basin Valleys	1.0105 (0.9736, 1.0488)	0.582	1.0095 (0.9694, 1.0513)	0.647
North Central Coast	0.9973 (0.9855, 1.0092)	0.652	0.9962 (0.9853, 1.0073)	0.503
Lake County	1.0058 (0.9891, 1.0228)	0.500	0.9980 (0.9792, 1.0172)	0.838
Lake Tahoe	1.0239 (0.9756, 1.0745)	0.339	1.0024 (0.9553, 1.0519)	0.922
Northeast Plateau	1.0016 (0.9861, 1.0174)	0.842	1.0053 (0.9861, 1.0249)	0.593

California's most acutely dust-impacted region, the Salton Sea, shows no detectable case-crossover signal in either direction (CVD OR 1.0007, p = 0.77; RESP OR 0.9995, p = 0.82). This is most simply explained as a design limitation, the case-crossover relies on within-stratum day-to-day exposure variation, and in a chronically saturated dust regime that variation is small relative to the absolute exposure level. We do not interpret this as evidence of no biological effect, the appropriate designs for that question are cohort and long-term ecological studies, and we treat the null as motivating one of the most important future directions of the thesis (Section 7.1).

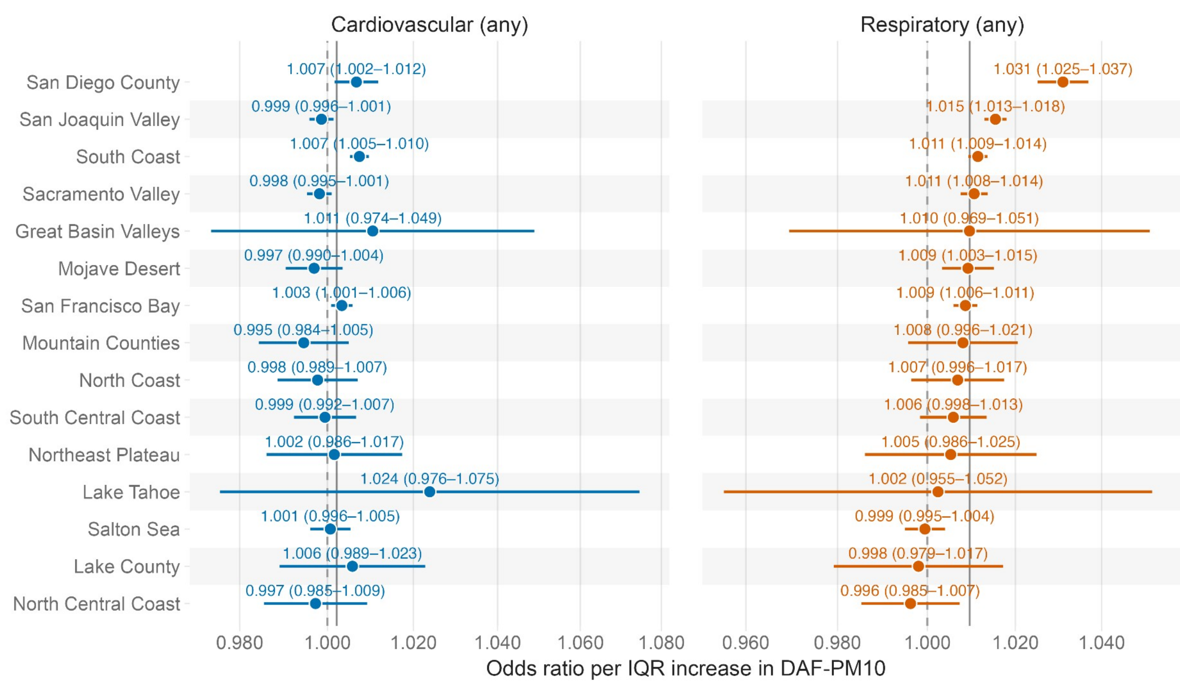

Figure 3.10. Per-IQR DAF-PM_{10} odds ratios by official CARB air basin, for broad cardiovascular (top) and respiratory (bottom) outcomes (combined ED+PDD). Respiratory amplification is positive across populated basins but varies about three-fold (San Diego strongest); cardiovascular effects are concentrated in coastal metropolitan basins; the Salton Sea is null for both.

The same basin-level respiratory estimates are mapped in Figure 3.11, where the choropleth makes the geographic reach of the association immediate, the amplification is not confined to the dust-source interior but extends across the populated coast, with San Diego the most pronounced.

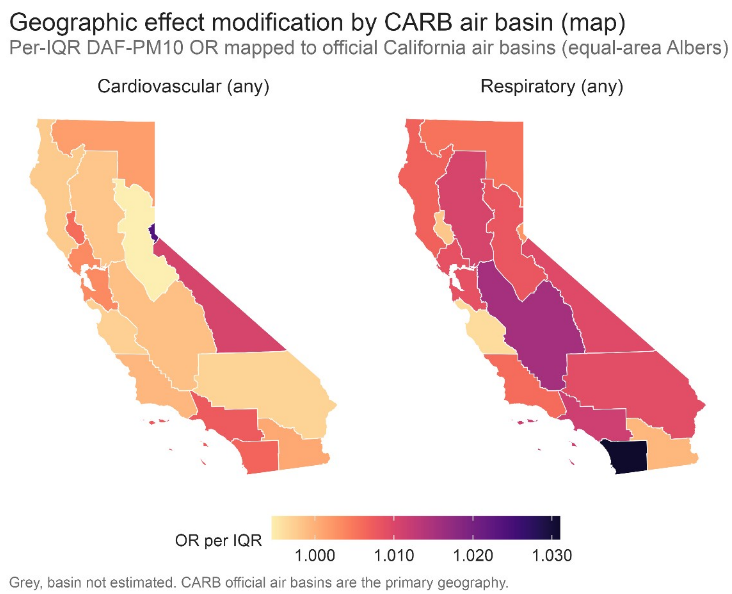

Figure 3.11. Geographic distribution of basin-specific per-IQR respiratory DAF-PM_{10} odds ratios across California CARB air basins (choropleth). Warmer shading denotes stronger amplification. The map emphasises the statewide reach of the respiratory association and the prominence of San Diego and the populated interior valleys.

4.7 Season heterogeneity

Season-stratified estimates (Table 3.15; Figure 3.12) show no seasonal heterogeneity for cardiovascular disease (Cochran-Q $p = 0.33$) but strong heterogeneity for respiratory disease ($p < 10^{-5}$), with summer largest, then winter, fall, and spring. The summer-largest pattern points in the same direction as the Tmax-amplification finding from the heat-metric decomposition, so the season-stratified analysis is independently consistent with peak daytime heat as the active modifier on the respiratory side.

Table 3.15. Season-stratified per-IQR DAF-PM_{10} odds ratios with Cochran-Q heterogeneity test (combined ED+PDD).

Season	CVD_any OR (95% CI)	p	RESP_any OR (95% CI)	p
Winter	1.0009 (0.9984,	0.478	1.0106 (1.0084,	1.4×10^{-20}

Season	CVD_any OR (95% CI)	p	RESP_any OR (95% CI)	p
	1.0035)		1.0129)	
Spring	0.9996 (0.9969, 1.0023)	0.781	1.0054 (1.0028, 1.0081)	4.8×10^{-5}
Summer	1.0019 (0.9996, 1.0042)	0.106	1.0143 (1.0116, 1.0170)	2.7×10^{-25}
Fall	1.0025 (1.0008, 1.0042)	0.005	1.0075 (1.0057, 1.0093)	9.2×10^{-17}
Cochran Q (df = 3)	3.42	n/a	26.41	n/a
p (heterogeneity)	0.332	n/a	7.8×10^{-6}	n/a

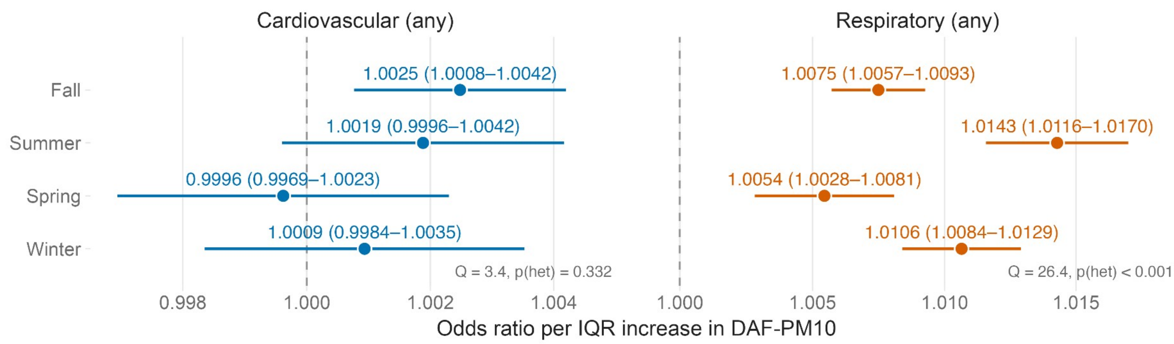

Figure 3.12. Season-stratified per-IQR DAF-PM_{10} odds ratios for broad cardiovascular and respiratory outcomes (combined ED+PDD). Cardiovascular estimates are homogeneous across seasons; respiratory estimates are strongly heterogeneous, with summer largest, consistent with peak-daytime-heat amplification.

5. Internal Verification, Robustness, and Audit

This section consolidates the checks that bear on the credibility of the primary findings. We are transparent about analytic corrections made during the project; each is described with its substantive consequence.

5.1 Specification sensitivities (time-varying confounding)

The primary direction and significance are preserved across all five time-varying-confounding specifications (Table 3.16; Figure 3.13). Adding nonlinear temperature slightly attenuates the cardiovascular estimate (roughly 25% attenuation in the log-odds magnitude) but does not cross the null; excluding the wildfire season slightly lowers the cardiovascular lower bound (to 1.00001) without changing inference. We retain the primary linear-Tmax specification as the substantive estimate and report the nonlinear-temperature version as the most informative sensitivity.

Table 3.16. Per-IQR DAF-PM_{10} odds ratios under the time-varying-confounding sensitivity suite (combined ED+PDD).

Specification	**CVD_any OR (95% CI)**	**RESP_any OR (95% CI)**
Primary (linear Tmax)	1.00213 (1.00104, 1.00323)	1.00957 (1.00846, 1.01068)
Nonlinear Tmax (spline, df = 4)	1.00166 (1.00056, 1.00276)	1.01003 (1.00891, 1.01115)
Nonlinear Tmax + extreme-heat indicator	1.00165 (1.00055, 1.00276)	1.01005 (1.00893, 1.01117)
Exclude wildfire season (Aug 15 to Dec 15)	1.00158 (1.00001, 1.00315)	1.01045 (1.00889, 1.01201)
Exclude respiratory-virus season (Dec to Mar)	1.00208 (1.00083, 1.00333)	1.00889 (1.00755, 1.01023)
Two-day Tmax moving average	1.00185 (1.00074, 1.00295)	1.00916 (1.00804, 1.01028)

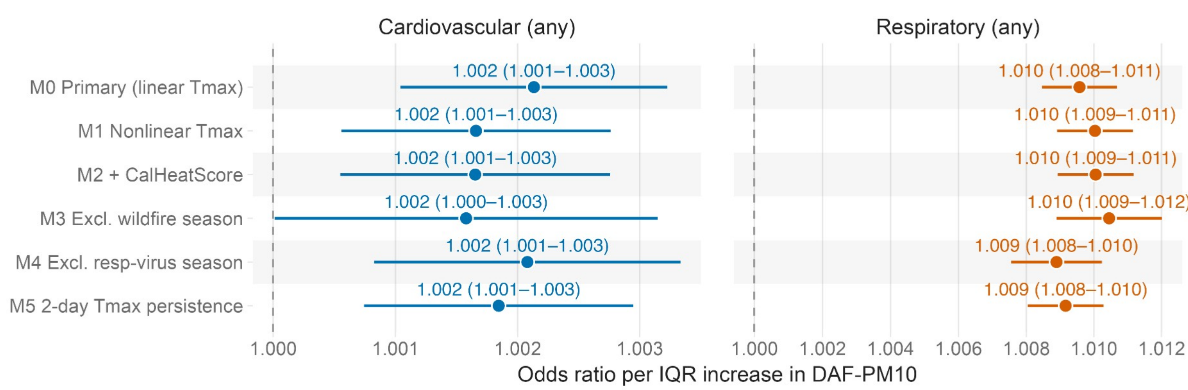

Figure 3.13. Per-IQR DAF-PM_{10} odds ratios for broad cardiovascular and respiratory outcomes across the time-varying-confounding specifications. All specifications preserve the direction and significance of the primary estimate.

5.2 Heat-definition sensitivities and a corrected consistency summary

The temperature-only threshold sensitivity tested twelve heat definitions (four geographies crossed with three percentiles) under both linear-Tmax and nonlinear-Tmax adjustment.

summary is the one reported in Section 4.5 (Table 3.11). The interpretation is that respiratory amplification is consistent across all six central temperature-only models on the powered analyses, while cardiovascular attenuation is consistent on four of six central models under linear-Tmax adjustment but is threshold-family dependent and fragile under nonlinear-Tmax adjustment.

5.3 Temperature-only consistency at scale

The respiratory amplification reproduces across the temperature-only family with substantially greater statistical significance than the extreme-heat indicator itself. For example, the ZIP-overall 95th-percentile linear-Tmax model gives a respiratory interaction odds ratio of 1.0122 (1.0089, 1.0155), $p = 3.2 \times 10^{-13}$, on the emergency-department data, and the ZIP-overall 97.5th-percentile model gives 1.0155 (1.0110, 1.0201), $p = 1.6 \times 10^{-11}$. These estimates are independent of the health-informed heat index (the CalHeatScore) and confirm the respiratory amplification on purely temperature-defined grounds.

5.4 Permutation negative control

Twenty permutations on a random 100,000-stratum subset of respiratory case strata produced a null distribution properly centred near zero on the coefficient scale (mean coefficient 2.45×10^{-4}; standard deviation 5.37×10^{-4}; 95% null interval for the per-IQR odds ratio 0.9914 to 1.0233). Because the permutation ran at roughly 1/230th of the analytic sample, its null is about fifteen times wider than the full-sample sampling distribution and cannot by itself falsify the primary estimate; the primary full-sample respiratory coefficient (5.75×10^{-4}) is 17 standard errors above zero on its own sampling distribution. This is a pipeline-validation diagnostic at reduced sample size rather than a formal falsification at full sample size, the full-sample permutation would have required several weeks of additional computation and was outside the project timeline.

5.5 The person-identifier limitation

The analysis-ready person identifier has only about 33,000 unique values across 42 million case-days, with the same identifier spanning emergency-department and patient-discharge records and the full study period. It is therefore not a globally unique person identifier, and three analyses that were originally planned (person-clustered robust standard errors, a first-event-per-person restriction, and a referent-overlap diagnostic) cannot be implemented as intended. The substantive primary estimate is not affected, because the matched strata are independent by design. Within-person sensitivities are identified as a high-priority future direction (Section 7.8).

5.6 Temperature collinearity and the temperature-adjustment sensitivity

Because temperature was one of the predictors used to generate the Aim 2 exposure surface (Section 2.3), a natural question is whether adjusting the health model for temperature induces collinearity instability or over-adjusts away part of the dust signal. The diagnostics show that the model is well conditioned, variance-inflation factors in the primary covariate set are 1.77 for DAF-PM_{10}, 1.83 for Tmax, 2.01 for ozone, 2.83 for nitrogen dioxide, 1.81 for sulphur dioxide, and 1.69 for relative humidity, all far below conventional warning levels, and the condition number of the standardised covariate matrix is 3.39, well under the ill-conditioning threshold near 30. The dust-temperature association is moderate (overall Spearman $\rho = 0.524$), leaving roughly 73% of DAF-PM_{10} variance unshared with Tmax, which is ample independent contrast to identify a distinct dust coefficient.

Refitting the primary same-day model with no temperature term, with linear Tmax (the primary), and with nonlinear Tmax shows that the estimate is significant and correctly signed under all three treatments (Table 3.17). The unadjusted estimate is the largest, and adding temperature reduces it. Because Tmax helped build the exposure surface, the temperature term absorbs part of the dust signal, so the primary estimate is the conservative one and the meteorological double-counting biases toward the null.

Table 3.17. DAF-PM_{10} per-IQR odds ratio under three temperature-adjustment specifications (combined ED+PDD).

Outcome	No temperature	Linear Tmax (primary)	Nonlinear Tmax
CVD_any	1.00434 (1.00326, 1.00543)	1.00213 (1.00104, 1.00323)	1.00179 (1.00069, 1.00289)
RESP_any	1.01054 (1.00944, 1.01164)	1.00957 (1.00846, 1.01068)	1.00968 (1.00857, 1.01080)

The variance-inflation and condition-number diagnostics are computed on the final covariate values and are blind to the dependence created at exposure-training time. These diagnostics are therefore a necessary but not a sufficient check, so we base the conservativeness argument on the direction of the change when meteorological adjustment is relaxed rather than on the diagnostics alone. The definitive remedy is to re-derive a temperature-free exposure surface upstream (Section 8.8), further downstream temperature adjustment is not indicated, as it would only deepen the toward-null bias.

5.7 Humidity-adjustment sensitivity

Humidity is the natural companion to the temperature check, because specific humidity was also among the more influential predictors of the Aim 2 surface (Section 2.3) while the health model adjusts for relative humidity. Four nested specifications were compared (Table 3.18; Figure 3.14): the primary model (linear temperature plus linear relative humidity); the same model with relative humidity dropped; the same model with relative humidity modelled nonlinearly; and a fully meteorology-stripped model that drops both temperature and relative humidity.

Table 3.18. DAF-PM_{10} per-IQR odds ratio under four humidity and meteorological-adjustment specifications (combined ED+PDD).

Outcome	Primary (linear temp + linear RH)	Drop RH, keep temp	Temp + nonlinear RH	Drop temp and RH
CVD_any	1.00213 (1.00104, 1.00323)	1.00105 (0.99998, 1.00211)	1.00178 (1.00068, 1.00288)	1.00596 (1.00493, 1.00698)
RESP_any	1.00957 (1.00846, 1.01068)	1.01127 (1.01019, 1.01235)	1.01024 (1.00913, 1.01135)	1.01480 (1.01377, 1.01583)

Four readings reinforce Section 5.6. First, the fully meteorology-stripped model gives the largest odds ratio for both outcomes (respiratory 1.01480; cardiovascular 1.00596), above even the temperature-only-removed estimate of Table 3.17, so stripping meteorological adjustment raises the dust signal in the same toward-null direction established for temperature alone. Second, the respiratory ladder is fully monotone, every step that relaxes meteorological adjustment lifts the odds ratio. Third, the one non-monotone cell (cardiovascular with relative humidity dropped but temperature retained) is a small, explicable consequence of the moderate negative Tmax-relative-humidity correlation (approximately −0.57), removing relative humidity lets the retained temperature term re-absorb a little more shared meteorological signal, and the fully stripped model restores the expected ordering. Fourth, humidity functional form is nearly a non-event, replacing linear with spline relative humidity leaves odds ratio, sign, and significance intact.

Two construct points keep this in proportion. The humidity overlap is looser than the temperature overlap, because the exposure model used specific humidity whereas the health model adjusts for relative humidity, so the two are not the identical field. And the remedy is the same as for temperature and lies upstream, in a meteorology-decoupled re-derivation of the exposure (Section 8.8), no additional downstream humidity adjustment is indicated, because it would only deepen the toward-null bias. Taken together, Sections 5.6 and 5.7 are two halves of one meteorological-overlap diagnostic, and both halves point the same way, the reported dust associations are conservative, and the fully meteorology-stripped models bound their upper end.

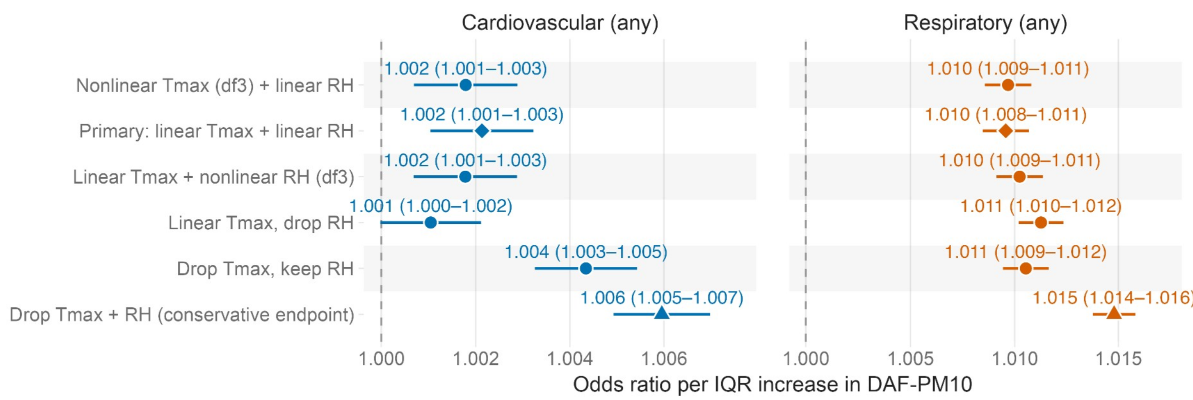

Figure 3.14. DAF-PM_{10} per-IQR odds ratios under progressively relaxed meteorological adjustment, for broad cardiovascular and respiratory outcomes (combined ED+PDD). Relaxing temperature and humidity adjustment raises the dust signal, indicating that the primary, fully adjusted estimate is conservative and that the meteorological double-counting biases toward the null.

6. Discussion

6.1 Principal findings

The same-day primary estimates establish small but precisely estimated short-term effects, a one-IQR (+16.558 μg/m³) increase in dust-attributable PM_{10} raises the odds of a broad respiratory encounter by approximately 1.0% and the odds of a broad cardiovascular encounter by approximately 0.2% on the same day, with the respiratory effect roughly five times the cardiovascular effect per IQR on the log-odds scale. Both estimates survive an exposure-metric sensitivity across three operationalisations, a five-specification time-varying-confounding suite, and a permutation null at a 100,000-stratum subsample.

The cause-specific decomposition pinpoints both signals to biologically plausible subgroups. The respiratory signal is largest for asthma, followed by acute respiratory infection and COPD, with pneumonia null, the cardiovascular signal is essentially the hypertensive-disease subgroup, with ischemic heart disease borderline positive and cerebrovascular and ischemic-stroke estimates null. The pneumonia null is the cleanest mechanistic check available in these data, bacterial pneumonia follows a multi-day clinical course that is poorly matched to a same-day case-crossover, whereas asthma exacerbations and acute respiratory infections evolve over hours and are well matched to it. The distributed-lag analysis sharpens this picture the respiratory same-day estimate undercounts the total respiratory burden by roughly 40%, because the effect peaks one day after exposure, persists through lag 4, and accumulates to a seven-day cumulative odds ratio of 1.0136. The concentration-response analysis then shows that the respiratory exposure-response is strongly nonlinear and U-shaped, which materially changes the attributable-fraction calculation, while the cardiovascular curve is near-monotone.

The heat-interaction analysis progresses through four stages of refinement that together convert a single binary interaction into mechanism-specific findings. A pre-specified extreme-heat interaction shows respiratory amplification and cardiovascular attenuation, a backcast sensitivity confirms that neither depends on the threshold backcast, a temperature-only threshold sensitivity confirms that the respiratory amplification reproduces under purely temperature-defined heat, resolving the circularity concern for the respiratory side, and the heat-metric decomposition attributes the respiratory amplification to peak daytime heat (Tmax), the cardiovascular attenuation to overnight heat (Tmin), and a reversal of both to humid, stagnant low-diurnal-range conditions. A four-group joint exposure model shows that combined high-dust and high-heat exposure is slightly sub-additive on the log-odds scale for respiratory disease, which is the relevant scale for projecting joint burden. Geographic stratification shows respiratory uniformity across populated basins and cardiovascular concentration in coastal metropolitan basins, and the season-stratified analysis independently corroborates the Tmax mechanism by placing the largest respiratory effect in summer.

6.2 Comparison with the existing literature

The respiratory findings sit comfortably within the international dust-health evidence base, which is most consistent and largest for respiratory morbidity. Meta-analysis of short-term Asian-dust exposure reports asthma admissions elevated by roughly 15% and respiratory hospitalisations by roughly 9% on dust days relative to non-dust days (Hashizume et al., 2020), and the broader desert-dust meta-analysis reports respiratory morbidity elevated by roughly 7% while pooled respiratory mortality is null (Tobías et al., 2025). Single-setting studies that use designs close to ours reinforce the pattern: a time-stratified case-crossover of paediatric asthma admissions in Japan found odds ratios of 1.7 to 1.9 on heavy-dust days (Kanatani et al., 2010), a conditional-logistic case-crossover of paediatric asthma emergency visits during Saharan-dust transport to the Caribbean found excess risk concentrated on dust days (Cadelis et al., 2014), a case-crossover of emergency-room visits in Taiwan reported elevated cardiovascular and respiratory visits on Asian-dust days (Liu et al., 2017), and the extreme 2009 Sydney dust storm raised respiratory and asthma emergency visits without a corresponding rise in cardiovascular visits (Merrifield et al., 2013). Our state-wide, ZIP-day-resolved respiratory estimates are directionally concordant with all of these. The magnitudes are not directly comparable, because the prior studies contrast entire dust-storm days against non-dust days, a large categorical exposure contrast, whereas we estimate a continuous slope per interquartile-range increment of a dust-informed total-PM_{10} surface, a per-IQR slope is necessarily smaller than a whole-event contrast. The observation that respiratory effects are larger and more lag-persistent at the morbidity level than at the mortality level is itself consistent with the prior literature, in which dust-triggered exacerbations are common but frequently survivable with treatment (Tobías et al., 2025).

The cardiovascular findings are smaller and, as expected from the literature, more heterogeneous. Meta-analysis attributes a roughly 2% increase in cardiovascular mortality to dust days and a borderline increase in cardiovascular morbidity (Dominguez-Rodriguez et al., 2021; Tobías et al., 2025), and a nationwide Chinese study reports large excess risks for cause-specific endpoints such as ischaemic stroke and COPD on sand-and-dust-storm days (Zhang et al., 2023). Our concentration of the cardiovascular signal in hypertensive disease, with ischemic-stroke and cerebrovascular estimates null, is consistent with the mixed prior cardiovascular pattern and with evidence that dust-cardiovascular effects are modified by host susceptibility, sex, and age (Vodonos et al., 2015) and may be specific to particular myocardial-infarction subtypes (Lee et al., 2020). The contrast with the large positive stroke estimate reported nationwide in China (Zhang et al., 2023) is most plausibly explained by the difference in exposure contrast (a whole dust-storm-day contrast versus our continuous per-IQR slope), by differences in dust composition between source regions, and by the design difference between mortality time-series and morbidity case-crossover. Biological plausibility for the cardiovascular signal is supported by experimental and clinical evidence that desert dust provokes systemic inflammation, coagulation changes, and endothelial perturbation, including airway inflammation in patients with ischemic heart disease (Dominguez-Rodriguez et al., 2020; Fussell & Kelly, 2021).

For the pollution-by-heat interaction, the respiratory amplification we report is consistent with the general expectation that heat and particulate exposure act in the same direction on the respiratory system, and the summer-largest seasonal pattern corroborates it. This direction is consistent with multicity evidence that high temperature modifies the short-term particulate-health association and that the two exposures are mutually confounding modifiers (Stafoggia et al., 2008; Analitis et al., 2014). The more unusual result is the cardiovascular attenuation, which most prior particulate-cardiovascular work has not reported, partly because most prior work uses a single binary heat indicator that averages across distinct thermal regimes. Our heat-metric decomposition is informative here, it shows that the attenuation is specific to high-overnight-temperature (Tmin) days, while humid, stagnant low-diurnal-range days exhibit the more conventional amplification. A decomposition of this kind is rare in the dust-specific literature, and the opposing Tmin-versus-low-DTR directions would be invisible to an analysis that used a single heat threshold. The relevance of this distinction to California is direct, because compound heat, drought, and dust events are projected to become more frequent in the state (Pu et al., 2022). Finally, the four-group joint exposure analysis shows that dust and extreme heat combine sub-additively rather than multiplicatively for respiratory disease, which qualifies the common interpretation of a positive multiplicative interaction term as evidence of synergy and provides the additive-scale check that is most appropriate for projecting joint burden.

6.3 What this work adds

This chapter contributes to the dust-health literature in four distinct ways that prior work has not provided, individually or collectively. First, scale: roughly 23 to 24 million case-crossover strata per broad outcome yield per-IQR confidence intervals approximately 0.2% wide on the odds-ratio scale, narrower than any prior dust-oriented PM_{10} case-crossover, which is what allows detection of a 0.2%-per-IQR cardiovascular signal that would be undetectable in a smaller study. This is not merely a larger study but a structurally better-powered one, because the time-stratified design conditions each encounter on itself, power scales with the number of acute events rather than with the number of monitored locations, so coupling a daily statewide exposure surface to a near-complete administrative census of encounters recovers precision that sparse fixed-site dust networks cannot reach. Second, a dust-informed exposure surface at ZIP-day resolution and statewide scale that has not previously been linked to morbidity in California, although this surface is not an explicit source-apportioned dust mass, it is shaped by dust-relevant covariates and resolves spatial heterogeneity that fixed-monitor designs cannot (Section 2.3). Third, a formal likelihood-ratio test of linear versus spline exposure that demonstrates, at scale, a strongly nonlinear and U-shaped dust-respiratory concentration-response, a property that binary dust-day designs cannot characterise and that prior meta-analysis identified as a structural limitation of the evidence base (Tobías et al., 2025). Fourth, a heat-metric decomposition of the dust-by-heat interaction into Tmax, Tmin, and diurnal-temperature-range effects, which identifies an opposing-direction cardiovascular pattern that conventional binary heat indicators miss. Relative to the strongest prior source-apportioned designs, which separate desert from non-desert PM_{10} in European multi-city data (Stafoggia et al., 2016), this chapter trades explicit chemical apportionment for very high spatial and temporal resolution and an order-of-magnitude larger sample.

6.4 What this work does not establish

We are equally explicit about what this work does not show. It does not establish a causal mechanismm it estimates short-term associations conditional on the matched-stratum design. It does not quantify a chronic-exposure burden, it estimates the effect of day-to-day variation per interquartile-range increment. It does not generalise outside California, because the exposure model is California-trained and dust mineralogy and toxicity differ across source regions

(Goudie, 2014; Karanasiou et al., 2012). It does not isolate a pure mineral-dust mass, because the Aim 2 surface is a dust-informed total-PM_{10} prediction with no explicit source apportionment, the effects reported here are therefore most accurately attributed to a dust-related, dust-dominant particulate exposure rather than to chemically isolated dust, a caveat that applies to most of the dust-health literature, in which only a minority of studies cleanly separate dust-PM from total PM (Karanasiou et al., 2012; Stafoggia et al., 2016; Tobías et al., 2025). It does not explicitly adjust for non-dust PM_{10} or $PM_{2.5}$ as a co-pollutant, which is a high-priority future analysis (Morakinyo et al., 2016). Because temperature and humidity were predictors of the exposure surface as well as confounders in the health model, the dust estimate carries a meteorological double-adjustment that biases it toward the null (Sections 5.6 and 5.7), the estimates we report are therefore conservative, and the fully meteorology-stripped models bound their upper end, but the cleanest remedy lies upstream in a meteorology-decoupled re-derivation of the exposure (Section 8.8). Finally, it does not contain a within-person sensitivity, because the person identifier is not unique.

7. Limitations

7.1 Case-crossover limitation in chronically saturated exposure regions

The case-crossover relies on day-to-day within-stratum exposure variation. In a chronically saturated dust regime, most notably the Salton Sea basin in Imperial County, where the receding playa sustains high exposure on essentially every day (Frie et al., 2019; Parajuli & Zender, 2018), the contrast between case-day and matched referent days is small relative to the absolute exposure level. The Salton Sea null is therefore most simply interpreted as a design limitation rather than as evidence of no effect. The appropriate next step is a cohort or time-series design that compares Salton Sea resident populations against comparable non-saturated populations, or that uses absolute rather than within-stratum-contrasted exposure. This is identified as a top future direction (Section 8.1).

7.2 Exposure misclassification

ZIP-day DAF-PM_{10} exposures carry their own uncertainty. The three operationalisations agree in direction and significance, which rules out gross misallocation between metrics, but agreement across metrics is not a substitute for direct quantification of measurement error within a metric. Non-classical measurement error in a case-crossover generally biases toward the null, and the choice of exposure model in short-term dust-health studies is itself a recognised source of variability (Tobías & Stafoggia, 2020).

7.3 Residual co-pollutant confounding

The primary model adjusts for refined ZIP-day ozone, nitrogen dioxide, and sulphur dioxide and for relative humidity, but not for non-dust PM_{10} or $PM_{2.5}$ as a co-pollutant. As detailed in Section 2.3, the exposure surface predicts total dust-relevant PM_{10}without explicit source apportionment, so DAF-PM_{10} is best read as a dust-related, dust-dominant particulate metric rather than chemically isolated dust. To the extent it carries covariation with non-dust PM_{10} or $PM_{2.5}$, the primary estimate could partially reflect total particulate rather than dust alone (Morakinyo et al., 2016). The wildfire-season exclusion sensitivity shows that removing the smoke-heavy months does not materially change the primary estimate, but this is not a substitute for an explicit multi-PM co-pollutant model, which is identified as a high-priority future direction (Section 8.5).

7.4 Meteorological double-adjustment between exposure generation and the health model

A subtler dependence than ordinary co-pollutant confounding arises because meteorology enters the analytic chain twice, temperature and humidity are among the predictors that generated the exposure surface (Section 2.3), and they re-enter the case-crossover model as confounders. This double-counting is real and cannot be removed post hoc within Aim 3. Temperature is the sharper case, because the identical Tmax field sits on both sides, humidity is looser, because the exposure model used specific humidity while the health model adjusts for relative humidity. Three considerations bound the impact. The collinearity it induces within the health model is mild (all variance-inflation factors at or below 2.83, condition number 3.39; Section 5.6), so the coefficients are jointly identifiable without numerical instability. More importantly, the bias has a known direction: relaxing the meteorological adjustment raises the dust odds ratio (Sections 5.6 and 5.7), so the adjusted estimate we report is conservative and the meteorology-stripped models bound the upper end of the plausible range. Finally, the collinearity diagnostics are a necessary but not sufficient reassurance, because they are computed on the final covariate values and are blind to the training-time dependence; the clean remedy is upstream (Section 8.8), and additional downstream temperature or humidity adjustment is not needed, as it would deepen rather than correct the toward-null bias.

7.5 Circularity for the health-informed heat index

The extreme-heat indicator was developed using the same emergency-department data source used in this aim. The strongest circularity concern arises when such an index is used as a main heat exposure, an interaction (effect modification of a separate exposure) is less directly problematic but not fully independent. The temperature-only sensitivity (Section 4.5) is the appropriate fix and resolves the concern for the respiratory side, the cardiovascular attenuation remains threshold-family dependent and is reported as such.

7.6 Patient-discharge power

The patient-discharge-only analyses have between 0.19 and 0.31 million strata, two orders of magnitude smaller than the emergency-department analyses. Their confidence intervals span the null even for effects that are highly significant in the combined data, so we treat them as an underpowered side-analysis rather than as an independent test of the primary finding.

7.7 ZIP-centroid geographic assignment

Geographic stratification uses population-weighted ZIP centroids joined to CARB air-basin polygons. Misassignment to a neighbouring basin is possible at basin boundaries, but using the population-weighted centroid rather than the polygon centroid should minimise this, and the quality-assurance checks confirm no unclassified rows. We do not attempt a finer spatial resolution, such as census tract, because the upstream exposure surface is ZIP-day-resolved.

7.8 Generalisability and the within-person sensitivity

The exposure model is California-trained, so the estimates generalise most reliably within California and to settings with comparable dust composition, we caution against direct extrapolation to regimes with very different mineralogy, such as trans-Atlantic Saharan transport or East Asian dust (Goudie, 2014). Separately, because the person identifier is not globally unique (Section 5.5), three planned within-person analyses cannot be implemented, this is acknowledged as a project limitation and identified as a future direction.

8. Gaps and Future Directions

8.1 Cohort design for chronically saturated exposure

The single most important future direction is a Salton Sea and Imperial County cohort design that compares the chronically dust-exposed population against a matched, non-saturated comparison population on a between-population rather than a within-day exposure contrast. The case-crossover null in Section 4.6 reflects a design limitation that a cohort or long-term ecological design would address directly.

8.2 Cause-specific stratification by region and heat

The current chapter reports cause-specific results, region, and season only at the broad level. A finer cross of cause-specific outcome by air basin by season by heat regime was deliberately not run, to avoid an unwieldy multiple-testing burden, but a pre-registered subset of such crosses (for example, asthma in the San Joaquin Valley in summer under peak-daytime heat) would be a productive next step.

8.3 Dust composition and speciation

The exposure surface is a dust-informed total-PM_{10} prediction rather than an apportioned dust fraction, and the dust-relevant portion is itself a heterogeneous mixture of silicate, metal-bearing, bioaerosol, and salt-encrusted material whose toxicity varies by component (Fubini & Fenoglio, 2007). Both an explicit dust versus non-dust source-apportionment step and a speciation-resolved exposure variable would sharpen interpretation, the former separating dust from co-emitted non-dust PM, the latter distinguishing, for example, saline Salton Sea dust from Mojave silicate dust (Frie et al., 2019). Chemical speciation is rarely reported in the dust-health literature and remains a leading gap (Karanasiou et al., 2012).

8.4 Full-sample permutation negative control

The reduced-sample permutation null validates the pipeline but does not provide a full-sample falsification. A full-sample permutation would provide that falsification at the cost of several weeks of additional computation. It is identified as a future direction but is not a high-priority stability test, given that the primary respiratory coefficient is 17 standard errors above zero on its own sampling distribution.

8.5 Co-pollutant multi-PM model

An explicit multi-PM model that adds non-dust PM_{10} and $PM_{2.5}$ as co-pollutants would directly test whether the DAF-PM_{10} signal persists after accounting for total particulate, and would move the design closer to the source-apportioned gold standard in the European multi-city literature (Stafoggia et al., 2016). This is a high-priority future direction.

8.6 Climate-projection coupling

The dust-by-heat interaction, and especially the four-group joint exposure model, is directly relevant to climate projection, because the rate at which extreme-heat days and high-dust days co-occur in California is a function of both general circulation and regional dust-source dynamics, and such compound events are projected to increase (Pu et al., 2022). Coupling the dust-by-heat interaction estimates to downscaled climate-projection ensembles would generate California-specific projected burden estimates for the mid-century and end-of-century horizons.

8.7 Outcome-based negative controls

Reprocessing the upstream raw diagnosis fields to reconstruct outcomes such as injury, appendicitis, and fracture would enable outcome-based negative controls. Given the strength of the primary association relative to the permutation null, this is a lower-priority stability check.

8.8 A meteorology-decoupled exposure surface

The single cleanest methodological improvement lies upstream of Aim 3 entirely. Because the Aim 2 ensemble used temperature and humidity among its more influential predictors (Section 2.3), the exposure surface and the health-model meteorological confounders share a training-time dependence that no downstream adjustment can fully resolve and that biases the dust estimate toward the null (Sections 5.6 and 5.7). Re-deriving the surface from a predictor set that excludes the meteorological fields, or that separates out the dust signal with respect to them, and ideally adding an explicit dust versus non-dust apportionment step, would break that dependence at its source. The resulting exposure could then be carried through the identical case-crossover pipeline, and comparing the meteorology-decoupled and meteorology-informed estimates would directly quantify the toward-null bias rather than merely bounding its direction. We regard this as the highest-value exposure-side future direction.

9. Conclusion

This chapter has presented the most comprehensive California-wide, dust-specific, ZIP-day-resolved case-crossover analysis of short-term cardiovascular and respiratory morbidity conducted to date. Across roughly 23 million case-crossover strata per outcome, a one-IQR (+16.558 μg/m³) increase in dust-attributable PM_{10} raised the odds of a broad respiratory encounter by approximately 1% on the same day, peaking one day after exposure and accumulating over the following week to approximately 1.4%, and raised the odds of a broad cardiovascular encounter by approximately 0.2% on the same day, with no further accumulation. The respiratory effect was strongly nonlinear, concentrated in asthma and COPD, amplified specifically by peak daytime heat, and remarkably uniform across populated air basins, with San Diego strongest. The cardiovascular effect was near-monotone, concentrated in hypertensive disease, attenuated specifically by overnight heat under certain temperature-threshold choices, concentrated in coastal metropolitan basins, and absent from the chronically saturated Salton Sea basin, which we interpret as a design limitation rather than a biological null. The heat-metric decomposition refined a single binary heat interaction into mechanism-specific results, and the four-group joint analysis showed slightly sub-additive combined effects on respiratory disease. All primary findings were stable across five time-varying-confounding specifications, three exposure-metric operationalisations, three heat-backcast scenarios, and a permutation null. The chapter contributes to the dust-health literature in four distinct ways, namely scale, a dust-informed high-resolution exposure, a formal nonlinearity test, and a heat-metric decomposition. It identifies nine concrete future directions that follow directly from its limitations. The integrated finding is that dust-attributable PM_{10} is a small but precisely estimated, mechanistically plausible, geographically heterogeneous, heat-modified, and policy-relevant short-term cardiopulmonary morbidity exposure in California, and that the burden it imposes is most accurately captured by a U-shaped respiratory concentration-response that accumulates over a week and is amplified specifically by peak daytime heat.

References

Achakulwisut, P., Shen, L., & Mickley, L. J. (2017). What controls springtime fine dust variability in the western United States? Investigating the 2002–2015 increase in fine dust in the U.S. Southwest. *Journal of Geophysical Research: Atmospheres, 122*(22), 12449–12467. https://doi.org/10.1002/2017JD027208

Adebiyi, A. A., Kibria, M. M., Abatzoglou, J. T., Ginoux, P., & Pandey, P. (2025). Fallowed agricultural lands dominate anthropogenic dust sources in California. *Communications Earth & Environment, 6*, Article 324. https://doi.org/10.1038/s43247-025-02306-0

Analitis, A., Michelozzi, P., D'Ippoliti, D., de'Donato, F., Menne, B., Matthies, F., Atkinson, R. W., Iñiguez, C., Basagaña, X., Schneider, A., Lefranc, A., Paldy, A., Bisanti, L., & Katsouyanni, K. (2014). Effects of heat waves on mortality: Effect modification and confounding by air pollutants. *Epidemiology, 25*(1), 15–22. https://doi.org/10.1097/EDE.0b013e31828ac01b

Armstrong, B. G., Gasparrini, A., & Tobías, A. (2014). Conditional Poisson models: A flexible alternative to conditional logistic case cross-over analysis. *BMC Medical Research Methodology, 14*, Article 122. https://doi.org/10.1186/1471-2288-14-122

Bell, M. L., Zanobetti, A., & Dominici, F. (2013). Evidence on vulnerability and susceptibility to health risks associated with short-term exposure to particulate matter: A systematic review and meta-analysis. *American Journal of Epidemiology, 178*(6), 865–876. https://doi.org/10.1093/aje/kwt090

Bhattachan, A., Okin, G. S., Zhang, J., Vimal, S., & Lettenmaier, D. P. (2019). Characterizing the role of wind and dust in traffic accidents in California. *GeoHealth, 3*(10), 328–336. https://doi.org/10.1029/2019GH000212

Cadelis, G., Tourres, R., & Molinie, J. (2014). Short-term effects of the particulate pollutants contained in Saharan dust on the visits of children to the emergency department due to asthmatic conditions in Guadeloupe (French archipelago of the Caribbean). *PLOS ONE, 9*(3), e91136. https://doi.org/10.1371/journal.pone.0091136

Crooks, J. L., Cascio, W. E., Percy, M. S., Reyes, J., Neas, L., & Hilborn, E. D. (2016). The association between dust storms and daily non-accidental mortality in the United States, 1993–2005. *Environmental Health Perspectives, 124*(11), 1735–1743. https://doi.org/10.1289/EHP216

Dominguez-Rodriguez, A., Rodríguez, S., Baez-Ferrer, N., Abreu-Gonzalez, P., Abreu-Gonzalez, J., Avanzas, P., Carnero, M., Moris, C., López-Darias, J., & Hernández-Vaquero, D. (2020). Impact of Saharan dust exposure on airway inflammation in patients with ischemic heart disease. *Translational Research, 224*, 16–25. https://doi.org/10.1016/j.trsl.2020.05.011

Dominguez-Rodriguez, A., Baez-Ferrer, N., Abreu-Gonzalez, P., Rodríguez, S., Díaz, R., Avanzas, P., & Hernández-Vaquero, D. (2021). Impact of desert dust events on the cardiovascular disease: A systematic review and meta-analysis. *Journal of Clinical Medicine, 10*(4), 727. https://doi.org/10.3390/jcm10040727

Frie, A. L., Garrison, A. C., Schaefer, M. V., Bates, S. M., Botthoff, J., Maltz, M., Ying, S. C., Lyons, T., Allen, M. F., Aronson, E., & Bahreini, R. (2019). Dust sources in the Salton Sea basin: A clear case of an anthropogenically impacted dust budget. *Environmental Science & Technology, 53*(16), 9378–9388. https://doi.org/10.1021/acs.est.9b02137

Fubini, B., & Fenoglio, I. (2007). Toxic potential of mineral dusts. *Elements, 3*(6), 407–414. https://doi.org/10.2113/gselements.3.6.407

Fussell, J. C., & Kelly, F. J. (2021). Mechanisms underlying the health effects of desert sand dust. *Environment International, 157*, 106790. https://doi.org/10.1016/j.envint.2021.106790

Gasparrini, A. (2014). Modeling exposure-lag-response associations with distributed lag non-linear models. *Statistics in Medicine, 33*(5), 881–899. https://doi.org/10.1002/sim.5963

Gasparrini, A., Armstrong, B., & Kenward, M. G. (2010). Distributed lag non-linear models. *Statistics in Medicine, 29*(21), 2224–2234. https://doi.org/10.1002/sim.3940

Gasparrini, A., & Leone, M. (2014). Attributable risk from distributed lag models. *BMC Medical Research Methodology, 14*, Article 55. https://doi.org/10.1186/1471-2288-14-55

Goudie, A. S. (2014). Desert dust and human health disorders. *Environment International, 63*, 101–113. https://doi.org/10.1016/j.envint.2013.10.011

Hashizume, M., Kim, Y., Ng, C. F. S., Chung, Y., Madaniyazi, L., Bell, M. L., Guo, Y. L., Kan, H., Honda, Y., Yi, S.-M., Kim, H., & Nishiwaki, Y. (2020). Health effects of Asian dust: A systematic review and meta-analysis. *Environmental Health Perspectives, 128*(6), 066001. https://doi.org/10.1289/EHP5312

Janes, H., Sheppard, L., & Lumley, T. (2005). Case-crossover analyses of air pollution exposure data: Referent selection strategies and their implications for bias. *Epidemiology, 16*(6), 717–726. https://doi.org/10.1097/01.ede.0000181315.18836.9d

Kanatani, K. T., Ito, I., Al-Delaimy, W. K., Adachi, Y., Mathews, W. C., & Ramsdell, J. W. (2010). Desert dust exposure is associated with increased risk of asthma hospitalization in children. *American Journal of Respiratory and Critical Care Medicine, 182*(12), 1475–1481. https://doi.org/10.1164/rccm.201002-0296OC

Karanasiou, A., Moreno, N., Moreno, T., Viana, M., de Leeuw, F., & Querol, X. (2012). Health effects from Sahara dust episodes in Europe: Literature review and research gaps. *Environment International, 47*, 107–114. https://doi.org/10.1016/j.envint.2012.06.012

Lee, S., Lee, W., Lee, E., Jeong, M. H., Rha, S.-W., Kim, C.-J., Chae, S. C., Kim, H.-S., Gwon, H.-C., & Kim, H. (2020). Effects of Asian dust-derived particulate matter on ST-elevation myocardial infarction: Retrospective, time series study. *BMC Public Health, 20*, Article 1820. https://doi.org/10.1186/s12889-020-10067-y

Lei, H., & Wang, J. X. L. (2014). Observed characteristics of dust storm events over the western United States using meteorological, satellite, and air quality measurements. *Atmospheric Chemistry and Physics, 14*(15), 7847–7857. https://doi.org/10.5194/acp-14-7847-2014

Liu, S.-T., Liao, C.-Y., Kuo, C.-Y., & Kuo, H.-W. (2017). The effects of $PM_{2.5}$ from Asian dust storms on emergency room visits for cardiovascular and respiratory diseases. *International Journal of Environmental Research and Public Health, 14*(4), 428. https://doi.org/10.3390/ijerph14040428

Lu, Y., & Zeger, S. L. (2007). On the equivalence of case-crossover and time series methods in environmental epidemiology. *Biostatistics, 8*(2), 337–344. https://doi.org/10.1093/biostatistics/kxl013

Maclure, M. (1991). The case-crossover design: A method for studying transient effects on the risk of acute events. *American Journal of Epidemiology, 133*(2), 144–153. https://doi.org/10.1093/oxfordjournals.aje.a115853

Merrifield, A., Schindeler, S., Jalaludin, B., & Smith, W. (2013). Health effects of the September 2009 dust storm in Sydney, Australia: Did emergency department visits and hospital admissions increase? *Environmental Health, 12*, Article 32. https://doi.org/10.1186/1476-069X-12-32

Morakinyo, O. M., Mokgobu, M. I., Mukhola, M. S., & Hunter, R. (2016). Health outcomes of exposure to biological and chemical components of inhalable and respirable particulate matter. *International Journal of Environmental Research and Public Health, 13*(6), 592. https://doi.org/10.3390/ijerph13060592

Parajuli, S. P., & Zender, C. S. (2018). Projected changes in dust emissions and regional air quality due to the shrinking Salton Sea. *Aeolian Research, 33*, 82–92. https://doi.org/10.1016/j.aeolia.2018.05.004

Pu, B., Jin, Q., Ginoux, P., & Yu, Y. (2022). Compound heat wave, drought, and dust events in California. *Journal of Climate, 35*(24), 8133–8152. https://doi.org/10.1175/JCLI-D-21-0889.1

Rowan, C., D'Souza, R. R., Zheng, X., Crooks, J., Hohsfield, K., Tong, D., Chang, H. H., & Ebelt, S. (2024). Dust storms and cardiorespiratory emergency department visits in three southwestern United States: Application of a monitoring-based exposure metric. *Environmental Research: Health, 2*(3), 031003. https://doi.org/10.1088/2752-5309/ad5751

Rublee, C. S., Sorensen, C. J., Lemery, J., Wade, T. J., Sams, E. A., Hilborn, E. D., & Crooks, J. L. (2020). Associations between dust storms and intensive care unit admissions in the

United States, 2000–2015. *GeoHealth, 4*(8), e2020GH000260. https://doi.org/10.1029/2020GH000260

Schwartz, J. (2000). Harvesting and long term exposure effects in the relation between air pollution and mortality. *American Journal of Epidemiology, 151*(5), 440–448. https://doi.org/10.1093/oxfordjournals.aje.a010228

Stafoggia, M., Schwartz, J., Forastiere, F., Perucci, C. A., & the SISTI Group. (2008). Does temperature modify the association between air pollution and mortality? A multicity case-crossover analysis in Italy. *American Journal of Epidemiology, 167*(12), 1476–1485. https://doi.org/10.1093/aje/kwn074

Stafoggia, M., Zauli-Sajani, S., Pey, J., Samoli, E., Alessandrini, E., Basagaña, X., Cernigliaro, A., Chiusolo, M., Demaria, M., Díaz, J., Faustini, A., Katsouyanni, K., Kelessis, A. G., Linares, C., Marchesi, S., Medina, S., Pandolfi, P., Pérez, N., Querol, X., … Forastiere, F. (2016). Desert dust outbreaks in Southern Europe: Contribution to daily PM_{10} concentrations and short-term associations with mortality and hospital admissions. *Environmental Health Perspectives, 124*(4), 413–419. https://doi.org/10.1289/ehp.1409164

Tobías, A., & Stafoggia, M. (2020). Modeling desert dust exposures in epidemiologic short-term health effects studies. *Epidemiology, 31*(6), 788–795. https://doi.org/10.1097/EDE.0000000000001255

Tobías, A., Karanasiou, A., Amato, F., Roqué, M., Suu Lwin, K., Yuan, L., Ith, S., Zin Wai, H., Lester Chua, P., Solá, I., Renzi, M., Stafoggia, M., Querol, X., & Hashizume, M. (2025). Short-term exposure to desert dust and sandstorms and all-cause and cause-specific mortality and morbidity: A systematic review and meta-analysis. *Environment International, 196*, 109277. https://doi.org/10.1016/j.envint.2025.109277

Tong, D. Q., Dan, M., Wang, T., & Lee, P. (2012). Long-term dust climatology in the western United States reconstructed from routine aerosol ground monitoring. *Atmospheric Chemistry and Physics, 12*(11), 5189–5205. https://doi.org/10.5194/acp-12-5189-2012

Vodonos, A., Friger, M., Katra, I., Krasnov, H., Zahger, D., Schwartz, J., & Novack, V. (2015). Individual effect modifiers of dust exposure effect on cardiovascular morbidity. *PLOS ONE, 10*(9), e0137714. https://doi.org/10.1371/journal.pone.0137714

Zhang, X., Zhao, L., Tong, D. Q., Wu, G., Dan, M., & Teng, B. (2016). A systematic review of global desert dust and associated human health effects. *Atmosphere, 7*(12), 158. https://doi.org/10.3390/atmos7120158

Zhang, Y., Ma, R., Ban, J., Lu, F., Guo, M., Zhang, Y., Liu, J., Wang, R., Wang, K., Lin, S., Li, X., Niu, Y., & Li, T. (2023). Risk of cause-specific death associated with sand and dust storms in China: A nationwide multicity time-series study. *Nature Communications, 14*, Article 6800. https://doi.org/10.1038/s41467-023-42530-w

Chapter 4. Translating Dust Storm Evidence into Action: A Risk Communication and Community Engagement Strategy

Chapter Overview

The preceding chapters built a chain of evidence about dust storms, from how dust-attributable particulate matter can be identified and quantified, to whether that exposure is associated with acute illness, to what the global record of early warning systems reveals about effective design. This chapter completes the chain by addressing the step that separates scientific knowledge from public health benefit, the translation of evidence into resources that agencies and communities can use. It presents three products developed for this purpose. The first is a dust storm risk communication and community engagement plan, applied to Riyadh, Saudi Arabia as a high-burden setting. The second is a public-facing infographic. The third is a short educational video. Each was built on the established peer-reviewed evidence base and on best practices in risk and science communication, and each is reinforced by the systematic review of dust storm early warning systems conducted in Aim 1. The chapter sets out the conceptual and empirical foundations of the products, describes and presents each one, and explains how the broader dissertation positions this work to inform future operational warning systems.

Introduction

A dust storm forecast has no public health value on its own. Value is created only when a warning reaches the people at risk, is understood, is believed, and is converted into protective action in the short interval before exposure occurs. This final step, often called the last mile of an early warning system, is where many otherwise sound systems fail (Basher, 2006; Sättele et al., 2016). The first three aims of this dissertation strengthened the scientific basis for warning about dust. The systematic review of global dust storm early warning systems (Aim 1) characterized the data sources, designs, and gaps in current practice. The exposure work (Aim 2) produced a health-relevant metric for the dust-attributable fraction of PM_{10} (DAF-PM_{10}). The epidemiologic analysis (Aim 3) examined whether that exposure is associated with acute cardiovascular and respiratory morbidity and whether extreme heat modifies the association. What remains is to carry evidence of this kind across the last mile.

That task is not a matter of writing better press releases. Dust storms are a recurring, climate-sensitive hazard whose burden falls unevenly across a population, and the people most exposed are frequently the least reached by conventional messaging (Bell et al., 2013; D'Antoni et al., 2017). Effective translation requires a deliberate strategy that is grounded in communication science, designed around the needs of specific audiences, and embedded in a coordinated operational structure rather than improvised during each event. The discipline that addresses this problem is risk communication, and its current global standard treats communication as a two-way relationship with affected communities rather than a one-directional broadcast (Reynolds & Seeger, 2005; World Health Organization, 2017).

This chapter pursues the objective of Specific Aim 4, to synthesize the dissertation's contribution and develop examples of evidence-based resources that enhance dust storm preparedness and response. The framing of these resources matters and should be stated plainly. The three products are public health communication tools built from the general scientific and communication evidence base. The Aim 1 systematic review and its findings on how dust early warning systems are designed, and on where they fall short, reinforce the content and the priorities of the products, particularly the emphasis on surveillance, warning, and communication. The exposure metric and the epidemiologic results contribute in a different way, by strengthening the scientific case for investing in dust warning and by contributing toward the establishment of criteria, thresholds, and timing that future operational systems will need. The chapter develops this argument in four moves, it sets out the conceptual and empirical foundations, describes how each product was developed, presents the products themselves, and discusses their translational implications and limitations.

Background and Conceptual Foundations

Why dust storms demand a communication strategy

The health rationale for warning the public about dust storms is well established. Reviews of the global evidence link desert dust exposure to increased cardiovascular and respiratory morbidity and mortality, with effects that are biologically plausible given the size, mineral composition, and biological load of dust particles (Goudie, 2013; Lwin et al., 2023; Zhang et al., 2016). Time-series and case-crossover studies across several regions reinforce this picture. Dust events have been associated with increases in non-accidental mortality in the United States (Crooks et al., 2016) and with cause-specific mortality in Asia (Kashima et al., 2012), with emergency department visits for cardiovascular and respiratory disease (Liu et al., 2017), and with intensive care unit admissions (Rublee et al., 2020). A systematic review of cardiovascular and respiratory outcomes reaches consistent conclusions (Farrokhi et al., 2021). Specific outcomes documented in the literature include asthma exacerbation and hospitalization, particularly in children and in patients with existing asthma (Kanatani et al., 2010; Watanabe et al., 2015), and dust-associated infectious disease (Vergadi et al., 2022). Dust also harms health indirectly by reducing visibility and raising the risk of road traffic collisions (Bhattachan et al., 2019; Islam et al., 2019). These risks are not evenly distributed. People with chronic cardiorespiratory disease, older adults, and young children are more susceptible to particulate exposure, so a strategy that reaches only the general population will miss those who need it most (Bell et al., 2013; Vodonos et al., 2015).

The hazard is also intensifying. Climate change is expected to increase the frequency and intensity of dust events in arid and semi-arid regions (Intergovernmental Panel on Climate Change, 2023; Middleton, 2019), and dust increasingly co-occurs with other climate-sensitive hazards. Analysis of compound heat wave, drought, and dust events in California shows that these stressors cluster in time (Pu et al., 2022), and the epidemiologic chapter of this dissertation examined whether extreme heat modifies the health effect of dust exposure and found that extreme heat amplified the respiratory association with dust-attributable PM_{10} and attenuated the cardiovascular association, with the respiratory amplification driven specifically by peak daytime temperature (Chapter 3). The burden is acute in the Arabian Peninsula and the broader Middle East, where dust is a dominant source of particulate air pollution and where source apportionment and exposure studies document high PM_{10} concentrations and a large mineral dust contribution (Al-Hemoud et al., 2018; Khodeir et al., 2012). This regional context is the reason the communication plan in this chapter is grounded in Riyadh.

Risk communication as the bridge from evidence to action

Risk communication provides the theoretical structure for the last mile. The Crisis and Emergency Risk Communication (CERC) model organizes communication around the lifecycle of an event, from a pre-crisis stage of preparation and public education, through the initial event, to resolution and evaluation, and it treats the management of rumor and the building of credibility as core functions (Reynolds & Seeger, 2005). The World Health Organization's guideline on emergency risk communication, derived from a series of systematic reviews, reaches a complementary conclusion: trust is the central currency of emergency communication, built through transparency, community engagement, and coordination across the agencies that speak to the public (World Health Organization, 2017). Long-standing best-practice syntheses converge on the same points, emphasizing that communicators should accept the public as a legitimate partner, listen before broadcasting, and communicate with clarity and compassion (Covello, 2003).

Two further principles shape the design of the products. The first concerns the structure of a persuasive warning. The extended parallel process model predicts that messages which raise concern about a threat will motivate protective action only when they are paired with clear, achievable guidance on how to reduce the risk, threat without efficacy tends to produce avoidance rather than action (Witte, 1992). A dust message that describes the hazard but does not tell people precisely what to do is therefore likely to underperform. The second principle concerns tailoring. A meta-analysis of tailored health interventions found that messages adapted to the characteristics and needs of specific audiences produce larger behavioral effects than generic messages (Noar et al., 2007). The empirical record on environmental warnings shows why these principles matter. A systematic review of adherence to air quality warning systems found that protective behavior is often suboptimal and is shaped by whether people believe the pollution is harmful, whether they believe protective action is worthwhile, and whether the advice comes from a trusted source (D'Antoni et al., 2017). Research on the path from warning messages to preparedness behavior similarly emphasizes the mediating role of risk perception and of the way information is exchanged (Guo et al., 2022). Together, these findings argue for warnings that are action-first, audience-aware, and delivered through trusted channels.

Visual and audiovisual science communication

Reaching a diverse public requires formats that do not depend on dense text, which is the rationale for an infographic and a video. A review of the role of pictures in health communication found that adding pictures to written or spoken information improves attention, comprehension, recall, and adherence, and that the benefit is largest for audiences with low literacy (Houts et al., 2006). This supports icon-based, plain-language materials in a city with a linguistically diverse population. For audiovisual media, a systematic review of video-based education found that video can change health behavior by combining knowledge transfer with visual modeling of the desired action, and that materials can be revisited for reinforcement (Tuong et al., 2014). These findings support a static infographic and a short video as complementary formats, each playing to its strengths.

Misinformation and the infodemic

A modern communication strategy must plan for the information environment it enters. During emergencies, inaccurate information can spread quickly and crowd out official guidance. The cognitive science of misinformation shows that false beliefs persist even after correction, a continued-influence effect that makes simple retraction unreliable, and it identifies approaches that pre-empt and repeatedly counter false claims (Lewandowsky et al., 2012). One such approach is inoculation, or prebunking, in which audiences are exposed in advance to the techniques used to mislead them, controlled evaluation has shown that this can build durable resistance to misinformation across diverse populations (Roozenbeek & van der Linden, 2019). At the system level, the World Health Organization and partners have formalized infodemic management as a defined public health function with planning, monitoring, and response components (Tangcharoensathien et al., 2020). The communication plan adopts this stance, treating misinformation as a managed risk.

People-centered early warning

The products are framed within the international consensus on what makes an early warning system effective. The defining feature of a successful system is that it is people-centered, it connects hazard detection to the knowledge, communication channels, and response capacity of the communities at risk (Basher, 2006). The Sendai Framework for Disaster Risk Reduction sets the reduction of exposure and the expansion of multi-hazard early warning as global targets (United Nations Office for Disaster Risk Reduction, 2015), and the United Nations Early Warnings for All initiative has made universal coverage an explicit goal, with communication and response recognized as the components most in need of strengthening (World Meteorological Organization, 2022). For dust specifically, the World Meteorological Organization's Sand and Dust Storm Warning Advisory and Assessment System and the global assessment of sand and dust storms provide the forecasting backbone on which national warning and communication functions can be built (United Nations Environment Programme et al., 2022). Forecasting capability is advancing in the region of interest, machine learning models now predict dust storm frequency for Saudi cities, including Riyadh, using meteorological and environmental features (Alshammari et al., 2024). The products in this chapter are designed to occupy the communication and response layer that these initiatives identify as decisive, and the Aim 1 systematic review is what links that international evidence to the specific content of the products.

Approach: Developing the Products

The three products were developed as a coordinated set rather than as independent items, so that a person encountering the video, the infographic, and an official alert would meet one consistent message in three complementary forms. They share a common evidence foundation and a common purpose, which is to make the science of dust storms usable by the public and by decision-makers.

Evidence foundation

The content of all three products rests on the established peer-reviewed evidence base and on best practices in risk and science communication, and it is reinforced by the dissertation's own systematic review of dust storm early warning systems (Aim 1). The infographic and video predate and are independent of the exposure and epidemiologic results in Aims 2 and 3, they were assembled from the general literature on what dust storms are, how they form, how they affect health, and how their impacts can be reduced. The Aim 1 review strengthens this foundation in two ways. It confirms that detection and forecasting are comparatively well developed while warning dissemination, community communication, and response remain the weaker components of dust early warning, which is the gap the products are meant to help fill in the early warning literature, where nineteen of the twenty-two studies reviewed in Chapter 1 addressed hazard monitoring and forecasting while only three addressed risk knowledge, warning communication, or response capacity. It also identifies the surveillance and warning recommendations that appear in the infographic and the broader plan as evidence-supported priorities. The exposure metric from Aim 2 and the health findings from Aim 3 are not inputs to the educational products, but they add scientific weight to the case the products make and inform the operational warning logic discussed later in the chapter.

The risk communication and community engagement plan

The plan was built on the CERC event lifecycle (Reynolds & Seeger, 2005) and updated to current World Health Organization practice on risk communication and community engagement, which foregrounds two-way listening, equity, trusted messengers, and coordination (World Health Organization, 2017). Consistent with the design of Aim 4, the plan is informed by the Aim 1 systematic review and by established risk communication strategies. Six principles govern it, two-way engagement, equity and reach, accessibility, trust and transparency, action over information, and coordination. The plan specifies a single coordinating body, a defined crisis communication team, an audience and partner map, a tiered alert and messaging structure, an infodemic management function, an operational sequence across the pre-event, event, and post-event phases, a preparedness timeline keyed to the spring dust season, and a monitoring and evaluation framework. The full operational plan is provided in Appendix A; this chapter presents its structure and rationale.

The infographic

The infographic was designed to communicate the essential science of dust storms to a general audience in a single, scannable visual. Its design follows the evidence that icon-based, plain-language visuals improve comprehension, recall, and adherence, especially among readers with low literacy (Houts et al., 2006), and the principle that protective guidance must be specific and achievable to motivate action (Witte, 1992). The artifact takes the form of two versions that share an identical educational core and differ only in their closing section, which lets the same visual language serve two communicative purposes. Its content and structure are presented below and reproduced in Appendix B.

The video

The video was designed as a short explainer video for a broad audience, consistent with the goal of communicating complex science in an accessible form. Its structure follows evidence that video can change behaviour by pairing clear explanation with visual demonstration (Tuong et al., 2014); here, it shows how stabilising the land surface and planting vegetation can lower dust-storm occurrence. The narrative was assembled from the general evidence on dust storm formation, impacts, and mitigation, and it was directed, produced, and edited by the author. Its content is presented below and reproduced as a storyboard in Appendix C. Formal audience evaluation of both the infographic and the video is identified as future work.

The Products

The dust storm risk communication and community engagement plan

The plan converts the principles above into an operational structure. Its purpose is to deliver timely, trusted, and actionable guidance to the population of Riyadh across the full cycle of an event, and to strengthen local capacity to do so over successive dust seasons. It separates the audiences it protects from the partners that act during an event, and it gives priority to people with chronic cardiovascular or respiratory conditions, older adults, families with young children, people with disabilities, outdoor workers, and the city's large non-Arabic-speaking resident population. Reach and credibility are treated as distinct problems, alongside official channels, the plan works through trusted messengers such as clinicians and pharmacists, schools, mosque networks, and employers (Covello, 2003; World Health Organization, 2017).

Coordination is vested in a single designated body that connects these partners and houses a crisis communication team with defined roles, including a risk communication lead, a content and message coordinator, a media coordinator, a community engagement and outreach coordinator who also gathers community feedback, a partner and stakeholder coordinator, and a social listening and infodemic management coordinator. Public messaging is organized around a tiered alert structure that ties the level of warning to forecast confidence and expected severity, that leads with the action before the explanation (Witte, 1992), and that reserves an enhanced level for compound high-heat, high-dust days. The plan embeds infodemic management as a standing function (Lewandowsky et al., 2012; Tangcharoensathien et al., 2020) and follows the CERC lifecycle through readiness, response, and recovery phases, supported by a preparedness timeline and a monitoring and evaluation framework that tracks the time from forecast to warning, the uptake of protective behavior, the event's measured health and safety impact, and community feedback. Table 4.1 maps the principal elements of the plan to the risk communication principles and the evidence that supports them.

The dust storm infographic

The infographic, titled simply "Dust Storms" and designed by the author, presents the science of dust storms in four illustrated sections built around a consistent set of icons and a single color palette. It exists in two versions that are identical except for their final section, an arrangement that allows one educational core to be pointed at two different communicative purposes. Both versions are reproduced in Appendix B (Figure B1, the Solutions version, and Figure B2, the How To Be Safe version).

The shared core covers three topics. The **health impacts** section states that exposure to wind-blown dust can cause premature death from cardiovascular and respiratory disease and can worsen health through asthma exacerbations, allergic reactions, and skin and eye irritation, and it notes the association with serious infections such as valley fever and meningococcal meningitis. These statements are well supported by the peer-reviewed literature on dust and cardiorespiratory outcomes (Crooks et al., 2016; Farrokhi et al., 2021; Goudie, 2013; Liu et al., 2017; Rublee et al., 2020) and on dust-associated infection (Vergadi et al., 2022). The **causes** section gives the World Meteorological Organization definition of a dust storm as surface winds raising large quantities of dust and reducing visibility to less than 1,000 meters, and it distinguishes natural drivers, such as wind speed and direction and the global abundance of sand and desert, from anthropogenic drivers such as land use and desertification. The attribution to both natural and human land-surface sources is consistent with the dust-source literature (Ginoux et al., 2012;

Lambert et al., 2020). The **frequency** section explains that dust storms are most common in arid and semi-arid areas and concentrate in the northern-hemisphere dust belt, accompanied by a global dust map, this distribution is well documented (Goudie & Middleton, 2001; Wang, 2015).

The two versions diverge in their fourth section, and the difference is the design's main idea. One version closes with a **Solutions** section that addresses mitigation and policy, decreasing anthropogenic emissions from land use, increasing reforestation, preventing further desertification, and implementing surveillance and warning systems to reduce exposure. The mitigation content reflects evidence that land-surface disturbance increases dust emission, so that reducing it does the reverse (Ginoux et al., 2012; Lal, 2012; Lambert et al., 2020), and the call for surveillance and warning systems is precisely the priority reinforced by the Aim 1 review (Basher, 2006; Sättele et al., 2016). The other version closes with a **How To Be Safe** section that addresses individual protective action, staying indoors during an event, covering the nose and mouth with a mask and protecting the eyes when outdoors, keeping an inhaler on hand for people with asthma, and not driving during a storm but pulling over and calling emergency services. This protective guidance is action-first in the sense the extended parallel process model recommends (Witte, 1992) and aligns with the protective behaviors examined in the air quality adherence literature (D'Antoni et al., 2017).

The two versions therefore instantiate two communicative intentions from one evidence-based template. The Solutions version is oriented toward awareness and mitigation and is suited to policymakers, planners, and the general public outside of an active event. The How To Be Safe version is oriented toward protective action and is suited to the general public and at-risk groups when an event threatens or occurs. This realized pairing is the concrete demonstration of audience-and-intention tailoring, and the same template could be extended to additional audiences, such as outdoor workers or schools, or additional intentions, such as preparedness checklists, without redesigning the visual system. Table 4.2 maps the infographic's content to its supporting evidence, and Table 4.3 sets out the two versions against their intention, audience, and timing.

The educational video

The educational video is a short film of about three minutes and fifty seconds that explains dust storms through original desert cinematography overlaid with concise on-screen text. Its narrative arc moves from definition to formation to mitigation. It opens by establishing that dust storms occur in arid and semi-arid environments and carry weather, health, and economic implications. It then explains formation, with the effect of wind, loose dust can be lifted and transported to different locations, often originating from loose sand deserts, and the erodibility of that surface material is potentiated by hot and dry weather, which makes sand looser and easier for wind to lift. This account is consistent with the physical literature on dust emission, in which low surface moisture and wind shear govern how readily particles are mobilized (Csavina et al., 2014; Neuman & Sanderson, 2008). The film then turns to mitigation, arguing that limiting desertification and increasing vegetation through reforestation reduce dust storm frequency, a mechanism supported by evidence that land-surface disturbance increases dust emission (Ginoux et al., 2012; Lambert et al., 2020).

The video closes with a national example. It reports that Saudi Arabia expanded vegetation cover after 2017 through a number of national projects, with reported increases in vegetation cover and an associated reduction in dust storm frequency. These specific figures derive from national environmental-initiative reporting rather than from a peer-reviewed study and are not yet documented in the peer-reviewed record; they should be attributed to a verifiable source, or the percentages removed, before the video is presented as a scholarly artifact. The video is original work, filmed entirely in Saudi Arabia, with music credited on screen. A storyboard of its principal scenes and on-screen text is provided in Appendix C (Figure C1). As an awareness tool it complements the infographic and the plan, and its emphasis on surveillance and mitigation is again the point at which the Aim 1 review reinforces its content.

Discussion

Principal contribution

This chapter converts the dissertation's contribution into a coordinated set of communication resources and, in doing so, addresses the step at which dust early warning most often fails. The contribution is not a new measurement or estimate but a defensible pathway from evidence to action, a risk communication and community engagement plan built on the CERC lifecycle and current World Health Organization practice, an evidence-based infographic that exists in a mitigation-oriented and a protection-oriented version, and a short educational video on dust storm causes, impacts, and mitigation. All three are anchored in the peer-reviewed literature and in established communication science, and all three are reinforced by the Aim 1 systematic review, which identifies communication and warning as the weak components of existing dust early warning systems.

Implications for future operational warning systems

The dissertation's exposure and health results inform the design of the warning systems these products are meant to serve, even though they are not inputs to the products themselves. Identifying which dust events and which exposure levels are most predictive of adverse health outcomes informs the criteria for issuing public alerts, the thresholds at which each alert level is triggered, and the timing of protective recommendations. The DAF-PM_{10} metric offers a health-relevant target variable for such thresholds, and the analysis of effect modification by heat, together with evidence that heat, drought, and dust co-occur (Pu et al., 2022), supports a compound-hazard tier that combines dust and meteorological criteria. Operational forecasting to drive such a system is already developing in the region, including machine learning prediction of dust storm frequency for Riyadh and other Saudi cities (Alshammari et al., 2024). In this sense the products specify what an operational system must be able to say, to whom, and when, and the exposure and epidemiologic chapters help specify the thresholds and timing that the detection logic must deliver. This positions the dissertation to contribute to the communication and response layer that international initiatives identify as the weakest link in dust early warning (United Nations Environment Programme et al., 2022; World Meteorological Organization, 2022).

Integration across the dissertation

The chapter closes the arc that the dissertation set out to trace. Aim 1 established what effective dust warning systems look like and where they fall short, and it is the dissertation result that most directly reinforces the products developed here. Aim 2 produced a health-relevant, dust-informed particulate exposure metric that discriminates high-dust days against an independent composition-based reference (AUC 0.84), without performing explicit source apportionment. Aim 3 tested whether that exposure is associated with acute illness and how heat modifies the association. Aim 4 carries the broader evidence base to the point of use and shows how the dissertation's own findings strengthen the case for, and the future design of, operational dust warning. Figure 4.1 renders this integration as a single conceptual framework, linking the four aims to the Aim 4 products and to the operational warning layer they are designed to support.

Strengths

The approach has three strengths. It is evidence-anchored, in that every design choice traces to a peer-reviewed source or an established best practice. It is equity-oriented, in that the plan prioritizes the audiences who are both most exposed and least reached and builds accessibility into every product (Bell et al., 2013; Houts et al., 2006). And it is coherent, in that the plan, infographic, and video share one message and a common visual logic, which reduces the risk of the contradictory guidance that erodes trust during real events.

Limitations

Several limitations bound the claims of this chapter. The products have been designed and grounded in evidence but have not yet been evaluated with their intended audiences, so their effect on knowledge, trust, and protective behavior is expected rather than demonstrated, the air quality literature is a caution here, since adherence to health advice is frequently suboptimal even when warnings are issued (D'Antoni et al., 2017). The epidemiologic evidence that strengthens the case for these products was generated in California, whereas the communication plan is applied to Riyadh; although the health relationships and the communication principles are expected to transfer, the specific thresholds, audiences, and channels require local validation. The dissertation does not implement a real-time operational warning system, and the products assume an institutional home and a level of coordination that would have to be established through policy.

Future directions

These limitations define a clear program of future work. The most immediate need is evaluation, pretesting the infographic and video with members of the target audiences and assessing comprehension, trust, and intended action, followed by measurement of real protective behavior and of health and safety outcomes after events using the indicators specified in the plan. A second priority is local adaptation and partnership, including confirmation of audiences, languages, thresholds, and channels with the responsible agencies, the sourcing of national statistics used in the video, and the negotiation of the standing agreements the plan requires. A third priority is integration with operational forecasting, connecting the DAF-PM_{10} metric and the compound-hazard logic to the dust prediction systems coordinated through the World Meteorological Organization and to the regional machine learning forecasts now emerging (Alshammari et al., 2024). Pursued together, these steps would move the resources from a designed strategy toward an evaluated, locally owned, and operational warning capability.

Conclusion

Scientific advances in dust exposure assessment and epidemiology protect health only when they change what people do before a storm arrives. This chapter built the bridge from evidence to that behavior. It presented a risk communication and community engagement plan, an evidence-based infographic in two purpose-specific versions, and an educational video, each grounded in the peer-reviewed literature and in the established science of risk and health communication, and each reinforced by the dissertation's systematic review of dust storm early warning systems. The strategy is offered as a practical, defensible, and adaptable contribution, a way to ensure that the dissertation's science does not end at publication but continues into the streets, homes, and workplaces where dust storms do their harm.

Tables

Table 4.1. Elements of the dust storm risk communication plan mapped to risk communication principles and supporting evidence.

Plan element	Guiding principle	Supporting evidence
Pre-event, event, and post-event operational sequence	Manage the full event lifecycle	Reynolds & Seeger (2005); Centers for Disease Control and Prevention (2018)
Single coordinating body and defined team roles	Coordination and consistent messaging	Covello (2003); World Health Organization (2017)
Audience and partner map with	Equity and reach	Bell et al. (2013); D'Antoni et al.

Plan element	Guiding principle	Supporting evidence
priority groups		(2017)
Trusted messengers alongside official channels	Trust and credibility	Covello (2003); World Health Organization (2017)
Tiered alerts with action-first messaging	Action over information	Guo et al. (2022); Witte (1992)
Compound high-heat, high-dust enhanced tier	Match warning to combined risk	Pu et al. (2022); Aim 3 analysis
Two-way listening and community feedback	Two-way engagement	Guo et al. (2022); World Health Organization (2017)
Infodemic management function	Manage misinformation	Lewandowsky et al. (2012); Tangcharoensathien et al. (2020)
Accessibility in language and format	Reach diverse audiences	Houts et al. (2006)
Monitoring and evaluation framework	Learn and improve	Basher (2006); Sättele et al. (2016)

Table 4.2. Content of the dust storm infographic mapped to supporting peer-reviewed evidence.

Infographic section	What it states	Supporting evidence
Health impacts	Premature cardiovascular and respiratory death; asthma, allergy, skin and eye irritation; valley fever and meningococcal meningitis	Crooks et al. (2016); Farrokhi et al. (2021); Goudie (2013); Liu et al. (2017); Rublee et al. (2020); Vergadi et al. (2022)
Causes (definition)	WMO definition; visibility reduced below 1,000 m; wind transports loose particles	Goudie & Middleton (2001); Middleton (2019)
Causes (drivers)	Natural drivers (wind, abundant sand) and anthropogenic drivers (land use, desertification)	Ginoux et al. (2012); Lambert et al. (2020)
Frequency	Concentrated in arid and semi-arid areas and the northern-hemisphere dust belt	Goudie & Middleton (2001); Wang (2015)
Solutions (version 1)	Reduce land-use emissions, reforest, prevent desertification, implement surveillance and warning systems	Basher (2006); Lal (2012); Lambert et al. (2020); Sättele et al. (2016)
How to be safe (version 2)	Stay indoors, use mask and eye protection, keep an inhaler, do not drive, call emergency services	D'Antoni et al. (2017); Witte (1992)

Table 4.3. The two infographic versions mapped to communicative intention, audience, and timing.

Version	Communicative intention	Primary audience	Timing
Solutions	Awareness and mitigation	Policymakers, planners, general public	Calm season and ongoing education
How to be safe	Protective action	General public and at-risk groups	Approaching or during an event

Figures

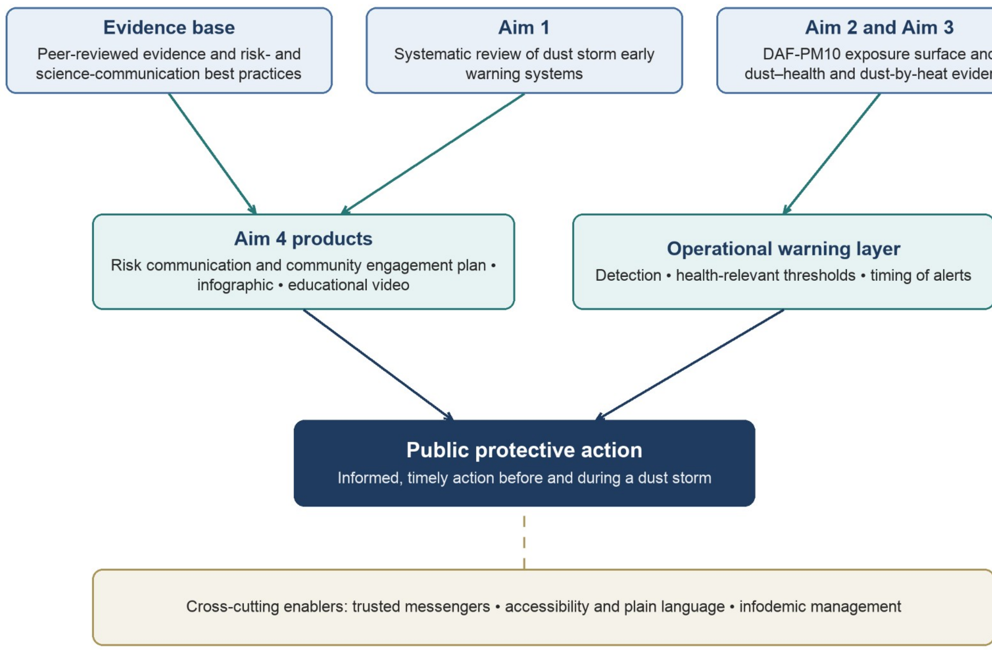

Figure 4.1. Conceptual framework linking the dissertation's aims to the Aim 4 products and to the operational warning layer that together support public protective action. The established evidence base and the Aim 1 review reinforce the Aim 4 products; the Aim 2 exposure surface and the Aim 3 health and dust-by-heat findings inform the future operational warning layer; and trusted messengers, accessibility, and infodemic management act as cross-cutting enablers.

Appendix A. Dust Storm Risk Communication and Community Engagement Plan

The full operational plan, including the audience and partner map, the crisis communication team roles, the tiered alert and messaging structure, the pre-event, event, and post-event activities, the preparedness timeline, and the monitoring and evaluation framework, is provided as a separate document that accompanies this dissertation (Risk Communication Plan).

Appendix B. The Dust Storm Infographic

The infographic exists in two versions that share an identical educational core and differ in their closing section. Version 1 closes with a Solutions section oriented toward mitigation and policy.

Version 2 closes with a How To Be Safe section oriented toward individual protective action.

Both versions were created by the candidate.

Figure B1. Dust storm infographic, version 1 (Solutions).

HEALTH IMPACTS

Exposure to wind blown dust has many consequences on the human health. It can cause premature death from cardiovasuclar and respiratory diseases.

It can also worsen the well-being through exacerbations asthma and allergic reactions as well as causeing skin and eye irritations. It has also been reported that it can cause deadly diseases like vally fever and meningiococcal meningiti.

Dust storms are defined by the WMO as the result of surface winds raising large quantities of dust into the air and reducing visibility at eye level to to less than 1000 m. In many cases wind transports loose particles from one location to another.

CAUSES

Dust storms are mainly caused by natural causes such as wind speed, wind direction as well as vast amount of sand and dessert globally. However there are also many anthropogenic causes to dust storms such as land-use, and desertification.

FREQUENCY

Dust storms are very frequent in arid and semi-arid areas. However they can occur in any place where dry sediments exist. They mainly happened in the northan hemisphere in what is known as the dust belt

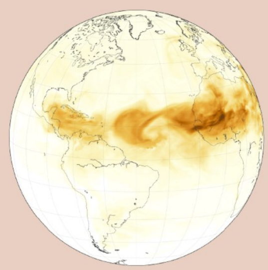

SOLUTIONS

Decreasing anthropogenic sources of emissions such as those from land-use.

Increasing reforestation efforts to mitigate the magnitude of dust storms.

Preventing further desertification

Implementing survailliance and warning systems to prevent more exposure to dust particles

Made by Saud Alomair

Figure B2. Dust storm infographic, version 2 (How To Be Safe).

HEALTH IMPACTS

Exposure to wind blown dust has many consequences on the human health. It can cause premature death from cardiovasuclar and respiratory diseases.

It can also worsen the well-being through exacerbations asthma and allergic reactions as well as causeing skin and eye irritations. It has also been reported that it can cause deadly diseases like vally fever and meningiococcal meningiti.

Dust storms are defined by the WMO as the result of surface winds raising large quantities of dust into the air and reducing visibility at eye level to to less than 1000 m. In many cases wind transports loose particles from one location to another.

CAUSES

Dust storms are mainly caused by natural causes such as wind speed, wind direction as well as vast amount of sand and dessert globally. However there are also many anthropogenic causes to dust storms such as land-use, and desertification.

FREQUENCY

Dust storms are very frequent in arid and semi-arid areas. However they can occur in any place where dry sediments exist. They mainly happened in the northan hemisphere in what is known as the dust belt

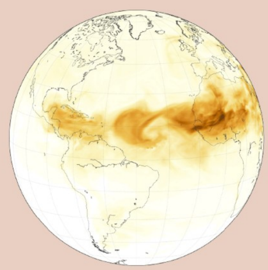

HOW TO BE SAFE

Stay indoors and avoid going outdoors during a dust storm event.

Cover your mouth and nose with a mask and protect your eyes if you have to be outdoors

Asthmatic patients should always keep an inhaler at hand in case of an asthma exacerbation

Do not drive during a dust storm event. stop the car on the side until the storm subsides. Call 112 in case of an emergency

Made by Saud Alomair

Appendix C. The Educational Video

The educational video is a narrated short film of approximately three minutes and fifty seconds that explains dust storm formation, impacts, and mitigation through original desert cinematography and concise on-screen text. The full video file accompanies this dissertation as supplementary media (file: "Dust storm (low size).mp4") and will be uploaded as a separate supplementary media file during electronic thesis submission. The storyboard below reproduces its principal scenes and on-screen text in narrative order.

Figure C1. Storyboard of the educational video, showing principal scenes and on-screen text with approximate timestamps.

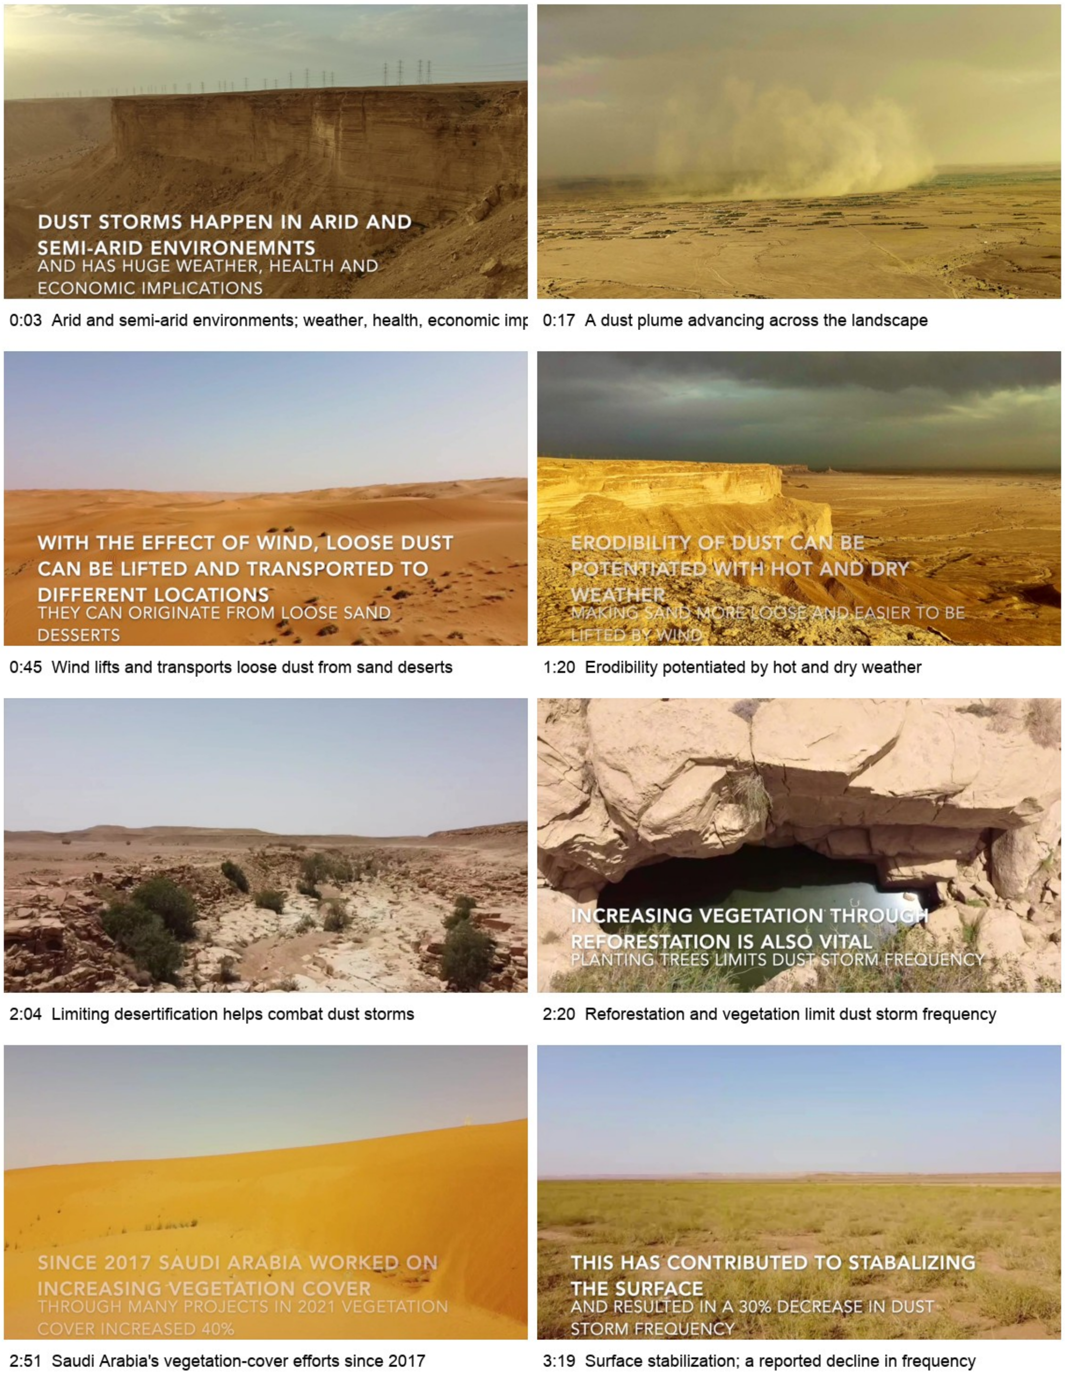

0:03 Arid and semi-arid environments; weather, health, economic imp

0:17 A dust plume advancing across the landscape

0:45 Wind lifts and transports loose dust from sand deserts

1:20 Erodibility potentiated by hot and dry weather

2:04 Limiting desertification helps combat dust storms

2:20 Reforestation and vegetation limit dust storm frequency

2:51 Saudi Arabia's vegetation-cover efforts since 2017

3:19 Surface stabilization; a reported decline in frequency

References

Al-Hemoud, A., Al-Dousari, A., Al-Shatti, A., Al-Khayat, A., Behbehani, W., & Malak, M. (2018). Health impact assessment associated with exposure to PM_{10} and dust storms in Kuwait. Atmosphere, 9(1), 6. https://doi.org/10.3390/atmos9010006

Alshammari, R. K., Alrwais, O., & Aksoy, M. S. (2024). Machine learning forecast of dust storm frequency in Saudi Arabia using multiple features. Atmosphere, 15(5), 520. https://doi.org/10.3390/atmos15050520

Basher, R. (2006). Global early warning systems for natural hazards: Systematic and people-centred. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 364(1845), 2167–2182. https://doi.org/10.1098/rsta.2006.1819

Bell, M. L., Zanobetti, A., & Dominici, F. (2013). Evidence on vulnerability and susceptibility to health risks associated with short-term exposure to particulate matter: A systematic review and meta-analysis. American Journal of Epidemiology, 178(6), 865–876. https://doi.org/10.1093/aje/kwt090

Bhattachan, A., Okin, G. S., Zhang, J., Vimal, S., & Lettenmaier, D. P. (2019). Characterizing the role of wind and dust in traffic accidents in California. GeoHealth, 3(10), 328–336. https://doi.org/10.1029/2019GH000212

Centers for Disease Control and Prevention. (2018). Crisis and emergency risk communication (CERC): 2018 update. U.S. Department of Health and Human Services.

Covello, V. T. (2003). Best practices in public health risk and crisis communication. Journal of Health Communication, 8(Suppl. 1), 5–8. https://doi.org/10.1080/713851971

Crooks, J. L., Cascio, W. E., Percy, M. S., Reyes, J., Neas, L. M., & Hilborn, E. D. (2016). The association between dust storms and daily non-accidental mortality in the United States, 1993–2005. Environmental Health Perspectives, 124(11), 1735–1743. https://doi.org/10.1289/EHP216

Csavina, J., Field, J., Félix, O., Corral-Avitia, A. Y., Sáez, A. E., & Betterton, E. A. (2014). Effect of wind speed and relative humidity on atmospheric dust concentrations in semi-arid climates. Science of the Total Environment, 487, 82–90. https://doi.org/10.1016/j.scitotenv.2014.03.138

D'Antoni, D., Smith, L., Auyeung, V., & Weinman, J. (2017). Psychosocial and demographic predictors of adherence and non-adherence to health advice accompanying air quality warning systems: A systematic review. Environmental Health, 16, 100. https://doi.org/10.1186/s12940-017-0307-4

Farrokhi, M., Sadeghimoghaddam, A., Khankeh, H., Norozi, M., & Fateh, S. (2021). Investigating the effects of dust storms on morbidity and mortality due to cardiovascular and respiratory diseases: A systematic review. Journal of Education and Health Promotion, 10, 191. https://doi.org/10.4103/jehp.jehp_1272_20

Ginoux, P., Prospero, J. M., Gill, T. E., Hsu, N. C., & Zhao, M. (2012). Global-scale attribution of anthropogenic and natural dust sources and their emission rates based on MODIS Deep Blue aerosol products. Reviews of Geophysics, 50(3), RG3005. https://doi.org/10.1029/2012RG000388

Goudie, A. S. (2013). Desert dust and human health disorders. Environment International, 63, 101–113. https://doi.org/10.1016/j.envint.2013.10.011

Goudie, A. S., & Middleton, N. J. (2001). Saharan dust storms: Nature and consequences. Earth-Science Reviews, 56(1–4), 179–204. https://doi.org/10.1016/S0012-8252(01)00067-8

Guo, Y., An, S., & Comes, T. (2022). From warning messages to preparedness behavior: The role of risk perception and information interaction in the COVID-19 pandemic. International Journal of Disaster Risk Reduction, 73, 102871. https://doi.org/10.1016/j.ijdrr.2022.102871

Houts, P. S., Doak, C. C., Doak, L. G., & Loscalzo, M. J. (2006). The role of pictures in improving health communication: A review of research on attention, comprehension, recall, and adherence. Patient Education and Counseling, 61(2), 173–190. https://doi.org/10.1016/j.pec.2005.05.004

Intergovernmental Panel on Climate Change. (2023). Climate change 2023: Synthesis report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change (Core Writing Team, H. Lee, & J. Romero, Eds.). IPCC. https://doi.org/10.59327/IPCC/AR6-9789291691647

Islam, M. M., Alharthi, M., & Alam, M. M. (2019). The impacts of climate change on road traffic accidents in Saudi Arabia. Climate, 7(9), 103. https://doi.org/10.3390/cli7090103

Kanatani, K. T., Ito, I., Al-Delaimy, W. K., Adachi, Y., Mathews, W. C., & Ramsdell, J. W. (2010). Desert dust exposure is associated with increased risk of asthma hospitalization in children. American Journal of Respiratory and Critical Care Medicine, 182(12), 1475–1481. https://doi.org/10.1164/rccm.201002-0296OC

Kashima, S., Yorifuji, T., Tsuda, T., & Eboshida, A. (2012). Asian dust and daily all-cause or cause-specific mortality in western Japan. Occupational and Environmental Medicine, 69(12), 908–915. https://doi.org/10.1136/oemed-2012-100797

Khodeir, M., Shamy, M., Alghamdi, M., Zhong, M., Sun, H., Costa, M., Chen, L.-C., & Maciejczyk, P. (2012). Source apportionment and elemental composition of $PM_{2.5}$ and PM_{10} in Jeddah City, Saudi Arabia. Atmospheric Pollution Research, 3(3), 331–340. https://doi.org/10.5094/APR.2012.037

Lal, R. (2012). Climate change and soil degradation mitigation by sustainable management of soils and other natural resources. Agricultural Research, 1(3), 199–212. https://doi.org/10.1007/s40003-012-0031-9

Lambert, A., Hallar, A. G., Garcia, M., Strong, C., Andrews, E., & Hand, J. L. (2020). Dust impacts of rapid agricultural expansion on the Great Plains. Geophysical Research Letters, 47(20), e2020GL090347. https://doi.org/10.1029/2020GL090347

Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. Psychological Science in the Public Interest, 13(3), 106–131. https://doi.org/10.1177/1529100612451018

Liu, S.-T., Liao, C.-Y., Kuo, C.-Y., & Kuo, H.-W. (2017). The effects of $PM_{2.5}$ from Asian dust storms on emergency room visits for cardiovascular and respiratory diseases. International Journal of Environmental Research and Public Health, 14(4), 428. https://doi.org/10.3390/ijerph14040428

Lwin, K. S., Tobías, A., Chua, P. L., Yuan, L., Thawonmas, R., Ith, S., Çelebi Sözener, Z., Navaratnam, A. M. D., & Hashizume, M. (2023). Effects of desert dust and sandstorms on human health: A scoping review. GeoHealth, 7(3), e2022GH000728. https://doi.org/10.1029/2022GH000728

Middleton, N. (2019). Dust storm hazards. E3S Web of Conferences, 99, 04001. https://doi.org/10.1051/e3sconf/20199904001

Neuman, C. M., & Sanderson, S. (2008). Humidity control of particle emissions in aeolian systems. Journal of Geophysical Research: Earth Surface, 113, F02S14. https://doi.org/10.1029/2007JF000780

Noar, S. M., Benac, C. N., & Harris, M. S. (2007). Does tailoring matter? Meta-analytic review of tailored print health behavior change interventions. Psychological Bulletin, 133(4), 673–693. https://doi.org/10.1037/0033-2909.133.4.673

Pu, B., Jin, Q., Ginoux, P., & Yu, Y. (2022). Compound heat wave, drought, and dust events in California. Journal of Climate, 35(24), 8133–8152. https://doi.org/10.1175/JCLI-D-21-0889.1

Reynolds, B., & Seeger, M. W. (2005). Crisis and emergency risk communication as an integrative model. Journal of Health Communication, 10(1), 43–55. https://doi.org/10.1080/10810730590904571

Roozenbeek, J., & van der Linden, S. (2019). Fake news game confers psychological resistance against online misinformation. Palgrave Communications, 5, 65. https://doi.org/10.1057/s41599-019-0279-9

Rublee, C. S., Sorensen, C. J., Lemery, J., Wade, T. J., Sams, E. A., Hilborn, E. D., & Crooks, J. L. (2020). Associations between dust storms and intensive care unit admissions in the United States, 2000–2015. GeoHealth, 4(8), e2020GH000260. https://doi.org/10.1029/2020GH000260

Sättele, M., Bründl, M., & Straub, D. (2016). Quantifying the effectiveness of early warning systems for natural hazards. Natural Hazards and Earth System Sciences, 16(1), 149–166. https://doi.org/10.5194/nhess-16-149-2016

Tangcharoensathien, V., Calleja, N., Nguyen, T., Purnat, T., D'Agostino, M., Garcia-Saiso, S., Landry, M., Rashidian, A., Hamilton, C., AbdAllah, A., Ghiga, I., Hill, A., Hougendobler, D., van Andel, J., Nunn, M., Brooks, I., Sacco, P. L., De Domenico, M., Mai, P., … Briand, S. (2020). Framework for managing the COVID-19 infodemic: Methods and results of an online, crowdsourced WHO technical consultation. Journal of Medical Internet Research, 22(6), e19659. https://doi.org/10.2196/19659

Tuong, W., Larsen, E. R., & Armstrong, A. W. (2014). Videos to influence: A systematic review of effectiveness of video-based education in modifying health behaviors. Journal of Behavioral Medicine, 37(2), 218–233. https://doi.org/10.1007/s10865-012-9480-7

United Nations Environment Programme, World Meteorological Organization, & United Nations Convention to Combat Desertification. (2022). Global assessment of sand and dust storms. United Nations Environment Programme. https://www.unccd.int/resources/publications/sand-and-dust-storms-compendium

United Nations Office for Disaster Risk Reduction. (2015). Sendai framework for disaster risk reduction 2015–2030. United Nations. https://www.undrr.org/publication/sendai-framework-disaster-risk-reduction-2015-2030

Vergadi, E., Rouva, G., Angeli, M., & Galanakis, E. (2022). Infectious diseases associated with desert dust outbreaks: A systematic review. International Journal of Environmental Research and Public Health, 19(11), 6907. https://doi.org/10.3390/ijerph19116907

Vodonos, A., Friger, M., Katra, I., Krasnov, H., Zahger, D., Schwartz, J., & Novack, V. (2015). Individual effect modifiers of dust exposure effect on cardiovascular morbidity. PLOS ONE, 10(9), e0137714. https://doi.org/10.1371/journal.pone.0137714

Wang, J. X. L. (2015). Mapping the global dust storm records: Review of dust data sources in supporting modeling/climate study. Current Pollution Reports, 1(2), 82–94. https://doi.org/10.1007/s40726-015-0008-y

Watanabe, M., Noma, H., Kurai, J., Shimizu, A., Sano, H., Kato, K., Mikami, M., Ueda, Y., Tatsukawa, T., Yamasaki, A., Igishi, T., & Shimizu, E. (2015). Association of sand dust particles with pulmonary function and respiratory symptoms in adult patients with asthma in western Japan using light detection and ranging: A panel study. International Journal of Environmental Research and Public Health, 12(10), 13038–13052. https://doi.org/10.3390/ijerph121013038

Witte, K. (1992). Putting the fear back into fear appeals: The extended parallel process model. Communication Monographs, 59(4), 329–349. https://doi.org/10.1080/03637759209376276

World Health Organization. (2017). Communicating risk in public health emergencies: A WHO guideline for emergency risk communication (ERC) policy and practice. World Health Organization. https://www.who.int/publications/i/item/9789241550208

World Meteorological Organization. (2022). Early warnings for all: Executive action plan 2023–2027. World Meteorological Organization.

Zhang, X., Zhao, L., Tong, D. Q., Wu, G., Dan, M., & Teng, B. (2016). A systematic review of global desert dust and associated human health effects. Atmosphere, 7(12), 158. https://doi.org/10.3390/atmos7120158

Conclusion

This dissertation set out to build, in sequence, the evidence that a dust storm warning system needs, and to do so for a state where dust is an under-characterized but growing public health concern. The four aims trace a single pathway from hazard to harm to action. A systematic review established what dust early warning systems do well and where they fail, an exposure surface made dust-related particulate conditions measurable across populated California at a daily resolution, an epidemiologic analysis tested whether that exposure is associated with acute cardiopulmonary illness and under what conditions, and a translational chapter carried the accumulated evidence toward the communities the warning is meant to protect.

The four aims reinforce one another in specific ways. The review (Chapter 1) found that the dust warning literature has developed one component of a people-centered system in depth, the detection and forecasting of the hazard, while leaving risk knowledge, communication, and response thinly studied across more than two decades of work. That imbalance framed the rest of the dissertation, it justified investing in exposure science and health evidence that could give the neglected components a rationale, and it identified communication and response as the links most in need of strengthening. The exposure work (Chapter 2) answered the first prerequisite by producing a daily, statewide DAF-PM_{10} surface at population-weighted centroids, externally anchored against composition-based dust measurements and accompanied by a transparent rule-based comparator. That surface preserved meaningful rank-order agreement with the independent reference and degraded smoothly with monitor-to-centroid distance, a pattern that points to geometric separation rather than model gap as a leading source of apparent disagreement, and it localized the heaviest chronic burden to the southern San Joaquin and Imperial Valleys, where elevated concentrations coincide with dense population.

The health analysis (Chapter 3) then carried that surface into the largest dust-oriented PM_{10} case-crossover conducted to date, roughly twenty-four million strata per broad outcome. A one-interquartile-range increase in dust-attributable PM_{10} raised the odds of a broad respiratory encounter by about one percent on the same day, peaking one day after exposure and accumulating over the following week to roughly one and a half times that figure, while the cardiovascular effect was smaller and essentially same-day. The respiratory concentration-response was strongly nonlinear, concentrated in asthma and chronic obstructive pulmonary disease, amplified specifically by peak daytime heat, and positive across nearly every populated air basin though varying about three-fold in magnitude; the cardiovascular effect was near-monotone, concentrated in hypertensive disease, and confined to coastal metropolitan basins. Both effects were null in the chronically saturated Salton Sea basin, which the design cannot evaluate well. These results matter for the warning system the dissertation is ultimately about, because they begin to identify the exposure levels, lag windows, and compound heat conditions at which a dust alert would protect the most health. The translational chapter (Chapter 4) used exactly that logic, proposing a coordinated risk communication plan, an evidence-based infographic, and an educational video, and showing how the exposure metric and the dust-by-heat findings can inform the criteria, thresholds, and timing of an operational system, with the systematic review's findings reinforcing the priority placed on communication and response.

Taken together, the chapters add something that none provides alone, an exposure-to-action framework for a climate-sensitive hazard that existing tools serve poorly. The dissertation moves California-relevant dust science from event documentation and source characterization toward continuous, population-relevant exposure assignment, it links that exposure to morbidity at a scale and spatial resolution not previously available in the state, and it connects the resulting evidence to the communication and response components that the warning literature has most neglected. Several cross-cutting themes emerge from reading the aims as a whole. Heat is not a secondary variable but an active modifier of the dust-respiratory association, which means that dust and heat warnings are most useful when designed together, a conclusion strengthened by evidence that compound heat, drought, and dust events are projected to become more frequent in California (Pu et al., 2022). The burden of dust is as much chronic as episodic, the communities of the southern San Joaquin and Imperial Valleys experience sustained high exposure rather than occasional spikes, which is the more actionable distinction for surveillance and intervention, and which also marks the limit of a case-crossover design built on day-to-day contrast. And the value of an exposure surface for epidemiology lies in its rank ordering across space and time, not in reproducing absolute concentrations, a property the validation was designed to test directly.

The dissertation is deliberate about what it does not establish. The exposure surface is evaluated against the strongest available comparator rather than an ideal truth standard, and it is most accurately described as a dust-informed particulate surface rather than a chemically isolated dust mass, this framing is carried consistently into the health analysis, where the estimated associations are attributed to a dust-related, dust-dominant exposure rather than to mineral dust alone. The health estimates are short-term associations conditional on the matched design, not causal effects, and because temperature and humidity helped generate the exposure surface and also enter the health model as confounders, the reported associations carry a meteorological double-adjustment that biases them toward the null. They are, in that precise sense, conservative. The communication products are grounded in evidence and in established communication science but have not yet been evaluated with their intended audiences, so their effect on knowledge, trust, and protective behavior is expected rather than demonstrated. These limitations are not incidental, they define the agenda that the dissertation hands to future work.

The chronically saturated Salton Sea regime, which the case-crossover cannot assess, calls for a cohort or long-term ecological design that compares exposed populations against matched, non-saturated populations on a between-population contrast (Frie et al., 2019; Parajuli & Zender, 2018). An explicit multi-pollutant model that adds non-dust PM_{10} and fine particulate would move the design closer to the source-apportioned standard of the European multi-city literature (Stafoggia et al., 2016). Coupling the dust-by-heat estimates to downscaled climate projections would translate the present associations into California-specific projected burden under mid-century and end-of-century conditions (Pu et al., 2022). And the translational program needs evaluation, pretesting the infographic and video with target audiences, measuring real protective behavior and post-event health and safety outcomes, and connecting the DAF-PM_{10} metric and the compound-hazard logic to operational dust forecasting and to the international Early Warnings for All effort (World Meteorological Organization, 2022).

Dust is a hazard that warming, drought, and land disturbance are making more common, and its health burden falls hardest on the people least reached by conventional warning. This dissertation has tried to make that hazard more measurable, more clearly tied to acute illness, and more actionable in the places where it does its harm. The next dust season will arrive in California whether or not the science is ready for it. The aim of this dissertation has been to make the science, and the warning it can support, a little readier than before.

References

Frie, A. L., Garrison, A. C., Schaefer, M. V., Bates, S. M., Botthoff, J., Maltz, M., Ying, S. C., Lyons, T., Allen, M. F., Aronson, E., & Bahreini, R. (2019). Dust sources in the Salton Sea basin: A clear case of an anthropogenically impacted dust budget. Environmental Science & Technology, 53(16), 9378–9388. https://doi.org/10.1021/acs.est.9b02137

Parajuli, S. P., & Zender, C. S. (2018). Projected changes in dust emissions and regional air quality due to the shrinking Salton Sea. Aeolian Research, 33, 82–92. https://doi.org/10.1016/j.aeolia.2018.05.004

Pu, B., Jin, Q., Ginoux, P., & Yu, Y. (2022). Compound heat wave, drought, and dust events in California. Journal of Climate, 35(24), 8133–8152. https://doi.org/10.1175/JCLI-D-21-0889.1

Stafoggia, M., Zauli-Sajani, S., Pey, J., Samoli, E., Alessandrini, E., Basagaña, X., et al. (2016). Desert dust outbreaks in Southern Europe: Contribution to daily PM_{10} concentrations and short-term associations with mortality and hospital admissions. Environmental Health Perspectives, 124(4), 413–419. https://doi.org/10.1289/ehp.1409164

Tobías, A., Karanasiou, A., Amato, F., Roqué, M., Lwin, K. S., Yuan, L., Ith, S., Wai, H. Z., Chua, P. L., Solà, I., Renzi, M., Stafoggia, M., Querol, X., & Hashizume, M. (2025). Short-term exposure to desert dust and sandstorms and all-cause and cause-specific mortality and morbidity: A systematic review and meta-analysis. Environment International, 196, 109277. https://doi.org/10.1016/j.envint.2025.109277

World Meteorological Organization. (2022). Early warnings for all: Executive action plan 2023–2027. World Meteorological Organization.