Environmental and Ecological Statistics
https://doi.org/10.1007/s10651-026-00731-8

Climate-informed forecasting of the respiratory pharmaceutical demand in Greece from a spatio-temporal panel dataset

Viviana Schisa[1] · **Matteo Farnè**[1]

Received: 12 December 2025 / Revised: 25 March 2026 / Accepted: 30 April 2026

Abstract
Respiratory medicines are among the first lines of defence when weather and climate push vulnerable lungs past their limits. This study quantifies how atmospheric conditions shape weekly prescription volumes and examines what continued global warming implies for pharmaceutical planning in Greece. A national retail panel of prescription respiratory sales for 20 regions (2016–2023) is combined with high-resolution meteorological reanalysis to estimate two classes of models: a Spatial Lag of X panel with region and week-of-year fixed effects, and a climate-augmented fixed-effects distributed lag forecaster. The spatial specification shows that contemporaneous conditions in neighbouring regions, especially warmer temperatures and stronger winds, exert more systematic effects on local demand than purely local shocks—positive for temperature, negative for wind—with spillovers concentrated within distance bands of a few hundred kilometres. The forecasting model couples short distributed lags of climatic variables with autoregressive dynamics and attains one-year-ahead accuracy with mean absolute percentage errors around 11%. A purely exogenous variant driven only by climatic regressors is then run on weekly projections from global climate models under alternative forcing pathways, yielding scenario-conditioned trajectories of national respiratory pharmaceutical needs. By the late 2020s, these projections indicate peak-season demand roughly 25–35% above recent (2022–2023) levels and about 45–55% above pre-pandemic volumes. Taken together, the results support climate-aware, spatially resolved procurement strategies and underscore the value of integrating routine pharmaceutical data into health-system adaptation to a warming atmosphere.

Keywords Respiratory pharmaceuticals · Climate change · Spatial spillovers · Pharmaceutical supply planning

Extended author information available on the last page of the article

Published online: 21 May 2026

1 Introduction

Climate change is widely regarded as one of the defining challenges of recent decades, exerting deep and tightly interlinked effects on environmental systems, population health, and economic activity worldwide (Barros et al. 2014; Carleton et al. 2022; Kotz et al. 2024). In addition to gradual shifts in mean temperature and precipitation, it modifies the frequency, intensity, and spatial distribution of extreme events, thereby altering patterns of exposure and vulnerability across regions and population groups.

Climate variability has long been recognised as a key determinant of human health operating through multiple, interacting pathways (Haines et al. 2006). These mechanisms include direct injuries and mortality associated with extreme weather events, cardiovascular and respiratory strain during heatwaves, mental health consequences arising from acute climate-related disasters and chronic climate-related distress, and changes in the transmission dynamics of vector- and water-borne diseases (Rocklöv and Dubrow 2020; Cianconi et al. 2020; Khraishah et al. 2022). Climate change intensifies these mechanisms by shifting the distribution of environmental conditions—rising near-surface temperatures, more frequent heatwaves, altered humidity and precipitation regimes, and a higher occurrence of compound extremes.

These shifts modify the timing and intensity of demand for medicines, transmitting environmental variability into healthcare utilisation and pharmaceutical consumption (Redshaw et al. 2013). This paper focuses on Greece in the Eastern Mediterranean, where recent decades have seen pronounced warming, more frequent and persistent heatwaves, and shifts in precipitation that increase drought and fire-weather risk (Zerefos et al. 2011; Tolika et al. 2012). The combination of these climatic pressures with heterogeneous geography and population exposure provides a highly volatile environment in which to identify and forecast climate-sensitive respiratory pharmaceutical demand.

Health shocks therefore transmit to the demand side of the healthcare system, affecting public budgets, procurement planning, and private supply chains, where decisions must be taken under uncertainty and with adjustment costs. At the same time, pharmaceutical supply chains are themselves exposed to climate stressors: from sourcing to manufacturing and distribution, each stage can be perturbed by heatwaves, cold snaps, floods, or transport disruptions. Recent years have shown how unexpected emergencies can impede access to essential goods, as starkly illustrated during the COVID-19 pandemic (Ranney et al. 2020; Toković 2023). In such settings, seasonal regularities justify medium-term planning, whereas episodic surges require rapid responses that are feasible only if risk is anticipated and conveyed through calibrated predictive intervals.

These considerations are particularly salient in fiscally constrained healthcare systems, especially where heterogeneous geography intersects with climate exposure and amplifies risk; in Greece, health-system constraints and procurement institutions further raise the penalty for forecast error (Economou et al. 2017). In this context, geographically resolved forecasts that integrate climate information and

quantify uncertainty are valuable inputs for procurement, inventory management, and spatially targeted preparedness.

A still limited but recently growing literature links climate variability and change to respiratory pharmaceutical use and situates demand forecasting within fragile medicine supply chains. Abir et al. (2025), for example, simulate future climate scenarios and estimate that demand for asthma medications such as albuterol could rise by 1.5 to 3% by 2040, driven by higher asthma prevalence projected across nearly all age groups. For Greece, Schisa and Farnè (2025) provide one of the first comprehensive analyses combining frequency-domain inference, machine learning, and panel designs to map climate–pharmaceutical relationships. The present study builds on that framework by exploiting the full 20-region weekly panel in a spatially explicit setting, embedding climate covariates in two-way fixed-effects forecasters, and linking these models directly to CMIP6 climate projections to obtain scenario-conditioned trajectories of pharmaceutical demand.

Thus, while a substantial literature links climate and environmental exposures to morbidity and mortality, evidence that *directly* quantifies and rigorously *forecasts* climate-driven variation in pharmaceutical demand remains limited, particularly at weekly resolution and in a spatially explicit framework. From a statistical perspective, there is a need for spatio-temporal panel models that (i) exploit cross-sectional and seasonal structure, (ii) accommodate spatial spillovers in climate exposures, and (iii) can be embedded in a scenario-based forecasting pipeline driven by climate projections.

Against this backdrop, the paper develops a spatio-temporal, weekly frequency analysis of respiratory pharmaceutical demand in Greece and pursues three complementary objectives. First, it conducts an exploratory spatial econometric analysis to characterise how contemporaneous and synoptic-scale climate conditions, together with their spatial spillovers, are associated with regional respiratory pharmaceutical demand. This is achieved by estimating a Spatial Lag of X (SLX) panel model with two-way fixed effects, which disentangles local from cross-regional climate channels and provides an interpretable description of how climate variability propagates into regional pharmaceutical sales. Second, it develops a two-way fixed-effects panel model tailored for one-year-ahead forecasting of regional demand, combining autoregressive dynamics with climate covariates summarised through parsimonious distributed-lag polynomials. Third, it uses a purely exogenous variant of this panel forecaster as a scenario engine for medium-term projections, linking weekly demand to climate trajectories from the Coupled Model Intercomparison Project Phase 6 (CMIP6) under alternative Shared Socioeconomic Pathways.

All in all, this threefold approach allows historical data to be fully exploited to identify and disentangle short-term impacts onto respiratory medicine demand in Greece, so that we can model and forecast respiratory medicine consumptions in Greek regions in the short-run, also incorporating spatial dependence patterns beyond temporal and climatic ones. On top of that, by integrating CMIP6 climate projections with our validated model, we also manage to extrapolate medium-run plausible trajectories for respiratory medicine demand in Greece. By this approach, we empirically prove that climate-informed forecasts can support procurement and budgeting decisions by providing probabilistic, geographically resolved, and

operationally credible signals of upcoming demand pressure. This is particularly relevant for Greece, where health-system constraints and procurement institutions further raise the penalty for forecast error, and spending adjustments in strategic demand management have not always been grounded in a needs-based assessment (Economou et al. 2017; Deloitte Business Solutions 2025).

The remainder of the paper is organised as follows. Section 2 describes the pharmaceutical and climate data, as well as the construction of scenario-consistent climate drivers. Section 3 presents the econometric framework, introducing both the climate-augmented two-way fixed-effects panel forecaster (FE–CLIM–AR and its exogenous variant FE–CLIM–X) and the SLX panel model for spatial spillovers and distance-band effects. Section 4 reports empirical results on spatial spillovers, forecasting performance, and scenario-based projections under alternative climate pathways. Finally, Sect. 5 discusses the main findings and their implications for climate-sensitive pharmaceutical planning and outlines directions for future work.

2 Data

The empirical setting is a weekly panel for Greece spanning January 2016 to June 2023 and covering 20 pharmaceutical regions. Weekly respiratory pharmaceutical sales and climate variables are used to develop econometric models that (i) quantify spatial spillovers in demand and (ii) support forecasting and scenario-based projections to inform climate-sensitive healthcare planning.

2.1 Pharmaceutical data

The pharmaceutical demand data (y), expressed in terms of packs (single dispensed packages) independently of pack size, dosage strength, or route of administration, were provided by *Alira Health*[1] and originate from a retail *sell-out* panel that the company designs and maintains. The panel comprises approximately 1,200 pharmacies distributed across Greece, sampled to ensure geographic representativeness with respect to the 20 pharmaceutical areas. Data are collected automatically via software integration with pharmacy enterprise resource planning (ERP) systems, which acquire de-identified transaction records (volumes, values, de-identified e-prescription metadata, and product codes) on each business day. The process is fully automated and timestamped, enabling routine refreshes and auditable traceability while enforcing GDPR-compliant anonymisation and standardised quality controls.

The geographical partition is fixed over the study period and broadly aligns with major urban hubs, mainland macro-areas, and island groups (see Section S1 of the Supplementary Material for regional centroids and mapping to the retail panel). Consistent with Greece's settlement pattern, population and pharmacy density are highest in the Athens and Thessaloniki metropolitan areas, while the island and

[1] Alira Health, https://www.alirahealth.com/

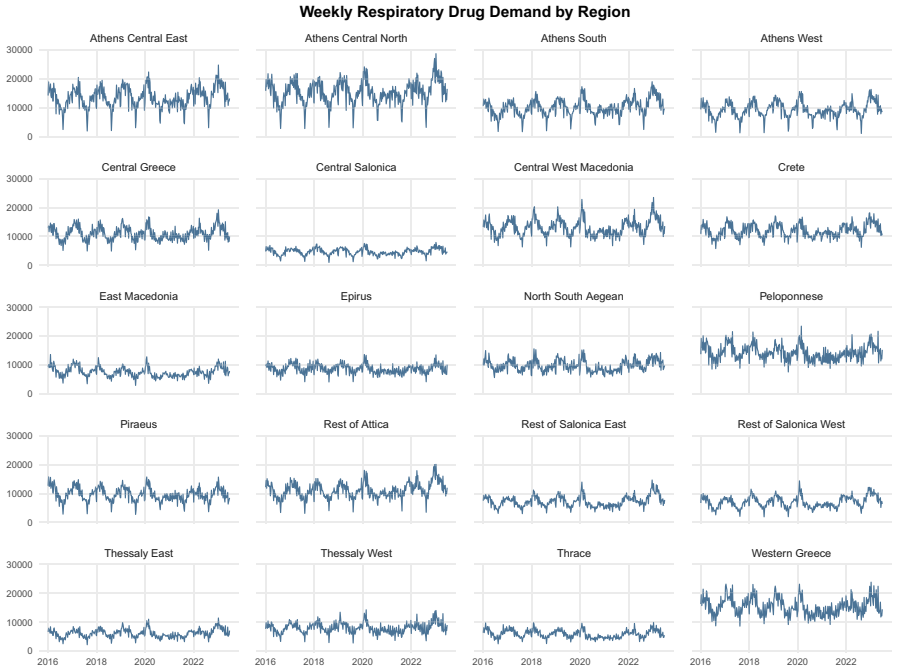

Fig. 1 Weekly respiratory pharmaceutical demand for 20 pharmaceutical regions (2016–2023)

peripheral territories are more sparsely populated but remain covered by the retail panel.

Although transaction records are available down to ATC Level 5, the analysis is conducted at Level 1[2] because the objective is to quantify the total climate-related impact on respiratory demand, for which therapeutic substitution across subclasses (e.g. between bronchodilators and anti-inflammatories) is part of the relevant behavioural response. Moreover, only prescription medications are considered, as they are subject to stricter regulatory control and monitoring, thus offering more reliable data for analysis.

In Fig. 1, the weekly series for all twenty pharmaceutical regions display a clear winter-dominant seasonal pattern: peaks recur in late autumn and winter, while troughs fall in late spring and summer, in line with respiratory demand dynamics. Within each panel, week-to-week movements are smooth rather than erratic, indicating strong short-run persistence and motivating autoregressive dynamics in subsequent models. The COVID-19 period leaves a visible imprint across most regions:

[2] The Anatomical Therapeutic Chemical (ATC) system (https://www.whocc.no/atc_ddd_index/), widely used in pharmacoepidemiology, provides a hierarchical therapeutic classification: Level 1 identifies the anatomical main group, Level 2 the therapeutic main group, Levels 3–4 the pharmacological and chemical subgroups, and Level 5 the chemical substance. The focus here is ATC Level 1 = R (Respiratory system), which aggregates respiratory agents regardless of specific mechanism or molecule.

during 2020–2021, levels are systematically lower and seasonal peaks appear compressed, followed by a rebound from late 2021 onward with sharper winter maxima.

At the same time, Fig. 1 reveals pronounced cross-regional heterogeneity in both average levels and in the amplitude and sharpness of seasonality. Volumes are highest in the Athens and Thessaloniki metropolitan areas and in densely populated mainland regions, whereas island and peripheral areas operate at markedly lower levels. Some territories exhibit very sharp winter spikes, others smoother and lower-amplitude cycles, and volatility is not uniform across space, with certain regions showing more frequent extreme excursions than others. Such spatial non-homogeneity cautions against reliance on aggregates alone: meaningful within-region dynamics may be masked, parameter estimates may be distorted by aggregation, and climate–demand linkages may be confounded by unobserved, time-invariant differences across areas. This motivates a regional panel design with two-way fixed effects.

2.2 Climate data

Climate inputs are drawn primarily from ERA5, the global reanalysis produced by the European Centre for Medium-Range Weather Forecasts (ECMWF) (Hersbach et al. 2020). The analysis uses ERA5 near-surface variables: 2 m air temperature (τ), 10 m wind components used to derive wind speed, total precipitation (r), total cloud cover (c), and specific humidity (q). The horizontal wind is represented by its 10 m components, with u_{10} denoting the eastward (zonal) flow and v_{10} the northward (meridional) flow, both in m s^{-1}. Hourly wind speed is then obtained as

$$w_{i,h} = \sqrt{u_{10,i,h}^2 + v_{10,i,h}^2}\,. \tag{1}$$

In addition, weekly temperature variability (T_v) is defined as the standard deviation of daily mean temperature within each week. For region i and week t,

$$T_{v;(i,t)} = \sqrt{\frac{1}{7}\sum_{d=1}^{7}\left(\tau_{i,d} - \overline{\tau}_{i,t}\right)^2}, \tag{2}$$

where $\tau_{i,d}$ is the daily mean 2 m temperature on day d in region i, and $\overline{\tau}_{i,t}$ is the corresponding weekly mean.

All climate series are sampled at point locations corresponding to the most populous city within each of the twenty territories (see Table S1 in the Supplementary Material). For each coordinate, the nearest ERA5 land grid point is selected and a land–sea mask is applied in coastal areas to avoid contamination from ocean cells. This point-based sampling at population-dense urban locations emphasises conditions where demand is concentrated and avoids diluting urban signals with sparsely inhabited rural or mountainous areas, which would dominate an area-weighted polygon mean despite contributing less to retail utilisation.

All variables are handled in units and with aggregation rules that preserve their physical meaning—state variables are averaged in time, fluxes are summed—and

are converted, where appropriate, to conventions common in the climate–health literature. A compact reference with variable names, symbols, native ERA5 units, and analysis units after conversion is provided in Table S2 of the Supplementary Material.

For projecting pharmaceutical demand, climate projections from the Coupled Model Intercomparison Project Phase 6 (CMIP6), distributed via the Copernicus Climate Change Service Climate Data Store (2021), are used to extend the forecasting framework under future climate scenarios. Two Shared Socioeconomic Pathways (SSPs) are considered to bracket a plausible range of future forcing: SSP2-4.5 (intermediate) and SSP5-8.5 (very high). The variables ingested are near-surface air temperature and precipitation at daily resolution.

To represent structural uncertainty across modelling systems, projections are taken from four widely used CMIP6 models with complementary physics and resolutions. For each model and scenario, daily fields are extracted for the European domain, then mapped to the twenty territories using the same point-sampling strategy as for ERA5: for each territory, the nearest land grid cell to the most populous city is selected, and temporal aggregation mirrors the historical pipeline to obtain weekly covariates. For each week, region, variable, and scenario, the multi-model ensemble mean is computed to obtain a single scenario-consistent driver. To preserve transparency, projections are driven by CMIP6 weekly covariates processed with the same temporal aggregation as ERA5. Results are interpreted as scenario-conditioned trajectories; absolute levels may be affected by model mean-state differences. The focus is on relative changes with respect to recent baselines, and report uncertainty bands. Further details on the CMIP6 models, domains, and pre-processing steps are reported in Section S1 of the Supplementary Material.

3 Methods

The empirical strategy combines a spatial econometric component with a climate-augmented panel forecasting framework. The first part uses a Spatial Lag of X (SLX) specification with two-way fixed effects to characterise how contemporaneous local and neighbouring climate conditions co-move with regional respiratory pharmaceutical demand, and to quantify distance-dependent spillovers. The second part introduces two-way fixed-effects distributed-lag models that embed short climate histories in autoregressive panels designed for one-year-ahead prediction and for scenario-based projections under alternative climate pathways. The two components are designed for different purposes and therefore need not use identical climate covariate sets: the SLX is intended to characterise contemporaneous local and neighbouring co-movements and spatial spillover channels, whereas the forecasting specification is deliberately parsimonious and retains only the climate drivers that proved most robust and operationally coherent for out-of-sample prediction and scenario-based projection. In all cases, inference is based on wild cluster bootstrap procedures that are robust to serial correlation and cross-sectional dependence. Throughout, regions are indexed by $i = 1, \dots, N$ and weeks by $t = 1, \dots, T$;

$s(t) \in \{1, \ldots, 52\}$ denotes the week-of-year corresponding to week t, and $\delta_{s(t)}$ are week-of-year fixed effects.

3.1 SLX model for spatial spillovers

We capture these linkages with a Spatial Lag of X (SLX) specification adapted to a panel setting with two-way fixed effects[3] (Halleck Vega and Elhorst 2015; Baltagi 2005). Let $y_{i,t}$ denote weekly pharmaceutical demand in region i at week t, and let $\mathcal{C}_{\text{spat}} = \{\tau, T_v, q, w, r, c\}$ collect 2 m temperature, intra-week temperature variability, specific humidity, 10 m wind speed, total precipitation, and cloud cover. The SLX panel with two-way fixed effects for $i = 1, \ldots, N, \quad t = 1, \ldots, T$ is

$$y_{i,t} = \alpha_i + \delta_{s(t)} + \sum_{C \in \mathcal{C}_{\text{spat}}} \beta_C \, C_{i,t} + \sum_{C \in \mathcal{C}_{\text{spat}}} \theta_C \, C^{\text{SL}}_{i,t} + u_{i,t}, \tag{3}$$

where α_i and $\delta_{s(t)}$ are region and week-of-year fixed effects, respectively, and $u_{i,t}$ is an idiosyncratic error term. The spatial lag of covariate C is

$$C^{\text{SL}}_{i,t} = \sum_{m \neq i} w_{im} \, C_{m,t}, \tag{4}$$

with $W = \{w_{im}\}$ a nonnegative row-standardised weight matrix, $w_{ii} = 0$ and $\sum_m w_{im} = 1$ for all i. Here and throughout, *spillover effects* refer to the partial association between $y_{i,t}$ and contemporaneous climate exposures in *other* regions, operationalised through spatially lagged covariates $C^{\text{SL}}_{i,t}$. Accordingly, θ_C captures how variation in the W-weighted-average exposure of neighbouring regions is associated with demand in region i, conditional on own-region exposures, region fixed effects, and week-of-year fixed effects. The local component $\sum_{C \in \mathcal{C}_{\text{spat}}} \beta_C \, C_{i,t}$ is included alongside the spatial lags to isolate spillover channels and mitigate omitted-variable bias, since $C_{i,t}$ and $C^{\text{SL}}_{i,t}$ are typically correlated under spatial persistence. Coefficients β_C measure within-region associations, net of fixed effects. Positive (negative) entries of $\boldsymbol{\theta} = (\theta_C)_{C \in \mathcal{C}_{\text{spat}}}$ indicate that higher neighbouring exposures are associated with higher (lower) local demand, consistent with cross-regional mobility, integrated supply chains, or supra-regional policy coordination. A spatially homogeneous one-unit increase in C (i.e. $C_{i,t}$ and $C^{\text{SL}}_{i,t}$ both rise by one due to row-standardisation) yields a "total effect" equal to $\beta_C + \theta_C$. In this linear-in-levels specification, θ_C can be read as the marginal change in weekly demand (packs) associated with a one-unit increase in the weighted-average exposure of neighbouring regions, holding own-region exposures and fixed effects constant.

[3] Two-way fixed effects are preferred to random effects since unobserved, time-invariant regional factors are plausibly correlated with the regressors. Random-effects consistency requires $\mathbb{E}[\alpha_i \mid X_{i\cdot}] = 0$, an assumption unlikely in this setting; by conditioning on α_i and week-of-year fixed effects, Fixed-Effects Ordinary Least Squares (FE-OLS, Bergé (2018)) remains consistent.

The main analysis uses an inverse-distance-squared kernel for W, built from great-circle (Haversine) distances between regional centroids, so that nearby regions exert stronger influence and weights taper rapidly with distance. Alternative weight constructions, including inverse-distance and k-nearest neighbours, are reported in Section S2.1 of the Supplementary Material, while distance-band decompositions of spillovers are discussed in Section S2.2.

Parameters are estimated by OLS with two-way fixed effects under a strict exogeneity assumption on $u_{i,t}$. Inference relies on a two-way wild cluster bootstrap (Cameron et al. 2008; Roodman et al. 2019; MacKinnon et al. 2023), clustering by region and calendar week to accommodate serial correlation within regions and cross-sectional dependence within weeks; further implementation details are provided in Section S2.

3.2 Two-way fixed-effects distributed lag models for forecasting and future projections

Figure 1 reveals pronounced cross-regional heterogeneity in both levels and in the amplitude and sharpness of seasonality, cautioning against reliance on aggregates alone: meaningful within-region dynamics may be masked, parameter estimates distorted by aggregation, and climate–demand linkages confounded by unobserved, time-invariant differences across areas. A regional panel design with two-way fixed effects is therefore warranted.

Let $y_{i,t}$ denote weekly pharmaceutical demand in region $i = 1, \dots, N$ and week $t = 1, \dots, T$. The baseline specification is a two-way fixed-effects distributed lag model,

$$y_{i,t} = \alpha_i + \delta_{s(t)} + \sum_{\ell=1}^{p_y} \varphi_\ell \, y_{i,t-\ell} + \sum_{C \in \mathcal{C}_{\text{reg}}} \sum_{\ell=1}^{L_C} \beta_{C,\ell} \, C_{i,t-\ell} + u_{i,t}, \tag{5}$$

where α_i and $\delta_{s(t)}$ are region and week-of-year fixed effects, respectively, p_y denotes the number of autoregressive lags of the outcome, and $\mathcal{C}_{\text{reg}} = \{\tau, T_v, p\}$ collects temperature, intra-week temperature variability, and precipitation. The use of this restricted climate set reflects the specific forecasting objective of the panel model. In particular, we retain a parsimonious subset of climate drivers—temperature, intra-week temperature variability, and precipitation—that emerged as the most robust predictors in the forecasting pipeline while remaining consistent with the CMIP6-based scenario projections. Variables such as specific humidity and wind speed are instead retained in the SLX, where the objective is to characterise contemporaneous spatial co-movements and spillover channels and to disentangle physically distinct pathways of neighbouring climate exposure. The coefficients $\beta_{C,\ell}$ capture within-region, week-to-week associations net of time-invariant regional heterogeneity and of the common seasonal profile (Baltagi 2005). Conditioning on $(\alpha_i, \delta_{s(t)})$ forces identification to come from deviations around region-specific levels and the national annual cycle. Accordingly, climate covariates are reported in physical units rather than standardised units. In this panel setting, standardization is not unique, since it

may be defined globally, region-wise, or over different time windows; the variables are therefore retained in levels so as to preserve the direct interpretability of coefficients in degrees Celsius, millimetres, metres per second, and related physical units.

Panel unit-root tests of Im–Pesaran–Shin and Maddala–Wu reject the null of a unit root in regional respiratory pharmaceutical demand, indicating that the panel is stationary in levels. This supports estimating the models in levels rather than differences, thereby preserving the low-frequency climatic signal that is relevant for climate-change–driven shifts in demand.

Two concrete implementations of (5) are used. The first, denoted FE–CLIM–AR, includes four autoregressive lags of the outcome and short distributed lags of climatic covariates. It is estimated by FE–OLS and evaluated under expanding-window protocols for one-year-ahead prediction, with the lag structure for each climate variable selected on a fixed grid that balances persistence, climatic information, and parsimony. The candidate grid allowed each climate variable either to be excluded, to enter through unrestricted lag blocks, or to enter through Almon polynomial distributed-lag blocks with lag spans up to six weeks and polynomial degrees from 0 to 3. The preferred specification was then selected by minimising average out-of-sample RMSE across expanding-window forecast origins over a 52-week horizon. This choice was intended to make the lag structure data-driven rather than discretionary, while aligning model selection with the operational one-year forecasting objective. Details on the model grid and expanding-window validation protocol are reported in Section S3.1 of the Supplementary Material, while the polynomial distributed-lag representation is discussed in Section S3.2.

The Almon representation was adopted because, at weekly frequency, adjacent lags of the same climate variable are typically highly autocorrelated, so unrestricted distributed-lag specifications are prone to multicollinearity and unstable lag coefficients. Restricting the lag profile to low-order polynomials provides a parsimonious and interpretable approximation, consistent with standard distributed-lag practice in which only smooth low-degree lag shapes are usually entertained unless the data provide strong evidence in favour of more complex profiles. At the same time, the search space was anchored to short lag horizons. This is consistent with the climate–health literature, where short-term environmental effects are typically studied over the first days or weeks after exposure. Because our outcome is weekly pharmaceutical demand rather than same-day morbidity or mortality, we allowed lag spans up to six weeks so as to cover not only short-run physiological responses, but also the delay between symptom onset, care-seeking, prescribing, and dispensing.

For scenario-conditioned projections, a second specification drops the autoregressive terms and relies solely on exogenous climatic regressors entering through the same short lag blocks, while retaining regional and week-of-year fixed effects. This purely exogenous climate-driven panel is referred to as *FE–CLIM–X*. Removing lagged outcomes prevents error propagation across horizons and makes the link between climate scenarios and projected demand more transparent. Inference for both panel models is based on a two-way wild cluster bootstrap with clustering by region and calendar week (see Section S3 for implementation details).

4 Results

The empirical results unfold in three steps. First, SLX estimates are used to quantify how contemporaneous local and neighbouring climate conditions co-move with regional respiratory pharmaceutical demand, and how these spillovers decay with distance. Next, the out-of-sample performance of the climate-augmented fixed-effects forecaster (FE–CLIM–AR) is assessed on a one-year evaluation window, focussing on national-level accuracy and error profiles. Finally, the purely exogenous variant (FE–CLIM–X) is evaluated as a scenario engine and used to generate medium-term projections of national respiratory pharmaceutical demand under alternative CMIP6 climate pathways.

4.1 Spatial spillovers

Spillover coefficients θ_C measure the association between $y_{i,t}$ and a one-unit change in the spatially lagged covariate $C^{\mathrm{SL}}_{i,t}$, i.e. the W-weighted-average exposure in other regions; 95% confidence intervals and p-values are from a two-way wild cluster bootstrap (Rademacher, $B = 9{,}999$, finite-sample correction), clustered by region and calendar week. Significance codes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$

In SLX, specific humidity is retained alongside temperature to decompose local and spillover channels of moisture independently of thermal conditions. This choice helps attribute contemporaneous effects to physically distinct pathways. Wind speed is likewise retained because, although it is not part of the final forecasting specification, it is informative for contemporaneous spatial transmission, especially through dispersion and atmospheric-mixing mechanisms that may operate across neighbouring regions within the same week.

Estimates from the SLX in Table 1 indicate that neighbouring climatic conditions exert economically and statistically meaningful spillovers on local pharmaceutical demand. Goodness-of-fit is high in levels ($R^2 = 0.822$), driven by region and week-of-year fixed effects (namely, the common annual seasonal profile) that absorb persistent cross-sectional heterogeneity and common seasonal shocks; by contrast, the within R^2 is modest (0.046), indicating that contemporaneous climatic covariates explain a limited share of within-region weekly fluctuations. Inference is based on a two-way wild cluster bootstrap with $B = 9{,}999$ resamples, clustering by region and calendar week. This procedure is robust to arbitrary serial correlation within regions and cross-sectional dependence within weeks and therefore provides size-robust p-values and confidence intervals in panels with a limited number of clusters.

Table 1 reports that most local coefficients are small and not statistically distinguishable from zero at conventional levels, when considering the traditional significance level $\alpha = 0.05$. In particular, the local effects of temperature, wind speed, specific humidity and cloud cover are imprecise; the only patterns approaching significance are a weakly negative local coefficient for temperature variability and a borderline negative coefficient for precipitation. This indicates that, once region and

Table 1 SLX panel estimates with two-way fixed effects (3): local, spillover, and total effects on weekly pharmaceutical demand

Variable	Estimate	95% CI	p
Local effects			
Temperature	−103.1	[−232.4; 33.6]	0.118
Temperature variability	−167.1	[−347.0; 15.2]	0.065$^{\cdot}$
Wind speed	44.4	[−33.2; 123.8]	0.209
Specific humidity	−41.0	[−227.6; 144.8]	0.619
Precipitation	−3.5	[−7.1; 0.0]	0.051$^{\cdot}$
Cloud cover	233.4	[−641.2; 1124.8]	0.568
Spillover effects			
Temperature	264.0	[57.2; 464.1]	0.021*
Temperature variability	−131.2	[−452.7; 194.8]	0.390
Wind speed	−419.6	[−824.9; − 18.6]	0.042*
Specific humidity	−208.9	[−531.5; 115.2]	0.170
Precipitation	−0.7	[−14.9; 13.5]	0.916
Cloud cover	401.1	[−1084.7; 1883.4]	0.567
Total effects			
Temperature	160.9	[45.5; 277.5]	0.013*
Temperature variability	−298.3	[−585.2; − 7.4]	0.046*
Wind speed	−375.2	[−739.6; − 12.2]	0.042*
Specific humidity	−249.9	[−493.2; − 2.2]	0.048*
Precipitation	−4.2	[−17.1; 8.6]	0.504
Cloud cover	634.5	[−452.2; 1710.6]	0.234

Temperature is measured in °C; wind speed in m s^{-1}; precipitation in mm/week; specific humidity in kg/kg; cloud cover is dimensionless in [0, 1]. Spillover regressors are spatially lagged covariates $C_{i,t}^{\mathrm{SL}} = \sum_{m \neq i} w_{im} C_{m,t}$

week-of-year fixed effects are absorbed, contemporaneous *within-region* climatic fluctuations explain only a modest share of the week-to-week variation in demand.

By contrast, spillover coefficients isolate substantively meaningful cross-regional transmission. Temperature exhibits a positive and significant spillover, indicating that contemporaneous warming in neighbouring regions is associated with higher local demand—consistent with spatial interdependencies operating through patient mobility, shared supply chains, and synchronous meteorological forcing typical of regional weather systems. Wind speed displays a negative and significant spillover, consistent with enhanced dispersion and atmospheric mixing that lower pollutant and allergen concentrations and, in turn, reduce respiratory symptoms locally. Spillovers for temperature variability, specific humidity, precipitation and cloud cover are not significant in this specification.

Aggregating local and spillover components yields several statistically meaningful totals. Temperature has a positive, significant and spillover-driven total effect, in line with the idea that widespread warming raises symptomatic load and care-seeking in a spatially coordinated manner. Temperature variability shows a negative and

significant total effect, consistent with transitional periods (e.g. shoulder seasons) in which fluctuations around milder averages attenuate respiratory stress relative to persistently cold conditions. Wind speed yields a negative, significant and spillover-driven total effect, aligning with the epidemiological mechanism whereby stronger winds dilute irritants and pathogens and shorten exposure durations. Specific humidity also delivers a negative and significant total effect, consistent with clinical and aerosol evidence that moister air improves mucosal hydration and reduces airborne persistence of respiratory pathogens. By contrast, total effects for precipitation and cloud cover remain not statistically different from zero. These patterns—spillovers dominating local effects for temperature (positive) and wind (negative)—are robust in sign and interpretation and remain qualitatively stable under alternative distance-weighting schemes (see Section S2 in Supplementary Material), reinforcing the conclusion that climate-driven spatial interdependencies are a systematic feature of pharmaceutical demand.

Spatial-lag coefficients quantify how climate conditions in adjacent areas co-move with local demand after controlling for own-area drivers and fixed effects. Positive spillovers indicate that warmer conditions in neighbouring territories are associated with higher local demand, while negative spillovers linked to wind speed suggest dispersion effects. In this linear-level specification, coefficients can be read as the marginal change in local demand (in packs per week) associated with a one-unit change in neighbouring climate exposures, conditional on fixed effects and local covariates. Throughout, coefficients are reported in levels, so estimates can be read as changes in weekly demand (packs) associated with a one-unit increase in the corresponding climate variable (in its physical unit), holding fixed effects and the remaining covariates constant. The dominance of the spillover over the direct component for several covariates supports a spatial view of demand formation in which mobility, logistics, and shared atmospheric forcing transmit shocks across administrative boundaries.

The single-neighbourhood SLX is reparameterized into the five distance bands defined in Sect. 3.1 by replacing W with $\{W^{(j)}\}_{j=1}^{5}$ to characterise the spatial reach of spillovers. Band-specific SLX estimates appear in Table 2, and the corresponding cumulative marginal effects are shown in Fig. 2. For improved readability, a version with panel-specific y-axis scaling is provided in Figure S2 reported in Section S2.3 of the Supplementary Material. Cloud cover shows isolated band-specific spillovers (bands 1, 3, and 5), with alternating signs and large uncertainty; the pattern is patchy rather than monotonic and does not yield a clear cumulative effect. Precipitation remains close to zero across bands, with confidence regions including zero throughout, indicating no detectable cumulative spillover. Specific humidity is negative at short ranges and moves towards zero as farther neighbours are added; intervals narrow around the origin, yielding at most borderline evidence at intermediate bands. Temperature shifts from near-zero or negative values at short range to positive once mid-distance neighbours are included; upper bands deliver a positive cumulative effect with confidence regions that approach or exclude zero. Temperature variability declines monotonically with distance: the cumulative effect becomes increasingly negative and is statistically distinct from zero for the largest bands, suggesting

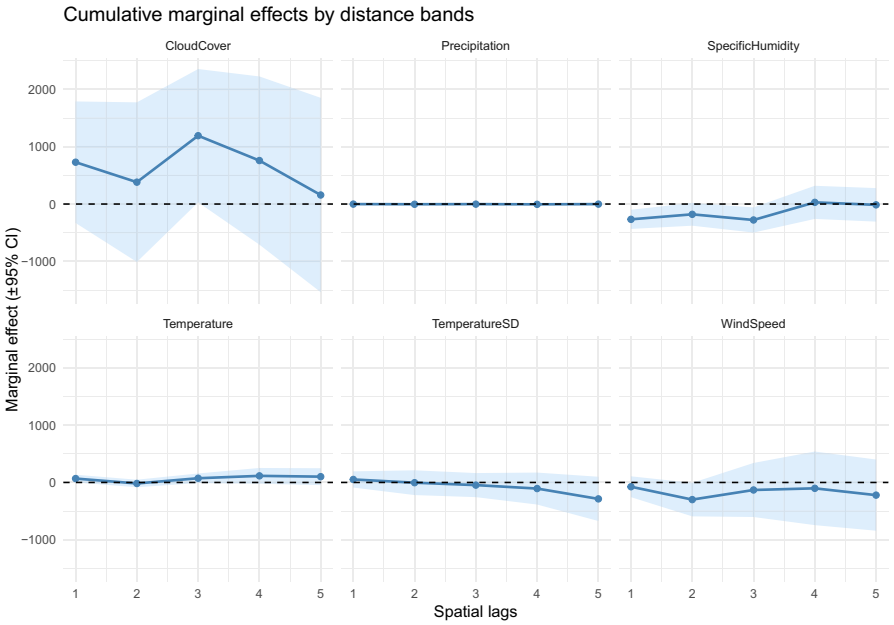

Fig. 2 Cumulative marginal effects on weekly pharmaceutical demand by distance bands for the SLX specification. Points are cumulative estimates $\hat{\theta}^{\text{cum}}_{C,k} = \sum_{j=1}^{J} \hat{\theta}_{C,j}$. Cumulative spillovers aggregate the estimated θ's across distance bands, so values below (above) zero indicate that higher exposures in surrounding regions are associated with lower (higher) demand in region i. Shaded areas are 95% confidence intervals obtained with a two-way wild cluster bootstrap (Rademacher weights, $B = 9{,}999$, finite–sample correction), clustered by region and calendar week. The horizontal dashed line marks zero. All panels are displayed on a common y-axis scale to facilitate cross-variable comparison

adverse spillovers that build with spatial reach. Wind speed shows negative spillovers concentrated at shorter ranges (the cumulative effect dips below zero by $k = 2$), partial attenuation at intermediate bands, and a still-negative cumulative impact at the outer band, with intervals that are borderline to significant depending on k. Overall, the evidence in Fig. 2 complements Table 2: negligible cumulative spillovers for precipitation and cloud cover, short-to-midrange propagation for temperature and wind, and a robust, distance-amplified negative spillover for temperature variability.

The SLX and the climate-augmented panel forecaster speak to different components of the data-generating process and therefore need not deliver identical signs for temperature. Three elements help reconcile the findings. First, the SLX identifies contemporaneous, cross-sectional co-movements: the positive, spillover-driven total effect suggests that warmer-than-usual conditions occurring simultaneously across neighbouring areas are associated with higher same-week demand, plausibly reflecting short-horizon channels (heat stress, ozone, pollen bursts, and correlated smoke risk in warm spells) that do not require propagation through infections. Second, the SLX explicitly decomposes local and neighbourhood exposures, and results indicate that the temperature signal is primarily spillover-driven. This aligns with spatially synchronous weather regimes and mobility/logistics links: when several adjacent

Table 2 Distance-band SLX panel with two-way fixed effects

	Local	SL_1	SL_2	SL_3	SL_4	SL_5
τ	−11.10	68.02***	−85.45***	91.10***	41.57	−13.86
	(37.83)	(33.45)	(25.83)	(39.10)	(65.79)	(45.19)
T_v	−79.06	52.57	−56.47	−42.13	−60.24	−180.32**
	(60.03)	(66.17)	(63.65)	(114.48)	(81.97)	(155.30)
q	−136.56***	−266.82***	87.39*	−98.76**	307.62***	−43.67
	(69.12)	(84.17)	(53.98)	(86.61)	(84.94)	(84.46)
w	−3.22	−74.53	−224.52***	167.83*	28.88	−118.56***
	(27.32)	(89.11)	(97.31)	(170.63)	(98.64)	(68.62)
r	−4.06**	−1.22	−2.50	2.01	−3.32·	3.73
	(1.55)	(4.45)	(2.82)	(5.08)	(4.17)	(5.76)
c	19.76	730.10*	−348.22	810.81*	−433.77	−601.16*
	(321.19)	(540.89)	(361.32)	(691.40)	(523.39)	(749.85)

For any regressor $C \in \{\tau, T_v, q, w, r, c\}$, where τ denotes 2 m air temperature, T_v intra-week temperature variability, q specific humidity, w wind speed, r total precipitation, and c total cloud cover, C^{SL_j} denotes the j-th distance–band spillover. Standard errors are reported in parentheses. Significance codes are based on a two-way wild cluster bootstrap (Rademacher weights, $B = 9{,}999$, finite-sample correction), clustered by region and calendar week: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$

regions warm together, the regional healthcare system faces coordinated shocks even if each region's own contemporaneous temperature has a small net effect after fixed effects. Finally, week-of-year fixed effects remove the mean seasonal cycle, so coefficients load on deviations from the intra-annual norm. In winter, positive temperature deviations tend to lower infection pressure and hence subsequent demand, whereas in the shoulder and summer periods positive deviations often co-occur with heat/ozone/pollen episodes that raise same-week demand in a spatially coordinated way. Pooling across weeks can therefore yield negative delayed responses in climate-augmented distributed-lag panels and positive contemporaneous spillovers in SLX without contradiction. These pieces are best viewed as complementary: the distributed-lag panel quantifies delayed, infection-mediated responses, while the SLX quantifies immediate, spatially propagating weather impacts.

4.2 Forecasting model performance

Model choice over the training grid yields a stable specification across expanding origins, with region fixed effects α_i absorbing persistent cross-sectional differences and week-of-year effects $\delta_{s(t)}$ capturing the common seasonal profile. The resulting dynamic climate-augmented panel forecaster, henceforth denoted *FE–CLIM–AR*, is given by

$$y_{i,t} = \alpha_i + \delta_{s(t)} + \sum_{p=1}^{4} \phi_p \, y_{i,t-p} + \theta_{\tau,0} \, Z^{(\tau)}_{0,i,t} + \theta_{T_v,0} \, Z^{(T_v)}_{0,i,t} + \theta_{\mathrm{r},0} \, Z^{(\mathrm{r})}_{0,i,t} + \theta_{\mathrm{r},1} \, Z^{(\mathrm{r})}_{1,i,t} + u_{i,t}. \tag{6}$$

Climate terms enter parsimoniously through short Almon polynomial distributed–lag blocks: the lag spans are $K_\tau = 6$, $K_{T_v} = 6$, $K_{\mathrm{r}} = 4$. These lag spans emerged from the expanding-window selection procedure described in Section S3.1, which favoured smooth low-order lag structures as the best compromise between predictive accuracy and parsimony. For temperature and temperature variability, degree–0 Almon blocks were used, meaning equal–weight sums of the first six lags:

$$Z_{0,i,t}^{(\tau)} = \sum_{\ell=1}^{K_\tau} \tau_{i,t-\ell}, \tag{7}$$

$$Z_{0,i,t}^{(T_v)} = \sum_{\ell=1}^{K_{T_v}} T_{v,\,i,t-\ell}. \tag{8}$$

For precipitation, a degree–1 Almon block was selected, i.e. level and slope over the first four lags on an *orthogonal* polynomial basis:

$$Z_{0,i,t}^{(\mathrm{r})} = \sum_{\ell=1}^{K_{\mathrm{r}}} p_{i,t-\ell}, \tag{9}$$

$$Z_{1,i,t}^{(\mathrm{r})} = \sum_{\ell=1}^{K_{\mathrm{r}}} b_1(\ell)\, p_{i,t-\ell}, \tag{10}$$

where the degree–1 orthogonal polynomial is the centred and normalised lag index

$$b_1(\ell) \;=\; \frac{\ell - \bar{\ell}}{\sqrt{\sum_{j=1}^{K_{\mathrm{r}}} (j - \bar{\ell}\,)^2}}, \qquad \bar{\ell} \;=\; \frac{K_{\mathrm{r}} + 1}{2}. \tag{11}$$

By construction $\sum_{\ell=1}^{K_{\mathrm{r}}} b_1(\ell) = 0$, which ensures orthogonality to the constant and implies that $\theta_{\mathrm{r},0}$ governs the four–week cumulative effect, whereas $\theta_{\mathrm{r},1}$ *tilts* the profile across lags without changing the total.

The selected FE–CLIM–AR specification in (6) and reported in Table 3 combines region fixed effects (α_i), week-of-year fixed effects ($\delta_{s(t)}$), an AR(4) outcome block, and short distributed-lag climate terms summarised via Almon polynomials. Identification therefore comes from *within–region, week–to–week* deviations after removing time-invariant heterogeneity and the common national seasonal profile. Weekly demand exhibits strong short-horizon inertia: all four autoregressive terms are positive and precisely estimated, confirming the material propagation of shocks over roughly one month.

The term $Z_{0,i,t}^{(\tau)}$ enters negatively and significantly. Because $Z_{0,i,t}^{(\tau)}$ is a sum, the coefficient scales linearly with exposure: a +1°C anomaly for one week is associated with an effect of about −15 units (holding other lags fixed), whereas a sustained +1°C for six consecutive weeks maps to a cumulative effect near −91.8 units. The term $Z_{0,i,t}^{(T_v)}$ is negative with borderline evidence. Point estimates suggest that more

Table 3 Two-way FE–OLS coefficients with 95% CIs. Wild cluster bootstrap p-values. Significance codes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$

Regressor	Estimate [95% CI]	p-value
y_{t-1}	0.242 [0.113; 0.374]	0.0088^{**}
y_{t-2}	0.123 [0.007; 0.240]	0.0402^{*}
y_{t-3}	0.152 [0.068; 0.236]	0.0076^{**}
y_{t-4}	0.204 [0.087; 0.322]	0.0078^{**}
$Z^{(\tau)}_{0,i,t}$	−15.305 [−27.971; − 2.757]	0.0240^{*}
$Z^{(T_\nu)}_{0,i,t}$	−53.874 [−120.520; 11.908]	$0.0998^{\cdot}$
$Z^{(\mathrm{r})}_{0,i,t}$	−0.436 [−2.364; 1.496]	0.6248
$Z^{(\mathrm{r})}_{1,i,t}$	5.698 [0.339; 11.465]	0.0386^{*}

Two-way fixed effects for region and week-of-year. CIs and p-values use a two-way wild cluster bootstrap (Rademacher weights, B=9,999; clustering by region and by calendar week)

stable thermal conditions are associated with slightly lower demand, though statistical support is weaker than for mean temperature. Precipitation enters through a level and a slope synthesised on an *orthogonal* polynomial basis. Empirically, the level term is small and not significant, while the slope is positive and significant. Two implications follow: (i) the near-zero level indicates a small net four-week cumulative effect; (ii) the positive slope *tilts* the lag profile—early lags contribute less and later lags more—yielding a *back-loaded* association that is modest at very short delays and larger a few weeks out. Because $Z^{(\mathrm{r})}_{1,i,t}$ is orthogonal to $Z^{(\mathrm{r})}_{0,i,t}$, the slope reallocates weight across lags without altering the total.

In conclusion, after absorbing regional heterogeneity and the common seasonal cycle, weekly demand is dominated by short-run persistence; higher recent temperatures are associated with lower near-term demand; thermal variability points in the same direction with weaker evidence; and precipitation displays *timing* rather than bulk effects, with a back-loaded lag profile and near-zero cumulative impact over four weeks. All coefficients are *within–region, week–to–week* associations rather than causal effects, and inference is robust to serial correlation and common shocks under the two-way wild cluster bootstrap.

Aggregating dynamic predictions to the national level yields $MAPE = 11.16\%$, $RMSE = 37{,}351$, $RSR = 0.748$, and $WAPE = 12.24\%$. The error profile combines frequent under-prediction ($UR = 0.75$) with a negative mean error ($ME < 0$), i.e. many small under-predictions punctuated by fewer, larger over-predictions that dominate the average. Formal Clark–West tests of the incremental predictive value of climate covariates relative to a purely autoregressive benchmark are reported in Section S3.3 of the Supplementary Material.

See Section S4 in Supplementary Material for residual diagnostics—serial and cross–sectional dependence and heteroskedasticity—and for coefficient stability based on a leave–one–region–out analysis.

4.3 Future projections

Climatic regressors enter as exogenous trajectories aligned with scenarios: the weekly dynamics are anchored by the estimated seasonal structure, and the accumulation of forecast error is limited to deviations of the exogenous paths from the scenario.

Projections are based on the FE–CLIM–X specification, a two–way fixed-effects model with short Almon blocks that summarise delayed climatic effects:

$$y_{i,t} = \alpha_i + \delta_{s(t)} + \theta_{\tau,0} Z^{(\tau)}_{0,i,t} + \theta_{T_v,0} Z^{(T_v)}_{0,i,t} + \theta_{\mathrm{r},0} Z^{(\mathrm{r})}_{0,i,t} + \theta_{\mathrm{r},1} Z^{(\mathrm{r})}_{1,i,t} + u_{i,t}. \tag{12}$$

On the 52-week hold-out the CLIM–X model attains MAPE=11.14%, RMSE=37,825, and RSR=0.757, in line with the panel that includes autoregression. The absence of output lags eliminates feedback of errors across horizons, preserves coherent aggregation from regional to national series, and maintains a transparent link between climatic scenarios and predicted demand. With comparable accuracy and reduced channels for error accumulation, the FE–OLS specification without autoregression offers a robust platform for medium-term projections. Figure 3 reports observed and model-based forecasts at the national level from FE–CLIM–X

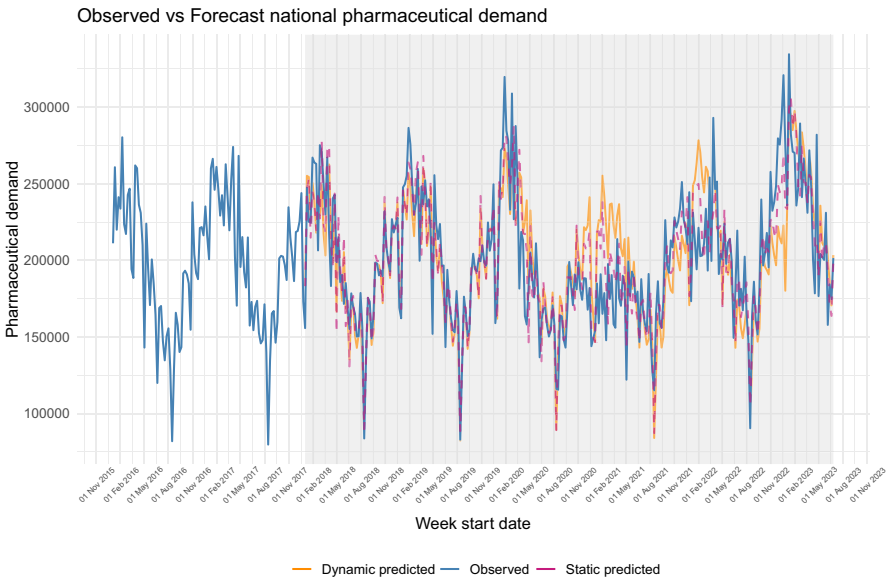

Fig. 3 Observed vs. one-year-ahead forecasts at the national level from the FE–CLIM–AR specification under an *expanding-window* out-of-sample protocol. The initial training sample covers two years and is then expanded by one year at each fold; at every re-estimation, the model produces a 52-week-ahead forecast path for the subsequent year, conditioning on realised weekly climate covariates. The *static* forecasts treat each week in the evaluation year as a separate one-step-ahead problem and are obtained by plugging the *observed* variables in all lagged terms. The *dynamic* forecasts instead generate a genuine 52-step path: starting from the last observed lags at the forecast origin, the model is iterated forward and, for all subsequent weeks, the lagged y terms are filled with past predictions, so forecast errors propagate recursively through the autoregressive structure

under an expanding-window evaluation. Additional expanding-window performance metrics under alternative re-estimation cadences (weekly and annual retraining) are summarised in Section S3.4 of the Supplementary Material. The initial training sample covers two years; at each step, the model is re-estimated on all data available so far and asked to predict the following 52 weeks. The fit is tight overall, with only marginal deviations; the largest gaps occur during the pandemic phase, when non-pharmaceutical interventions flattened demand and dampened the seasonal amplitude.

Forecasts are then generated by re-feeding the model with CMIP6 weekly drivers aligned to ISO weeks under SSP2-4.5 (intermediate forcing) and SSP5-8.5 (very high forcing), building the same short Almon polynomials on the future covariates using the historical tail to initialise lags. Predictive uncertainty is quantified via a moving-block bootstrap of national residuals, which preserves temporal dependence in shocks. The bootstrap procedure and quantitative summaries of scenario-conditioned prediction intervals are reported in Section S3.5 of the Supplementary Material. Point forecasts are reported as medians across draws and 95% predictive intervals are given by the empirical 2.5th and 97.5th percentiles at each horizon.

Figure 4 shows that seasonality is sustained by the week-of-year fixed effects, yielding the familiar annual oscillation and uncertainty widens around seasonal peaks. Both scenarios imply a higher mean level relative to the historical tail, reflecting the combination of the estimated seasonal structure and the projected climate covariates. SSP5-8.5 tends to sit modestly above SSP2-4.5 at demand peaks, while troughs are often similar, reflecting warmer/wetter trajectories under the higher-forcing pathway. The segment for the last semester of observed 2023 lies largely within the predictive bands, indicating well-calibrated uncertainty at the one-year horizon, and intervals widen gradually with lead time as shocks accumulate. Overall, the

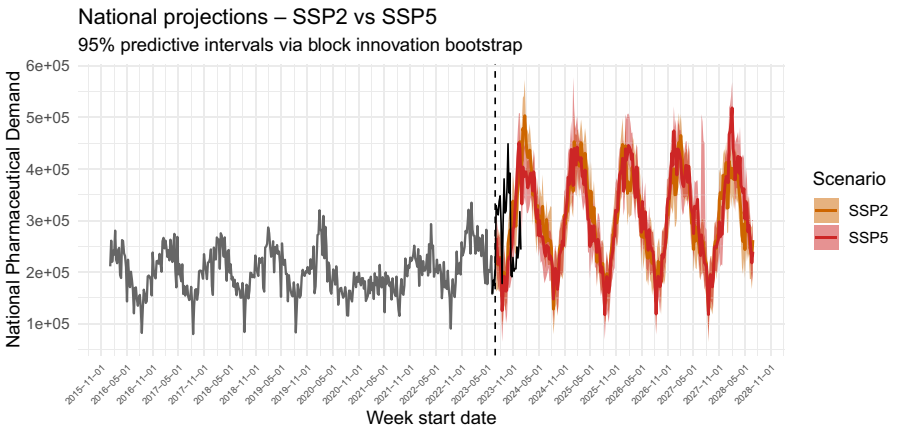

Fig. 4 National projections from the CLIM-X model in (12) driven by CMIP6 weekly covariates under SSP2-4.5 and SSP5-8.5. Grey shows the historical national series, while black shows the June–December 2023 representing the last observed period for the time series; the dashed line marks the end of the estimation sample. Coloured lines are median scenario–conditioned forecasts; shaded envelopes are 95% predictive intervals obtained via a moving–block bootstrap of national innovations

pipeline delivers scenario-conditioned national projections from FE–CLIM–X with time-dependent, serial–correlation–aware uncertainty that is consistent with the model's out-of-sample accuracy profile. By the mid-2028 horizon, national demand is projected to lie roughly 45–55% above pre-pandemic levels under SSP2-4.5 and SSP5-8.5, respectively, with somewhat smaller yet still substantial uplifts (around 25–35%) when using a post-pandemic 2022–2023 baseline. SSP5-8.5 consistently exceeds SSP2-4.5, and the uplift is driven primarily by higher seasonal peaks rather than uniformly elevated troughs (full numeric summaries and predictive intervals are reported in Table S6 in Section S3.5 of the Supplementary Material).

Taken together, these results are robust to the baseline choice: SSP5-8.5 consistently exceeds SSP2-4.5, and the uplift is driven primarily by higher seasonal peaks rather than uniformly elevated troughs. Operationally, this points to earlier procurement, scenario-aware capacity planning, and tighter peak-season buffers; even under the milder pathway, mid-2028 volumes materially exceed historical norms.

These projections should be interpreted as climate-conditioned extensions of the recent demand regime, rather than point predictions of future utilisation. They are intended to inform stress-testing and planning under plausible climate futures, not to replace context-specific judgement by health authorities. More specifically, the projection framework is based on a validated panel-data specification designed to recover short-run associations from historical data. Accordingly, the medium-run extrapolations reported here should not be interpreted causally, but rather as scenario-conditioned tools to support informed health-policy planning.

A supplementary robustness exercise augments the historical forecasting specification with common pandemic-regime dummies identified from the main structural breaks in the national aggregate. Although these controls improve historical fit, they are not retained in the main projection framework because they encode exceptional non-climatic shocks whose future path is not part of the climate-scenario exercise (see Section S5 of the Supplementary Material).

5 Conclusions

Climate change acts as a pervasive stressor on health systems, with respiratory conditions particularly exposed because the airways constitute the primary interface with ambient conditions. In Greece, recent warming, more frequent heatwaves, altered precipitation regimes, and a higher incidence of fire-weather days combine with heterogeneous geography and health-system constraints, generating a demanding environment for planning respiratory pharmaceutical supply. Using a weekly panel for 20 pharmaceutical regions over 2016–2023, this study quantifies how climate variability and its spatial propagation shape respiratory pharmaceutical demand and develops a panel-based forecasting framework for medium-term projections under alternative climate scenarios.

Relative to previous work for Greece—in particular Schisa and Farnè (2025), which analysed climate-sensitive respiratory demand using a national aggregate series and compared alternative forecasting algorithms—the present study advances the evidence base in three directions. First, it adopts a fully spatially explicit panel

design, leveraging the 20-region structure of the retail data and quantifying distance-dependent spillovers through SLX specifications with two-way fixed effects. Second, it recasts climate-sensitive forecasting in a parsimonious panel framework (FE–CLIM–AR and FE–CLIM–X) that pools information across regions, summarises delayed climatic effects via short distributed-lag polynomials, and remains transparent enough to support structural interpretation. Third, by embedding the exogenous climate-driven panel forecaster into a CMIP6-based scenario pipeline, it moves from "forecast comparison under historical climate" to explicit scenario-conditioned projections of pharmaceutical demand under alternative warming pathways. Taken together, the two studies are complementary: Schisa and Farnè (2025) map the climatic signal in national-level demand and compare flexible forecasting tools, whereas the present work embeds that signal in a spatially resolved panel structure and projects it forward under explicit climate scenarios.

The analysis highlights a demand process characterised by pronounced winter-dominant seasonality, strong short-run persistence, and non-negligible structural breaks around the COVID-19 period. Within this backdrop, the spatial econometric results show that contemporaneous climate signals do carry predictive content, but mostly through spatially propagated channels rather than purely local shocks. Two-way fixed-effects SLX specifications indicate that once regional heterogeneity and the common seasonal profile are absorbed, contemporaneous within-region fluctuations in climate explain only a modest share of week-to-week variation. By contrast, neighbouring conditions—especially temperature and wind speed, with additional contributions from specific humidity and temperature variability—exert economically meaningful spillovers. Distance-band decompositions reveal that these spillovers have a finite geographical reach, with temperature and wind effects concentrated at short to mid ranges and negligible cumulative impacts for precipitation and cloud cover. These results support a view of respiratory pharmaceutical demand as shaped by synoptic-scale regimes, mobility patterns, and integrated supply chains rather than isolated local exposures.

On the temporal side, the two-way fixed-effects forecaster with short Almon blocks on temperature, temperature variability, and precipitation parsimoniously summarises delayed climate effects while exploiting cross-sectional and week-of-year structure. In its predictive version (FE–CLIM–AR), which combines autoregressive dynamics with climate polynomials, the model attains out-of-sample accuracy on the annual hold-out that is comparable to more flexible benchmarks when aggregated to the national level. A climate-only variant without autoregressive terms (FE–CLIM–X) delivers similar one-year-ahead performance while eliminating mechanical error propagation across horizons and maintaining a transparent mapping between exogenous climate trajectories and predicted demand, making it well suited as a scenario engine.

Embedding FE–CLIM–X in a CMIP6-driven projection pipeline yields medium-term projections of national respiratory pharmaceutical demand under SSP2-4.5 and SSP5-8.5. Scenario-conditioned covariates preserve the familiar seasonal oscillation through the week-of-year effects, while warming and associated changes in variability translate into higher seasonal peaks and elevated annual means. Although uncertainty bands widen with lead time and remain wider than the gap between the

two scenarios over much of the horizon, both pathways imply substantial increases in average demand by the mid-2020s relative to historical baselines, with SSP5-8.5 consistently sitting above SSP2-4.5 and differences most pronounced around peak weeks. These findings suggest that climate-driven pressure on respiratory pharmaceutical demand is not a marginal adjustment but a structural shift, and that neglecting climate information risks systematic under-provision of respiratory pharmaceuticals during high-risk seasons. These scenario-conditioned trajectories are not intended as exact point predictions of future utilisation, but as climate-informed extensions of the recent demand regime that can support stress-testing and medium-term planning by health authorities.

From an operational perspective, the results argue for forecasting pipelines that are both spatially explicit and climate-aware. National aggregates can mask pockets of vulnerability, and procurement rules based on historical averages risk being systematically too low in regions that face both higher baseline demand and stronger climatic amplification. Scenario-conditioned projections can be used as stress tests for the pharmaceutical supply chain, helping ministries of health, national insurance funds, and regional procurement agencies to size seasonal buffers, prioritise cold-season capacity in high-risk regions, and evaluate how robust existing procurement contracts are to warmer and more variable winters. Because the underlying tools rely on routinely collected retail data and publicly available climate information, the approach is transferable to other therapeutic classes and to other Mediterranean or climate-vulnerable health systems, offering a practical template for climate-informed pharmaceutical planning.

These findings have direct implications for health and economic policy. The spatial evidence indicates that respiratory pharmaceutical demand is shaped not only by local climate conditions, but also by coordinated signals across neighbouring regions. Accordingly, procurement should reflect the regional mosaic of climate signals rather than single-area indicators or national aggregates alone. This is especially relevant in fiscally constrained health systems, where geographically uneven shocks can translate into local shortages, inefficient stock allocation, and avoidable pressure on procurement budgets (Ranney et al. 2020; Toković 2023).

Our scenario-based projections further reinforce this message. Although uncertainty bands widen with the forecast horizon, the differences between scenarios become economically meaningful around seasonal peaks, where the upper tail under SSP5-8.5 implies higher capacity and inventory needs. In practical terms, the higher-forcing pathway implies more pronounced peak-season surges in respiratory pharmaceutical demand, so that the scenario gap becomes most consequential precisely in the weeks when health systems and supply chains are already under the greatest seasonal strain. This, in turn, places greater pressure on procurement timing, inventory requirements, and the short-run logistical capacity needed to prevent shortages during peak periods. Operationally, these magnitudes argue for earlier procurement, scenario-aware capacity planning, and tighter peak-season buffers to maintain service levels in a warming climate. Taken together, these considerations

indicate that integrating climate information into pharmaceutical-demand forecasting can improve preparedness not only by strengthening predictive performance, but also by supporting more timely procurement, more efficient stock allocation, and more resilient supply-chain management under growing climatic stress.

Several limitations suggest avenues for further work. The analysis relies on seven and a half years of weekly data, which are well suited to characterising seasonal dynamics and medium-term variability but too short to extract trends on decadal climate scales. Extending the historical window, where data permit, would strengthen the attribution of slow-moving changes to climate rather than to policy, demographic, or reimbursement shocks. The focus on aggregated respiratory prescription demand at ATC Level 1 trades therapeutic detail for robustness; future studies could complement this perspective with molecule- or class-specific analyses to distinguish climate-sensitive segments from those driven primarily by other factors. At the climate side, incorporating additional indicators (e.g. air quality, pollen indices, or fire activity) and examining alternative downscaling and bias-adjustment schemes for CMIP6 would refine the link between large-scale forcing and local exposures. Furthermore, a full bias-adjustment or delta-method downscaling is left for future work. This would refine absolute levels; the scenario comparisons should therefore be interpreted primarily in terms of relative changes and timing rather than precise levels. Importantly, our scenario-conditioned projections cannot be interpreted as long-run impacts, as they mainly come from a validated identification and disentanglement of short-run associations from historical data. Nevertheless, they can still be useful for catching potential shift from current seasonal patterns, able to help design future health policies in effective and informed way.

Overall, the evidence assembled here indicates that climate—and its ongoing change—is a material and spatially structured driver of respiratory pharmaceutical demand in Greece. Two-way fixed-effects panel models, equipped with climate-sensitive distributed lags and spatial spillovers, offer an interpretable and operationally credible framework for quantifying these links and for generating scenario-conditioned projections. Embedding such tools in routine forecasting practice can support earlier warnings, better targeted stocks, and more resilient pharmaceutical supply in a warming climate.

6 Supplementary Information

Supplementary Material provides additional technical details and robustness analyses that complement the main text. Section S1 documents the construction of the weekly spatio-temporal panel, including the definition of the twenty pharmaceutical regions, the choice of reference coordinates, and the processing of ERA5 reanalysis variables and CMIP6 climate projections. Section S2 describes the spatial weighting schemes used in the SLX panel model and reports robustness checks based on alternative distance kernels, k-nearest-neighbour matrices, and distance-band

decompositions of climatic spillover effects. Section S3 presents the specification grid and expanding-window validation strategy for the two-way fixed-effects forecasting panels, together with the polynomial distributed-lag representation of climate covariates, formal Clark–West tests for the incremental predictive value of climate information, and additional expanding-window forecast metrics. Section S4 reports diagnostic checks for the FE–CLIM–AR model, including serial correlation and cross-sectional dependence. Section S5 reports a robustness check for pandemic-related regime shifts in the forecasting model, based on common COVID-period dummies and their effects on predictive performance and coefficient stability.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s10651-026-00731-8.

Acknowledgements The authors sincerely thank Dr. Filippo Benetti and Dr. Rosa Falotico of Alira Health Srl for their support in the acquisition and interpretation of the pharmaceutical data.

Author contributions V.S. conceived the study, curated the data, performed the analyses, prepared the figures, and wrote the main manuscript text. M.F. contributed to the study design and methodology, supervised the work, and revised the manuscript. All authors reviewed and approved the final manuscript.

Funding Open access funding provided by Alma Mater Studiorum - Università di Bologna within the CRUI-CARE Agreement. Viviana Schisa acknowledges financial support from Alira Health Srl and from the European Union – NextGenerationEU, under the Italian National Recovery and Resilience Plan (PNRR), Mission 4, Component 2, Investment 3.3, as part of the research fellowship Ex D.M. 352/2022.

Data availability A pharmaceutical proprietary dataset, kindly provided by A pharmaceutical proprietary dataset, kindly provided byAlira Health Srl, was analysed during the current study.

Declarations

Conflict of interest The authors declare no conflict of interest.

Ethical approval This study is based on aggregated and anonymised pharmaceutical dispensing data and publicly available meteorological data. It does not involve research with human participants or animals; therefore, ethics approval was not required.

Consent to participate Not applicable.

Consent for publication Not applicable.

References

Abir M, Vardavas R, Tariq ZH, Hoch E, Lawson E, Cortner S (2025) Impact of climate change on health and drug demand. RAND Health Q 12(2):13

Baltagi BH (2005) Econometric Analysis of Panel Data, 3rd edn. John Wiley & Sons, Chichester

Bergé L (2018) Efficient estimation of maximum likelihood models with multiple fixed-effects: the r package fenmlm. Technical report, Department of Economics at the University of Luxembourg

Barros V, Field CB, Dokken DJ, Mastrandrea MD, Mach KJ, Bilir TE, Chatterjee M, Ebi KL, Estrada YO, Genova RC, Girma B, Kissel ES, Levy AN, MacCracken S, Mastrandrea PR, White LL (eds) (2014) Climate change 2014: impacts, adaptation, and vulnerability. Part B: regional aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA

Cianconi P, Betrò S, Janiri L (2020) The impact of climate change on mental health: a systematic descriptive review. Front Psychol 11:490206

Cameron AC, Gelbach JB, Miller DL (2008) Bootstrap-based improvements for inference with clustered errors. Rev Econ Stat 90(3):414–427

Carleton T, Jina A, Delgado M, Greenstone M, Houser T, Hsiang S, Hultgren A, Kopp RE, McCusker KE, Nath I et al (2022) Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. Q J Econ 137(4):2037–2105

Copernicus Climate Change Service: CMIP6 climate projections (dataset). Copernicus Climate Data Store (CDS). Available from the Climate Data Store: https://cds.climate.copernicus.eu/datasets/projections-cmip6 (2021)

Deloitte Business Solutions S.A. (2025) The way forward: A roadmap for greece's pharmaceutical policy. Report, Deloitte Business Solutions S.A. Prepared on behalf of SFEE (National Pharmaceutical Industry Association). https://www.sfee.gr/wp-content/uploads/2025/12/SFEE-Deloitte-A-Roadmap-Towards-Greeces-Pharma-Policy_Full-Report.pdf Accessed 2026-01-24

Economou C, Kaitelidou D, Karanikolos M, Maresso A (2017) Health system review. Health. 19(5)

Hersbach H, Bell B, Berrisford P, Hirahara S, Horányi A, Muñoz-Sabater J, Nicolas J, Peubey C, Radu R, Schepers D et al (2020) The era5 global reanalysis. Q J R Meteorol Soc 146(730):1999–2049

Haines A, Kovats RS, Campbell-Lendrum D, Corvalán C (2006) Climate change and human health: impacts, vulnerability, and mitigation. Lancet 367(9528):2101–2109

Halleck Vega S, Elhorst JP (2015) The slx model. J Reg Sci 55(3):339–363

Khraishah H, Alahmad B, Ostergard RL Jr, AlAshqar A, Albaghdadi M, Vellanki N, Chowdhury MM, Al-Kindi SG, Zanobetti A, Gasparrini A et al (2022) Climate change and cardiovascular disease: implications for global health. Nat Rev Cardiol 19(12):798–812

Kotz M, Levermann A, Wenz L (2024) The economic commitment of climate change. Nature 628(8008):551–557

MacKinnon JG, Nielsen MØ, Webb MD (2023) Fast and reliable jackknife and bootstrap methods for cluster-robust inference. J Appl Economet 38(5):671–694

Rocklöv J, Dubrow R (2020) Climate change: an enduring challenge for vector-borne disease prevention and control. Nat Immunol 21(5):479–483

Ranney ML, Griffeth V, Jha AK (2020) Critical supply shortages–the need for ventilators and personal protective equipment during the COVID-19 pandemic. N Engl J Med 382(18):41

Roodman D, Nielsen MØ, MacKinnon JG, Webb MD (2019) Fast and wild: bootstrap inference in Stata using boottest. Stand Genomic Sci 19(1):4–60

Redshaw CH, Stahl-Timmins WM, Fleming LE, Davidson I, Depledge MH (2013) Potential changes in disease patterns and pharmaceutical use in response to climate change. J Toxicol Environ Health B Crit Rev 16(5):285–320

Schisa V, Farnè M (2025) The impact of climatic factors on respiratory pharmaceutical demand: a comparison of forecasting models for Greece. Environmetrics 36(7):70041

Toković K (2023) Impact of disruptions in the pharmaceutical industry–how to build more resilient sales teams. In: International Scientific Conference on Economy, Management and Information Technologies, vol. 1, pp. 213–218

Tolika C, Zanis P, Anagnostopoulou C (2012) Regional climate change scenarios for greece: future temperature and precipitation projections from ensembles of rcms.

Zerefos C, Repapis C, Giannakopoulos C, Kapsomenakis J, Papanikolaou D, Papanikolaou M, Poulos S, Vrekoussis M, Philandras C, Tselioudis G, et al. (2011) The climate of the eastern mediterranean and greece: Past, present and future. The environmental, economic and social impacts of climate change in Greece, 1–126

Authors and Affiliations

Viviana Schisa[1] · Matteo Farnè[1]

✉ Viviana Schisa
viviana.schisa2@unibo.it

Matteo Farnè
matteo.farne@unibo.it

[1] Department of Statistical Sciences, University of Bologna, Via delle Belle Arti, 41, 40126 Bologna, Italy