The Economic Impact of Inaccurate Rain Forecasts: Evidence from U.S. Businesses

Tim Sels
UC Berkeley

Weilong Wang
St. John's University

Jacob Snyder
NBER

Vaibhav Anand
St. John's University

Abhishek Nagaraj
*UC Berkeley**

March 26, 2026

(Working Draft– Please do not cite or circulate)

Weather forecasts shape routine economic decisions, yet little is known about the economic consequences of errors in everyday forecasts. This paper examines the impact of one-day-ahead rainfall forecast errors on consumer spending. Using zip-code–level data on daily precipitation forecasts matched to consumer transactions, we show that forecast errors generate economically meaningful distortions in spending. When rainfall is overpredicted, same-day consumer spending falls by about 3 percent relative to days with accurate forecasts. These declines persist for at least seven days, indicating foregone economic activity rather than intertemporal substitution. By contrast, underpredicted rainfall produces small and short-lived increases in spending that do not persist over time. Our estimates imply that inaccurate rainfall forecasts reduce annual consumer spending by approximately 1.7 percent. Aggregated nationally, this corresponds to roughly $123 billion in lost consumer spending each year. The results demonstrate that routine forecast errors have large and persistent effects on real economic activity, implying that standard measures of forecast accuracy miss important economic consequences.

1 Introduction

Weather forecasts play an important role in routine decision-making by households and firms. Forecast information shapes a wide range of everyday economic decisions, from consumers' time use and consumption choices to firms' production and investment planning. When forecasts are inaccurate, these decisions may reflect misinformed adjustments based on incorrect expectations about realized weather conditions. Despite substantial investment in improving forecast quality, the evaluation of weather forecasts remains largely focused on technical measures of predictive skill, with limited attention to the economic consequences of forecast errors. A growing literature examines the social and economic impacts of forecast accuracy in the context of seasonal weather or extreme weather events (Downey, Lind and Shrader 2023, Shrader 2023, Molina and Rudik 2022). However, much less is known about the economic implications of errors in routine forecasts. Yet routine forecasts, such as daily precipitation forecasts, constitute the core output of modern weather forecasting and shape a wide range of everyday economic decisions (Shrader, Bakkensen and Lemoine 2023, Anand 2025). Understanding how errors in these forecasts influence economic behavior is therefore essential for evaluating the economic relevance of forecast accuracy.

*All errors are our own. This research is supported by the National Oceanic and Atmospheric Administration (NOAA) under award number NOAA-NFA-NFAPO-2021-2006626. The authors have no financial relationships or other potential conflicts of interest relevant to this work.

This paper studies how routine weather forecast errors influence consumer spending, focusing on one-day-ahead rainfall forecasts. Consumer spending is a central component of local economic activity, and even small changes in daily spending can aggregate into meaningful economic effects. Rainfall is a common but salient weather event that shapes everyday decisions, including shopping, commuting, and leisure activities, making it a natural setting for studying the economic implications of routine forecast errors. Households often adjust spending plans in response to expected weather conditions (Roth Tran 2023). When rainfall forecasts are inaccurate, these plans may reflect misinformed adjustments that do not align with realized conditions, for example, reducing shopping trips when rain is overpredicted or proceeding with planned activities when rain is underpredicted. Using high-frequency daily data on consumer spending, realized precipitation, and rainfall forecasts, this paper quantifies how forecast errors affect consumer spending behavior.

Using zip-code–level data on daily weather and consumer spending for over one million U.S. businesses from 2022–2024, we estimate how rainfall forecast errors affect consumer spending. Specifically, we examine how overprediction and underprediction of one day ahead rainfall shift spending patterns and quantify the associated economic impacts. Our empirical strategy exploits naturally occurring variation in the accuracy of one-day-ahead precipitation forecasts while holding realized weather constant. We compare daily consumer spending across instances in which zip codes experience similar weather conditions but differ in forecast accuracy, using anonymized transaction data from SafeGraph covering a wide range of industries. By including date, industry, and location-by-month-by-year fixed effects and controlling for realized weather, we ensure that variation in forecast error reflects differences in forecast quality rather than weather severity, local seasonality, or underlying economic conditions.

Our results show that forecast errors meaningfully affect consumer spending, with large and persistent effects for overprediction: when forecasts exceed actual rainfall by more than 10 mm, same-day spending falls by about 3% relative to days when forecasts are within 3 mm of realized rain. Moderate overpredictions (3–10 mm) reduce same-day spending by 1.2%. These effects persist over time: in the seven days following an overprediction event, spending remains lower by roughly 4.5% and 1.75% for large and moderate errors, respectively. The size of these declines is comparable to the direct effect of rainfall itself, as light to heavy rain reduces same-day spending by 1.4% to 5.1% in our sample.

Underprediction errors, in contrast, have smaller and opposite effects. When forecasts underpredict rainfall by more than 10 mm, same-day spending increases by about 1% relative to accurate forecasts for similar rainfall levels. Moderate underpredictions of 3–10 mm raise spending by about 1%. Some of these positive effects diminish over time: over the following week, large underpredictions are associated with roughly a 1% increase in spending, while moderate underpredictions have no meaningful or statistically significant longer-run impact.

We find that the effect of forecast errors on spending is broadly consistent across industries; however the magnitude of this effect varies. Sectors that rely on discretionary trips and in-person services are likely most exposed. We find that large overpredictions of rain reduce same-day spending for services, wholesale trades, and transportation by about 4–6% , compared with declines of roughly 3% in retail and arts and entertainment, and about 2% in hospitality and healthcare. Moderate overpredictions show a similar pattern, with effect size varying from 3% in services to nearly no effect on education spending. Underprediction errors, by contrast, generate smaller spending increases of around 1–2% across sectors, with weak response in services. These results suggest that sectors dependent on discretionary household trips are likely most sensitive to forecast errors.

Our preliminary calculations show that rainfall forecast errors reduce daily consumer spending by about

$1.57 billion in our sample, which translates to a decrease of nearly 1.7% in total spending in our sample.[1] Scaled to the national level, this corresponds to an annual reduction of roughly $123 billion in consumer spending.[2] These findings indicate that even modest improvements in short-term precipitation prediction can yield measurable gains in consumer activity and local business performance.[3]

This paper contributes to the literature on the economic value of weather forecasts in decision-making. Specifically, we examine the economic costs of errors in routine weather forecasts in the context of consumer spending and rainfall. Existing literature has primarily focused on the social and economic value of forecasts for extreme weather events, such as hurricanes (Molina and Rudik 2022), or on long-horizon and seasonal forecasts, particularly in the context of production decisions by firms and farmers (Shrader 2023, Downey, Lind and Shrader 2023, Burlig et al. 2024). In contrast, there is limited evidence on the economic consequences of errors in routine, short-term forecasts, which constitute the core output of modern forecasting systems and are most frequently used by households and businesses in everyday decision-making. We provide evidence that errors in one-day-ahead rainfall forecasts meaningfully affect consumer spending. Overprediction and underprediction errors have asymmetric effects, with overprediction leading to large and persistent reductions in spending. These results suggest that forecast errors generate systematic distortions of household activity, rather than merely shifting consumption across days.

Our paper is related to work studying the benefits of routine weather forecasts for adaptation: Shrader, Bakkensen and Lemoine (2023) examine the effect of daily temperature forecasts on mortality, Anand (2025) examines the effect of winter weather warnings on vehicle crashes, and Song (2024) examines the effect of daily temperature forecasts on labor participation. We differ from this literature by focusing on consumer spending as a key economic margin through which routine forecast errors affect local economic activity. Using high-frequency transaction data, we show that forecast errors distort everyday consumption decisions even when realized weather is held constant, highlighting the economic costs of misinformation rather than weather itself. By using consumer spending as a context, our findings demonstrate that routine forecast errors can generate meaningful and persistent economic losses. These results suggest that evaluating weather forecasts solely on technical accuracy may overlook economically relevant costs, with implications for how short-term forecasts are assessed and prioritized for improvement.

There are some limitations to our study. We focus on quantifying the economic cost of forecast errors through reduced consumer spending, but do not estimate the associated benefits of acting on forecasts. Overpredictions may reduce traffic or safety risks as well, but our data do not allow us to measure these benefits. While prior work has focused on those benefits, our aim is to document the cost side of the analysis. Second, we measure lost point-of-sale spending but do not account for substitution to other channels, such as online purchases. Such substitution is likely in the retail sector; however, we find consistent effects across sectors, and even stronger effects for services and transportation, where substitution is less relevant. A related limitation is that some observed declines may reflect changes in business operations—such as reduced hours—rather than consumer demand alone. Our design cannot separate demand responses from supply-side adjustments. However, we believe both channels contribute to the economic costs we aim to measure. Third, we quantify costs in terms of dollar spending rather than welfare and therefore do not translate these effects into utility losses. These dollar-based estimates nonetheless provide a meaningful and

[1] This includes a decrease of $2.44 billion from overprediction errors and an increase of $0.87 billion from underprediction errors. Mean annual spending in our sample is $93 billion.

[2] Annual consumer spending in the United States is approximately $18 trillion as of 2025, of which nearly 40% is discretionary and likely to be affected by forecast errors. See section 4 for details.

[3] In comparison, prior studies estimate that the annual benefits of accurate precipitation forecasts in the United States range from $250 million to $1.5 billion (Farkas et al. 2025).

policy-relevant measure of the economic activity directly affected by forecast errors. Finally, our estimates reflect the costs of forecast errors under current forecasting technology. We do not offer recommendations on the optimal accuracy target for the National Weather Service; instead, we provide estimates of the economic costs that can inform assessments of forecast improvements. Our paper provides the first set of forecast error cost estimates using large-scale data at the national level, and future research can extend this work by incorporating additional behavioral channels and cost-benefit measures.

The paper is structured as follows: Section 2 describes our data and sample construction and section 3 discusses the empirical design. Section 4 presents our main results and section 5 concludes.

2 Data

We construct a comprehensive dataset that links high-frequency weather conditions, weather forecasts, and consumer economic activity across the continental United States. Our unit of observation is the ZIP code $\times$ calendar day $\times$ industry. The weather data provide realized precipitation and temperature at fine spatial resolution, allowing us to hold actual conditions constant when assessing the economic effects of forecast errors. Forecast data provide the predicted values against which we measure those errors. Consumer spending and foot-traffic datasets offer high-frequency measures of economic behavior at the establishment level, which we aggregate to the ZIP–day–industry level to match the weather data. We additionally merge ZIP-code-level demographic and socioeconomic characteristics from the U.S. Census to capture heterogeneity in local responses. Together, these sources form a panel that enables comparisons across locations experiencing the same realized weather but different forecast accuracy, forming the basis of our empirical strategy.

2.1 Actual and Forecasted Weather

We measure realized daily precipitation (in millimeters) using NCEP Stage IV data from the National Oceanic and Atmospheric Administration (NOAA), the U.S. federal agency responsible for monitoring and forecasting weather, climate, and environmental conditions. The NCEP data combine radar- and gauge-based measurements to generate high-resolution precipitation estimates on a 4×4 km resolution across the continental United States.

Because our outcome—economic consumer activity—is driven by people rather than land area, we weight precipitation by local population density so that more densely populated areas receive greater weight within each ZIP. This approach prevents large unpopulated regions from disproportionately influencing ZIP-level weather measures despite having limited economic relevance. In practice, we use 1×1 km census population-density data and overlay these grids with the 4×4 km precipitation rasters, calculating an area-weighted precipitation value for each 1 km cell based on the exact spatial overlap of the two grids. We then convert this 1×1 km population-weighted precipitation surface into ZIP-level daily precipitation measures by overlaying ZIP-polygon boundaries and computing the area-weighted average for each ZIP. Cells that only partially overlap a ZIP contribute proportionally to the ZIP-level value. This procedure yields population-weighted precipitation estimates for every ZIP–day.

We measure realized daily temperature (in degrees Celsius) using data from NOAA's National Climatic Data Center (NCDC). The underlying temperature fields are provided as 25×25 km raster grids. We harmonize these data to our ZIP–day panel using the same population-weighted interpolation procedure applied to precipitation: we overlay each temperature grid with a 1×1 km population-density raster to compute

population-weighted temperature at the grid-cell level, and then aggregate these values to ZIP-code boundaries. This procedure yields population-weighted daily temperatures for each ZIP code in our sample.

Gridded forecasts of precipitation and temperature come from the National Digital Forecast Database (NDFD), which issues updated predictions in four 6-hour cycles each day. We average across these forecast windows to construct a single day-ahead prediction for each ZIP code. The NDFD data are processed using the same population-weighted area-overlap procedure applied to the realized weather data, yielding consistent ZIP–day forecasts for both precipitation and temperature.

Forecast accuracy is defined as the difference between the realized and predicted values of precipitation (in millimeters), with positive values indicating underprediction. Because both the forecast and realized weather grids are processed using the same population-weighted ZIP-level interpolation procedure, the resulting forecast errors capture genuine predictive inaccuracies rather than differences in spatial resolution or aggregation.

We construct daily ZIP-code–level measures of forecasted precipitation using high-resolution data from the National Digital Forecast Database (NDFD). The NDFD produces gridded forecasts at 2.5-km resolution and publishes them in GRIB2 format through a public NOAA S3 bucket. Our goal is to convert these raw meteorological grids into a ZIP-day panel aligned with our spending data while preserving the forecast horizons relevant for measuring forecast errors.

For each date in our sample, we query the NOAA S3 directory structure for the GRIB2 files associated with quantitative precipitation forecasts (QPF) and snowfall, identified respectively by their WMO headers (e.g., YIUZ98, YSUZ98). We download the first available forecast package for each day, ensuring consistent timing across the sample. All data are retrieved directly from the public bucket (noaa-ndfd-pds) using unsigned AWS requests. We then use the cfgrib engine in xarray to decode each GRIB2 file into a structured dataset containing forecast values, grid-cell coordinates, and forecast horizons. NDFD products include multiple short-term horizons; we extract the four consecutive 6- to 12-hour steps that correspond to the precipitation windows used in our empirical design. For each forecast step, we convert the raster into a tidy grid-cell–level dataset with latitude, longitude, and predicted precipitation.

Because NDFD forecast grids do not align with ZIP-code boundaries, we map each ZIP centroid to its nearest forecast grid cells. We compute great-circle distances using a Ball Tree with the Haversine metric, matching each ZIP to its ten nearest NDFD grid points. This approach captures the local weather conditions most relevant to ZIP-code populations while avoiding interpolation artifacts. Using these matched indices, we merge the forecasted precipitation values onto the ZIP-level dataset and compute, for each ZIP $\times$ date $\times$ horizon, the mean forecast amount across matched grid cells. We repeat this procedure for every forecast horizon and for both QPF and snowfall products. The resulting dataset includes, for each ZIP-date, up to four short-term precipitation forecasts and the corresponding realized precipitation measures. We compile the full 2022–2024 panel and export the final dataset as a ZIP-day-horizon file. This dataset serves as the basis for our empirical strategy, which contrasts ZIP-days with similar realized weather but differing forecasts to isolate the causal effects of forecast errors on economic activity.

Additionally, we incorporate polygon-based weather warnings layers issued by the National Weather Service (NWS) to capture salient, publicly communicated risk assessments that may influence consumer expectations independently of realized conditions. These warnings are mapped to ZIP codes using the same population-weighted procedure. Although warnings are not central to our main specification, they allow us to examine whether the behavioral effects of forecast errors vary with the presence of official hazard communication.

2.2 Economic Consumer Activity

To measure consumer spending, we use transaction-level data from SafeGraph, which report anonymized daily dollar expenditures at individual points of interest (POIs) across the United States. Each POI corresponds to a distinct business location and is linked to a NAICS industry code and precise geographic coordinates. The spending measure reflects total in-person card-based transactions recorded at the establishment and includes credit and debit card payments across major providers; online transactions are not included. For each POI, SafeGraph reports the total daily dollar volume of transactions, not the number of transactions, so all spending variables are measured in dollars. We aggregate POI-level spending to the ZIP–day–industry level by summing across all establishments within a ZIP code and NAICS industry. Because POIs vary substantially in size, the aggregation is unweighted—larger establishments naturally contribute more spending. Industry coverage varies across ZIP codes and days, but the dataset includes broad representation across retail, food services, transportation, entertainment, and other consumer-facing sectors.

To measure consumer mobility, we use Advan Foot Traffic data, which provide anonymized daily counts of visits to millions of commercial establishments based on aggregated mobile-device location signals. Each establishment is geolocated and assigned an industry category, allowing us to aggregate visit counts to the ZIP–day–industry level. We use foot traffic primarily as a complementary behavioral measure to distinguish between two economic mechanisms: (i) changes in the number of people who decide to visit commercial areas and (ii) changes in spending behavior among those who do visit. In our empirical analysis, foot traffic serves both as an auxiliary outcome and as a validation tool confirming that forecast errors affect spending intensity rather than local mobility. The data are not used as a control variable in the baseline specification but provide an important robustness check on the underlying consumer response to forecast inaccuracies.

In addition, we merge ZIP-code–level demographic and socioeconomic characteristics from the U.S. Census Bureau. These data provide information on income, age, race, transportation access, and other local attributes.

Summary Statistics

Table 1 reports summary statistics for all key variables, including precipitation, temperature, forecast errors, consumer spending, foot traffic, industries, businesses, and ZIP-level demographics. The final panel contains 27,775 unique ZIP codes and more than 150 million ZIP–day–industry observations from January 1, 2022 through December 31, 2024. The distributions of precipitation, forecast errors, and temperature bins used in our main specification are shown in Figure 1.

Table 1 and Figure 1 summarize the distributions of realized weather, forecasted weather, forecast accuracy, and spending across our sample of ZIP–day–industry observations. Daily precipitation is highly right-skewed, with a median of zero rainfall and a mean of 2.8 mm, indicating that most days experience no rain while a small number of days see heavy precipitation. Forecasted precipitation closely tracks realized rainfall on average, with a mean forecast error near zero, but the distribution of errors is wide and fat-tailed, reflecting frequent and sometimes substantial inaccuracies in one-day-ahead forecasts. Both overprediction and underprediction errors occur regularly, with moderate errors (3–10 mm) more common than large errors (exceeding 10 mm). Table 1 shows that consumer spending is also highly skewed, with many ZIP–day–industry cells recording low or zero spending and a long right tail driven by large commercial areas. These descriptive patterns highlight two key features of the data that motivate our empirical approach: forecast errors are common and sizable even conditional on similar realized weather, and daily consumer spending

exhibits substantial variation that could plausibly respond to forecast-driven behavioral adjustments.

3 Empirical Strategy

To estimate how short-term precipitation forecast errors affect daily consumer spending at the local level, we exploit variation in forecast errors while holding realized weather conditions constant at the ZIP–day–industry level. Panel A in Figure 1 displays the distribution of rainfall forecast errors. Approximately 31.5 percent of ZIP-day observations have forecast errors smaller than 3 mm, while 9 percent underestimate rainfall by more than 3 mm, and 10.5 percent overestimate rainfall by more than 3 mm. We compare ZIP-day-industry observations with small forecast errors to those with large forecast errors to estimate how inaccuracies in precipitation forecasts influence consumer spending.

There are several potential challenges to the estimation strategy. A primary concern is that forecast errors are not random and may be systematically related to realized weather conditions. Large forecast errors often occur on days with volatile or severe weather, creating potential correlations between the magnitude of the error and the underlying atmospheric conditions. Panel B in Figure 1 plots the distribution of underprediction and overprediction errors (box plot on y-axis) across bins of realized rainfall (x-axis), showing that forecasts are more likely to overpredict rain on light rainfall days and underpredict on heavy rainfall days. Such relationships can bias our estimates. To mitigate this concern, we control for realized rainfall and temperature at the ZIP-date level.

Another potential challenge is that ZIP codes differ in both forecast quality and the severity of the weather they experience. Locations that tend to receive forecasts with larger errors may also be those that experience more extreme weather or have weaker infrastructure, which could amplify the economic impact of rainfall. Similarly, certain periods of the year may systematically feature both larger forecast errors and more severe weather conditions. To address these concerns, we identify our estimates using variation within ZIP-month-year cells. After controlling for realized weather at the ZIP-date level, the remaining within–ZIP–month–year variation in forecast errors is likely uncorrelated with other factors that could influence consumer spending.

Our primary specification is

$$y_{z,i,t} = \sum_b \beta_b error_{z,t,b} + \sum_b \rho_b precipitation_{z,t,b} + \sum_b \tau_b temperature_{z,t,b} + \Phi_{z,m,y} + \Theta_t + \Gamma_i + \epsilon_{z,i,t} \quad (1)$$

where y_{zit} is the log of spending in ZIP z on date t in industry i. The key variable of interest is $error_{ztb}$, an indicator equal to 1 if the rainfall forecast error on date t for ZIP z falls in bin b, and 0 otherwise. We include five bins of forecast error: Large Overprediction (forecast exceeds realized rainfall by more than 10 mm), Slight Overprediction (forecast exceeds realized rainfall by 3-10 mm), Accurate (forecast error within 3 mm), Slight Underprediction (realized rainfall exceeds the forecast by 3-10 mm), and Large Underprediction (realized rainfall exceeds the forecast by more than 10 mm). Panel A in Figure 1 shows the distribution of observations across these bins.[4] In our specification, the "Accurate" category serves as the omitted reference group. Thus, each coefficient β_b captures the change in log spending on days in which the forecast error falls in bin b, relative to days with Accurate forecasts (i.e., forecast errors within 3 mm of realized rainfall).

The indicator variables $precipitation_{ztb}$ and $temperature_{ztb}$ control for the effects of average realized rainfall and temperature. Each indicator equals 1 if the observed rainfall (or temperature) on date t in ZIP z falls into bin b, and 0 otherwise. We include five bins of observed rainfall: No Rain, Very Light (0-2.5 mm),

[4] We include days with no rainfall and no error in the accurate category.

Light (2.5-7.5 mm), Heavy (7.5-20 mm), and Very Heavy (more than 20 mm); and five bins of observed average temperature: Very Cold (below 0°C), Cold (0-10°C), Moderate (10-20°C), Warm (20-30°C), and Hot (above 30°C). The distributions for these bins are shown in Panels C and D of Figure 1.

$\Phi_{z,m,y}$ are ZIP-month-year fixed effects, which ensure that identification comes from variation in forecast errors and spending within each ZIP, month, and year combination. Θ_t and Γ_i are date and industry fixed effects that control for time-specific and industry-specific factors—such as holidays, weekdays versus weekends, and broader macroeconomic conditions—that may influence spending. We also control for the number of point-of-sale devices reporting data within each industry for each observation. All regressions are weighted by ZIP population. Standard errors are clustered at the ZIP–date level, which corresponds to the level of treatment assignment because forecast errors are identical across industries within a given ZIP–date (Abadie et al. 2023).

4 Effect of Forecast Errors on Spending

Preliminary results

We begin with raw averages across all zip codes. Figure 2 plots demeaned log spending by rain and forecast error categories. Spending is demeaned at the zip–month–year level. Days are classified by whether realized rainfall is below 0.01 inches (dry) or above, and whether the forecast is "less than or equal to" actual rain or "greater than" actual rain. The figure shows that overprediction is associated with lower spending on both dry and rainy days, with a larger drop on dry days.

Table 2 formalizes this in a zip–date–industry panel regression with controls for temperature bins, business PoI count, and industry-, date-, and zip–month–year- fixed effects. The omitted category is dry days with forecast correctly predicting no rain. Relative to this baseline, overprediction on dry days reduces spending by approximately 1.24 percent. On rainy days, spending is lower overall: underprediction is associated with a 2.71 percent decline, while overprediction is associated with a 3.50 percent decline. This shows that, conditional on rain occurring, overprediction reduces spending by about 0.79 percent relative to underprediction. This preliminary analysis suggests that overprediction errors in rainfall forecast error reduce spending, with stronger effects observed on days with no rain.

Fixed effects analysis

Table 3, Column 4 reports estimates from our baseline specification in equation (1). The results show that precipitation forecast errors have economically meaningful and statistically significant effects on consumer spending, even after flexibly controlling for realized weather and rich fixed effects. Overprediction errors are associated with reductions in spending, while underprediction errors lead to increases. Specifically, when forecasted rainfall exceeds realized rainfall by more than 10 mm, same-day spending falls by approximately 3 percent relative to days with accurate forecasts. More moderate overpredictions of 3–10 mm reduce spending by about 1.2 percent. In contrast, large underpredictions—when realized rainfall exceeds the forecast by more than 10 mm—increase spending by roughly 1 percent, while moderate underpredictions raise spending by about 0.5 percent.

The table also reports the direct effect of realized precipitation on spending. Consistent with intuition, rainfall depresses consumer spending, with larger declines at higher rainfall intensities. Very light rain (0–2.5 mm) reduces spending by roughly 1.4 percent, while very heavy rain (exceeding 20 mm) lowers spending by

about 5.1 percent. Notably, the magnitude of the spending response to forecast errors is comparable to the direct effect of realized rainfall itself.

Taken together, these results yield two key insights. First, forecast inaccuracies meaningfully distort economic activity, independently of actual weather conditions. Second, the effects of forecast errors are strongly asymmetric: overpredictions generate substantially larger spending losses than the gains associated with underpredictions. This asymmetry suggests that false alarms—forecasts predicting adverse weather that does not materialize—are particularly costly in terms of foregone consumer spending.

Heterogeneity effects by Industry

We next examine how the spending response to precipitation forecast errors varies across industries. Figure 3 plots the estimated coefficients from our baseline specification separately by NAICS sectors. While the qualitative pattern is consistent across industries—overpredictions reduce spending and underpredictions increase it—the magnitude of these effects varies substantially.

Industries that rely heavily on discretionary, in-person trips exhibit the strongest responses to forecast errors. Large overpredictions of rainfall reduce same-day spending by approximately 4–6 percent in services, wholesale trade, and transportation, compared with declines of roughly 2–3 percent in retail, arts and entertainment, and accommodation and food services. Moderate overpredictions generate similar but attenuated effects. These patterns suggest that forecast-induced cancellations or postponements of planned trips disproportionately affect sectors where consumption requires physical presence and flexible timing.

Underprediction errors produce smaller and more muted responses across industries. Large underpredictions are associated with spending increases of around 1–2 percent in most sectors, with particularly weak responses in services. This asymmetry mirrors the aggregate results and indicates that unexpected favorable weather generates limited upside, even in sectors where activity is weather-sensitive.

Overall, the heterogeneity results reinforce the interpretation that forecast errors primarily operate through consumer trip-planning and activity avoidance. Sectors dependent on discretionary travel and in-person engagement are most exposed to false alarms, while sectors with more routine or necessity-driven demand show comparatively muted responses. These findings underscore that the economic costs of inaccurate precipitation forecasts are unevenly distributed across the economy, with implications for local business resilience and sector-specific exposure to forecast quality.

Dynamic effects of forecast errors

Forecast errors may affect not only same-day spending but also its timing, potentially shifting consumption to earlier or later days rather than reducing it outright. To examine these dynamics, we estimate distributed lead–lag effects of precipitation forecast errors following (Heutel, Miller and Molitor 2021). Specifically, we augment equation (1) with indicators for forecast errors occurring from one day prior $(t-1)$ through seven days after the event $(t+7)$, allowing spending responses to unfold over time. We then compute cumulative effects by aggregating the estimated lead and lag coefficients using the delta method.

Figure 4 presents the resulting dynamic patterns separately for each forecast error category. The top panels show the estimated lead–lag coefficients, while the bottom panels plot the corresponding cumulative effects. Several clear patterns emerge. First, overprediction errors generate large and immediate declines in spending on the event day. For large overpredictions, same-day spending falls by roughly 3 percent, consistent with the baseline estimates. Importantly, this decline is not offset by higher spending in subsequent days. Instead, the cumulative effect remains negative and statistically significant for at least seven days at about

4.1 percent, indicating that the lost spending is not merely postponed but represents a persistent contraction in demand.

Second, moderate overpredictions exhibit a similar but attenuated pattern. Spending declines on impact and remains below baseline over the following days at about 2.7 percent, with little evidence of rebound. This persistence suggests that false alarms lead consumers to cancel or forgo activities altogether rather than reschedule them shortly thereafter.

In contrast, underprediction errors display a markedly different dynamic profile. Especially large underpredictions but also slight underpredictions are associated with a modest increase in same-day spending, but this effect dissipates quickly. The cumulative response flattens within a few days and is close to zero over longer horizons, indicating that unexpected favorable weather generates limited and short-lived increases in economic activity.

Taken together, the dynamic results reinforce the asymmetric nature of forecast-error effects. Overpredictions cause immediate and lasting reductions in spending, while underpredictions yield only transitory gains. These findings provide strong evidence against simple intertemporal substitution and suggest that inaccurate forecasts—particularly false alarms—destroy economic activity rather than merely shifting it across days.

Economic scale of effects

We next translate our estimated spending responses into aggregate economic impacts. Figure 5 summarizes the implied annual change in consumer spending associated with precipitation forecast errors by error category, combining the estimated same-day effects with the frequency of each error type in the data. The accurate-forecast category (within ±3 mm) serves as the baseline and, by construction, contributes no aggregate spending change.

Overprediction errors generate economically meaningful losses. Slight overpredictions (–10 mm to –3 mm) reduce consumer spending by approximately 1.25 billion per year. Large overpredictions (greater than 10 mm) account for an additional 1.19 billion annual reduction in spending. Despite their lower frequency, large overprediction errors contribute nearly as much to aggregate losses as moderate overpredictions, due to their larger per-day spending impacts. Underprediction errors partially offset these losses but to a much smaller extent. Slight underpredictions (3–10 mm) increase annual spending by about 0.41 billion, while large underpredictions raise spending by approximately 0.46 billion per year. These gains reflect the modest and short-lived spending increases documented in the dynamic analysis.

Aggregating across all forecast-error categories, precipitation forecast inaccuracies reduce consumer spending by approximately 1.57 billion per year. In our sample, total annual spending is $93 billion. This means forecast errors reduce consumer spending by roughly 1.68%. We use a back-of-the-envelope calculation to scale this impact to the national level using Bureau of Economic Analysis (BEA) Personal Consumption Expenditures (PCE) data. Given that total estimated annual discretionary consumer spending in the United States as of 2025 is approximately $7.3 trillion, our calculations suggest that forecast errors may reduce national consumer spending by about $123 billion per year, equivalent to roughly 0.4% of U.S. GDP.[5] These losses are sizable in absolute terms and accrue from routine, short-term forecast errors rather than rare extreme events. The results underscore that even marginal improvements in day-ahead precipitation forecast

[5]Total annual personal consumption expenditures (PCE) reported is approximately $18 trillion in 2025. However, weather shocks are more likely to affect discretionary, non-contractual consumption rather than total expenditures. To account for discretionary expenditure, we exclude spending on durable goods (such as motor vehicles and appliances), housing services, household utilities, health care, financial services and insurance, communication services, education, professional services, child care, and household maintenance. Together, these categories account for roughly 60% of total PCE.

accuracy can yield meaningful economic returns through increased consumer activity.

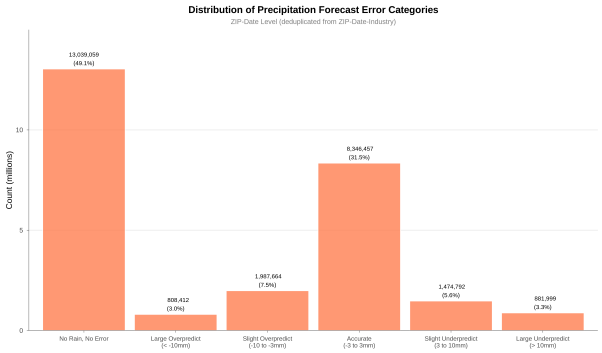
5 Conclusion

Weather forecasts play an important role in routine economic decision-making by households and firms. Acting on these forecasts often requires costly adjustments, such as reallocating time or foregoing consumption. When forecasts are inaccurate, these adjustments may be misaligned with realized weather, generating economic costs rather than improving outcomes. Using high-frequency transaction data and holding realized precipitation constant, this paper shows that errors in one-day-ahead rainfall forecasts meaningfully affect consumer spending.

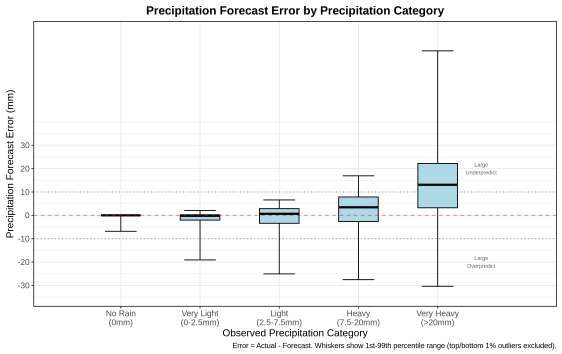
We find that overprediction and underprediction errors have asymmetric effects. Overpredicted rainfall leads to reductions in spending that persist for at least a week, while underpredicted rainfall results in smaller and shorter-lived increases. The persistence of spending declines following overprediction suggests that forecast errors generate foregone activity, such as cancelled trips or missed purchases, rather than simply shifting consumption across days. These effects are evident across industries and are larger in sectors that depend on discretionary, in-person activity.

These results indicate that errors in routine forecasts have economically relevant consequences, even outside extreme weather events. They also suggest that evaluating forecasts solely on technical accuracy may miss important dimensions of economic value. Incorporating the costs of routine forecast errors—particularly false alarms—into forecast evaluation may therefore be important for the efficient production of forecasts.

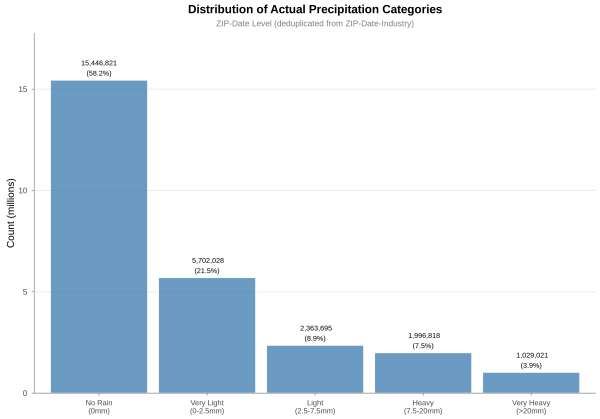
Figure 1: Distribution of precipitation, temperature, and precipitation forecast errors



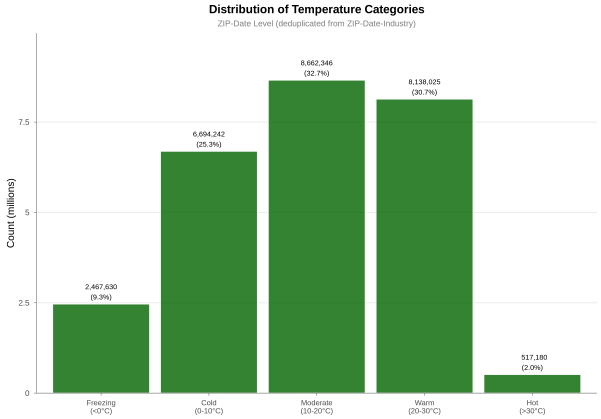
(a) Rainfall forecast errors



(b) Rainfall forecast errors by observed rain

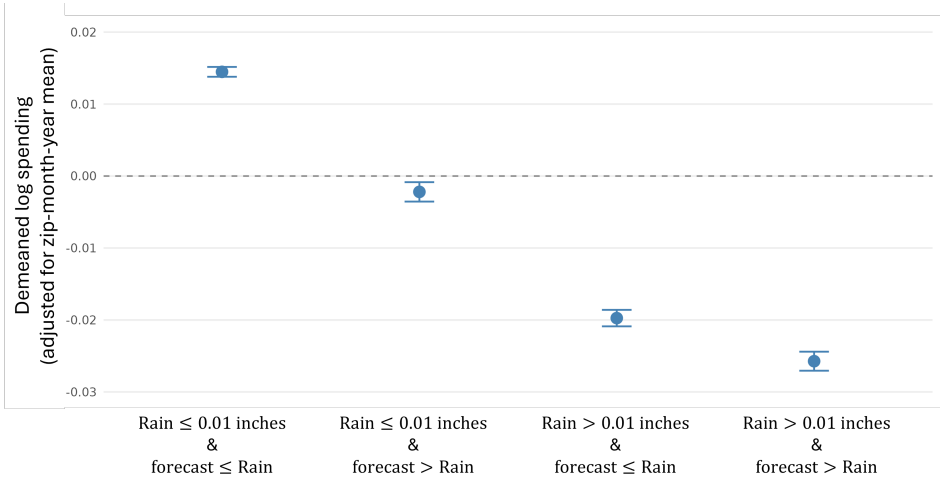


(c) Observed rainfall



(d) Observed temperature

Figure 2: Average Demeaned Log Spending by Rain and Forecast Error Category



Notes: The figure plots average log total spending demeaned at the zip–month–year level. Days are classified by realized rainfall (Rain $\leq$ 0.01 inches vs. Rain $>$ 0.01 inches) and whether the forecast is less than or equal to realized rainfall (Forecast $\leq$ Actual) or greater than realized rainfall (Forecast $>$ Actual). Each point represents the pooled average across all zip codes within the corresponding rain–forecast category.

Figure 3: Effect of forecast error: Heterogeneity by industry

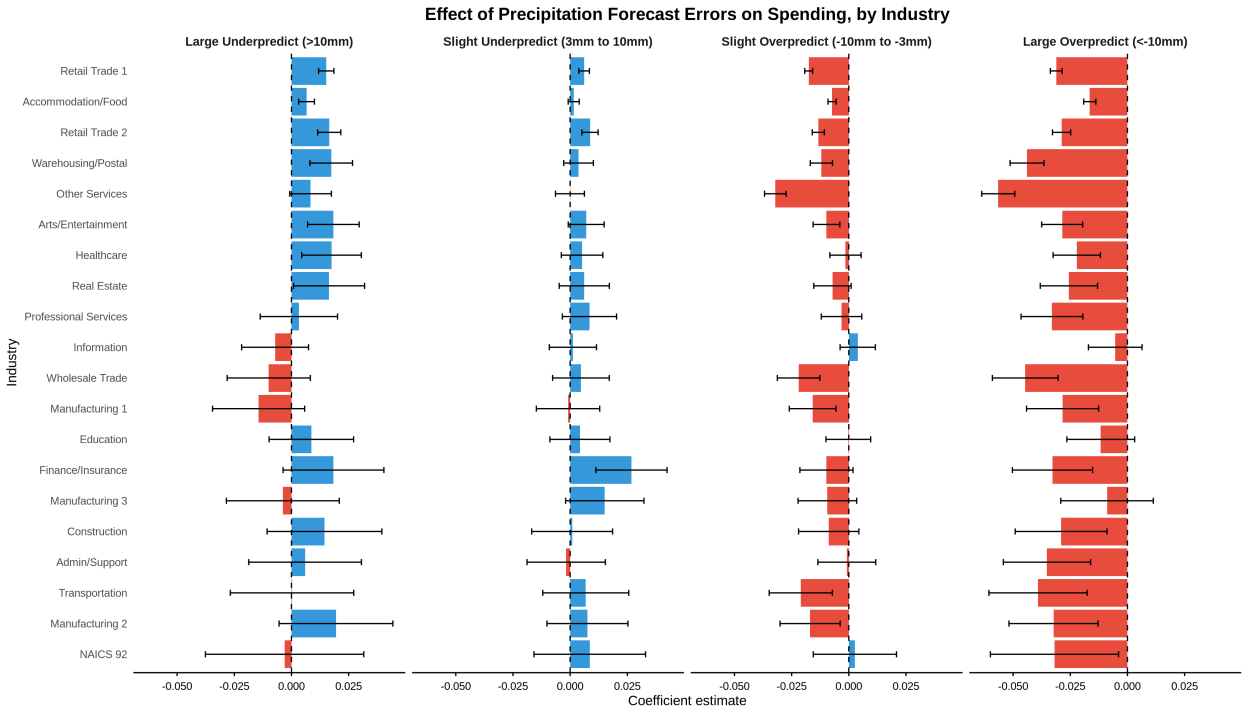
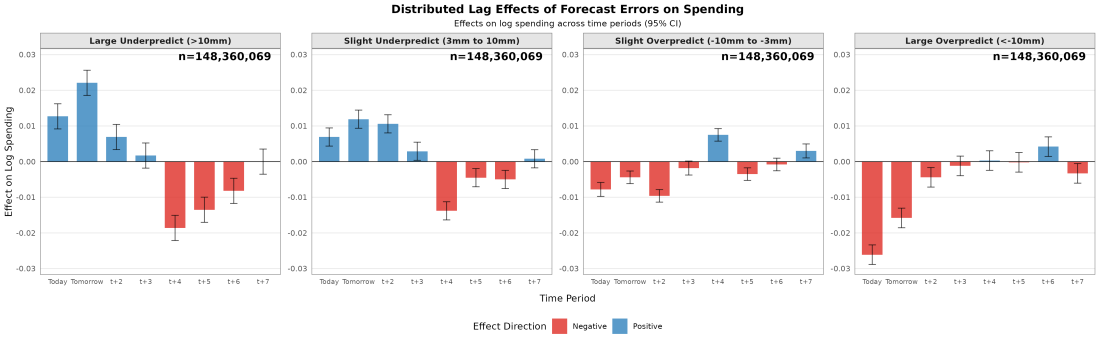
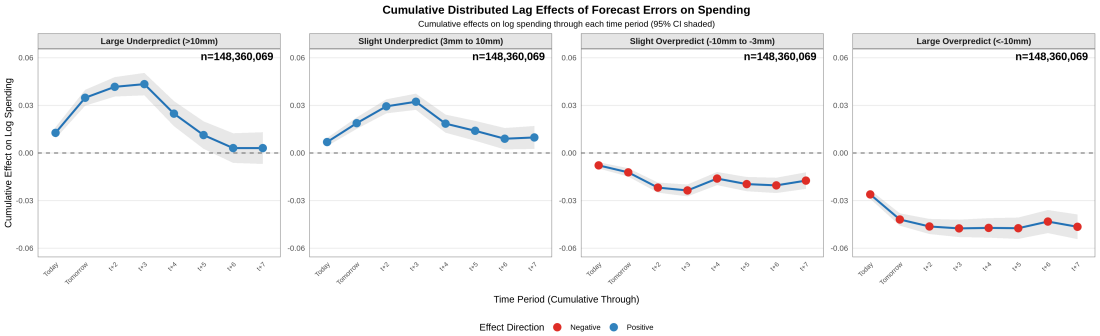


Figure 4: Dynamic effect of forecast errors on spending



(a) Lag effect of forecast errors



(b) Cumulative effect of forecast errors

Figure 5: Total annual impact of forecast errors on consumer spending in our sample

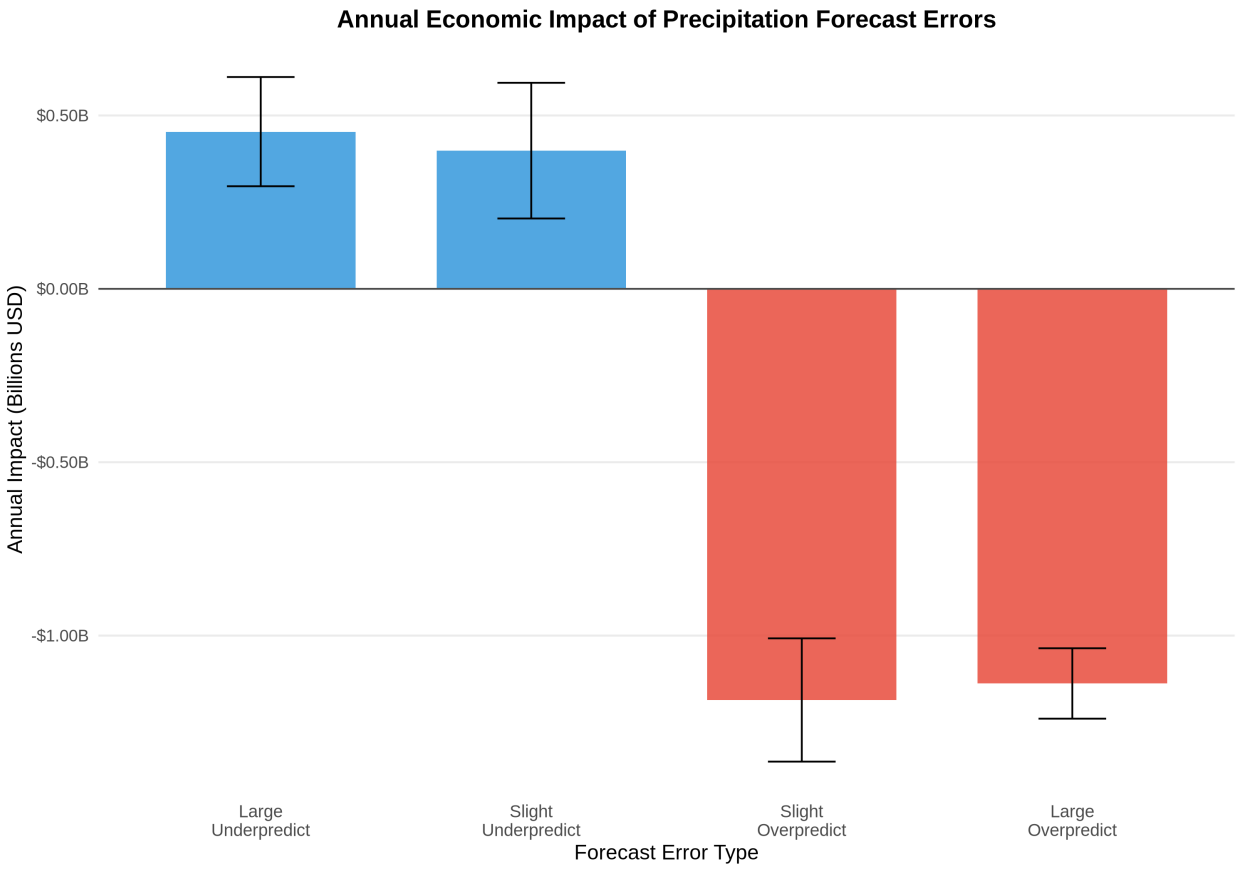


Table 1: SUMMARY STATISTICS

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Rain Actual (mm)	150,323,294	2.84	8.23	0.00	0.00	0.00	1.23	525.56
Rain Forecast (mm)	150,323,294	2.82	7.16	0.00	0.00	0.00	2.00	317.98
Rain Forecast Error (mm)	150,323,294	0.02	6.12	−281.73	−0.25	0.00	0.00	509.59
Temp Actual (C)	150,323,294	14.63	9.78	−36.32	7.58	15.83	22.34	41.55
Temp Forecast (C)	150,323,294	14.89	9.90	−36.95	7.70	16.20	22.75	41.13
Total Spending ($)	150,323,294	1,865.61	12,307.12	0	0	33	468	68,800,950
Log Total Spending	150,323,294	3.40	3.33	0.00	0.00	3.53	6.15	18.05
Population	150,323,294	21,929.01	18,512.65	0	6,211	18,073	32,888	137,213
POI Count (Total)	150,323,294	7.14	13.84	1	1	2	6	300

Table 2: Effect of Rain and Forecast Errors on Spending

	ln(total spending)
	(1)
Rain $\leq$ 0.01in, Forecast $>$ Actual	-0.0124***
	(0.0007)
Rain $>$ 0.01in, Forecast $\leq$ Actual	-0.0271***
	(0.0007)
Rain $>$ 0.01in, Forecast $>$ Actual	-0.0350***
	(0.0008)
temp cat binCold (0–10C)	0.0444***
	(0.0014)
temp cat binModerate (10–20C)	0.0680***
	(0.0017)
temp cat binWarm (20–30C)	0.0720***
	(0.0019)
temp cat binHot ($>$ 30C)	0.0748***
	(0.0029)
POI Count	0.0291***
	(0.0001)
Observations	150,087,469
R^2	0.64805
Within R^2	0.02255
Industry type fixed effects	X
zip–month–year fixed effects	X
date fixed effects	X

Standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Effect of forecast error on spending

Dependent Variables:	ln(Total Spending)					ln(TS/100k Pop)	Total Spending		ln(Total Spending)
Model:	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) OLS	(7) OLS	(8) Poisson	(9) OLS
Variables									
Large Overpredict (<-10mm)	-0.0366***	-0.0302***	-0.0302***	-0.0303***	-0.0303***	-0.0348***	-62.64***	-0.0251***	-0.0263***
	(0.0011)	(0.0010)	(0.0010)	(0.0014)	(0.0087)	(0.0016)	(7.766)	(0.0024)	(0.0011)
Slight Overpredict (-10 to -3mm)	-0.0215***	-0.0144***	-0.0144***	-0.0117***	-0.0117*	-0.0138***	-31.02***	-0.0109***	-0.0129***
	(0.0008)	(0.0007)	(0.0007)	(0.0009)	(0.0050)	(0.0011)	(5.571)	(0.0016)	(0.0007)
Slight Underpredict (3 to 10mm)	0.0088***	0.0049***	0.0049***	0.0052***	0.0052	0.0064***	3.908	0.0027	0.0061***
	(0.0010)	(0.0009)	(0.0009)	(0.0013)	(0.0052)	(0.0014)	(17.49)	(0.0055)	(0.0010)
Large Underpredict (>10mm)	0.0155***	0.0100***	0.0100***	0.0102***	0.0102	0.0117***	3.433	-0.0019	0.0131***
	(0.0014)	(0.0013)	(0.0013)	(0.0018)	(0.0089)	(0.0020)	(11.94)	(0.0036)	(0.0014)
Precipitation Category									
Very Light (0–2.5mm)	-0.0174***	-0.0155***	-0.0155***	-0.0139***	-0.0139**	-0.0172***	-17.52***	-0.0081***	-0.0144***
	(0.0006)	(0.0004)	(0.0004)	(0.0006)	(0.0048)	(0.0007)	(3.619)	(0.0011)	(0.0005)
Light (2.5–7.5mm)	-0.0355***	-0.0307***	-0.0307***	-0.0269***	-0.0269***	-0.0325***	-17.86*	-0.0130***	-0.0295***
	(0.0008)	(0.0007)	(0.0007)	(0.0010)	(0.0069)	(0.0011)	(7.736)	(0.0021)	(0.0007)
Heavy (7.5–20mm)	-0.0514***	-0.0416***	-0.0416***	-0.0386***	-0.0386***	-0.0460***	-47.79***	-0.0163***	-0.0416***
	(0.0010)	(0.0008)	(0.0008)	(0.0012)	(0.0086)	(0.0013)	(10.03)	(0.0031)	(0.0009)
Very Heavy (>20mm)	-0.0750***	-0.0542***	-0.0542***	-0.0507***	-0.0507***	-0.0592***	-76.35***	-0.0237***	-0.0456***
	(0.0014)	(0.0012)	(0.0012)	(0.0017)	(0.0127)	(0.0020)	(11.62)	(0.0035)	(0.0013)
Temperature Category									
Cold (0–10°C)	0.1003***	0.0409***	0.0409***	0.0460***	0.0460**	0.0563***	25.31*	0.0114**	0.0390***
	(0.0020)	(0.0009)	(0.0009)	(0.0014)	(0.0143)	(0.0016)	(10.40)	(0.0040)	(0.0011)
Moderate (10–20°C)	0.1462***	0.0644***	0.0644***	0.0700***	0.0700***	0.0849***	10.02	0.0130**	0.0588***
	(0.0026)	(0.0012)	(0.0012)	(0.0017)	(0.0199)	(0.0019)	(11.48)	(0.0042)	(0.0013)
Warm (20–30°C)	0.1403***	0.0688***	0.0688***	0.0745***	0.0745***	0.0894***	67.07***	0.0270***	0.0645***
	(0.0026)	(0.0013)	(0.0013)	(0.0019)	(0.0225)	(0.0022)	(12.33)	(0.0045)	(0.0014)
Hot (>30°C)	0.1233***	0.0653***	0.0653***	0.0759***	0.0759**	0.0894***	84.29***	0.0274***	0.0624***
	(0.0037)	(0.0021)	(0.0021)	(0.0029)	(0.0265)	(0.0032)	(14.97)	(0.0054)	(0.0022)
POI count	0.0510***	0.0479***	0.0479***	0.0304***	0.0304***	0.0228***	228.6***	0.0053***	
	(0.0008)	(0.0001)	(0.0001)	(0.0001)	(0.0007)	(0.0001)	(1.629)	(0.0002)	
Fixed-effects									
naics2	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
zip	Yes								
date	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
zip_month		Yes							
zip_month_year			Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics									
Observations	150,401,558	150,401,558	150,401,558	150,087,469	150,087,469	150,087,469	150,087,474	150,087,256	26,298,042
R^2	0.61066	0.62644	0.62644	0.65129	0.65129	0.62319	0.17770		0.91667

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, †: 0.1*
Notes: Reference categories: Accurate forecast (-3mm to 3mm), No Rain (0mm), Very Cold (<0°C).
Models 4–8 weighted by ZIP population. Model 5 uses two-way clustering (zip + date). Model 9 aggregates to zip-date level.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. 2023. "When should you adjust standard errors for clustering?" *The Quarterly Journal of Economics*, 138(1): 1–35.

Anand, Vaibhav. 2025. "Does getting forecasts earlier matter? Evidence from winter advisories and vehicle crashes." *American Economic Journal: Economic Policy*, 17(4): 106–134.

Burlig, Fiona, Amir Jina, Erin M Kelley, Gregory Lane, and Harshil Sahai. 2024. "The value of forecasts: Experimental evidence from India." *NBER Working Paper No*, 32173.

Downey, Mitch, Nelson Lind, and Jeffrey G Shrader. 2023. "Adjusting to rain before it falls." *Management Science*.

Farkas, Hannah, Manuel Linsenmeier, Marta Talevi, Paolo Avner, Bramka Arga Jafino, and Moussa Sidibe. 2025. "The Economic Value of Weather Forecasts: A Quantitative Systematic Literature Review." *Policy Research Working Paper Series*.

Heutel, Garth, Nolan H Miller, and David Molitor. 2021. "Adaptation and the mortality effects of temperature across US climate regions." *Review of Economics and Statistics*, 103(4): 740–753.

Molina, Renato, and Ivan Rudik. 2022. "The Social Value of Predicting Hurricanes." *CESifo Working Paper*.

Roth Tran, Brigitte. 2023. "Sellin'in the rain: Weather, climate, and retail sales." *Management Science*, 69(12): 7423–7447.

Shrader, Jeffrey. 2023. "Improving climate damage estimates by accounting for adaptation." *Available at SSRN 3212073*.

Shrader, Jeffrey G, Laura Bakkensen, and Derek Lemoine. 2023. "Fatal errors: The mortality value of accurate weather forecasts." National Bureau of Economic Research.

Song, Yuqi. 2024. "The value of weather forecasts: evidence from labor responses to accurate versus inaccurate temperature forecasts in China." *Journal of Environmental Economics and Management*, 125: 102970.