The Social Value of Hurricane Forecasts

Renato Molina and Ivan Rudik*

May 26, 2023

Abstract

What is the impact and value of hurricane forecasts? We study this question using newly-collected forecast data for major US hurricanes since 2005. We find that higher wind speed forecasts increase pre-landfall protective spending, but erroneous under-forecasts increase post-landfall damage and rebuilding costs. We develop a theoretically-grounded approach for estimating the marginal value of forecast improvements and find that the average annual improvement reduces total per-hurricane costs by over \$400,000/county. Improvements since 2007 reduced costs by 18%, totalling billions of dollars per hurricane. This exceeds the annual budget for all federal weather forecasting in the US.

JEL: Q54, Q58, C53

Keywords: extreme weather, natural disasters, hurricanes, tropical cyclones, forecasts, information, climate change

*Molina: University of Miami, Rosenstiel School of Marine, Atmospheric, and Earth Science and Miami Herbert Business School (renato.molina@earth.miami.edu). Rudik: Cornell University, Dyson School of Applied Economics and Management (irudik@cornell.edu). We thank Diego Cardoso, Luc Esprabens, Mike Huang, Steven Koller, Brian McNoldy, and Juan Carlos Villaseñor-Derbez for excellent research assistance. We are also grateful for the guidance and support provided by Andrea Schumacher and the staff at the Hurricane Forecast Improvement Program and the Weather Program Office at the National Oceanic and Atmospheric Administration, especially to Gina Eosco and Frank Marks. This manuscript benefited from discussions by Jeff Shrader, and comments from Christopher Costello, Gabriel Lade, Derek Lemoine, Cynthia Lin-Lawell, Antony Millner, Christopher Parmeter, Christopher Timmins, and Jinhua Zhao, as well as from feedback by seminar participants at Cornell University, St. John's University, the University of Miami, the Colorado Environmental Economics Workshop, the Kansas City Fed, the Occasional Workshop, the Seminar Series of the National Oceanic and Atmospheric Administration, and the Summer Conference of the Association of Environmental and Resource Economists. Funding for this project was provided by Grant NA20OAR4320472 from the National Oceanic and Atmospheric Administration.

Extreme weather like hurricanes, flooding, and extreme heat has devastated regions around the world. In the United States alone, these events have caused over $700 billion in damage since 2017, and trillions of dollars of damage since 1980, with the majority caused by hurricanes (Weinkle et al., 2018; NOAA Office for Coastal Management, 2022; NOAA National Centers for Environmental Information, 2022). One of the key levers for mitigating the destructive impacts of extreme weather, and especially hurricanes, is forecasting. Forecasts provide information on the expected strength, location, and timing of the event, allowing households and government actors to make better preparation decisions. Despite their importance and ubiquity, however, there is limited evidence on the historical value of hurricane forecasts or the potential value of future forecasting improvements.

In this paper, we investigate the value and economic impact of hurricane forecasts in the US. Using the actual models underpinning the national hurricane forecast system, we develop a new county-level dataset of forecasts and realizations of wind speed and precipitation, as well as the *ex ante* uncertainty embedded in the forecasts. Our dataset accounts for all major hurricanes in the US from 2005–2020, including all hurricanes that were classified at Category 3 or above (maximum wind speeds greater than 50 meters per second [m/s]), or generated at least $20 billion in damage. In total, our dataset accounts for 90% of hurricane-caused property damage and deaths in the mainland US.[1] We use these new data to (1) estimate how emergency federal expenditures for protecting against hurricanes respond to forecast information, (2) estimate the costs of forecast errors in terms of damages and increased expenditures for after-storm recovery, and (3) our primary contribution, develop a theoretically-grounded approach for estimating the *ex ante* marginal value of a forecast improvement. Our method accounts for unobserved protective actions taken prior to landfall, and it is flexible enough to be applied to other kinds of hazardous weather forecasts. We use this method to value the dramatic improvements in wind speed forecast accuracy since the 2000s.

The value of hurricane forecasts comes from how they help agents make better protective decisions. We start our analysis by estimating how hurricane forecasts affect the allocation of federal emergency protective funding across counties in the days before a hurricane reaches land. We find federal protective funding is targeted toward areas predicted to experience higher wind speeds. Counties forecast to experience hurricane-force winds receive $30 million more in protective funding than counties forecast to have lower, sub-hurricane force winds. This is equivalent to 0.5% more funding as a share of GDP or $200 more per person.

We then estimate the consequences of forecast errors. Conditional on hurricane intensity,

[1]The hurricanes in our study also account for over 70% of property damage and 40% of deaths for all environmental hazards.

forecast errors only matter if protective actions respond to forecasts and mitigate hurricane impacts. We find that there are economically significant increases in damages and post-landfall federal disaster recovery costs from underestimating hurricane wind speed. Relative to a perfect forecast, underestimating wind speed by 10 m/s – an error that would be a misclassification by up to two categories on the commonly used Saffir-Simpson scale – increases county-specific damages by $500 million and after-landfall federal emergency spending for recovery by $30 million. For damages, this is about 40% of county GDP or $15,000/person.

Finally, our main contribution is a new theoretically-grounded approach to estimate the expected total cost reduction from a marginal decrease in a forecast's *ex ante* standard deviation. We call this estimand the *value of a forecast improvement*. Lower standard deviation forecasts have smaller *ex post* errors, which means agents are less likely to uptake excess protective costs from an over-forecast, or incur excess damages and recovery costs from an under-forecast. We show that the value of a forecast improvement can be identified by first regressing the sum of damages and recovery spending on the *ex post* squared error in the forecast, and then multiplying the estimate by the baseline *ex ante* standard deviation at which we are valuing the marginal improvement. This approach does not require observing pre-landfall protective actions, so we are able to establish the value of a forecast improvement without having to track how agents might protect themselves against a hurricane.

Overall, we find that a one standard deviation forecast improvement reduces total protective spending, damages, and recovery spending costs in a county by $20 million per hurricane, equivalent to 2% of local GDP or over $700 per person. This value is driven entirely by counties that experience hurricane-force winds. We use our estimates to value the improvements in forecasting over 2007-2020 and find that they have led to an 18% reduction in total hurricane costs during this timespan. The total cost of the average hurricane in our dataset was $5 billion lower because of historical forecast improvements. The average benefit per hurricane is larger than the budget for all federal weather forecasting in the US in 2015 (Congressional Research Service, 2015).

Our paper adds to a sparse and relatively new literature on environmental forecasts. Some of the earliest work studied the role of weather forecasts in agriculture and shipping (Lave, 1963; Craft, 1998). More recently, researchers have studied the economic effects of precipitation forecasts in construction and automobile accidents (Downey, Lind and Shrader, Forthcoming; Anand, 2022), as well as how forecasts can be used to measure climate damages accounting for adaptation (Shrader, 2021). Two recent papers on temperature and pollution are closest to ours in spirit. Barwick, Li, Lin and Zou (2020) estimates the value of air pollution monitoring in China – accounting for some adaptation costs by directly estimating them – and finds that the benefits of the monitoring system exceed the costs by an order of magnitude. Shrader, Bakkensen and Lemoine (2022) evaluates the consequences of errors in routine temperature forecasts and finds that cutting errors in half would save thousands of lives per year, generating benefits of billions of dollars.

We contribute to this literature in several ways. First, we provide a novel overall assessment of the US hurricane forecast system and the improvements in its accuracy.[2] Second, we provide a general method to value any kind of hazard forecast, inclusive of all *ex ante* adaptation or protective costs. Third, our approach can value arbitrary improvements in forecast systems. This allows us to go beyond aggregate cost-benefit analysis and do marginal analyses that can speak to valuing historical improvements and determining optimal investment levels to improve forecasts.

This paper also contributes to a broader literature on the economic impacts of hurricanes and natural disasters. Hurricanes and tropical cyclones have been shown to be strongly associated with negative impacts on industrial production, national income, municipal financing, and welfare (Noy, 2009; Hsiang, 2010; Strobl, 2011; Hsiang and Jina, 2014; Bakkensen and Barrage, Forthcoming; Auh, Choi, Deryugina and Park, 2022; Jerch, Kahn and Lin, 2023). Historically, the US has suffered abnormally high damages due to hurricanes, and climate change is expected to amplify them further (Mendelsohn et al., 2012; Kossin et al., 2020).[3] Recent research suggests that about a third of the climate change-induced damages in the US could be offset by appropriate investments into long-run adaptation capital (Fried, Forthcoming).

We add to this literature by studying the role of information. Because the US has made only limited long-run hurricane adaptation investments, accurate forecasts are even more critical to reduce the impacts of hurricanes. Good forecasts help households and governmental agencies better allocate the necessary adaptive resources in the short window of time between the formation of a hurricane and its landfall. Our estimates suggest the avoided costs from actual hurricane forecast improvements since 2007 are about 10% of those achieved from the entire US adaptive capital stock (Fried, Forthcoming).

Finally, our findings also add to a limited stated-preference literature, which finds that, in the aggregate, households in hurricane-vulnerable areas consider the value of recent improvements in forecasts to be more than $300 million per year (Molina et al., 2021). We find that the actual value of hurricane forecast improvements is significantly larger.

[2]Martinez (2020) performs a similar exercise but only for 12 hour ahead forecasts of hurricane track, and using less than 100 observations of outcomes aggregated to the hurricane level.

[3]Hurricanes have recently been both moving slower across space while also intensifying much more rapidly (Kossin, 2018; Bhatia et al., 2019), potentially leading to their observed rising destructiveness in recent decades (Emanuel, 2005; Grinsted et al., 2019).

1 Data

Our analysis focuses on a county-hurricane as the unit of observation (e.g., Kings County, NY and Hurricane Sandy), and it uses data on hurricane forecasts, protective spending, recovery spending, and damages at the county-level for landfalling hurricanes from 2005 to 2020. Our data cover 18 out of the 29 total hurricanes from 2005–2020.

1.1 Forecasts

For our analysis, we reconstruct the National Hurricane Center (NHC) forecast products from their raw data and models to replicate the official National Oceanic and Atmospheric Administration (NOAA) forecast using the surface wind swath model described in Anderson et al. (2020). Here we roughly outline the data construction procedure. We follow DeMaria et al. (2009) and DeMaria et al. (2013) to produce a distribution of potential storm tracks, based on 1,000 different track forecasts. We then generate 1,000 forecasts of the hurricane's maximum wind speed which account for factors such as the previously realized wind speed and distance inland. Finally, we generate 1,000 wind speed forecast swaths by combining the 1,000 forecasts of track and maximum wind speed with the radii-climatological and persistence (CLIPER) model. For the purpose of this study, we focus on the maximum sustained wind in a county, which is defined as the maximum average wind speed over one minute. From here on, we call this measure "wind speed."

This process generates 1,000 swaths of wind speed forecasts across the entire US, where the variability across the swaths captures errors and uncertainties that are specific to each storm because of the recent history of forecast accuracy, the hurricane's movement and location, and the local climate.[4] The forecast wind speed mean and standard deviation are calculated across all 1,000 swaths for forecasts one, two, and three days prior to landfall and are identical to the official predictions by NOAA for the storms in our sample.

Observed wind speed is obtained by evaluating the observed storm track and wind speed in the wind swath model. Forecast errors are thus the difference between observed wind speed and the forecast ensemble mean, and the standard deviation of the forecast error is the same as the standard deviation of the forecast itself. For each hurricane, we aggregate these forecast statistics and errors to the county-level.

Precipitation is handled similarly to Molina et al. (2021). Using the same Monte Carlo ensembles as those for wind speed, we create 1,000 rainfall forecast swaths for each lead time

[4] We note that most economics research in this area only uses aggregated data instead of the full distribution of outcomes. For example, previous work has used probabilities of hurricane force winds (Kruttli et al., 2021) or fluctuations in the ENSO phenomenon (Downey et al., Forthcoming).

and hurricane using the probabilistic version of the Parametric Hurricane Rainfall Model (PHRaM) (Lonfat et al., 2007; Marks et al., 2020). To match observed precipitation with the forecast, only rainfall within 500 km of the storm center is considered. As with wind speed, precipitation is aggregated up to the county-hurricane level.

Our final dataset reports forecasts, realizations, and errors for wind speed and precipitation by county-hurricane pairs one, two, and three days prior to landfall.[5] For our analysis, we use the average forecast over this three day lead period. Some protective actions, like temporary levees, may take more than 1 day to complete, while others, like flying in generators for hospitals may be able to be done on shorter notice. Averaging the forecasts allows us to pick up all the different protective mechanisms in a parsimonious way.

Panels A-C in Figure 1 plot an example of our forecast data for hurricane Katrina. The black dashed line is the storm track that was forecast 72 hours before landfall, while the blue line is the actual storm track. Panel A, B, and C map the average wind speed forecast, its *ex ante* standard deviation, and the average differences between the predicted and actual wind speeds.

The figure shows that Katrina's 72 hour forecast was most uncertain around the predicted point of landfall, because of uncertainty about the degree of intensification before the storm's arrival. The forecast had errors in both directions because of the track prediction error, but the underestimates are larger than the overestimates because Katrina also grew to be stronger than expected.

Panel D in Figure 1 shows a time series of the annual average wind speed error for *all* hurricanes since 1990. The trend shows improvement over time, with a significant change in the rate of progress around 2007. This coincides with the year that Congress mandated the creation of the Hurricane Forecast Improvement Project (HFIP) due to the catastrophic 2004 and 2005 hurricane seasons. Prior to the HFIP in 2007, hurricane wind speed forecast errors were declining by 0.3 meters per second each year, or about a 0.4% annual improvement. Since 2007, hurricane wind speed forecasts errors have been declining by 0.21 m/s each year since 2007, or 3% annually, leading to a nearly 50% decline in wind speed errors. In our valuation exercises, we will estimate the value of this change in the rate of forecast improvement.

Appendix E contains several additional figures highlighting the distribution of forecast outcomes. There, we show that errors are correlated with intensity, that wind speed and precipitation are positively correlated, and that the *ex ante* uncertainty in in the forecast is a close match to the *ex post* error. It also shows that our data cover a wide range of intensities

[5] We do not explicitly study storm track forecasts, but track prediction, and its progress over the sample period, directly affects wind and precipitation forecasts.

of wind speeds – from no wind to nearly category 5 winds – even though our dataset does not have the universe of hurricanes during this time period.

1.2 Expenditures for Pre-Storm Protection and Post-Storm Recovery

Using a Freedom of Information Act (FOIA) request, we obtain data on measures and funding for public protection under the Public Assistance Grant Program (PAGM) (Kousky et al., 2016). PAGM is administered by the Federal Emergency Management Agency (FEMA) and provides grant assistance for eligible disasters (Kousky et al., 2016), thereby supporting efforts such as removal of debris, establishment of shelters and temporary levees, and emergency power generation. These funds are for efforts prior to a hurricane's landfall, and they are aimed at reducing overall storm impacts. We call these kinds of expenditures *protective expenditures*. PAGM also assists with rebuilding areas after a disaster. For example, it funds the repair, replacement, or restoration of disaster-damaged, publicly-owned facilities and facilities owned by certain nonprofit organizations, as well as the administrative expenses associated with these grants. These funds are provided for restoring an already damaged area so we call them *recovery expenditures*.

Typical beneficiaries of the PAGM include local governments and nonprofit organizations. The federal government provides a minimum of 75% of the cost of eligible assistance for these entities, and from FY2000 to FY2013, more than 90% of major disaster declarations received some assistance through the PAGM.

1.3 Economic Damages

Data on hurricane damages come from the Spatial Hazards Event and Losses Database for the United States (SHELDUS). SHELDUS provides information on the year and month of the event, the affected US counties, and the direct losses that stem from fatalities, injuries, and damages to property and crops at the county-level for each hurricane. Following the Environmental Protection Agency's guidelines, we estimate the losses from deaths using a value of a statistical life of $9.39 million in 2019 dollars (US EPA, 2022). Because we do not observe the type of injuries incurred, and have no way to clearly monetize them, we ignore injuries in our analysis. SHELDUS obtains the storm data from National Centers for Environmental Information (NCEI) and the estimates of mortality, injuries, and losses to crops and properties from various authorities, including insurance companies, the US Geological Survey, and the US Department of Agriculture (USDA). All damages are in 2019 dollars.

SHELDUS is widely-used, and is typically thought of as the best available dataset for measuring damages at a county-level (Gallagher, 2021; Auh et al., 2022). One concern with SHELDUS, however, is that it may under-report whether a county was affected by an extreme weather event (Gallagher, 2014), although misreporting error seems to be less of an issue for the most extreme events (Gallagher, 2021; Auh et al., 2022). For example, Gallagher (2014) shows that SHELDUS is missing records for over 90% of flooding events that had a presidential disaster declaration (PDD). PDDs are essentially a sufficient condition for there being damage. For our set of hurricanes, we find that SHELDUS contains records for every PDD listed by FEMA, which helps assuage concerns about data quality. Section B in the appendix further explores the robustness of using data from SHELDUS, where we show that for our hurricanes, SHELDUS slightly under-reports damages, but not deaths.

1.4 Summary Statistics

Table 1 shows summary statistics for the 18 storms in our sample. Note that these statistics are generated using only counties that experienced positive wind from the hurricane. The wind speed and precipitation columns are averages across all counties with the standard deviation in parentheses, while the damages and expenditures columns are summed across all counties. The table shows that there is substantial heterogeneity in mean wind speed, precipitation, and forecast errors across storms, and also across counties within the same storm. The total costs from the losses associated with all storms is half a trillion dollars, which highlights the economic importance of hurricanes in the continental US. Total emergency spending is about $30 billion, under one-tenth of the reported damages.

2 Methods and Results

2.1 Does FEMA Respond to Forecasts?

First, we estimate how FEMA's pre-landfall, protective emergency expenditures respond to the wind speed forecast.[6] We use the following flexible model for our main results:

$$\begin{aligned} f\left(\text{Protective FEMA Spending}_{csh}\right) = & \sum_{b \in \mathcal{B}_w} \beta_b^w 1(\text{Wind Forecast}_{csh} \in b) \\ & + \sum_{b \in \mathcal{B}_p} \beta_b^p 1(\text{Precip Forecast}_{csh} \in b) \\ & + \gamma_c + \eta_{sh} + \varepsilon_{csh}. \end{aligned} \tag{1}$$

$f\left(\text{Protective FEMA Spending}_{csh}\right)$ is either protective FEMA spending, protective FEMA spending as a share of county GDP, or protective FEMA spending per capita. The latter two outcomes adjust for how protective spending may be directed toward areas with larger economies or more people. $\mathcal{B}_w$ is a set of 5 m/s bins of wind speed forecasts up to 35 m/s, with forecasts of 0-5 m/s as the omitted category. $\mathcal{B}_p$ is a set of 20 mm bins of precipitation forecasts up to 200 mm, where forecasts of 0-20 mm are the omitted category. Recall that these forecasts are averages of the 1-3 day prior to landfall forecasts. We include both wind and precipitation forecasts in the same regression as they are positively correlated (Appendix E), and omitting one may result in omitted variable bias.

All of our specifications in Section 2 use county fixed effects, γ_c, and state-by-hurricane fixed effects, η_{sh}. γ_c controls for time-invariant factors that vary across counties that may drive protective spending and forecast hurricane intensity, like distance to the coast or elevation. η_{sh} addresses shocks that vary across states for the same hurricane, such as the political composition of the state government, and whether states used emergency declarations to marshall local resources. Following other papers in the literature (Hsiang, 2010; Deryugina, 2017), we compute spatial heteroskedasticity and autocorrelation consistent (HAC) standard errors using the approach documented by Conley (1999). Our standard errors account for arbitrary serial correlation within a county, and spatial correlation across all other counties that are within 400 km of a county's centroid. We note that this radius is about double the values used in this prior literature, and thus more conservative. The area traced out by this

[6] One channel through which FEMA spending is able to respond rapidly to new forecast information is the Hurricane Liaison Team (HLT). Its purpose is to connect local and federal officials with scientists and meteorologists at the National Hurricane Center. The HLT assists with properly communicating the forecast in order to better guide response operations, including evacuations, sheltering, and mobilizing manpower and equipment (Cannon, 2008).

radius is larger than Florida, Georgia, and Alabama combined. Figure 1 Panel C plots this radius along with a 600 km radius we use as a robustness check in Appendix D.3.

Figure 2 plots the wind speed estimates from equation (1). Panel A shows the effect of wind speed forecasts on pre-landfall protective expenditures. The results indicate that the effects of the wind speed forecast are negligible until above 20 m/s and increase rapidly up to 35 m/s, about the threshold for category 1 hurricane winds. Relative to counties forecast to have winds of 0-5 m/s, counties predicted to experience wind speeds of 30-35 m/s receive \$30 million more, while counties predicted to only experience wind speeds of 20-25 m/s – a low-end tropical storm forecast – receive only \$4 million more. Overall, these estimates show that protective emergency spending increases monotonically with the anticipated amount of wind, and that protective expenditures are targeted toward storms above the threshold to be designated as a hurricane.

Panels B and C plot estimates for spending as a share of GDP and per capita. Relative to 15-20 m/s or lower forecasts, spending increases by 0.1% of GDP or \$30/person for wind speed forecasts of 25-30 m/s, and by over 0.5% of GDP or \$200 per person for hurricanes forecast above 30 m/s.

2.2 Does Forecast Accuracy Matter?

Next, we test whether forecast errors affect hurricane damages and after-landfall recovery expenditures. As in Figure 1, we define forecast error as how much the forecast *underestimated* realized wind speeds. We estimate the effect of forecast errors on damages and FEMA recovery spending using the following flexible model:

$$\begin{aligned} f\left(\mathrm{Y}_{csh}\right) = & \sum_{b \in \mathcal{E}^w} \beta_b^w 1(\text{Wind Error}_{csh} \in b) + \sum_{b \in \mathcal{E}^p} \beta_b^p 1(\text{Precip Error}_{csh} \in b) \\ & + \sum_{b \in \mathcal{E}_i^w} \gamma_b^w 1(\text{Wind Realization}_{csh} \in b) + \sum_{b \in \mathcal{E}_i^p} \gamma_b^p 1(\text{Precip Realization}_{csh} \in b) \\ & + \gamma_c + \eta_{sh} + \varepsilon_{csh}. \end{aligned} \tag{2}$$

$f\left(\mathrm{Y}_{csh}\right)$ is the same set of functions as in equation (1). $\mathcal{E}^w$ and $\mathcal{E}^p$ are sets of bins of forecast errors (realization minus forecast) and $\mathcal{E}_i^w$ and $\mathcal{E}_i^p$ are sets of bins of intensity realizations. The omitted error bin categories are (-2,0] for wind and (-20,0] for precipitation. We flexibly control for hurricane realizations to make sure we are picking up the effect of forecast errors and not just that more intense storms tend to have larger errors as shown in Figure E.1 in the Appendix. Y_{csh} is either damages caused by the hurricane, or FEMA's post-landfall spending aimed at recovering the damaged area. The fixed effects and standard errors are identical to equation (1).

Figure 3 plots the results. Panels A and C plot the effect of wind speed forecast underestimates on damages and after-landfall recovery spending, Panels B and D plot the effect in terms of share of county GDP, while panels C and F plot the effects in per capita terms. All six panels show an increasing relationship between the outcome and wind speed underestimates. County damages are \$200 million higher if wind speed is underestimated by 6 m/s, and over \$500 million higher if underestimated by 10 m/s.[7] In GDP or per capita terms, a 10 m/s underestimate increases damages by 40% of GDP or \$15,000/person. The effects on recovery spending follow the same pattern: underestimating wind speed by 6 m/s increases recovery spending by \$10 million, while underestimating by 10 m/s increases recovery spending \$30 million. The latter error is equivalent to an increase spending by 0.75% of GDP or \$400/person. We also find that overestimates reduce damages relative to a perfect forecast, but the magnitude of the effect of overestimates is smaller than that of underestimates.[8]

2.3 What is the *Ex Ante* Value of Improving Hurricane Forecasts?

Figures 2 and 3 provide evidence for how the information in forecasts generates social value. Figure 2 shows that higher forecasts marshall more adaptive resources to an area. Figure 3 shows that, conditional on realized storm intensity, overestimating intensity (through higher forecasts) reduces *ex post* costs. We now formalize the *ex ante* value of improving hurricane forecasts.

2.3.1 Theoretical Foundation

Suppose a risk-neutral representative agent faces a future hurricane with total after-landfall costs from damages and recovery spending $D(x, a, \mathbf{i}, \mathbf{t})$. To be concise, we will call D damages from hereon. x is the hurricane's intensity (e.g., wind speed, precipitation); $\tilde{x}$ is the forecast of this intensity; a is the agent's continuous choice of before-landfall protective actions to reduce damages (e.g., sandbags, evacuations, structure hardening), which is a function of the forecast $\tilde{x}$ and has associated continuous cost function $C(a)$; $\mathbf{i}$ is a vector of time-invariant features of the agent's location i (e.g., elevation, proximity to the coast, long-lived capital structures); and $\mathbf{t}$ is a vector of common features across locations in period t. D is continuous in a. The agent has access to a forecast $\tilde{x}$ of the realized storm intensity x at time t, specific

[7] A 10 m/s error would result in misclassifying a storm by 1-2 categories.

[8] Figures 2 and 3 highlight a key cost tradeoff in reporting forecasts. Higher wind speed forecasts make underestimates less likely, which results in lower damages and recovery costs. However, this is not costless: higher forecasts increase federal protective expenditures, and presumably affect other costly protective actions we do not observe, and there may also be dynamic incentives to maintain public trust in forecasts.

to location i. The forecast is a noisy signal with normally distributed error: $e \sim \mathcal{N}(\mu, \sigma)$. We can write the forecast as a function of realized intensity and the error as $\tilde{x} = x + e$.

Our agent's objective is to minimize her total expected costs:

$$\mathcal{C}(\tilde{x}, \mu, \sigma, \mathbf{i}, \mathbf{t}) = \min_{a} \mathbb{E}_{x|\tilde{x}} \left[D(x, a, \mathbf{i}, \mathbf{t}) \right] + C\left(a\right).$$

where the expectation is over intensity given the forecast. We are interested in the marginal effect of a change in the standard deviation of the forecast error on the minimized total cost, which is equivalent to a change in minimized total costs from a change in the forecast error standard deviation. We define the value of a forecast improvement as the reduction in *ex ante* minimized expected cost – inclusive of both after-landfall damages and recovery spending, and before-landfall protective spending – from a marginal reduction in the standard deviation of forecast error. Proposition 1 provides an intuitive closed-form expression for this quantity.

Proposition 1 *The value of a forecast improvement is:*

$$\frac{\mathrm{d}\mathcal{C}(\tilde{x}, \mu, \sigma, \mathbf{i}, \mathbf{t})}{\mathrm{d}\sigma} = \frac{1}{\sigma^3} \operatorname{cov}_{x|\tilde{x}} \left(D(x, a^*(\tilde{x}, \mu, \sigma), \mathbf{i}, \mathbf{t}), (e - \mu)^2 \right) \tag{3}$$

$$= 2\sigma\beta_2. \tag{4}$$

Where $a^*(\tilde{x}, \mu, \sigma)$ is the optimized protective action choice, and β_2 is the coefficient from a regression of observed damages $D(x, a^*(\tilde{x}, \mu, \sigma), \mathbf{i}, \mathbf{t})$ on the observed squared demeaned intensity error $(e - \mu)^2$.

Proof: See Appendix C.1. □

Proposition 1 shows that the marginal value of a forecast improvement is proportional to a covariance between realized damages at the optimized protective actions and the the squared demeaned forecast error.[9] The value of an improvement and the covariance is positive if damages tend to be higher when the squared demeaned error is higher. Figure 3 provides evidence of the sign of the covariance: damages are increasing and convex in errors conditional on intensity, so the covariance between damages and squared errors will be positive. Better forecasts help the policymaker to reduce the difference between the *ex ante* optimized level of protective spending and the protective spending the policymaker would have chosen if she could observe realized hurricane intensity when making her decision.

The second line of Proposition 1 shows the *ex ante* value of a forecast improvement, inclusive of any protective actions, can be recovered by simply regressing total after-landfall costs on squared demeaned errors, and evaluating it at some reference forecast standard

[9]Section C.2 in the appendix formalizes a risk-averse agent.

deviation. A higher standard deviation baseline, reflecting more *ex ante* uncertainty, tends to raise the value of a forecast improvement. Unlike prior work, our new dataset reports the standard deviation of the forecast and forecast error. This is a necessary piece of data that is required to properly calculate the value of a forecast improvement.

In our empirical results, we derive estimates of β_2 with the following model:

$$\begin{aligned} f(\mathrm{D}_{csh}) \quad &= \quad \beta_2^w \left(e_{csh}^w - \mu_{csh}^w\right)^2 + \beta_2^p \left(e_{csh}^p - \mu_{csh}^p\right)^2 \\ &\quad + \sum_{b \in \mathcal{E}_i^w} \gamma_b^w 1(x_{csh}^w \in b) + \sum_{b \in \mathcal{E}_i^p} \gamma_b^p 1(p_{csh}^w \in b) \\ &\quad + \gamma_c + \eta_{sh} + \varepsilon_{csh}. \end{aligned} \tag{5}$$

D_{csh} is observed damages, $\left(e_{csh}^w - \mu_{csh}^w\right)^2$ is the observed squared demeaned error in wind speed, and $\left(e_{csh}^p - \mu_{csh}^p\right)^2$ is the observed squared demeaned error in precipitation. On the second line, we control for realizations, so the identifying variation is driven by forecasts conditional on storm intensity as with equation (2). The theoretical model assumes that the forecast distributional parameters were constant, but Figure E.1 shows that errors increase in storm intensity. Including flexible binned intensity controls ensures we do not confound the effect of higher-error forecasts with purely stronger storms. The fixed effects and standard errors follow the previous results. We note that we would obtain similar theoretically-grounded estimating equations given objectives of minimizing costs as a share of GDP or minimizing costs per person. We will explore these alternative objectives in our empirical analysis.

It is important to highlight the a key assumption that makes this approach work: forecast errors are normally distributed. This parametric assumption on beliefs lets us quantify how beliefs change as σ changes. In Appendix C.3 we plot the error distribution which is roughly normal.

2.3.2 Estimation Results

Table 2 reports our results corresponding to Proposition 1. The first panel shows the results assuming the planner is minimizing total costs, while the second and third panels show results if the agent is minimizing costs as a share of GDP or per capita costs. Within each panel, we report the coefficient estimate on squared demeaned wind errors. The sample average forecast standard deviation is 1.5, so the marginal value of an improvement of the average forecast is three times the coefficient. Because this is the main contribution of the paper, in the main text we show robustness of our results to a variety of specification choices. These include county-by-month of year effects which allow for county-specific seasonality in how hurricanes may cause damage throughout the year, county-by-year effects which controls for things like prior hurricane experience and damage that may change how forecast errors affect current damages, date effects which finely control for common shocks, as well as linear forecast errors. Our preferred specification is in column 3, which has our base fixed effects along with controls for intensity realizations and linear forecast errors.

The first panel shows that a one unit increase in the squared error of wind speed increases damages. For the sample average σ, the value of a forecast improvement is \$20 million per county in our preferred specification. A forecast improvement of 0.035 standard deviations, about 2% of the sample mean and an improvement that occurs just over every year in our sample, reduces total costs by \$400,000 for the average affected county. This result suggests that every year, forecast improvements are generating tens of millions of dollars of benefits per hurricane when aggregated over the entire US.

The second and third panels show that the value of a 1 standard deviation forecast improvement is \$740 in per-person terms, or 2.1% of county GDP. Using the same thought experiment as in the top panel, the annual average forecast improve would reduce costs by \$20 per person or 0.05% of GDP.

Column 8 replicates our preferred specification but allows for heterogeneous effects depending on whether a location experiences hurricane-force winds above 33 m/s. Hurricane-force winds are typically when structural damage such as fallen trees and broken windows occurs. The estimates in Column 8 show that the value of improvements comes entirely from places experiencing hurricane-force winds.

2.4 The Value of Historical Forecast Improvements

We now use our estimates in Table 2 to value historical improvements in forecast accuracy. Specifically, we estimate the value of the sudden increase in the rate of forecast improvement in 2007, as depicted in Figure 1.[10] For each storm after 2007, we compute its counterfactual forecast uncertainty if forecasts had continued to follow only the pre-HFIP 0.4% annual improvement, and then use the estimate in Column 8 of the top panel of Table 2 to value the increase in costs compared to the actual forecast uncertainty.

Our findings suggest that accelerated improvements in forecast accuracy since 2007 reduced hurricane costs – damages, recovery spending, and protective spending – by 18% or \$72 million per hurricane. How large is this value? \$5 billion is nearly the entire NOAA budget, five times the 2015 budget of the National Weather Service, the weather forecasting

[10] Figure 1 shows the decline in absolute wind speed error which is not quite the same as wind speed uncertainty. The wind speed standard deviation shows a 7% annual decline since the first storm in our dataset, however since we only have one year prior to 2007 we cannot construct a counterfactual trend.

arm of NOAA; and more than ten times cumulative budget of the HFIP since its inception in 2007, which was tasked with accelerating forecast improvements. The total avoided costs for the hurricanes in our sample is \$5 billion. Since the average per-hurricane benefit is larger than the annual budget of the National Weather Service, which is tasked with handling all forms of weather, the benefits likely exceed all the costs of incurred to generate the improvements.

2.5 Robustness

Section D contains a number of robustness checks of our main results that we summarize here. First, we show all of our results are robust to using a 600 km radius for spatial correlation, as well as inclusion of the additional fixed effects used in Table 2. Second, we show our results are robust to restricting our sample to only states on the Atlantic coast and Gulf of Mexico, and to dropping counties that were reported as having no damage in SHELDUS, but were in states that declared emergencies in order to receive recovery funding. Third, we show that our results are robust to using the inverse hyperbolic sine transformation. Fourth, we show forecast errors have larger effects for hurricane-force winds compared to sub-hurricane-force winds. Fifth, we show that errors increase all three types of damage, but mostly property damage.

We also replicate the analysis for precipitation forecasts, for which results tend to be less robust. This may be because the widely-used Saffir-Simpson categories for classifying hurricanes are based entirely on wind speed, and also the way in which hurricane strength has historically been communicated (Kantha, 2006; Murnane and Elsner, 2012).

3 Conclusion

In this paper we estimate the economic impact of hurricane forecasts and the value of improving them. We find forecasts are major determinants of the allocation of emergency resources, both before and after the storm. Counties projected to face the strongest wind speeds receive millions more in protective funding, while those that experienced the largest forecast underestimates had several times higher after-landfall recovery spending. We also find that forecasts affect hurricane damages. Conditional on wind speed, an under-forecast increases damages by hundreds of millions of dollars compared to an accurate one. These results suggest that forecasts direct valuable protective resources and actions.

Our main contribution is an estimate of the marginal value of reducing forecast uncertainty, inclusive of observed hurricane damages, observed after-landfall recovery costs, and unobserved before-landfall protective costs. Per-hurricane benefits from forecast improvements since 2007 amount to \$5 billion – a figure that far exceeds the total budget for all federal weather forecasting.

We conclude with several limitations that we leave for future work. First, our data do not capture all forms of damage and recovery costs. Accounting for these additional factors, such as longer-run social insurance costs (Deryugina, 2017), would only increase the value of a forecast improvement. Second, our estimates in the main text only cover the value generated by wind speed forecasts. While wind speed is arguably one of the leading attributes when it comes to hurricane damage (Murnane and Elsner, 2012), flooding and storm surge are important as well. Storm surge forecasting is in its infancy and likely less accurate compared to predicting storm track and wind speed, so there may be significant gains from further forecasting improvements along these additional dimensions of a hurricane.

16

Figure 1: Forecasts, Forecast Uncertainty, and Forecast Errors for Hurricane Katrina 72 Hours Before Landfall.

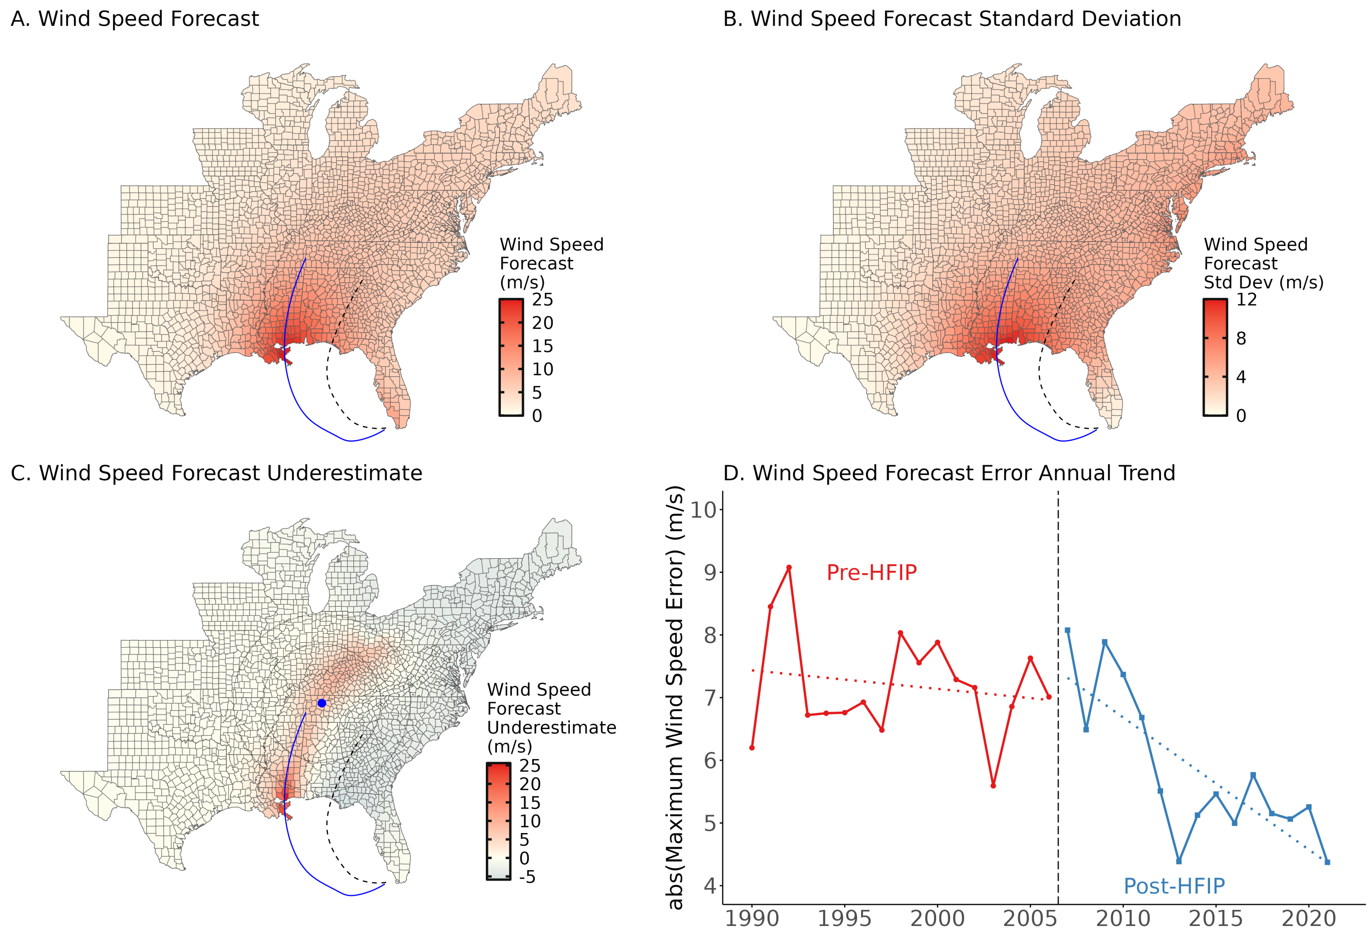

Note: Panels A, B, and C show Hurricane Katrina's 1-3 day ahead average landfall forecast wind speed, the forecast's *ex ante* standard deviation, and the forecast's errors. Positive values in Panel C are underestimates of the actual wind speed. The dotted circles in Panel C display radii of 400 km and 600 km. For our empirical results we use Conley Spatial HAC standard errors with a distance radius of 400 km. In the appendix we show robustness of our main results to the alternative radius. Panel D shows the absolute values of forecasting errors for wind speed, averaged across 1, 2, and 3 day ahead forecasts for all Atlantic tropical cyclones in a given year using aggregate data reported by the National Hurricane Center. Dotted lines are the best linear fits to the time series before and after 2007. The vertical dashed line is when the Hurricane Forecast Improvement Program was implemented in 2007, which expanded funding for forecast research and development. The archive for official historical records data is available at: `https://www.nhc.noaa.gov/verification` [Last visited on July 26, 2022].

Table 1: Summary Statistics by Storm.

Hurricane	Year	Wind Speed (m/s)	abs(Wind Error) (m/s)	Precip. (mm)	abs(Precip. Error) (mm)	Total Damage (Billion USD)	Protective Exp. (Billion USD)	Recovery Exp. (Billion USD)
Dennis	2005	6.03 (4.95)	1.73 (1.31)	23.81 (23.94)	13.19 (13.59)	2.35	0.02	0.13
Katrina	2005	8.88 (6.45)	2.82 (2.78)	30.42 (35.79)	23.81 (25.82)	214.83	1.46	11.39
Rita	2005	7.07 (4.89)	2.33 (2.01)	27.68 (35.46)	22.19 (23.06)	34.65	0.14	0.42
Wilma	2005	10.85 (9.37)	3.04 (5.16)	24.16 (36.09)	16.64 (26.36)	13.87	0.17	1.11
Ike	2008	11.37 (6.2)	6.94 (5.61)	17.05 (26.28)	14.58 (20.44)	37.39	0.24	1.18
Sandy	2012	10.62 (6.63)	2.63 (2.53)	29.35 (30.69)	18.28 (20.46)	58.22	1.19	3.63
Harvey	2017	9.92 (8.29)	1.76 (2.82)	76.48 (108.21)	60.84 (85.09)	94.24	0.44	1.79
Irma	2017	6.92 (6.51)	1.99 (1.58)	39.23 (44.98)	23.03 (29.47)	11.05	0.41	1.77
Florence	2018	7.23 (5.51)	0.8 (0.76)	25.21 (55.78)	15.1 (29.38)	5.15	0.15	0.53
Michael	2018	12.82 (8.8)	2.83 (3.67)	29.1 (30.09)	20.45 (22.44)	42.17	0.21	1.29
Barry	2019	6.34 (5.17)	1.76 (1.38)	25.83 (38.42)	17.82 (22.38)	0.02	0.02	0.02
Dorian	2019	10.45 (4.94)	0.41 (0.33)	13.87 (28.68)	8.23 (12.18)	0.04	0.05	0.13
Delta	2020	6.55 (5.42)	1.22 (0.84)	20.93 (27.49)	12.54 (16.94)	7.08	0.02	0.02
Hanna	2020	9.36 (6.42)	0.91 (2.15)	17.3 (30.64)	14.39 (20.09)	0.25	0.00	0.00
Isaias	2020	13.57 (7.7)	4.74 (4.69)	19.27 (24.63)	17.52 (21.64)	0.20	0.01	0.08
Laura	2020	7.54 (6.46)	2.28 (2.33)	15.47 (24.99)	10.83 (14.41)	24.57	0.21	0.60
Sally	2020	8.43 (5.28)	3.06 (2.82)	42.12 (52.78)	41.7 (42.33)	1.21	0.02	0.26
Zeta	2020	10.06 (7.56)	3.62 (3.78)	14.61 (12.64)	10.63 (8.58)	7.08	0.00	0.00

Note: Wind speed, precipitation, and their associated errors are averaged across counties that experienced positive wind speeds only. Standard deviations for these variables are reported in parentheses. Damages and expenditures are summed across counties for each storm. “Precip” is short for precipitation, and “Exp” is short for Expenditure. Wind speed is the maximum sustained wind speed in meters per second (m/s).

Figure 2: FEMA Protective Spending Responses to Forecasts.

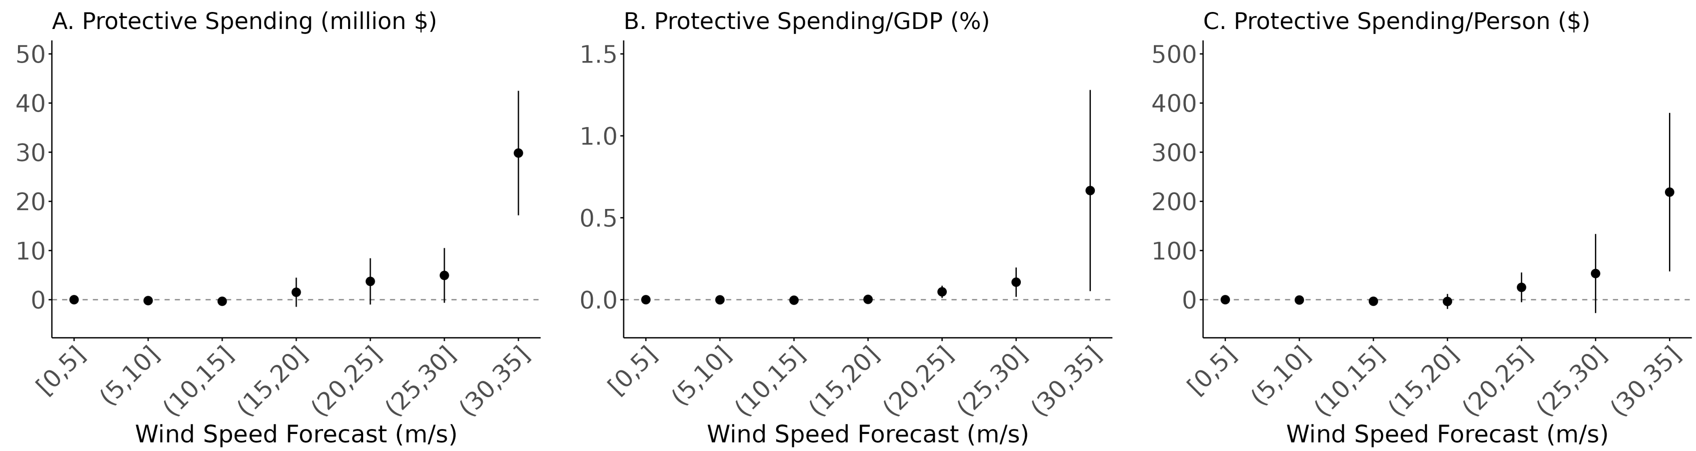

Note: The estimates correspond to equation (1). Points are point estimates and the bars are the 95% confidence intervals. The omitted category for each panel is [0,5]. All panels control for bins for the precipitation forecast, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

18

Figure 3: Forecast Errors, Damages, and *Ex Post* Recovery Spending.

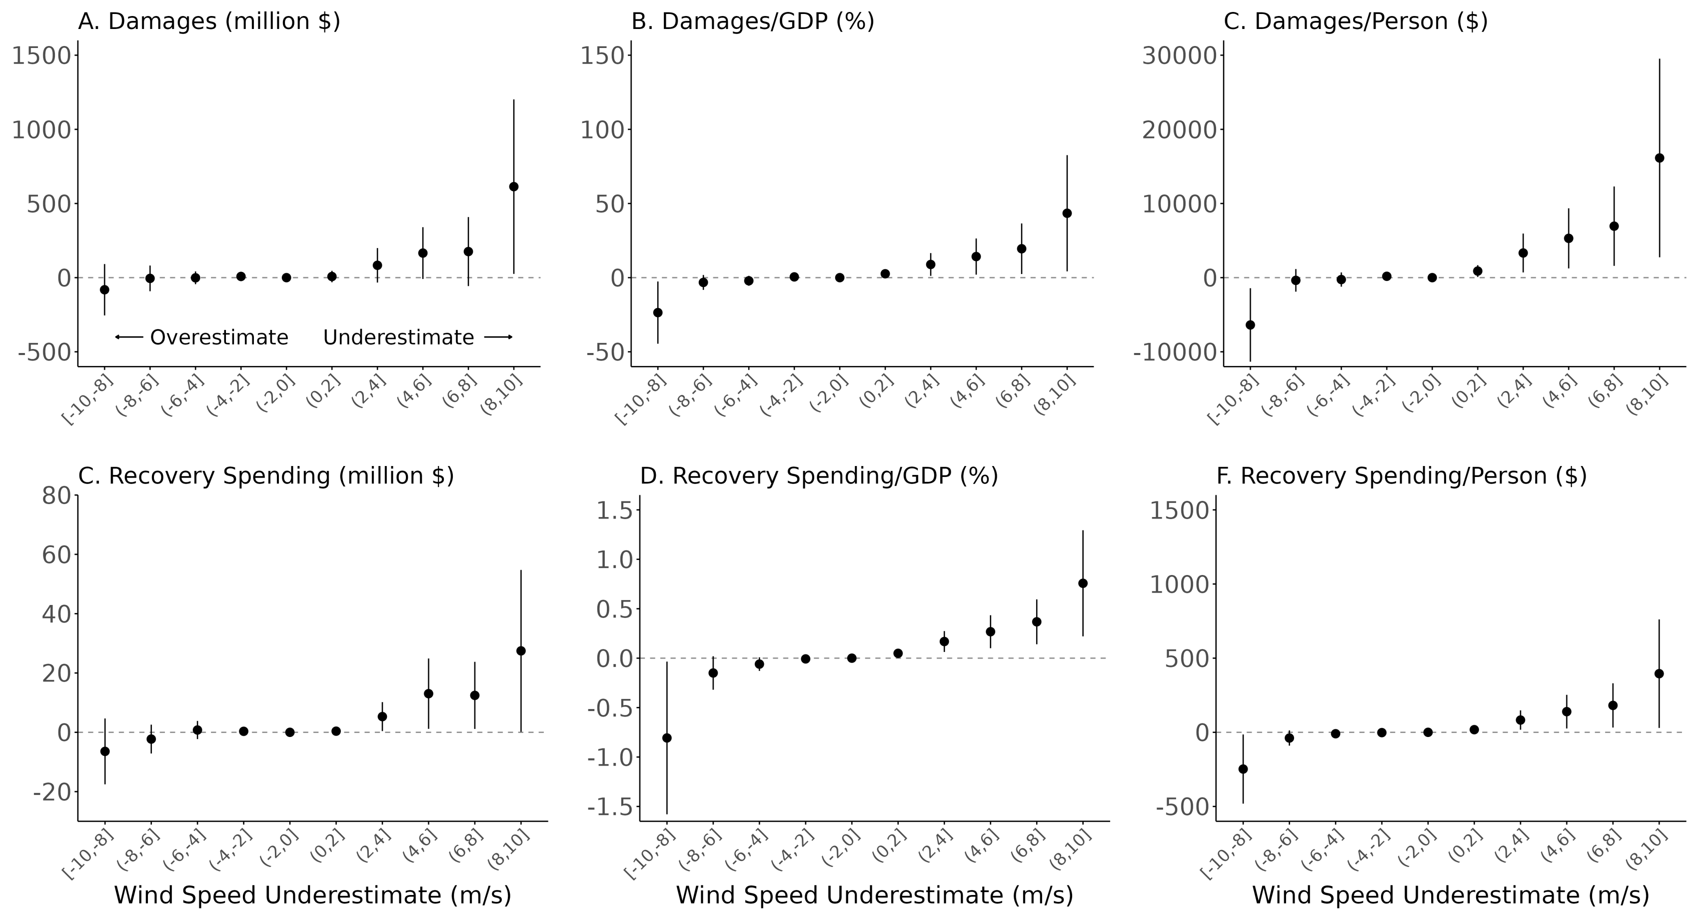

Note: The points are point estimates, and the bars are the 95% confidence intervals. The omitted category is $(-2, 0]$. All panels control for binned precipitation errors, binned realized wind speed, binned realized precipitation, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

19

Table 2: The Value of a Wind Speed Forecast Improvement.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Damages + Recovery Spending (million \$)							
$\beta_2 : (e-\mu)^2$	8.30***	8.37***	6.62***	6.52***	8.76**	8.77***	
	(3.20)	(3.23)	(2.54)	(2.35)	(3.44)	(3.02)	
Hurricane $\beta_2 : (e-\mu)^2$							9.65**
							(3.94)
Sub-Hurricane $\beta_2 : (e-\mu)^2$							-0.45
							(1.08)
(Damages + Recovery Spending) / GDP (%)							
$\beta_2 : (e-\mu)^2$	0.79***	0.82***	0.73***	0.74***	0.73***	0.72***	
	(0.30)	(0.31)	(0.27)	(0.24)	(0.28)	(0.22)	
Hurricane $\beta_2 : (e-\mu)^2$							1.03***
							(0.32)
Sub-Hurricane $\beta_2 : (e-\mu)^2$							0.01
							(0.03)
Damages + Recovery Spending Per Capita (\$/person)							
$\beta_2 : (e-\mu)^2$	273.30***	280.97***	247.68***	260.03***	284.53***	296.52***	
	(78.43)	(79.04)	(66.48)	(65.38)	(80.60)	(74.92)	
Hurricane $\beta_2 : (e-\mu)^2$							351.92***
							(82.54)
Sub-Hurricane $\beta_2 : (e-\mu)^2$							4.37
							(9.78)
Observations	55,350	55,350	55,350	55,350	55,350	55,350	55,350
Realized Wind/Precip Bins		✓	✓	✓	✓	✓	✓
Level Wind/Precip Error			✓	✓	✓	✓	✓
State-Hurricane FE	✓	✓	✓	✓	✓	✓	✓
County FE	✓	✓	✓				✓
County-Month of Year FE				✓		✓	
County-Year FE					✓	✓	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

20

References

Anand, Vaibhav (2022) "The value of forecast improvements: Evidence from advisory lead times and vehicle crashes," *Available at SSRN 4206910*.

Anderson, G Brooke, Joshua Ferreri, Mohammad Al-Hamdan, William Crosson, Andrea Schumacher, Seth Guikema, Steven Quiring, Dirk Eddelbuettel, Meilin Yan, and Roger D Peng (2020) "Assessing united states county-level exposure for research on tropical cyclones and human health," *Environmental health perspectives*, Vol. 128, No. 10, p. 107009.

Auh, Jun Kyung, Jaewon Choi, Tatyana Deryugina, and Tim Park (2022) "Natural disasters and municipal bonds,"Technical report, National Bureau of Economic Research.

Bakkensen, Laura and Lint Barrage (Forthcoming) "Climate shocks, cyclones, and economic growth: bridging the micro-macro gap," *The Economic Journal*.

Barwick, Panle Jia, Shanjun Li, Liguo Lin, and Eric Zou (2020) "From fog to smog: The value of pollution information,"Technical report, National Bureau of Economic Research.

Bhatia, Kieran T, Gabriel A Vecchi, Thomas R Knutson, Hiroyuki Murakami, James Kossin, Keith W Dixon, and Carolyn E Whitlock (2019) "Recent increases in tropical cyclone intensification rates," *Nature communications*, Vol. 10, No. 1, pp. 1–9.

Cannon, Glenn (2008) "Strengthening the federal emergency management agency's disaster response capabilities," *Journal of business continuity & emergency planning*, Vol. 2, No. 3, pp. 281–293.

Congressional Research Service (2015) "The national oceanic and atmospheric administration (noaa) budget for fy2016."

Conley, Timothy G (1999) "Gmm estimation with cross sectional dependence," *Journal of econometrics*, Vol. 92, No. 1, pp. 1–45.

Craft, Erik D (1998) "The value of weather information services for nineteenth-century great lakes shipping," *American Economic Review*, pp. 1059–1076.

Czajkowski, Jeffrey, Kevin Simmons, and Daniel Sutter (2011) "An analysis of coastal and inland fatalities in landfalling us hurricanes," *Natural hazards*, Vol. 59, No. 3, pp. 1513–1531.

DeMaria, Mark (1996) "A history of hurricane forecasting for the atlantic basin, 1920–1995," *Historical Essays on Meteorology 1919–1995*, pp. 263–305.

DeMaria, Mark, John A Knaff, Michael J Brennan, Daniel Brown, Richard D Knabb, Robert T DeMaria, Andrea Schumacher, Christopher A Lauer, David P Roberts, Charles R Sampson et al. (2013) "Improvements to the operational tropical cyclone wind speed probability model," *Weather and forecasting*, Vol. 28, No. 3, pp. 586–602.

DeMaria, Mark, John A Knaff, Richard Knabb, Chris Lauer, Charles R Sampson, and Robert T DeMaria (2009) "A new method for estimating tropical cyclone wind speed probabilities," *Weather and Forecasting*, Vol. 24, No. 6, pp. 1573–1591.

Deryugina, Tatyana (2017) "The fiscal cost of hurricanes: Disaster aid versus social insurance," *American Economic Journal: Economic Policy*, Vol. 9, No. 3, pp. 168–98.

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

Emanuel, Kerry (2005) "Increasing destructiveness of tropical cyclones over the past 30 years," *Nature*, Vol. 436, No. 7051, pp. 686–688.

Enten, Harry (2017) "Why hurricane maria surprised forecasters by getting so strong so fast."

Fried, Stephie (Forthcoming) "Seawalls and stilts: A quantitative macro study of climate adaptation," *The Review of Economic Studies*.

Gall, Robert, James Franklin, Frank Marks, Edward N Rappaport, and Frederick Toepfer (2013) "The hurricane forecast improvement project," *Bulletin of the American Meteorological Society*, Vol. 94, No. 3, pp. 329–343.

Gallagher, Justin (2014) "Learning about an infrequent event: evidence from flood insurance take-up in the united states," *American Economic Journal: Applied Economics*, pp. 206–233.

——— (2021) "Natural disasters that cause no damage: Accounting for the selective reporting of weather damage,"Technical report, Working Paper.

Grinsted, Aslak, Peter Ditlevsen, and Jens Hesselbjerg Christensen (2019) "Normalized us hurricane damage estimates using area of total destruction, 1900- 2018," *Proceedings of the National Academy of Sciences*, Vol. 116, No. 48, pp. 23942–23946.

Hsiang, Solomon M (2010) "Temperatures and cyclones strongly associated with economic production in the caribbean and central america," *Proceedings of the National Academy of sciences*, Vol. 107, No. 35, pp. 15367–15372.

Hsiang, Solomon M and Amir S Jina (2014) "The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones,"Technical report, National Bureau of Economic Research.

Jerch, Rhiannon, Matthew E Kahn, and Gary C Lin (2023) "Local public finance dynamics and hurricane shocks," *Journal of Urban Economics*, Vol. 134, p. 103516.

Kantha, Lakshmi (2006) "Time to replace the saffir-simpson hurricane scale?" *Eos, Transactions American Geophysical Union*, Vol. 87, No. 1, pp. 3–6.

Kossin, James P (2018) "A global slowdown of tropical-cyclone translation speed," *Nature*, Vol. 558, No. 7708, pp. 104–107.

Kossin, James P, Kenneth R Knapp, Timothy L Olander, and Christopher S Velden (2020) "Global increase in major tropical cyclone exceedance probability over the past four decades," *Proceedings of the National Academy of Sciences*, Vol. 117, No. 22, pp. 11975–11980.

Kousky, Carolyn, Brett Lingle, and Leonard Shabman (2016) "Fema public assistance grants: Implications of a disaster deductible," *Report, Resources for the Future*.

Kruttli, Mathias S, Brigitte Roth Tran, and Sumudu W Watugala (2021) "Pricing poseidon: Extreme weather uncertainty and firm return dynamics."

Lave, Lester B (1963) "The value of better weather information to the raisin industry," *Econometrica*, pp. 151–164.

Lonfat, Manuel, Robert Rogers, Timothy Marchok, and Frank D Marks Jr (2007) "A parametric model for predicting hurricane rainfall," *Monthly Weather Review*, Vol. 135, No. 9, pp. 3086–3097.

Marks, FD, Brian D McNoldy, Mu-Chieh Ko, and Andrea B Schumacher (2020) "Development of a probabilistic tropical cyclone rainfall model: P-rain," in *100th American Meteorological Society Annual Meeting*, AMS.

Marks, N. Kurkowski M. DeMaria, F. and M. Brennan (2019) "Hurricane forecast improvement project years ten to fifteen strategic plan," technical report, National Oceanic and Atmospheric Administration.

Martinez, Andrew B (2020) "Forecast accuracy matters for hurricane damage," *Econometrics*, Vol. 8, No. 2, p. 18.

Mendelsohn, Robert, Kerry Emanuel, Shun Chonabayashi, and Laura Bakkensen (2012) "The impact of climate change on global tropical cyclone damage," *Nature climate change*, Vol. 2, No. 3, pp. 205–209.

Molina, Renato, David Letson, Brian McNoldy, Pallab Mozumder, and Matthew Varkony (2021) "Striving for improvement: The perceived value of improving hurricane forecast accuracy," *Bulletin of the American Meteorological Society*, pp. 1–41.

Murnane, Richard J and James B Elsner (2012) "Maximum wind speeds and us hurricane losses," *Geophysical Research Letters*, Vol. 39, No. 16.

NOAA National Centers for Environmental Information (2022) "Billion Dollar Weather and Climate Disasters."

NOAA Office for Coastal Management (2022) "Fast Facts: Hurricane Costs."

Norcross, Bryan (2018) "Hurricane forecasts are great — as long as the storms themselves behave," *The Washington Post*.

Norton, Grady (1951) "Hurricanes of the 1950 season," in *Monthly Weather Review*, Citeseer.

Noy, Ilan (2009) "The macroeconomic consequences of disasters," *Journal of Development economics*, Vol. 88, No. 2, pp. 221–231.

Resnick, Brian (2018) "Hurricane michael and why it's so hard to predict storm intensity."

Sheets, Robert C (1990) "The national hurricane center—past, present, and future," *Weather and Forecasting*, Vol. 5, No. 2, pp. 185–232.

Shrader, Jeffrey (2021) "Improving climate damage estimates by accounting for adaptation," *Available at SSRN 3212073*.

Shrader, Jeffrey G, Laura Bakkensen, and Derek Lemoine (2022) "Fatal errors: The mortality value of accurate weather forecasts."

Strobl, Eric (2011) "The economic growth impact of hurricanes: Evidence from us coastal counties," *Review of Economics and Statistics*, Vol. 93, No. 2, pp. 575–589.

US EPA (2022) "Guidelines for preparing economic analyses."

Weinkle, Jessica, Chris Landsea, Douglas Collins, Rade Musulin, Ryan P Crompton, Philip J Klotzbach, and Roger Pielke (2018) "Normalized hurricane damage in the continental united states 1900–2017," *Nature Sustainability*, Vol. 1, No. 12, pp. 808–813.

Appendix

A Background

Officially sanctioned forecasts for hurricanes in the US date back to the late 1800s. Initially, forecasts and warnings were the responsibility of the US Weather Bureau, which relied on land-based weather stations and observations from vessels along the Atlantic coast and in the Gulf of Mexico (DeMaria, 1996). The detection of hurricanes and the ability to predict their paths significantly improved following World War II, with advances in the understanding of atmospheric processes, and access to aircraft reconnaissance and radar. These advances eventually led to the establishment of the Miami Hurricane Warning Office to provide yearly hurricane season summaries for the US (Norton, 1951).

Further federal commitment to hurricane forecasts came after a series of devastating storms in the 1954 and 1955 seasons, which led Congress to create the National Hurricane Research Project in 1956 (DeMaria, 1996). The eventual coordination and collocation of the Research Project, the Warning Office, and Aircraft Operations resulted in what is now known as the National Hurricane Center (Sheets, 1990).

The advent of computer modeling and meteorological satellites resulted in significant improvements in forecasting capabilities after 1970, thereby setting the foundation for modern forecasts (Sheets, 1990). Nonetheless, and while forecasts of hurricane tracks continued to improve gradually over the years, generating reliable forecasts of wind speed remained a challenge. These limitations became evident to US policy makers when the country experienced 13 hurricane landfalls during the 2002-2005 hurricane seasons – 10 of them in 2004 and 2005. The 2004 and 2005 hurricanes alone were responsible for at least 5,200 deaths and $229 billion in damages, underscoring the need for more aggressive forecast improvements (Czajkowski et al., 2011; Strobl, 2011).[11]

Following these catastrophic seasons, Congress mandated the creation of the Hurricane Forecast Improvement Project (HFIP) in 2007 by the National Oceanic and Atmospheric Administration (NOAA). The goal of the HFIP was to improve both storm track and wind intensity forecasts through coordinated efforts from the research and operational communities (Gall et al., 2013). Initially, the project was intended to continue for 10 years. It funded research and operations, and made significant investments in high-performance computing to support both these aims. The original 10-year goals were to reduce average track errors

[11] Hurricane Charley, which struck in 2004, was the strongest hurricane to reach land in the US since 1992. In 2005, Katrina struck, becoming one of the costliest hurricanes in US history. That same year, Rita and Wilma (two of the strongest Atlantic storms ever recorded at that time) also struck.

by 50%, and to reduce average wind speed errors by 50%. In addition, the project was also expected to improve the prediction of rapid intensification of hurricanes by increasing the probability of detection, reducing the false-alarm rate, and extending the forecast lead time from five to seven days.

In 2017 the project was given a new name, the Hurricane Forecast Improvement Program, and funding was renewed and extended through at least 2024. The goals of the extension include an emphasis on an advanced, unified-modeling system, probabilistic-hazard guidance, and improved communication of risk and uncertainty (Marks and Brennan, 2019). From 2009 to 2019, the HFIP budget for research and operations totaled approximately $250 million.

By any measure, these recent efforts to improve forecasts have been successful. Figure 1 shows that prior to the HFIP in 2007, hurricane wind speed forecast errors were declining by 0.03 meters per second each year, or about a 0.4% annual improvement. Since the inception of the HFIP in 2007, there has been a dramatic increase in the quality of the forecasts. Hurricane wind speed forecasts errors have been declining by 0.21 m/s each year since 2007, or 3% annually.[12]

B SHELDUS Data Quality

Previous work has raised concerns about the quality of the damages data reported in SHELDUS for earlier versions of the dataset (Gallagher, 2014), particularly for smaller events. Here we probe whether SHELDUS has quality issues in more recent versions of the data and for large hazards like hurricanes.

First, we compare SHELDUS to an alternative dataset provided by the National Center for Environmental Information (NCEI) at the National Oceanic and Atmospheric Administration (NOAA). The NCEI data reports estimates of damages for disasters where the estimated damage is over $1 billion. NCEI's damages include property damage, damage to agriculture, as well as time costs like business interruptions. Figure B.1 plots the raw SHELDUS property damage against the NCEI estimated damage for each hurricane. The figure shows that there is a very strong correspondence between the two data sets. When aggregating over all hurricanes in our sample, SHELDUS reports $530 billion in damage and 2,700 deaths while NCEI reports $596 billion in damage and 2,645 deaths. All monetary values are in 2019 dollars. This amounts to an 11% underestimate in aggregate property damages and 2% overestimate in aggregate deaths.

[12] Historically, it has been much more difficult to forecast the intensity a storm will have than to forecast the track it will follow (Resnick, 2018). This is due to a combination of many factors, including previously poor computational resolutions, and difficulties in predicting which storms will go under rapid intensification as they near landfall (Enten, 2017; Norcross, 2018).

Figure B.1: Aggregate Hurricane Damages and Deaths: SHELDUS Versus NCEI.

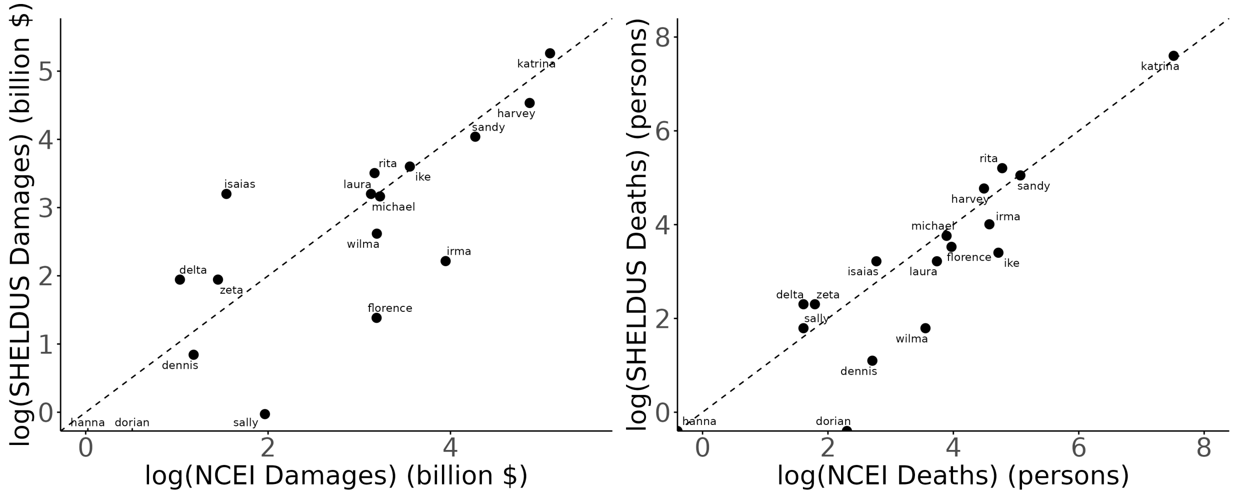

The y axis aggregates the county-level SHELDUS property damages (left) and deaths (right) to the hurricane-level for all hurricanes in our dataset. The x axis is the hurricane-level point estimates reported by the NCEI. Hurricane Barry is omitted since it is not reported in the NCEI data.

Figure B.1 shows that SHELDUS appears to do a relatively good job of measuring damages compared to NCEI. Another second way to probe data quality is to test the importance of counties that are not reported as having damage in SHELDUS, but were in states that issued a Presidential Disaster Declaration (PDD). These are counties that may have been erroneously coded as zero damage in SHELDUS. In Section D below, we test the sensitivity of our results to dropping these "error counties." We find that dropping these data points does not meaningfully affect our results.

One last point to note is that after Gallagher (2014) was published, the SHELDUS metadata shows significant retroactive corrections to the data and inclusion of additional records of damages. Each annual update to the database, there are typically several prior years which are retroactively updated. In some years, there are tens of thousands of records added and the corrections adjust annual damages by tens of billions of dollars, which may also explain the difference in data quality we observe from the earlier literature.

C Theoretical Foundation

C.1 Proof of Proposition 1

An agent is aiming to minimize the total costs of an incoming hurricane which consist of protective spending before the storm, and uncertain damages and *ex post* recovery costs after the storm. The agent has access to a forecast $\tilde{x}$ of the realized storm intensity x. The forecast is a noisy signal with error $e = x - \tilde{x}$ where we can write the intensity as the deviation from the forecast: $x = \tilde{x} + e$. As in the results shown in Figure 3, the error is how much the forecast underestimates the actual intensity. We assume that forecast errors are normally distributed: $e \sim \mathcal{N}(\mu, \sigma^2)$. We denote the probability density function as a function: $\Phi(\frac{e-\mu}{\sigma})$.

The agent uses the forecast to choose her level of protective spending a where protective spending decreases storm damages and *ex post* recovery costs. Damages are a function of realized intensity, the chosen level of protective spending, and location-specific and time-specific factors: $D(x, a, \mathbf{i}, \mathbf{t})$. We are interested in the reduction in minimized total costs from a reduction in the forecast error standard deviation. Our agent's objective is to minimize her expected total costs:

$$\mathcal{C}(\tilde{x}, \mu, \sigma, \mathbf{i}, \mathbf{t}) = \min_a \mathbb{E}_{x|\tilde{x}} \left[D(x, a, \mathbf{i}, \mathbf{t})\right] + C(a).$$

The cost of an increase in the standard deviation is:

$$\frac{\mathrm{d}\mathcal{C}(\tilde{x}, \mu, \sigma, \mathbf{i}, \mathbf{t})}{\mathrm{d}\sigma} = \frac{\partial \mathcal{C}(\tilde{x}, \mu, \sigma, \mathbf{i}, \mathbf{t})}{\partial \sigma}. \tag{C.1}$$

The envelope theorem gives us that:

$$\begin{aligned}\frac{\partial \mathcal{C}(\tilde{x}, \mu, \sigma, \mathbf{i}, \mathbf{t})}{\partial \sigma} &= \int D(x, a^*(\tilde{x}), \mathbf{i}, \mathbf{t}) \frac{\partial \Phi(\frac{e-\mu}{\sigma})}{\partial \sigma} de \\ &+ \int D_a(x, a^*(\tilde{x}), \mathbf{i}, \mathbf{t}) \frac{\partial a^*(\tilde{x})}{\partial \sigma} \Phi\left(\frac{e-\mu}{\sigma}\right) de \\ &+ C'(a) \frac{\partial a^*(\tilde{x})}{\partial \sigma},\end{aligned}$$

where the last two terms are zero from the agent's first order condition, and we write optimized protective actions explicitly as a function of the forecast for clarity. Taking the partial derivative inside the integral then gives:

$$\frac{\partial \mathcal{C}(\tilde{x}, \mu, \sigma, \mathbf{i}, \mathbf{t})}{\partial \sigma} = \int D(x, a^*(\tilde{x}), \mathbf{i}, \mathbf{t}) \left[\frac{(e-\mu)^2 - \sigma^2}{\sigma^3}\right] \Phi\left(\frac{e-\mu}{\sigma}\right) dx.$$

Since the normal density is still in the expression, it can go back into expectation notation as:

$$\frac{\partial \mathcal{C}(\tilde{x}, \mu, \sigma, \mathbf{i}, \mathbf{t})}{\partial \sigma} = \mathbb{E}_{x|\tilde{x}} \left\{ D(x, a^*(\tilde{x}), \mathbf{i}, \mathbf{t}) \left[\frac{(e-\mu)^2 - \sigma^2}{\sigma^3}\right] \right\},$$

where the expectation is again with respect to e. We can get a closed form solution by using the covariance identity:

$$\begin{aligned}
\frac{\partial \mathcal{C}(\tilde{x},\mu,\sigma,\mathbf{i},\mathbf{t})}{\partial \sigma} =& \mathbb{E}_{x|\tilde{x}}\left\{D(x,a^*(\tilde{x}),\mathbf{i},\mathbf{t}) \times \left[\frac{(e-\mu)^2-\sigma^2}{\sigma^3}\right]\right\} \\
=& \frac{1}{\sigma^3}\mathbb{E}_{x|\tilde{x}}\left\{D(x,a^*(\tilde{x}),\mathbf{i},\mathbf{t}) \times \left[(e-\mu)^2-\sigma^2\right]\right\} \\
=& \frac{1}{\sigma^3}\Bigg[\mathrm{cov}_{x|\tilde{x}}\left(D(x,a^*(\tilde{x}),\mathbf{i},\mathbf{t}),(e-\mu)^2\right) \\
& + \mathbb{E}_{x|\tilde{x}}\left\{D(x,a^*(\tilde{x}),\mathbf{i},\mathbf{t})\right\}\underbrace{\mathbb{E}_{x|\tilde{x}}\left\{(e-\mu)^2-\sigma^2\right\}}_{=0}\Bigg] \\
=& \frac{1}{\sigma^3}\mathrm{cov}_{x|\tilde{x}}\left(D(x,a^*(\tilde{x}),\mathbf{i},\mathbf{t}),(e-\mu)^2\right), \qquad \text{(C.2)}
\end{aligned}$$

where we use $x \sim \mathcal{N}(\mu,\sigma)$ so that $\mathbb{E}_{x|\tilde{x}}\{(e-\mu)^2\} = \sigma^2$. This result proves the first part of the proposition.

Next, we return to the last line in equation (C.2):

$$\frac{\mathrm{d}\mathcal{C}(\tilde{x},\mu,\sigma,\mathbf{i},\mathbf{t})}{\mathrm{d}\sigma} = \frac{1}{\sigma^3}\mathrm{cov}_{x|\tilde{x}}\left(D(x,a^*(\tilde{x}),\mathbf{i},\mathbf{t}),(e-\mu)^2\right).$$

First, compute the variance of $(e-\mu)^2$:

$$\begin{aligned}
\mathrm{var}\left((e-\mu)^2\right) &= \mathbb{E}_{x|\tilde{x}}\left[\left((e-\mu)^2 - \mathbb{E}_{x|\tilde{x}}\left[(e-\mu)^2\right]\right)^2\right] \\
&= \mathbb{E}_{x|\tilde{x}}\left[\left((e-\mu)^2-\sigma^2\right)^2\right] \\
&= \mathbb{E}_{x|\tilde{x}}\left[(e-\mu)^4\right] - 2\sigma^4 + \sigma^4 \\
&= 3\sigma^4 - 2\sigma^4 + \sigma^4 \\
&= 2\sigma^4, \qquad \text{(C.3)}
\end{aligned}$$

where the last line uses the fact that the fourth central moment of a normal variable x is $3\sigma^4$.

Use this to result to rewrite the last line in equation (C.2) as:

$$\frac{\mathrm{d}\mathcal{C}(\tilde{x},\mu,\sigma,\mathbf{i},\mathbf{t})}{\mathrm{d}\sigma} = 2\sigma\frac{\mathrm{cov}_{x|\tilde{x}}\left(D(x,a^*(\tilde{x}),\mathbf{i},\mathbf{t}),(e-\mu)^2\right)}{\mathrm{var}\left((e-\mu)^2\right)}. \qquad \text{(C.4)}$$

The covariance-variance ratio term is just a coefficient from a regression of damages on the squared demeaned error in wind speed. Denote this regression coefficient as β_2. The final expression is:

$$\frac{\mathrm{d}\mathcal{C}(\tilde{x}, \mu, \sigma, \mathbf{i}, \mathbf{t})}{\mathrm{d}\sigma} = 2\sigma\beta_2. \tag{C.5}$$

The marginal value of a forecast improvement is the product of this regression coefficient and the standard deviation of forecast errors at which we want to evaluate the marginal value.

We note two important features of the derivation. The first is that this result depends on assuming that the forecast error distribution has constant parameters that do not depend on the actual hurricane intensity. Figure E.1 shows that this is unlikely to be the case: forecast errors and squared demeaned forecast errors are both correlated with realized intensity This means that we need to flexibly condition on storm intensity when performing the regression over our full dataset, otherwise the coefficient on squared errors may just be picking up the fact that stronger storms cause more damage regardless of the forecast error.

The second feature is that we need to demean the error before squaring. Using e^2 instead of $(e-\mu)^2$ results in omitted terms $-2e\mu + \mu^2$ that are likely correlated with e^2 resulting in our estimates of the value of a forecast improvement being biased. What is this bias? The omitted terms pick up first-order effects of the forecast bias on observed squared error. By omitting these terms we may erroneously be confounding effects of the forecast bias μ with what we are interested in, the forecast uncertainty σ.

C.2 Risk Averse Agent

Suppose that now the agent is risk-averse with some continuous, increasing, and concave utility function U. The agent's utility is over consumption of some numeraire good Y, the costs of protective actions $C(a)$, and damages $D(x, a, \mathbf{i}, \mathbf{t})$. The agent maximizes her expected utility:

$$V(\tilde{x}, \mu, \sigma, Y, \mathbf{i}, \mathbf{t}) = \max_a \mathbb{E}_{x|\tilde{x}} \left\{ U\left(Y - [D(x, a, \mathbf{i}, \mathbf{t})] - C(a)\right) \right\}.$$

To simplify notation, let $U(x, a, Y, \mathbf{i}, \mathbf{t}) \equiv U\left(Y - [D(x, a, \mathbf{i}, \mathbf{t})] - C(a)\right)$. The value of a forecast improvement is:

$$-\frac{\mathrm{d}V(\mu, \sigma, Y, \mathbf{i}, \mathbf{t})}{\mathrm{d}\sigma} = -\frac{\partial V(\mu, \sigma, Y, \mathbf{i}, \mathbf{t})}{\partial \sigma}. \tag{C.6}$$

First, note that here we use a decrease in the standard deviation since we are maximizing utility instead of minimizing costs. Second, note that since the agent has a utility function over her (random) payoff, the value of a forecast improvement is in units of utility and will need to be translated back into dollar terms.

The envelope theorem gives us that:

$$\frac{\partial V(\tilde{x}, \mu, \sigma, Y, \mathbf{i}, \mathbf{t})}{\partial \sigma} = \int U(x, a^*(\tilde{x}), Y, \mathbf{i}, \mathbf{t}) \frac{\partial \Phi(\frac{e-\mu}{\sigma})}{\partial \sigma} de$$

recalling that $x = \tilde{x} - e$.

The rest of the proof follows identically to Proposition 1 except where $D(x, a^*(\tilde{x}), \mathbf{i}, \mathbf{t})$ is replaced by $U(x, a^*(\tilde{x}), Y, \mathbf{i}, \mathbf{t})$. We can get a closed form solution by using the covariance identity:

$$\begin{aligned}
-\frac{\partial V(\mu, \sigma, Y, \mathbf{i}, \mathbf{t})}{\partial \sigma} &= -\mathbb{E}\left\{ U(x, a^*(\tilde{x}), Y, \mathbf{i}, \mathbf{t}) \times \left[\frac{(e-\mu)^2 - \sigma^2}{\sigma^3}\right] \right\} \\
&= -\frac{1}{\sigma^3} \mathbb{E}\left\{ U(x, a^*(\tilde{x}), Y, \mathbf{i}, \mathbf{t}) \times \left[(e-\mu)^2 - \sigma^2\right] \right\} \\
&= -\frac{1}{\sigma^3} \Bigg[\mathrm{cov}\left(U(x, a^*(\tilde{x}), Y, \mathbf{i}, \mathbf{t}), (e-\mu)^2 \right) \\
&\qquad + \mathbb{E}\left\{ U(x, a^*(\tilde{x}), Y, \mathbf{i}, \mathbf{t}) \right\} \underbrace{\mathbb{E}\left\{ (e-\mu)^2 - \sigma^2 \right\}}_{=0} \Bigg] \\
&= -\frac{1}{\sigma^3} \mathrm{cov}\left(U(x, a^*(\tilde{x}), Y, \mathbf{i}, \mathbf{t}), (e-\mu)^2 \right).
\end{aligned} \tag{C.7}$$

If errors and utility are negatively correlated, then a decrease in the forecast standard deviation increases maximized utility. Since we do not observe utility like we do damages, to compute $\frac{\partial V(\mu,\sigma,Y,\mathbf{i},\mathbf{t})}{\partial \sigma}$ we will need to observe all the components of utility.

C.3 Model Assumption

The assumption in our theoretical model is that the hurricane intensity errors should be normally distributed. Figure C.1 plots the empirical distribution of wind speed and precipitation errors. Both appear to be roughly normal, although with a slight right skew indicating that the average forecast slightly underestimates intensity. For wind the average error is only 0.08 m/s.

Figure C.1: The Distribution of Realized Wind Speeds and Precipitation.

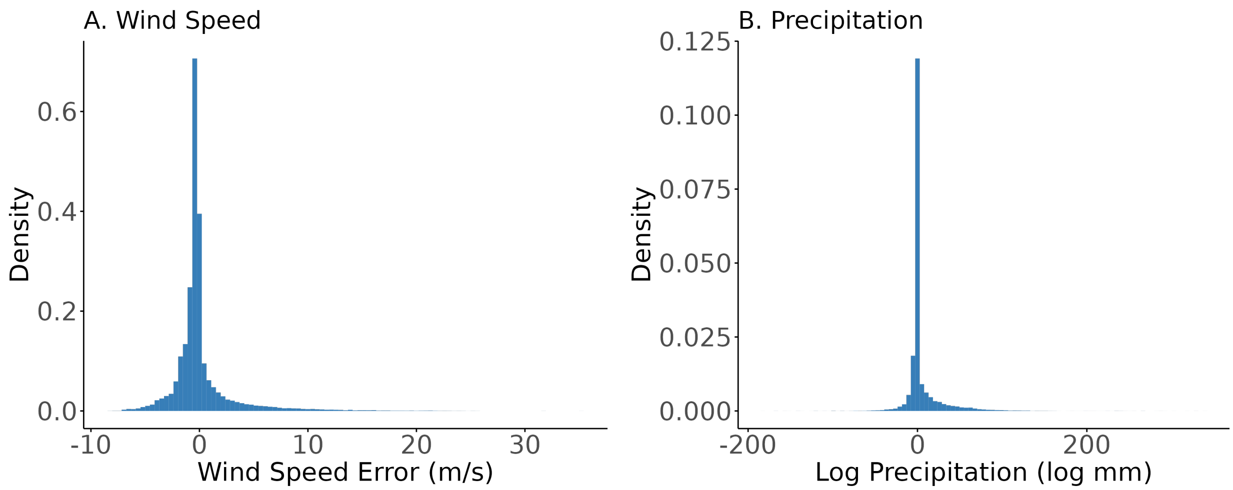

Note: Panel A shows the observed distribution of the realized wind speed error by county-hurricane. Panel B shows the observed distribution of the realized precipitation error by county-hurricane.

D Robustness Checks

Here we show robustness checks for all three sets of results: how forecasts affect emergency protective spending by FEMA, how forecast errors drive damages and emergency recovery spending by FEMA, and the value of a forecast improvement. The three sets of robustness checks largely follow the same pattern. First, we show robustness to different sets of fixed effects in tabular form. Second, we show robustness to a more conservative, 600 km spatial cutoff for computing Conley standard errors. This allows for spatial correlation over an area as large as the entire Southeastern US.[13] Third, we show our results are robust to dropping what we call "error counties." These are counties that are in a state that issued a presidential disaster declaration but did not have any reported damages in SHELDUS. Fourth, we show the results still hold even when focusing only on Atlantic and Gulf Coast states instead of the entire US. Fifth, we show robustness to the inverse hyperbolic sine transformation of the outcome variables for protective spending, damages, and recovery spending. Sixth, we show that forecast underestimate impact results are robust estimating impacts on property, crop, and mortality damages separately. Seventh, we show results for precipitation which tend to be less robust than wind speed. Finally, we show some additional robustness checks for the value of a forecast improvement results.

D.1 Does FEMA Respond to Forecasts?

Table D.1 presents estimates of the effect of the forecast wind speed and precipitation on before-landfall protective FEMA spending. Our binned estimates in Figure 2 are highly convex, so we include a quadratic term here to capture the convexity. The first column corresponds to the fixed effects in our main results. The second column adds date-of-year fixed effects to more finely address common time trends. The third column adds county-by-month-of-year fixed effects to account for potential location-specific seasonality. The fourth column adds county-by-year effects to flexibly account for variables trending over time but differentially across counties. The fifth column adds all three of these additional fixed effects. Consistent with Figure 2, we find that given a sufficiently high wind forecast, protective spending is increasing and convex in the forecast. Precipitation estimates are inconsistent and either small or noisy.

Figure D.1 increases our Conley cutoff to 600 km, allowing for spatial correlation over an area over twice as large. This has little effect on our standard errors.

[13]For example, the area of Alabama, Florida, Georgia, Kentucky, Mississippi, North Carolina, South Carolina, Tennessee, Virginia, and West Virginia is 1,026,065 km^2 while the area traced out by our 600 km cutoff is 1,130,972 km^2.

Figures D.2 and D.3 replicate Figure 2 but where we drop "error counties" or only include Atlantic and Gulf Coast states. These different sample restrictions have essentially no effect on our results.

Figure D.4 shows protective spending results, for wind speed and precipitation, when using an inverse hyperbolic sine transformation. Using this alternative outcome, we still find that forecasts of higher wind speeds spur more protective funding. This functional form also suggest greater precipitation forecasts increase protective funding as well.

Figure D.5 replicates Figure 2 but for precipitation. The plots show mixed, noisy results. Given the lack of a clear relationship between precipitation forecasts and emergency spending, we may not expect to find consistent effects for precipitation forecast errors or for the value of improving precipitation forecasts.

Table D.1: The Effect of Forecast Attributes on Before-Landfall FEMA Protective Spending.

	(1)	(2)	(3)	(4)
Protective Spending (million \$)				
Wind Forecast (m/s)	-0.4918**	-0.5093**	-0.4631*	-0.4529**
	(0.2236)	(0.2201)	(0.2674)	(0.2057)
Wind Forecast^2	0.0372***	0.0384***	0.0336**	0.0323***
	(0.0124)	(0.0120)	(0.0142)	(0.0115)
Precip Forecast (mm)	-0.0255	-0.0279	-0.0207	-0.0192
	(0.0374)	(0.0364)	(0.0273)	(0.0169)
Precip Forecast^2	-0.0001	-0.0001	-0.0000	-0.0001
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Protective Spending / GDP (%)				
Wind Forecast (m/s)	-0.0094**	-0.0104***	-0.0085**	-0.0103***
	(0.0039)	(0.0040)	(0.0039)	(0.0040)
Wind Forecast^2	0.0007***	0.0008***	0.0007***	0.0008***
	(0.0003)	(0.0003)	(0.0003)	(0.0002)
Precip Forecast (mm)	-0.0013**	-0.0012*	-0.0014**	-0.0014**
	(0.0006)	(0.0006)	(0.0006)	(0.0007)
Precip Forecast^2	0.0000*	0.0000	0.0000**	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Protective Spending / Person (\$)				
Wind Forecast (m/s)	-5.6281*	-6.1043*	-5.7773	-6.4521*
	(3.2130)	(3.2108)	(3.6887)	(3.3682)
Wind Forecast^2	0.3581**	0.3742***	0.3957**	0.4186***
	(0.1494)	(0.1423)	(0.1781)	(0.1533)
Precip Forecast (mm)	-0.1255	-0.0467	-0.3615	-0.2890
	(0.4442)	(0.4979)	(0.3830)	(0.4538)
Precip Forecast^2	0.0005	-0.0002	0.0016	-0.0004
	(0.0025)	(0.0029)	(0.0020)	(0.0035)
Observations	55,350	55,350	55,350	55,350
State-Hurricane FE	✓	✓	✓	✓
County FE	✓			
County-Month of Year FE		✓		✓
County-Year FE			✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

Figure D.1: FEMA Protective Spending Responses to Forecasts: 600 km Conley Cutoff.

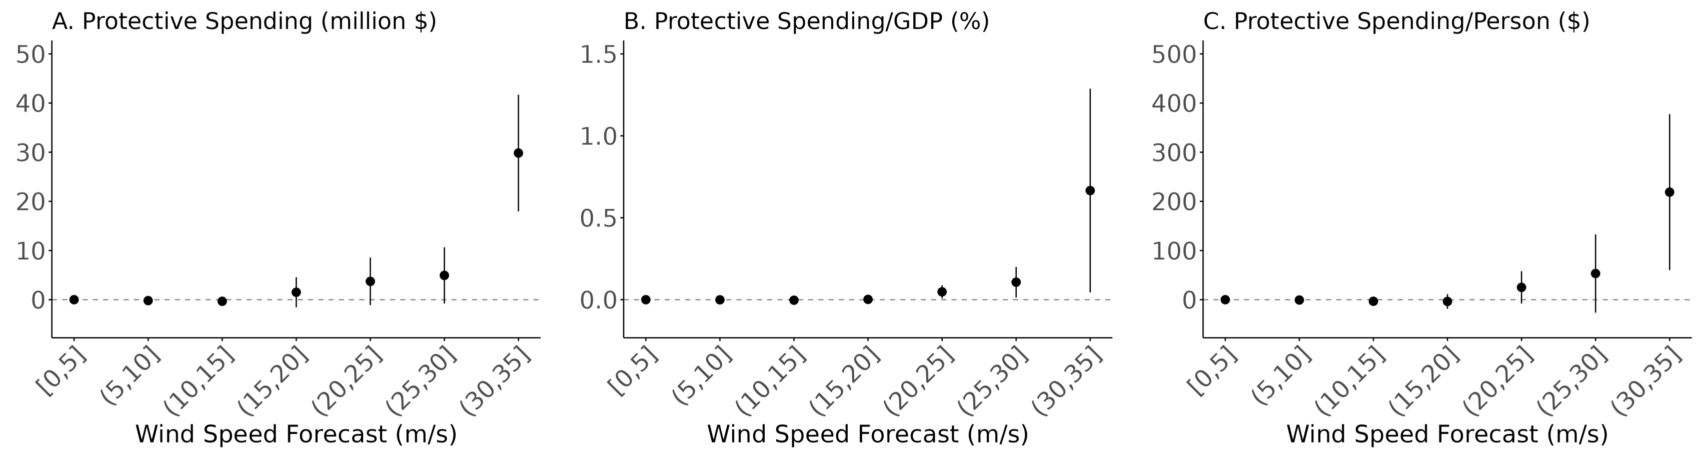

Note: The estimates correspond to equation (1). Points are point estimates and the bars are the 95% confidence intervals. The omitted category for each panel is [0,5]. All panels control for bins for the precipitation forecast, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 600 km for spatial correlation and arbitrary autocorrelation within counties.

37

Figure D.2: FEMA Protective Spending Responses to Forecasts: PDD Robustness.

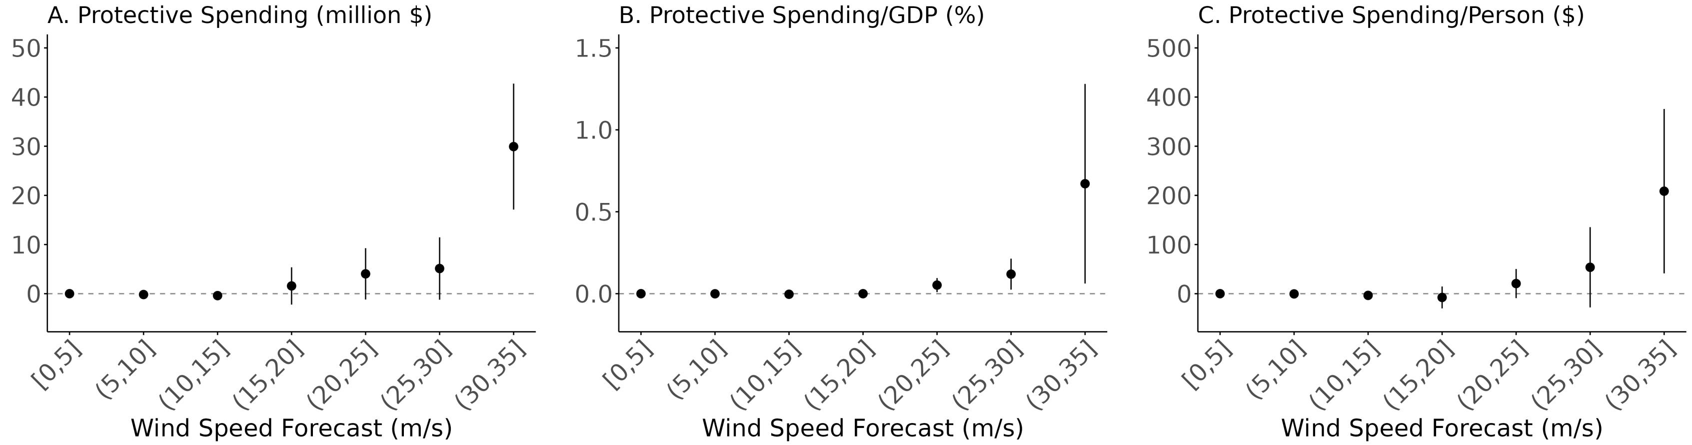

Note: The estimates correspond to equation (1). Points are point estimates and the bars are the 95% confidence intervals. The omitted category for each panel is [0,5]. All panels control for bins for the precipitation forecast, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties. The plots drop all "error counties" with a PDD but zero damage.

38

Figure D.3: FEMA Protective Spending Responses to Forecasts: Coastal States.

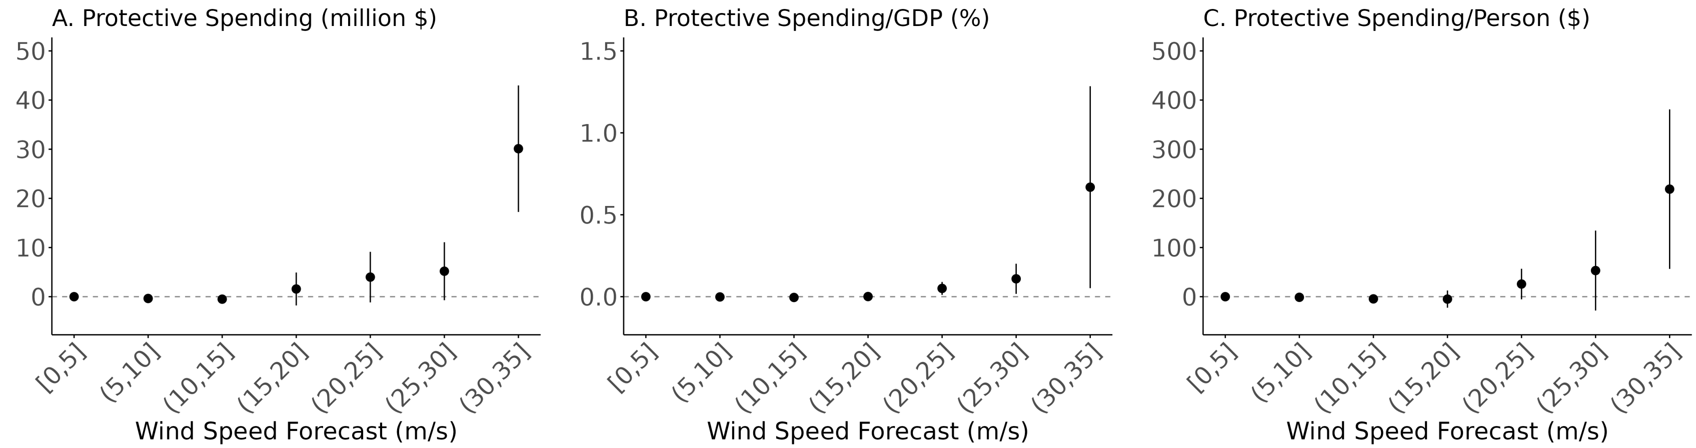

Note: The estimates correspond to equation (1). Points are point estimates and the bars are the 95% confidence intervals. The omitted category for each panel is [0,5]. All panels control for bins for the precipitation forecast, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties. Only the following states are included in the sample: Texas, Louisiana, Mississippi, Alabama, Georgia, Florida, South Carolina, North Carolina, Virginia, Maryland, New Jersey, Pennsylvania, Connecticut, Delaware, New York, Rhode Island, Massachusetts, New Hampshire, and Maine.

39

Figure D.4: FEMA Protective Spending Responses to Forecasts: Inverse Hyperbolic Sine.

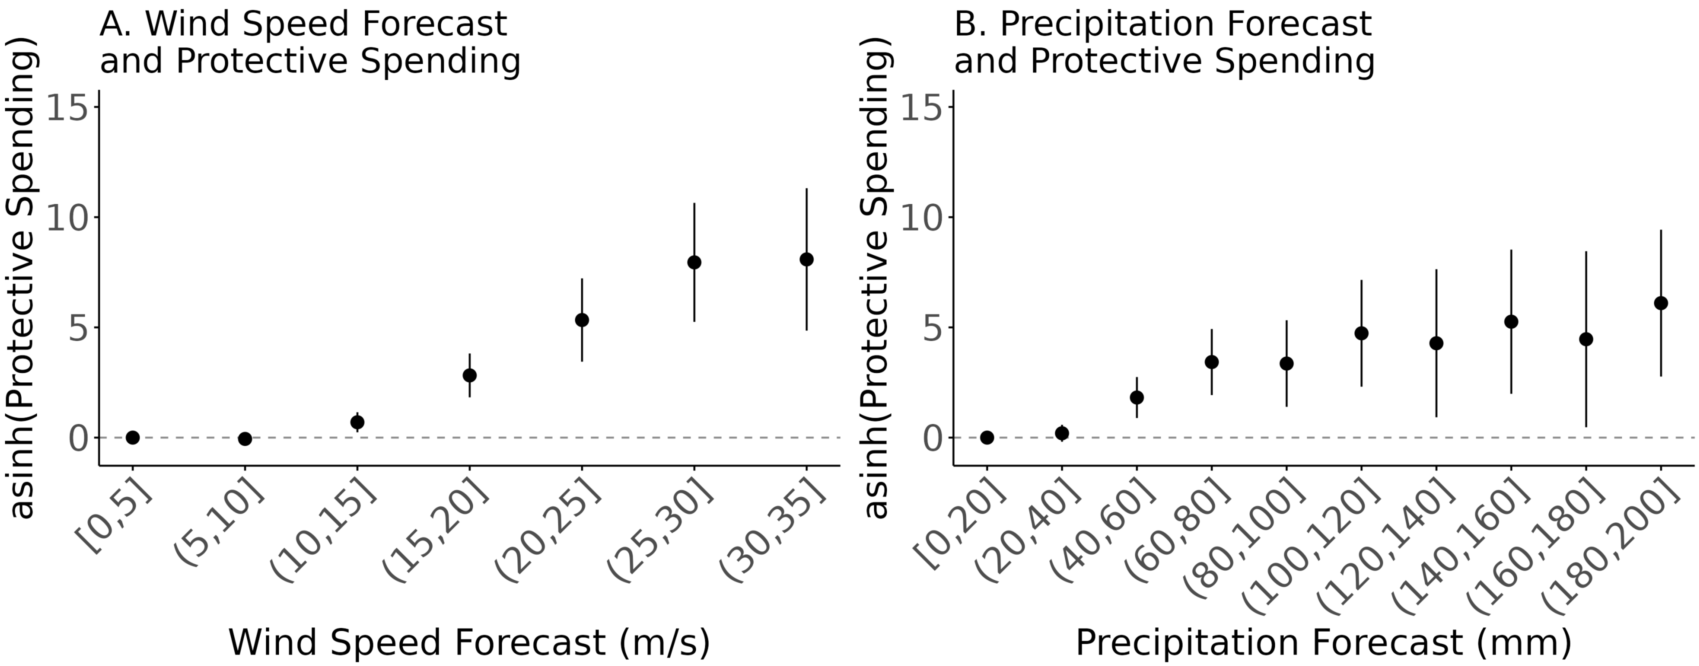

Note: The estimates correspond to equation (1). Points are point estimates and the bars are the 95% confidence intervals. The omitted category for wind speed is [0,5] and for precipitation is [0,20]. The estimates from both panels are from a single regression. All panels control for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

40

Figure D.5: FEMA Protective Spending Responses to Forecasts: Precipitation.

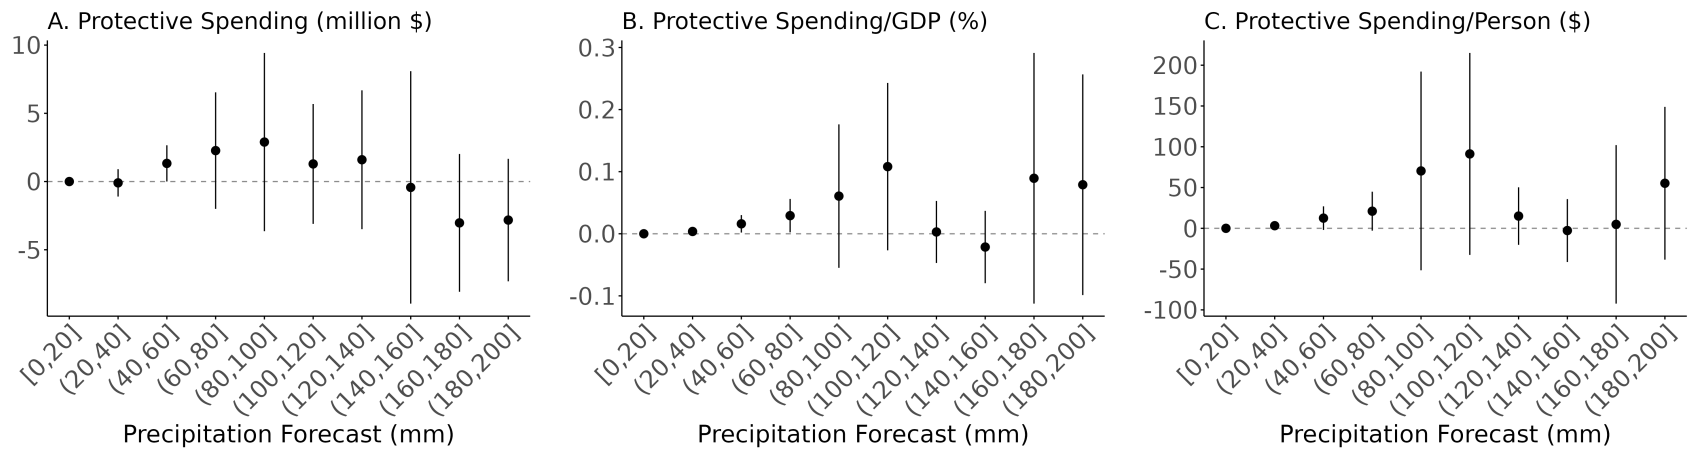

Note: The estimates correspond to equation (1). Points are point estimates and the bars are the 95% confidence intervals. The omitted category for each panel is [0,20]. All panels control for bins for the wind speed forecast, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

41

D.2 Does Forecast Accuracy Matter?

Table D.2 presents estimates of the effect of the forecast errors on damages and recovery spending. We sum the two after-landfall costs together to have more concise results. The columns correspond to the same sets of fixed effects as in Table D.1. All 20 specifications show that wind speed underestimates increase damages conditional on the realization of wind speed and precipitation. These damages are substantial: for a 1 m/s worse underestimate in a county, costs increase by almost two thousand dollars per person, or $70 million. Precipitation estimates are noisy and the sign of the effect changes depending on the specification.

Table D.3 reports the same estimates as Table D.2 but where we also interact the wind speed forecast error with an indicator variable for whether the wind speed was hurricane-force, or sub-hurricane force. This tests whether errors are more costly for higher-intensity storms. Across all specifications the interaction terms are positive: forecast errors are more costly for hurricane-force winds than for sub-hurricane-force winds.

Figure D.6 shows that these estimates are robust to the more conservative 600 km Conley spatial cutoff. Figures D.7 and D.8 show our wind speed forecast error results are robust to dropping "error counties" or only including Atlantic and Gulf Coast states, while Figure D.9 shows the wind speed results are robust to using an inverse hyperbolic sine transformation. Figure D.10 shows that wind speed forecast underestimates increase all of property damage, crop damage, and mortality damage independently in addition to increasing the aggregate cost. The plot makes clear that aggregate damage is driven by property losses.

Figure D.11 replicates our main results but for precipitation. Precipitation shows no strong pattern. This is consistent with our finding that precipitation forecasts do not have a consistent effect on protective expenditures. This may be because hurricane strength has historically been communicated through its wind speed (Kantha, 2006; Murnane and Elsner, 2012).

Table D.2: The Effect of Underestimating Wind and Precipitation on Damages and FEMA Recovery Spending.

	(1)	(2)	(3)	(4)
Damages + Recovery Spending (million \$)				
Wind Forecast Underestimate (m/s)	69.62**	72.50**	77.19*	79.00**
	(32.15)	(31.15)	(44.08)	(39.38)
Precip Forecast Underestimate (mm)	0.87	0.91	0.10	0.91
	(2.01)	(1.60)	(2.22)	(1.57)
(Damages + Recovery Spending) / GDP (%)				
Wind Forecast Underestimate (m/s)	5.20**	5.18**	5.23**	5.09***
	(2.50)	(2.19)	(2.42)	(1.92)
Precip Forecast Underestimate (mm)	-0.10	-0.10	-0.12	-0.13
	(0.09)	(0.08)	(0.09)	(0.09)
(Damages + Recovery Spending) / Person (\$)				
Wind Forecast Underestimate (m/s)	1852.37**	1817.06***	2066.40**	1927.85***
	(743.75)	(647.65)	(871.00)	(667.25)
Precip Forecast Underestimate (mm)	-38.79	-32.74	-53.22*	-41.06*
	(24.81)	(21.89)	(30.82)	(22.83)
Observations	55,350	55,350	55,350	55,350
Realized Wind/Precip Bins	✓	✓	✓	✓
State-Hurricane FE	✓	✓	✓	✓
County FE	✓			
County-Month of Year FE		✓		✓
County-Year FE			✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

Table D.3: The Marginal Effect of Underestimating Wind and Precipitation on Damages and FEMA Recovery Spending As a Function of Hurricane Intensity.

	(1)	(2)	(3)	(4)
Damages + Recovery Spending (million \$)				
Wind Forecast Underestimate (m/s): Hurricane	164.83**	167.58***	165.22**	169.20***
	(65.13)	(53.15)	(81.80)	(62.59)
Wind Forecast Underestimate (m/s): Sub-Hurricane	-7.31	-7.02	-8.27	-10.94
	(9.71)	(7.81)	(11.08)	(9.70)
(Damages + Recovery Spending) / GDP (%)				
Wind Forecast Underestimate (m/s): Hurricane	10.43**	10.44***	9.53**	9.36***
	(4.14)	(3.64)	(3.91)	(3.05)
Wind Forecast Underestimate (m/s): Sub-Hurricane	0.85*	0.97*	0.86	0.90**
	(0.52)	(0.51)	(0.53)	(0.45)
(Damages + Recovery Spending) / Person (\$)				
Wind Forecast Underestimate (m/s): Hurricane	3842.91***	3922.66***	3918.01***	3996.15***
	(1278.34)	(1169.47)	(1478.46)	(1260.22)
Wind Forecast Underestimate (m/s): Sub-Hurricane	243.98*	271.77**	268.96*	274.54**
	(128.76)	(123.13)	(150.21)	(133.29)
Observations	55,350	55,350	55,350	55,350
Precipitation Underestimate	✓	✓	✓	✓
Realized Wind/Precip Bins	✓	✓	✓	✓
State-Hurricane FE	✓	✓	✓	✓
County FE	✓			
County-Month of Year FE		✓		✓
County-Year FE			✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

44

Figure D.6: Forecast Errors, Damages, and *Ex Post* Recovery Spending: 600 km Conley Cutoff.

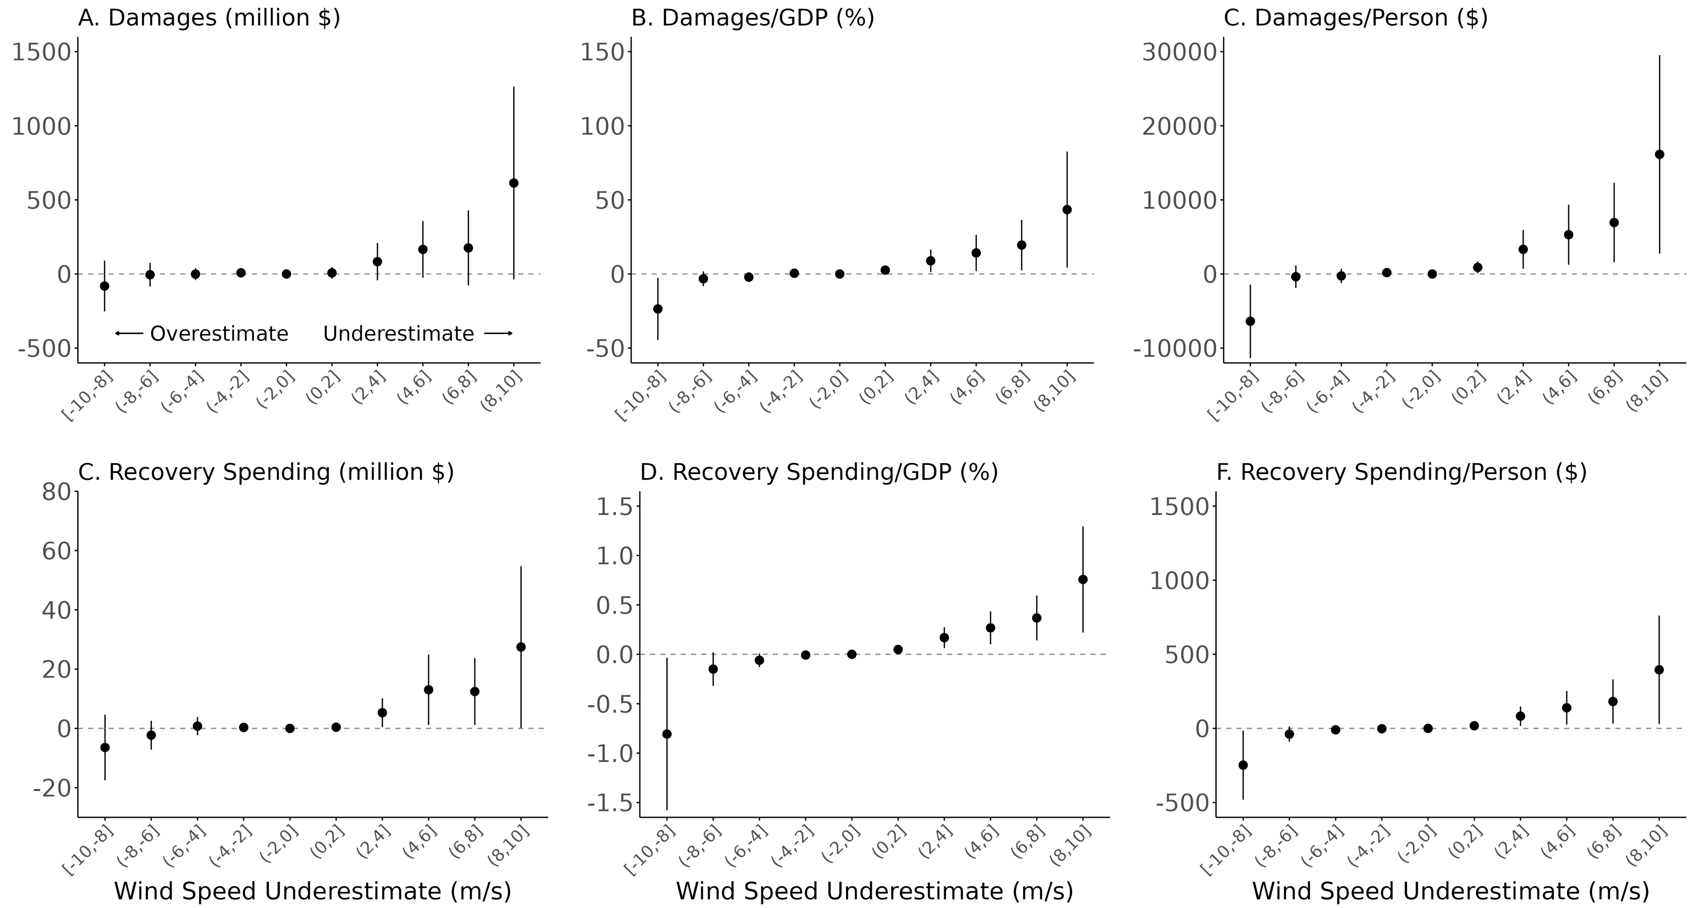

Note: The estimates correspond to equation (2). The points are point estimates, and the bars are the 95% confidence intervals. The omitted category is $(-1, 1]$. All panels control for binned precipitation errors, binned realized wind speed, binned realized precipitation, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 600 km for spatial correlation and arbitrary autocorrelation within counties.

45

Figure D.7: Forecast Errors, Damages, and *Ex Post* Recovery Spending: Coastal States.

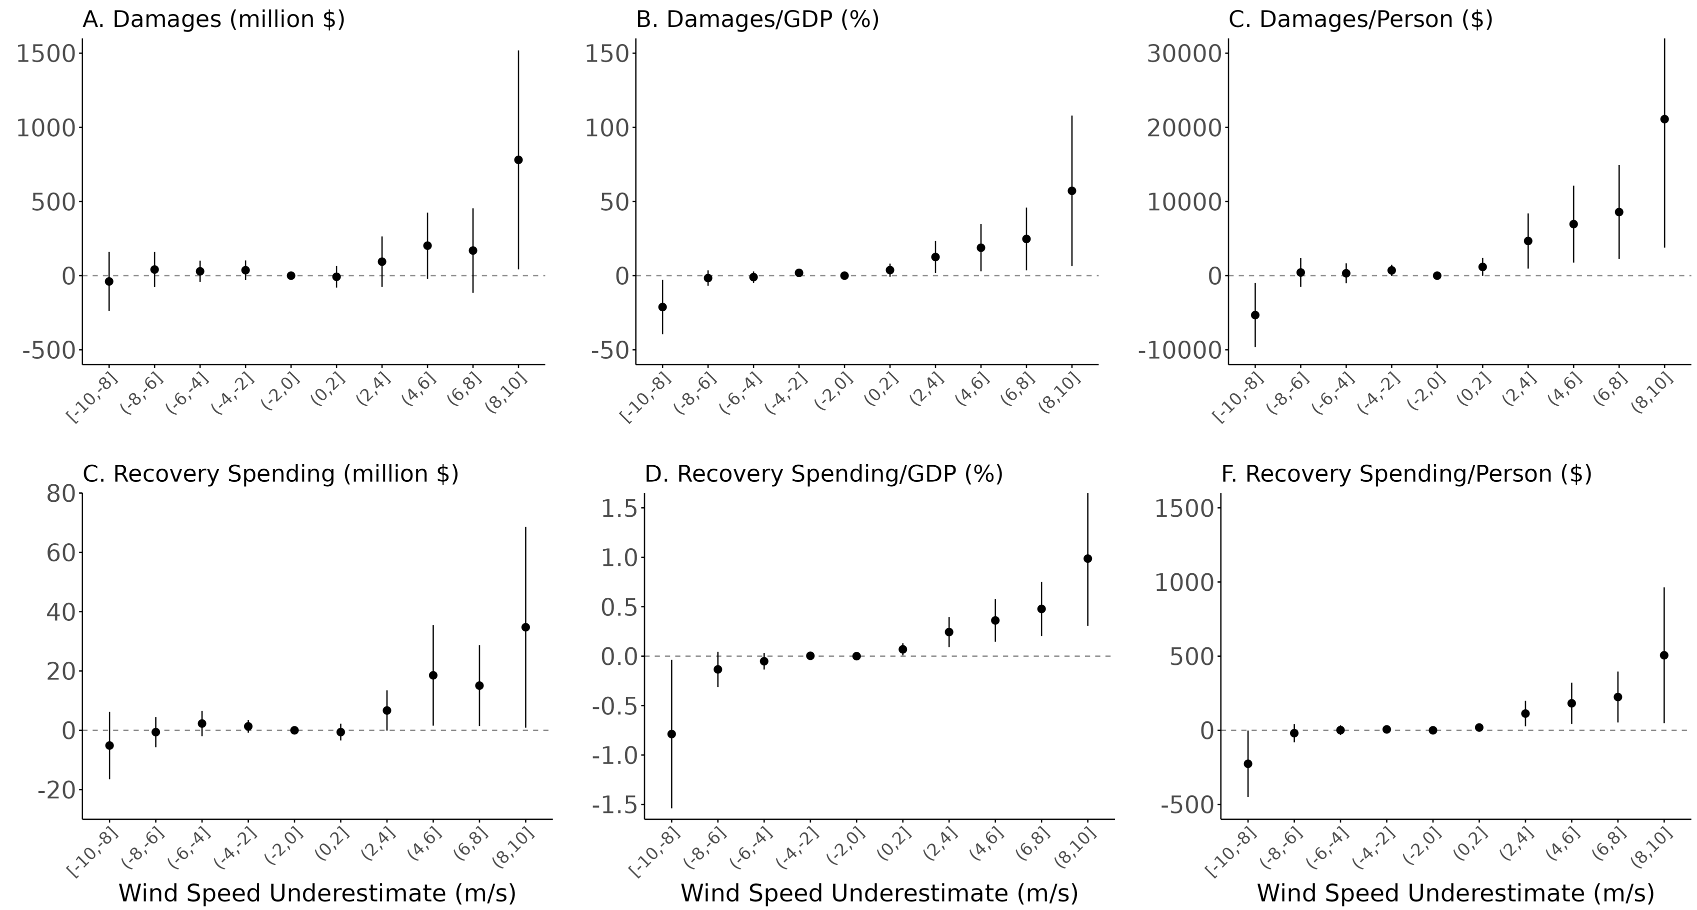

Note: The estimates correspond to equation (2). The points are point estimates, and the bars are the 95% confidence intervals. The omitted category is $(-1, 1]$. All panels control for binned precipitation errors, binned realized wind speed, binned realized precipitation, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties. Only the following states are included in the sample: Texas, Louisiana, Mississippi, Alabama, Georgia, Florida, South Carolina, North Carolina, Virginia, Maryland, New Jersey, Pennsylvania, Connecticut, Delaware, New York, Rhode Island, Massachusetts, New Hampshire, and Maine.

Figure D.8: Forecast Errors, Damages, and *Ex Post* Recovery Spending: PDD Robustness.

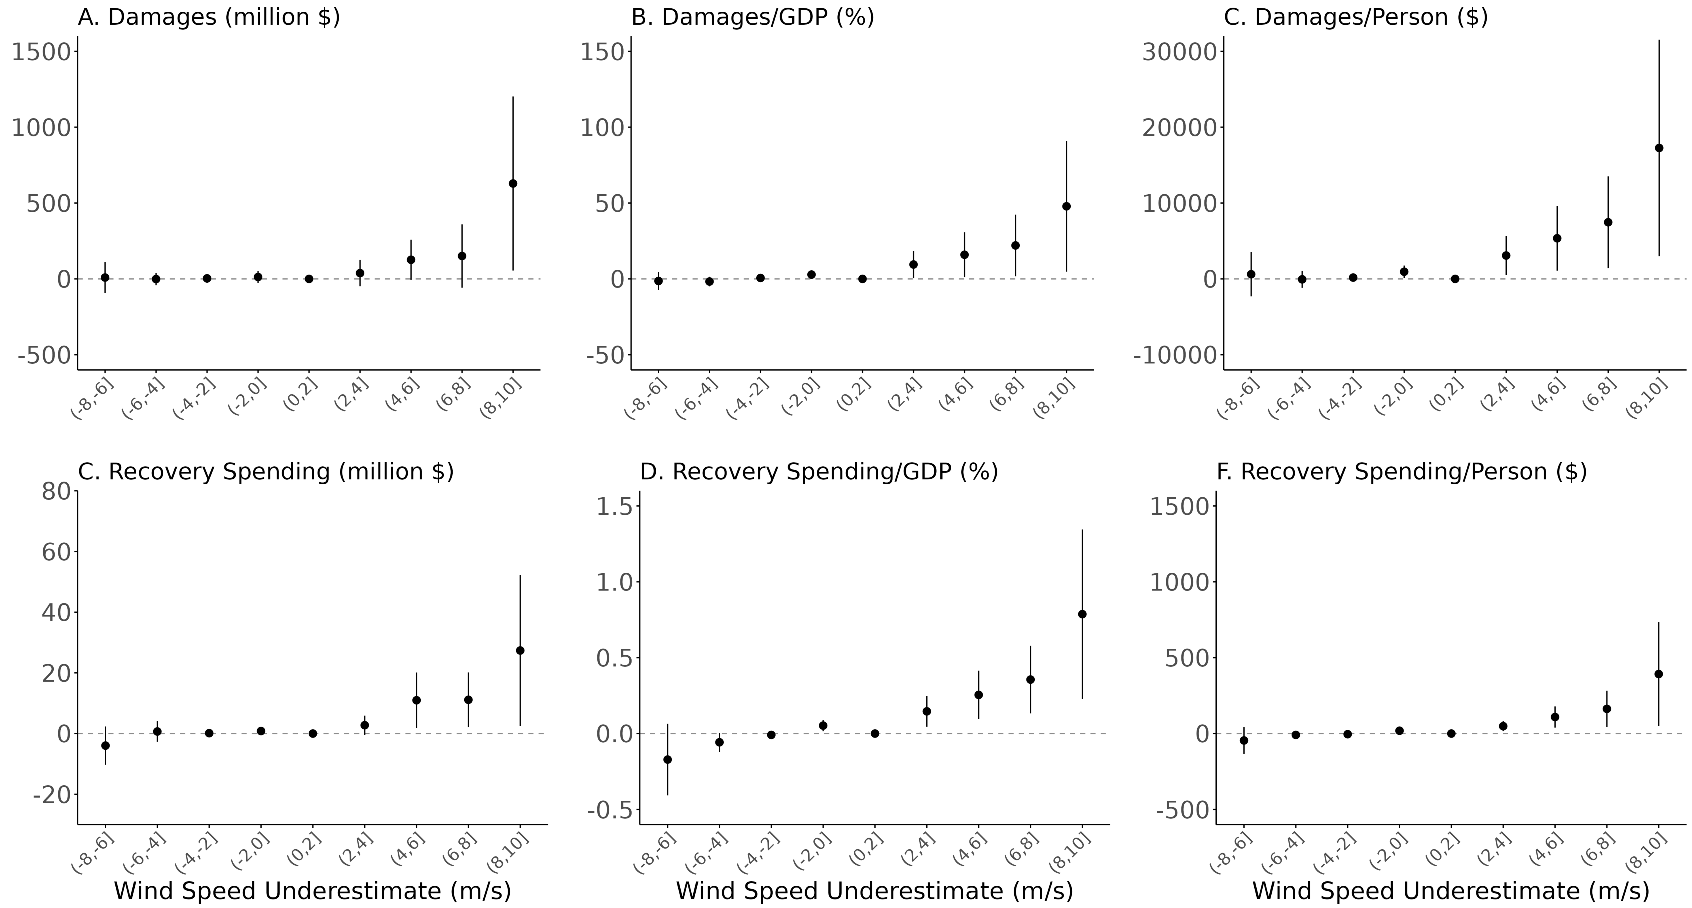

47

Note: The estimates correspond to equation (2). The points are point estimates, and the bars are the 95% confidence intervals. The omitted category is $(-1, 1]$. All panels control for binned precipitation errors, binned realized wind speed, binned realized precipitation, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties. The plots drop all "error counties" with a PDD but zero damage. Dropping error counties results in omitting the lowest bin.

Figure D.9: Forecast Errors, Damages, and *Ex Post* Recovery Spending: Inverse Hyperbolic Sine.

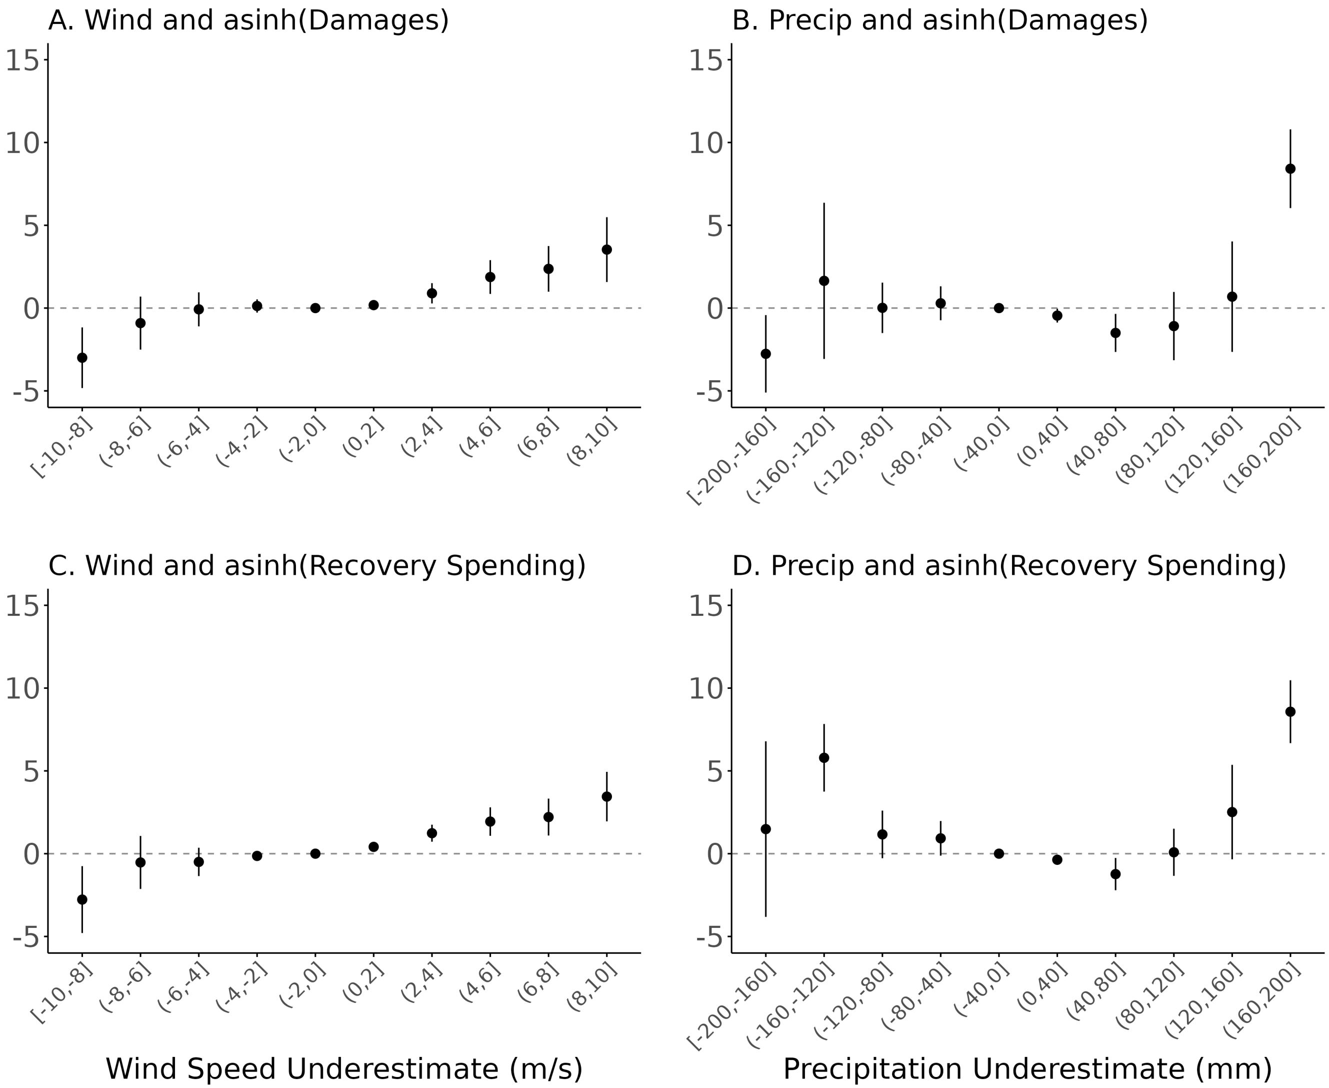

Note: The estimates correspond to equation (2). The points are point estimates, and the bars are the 95% confidence intervals. The omitted category is $(-1, 1]$ for wind speed and $(-20, 0]$ for precipitation. All panels control for binned precipitation errors, binned realized wind speed, binned realized precipitation, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

48

Figure D.10: Forecast Errors, Damages, and *Ex Post* Recovery Spending: By Damage Type.

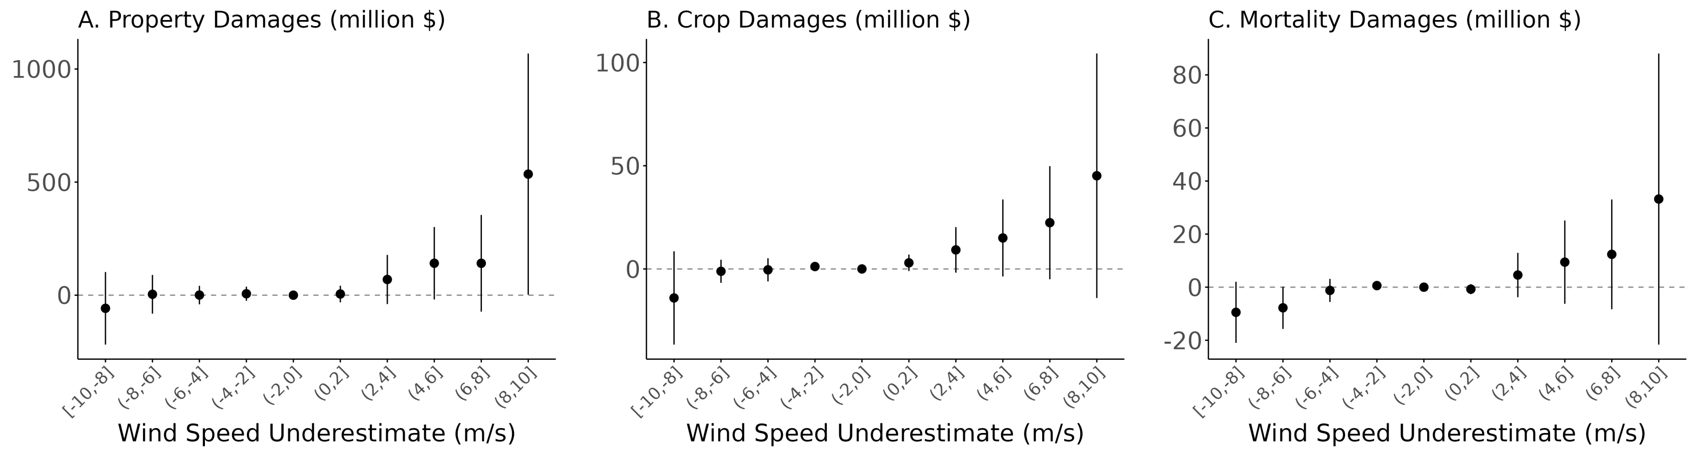

Note: The estimates correspond to equation (2). The points are point estimates, and the bars are the 95% confidence intervals. The omitted category is $(-1, 1]$ for wind speed and $(-20, 0]$ for precipitation. All panels control for binned precipitation errors, binned realized wind speed, binned realized precipitation, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

49

Figure D.11: Forecast Errors, Damages, and *Ex Post* Recovery Spending: Precipitation.

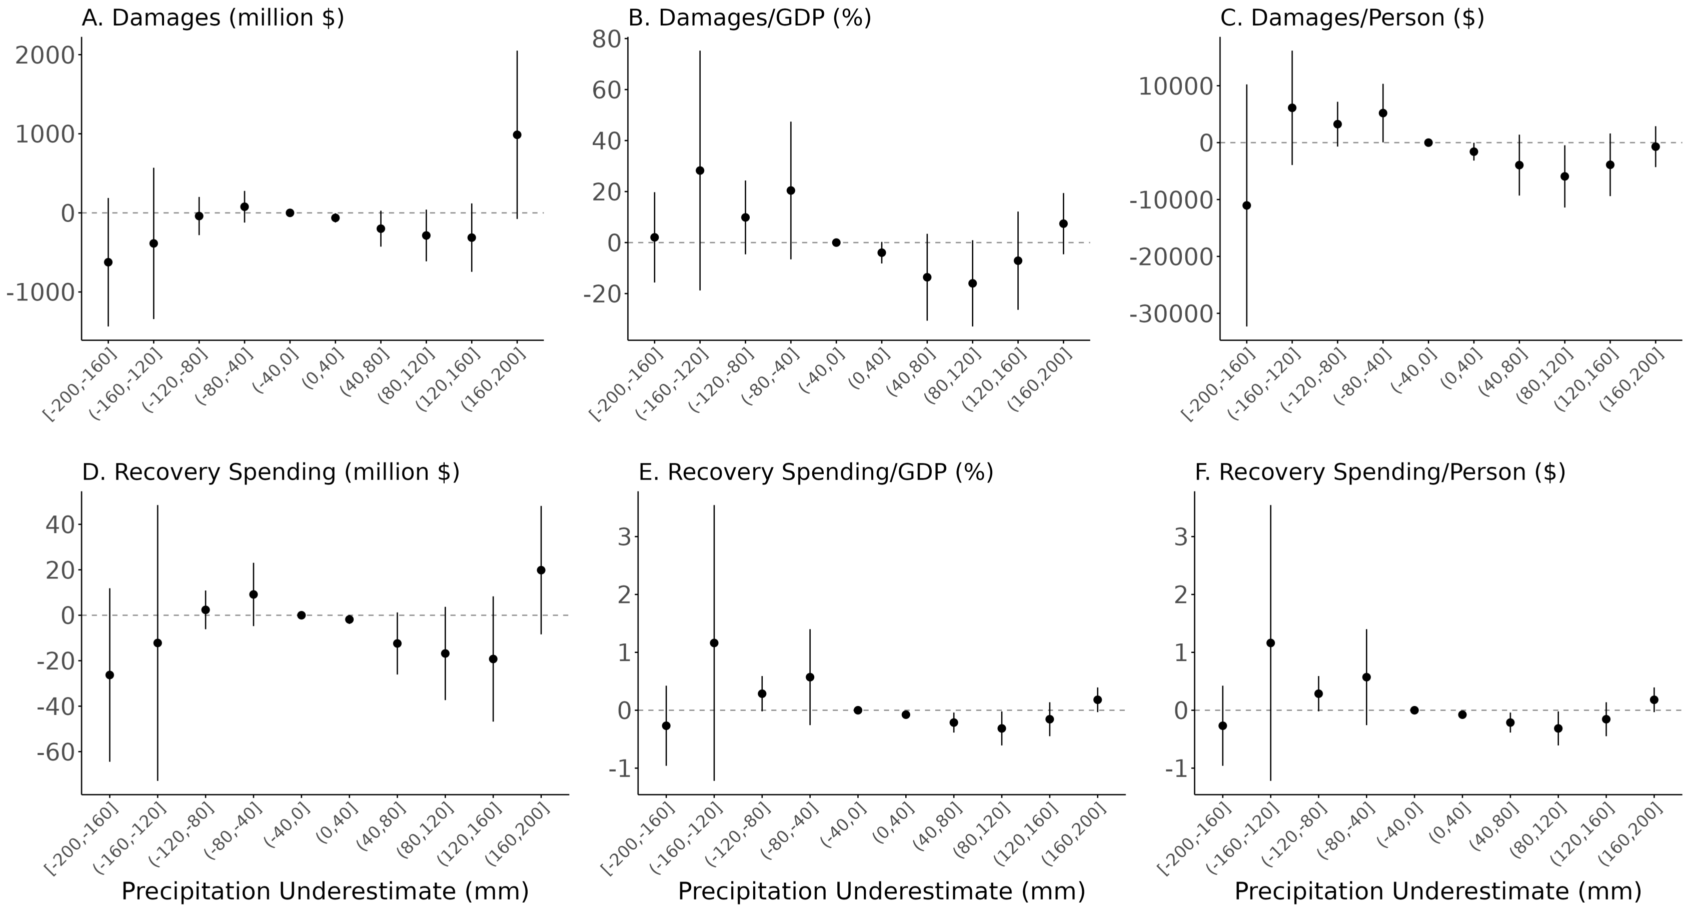

Note: The estimates correspond to equation (2). The points are point estimates, and the bars are the 95% confidence intervals. The omitted category is $(-20, 0]$. All panels control for binned wind speed errors, binned realized wind speed, binned realized precipitation, and for county and state-by-hurricane fixed effects. Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

50

D.3 What is the *Ex Ante* Value of Improving Hurricane Forecasts?

The fixed effects in each column again follow the previous sections. All specifications include binned wind speed and precipitation realizations as well as first-order forecast error terms. Table D.4 shows our results using the alternative 600 km Conley cutoff. The results are all still statistically significant. Tables D.5 and D.6 show our results are robust to the PDD and coastal county samples. Table D.7 performs the same exercise as Figure D.10 where we estimate impacts on different types of damage. The value of a forecast improvement is positive for all three, but it is driven by property damage, consistent with Figure D.10. Table D.9 shows our estimates when we do not demean the error. Not demeaning biases estimates toward zero.

Figure D.12 plots the t-statistic from the estimate in Column 3 of Table 2, but when we smoothly vary the distance cutoff for the Conley standard errors. The figure shows that our estimates are significant at the 95% level while allowing for spatial correlation up to over 1000 km away from the county centroid.

Figure D.13 shows the distribution of estimates corresponding to Column 1 of Table 2, but where we drop hurricanes from the sample, one-by-one. Most of the estimates are tightly clustered around the full sample estimate which is given by the dashed line. The large estimates are when we drop Ike and Michael, and the low estimate is when we drop Katrina.

Table D.4: The Value of a Wind Speed Forecast Improvement: Conley Robustness.

	(1)	(2)	(3)	(4)
Damages + Recovery Spending (million \$)				
$\beta_2 : (\log x - \mu)^2$	6.62**	6.52***	8.76**	8.77***
	(2.72)	(2.52)	(3.75)	(3.16)
(Damages + Recovery Spending) / GDP (%)				
$\beta_2 : (\log x - \mu)^2$	0.73**	0.74***	0.73**	0.72***
	(0.28)	(0.25)	(0.29)	(0.23)
Damages + Recovery Spending Per Capita (\$/person)				
$\beta_2 : (\log x - \mu)^2$	247.68***	260.03***	284.53***	296.52***
	(72.98)	(70.11)	(93.40)	(81.70)
Observations	55,350	55,350	55,350	55,350
Realized Wind/Precip Bins	✓	✓	✓	✓
Level Wind/Precip Error	✓	✓	✓	✓
State-Hurricane FE	✓	✓	✓	✓
County FE	✓			
County-Month of Year FE		✓		✓
County-Year FE			✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors are Conley Spatial HAC with a distance radius of 600 km for spatial correlation and arbitrary autocorrelation within counties.

Table D.5: The Value of a Wind Speed Forecast Improvement: PDD Robustness.

	(1)	(2)	(3)	(4)
Damages + Recovery Spending (million \$)				
$\beta_2 : (\log x - \mu)^2$	6.64**	5.82***	8.72**	7.40***
	(2.59)	(2.00)	(3.48)	(2.61)
(Damages + Recovery Spending) / GDP (%)				
$\beta_2 : (\log x - \mu)^2$	0.75***	0.76***	0.75***	0.72***
	(0.27)	(0.24)	(0.28)	(0.26)
Damages + Recovery Spending Per Capita (\$/person)				
$\beta_2 : (\log x - \mu)^2$	253.53***	259.41***	288.68***	249.32***
	(66.53)	(63.85)	(80.77)	(60.70)
Observations	54,613	54,613	54,613	54,613
Realized Wind/Precip Bins	✓	✓	✓	✓
Level Wind/Precip Error	✓	✓	✓	✓
State-Hurricane FE	✓	✓	✓	✓
County FE	✓			
County-Month of Year FE		✓		✓
County-Year FE			✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties. Counties issued a Presidential Disaster Declaration but without reported SHELDUS damage are dropped from the sample.

Table D.6: The Value of a Wind Speed Forecast Improvement: Coastal States.

	(1)	(2)	(3)	(4)
Damages + Recovery Spending (million \$)				
$\beta_2 : (\log x - \mu)^2$	6.60**	6.40***	8.67***	8.56***
	(2.62)	(2.42)	(3.36)	(2.93)
(Damages + Recovery Spending) / GDP (%)				
$\beta_2 : (\log x - \mu)^2$	0.74***	0.76***	0.73***	0.72***
	(0.27)	(0.24)	(0.27)	(0.22)
Damages + Recovery Spending Per Capita (\$/person)				
$\beta_2 : (\log x - \mu)^2$	253.62***	266.63***	285.23***	297.24***
	(64.69)	(64.45)	(78.91)	(73.81)
Observations	19,674	19,674	19,674	19,674
Realized Wind/Precip Bins	✓	✓	✓	✓
Level Wind/Precip Error	✓	✓	✓	✓
State-Hurricane FE	✓	✓	✓	✓
County FE	✓			
County-Month of Year FE		✓		✓
County-Year FE			✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties. Only the following states are included in the sample: Texas, Louisiana, Mississippi, Alabama, Georgia, Florida, South Carolina, North Carolina, Virginia, Maryland, New Jersey, Pennsylvania, Connecticut, Delaware, New York, Rhode Island, Massachusetts, New Hampshire, and Maine.

Table D.7: The Value of a Wind Speed Forecast Improvement by Damage Type.

	(1)	(2)	(3)	(4)
Property Damages (million \$)				
$\beta_2 : (e-\mu)^2$	5.53**	5.39**	7.58**	7.55***
	(2.43)	(2.25)	(3.21)	(2.80)
Crop Damages (million \$)				
$\beta_2 : (e-\mu)^2$	0.69*	0.69**	0.60	0.57*
	(0.39)	(0.34)	(0.39)	(0.32)
Mortality Damages (million \$)				
$\beta_2 : (e-\mu)^2$	0.20	0.22	0.28	0.32
	(0.17)	(0.16)	(0.24)	(0.22)
Observations	55,350	55,350	55,350	55,350
Realized Wind/Precip Bins	✓	✓	✓	✓
Level Wind/Precip Error	✓	✓	✓	✓
State-Hurricane FE	✓	✓	✓	✓
County FE	✓			
County-Month of Year FE		✓		✓
County-Year FE			✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

Table D.8: The Value of a Precipitation Forecast Improvement.

	(1)	(2)	(3)	(4)
Damages + Recovery Spending (million \$)				
$\beta_2 : (e-\mu)^2$	0.01	0.01*	0.01	0.02**
	(0.01)	(0.01)	(0.01)	(0.01)
(Damages + Recovery Spending) / GDP (%)				
$\beta_2 : (e-\mu)^2$	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)
Damages + Recovery Spending Per Capita (\$/person)				
$\beta_2 : (e-\mu)^2$	-0.05	-0.03	-0.09	0.10
	(0.09)	(0.06)	(0.10)	(0.12)
Observations	55,350	55,350	55,350	55,350
Realized Wind/Precip Bins	✓	✓	✓	✓
Level Wind/Precip Error	✓	✓	✓	✓
State-Hurricane FE	✓	✓	✓	✓
County FE	✓			
County-Month of Year FE		✓		✓
County-Year FE			✓	✓

* p < 0.1, ** p < 0.05, *** p < 0.01 Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties.

Table D.9: The Value of a Wind Speed Forecast Improvement without Demeaning.

	(1)	(2)	(3)	(4)
Damages + Recovery Spending (million \$)				
$\beta_2 : e^2$	3.57**	3.92***	6.60**	7.39***
	(1.48)	(1.43)	(2.81)	(2.36)
(Damages + Recovery Spending) / GDP (%)				
$\beta_2 : e^2$	0.32*	0.34**	0.46**	0.52***
	(0.17)	(0.15)	(0.21)	(0.16)
Damages + Recovery Spending Per Capita (\$/person)				
$\beta_2 : e^2$	97.91**	110.02***	163.34***	201.74***
	(43.99)	(39.66)	(58.33)	(49.40)
Observations	34,960	34,960	34,960	34,960
Realized Wind/Precip Bins	✓	✓	✓	✓
Level Wind/Precip Error	✓	✓	✓	✓
State-Hurricane FE	✓	✓	✓	✓
County FE	✓			
County-Month of Year FE		✓		✓
County-Year FE			✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors are Conley Spatial HAC with a distance radius of 400 km for spatial correlation and arbitrary autocorrelation within counties. The squared error terms are not demeaned.

Figure D.12: Conley Spatial HAC Distance Cutoff and Conley Spatial HAC T-Statistics.

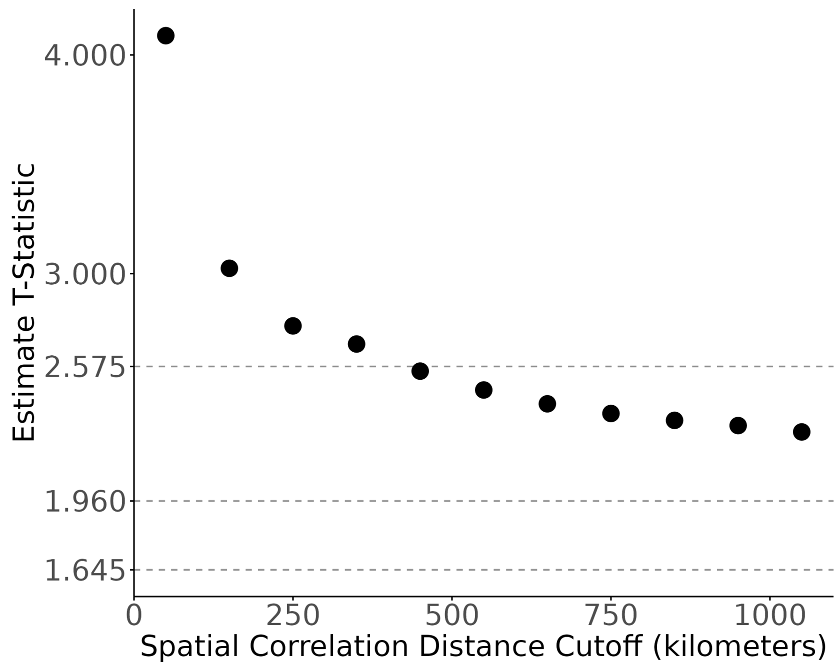

The figure plots t-statistics of the coefficient estimate from Table 2 Column 3, but using Conley (1999) standard errors that account for arbitrary autocorrelation within counties and spatial correlation up to 1,050 km in 100 km steps. Dashed lines correspond to 10%, 5%, and 1% levels of statistical significance.

Figure D.13: The Value of a Marginal Reduction in Wind Speed Forecast Uncertainty Dropping Individual Hurricanes.

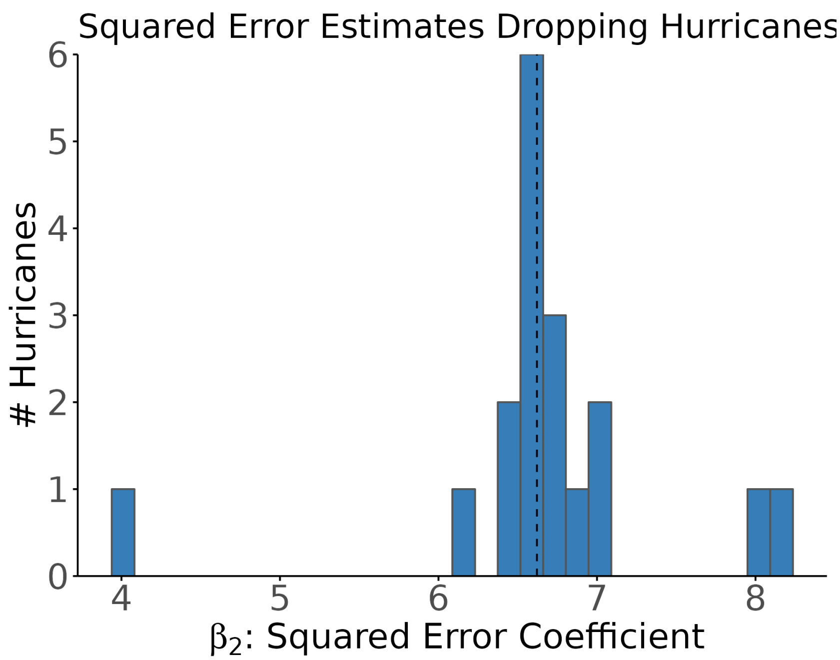

Note: The figure plots a histogram of the distribution of estimates of the value of a forecast improvement corresponding to Column 3 of Table 2 but where we drop individual hurricanes. The lowest value comes from dropping Katrina while the highest values come from dropping Ike and Michael.

E Additional Results

E.1 Correlations and Distributions

Figure E.1 presents correlations between storm and forecast attributes. Panels A and B show that higher-intensity storms tend to be under-forecast, while lower-intensity storms were over-forecast but to a lesser extent. Panels C and D show that this results in higher intensity storms having larger squared demeaned forecast errors, which is why we flexibly control for realized storm intensity in valuing forecast improvements. Panel E shows that more uncertain forecasts, in terms of the *ex ante* standard deviation, tend to result in larger forecast errors, showing why reductions in forecast standard deviations will result in more accurate forecasts ex post. Panel F shows that realized wind speed and realized precipitation are highly positively correlated. Thus, omitting one from a regression may result in omitted variable bias.

Figure E.2 plots the distribution of realizations and forecasts of wind speed in panel A and precipitation in panel B. The distributions are only over those with strictly positive values. The plots show that our data cover a large range of intensities. Most forecasts and realizations fall in the "tropical depression" category with wind speeds under 17 m/s. This is because most counties are not near the coast and end up not experiencing hurricane-force winds. However, our data do include counties experiencing wind speeds of up to 67 m/s, which would correspond to a high-end category 4 storm. Even though our dataset does not cover every hurricane, it still covers nearly the entire range of potential intensities.

Figure E.3 shows additional information about the hurricane forecast. Panel A plots the realized wind speed against the forecast wind speed using a 5 percentile binscatter. All the points are essentially on the 45 degree line: forecasts are quite accurate on average. Panel B plots the distribution of wind speed forecast errors. The average forecast error is only 0.08 m/s with a standard deviation of 2.82. The distribution is right-skewed: there are slightly more underestimates of wind speed than overestimates, likely driven by difficulties with forecasting rapidly intensifying storms.

Figure E.1: Relationships Between Different Forecast Attributes and Storm Attributes.

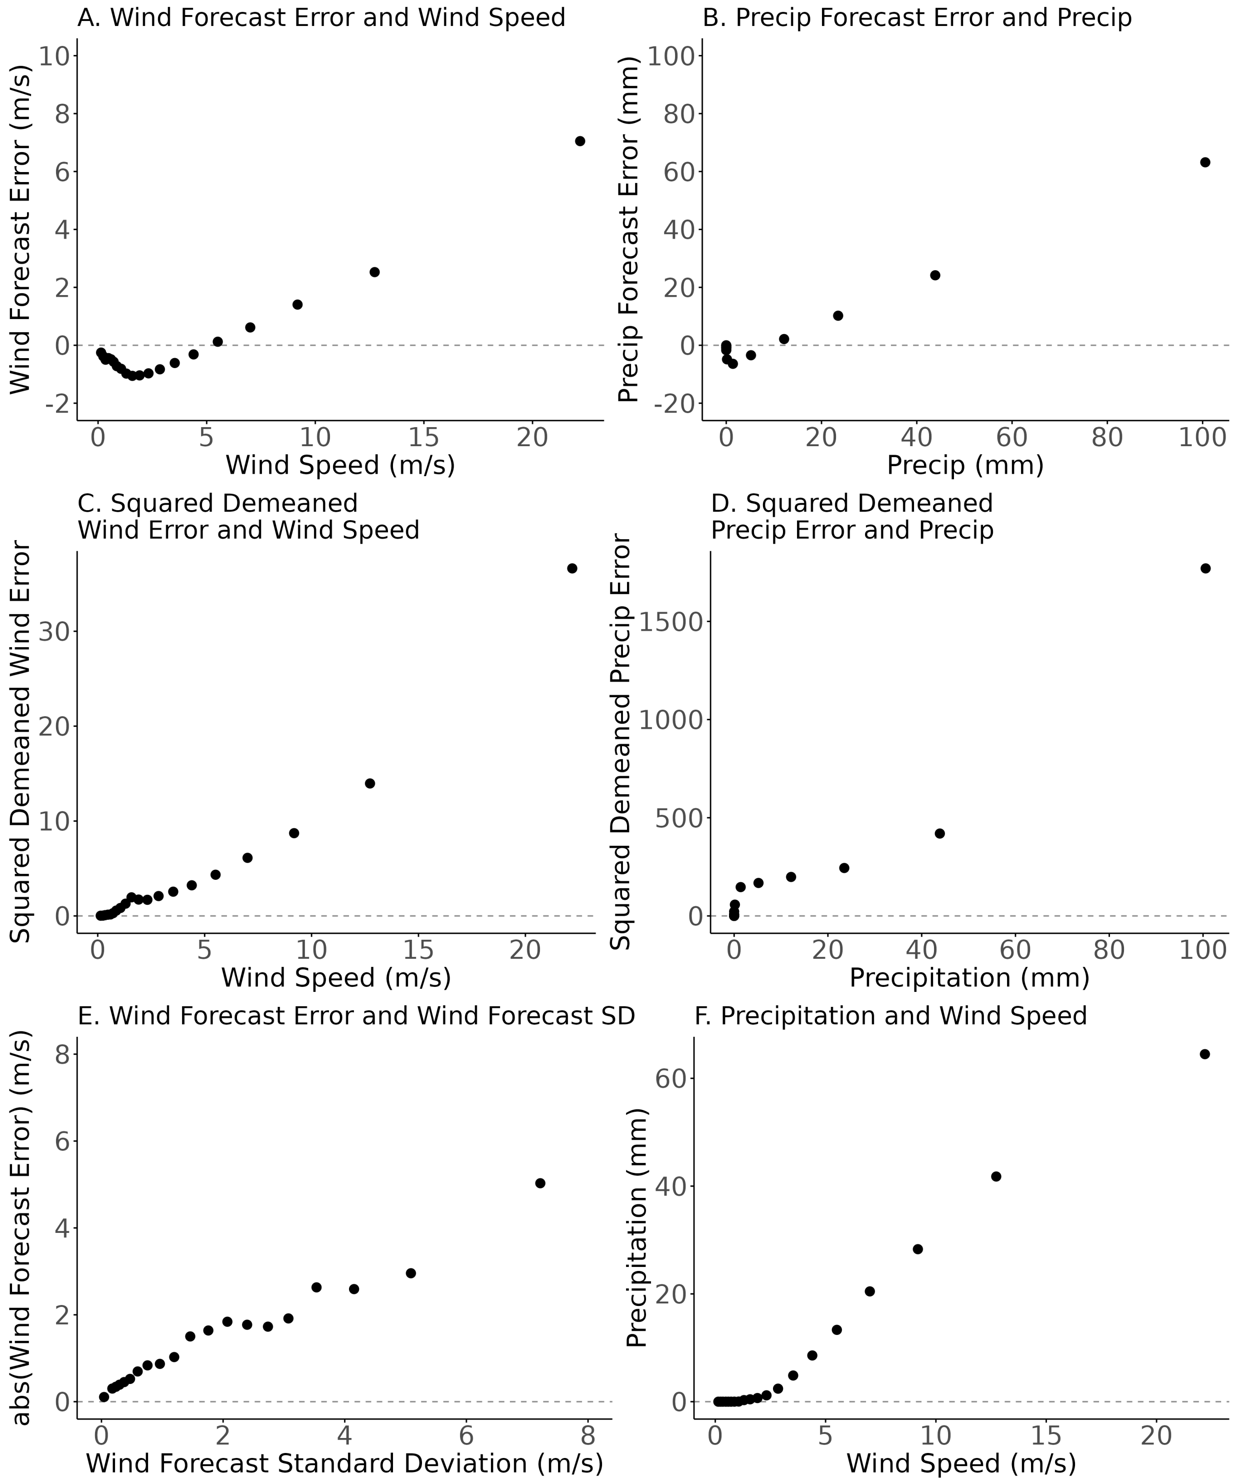

Note: Panel A plots the absolute error in the wind speed forecast (actual wind speed minus predicted wind speed) against the realized wind speed. Panel B plots the absolute error in the precipitation forecast against the realized precipitation. Panel C plots the squared demeaned error in the wind speed forecast against the realized wind speed. Panel D plots the squared demeaned error in the precipitation forecast against the realized precipitation. Panel E plots the absolute value of the wind speed forecast's error against the forecast's standard deviation. Panel F plots realized precipitation against realized wind speed. For all panels, each point is the mean of the x and y-axis variable within each viginile of the x-axis variable (i.e. a 20 bin binscatter).

Figure E.2: The Distribution of Realized Wind Speeds and Precipitation.

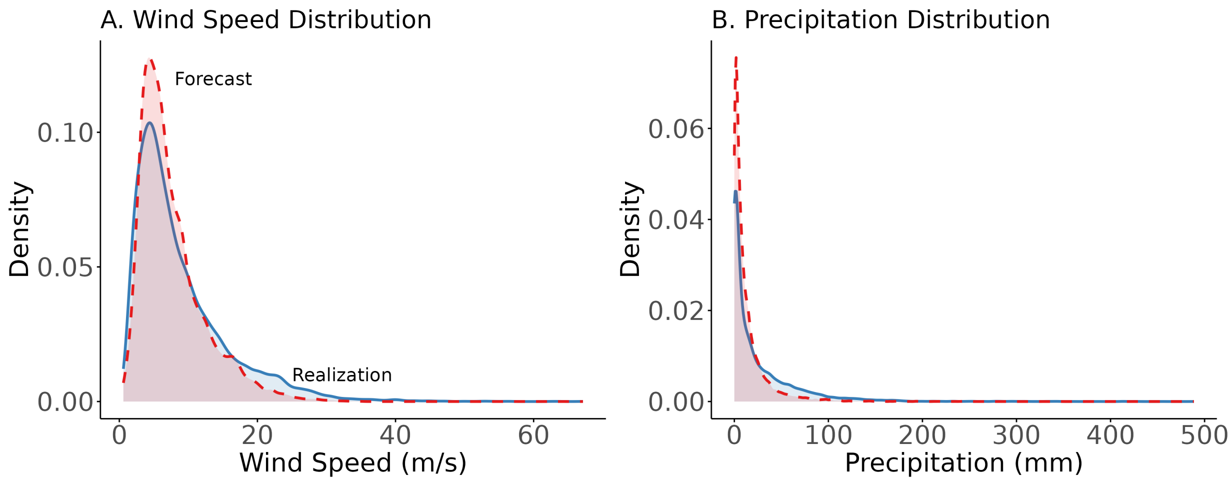

Note: Panel A shows the observed distribution of the realized and forecast wind speed by county-hurricane. Panel B shows the observed distribution of the realized and forecast precipitation by county-hurricane. The red dashed line is the distribution of the forecast and the blue line is the distribution of the realization. Values of 0 are omitted for clarity.

Figure E.3: The Distribution of Wind Speed Errors.

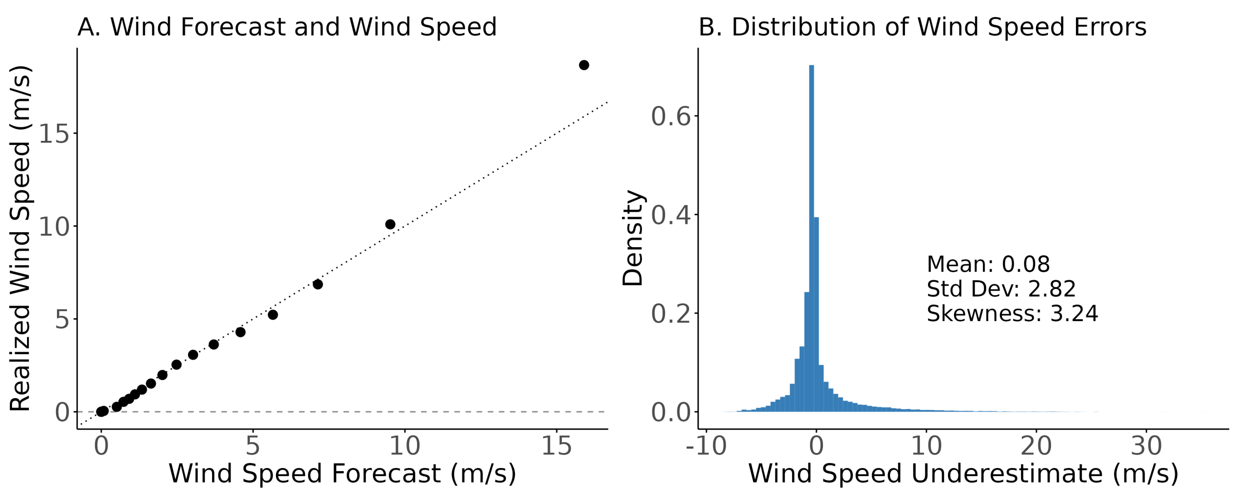

Note: Panel A plots a 20 bin binscatter of realized wind speed against the wind speed forecast. The dotted line is the 45 degree line. Panel B plots the underestimate of wind speed by a forecast. We omit observations where the forecast was for zero wind speed and the realized wind speed was zero.