Pools of rain and risk: How forecast skill affects enrollment in rainfall-indexed insurance

Anne Carroll[1] and Elinor Benami*[1,2]

[1]Department of Agricultural and Applied Economics, Virginia Tech
[2]NASA Harvest

August 12, 2025

Abstract

Index insurance, by design, mitigates some measure of adverse selection by tying payouts to objective proxies of losses (indices) rather than self-reported or manipulable measures of loss. However, when seasonal forecasts are available before enrollment deadlines, participants can rationally increase (forego) coverage in anticipation of unfavorable (favorable) conditions. Both types of selection can undermine actuarial performance. We examine this behavior in a U.S. program where payouts are determined solely by precipitation: the USDA Pasture, Rangeland, and Forage Rainfall Index Insurance Program, which is now the largest policy by acreage in the US Federal Crop Insurance Program portfolio. Using county-level data from 2017–2023, we link enrolled acreage to gridded rainfall forecasts issued before sign-up deadlines. Using variation in forecasted values and skill across space and time, we test whether enrollment responds to predicted low rainfall and whether the alignment between past forecasts and payouts accentuates or attenuates this response. We find increased enrollment following low-rainfall forecasts. When we account for forecast quality (measured against simulated payouts representing what would have been triggered under program rules), we find that past forecast errors reduce enrollment, but their effect is smaller than that of current forecasts. Furthermore, the direction of the error matters: forecasting a payout that never occurs has a much larger negative effect on future enrollment than the positive boost from an unpredicted payout. Finally, last year's payout amount has a stronger influence on enrollment than the current year's forecast. Taken together, these findings suggest that forecast availability enables strategic enrollment behavior, allowing potential policyholders to adjust their coverage in ways that exploit (somewhat) predictable payout patterns. Pricing and design should account for forecast skill and intertemporal adverse selection to maintain actuarial soundness and program sustainability.

Keywords seasonal forecasts, adverse selection, agricultural risk management, ENSO
JEL Codes D82, G22, Q12, Q14, Q18

*The authors thank the members of the Human Environment CompuTing and Ag (HECTA)Lab, Jennifer Ifft of Kansas State University, and Jisang Yu of Korea University for their constructive input on drafts of this work. All remaining errors are our own. We also acknowledge and are grateful for support from the Virginia Tech College of Agricultural and Life Sciences as well as the NASA Harvest Program (Award# 80NSSC23M0032).

1 Introduction

Seasonal climate forecasts have become an increasingly important input in economic decision-making, with evidence of their fingerprint across multiple sectors as well as broader financial markets (Carriquiry and Osgood 2012; Lemoine and Kapnick 2024; Nadolnyak and Vedenov 2013; Yu and Hendricks 2020). Their growing skill and accessibility raise the question: how does forecast information influence participation in insurance markets directly tied to weather conditions? This knowledge matters because the availability of skillful seasonal forecasts can influence coverage decisions, with implications for both program performance and financial sustainability.

The Pasture, Rangeland, and Forage Rainfall Index Insurance program (henceforth PRF) in the United States provides an important setting to examine this question. Since its launch in 2007, the PRF program has expanded steadily year over year to now become the largest program in the U.S. federal crop insurance portfolio by acreage. Indexed solely on precipitation data, PRF relies on gridded rainfall data to set premiums and activate payouts. This indexed design aims to offer comparatively quick payments to producers (relative to claims-based processes) by tying payments to observable adverse conditions that closely approximate the loss-inducing phenomenon of interest but are outside the control of either insurers or producers.

Even though the records of precipitation used to trigger payouts may not be (easily) manipulable,[1] in a world with skillful forecasts, the timing and scale of enrollment decisions can be. As one important case of this, the El Niño–Southern Oscillation (ENSO) is a major source of seasonal climate predictability for much of the United States. Driven by persistent sea surface temperature anomalies in the equatorial Pacific, its warm (El Niño) and cool (La Niña) phases shift atmospheric circulation in ways that make regional rainfall anomalies forecastable up to nine months in advance (Ehsan et al. 2024). While forecast skill varies over space and time,[2] in some PRF-relevant areas, it may be high enough to influence coverage decisions. This opens the door for producers to selectively join the program when poor rainfall is anticipated or forego coverage when more rainfall is expected. Because the rating process does not account for forecasts, either form of selection can distort the risk pool, raise costs for participants, drain federal resources due to the program's high subsidy rate, and, ultimately, undermine program sustainability.

This paper seeks to contribute to three main strands of economics literature. First, we build on theory-focused work modeling the demand for index insurance under basis risk and the availability of reliable forecasts. Prior models have shown how basis risk and forecast skill shape producer behavior, including insurance uptake and agricultural input usage (Carriquiry and Osgood 2012; Clarke 2016; Yu and Hendricks 2020), and early works on PRF's pilot years model how seasonal forecasts can both reduce producer exposure to climate risk *and* enable adverse selection in a selected set of Alabama counties (Nadolnyak and Vedenov 2013). Our analysis extends this work by broadening the spatiotemporal scope to span the diverse agroclimatic conditions of the continental US and the recent, extended period of major PRF program growth, when producer familiarity and participation patterns are meaningfully larger than the pilot phase of the program. We also seek to blend this literature with development economics literature on the differential impacts of forecast over- and under-prediction on behavior, which Lybbert et al. (2007)

[1] A notable exception occurred when some policyholders damaged rain gauges to manufacture agricultural insurance payouts, documented in the news *Wrecked rain gauges. Whistleblowers. Million-dollar payouts and manhunts. Then a Colorado crop fraud got really crazy* by Booth (2024).

[2] Clark and Johnson (2025) notes limited precipitation forecast skill across the US, but they have a different temporal horizon and use different benchmarks for performance than our specific policy application. Other recent work suggests that, on average, ENSO accounts for approximately 20-30% of seasonal precipitation variability observed in the US (Hobeichi et al. 2024).

suggest exhibit asymmetry due to risk aversion. Second, we seek to complement empirical studies of agricultural insurance demand. Prior work indicates more knowledge about indemnity probabilities (Cai and Song 2017) as well as recent payouts (Che et al. 2020; Stein 2016) meaningfully influence participation patterns. In the specific PRF context, producers have shifted toward insuring non-growing season periods with higher expected payouts (Goodrich et al. 2020; Goodrich and Davidson 2024), illustrating the potential for strategic financial decision-making behavior. Our analysis integrates forecasts into this empirical framework, allowing us to examine the influence of forecasts with the well-documented effects of recent payouts while building on the theories of incentives for adverse selection (Just et al. 1999). Third, we aim to contribute to the literature on the economic value of seasonal weather information. While seasonal forecasts have been shown to affect expectations in financial markets (Lemoine and Kapnick 2024) and agricultural decision-making (Cabrera et al. 2007; Doherty and Thistle 1996; Hansen et al. 2011; Meza et al. 2008), most studies rely on meteorological skill metrics or market-wide indicators to evaluate performance. By focusing on a rainfall-indexed insurance program, we evaluate the value of seasonal forecasts in a single-peril setting where payouts hinge directly on the forecasted variable, isolating a specific channel through which seasonal climate data can have direct financial (and fiscal) impacts.

In particular, this work links administrative data on rainfall index insurance enrollment (2017 – 2023) from across the continental US with spatially explicit, pre-enrollment ENSO-based forecasts to test whether and how forecasted precipitation influences insurance adoption. We quantify responses to both the probability and magnitude of forecasted rainfall anomalies and assess whether forecast quality—measured as the correspondence with rainfall records and, separately, payout determinations—amplifies or attenuates these estimated effects. The results provide empirical evidence on the extent of forecast-driven selection in a major U.S. index insurance program, with implications for the design of climate risk management tools and the role of forecast information in insurance pricing. The remainder of this manuscript is structured as follows. Section 2 introduces the background and institutional setting of the Pasture, Rangeland, and Forage (PRF) Rainfall Index Insurance program, outlining key program features relevant for our analysis. Section 3 describes the data sources and construction of variables used for our assessment. Section 4 details the empirical framework and estimation strategy. Section 5 presents the main results, including robustness checks and alternative specifications, and Section 6 discusses the implications of the findings for policy design and program evaluation.

2 Background

The PRF program seeks to aid livestock producers in managing the risk of feed losses caused by low rainfall, covering pasture, forage, and rangeland areas that account for roughly 55 percent of U.S. land area (USDA RMA 2025b). While many program features determine the magnitude of PRF payouts, we focus here on a subset most relevant to our analysis, and specifically the spatial and temporal scale at which those elements operate to activate payouts, along with the enrollment deadline that may permit intertemporal adverse selection. First, unlike traditional indemnity insurance, where producers submit claims for losses on their land, PRF is an indexed product that triggers payouts based on a single peril (anomalously low rainfall) measured over a given area (here, a grid) and time period (consecutive two-month index intervals). Second, policyholders must select at least two non-overlapping two-month "index intervals" over which to have coverage. Payouts occur when total rainfall within a selected interval falls below a threshold (typically 70–90% of the historical average for the grid(s) containing the enrolled acreage) cor-

responding to the producer's chosen coverage level (Figure 1). Third, when PRF began in 2007, it was as a pilot program offered in select U.S. counties. Since 2016, it has been offered across the continental United States. Our analysis focuses on the period beginning with the 2017 policy year through to 2023, when the most recent year of complete policy and indemnity data is available. Finally, the PRF enrollment deadline has typically fallen between mid-November to December 1, enabling producers to consider seasonal forecasts released as late as mid-November when making coverage decisions for the following year (USDA 2022; USDA RMA 2021, 2025a).[3]

3 Data and Descriptive Statistics

Our analysis relies on three primary sources of data: (1) gridded precipitation forecasts, (2) gridded precipitation observations, and (3) county-level administrative records on PRF enrollment and payouts. As the spatial resolution of each dataset differs, we harmonize them to match the county level for analysis. We also construct a forecast quality measure from the combination of the above inputs.

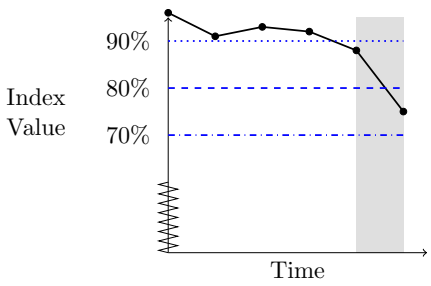


Figure 1: Conceptual diagram of index program payouts (shading) triggered when the index value dips below predefined thresholds (blue lines)

3.1 Gridded Precipitation Forecasts

We use two forecast-based variables in our analysis: the probability of below-normal precipitation and the predicted precipitation quantity.

The probability of below-normal precipitation comes from the International Research Institute for Climate and Society (IRI)'s Tercile Probability dataset, which reports monthly probabilities of below-normal, normal, and above-normal precipitation (summing to one) for lead times of up to four months. IRI generates these probabilities at a $1^\circ \times 1^\circ$ degree resolution (approximately 69×56 mile grid at the 35° latitude spanning most of the Southern U.S.) by applying extended logistic regression (ELR) to individual models from NOAA's North American Multi-Model Ensemble (NMME). For each location and forecast period, precipitation terciles (below-, near-, and above-normal) are defined from the historical distribution of observed precipitation over a baseline period, i.e., 1982–2010 through August 2021, updated to 1991–2020 thereafter (IRI 2021). Model probabilities for each tercile are then averaged and spatially smoothed to produce the final forecast (Kirtman et al. 2014).

Second, the predicted precipitation value comes from IRI's ELR-calibrated ensemble-mean precipitation product, which provides gridded forecasts of total monthly precipitation (mm) on a global $1^\circ \times 1^\circ$ grid for seasonal lead times of up to five months. Forecasts are produced by calibrating each participating NMME model using extended logistic regression (ELR) and then averaging the calibrated results to form a multi-model ensemble mean (IRI 2025).

[3]Future versions of this manuscript will include sensitivities to earlier forecasts, e.g., those starting in October, to accommodate the mid-November enrollment deadline that occurred before 2021. In the last several years, the NOAA Climate Prediction Center has maintained a seasonal production outlook release date of the 15th of every month, so it's plausible the most savvy producers would still be able to incorporate that information, although barring significant jumps in forecast performance from one month to the next at this time of year, we'd anticipate the difference to be minimal.

3.2 Gridded Precipitation Observations

To evaluate how forecasts align with policyholders' experiences in the PRF program, we use precipitation observations from the same precipitation dataset used to determine PRF payouts: the Climate Prediction Center's Unified Gauge-Based Analysis of Daily Precipitation (CPC 2025; USDA RMA 2025b). These data, generated from interpolating observations from a network of thousands of rainfall stations across the US, are publicly available on a 0.25 x 0.25 degree (approximately 17 × 17 mile grid) from 1948 to the present, providing a consistent and long-term record of observed precipitation across the continental United States.

To mirror the PRF program design, we then construct a simulated rainfall index by comparing each interval's two-month total to the historical average for the same interval and grid, calculated over the 1948 to the *t-2* period, i.e., excluding the two most recent years before the insurance year. This allows our observational measure to reflect the index used to activate PRF payouts. To integrate the gridded data with county-level administrative records, we also convert precipitation grid cell values to the county scale by computing an area-weighted average of all grid cells that overlap each county's boundaries.

3.3 County-Level Data on Enrollment and Payouts

We obtain administrative records on PRF participation and payouts from USDA's publicly available summary of business files. County-level annual and sub-annual acreage enrolled are drawn from the "Type/Practice/Unit Structure Data Files," while county-level counts of policies sold each year come from the "State/County/Crop/Coverage Level 1989–Present Data Files." Counties with zero enrollment are not always included in these datasets; accordingly, our sample is restricted to counties with nonzero enrollment in a given year. We further limit the analysis period to 2017–2023, when consistent forecast, rainfall, and enrollment records are available. U.S. county boundaries are the US Census Bureau 2022 TIGER/Line Shapefiles extracted via the *tigris* package (Walker 2025) in *R* (R Core Team 2025), with Connecticut counties sourced from 2021 due to recent changes to Connecticut administative boundaries.

3.4 Converting Gridded Data to the County Level

Because our administrative outcomes are reported at the county level, we transform all gridded precipitation forecasts and observations into county-level estimates prior to analysis. Forecasted precipitation is provided on a 1° by 1° grid (per Figure 2), while observed precipitation is available on a 0.25° by 0.25° resolution grid (∼ 21km east–west x 28km north–south, or about 13 x 17 miles). We convert forecast grid values to county-level estimates by computing area-weighted averages of all grid cells that overlap each county:

$$x_{im} = \sum_{j=1}^{n_i} w_{ji} r_{jm}, \qquad (1)$$

where x_{im} denotes the estimated rainfall in county i and month m, r_{jm} is the corresponding value for grid cell j in month m, and w_{ji} is

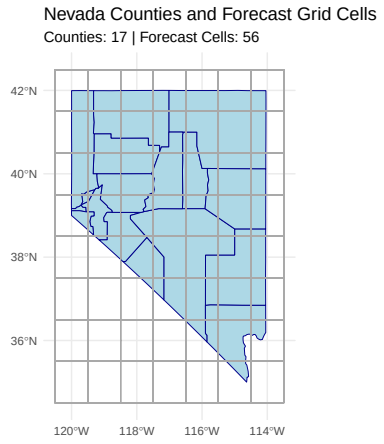


Figure 2: Forecast grid cells and county boundaries in Nevada, for scale.

the fraction of county i's area covered by grid j. For tercile probability forecasts, x_{im} represents the county-level probability of below-normal rainfall, and r_{jm} is the same probability for a grid cell. To align with the PRF program design parameters, we aggregate the gridded data into two-month totals corresponding to the duration of PRF index intervals. For observed and forecasted precipitation quantities, we sum the monthly values within each interval; for forecasted tercile probabilities, we take the mean of the monthly below-normal probabilities in each interval.

3.5 Forecasted Payout

Constructing a forecasted payout measure requires combining three pieces of data: forecasted precipitation for each PRF grid and interval, the corresponding historical rainfall baseline constructed from historical observations, and the producer-determined coverage threshold that triggers payments. In the PRF program, payouts occur when the rainfall index—i.e., the ratio of observed rainfall to its historical baseline, defined as the mean from 1948 until two years prior to the policy year—falls below the policyholder's coverage level. For comparability across policies, we fix a coverage threshold of 90%, meaning that any producer enrolled at or below this coverage level would expect to receive a payout if rainfall fell short of 90% of the baseline. Aside from being a conservative choice, this value is also the modal coverage level among enrolled producers.

Per section 3.1, the International Research Institute (IRI) provides seasonal ENSO forecasts on precipitation that are updated on a monthly basis. By mid-November, when producers make PRF enrollment decisions, probability forecasts are available through March of the following calendar (policy) year, and precipitation forecasts in millimeters are available through April. After aggregating gridded precipitation estimates to the county level using area weights (per subsection 3.4), we transform the temporal element of the data. Since our regressions are performed at the interval level, we sum forecasted precipitation over each two-month period for which the data are available at the time of decision-making (January-February, February-March, and March-April). We then compare these interval-level sums to their historical baselines to assign a binary forecasted payout indicator at the grid-interval level, equal to one if forecasted rainfall is less than 90 percent of the baseline (suggesting a payout is warranted) and zero otherwise.

3.6 Constructed Forecast Skill Metrics

In the scientific literature, ENSO forecast skill is typically assessed using metrics such as anomaly correlation coefficients and root-mean-square error on sea surface temperature indices (especially Niño 3.4 SST anomalies) over historical hindcasts spanning varying lead times and seasons Barnston et al. (2012), Ehsan et al. (2024), and Tippett et al. (2012)). These meteorological skill measures, and the emphasis on sea surface temperature, are important for scientists developing these models, but they are less relevant for our context, where the key concern is how well ENSO forecasts anticipate precipitation shortfalls that trigger payouts under PRF program parameters. To capture forecast skill in our setting, we construct several variables at the county–interval–year level. *Yesteryear Forecast Error* measures the difference between prior-year forecasted and realized precipitation in millimeters (positive = overestimate; negative = underestimate). We also define two categorical indicators of payout misprediction: *Forecasted Payout Miss* equals 1 if a payout was forecast in the prior year but none occurred; *Nonforecasted Payout* equals 1 if no payout was forecast in the prior year but realized precipitation triggered one.

3.7 Summary Statistics

Descriptions of the key variables used for analysis and descriptive statistics as well as maps characterizing these data in aggregate across the continental US are summarized as follows:

Variable (units)	Mean	Std. Dev.	Min	Max	N
Acreage enrolled (acres)	12,231	50,306	1	1,444,119	26,842
Policies sold	28.4	30.4	4	327	9,658
Payout indicator (1 if payout > 0)	0.43	0.50	0	1	26,842
Realized rainfall (mm)	124.7	103.0	0.3	941.4	26,842
Forecast absolute error (mm)	58.2	46.5	0.01	631.3	26,842
Centered forecasted rainfall (mm)	0.0	30.8	-148.0	348.73	26,842
Forecasted probability of below-normal rainfall (%)	34.9	1.4	31.8	37.3	18,901

Table 1: Descriptive statistics, averaged at the county–interval–year level across the 1,233 counties with nonzero enrollment in PRF in the continental US (2017–2023).

Variable	**Description**
Acres	Total PRF-enrolled area
Forecasted Precipitation (mm)	Forecast total, mean-centered by county and interval
Forecast Probability Below-Normal (%)	Probability that total forecasted precipitation will fall in the lowest tercile of the historical distribution
Observed Precipitation (mm)	Total recorded precipitation
Predicted Payout	Binary set to 1 if forecast predicts a payout; else 0.
Yesteryear Forecast Error (mm)	In the prior year, forecasted minus realized precipitation; positive = overestimate, negative = underestimate.
Forecasted Payout Miss (Overprediction)	Binary set to 1 if, in the prior year, a payout was forecast but none occurred; else 0.
Nonforecasted Payout (Underprediction)	Binary set to 1 if, in the prior year, forecasted precipitation predicted no payout, but realized precipitation resulted in a payout; else 0.
Yesteryear Payout ($)	In the prior year, total indemnities divided by total premiums

Note: All variables measured at the grid or county–year–interval level unless otherwise specified.

Table 2: Data Description for Selected Variables

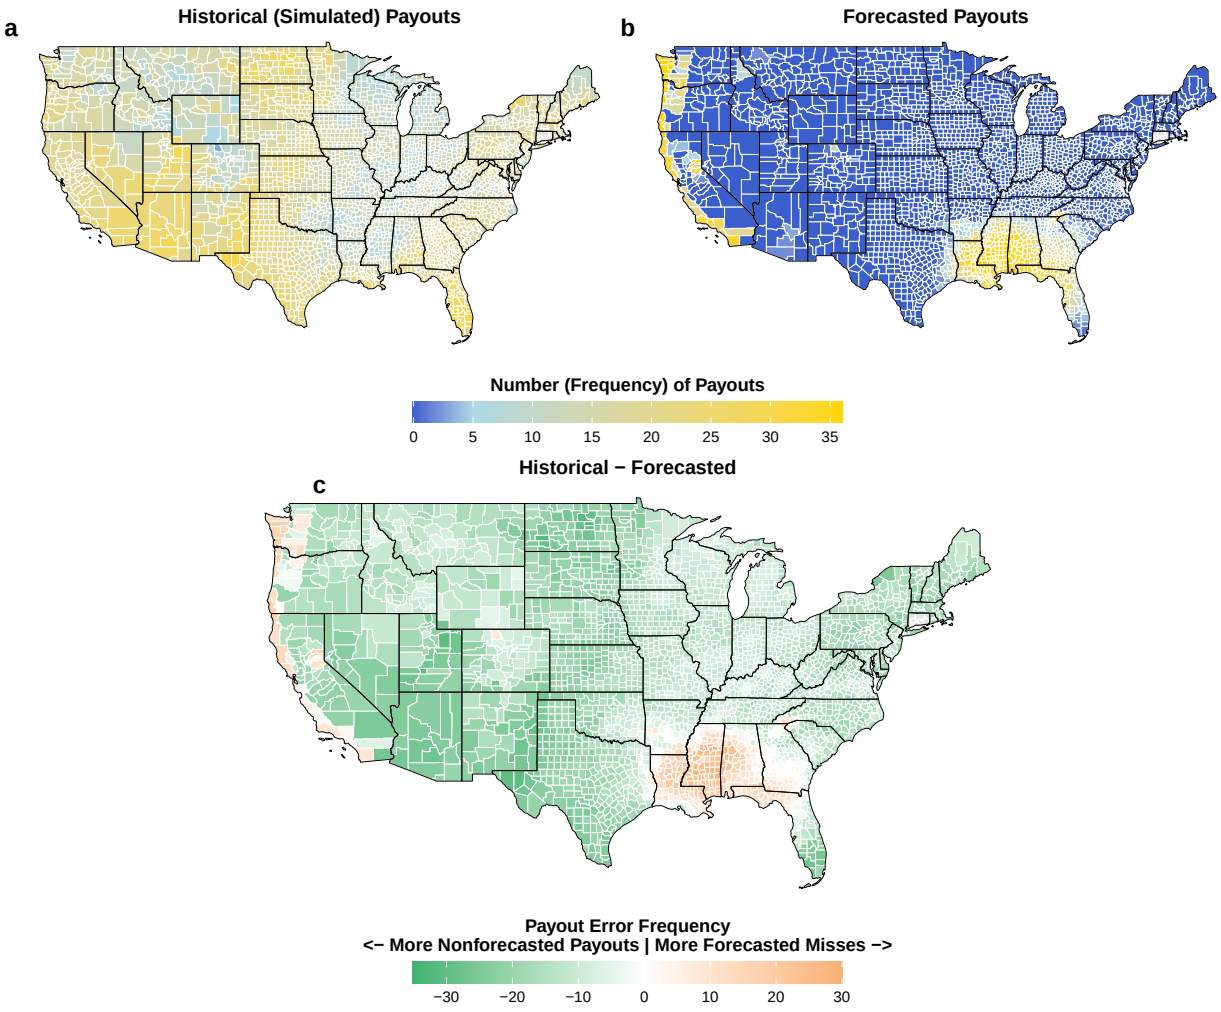


Figure 3: Frequency of payouts in the first three intervals of each year (2017–2023) from **a)** historical rainfall (simulated payouts), **b)** forecasted rainfall (projected payouts), and **c)** their difference, where negative (positive) values indicate over(under)estimated payouts.

4 Methods

4.1 Forecast Quantity and Low Rainfall Probability

Our foundational regression model takes the following form:

$$A_{itv} = \beta \mathbf{X}_{itv} + \alpha_i + \lambda_t + \epsilon_{itv} \tag{2}$$

where A_{itv} represents the (logged) acreage enrolled in PRF in county i, year t, and interval v. β is our coefficient of interest, representing the marginal effect of x_{itv}, which either represents the centered forecasted precipitation in millimeters or the probability of receiving below normal precipitation for a given county, year, and interval. County fixed effects (α_i) account for time-invariant factors within a given county, such as soil composition and geographic features, while year fixed effects (λ_t) account for shocks occurring consistently across all locations in a given year, i.e., national outreach campaigns and macroeconomic conditions. Finally, ϵ_{itv} is an error term with all the usual assumptions for a linear regression model. Standard errors are clustered by county, and regressions are performed in R using the fixest package (Berge 2024).

4.2 Forecast Quality

We also estimate how recent forecast quality influences enrollment. Given evidence for strong recency effects in insurance enrollment decisions Che et al. (2020), we focus on forecast quality in the most recent prior year. Let A_{itv} denote log enrolled acres in county i, year t, and interval v, with $\tilde{x}_{itv}$ a binary variable indicating whether a payout is forecasted for the current county, year, and interval. Forecast quality is first summarized by the yesteryear forecast error (forecasted minus realized precipitation), represented by $e_{iv,t-1}$, then split into overestimation and underestimation dummies. The forecasted payout miss variable $h_{iv,t-1}$ is a binary variable signifying that a payout was forecast but did not occur in the prior year for the same county and interval, while the nonforecasted payout variable $l_{iv,t-1}$ is a binary variable reflecting when a payout was activated but was not forecast in the prior year for the same county and interval. We estimate the following two equations:

$$A_{itv} = \alpha_i + \lambda_t + \beta_x \tilde{x}_{itv} + \beta_e e_{iv,t-1} + \epsilon_{itv} \tag{3}$$

$$A_{itv} = \alpha_i + \lambda_t + \beta_2 \tilde{x}_{itv} + \beta_h h_{iv,t-1} + \beta_l l_{iv,t-1} + \delta_h \left(\tilde{x}_{itv} \times h_{iv,t-1}\right) + \delta_l \left(\tilde{x}_{itv} \times l_{iv,t-1}\right) + \epsilon_{itv}, \tag{4}$$

with county and year fixed effects and standard errors clustered by county. Here, β_x is the forecast effect when recent forecast error is 0 or when recent forecasts have correctly identified payouts (or lack thereof); β_e is the average effect of forecast error; β_h and β_l capture level shifts linked to over- and underprediction of payouts in the prior year; and δ_h and δ_l capture how responsiveness to the current forecast varies with those histories.

4.3 Past Payouts

We also examine whether the history of *realized* PRF payouts influences future enrollment decisions, and how these effects compare to those of current forecasts and simulated forecast quality. As aforementioned, prior work has found that recent payouts can have substantial explanatory power for federal crop insurance program participation, more so even than historical weather patterns (Che et al. 2020). In addition, counties with persistently low forecasted rainfall may also experience consistently high payouts, making it useful to distinguish between these channels. Since the size of enrollment (and hence the potential size of payouts) varies across counties, we standardize the payout variable and estimate:

$$A_{itv} = \alpha_i + \lambda_t + \beta_x x_{itv} + \beta_p P^{\text{real}}_{iv,t-1} + \epsilon_{itv} \tag{5}$$

where x_{itv} is the centered forecasted precipitation in millimeters, and $P^{\text{real}}_{iv,t-1}$ is the ratio of the total dollar amount of indemnities paid over the total dollar amount of premiums earned for county and interval in the prior year. We also swap the centered forecasted precipitation for the forecasted payout dummy variable and include the prior year forecast skill variables as before:

$$A_{itv} = \alpha_i + \lambda_t + \beta_x \tilde{x}_{itv} + \beta_h h_{iv,t-1} + \beta_l l_{iv,t-1} + \delta_h \left(\tilde{x}_{itv} \times h_{iv,t-1}\right) + \delta_l \left(\tilde{x}_{itv} \times l_{iv,t-1}\right) + \beta_p P^{\text{real}}_{iv,t-1} + \epsilon_{itv}, \tag{6}$$

This specification allows us to assess whether historical payouts exert a larger influence on enrollment than contemporaneous forecasts or forecast quality.

5 Results

5.1 Low Rainfall Forecast: Probability and Quantity

Table 3 shows that the probability of below-normal precipitation and lower amounts of forecasted precipitation are both positively associated with enrollment. A one percent increase in the probability of below-normal rainfall is associated with an approximately 27% increase in enrolled acreage (Model 3, exponentiated due to the dependent variable log transformation). Forecasted precipitation appears to move in the opposite direction (Models 2, 4), but the variables capture different aspects of forecast conditions—the probability measure reflects a constrained range of variation in low rainfall likelihoods, while the forecasted precipitation analysis spans a wider range of absolute rainfall values. Both flavors, however, are directionally consistent: drier forecasts correspond to higher enrolled acreage. The apparent difference between raw and logged specifications reflects the influence of a few counties with very large acreages in the raw models (consistent with changes in the western US with large counties); our preferred specifications feature the logged dependent variable that downweights these outliers and more closely adheres to the assumptions of a linear model.

Dependent Variables:	Acres		Acres (Logged)	
Model:	(1)	(2)	(3)	(4)
Variables				
Prob of Low Precip	135.1		0.2445***	
	(323.4)		(0.0230)	
Forecasted Precip (mm)		-35.12***		-0.00046***
		(9.122)		(0.0008)
Fixed-effects				
County	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,901	26,842	18,901	26,842
R^2	0.83	0.82	0.83	0.82

Clustered (county) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 3: Response of enrolled acreage to the forecasted probability of below-normal precipitation and mean forecasted precipitation (mm, mean-centered by county and interval) on PRF-enrolled acres.

5.2 Quality of Forecast

Overall Forecast Error We next examine how forecast quality, both in the current year and in the prior year (or yesteryear), shapes enrollment behavior. Table 4 reports results for predicted payouts and prior-year forecast errors, allowing us to assess whether enrollment patterns reflect contemporaneous outlooks above and beyond the variation associated with recent forecast performance. In the logged-acre specification (Model 2), a predicted payout is associated with a 9.8 percent increase in enrolled acreage. As expected, lower forecast skill (here, prior-year forecast error) exhibits a negative relationship with enrollment: each additional millimeter of mispredicted rainfall corresponds to about a 0.08% decrease in acres ($p < 0.01$). Thus, the average error of 58.2 mm implies roughly a 4.6% decline in enrollment — smaller than the positive impact of a forecasted payout, but still statistically discernible and meaningful. Taken together, these results indicate that both current forecast conditions and recent forecast performance influence producer decisions, with past forecast errors dampening enthusiasm for participation in the following year.

Dependent Variables:	Acres	Acres (Logged)
Model:	(1)	(2)
Variables		
Predicted Payout Dummy	-489.4	0.0937**
	(331.7)	(0.0446)
Prior Year Forecast Error (mm)	-6.203***	-0.0008***
	(1.336)	(0.0001)
Fixed-effects		
County	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	26,842	26,842
R^2	0.82	0.82

Clustered (county) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 4: Acreage response to a predicted payout (current year) and yesteryear forecast errors.

Over- and Underestimation Table 5 examines how much the direction of forecast accuracy in the prior year influences enrollment above and beyond the variation associated with any given forecasted policy year payout. First, focusing on our preferred logged specification in Model 2, forecast overestimation (when the forecast suggests payout would be due but then it didn't materialize in the prior year) depresses enrollment. However, when there is also a payout forecast in the current year, this effect flips sign, yielding a small net increase and suggesting potential enrollees put more weight on current-year information than on prior missed payouts. Likewise, forecast underestimation—when a payout occurred unexpectedly in the prior year—boosts enrollment. However, this positive effect is smaller when there is also a payout forecast in the current year, perhaps reflecting reduced confidence in forecast skill despite the prior year's unexpected windfall.

Dependent Variables:	Acres	Acres (Logged)
Model:	(1)	(2)
Variables		
Forecasted Payout Dummy	-247.5	0.139**
	(367.2)	(0.058)
Prior Year Payout Overestimation Dummy	-2,294.3***	-0.336***
	(437.5)	(0.085)
Prior Year Payout Underestimation Dummy	2,139.2***	0.214***
	(276.3)	(0.015)
Forecasted Payout Dummy × Prior Year Payout Overest. Dummy	1,608.9***	0.239***
	(461.7)	(0.089)
Forecasted Payout Dummy × Prior Year Payout Underest. Dummy	-351.1	-0.146**
	(430.1)	(0.065)
Fixed-effects		
County	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	26,842	26,842
R^2	0.822	0.819

Clustered (county) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 5: Acreage response to prior year forecast accuracy. The forecasted payout dummy accounts for whether a payout is predicted for a given county, interval, and year, given forecasted precipitation amounts. The payout overestimation dummy represents whether a payout was predicted but did not occur in the prior year (*a forecast payout miss*), while the payout underestimation dummy signals that a payout was not predicted but did occur (*nonforecasted payout*).

5.3 Prior Payouts

Whereas the previous sections examined simulated payouts—modeled outcomes that align with features which would have triggered payouts based on CPC data—here we use actual historical payouts, as received by producers and documented in USDA's Summary of Business. Consistent with analyses showing the outsized impact of previous year payouts in affecting agricultural insurance enrollment compared to a policyholder's *prior* weather experience (Che et al. 2020; Stein 2016), table 6 shows that yesteryear's payouts have a strong, positive, and significant effect on enrollment in the subsequent year. In particular, a one-unit increase in standardized prior payouts (indemnities/premia) corresponds to roughly 1,140 more acres, or a 9.2% increase from our preferred logged specification. Lower forecasted precipitation is also linked to higher enrollment, though on the order of about 48 additional acres for each millimeter decrease or a 0.6% increase when using the logged specification. Overall, the results affirm that while producers respond to current-year forecasts, recent payout experience plays a far more important role in shaping enrollment decisions.

Dependent Variables:	Acres	Acres (Logged)
Model:	(1)	(2)
Variables		
Forecasted Precip (mm)	-48.40***	-0.006***
	(10.81)	(0.001)
Standardized Prior Payout ($)	1,140.3***	0.092***
	(152.3)	(0.006)
Fixed-effects		
County	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	23,421	23,421
R^2	0.0042	0.0146

Clustered (County) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 6: Regression results estimating the effects of forecasted precipitation and prior-year payouts on PRF enrollment. Forecasted precipitation is centered by county and interval; prior-year payouts are standardized by the dollar amount of indemnities per dollar of premiums.

6 Conclusion & Discussion

How do producers use seasonal rainfall forecasts when making insurance decisions, and what does this mean for the design and sustainability of index-based programs? Our analysis shows that forecasts show statistically distinguishable influences on enrollment patterns in rainfall index insurance: producers appear to respond to the information contained in precipitation forecasts, and increase (decrease) enrollment when forecasts suggest (no) payouts may be on the horizon. As one might expect, forecast error tends to erode confidence, though responses are asymmetric: overestimated payouts in the prior year reduce enrollment more than underestimated payouts increase it, consistent with greater sensitivity to downside risk. Since both of these error measures are binary and do not capture error magnitude, our results may reflect a tendency for forecasts in our sample to overpredict payouts to a greater degree than they underpredict them. Although the range of absolute error is wider for underestimated forecasts (overpredicted payouts), the median error is lower. Another possibility is that producers adjust enrollment more when negative expectations are confirmed, i.e., a predicted payout fails to occur, than when they receive an unpredicted payout. This asymmetry is consistent with Lybbert et al. (2007), who find pastoralists update expectations more readily in response to downside risk.[4] Such behavior benefits producers in the short run but poses long-run challenges for the financial sustainability of PRF and similar forecast-indexed insurance programs.

From the perspective of the producer, forecasts can serve to generate more accurate expectations of conditions in the following growing season. The producer can then make more informed insurance decisions, whether foregoing onerous premiums in good years or seeking additional protection in bad years. However, given PRF's rainfall index (as opposed to loss-based payouts), producers are also free to use PRF as a financial instrument rather than a risk management tool. Goodrich et al. (2020) also found evidence of enrollees using PRF for reasons other than its intended purpose — over time, many producers choose to enroll in early, non-growing season months with frequently high payouts as opposed to months in which rainfall is more relevant for yield but payouts are less guaranteed. In the short run, these unintended uses may not necessarily undermine the benefits of risk management for producers, who are ostensibly maximizing profits as rational actors. In addition, if producers enroll earlier in the year as discovered by Goodrich et al. (2020), they may also benefit from receiving payments well in advance of any potential loss, and it thus turns into a roundabout form of risk management tool. However, in the long run, producers may experience raised costs or reduced access to insurance due to an uneven distribution of the risk pool. Moreover, because payouts are tied to location-specific weather histories rather than actual losses, PRF can deliver windfall gains or losses unrelated to a producer's true exposure, making it a somewhat sophisticated yet uneven mechanism for managing risk, and one in which returns may hinge as much on if not more so on location-based quirks than on underlying need.

From the perspective of the insurance provider, this behavior is a clear manifestation of adverse selection, wherein producers with greater risk levels selectively enroll. The PRF program was designed to mitigate the risks of adverse selection present within more traditional indemnity insurance by accounting for the unique risk levels of producers in each grid, year, and interval (USDA RMA 2007). However, given PRF's indexing on the single, forecastable peril of rainfall, this program is particularly susceptible to adverse selection based on rainfall forecasts, unraveling the risk pool on which insurance fundamentally operates. Therefore, especially as forecast quality improves, insurers may need to account for the use of forecasts in producers' decision-making.

[4] In Lybbert et al. (2007), downside risk refers to forecasts of low rainfall. In our case, risk aversion manifests as lower enrollment under overestimated forecasts, a feature of index-based design.

Strategies for preserving actuarial soundness include pricing forecasts into premiums or setting earlier end-of-sales dates before forecasts are informative for the following year (Luo et al. 1994; Nadolnyak and Vedenov 2013). Carriquiry and Osgood (2012) recommend offering two different policies with their own premiums: one before the availability of skillful forecasts and one after. For producers who wish to enroll earlier, forecasts are not integrated into premiums, and policies primarily protect against weather risk as traditionally intended. For producers who seek coverage later in the year, premiums do incorporate forecasts, and policies primarily cover the risk of forecast error.

While this work provides empirical evidence of producers' use of forecasts in insurance decisions for a large and growing program in the US agricultural context, future research could complement this work by examining how forecast responsiveness varies with other risk management strategies that may complement or substitute for insurance, such as diversification, input adjustments, or access to alternative financial tools. Additional work could explore other commonly used sources of seasonal climate information, such as almanacs, and assess how they interact with our forecast-based estimates. Finally, research on complementarities between insurance and forecasts is ripe for revisiting. Carriquiry and Osgood (2012) and Osgood et al. (2008), for example, find that forecasts can guide input allocation and thereby increase both input use and insurance purchases when good weather is expected (and reduce both when poor weather is anticipated). Such informed allocation can reinforce program reliability and performance for producers and insurers alike, but in index-based contexts, the promise of payouts based on forecastable conditions may still redirect incentives towards adverse selection. Future work would benefit from examining what types of incentives can harness the joint value of insurance and forecasts while maintaining actuarial soundness, ensuring that adaptation and financial protection efforts work hand in hand to help induce responsible decisions in the face of risky business prospects.

References

Barnston, A. G., M. K. Tippett, M. L. L'Heureux, S. Li, and D. G. DeWitt (2012). "Skill of real-time seasonal ENSO model predictions during 2002–11: Is our capability increasing?" In: *Bulletin of the American Meteorological Society* 93.5, pp. 631–651.

Berge, L. (2024). *fixest: Fast Fixed-Effects Estimations*. R package version 0.12.1. URL: https://cran.r-project.org/web/packages/fixest/index.html.

Booth, M. (Sept. 2024). *Wrecked rain gauges. Whistleblowers. Million-dollar payouts and manhunts. Then a Colorado crop fraud got really crazy. The Colorado Sun*. Accessed 2025-08-07. URL: https://coloradosun.com/2024/09/08/patrich-esch-ed-dean-jagers-ranchers-fraud-insurance/.

Cabrera, V. E., D. Letson, and G. Podestá (2007). "The value of climate information when farm programs matter". In: *Agricultural Systems* 93.1-3, pp. 25–42.

Cai, J. and C. Song (2017). "Do disaster experience and knowledge affect insurance take-up decisions?" In: *Journal of Development Economics* 124, pp. 83–94.

Carriquiry, M. A. and D. E. Osgood (2012). "Index insurance, probabilistic climate forecasts, and production". In: *Journal of Risk and Insurance* 79 (1), pp. 287–300.

Che, Y., H. Feng, and D. A. Hennessy (2020). "Recency effects and participation at the extensive and intensive margins in the US Federal Crop Insurance Program". In: *The Geneva Papers on Risk and Insurance-Issues and Practice* 45.1, pp. 52–85.

Clark, J. P. and N. C. Johnson (2025). "Evaluating predictability limits of North American winter precipitation". In: *npj Climate and Atmospheric Science* 8.1, p. 254.

Clarke, D. J. (2016). "A theory of rational demand for index insurance". In: *American Economic Journal: Microeconomics* 8.1, pp. 283–306.

Doherty, N. A. and P. D. Thistle (1996). "Adverse selection with endogenous information in insurance markets". In: *Journal of Public Economics* 63 (1), pp. 83–102.

Ehsan, M. A., M. L. L'Heureux, M. K. Tippett, A. W. Robertson, and J. Turmelle (2024). "Real-time ENSO forecast skill evaluated over last two decades, with focus on the onset of ENSO events". In: *npj Climate and Atmospheric Science* 7.

Goodrich, B., J. Yu, and M. Vandeveer (2020). "Participation patterns of the rainfall index insurance for pasture, rangeland and forage programme". In: *The Geneva Papers on Risk and Insurance-Issues and Practice* 45.1, pp. 29–51.

Goodrich, B. K. and K. A. Davidson (2024). "Enrollment in Pasture, Rangeland, and Forage Rainfall Index Insurance: Awareness Matters". In: *Journal of Agricultural and Resource Economics* 49.2, pp. 311–331.

Hansen, J. W., S. J. Mason, L. Sun, and A. Tall (2011). "Review of seasonal climate forecasting for agriculture in sub-Saharan Africa". In: *Experimental agriculture* 47.2, pp. 205–240.

Hobeichi, S., G. Abramowitz, A. Sen Gupta, A. S. Taschetto, D. Richardson, N. Rampal, H. Ayat, L. V. Alexander, and A. J. Pitman (2024). "How well do climate modes explain precipitation variability?" In: *npj Climate and Atmospheric Science* 7.1, p. 295.

IRI (2021). *NMME Seasonal Climate Forecasts: Precipitation (ELR)*. https://iridl.ldeo.columbia.edu/maproom/Global/Forecasts/NMME_Seasonal_Forecasts/Precipitation_ELR.html. Baseline climatology 1982–2010 through Aug 2021; updated to 1991–2020 from Sep 2021 onward. Accessed: 2025-08-11.

— (2025). *NMME Seasonal Forecast (Delayed Mode) – ELR v2 Ensemble Mean Precipitation*. https://iridl.ldeo.columbia.edu/SOURCES/.IRI/.FD/.NMME_Seasonal_Forecast/.delayed/.ELR_coeffsv2/.fcstmean/.precip/. Gridded 1° ensemble-mean precipitation forecasts (mm) from ELR v2 applied to NMME models, up to 6 months lead time; delayed mode with quality-controlled inputs. Accessed: 2025-08-12.

Just, R. E., L. Calvin, and J. Quiggin (1999). "Adverse Selection in Crop Insurance: Actuarial and Asymmetric Information Incentives". In: *American Journal of Agricultural Economics* 81 (4), pp. 834–849.

Kirtman, B. P., D. Min, J. M. Infanti, J. L. Kinter III, D. A. Paolino, Q. Zhang, H. Van Den Dool, S. Saha, M. P. Mendez, E. Becker, et al. (2014). "The North American multimodel ensemble: phase-1 seasonal-to-interannual prediction; phase-2 toward developing intraseasonal prediction". In: *Bulletin of the American Meteorological Society* 95.4, pp. 585–601.

Lemoine, D. and S. Kapnick (2024). "Financial markets value skillful forecasts of seasonal climate". In: *Nature Communications* 15.1, p. 4059.

Luo, H., J. R. Skees, and M. A. Marchant (1994). "Weather information and the potential for intertemporal adverse selection in crop insurance". In: *Applied Economic Perspectives and Policy* 16.3, pp. 441–451.

Lybbert, T. J., C. B. Barrett, J. G. McPeak, and W. K. Luseno (2007). "Bayesian Herders: Updating of Rainfall Beliefs in Response to External Forecasts". In: *World Development* 35 (3), pp. 480–497.

Meza, F. J., J. W. Hansen, and D. Osgood (2008). "Economic Value of Seasonal Climate Forecasts for Agriculture: Review of Ex-Ante Assessments and Recommendations for Future Research". In: *Journal of Applied Meteorlogy and Climatology* 47 (5), pp. 1269–1286.

Nadolnyak, D. A. and D. Vedenov (2013). "Information Value of Climate Forecasts for Rainfall Index Insurance for Pasture, Rangeland,and Forage in the Southeast United States". In: *Journal of Agricultural and Applied Economics* 45 (1), pp. 109–124.

National Oceanic and Atmospheric Administration Climate Prediction Center (2025). *CPC Unified Gauge-Based Analysis of Daily Precipitation over CONUS*. Date reflects when accessed and not date of production. URL: https://www.psl.noaa.gov/data/gridded/data.unified.daily.conus.html.

Osgood, D. E., P. Suarez, J. Hansen, M. Carriquiry, and A. Mishra (2008). "Integrating Seasonal Forecasts and Insurance for Adaptation among Subsistence Farmers: The Case of Malawi". In: (Policy Research Working Paper 4651).

R Core Team (2025). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria. URL: https://www.R-project.org/.

Stein, D. (2016). "Dynamics of Demand for Rainfall Index Insurance: Evidence from a Commercial Product in India". In: *The World Bank Economic Review* 32 (3), pp. 692–708.

Tippett, M. K., A. G. Barnston, and S. Li (2012). "Performance of recent multimodel ENSO forecasts". In: *Journal of Applied Meteorology and Climatology* 51.3, pp. 637–654.

United States Department of Agriculture (2022). *Rainfall Index Insurance Standards Handbook*. URL: https://www.rma.usda.gov/sites/default/files/handbooks/2023-18150-Rainfall-Index-Handbook.pdf.

United States Department of Agriculture, Risk Management Agency (2007). *Pasture, Rangeland, Forage Rainfall Index Plan of Insurance*. uncertain date. URL: https://legacy.rma.usda.gov/policies/pasturerangeforage/PRF-RI-Presentation.pdf.

— (2021). *RMA Revises Dates for the Pasture, Rangeland and Forage and Apiculture Insurance Programs*. URL: https://old.rma.usda.gov/en/News-Room/Press/Press-Releases/2021-News/RMA-Revises-Dates-for-the-Pasture-Rangeland-and-Forage--and-Apiculture-Insurance-Programs.

— (2025a). *Pasture, Rangeland, Forage*. URL: https://www.rma.usda.gov/about-crop-insurance/frequently-asked-questions/pasture-rangeland-forage.

— (2025b). *Rainfall Index*. URL: https://www.rma.usda.gov/policy-procedure/general-policies/rainfall-index.

Walker, K. (2025). *tigris: Load Census TIGER/Line Shapefiles*. R package version 2.2.1. URL: https://cran.r-project.org/web/packages/tigris/index.html.

Yu, J. and N. P. Hendricks (2020). "Input use decisions with greater information on crop conditions: implications for insurance moral hazard and the environment". In: *American Journal of Agricultural Economics* 102.3, pp. 826–845.

Appendix

List of Figures

1 Conceptual diagram of index program payouts (shading) triggered when the index value dips below predefined thresholds (blue lines) 4
2 Forecast grid cells and county boundaries in Nevada, for scale. 5
3 Comparison of Historical (Simulated) and Forecasted Payouts for PRF Insurance (2017–2023) 8
A.1 The average absolute yearly change in millimeters of forecasted precipitation for an interval, calculated between 2017 and 2023. . . . . A-2

List of Tables

1 Descriptive statistics, averaged at the county–interval–year level across the 1,233 counties with nonzero enrollment in PRF in the continental US (2017–2023). 7
2 Data Description for Selected Variables 7
3 Response of enrolled acreage to forecasts 11
4 Acreage response to a predicted payout (current year) and yesteryear forecast errors. 12
5 Acreage response to prior year forecast accuracy 13
6 Regression of Prior Payouts 14
A.1 Regression results showing the effect of a forecasted payout as well as rolling 5-year payout overestimation and underestimation rates. 3

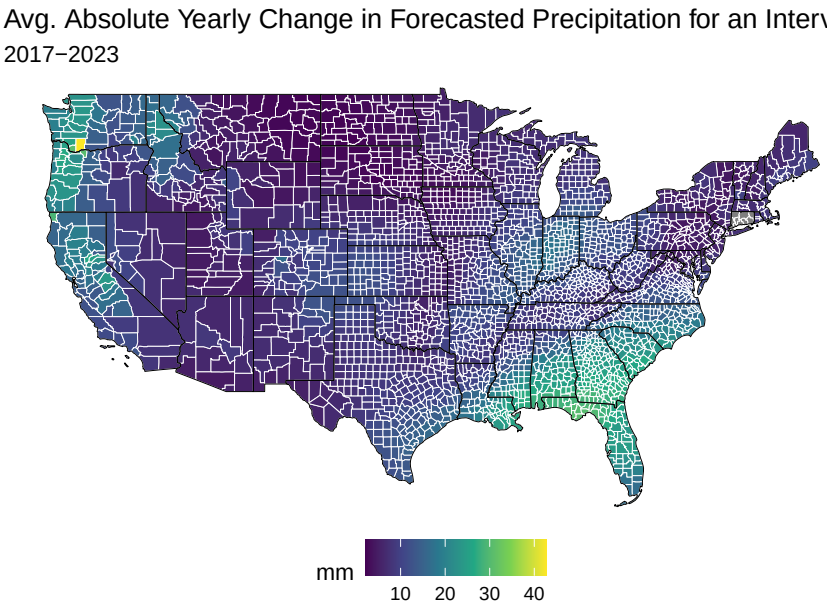


Figure A.1: The average absolute yearly change in millimeters of forecasted precipitation for an interval, calculated between 2017 and 2023.

.1 5 Year Rolling Over-Underprediction

Table A.1 shows the results of models with 5-year rolling over- and underestimation rates defined to capture forecast performance in the medium term. The payout rates refer to the percentage of times within the last 5 years when, for a given county and interval, payouts were predicted but did not occur, and when payouts were not predicted but did occur, respectively. Examining prior-year error against five-year rolling error rates shows the same signs but attenuated magnitudes, potentially reflecting that producers weigh recent forecast performance more heavily than medium-term information.

Dependent Variables:	Acres	Acres (Logged)
Model:	(1)	(2)
Variables		
Forecasted Payout Dummy	-327.7	0.2639*
	(2,513.6)	(0.1405)
Payout Overest. Rate	-114.9***	-0.0175***
	(18.61)	(0.0028)
Payout Underest. Rate	125.4***	0.0135***
	(15.13)	(0.0009)
Forecasted Payout Dummy × Payout Overest. Rate	24.79	-0.0008
	(32.91)	(0.0028)
Forecasted Payout Dummy × Payout Underest. Rate	-27.09	-0.0058**
	(51.84)	(0.0026)
Fixed-effects		
County	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	26,842	26,842
R^2	0.82219	0.82195

Clustered (county) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table A.1: Regression results showing the effect of a forecasted payout as well as rolling 5-year payout overestimation and underestimation rates.