NBER WORKING PAPER SERIES

LONG-RANGE FORECASTS AS CLIMATE ADAPTATION: EXPERIMENTAL EVIDENCE FROM DEVELOPING-COUNTRY AGRICULTURE

Fiona Burlig
Amir Jina
Erin M. Kelley
Gregory V. Lane
Harshil Sahai

Working Paper 32173
http://www.nber.org/papers/w32173

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2024

We are grateful to Susanna Berkouwer, Chris Blattman, Josh Dean, Xavier Gine, Rachel Glennerster, Michael Greenstone, Faraz Hayat, Koichiro Ito, Kelsey Jack, Florence Kondylis, Michael Kremer, Jeremy Magruder, Mushfiq Mobarak, Andrew Robertson, Elena Surovyatkina, Catherine Wolfram, Brian Wright, and seminar participants at the Coase Project, the Paris School of Economics, London School of Economics, EPIC Junior Workshop, Northwestern University, the UChicago Mini-Conference on Weather Advisory Services, the Yale Climate, Environment, and Economic Growth Conference, and Y-RISE for helpful comments and suggestions. We thank Manzoor Dar for excellent field support, Ramya Teeparthi for valuable project management, and Anjani Balu, Alina Gafanova, Sam Hsu, Meghna Singh, Prachi Shukla, Rathan Sudheer, and in particular Amrita Pal for outstanding research assistance. We thank the Becker Friedman Institute for Economics at the University of Chicago, J-PAL's Agricultural Technology Adoption Initiative and King Climate Action Initiative, and the World Bank for generously funding this project. This research has IRB approval from the University of Chicago (Protocol No. IRB20-1364), and is registered on the AEA RCT registry (Identification No. AEARCTR-0008846). All remaining errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Long-Range Forecasts As Climate Adaptation: Experimental Evidence From Developing-Country Agriculture
Fiona Burlig, Amir Jina, Erin M. Kelley, Gregory V. Lane, and Harshil Sahai
NBER Working Paper No. 32173
February 2024
JEL No. D81,O12,O13,Q12,Q54

ABSTRACT

Climate change increases weather variability, exacerbating agricultural risk in poor countries. Risk-averse farmers are unable to tailor their planting decisions to the coming season, and underinvest in profitable inputs. Accurate, long-range forecasts enable farmers to optimize for the season ahead. We experimentally evaluate monsoon onset forecasts in India, randomizing 250 villages into control; a forecast group receiving information well in advance of onset; and a benchmark index insurance group. Forecast farmers update their beliefs and their behavior: farmers who receive "bad news" relative to their priors substantially reduce land under cultivation and certain input expenditures, while those receiving "good news" significantly increase input expenditures. The forecast also impacts crop choice, as farmers tailor their investments. These investment changes meaningfully alter ex post outcomes. In contrast, insurance, which provides no information, increases investments but does not change crops. Our results demonstrate that forecasts are a promising tool for climate adaptation

Fiona Burlig
Harris School of Public Policy
University of Chicago
1307 East 60th Street
Chicago, IL 60637
and NBER
burlig@uchicago.edu

Amir Jina
Harris School of Public Policy
University of Chicago
1155 East 60th Street
Chicago, IL 60637
and NBER
amirjina@uchicago.edu

Erin M. Kelley
World Bank
erinmkelley@worldbank.org

Gregory V. Lane
Harris School of Public Policy
University of Chicago
1307 E. 60th Street
Chicago, IL 60637
and NBER
laneg@uchicago.edu

Harshil Sahai
Department of Economics
5757 S. University Ave.
Chicago, IL 60637
USA
harshil@uchicago.edu

A data appendix is available at
http://www.nber.org/data-appendix/w32173
A randomized controlled trials registry entry is available at
https://www.socialscienceregistry.org/trials/8846

1 Introduction

Climate change is disrupting weather patterns around the world (IPCC (2021)), with extreme temperatures occurring more frequently and rainfall patterns becoming less predictable (Bathiany et al. (2018); Wang et al. (2021)). Agriculture is particularly sensitive to climatic conditions (Hultgren et al. (2022)), putting the 65% of the world's working poor who depend on agriculture for their livelihoods in jeopardy (The World Bank (2022)). Two features make a variable climate particularly challenging for poor farmers to cope with: first, weather risk causes farmers to make fewer profitable investments (Rosenzweig and Binswanger (1993)); and second, when the growing season weather is difficult to predict, farmers cannot easily optimize their investments for the upcoming season. As a result, poor farmers are limited in their ability to adapt to climate change. New tools for adaptation are therefore essential but existing approaches, such as insurance, new seed varieties, or new infrastructure, have proven prohibitively costly (Donovan (2021); Emerick et al. (2016); Lybbert and Sumner (2012)).

In this paper, we use a cluster-randomized experiment to estimate the causal effects of a novel and cost-effective approach to improving farmer welfare in the face of a changing climate: accurate long-range (or "seasonal") forecasts. In theory, and in contrast to short-range (e.g., day-ahead) forecasts, these forecasts enable farmers to tailor their investment decisions to the upcoming growing season, making significant changes to their seasonal agricultural practices such as adjusting their crop mix and ordering inputs in advance (FAO (2019); Gine et al. (2015)). Our empirical results focus on a forecast that provides information about the onset of the Indian Summer Monsoon, which is part of a global phenomenon – almost two thirds of the global population is influenced by monsoonal climate systems (Wang et al. (2021)). Approximately 70–90 percent of total annual rainfall in the majority of India occurs during the monsoon season and the variability of both timing and quantity is large, making it difficult for farmers to predict (Kumar et al. (2013)).

We overcome this challenge with a novel forecast of the onset of the monsoon, maintained by the Potsdam Institute for Climate Impact Research (PIK) and first described in Stolbova et al. (2016).[1] While accurate forecasts could enable substantial behavior change, their use is not widespread in practice, in part because existing forecasts of this type have limited accuracy (Rosenzweig and Udry (2019); Mase and Prokopy (2014)).[2] In contrast, the PIK forecast is extremely accurate, locally-resolved, and can be provided to farmers well in advance of the monsoon's arrival. Released approximately 40 days before onset, the forecast enables farmers to make early decisions about key inputs such as crops, labor supply, and fertilizer purchases (Gine et al. (2015)). The PIK forecast has particular accuracy over Telangana, the site of our experiment: in this region, the forecasted

[1] The PIK forecast relies on recent improvements in weather modeling (e.g., Rajeevan et al. (2007)), and statistically identifies a "tipping point" that is relevant for rainfall onset in a particular location, rather than across the entire sub-continent.

[2] Long-range monsoon onset forecasts, which provide information about when the monsoon will arrive over a month in advance, are notably distinct from short-range forecasts, which typically provide information about day-ahead or week-ahead weather conditions (as studied in Fosu et al. (2018) and Fabregas et al. (2019)). In contrast to these short-run forecasts, which enable marginal behavioral changes, long-range forecasts allow farmers to make decisions at the level of the season, such as what crop to plant.

onset date has been accurate to within one week in each of the past 10 years.

We randomize 250 villages in Telangana into a control group, a group that receives a forecast offer, and a group that receives an index insurance offer to serve as a benchmark.[3] The forecast addresses risk by providing farmers with information about the upcoming growing season, allowing them to tailor their inputs accordingly. In contrast, insurance – the canonical market solution for addressing risk – enables farmers to shift consumption across states but provides no information, making it a useful comparison.[4] In order to ensure that farmers view the forecast as credible, we partner with the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT), a well-respected international organization based in Hyderabad.[5]

We ask four main research questions: how does the forecast (i) change farmers' stated beliefs about monsoon onset; (ii) impact farmers' *ex ante* (i.e., pre-harvest) agricultural behavior; (iii) affect *ex post* welfare metrics at the end of the growing season; and (iv) compare to insurance?

First, farmers update their beliefs in response to our forecast.[6] Farmers in forecast villages have posterior beliefs about the monsoon onset date that are 22% closer to the forecasted onset date than farmers in control. In addition to this concrete evidence on belief updating, farmers have significant demand for the forecast. Using a Becker et al. (1964) mechanism (henceforth "BDM") to elicit willingness-to-pay (WTP), we find that the average WTP for the forecast is comparable to the average WTP for our index insurance product.[7]

Second, farmers tailor their *ex ante* agricultural practices in response to the forecast. Section 3 presents a simple theoretical model of decision-making under variable weather, which demonstrates that a farmer's response to a monsoon forecast should depend on how the forecast compares to their prior. If the forecasted onset is later than the farmer's prior (henceforth "bad news"), the profit-maximizing farmer should respond by reducing their investment in risky agricultural production.[8] If the forecasted onset is in line with the farmer's prior ("neutral news"), they should not substantially adjust their cropping decisions. Finally, if the forecasted onset is earlier than the farmer's prior ("good news"), they should respond by increasing their investments.[9]

We find strong evidence in favor of these predictions in the data for key *ex ante* investments: land use, crop choice, and total input expenditures. Our point estimates imply that farmers who

[3]We sample 5-10 farmers per village for inclusion in the experiment. To avoid attenuation bias from spillovers, all main sample farmers in a given village receive the same treatment.

[4]Prior work has shown that while formal index insurance can improve outcomes substantially (Karlan et al. (2014)), demand is very low, even at actuarially fair rates (Cole and Xiong (2017)), and substantial subsidies are required to increase take-up (Mobarak and Rosenzweig (2014)).

[5]While we find evidence that the farmers did trust the forecast, as described below, it is possible that farmers did not fully believe our information, meaning that our treatment effects will be lower bounds.

[6]There is substantial within-village heterogeneity in priors, with 46% of the total variation in mean prior remaining after removing village fixed effects.

[7]We interpret our WTP results with some caution, as farmers could share information within the village, though we find no evidence of information sharing in practice (see Appendix Table A.8). Moreover, farmers with more diffuse priors have a higher WTP for the forecast, suggesting that this measure is capturing meaningful variation.

[8]To first order, an earlier monsoon – and therefore a longer growing season – is better for farmers, as delays are negatively associated with agricultural output (Mobarak and Rosenzweig (2014), Amale et al. (2023)).

[9]Given these opposite predictions, we are particularly interested in statistical comparisons between bad-news and good-news farmers, in addition to group-wise differences from zero.

received bad news substantially reduced their investments, driven by changes in the amount of land under cultivation (-22%). Those who received neutral news did not alter their investments. On the other hand, farmers who received good news increased their investments considerably: they were 16 percentage points more likely to plant a cash crop, and increased their total input expenditures by 34%.Summarizing these outcomes in an index, we find that farmers who receive bad news reduce investment by 0.12 relative to farmers *with similar priors* in the control group; we see no impact on neutral-news farmers' investments; and good-news farmers increase *ex ante* investment by a standardized effect of 0.29 (while the former is imprecisely estimated, the p-value on the difference between good news and bad news is 0.001).

Third, we see suggestive evidence that these changes in *ex ante* investments lead to changes in *ex post* agricultural outcomes.[10] Farmers who received bad news experience reductions in their agricultural output, in line with having reduced their *ex ante* investments. As a summary metric, these farmers saw a large decline in agricultural profits (-0.28 SD).[11] We find a noisy zero for profits amongst good-news farmers, despite suggestive evidence of increases in agricultural production.[12] Using a machine-learning approach (Chernozhukov et al. (2023)), we show that the average treatment effects on agricultural profits discussed above mask significant heterogeneity, which can be predicted using observable characteristics. Farmers with higher predicted treatment effects of the forecast on profits appear to be worse off at baseline, suggesting that the forecast may be a pro-poor approach to adapting to climate change. These farmers also responded to the forecast by increasing land under cultivation more, but had lower forecast treatment effects on input expenditures and the probability of planting cash crops, suggesting that there were positive returns to expanding land, but zero or negative returns to inputs and planting cash crops in this setting.[13] These results are consistent with our results on agricultural profits for bad-news and good-news farmers: the bad-news farmers in the forecast group responded by reducing land under cultivation, while the good-news farmers weakly expanded land and invested more in inputs.

We also estimate treatment effects on non-agricultural *ex post* outcomes. Bad-news farmers substantially increase net savings, largely driven by a 50% reduction in outstanding debt, and appear to have substituted away from agriculture and into entrepreneurship. In contrast, we find suggestive evidence that good-news farmers experience a slight increase in savings and weakly decrease their non-agricultural business activity, consistent with having devoted more effort to cultivation. Consumption falls for bad-news farmers and rises for good-news farmers, but neither effect is statistically significant and both are small.[14]

[10] While some of these outcomes, such as agricultural yields, are less directly impacted by farmers' decisions than the *ex ante* inputs, and therefore are subject to additional noise generated by the growing season, we nevertheless test the extent to which they are impacted by the forecast.

[11] We find declines in production, the value of agricultural output, and yields (though noisy) that are in line with the reduction in land under cultivation described above.

[12] We do not find corresponding increases in yields or the value of production.

[13] We corroborate this finding using data from the control group only. We hypothesize that the low returns to cash crops may be related to the costs associated with crop experimentation. Farmers in the control group who plant a new crop and/or a cash crop have lower yields on average.

[14] We see a worsening of mental health for the bad-news group, but no change for good-news farmers.

Fourth and finally, we compare the forecast to insurance. Our theoretical model predicts that insurance should cause all farmers to weakly raise investment levels, instead of – as forecasts do – allowing farmers to tailor their investments to the state. Our model also predicts that the effects of insurance should the vary based on farmers' priors. While insurance should weakly increase investment for all farmers, we expect "pessimistic" farmers with late priors (those for whom the forecast would have been good news) should not respond strongly to insurance, while "optimistic" (bad-news) farmers with earlier priors should respond to insurance by increasing investment. Our empirical evidence supports these predictions. On average, farmers in the insurance group substantially increase their overall investments (index treatment effect: 0.12 SD), in line with the prior literature on index insurance (e.g., Karlan et al. (2014)). While we cannot reject similar treatment effects on the investment index between good-news forecast farmers and insurance farmers (p-value 0.122), we find no evidence that insurance farmers changed their crop mix relative to the previous year, and reject that good-news forecast and insurance farmers changed their crops equally, highlighting the differing mechanisms behind forecasts and insurance.[15] Moreover, optimistic farmers see a 0.18 SD increase in the investment index, while pessimistic farmers with late priors see no change (point estimate 0.06 SD). Overall, insurance encourages the most optimistic farmers to invest more, but pessimistic farmers do not respond. In contrast, the forecast ultimately corrects these beliefs, reducing investment among optimistic farmers and increasing investment for pessimistic farmers.

Taken together, these results demonstrate that long-range monsoon forecasts are a promising technology for helping farmers cope with increasing agricultural risk in a changing climate. As a result, this study makes three primary contributions.

We begin by providing the first experimental evidence on the impact of a new climate adaptation technology – an accurate long-range monsoon forecast – on farmer behavior.[16] We identify a key determinant of farmer responses to the forecast: farmers' prior beliefs. We measure farmer priors over the upcoming monsoon's onset, and document substantial heterogeneity – even *within* village – at baseline. We therefore build heterogeneous priors into a simple theoretical model of farmer decision-making under risk to generate predictions about how farmers will respond to forecasts, and test these predictions using our randomized trial. Our treatment causes farmers to update their posterior beliefs in the direction of the PIK forecast, resulting in meaningful changes in both *ex ante* investment and *ex post* outcomes. Our results shed light on the mechanism through

[15] We also do not find evidence of changes in *ex post* agricultural output for the insurance group. The fact that increases in investments do not translate to changes in *ex post* outputs in either the good-news forecast or insurance group suggests that this is a more general phenomenon, rather than something particular to the forecast.

[16] See Meza et al. (2008) for a review of prior research in this area. As Rosenzweig and Udry (2019) write, prior to their own paper and "[despite the potential] importance of both weather outcomes and the existence of direct forecast effects on the overall economy in India, there is [sic] as of yet no rigorous assessments of the impact of long-term weather forecasts and improvements in weather forecast skill on the rural poor." There is a growing body of work on the impacts of short-run forecasts on agriculture (e.g., Fosu et al. (2018); Fabregas et al. (2019); Yegbemey et al. (2023)). Outside of the agricultural sector, a nascent literature in environmental economics uses quasi-experiments to estimate the value of (improving) short-range forecasts of hurricanes (Molina and Rudik (2023)), temperatures (Shrader (2023), Song (2023)), and pollution (Ahmad et al. (2023)), which highlights the value of forecasting under climate change.

which forecasts work: enabling farmers to tailor their behavior to the coming growing season. These findings demonstrate the value of considering prior beliefs in estimating the impacts of information, and illustrate the benefits of a high-quality forecast of the Indian Summer Monsoon. Our results build on important work by Rosenzweig and Udry (2019), who use a farmer fixed-effect design to study the Indian Meterological Department's (IMD) monsoon forecast, and argue that while the IMD's forecast has remarkably low accuracy over the country, an accurate long-range forecast of the Indian summer monsoon has the potential to be worth tens of billions of rupees.[17] Our results highlight the promise of such a forecast, which will only become more valuable as the climate changes.

Second, experimentally demonstrating that the forecast is effective at changing farmer behavior contributes to a broad literature on agricultural risk whose importance is increasing as low-income countries bear the brunt of global climate change (Hultgren et al. (2022)). Our results show that by providing information about the coming growing season, forecasts allow farmers to decide whether to plant at all, what to plant, and how to adjust inputs across crops. This demonstrates that the mechanism behind the effects of forecasts on farmer behavior differs from previous approaches in this literature. In the same experiment, we contrast the forecast with insurance, the most prominent risk-coping technology (Karlan et al. (2014); Cole and Xiong (2017)). We show that this canonical approach allows farmers to smooth risk across states of the world, but does not enable tailored investment.[18] We extend the insurance literature by showing that prior beliefs matter for determining farmer responses to insurance, and that the farmers with the most positive responses to insurance have the most negative responses to the forecast.[19] Finally, we find there is a demand for forecasts, suggesting their potential to be disseminated cheaply at scale.

Finally, by empirically demonstrating the effectiveness of a specific climate adaptation technology – the forecast – we directly advance the growing climate change economics literature. The majority of this work has focused on the economics of mitigation (see Nordhaus (1993) and Pindyck (2013) for reviews), or on the costs of climate change (e.g., Deschênes and Greenstone (2007); Hsiang et al. (2017); Carleton and Hsiang (2016)). We build on a much smaller body of newer work which highlights the importance of adaptation (e.g., Hultgren et al. (2022); Carleton et al. (2022)) but is unable to examine the role of specific adaptation strategies.[20] In contrast, we experimentally

[17]In India, the monsoon's onset is extremely important for the Indian economy (Rosenzweig and Binswanger (1993)) and that farmers' own predictions about monsoon onset shape their planting decisions (Gine et al. (2015)). Though a monsoon forecast would be extremely beneficial – and even more so under a changing climate – India's climatology is complex, which has made modeling and accurate forecasting difficult (Webster (2006); Wang et al. (2015)). Up until now, farmers have had very limited access to high-quality monsoon forecasts as a result.

[18]A nascent literature explores other up-front approaches to coping with risk, such as the adoption of high performing seed varieties and irrigation technologies (Emerick et al. (2016); Jones et al. (2022)). While these approaches are promising, they lock farmers into a particular technology, and technology adoption in low-income contexts has proven challenging (e.g., Duflo et al. (2008)).

[19]Prior work has focused on the low demand for insurance (Mobarak and Rosenzweig (2014); Jensen and Barrett (2017); Carter et al. (2017)), highlighting the large subsidies needed to increase take-up. Newer research aims to increase demand (e.g., through repeated relationships which allow for delayed premium payments (Casaburi and Willis (2018)), but the role of expectations about the coming growing season remains unexplored.

[20]A notable exception is Lane (2024), which demonstrates that an emergency credit product is an effective strategy for coping with flood risk. We build on this with an approach that does not require significant pre-existing financial

evaluate a promising and generally applicable adaptation strategy in the context of a population that is highly vulnerable to climate, and find that the forecast has substantial impacts on farmers' decision-making. While our focus is on one location and one forecast, the monsoon is a global phenomenon, and the timing of the beginning of the rainy seasons is a key input to farmer decisions.[21] As the climate changes further, costs to farmers will rise. Our results demonstrate that forecasts help farmers cope with variable weather, making them increasingly valuable under climate change.[22]

The remainder of this paper proceeds as follows. Section 2 provides relevant details about the research setting. Section 3 presents a simple theoretical model of farmer decision-making under risk. Section 4 describes our experimental design. Section 5 presents our analysis, including our regression specifications and results. Section 6 compares forecasts to insurance. Section 7 concludes.

2 Research context

2.1 Agriculture in Telangana

Our study takes place in Telangana, India. The state is home to 35 million people, and agricultural productivity per worker is low. While 55% of the labor force is employed in agriculture, the sector provides only 15% of the Gross State Value Added (Government of Telangana (2020)). The majority of farms are small, with the average landholding being 1 hectare. Rice is the main staple crop in the state, but Telangana also grows a number of important cash crops. In our research sample, 65% of farmers reported cultivating rice, 44% growing cotton, and 14% growing maize during the previous monsoon season.

Telangana, like much of central India, is dependent on the monsoon for agriculture with about 80% of the total annual rainfall occurring in the monsoon months from June to September. While the monsoon arrives in early–mid June on average, uncertainty over monsoon onset is high: between 1979 and 2019, the standard deviation of the onset date was approximately 20 days.

Weather risk is a substantial concern for agriculture in the state, as it rests in one of the most variable areas of the monsoonal region of India. Both formal and informal methods to smooth risk exist in Telangana. The Government of Telangana, through its *Rythu Bandhu* scheme, provides farmers with a number of pre-season incentives. Primary among these is the unconditional cash transfer of INR 5,000 for each acre planted for each season (Government of Telangana (2020)). This scheme also provides access to credit for farmers to spend on inputs including seeds and fertilizers. One notable national crop insurance program, Pradhan Mantri Fasal Bima Yojana (PMFBY), has

infrastructure and can be disseminated at low cost.

[21] Similar forecasts are showing promise in other parts of India, as well as in other countries (e.g., Potsdam Institute for Climate Impact Research (2021) in Tanzania), and recent advances in statistical-, machine learning-, and physics-based modeling of seasonal weather patterns promise to make long-range forecasts an important input for building climate resilience around the world (Schneider et al. (2022)).

[22] While changes in climatic variability will be damaging to farmers, they are not expected to reduce forecast accuracy, as the underlying physics of the atmosphere are unchanged.

ceased to operate in the state.[23] Private insurance exists, but is severely underutilised. At baseline, only 0.75% of farmers in our sample had heard of rainfall insurance (Appendix Figure A.2).

Information about the weather is also limited. While 65% of farmers in our sample report having received information about the upcoming Kharif season at baseline (conducted prior to planting in early May; see Figure 2), the reliability of these sources is unclear. Appendix Figure A.1 shows the breakdown of farmers' information by source. Very few farmers rely on information from the government (7.4%) or extension services (7.3%). Instead, a large share of farmers report receiving information from other farmers in their village (63.3%) or outside of their village (41.5%).

2.2 Forecasting the monsoon

We study a novel approach to reducing agricultural risk: long-range monsoon *onset* forecasts. These forecasts have the potential to substantially improve farmer welfare, because they enable farmers to materially alter their planting and other input decisions. Moreover, recent work by Mobarak and Rosenzweig (2014) demonstrates that onset timing matters for farmers: farmers are willing to pay for insurance against a delayed onset.

We rely on a novel long-range forecast of the monsoon's onset produced by PIK, and described in Stolbova et al. (2016).[24] This forecast uses climate data from the months leading up to the beginning of the monsoon to predict the timing of the monsoon's onset over specific regions of India, including Telangana.[25] The PIK model produces a probability distribution of potential onset dates, which can be summarized as a likely onset date range, making it easy for farmers to understand. The forecast is issued at least a month in advance of the monsoon onset, enabling farmers to substantively adjust their production decisions. In particular, a month-long period provides farmers with sufficient time to alter their crop selection, adjust the seeds they buy, redistribute their land among the chosen crops, and modify the inputs used along with the quantities purchased. Backcasting over the past 10 years, the PIK forecast was correct each year. When evaluated from 1965–2015, the forecast was correct for 73% of the years in the sample.This forecast is not yet widely available to farmers, leaving us with a unique opportunity to evaluate its impacts.

We prefer the PIK forecast to (i) existing monsoon onset forecasts; (ii) forecasts of monsoon rainfall quantity; and (iii) short-range weather forecasts. First, the PIK forecast represents a significant improvement over existing monsoon onset information. The IMD produces a monsoon onset forecast over Kerala rather than for specific locations around the country, and Moron and Robertson (2014) demonstrate that there is virtually no correlation between the monsoon's onset over Kerala and local onset anywhere else in India.[26] Moreover, the IMD forecast only arrives

[23]The program initially required all agricultural loan-holders to purchase insurance, but when the government subsequently made this condition voluntary, demand collapsed.

[24]See Appendix F for more details on monsoon forecasting.

[25]At the time of this writing, PIK provides three monsoon onset forecasts for India: Telangana, central India, and Delhi. We use the Telangana forecast as it covers one of the country's key agricultural regions.

[26]Unlike PIK, the IMD forecasts the monsoon's onset over Kerala.The IMD does not produce any other regional onset forecasts. However, the monsoon does not progress northwards from Kerala in a predictable manner – meaning that onset over Kerala carries little signal about onset timing over the rest of the country.

two weeks in advance of the monsoon's onset, which also limits its usefulness relative to the PIK forecast. Second, the PIK forecast provides a highly accurate forecast of onset *timing*, and there exist no corresponding accurate monsoon rainfall *quantity* forecasts. The most widely-available existing quantity forecast in India, produced by the IMD, is uncorrelated with actual rainfall in much of the country (Rosenzweig and Udry (2019)). Finally, the PIK monsoon forecast is distinct from the more common short-run "weather forecasts" that aim to predict exact weather conditions at a specific point in the upcoming week or two and cannot be used to make large-scale changes to *ex ante* inputs.[27]

3 Model

In this section we present a simple two-period model of farmers' decision-making under uncertainty, which we use to illustrate the effects of the monsoon forecasts and insurance product.[28] In period one, farmers decide how much to save (s), how much to consume (c_1), and how much to invest ($x \geq 0$) by forming expectations across monsoon onset states ϵ_i and a concave, risky agricultural production technology $f(x, \epsilon_i)$. In the period two, farmers consume (c_2^i) from production and savings.

Production The output from this production technology is modified by the state of the world ϵ_i for $i \in \{1, ..., S\}$, where ϵ_i are ordered so that for any $i > j$ we have higher production and a greater marginal product: $f(x, \epsilon_i) > f(x, \epsilon_j)$ and $f'(x, \epsilon_i) > f'(x, \epsilon_j)$ for all $x > 0$.[29] There is no product at zero investment regardless of the state: $f(0, \epsilon_i) = 0$ for all i. These states can be thought of as approximations for when the monsoon will arrive, with an earlier arrival being associated with greater returns to investment.[30]

Farmer decisions The farmer's prior belief over the probability distribution of ϵ for the coming agricultural season is given by $G(\cdot)$. They use these beliefs to weight possible future outcomes. The farmer therefore solves the following problem:

$$\begin{aligned} \max_{s,x} \quad & u(c_1) + \beta \textstyle\sum_{i=1}^{S} u(c_2^i|\epsilon_i) g(\epsilon_i) \\ \text{s.t.} \quad & c_1 = y - s - p \cdot x \ \& \ c_2^i = f(x, \epsilon_i) + s \end{aligned} \tag{1}$$

[27] Seasonal climate forecasts are a relatively new innovation (see Kirtman et al. (2014) for a review), and are typically physics-based models of the climate system linked to slower-moving conditions. In contrast, short-range weather forecasts use deterministic, numerical simulations of weather variables based on current conditions. Weather forecasting techniques, therefore, are not well-suited to forecasting beyond a short time window.

[28] We provide extended model details in Appendix B.

[29] For simplicity, we assume that monsoon onset is the only determinant of production and that output is monotonically decreasing in onset timing. Of course, in reality, agricultural output will depend on a variety of factors (e.g., temperature, the pest environment, etc), which can be thought of as an error term on the production function, and does not affect the results of the model. One such factor is monsoon rainfall *quantity*, which surely matters for production but has been shown to be largely orthogonal to onset timing (Moron and Robertson (2014)). While it is possible that extremely early rain could be detrimental to agricultural output, in general, delayed monsoons are associated with lower output (Amale et al. (2023)).

[30] The investment level x can also be interpreted as a continuum of crop choices, with varying productivities which depend on the state and are correlated with planting costs. In that sense, for any given state, there is an optimal crop choice x that would maximize production subject to budget constraints.

where $u(\cdot)$ is a concave utility function, c_1 is first period consumption, c_2^i is second period consumption in state i, $g(\epsilon_i)$ is the probability density of the farmer's prior over ϵ, y is starting wealth, s is risk-free savings (or interest free borrowing), p is the price of the input x, and β is the discount factor.

Appendix B.2 shows that, for sufficiently risk-seeking farmers, the optimal investment is an increasing function of their beliefs on the realization of ϵ. In other words, the higher a farmer's prior that it will be a good year, the more they will choose to invest.

Forecasts We now introduce a forecast, μ_f, which provides farmers with information on the likelihood of future states of the world. We assume that the forecast is unbiased (such that $\mu_f = \mathbf{E}[\epsilon]$), but has some noise ($\text{Var}(\mu_f) = \sigma_f^2$, with lower σ_f^2 indicating higher forecast accuracy). The farmer uses this prediction and combines it with their prior $G(\cdot)$ via Bayes' rule to calculate a posterior probability distribution for ϵ, say $G'(\cdot)$. The farmer's average posterior will fall between their prior and the forecast prediction, and will have a smaller standard deviation (less uncertainty) than their prior. How the farmer changes their behavior after receiving the forecast depends on both their priors and the realization of the forecast. Note that any given year will only have *one* such realization, but it is nevertheless valuable to consider treatment effects of different possible forecast realizations.

Figure 1 illustrates the key results of the model.Panel A plots the predicted treatment effect of a forecast on agricultural investment. Each curve represents a single forecast realization. The central line shows our predictions for an average monsoon – which is what occurred in the year of our intervention. Under an average monsoon, farmers with early (and therefore overly-optimistic) priors receive bad news from the forecast, and as a result, reduce their investments. Farmers with average (and therefore correct) priors receive neutral news from the forecast, and do not change their investment behavior strongly. Note, despite the lack of strong response, this information is still valuable to these farmers as it increases their confidence in their prior. Farmers with late (and therefore overly-pessimistic) priors receive good news from the forecast, and increase their investment.

The other two curves on this figure extend these predictions to highlight farmer responses to different forecasts. The top curve shows farmers' responses to a forecast of an early monsoon. Now, the early-prior farmers are correct, and do not update their behavior in response to the forecast, while the average- and late-prior farmers both receive information that they were likely too pessimistic, and invest more. The bottom curve shows responses to a forecast of a late monsoon. Here, early- and average- prior farmers receive a signal that the growing season will be later than they expected, so they reduce investments. The late-prior farmers receive corroborating information from the forecast, and do not adjust their behavior. This figure illustrates how forecasts help farmers to tailor their behavior to the coming growing season.

Panel B shows farmers' responses to an insurance product, which delivers a payout in sufficiently bad states. Regardless of farmers' priors on the upcoming growing season conditions, insurance (black) – which conveys no information about the state – causes all farmers to weakly increase their investment in agriculture. This reduces risk by shrinking the variance in consumption across states. In contrast, the forecast (purple) enables farmers to tailor their investments to the upcoming growing season realization, as shown in Panel A. This highlights the different mechanisms behind forecasts and insurance.

Figure 1: Investment choice with a forecast or insurance (model)

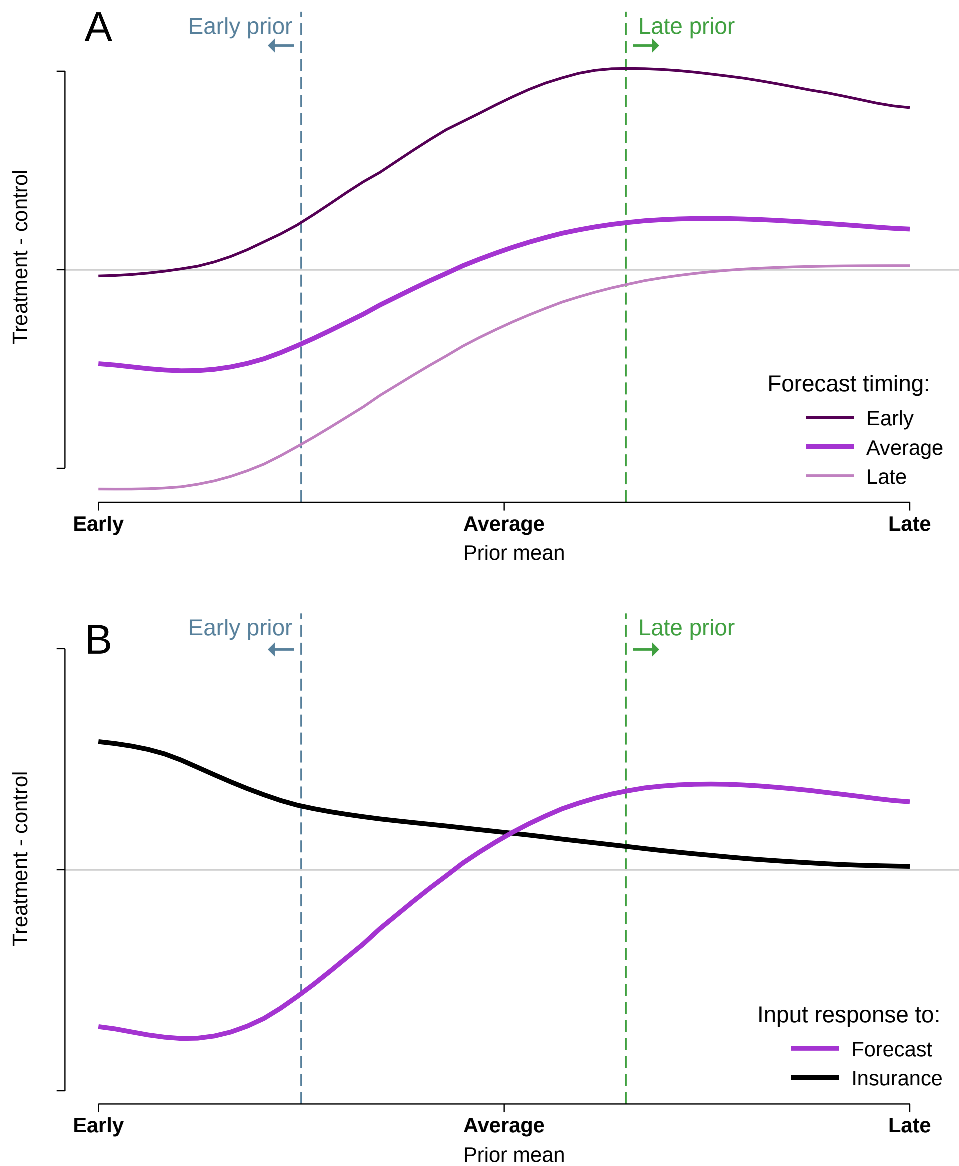

Notes: This figure plots the simulated relationship between the treatment effect on optimal investment and the farmer's prior that the good state of the world will be realized with a forecast or with insurance resulting from our model. The y-axis represents the difference between farmers who receive a treatment and those who do not. The grey horizontal line is centered at zero. The x-axis reflects whether farmers believe the monsoon will arrive early, at the average time, or late. Panel A indicates the investment response of farmers with different priors under different counterfactual realizations of the forecast. Responses to an early forecast realization are depicted by the dark line; responses to an average forecast realization (as was the case in our empirical setting) are depicted by the solid central line; and responses to a late forecast realization are depicted in the light bottom line. Panel B shows differential investment responses between the forecast (purple) and the insurance product (black) for farmers with different priors. See Appendix B.3 for simulation details.

4 Experimental design and data

4.1 Experimental design

Informed by our theoretical framework, we designed a randomized controlled trial to estimate the benefits of forecasts. We randomized 250 villages (sampling 5-10 farmers each) in Telangana into either a forecast group (100 villages), an insurance group (50 villages), or a control group (100 villages).Comparing the forecast group to control identifies the impact of forecasts relative to the control group. Comparing the insurance group to the control group identifies the impact of insurance relative to the forecast group, and allows us to benchmark the impact of forecast relative to another well known risk-mitigation technology.

We sampled villages in two districts in Telangana, Medak and Mahabubnagar, and restricted the sample by excluding villages with high penetration of irrigation, based on data from ICRISAT and the 2011 Indian Census, as these villages were already insulated from the variability of the monsoon. We also drew our sample with a distance buffer between villages, to prevent across-village information sharing. To increase statistical power and ensure balance, we stratified our randomization by district and an indicator for having an above-median number of farmers per acre – a measure of agricultural intensity. We then sampled households within each village for inclusion in our experiment. Each sampled household in a given village received the same treatment. In order to directly measure spillover effects on beliefs within villages, we also conducted a short survey on monsoon beliefs with 2-3 *untreated* households in the forecast villages.

We partnered with ICRISAT to implement this experiment. ICRISAT is an international organization headquartered in Hyderabad, Telangana, close to our study locations. They have over 50 years of experience in Telangana, and are known across the region for breeding and disseminating high-performance crops. They have become one of the most trusted partners for farmers and local extension services working in the area, with an extensive network of partners, which makes them uniquely positioned to deliver these technologies to those in need. Working with ICRISAT and their partners lent credibility to the forecasts for farmers who were encountering this information for the first time.

Forecasts Farmers were told about the forecast using the following text:

> *"In late May/early June each year, we can offer you a forecast which tells you which karte* [an approximately two-week local time step] *the monsoon will arrive in. In 37 of the past 50 years, this forecast has been within one week of the actual start of the rains. It has been better in the past recently: all of the past 10 years' forecasts have been correct."*

We also provided farmers with an information sheet to showcase the forecast's historical accuracy (Appendix Figure H.1). We offered farmers this forecast through a BDM mechanism to elicit farmer willingness-to-pay, which we describe in more detail below. If a farmer purchased a forecast, the enumerator would provide the farmer with the 2022 forecast:

"This year's forecast says that the monsoon is likely to start over Telangana between June 11th and June 19th, in Mrigashira karte. This is likely to be followed by a dry spell from June 20th to June 29th, in the first half of Aarudra karte. The continuous monsoon rainfall is expected after June 29th, in the second half of Aarudra karte."

After visiting the farmers in person to deliver this information, ICRISAT sent a follow-up SMS with the same text.

This forecast is for an average monsoon. Figure 3 shows that the forecast is approximately in the center of the distribution of farmers' prior beliefs. Per our theory, the fact that the forecast is for an average monsoon suggests that farmers with early priors should reduce investment in agriculture in response to the forecast, while farmers with late priors should increase investment.

Insurance Our insurance product provided farmers with financial protection against a late monsoon. We modeled this product directly on Mobarak and Rosenzweig (2014): farmers would receive a sliding-scale payout at harvest time if the monsoon onset was delayed, and not otherwise. We define a village-specific "on time" monsoon onset date based on the average monsoon onset date in that location, using reanalysis data from the ECMWF ERA-5 (Muñoz-Sabater et al. (2021)), and following the approach of Moron and Robertson (2014), as shown in Figure F.1. We installed rain gauges close to each village (approximately one rain gauge per 10 villages), and hired local staff to record their measurements throughout the growing season. For insurance payout purposes, we define onset conservatively (such that payouts are generous): when our rain gauges accumulated 30mm of rainfall over five days and this was not followed by a dry spell of 10 or more days with less than 1mm of rain per day (Mobarak and Rosenzweig (2014)).[31]

Farmers were informed that they would receive a low payout if the monsoon were 15-19 days late compared to the local "on time" onset date; a medium payout if the monsoon were 20-29 days late; and a large payout if the monsoon were 30 days late or later. The maximum payout was set to approximately $190 USD, and was designed to cover approximately 20 percent of the average farmer's agricultural revenues (Ministry of Statistics and Programme Implementation, Government of India (2013)).[32] Farmers in the insurance treatment arm received an information sheet covering these details (Figure H.2). As with the forecast product, we offered farmers this insurance product through a BDM mechanism to elicit willingness-to-pay, which we describe in more detail below. In September, households were notified about whether they would receive a payout, and the actual payments were disbursed in October.

Product offers and takeup In order to ensure high takeup of forecasts and insurance, while as an added benefit, allowing us to measure WTP, we offered these products to farmers through a BDM mechanism, with a price distribution set such that nearly all farmers with positive WTP would ultimately purchase the product (though this distribution was unknown to farmers).[33] We present take-up of the forecast and insurance product in Appendix Figure A.3 and Appendix Table A.4. We

[31] See Appendix Figure A.4 for our rain gauge data. In order to assess the accuracy of the forecast, we use a less strict measure, focused on whether measurable rainfall occurred within the forecasted onset range.

[32] For this calculation, as for all others in the paper, we use an exchange rate of $1 = INR 82.

[33] For more details on our BDM, which was modeled on Berkouwer and Dean (2022), see Appendix G.

Figure 2: Experimental timeline

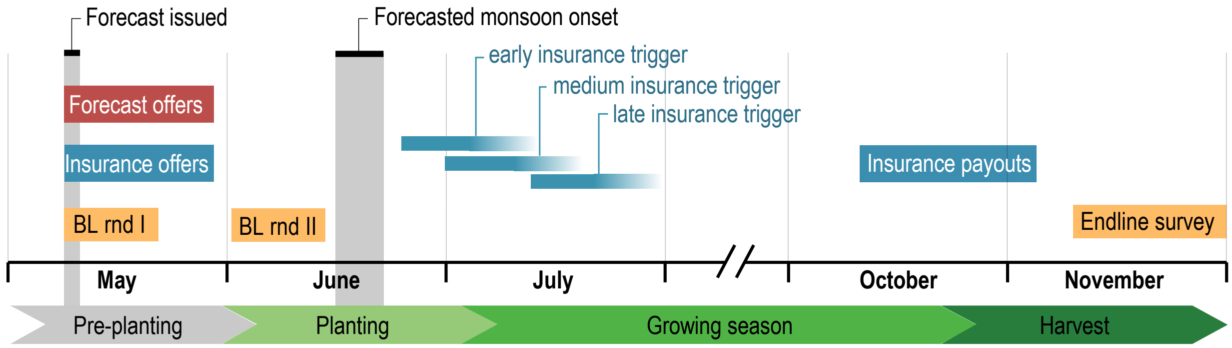

Notes: This figure presents the timeline of the first year of our experiment in relation to the agricultural cycle. The first year of the experiment took place during the 2022 Kharif season. We implemented the baseline survey, and provided treatment offers, and gave farmers the forecast in early May. We visited farmers in early June to collect posterior beliefs. Insurance payouts were triggered by monsoon onset timing, and insurance payouts occurred in October/November. We conclude with a November endline.

find that take-up is over 85 percent for both treatment groups.[34] The remaining farmers reported no interest in the product or declined to participate in the BDM.

Timeline Figure 2 presents the timeline for the experiment. We conducted a baseline survey in May 2022, timed such that we could deliver the PIK forecast at the end of the survey, but still several weeks before the IMD's forecast arrived. Households in the forecast and insurance villages were offered their respective products. For purchasing households in the forecast arm, the information was provided at the end of this visit. This was followed by another visit to households in June 2022, approximately two weeks after the baseline, where we collected data on farmer posterior beliefs about the monsoon. Finally, we conducted our endline survey in November 2022.

The realized monsoon As predicted, over Telangana, the monsoon rain arrived in Mrigashira karte (June 7 - June 20), followed by a dry spell, and then continuous rain beginning in Aarudra karte (June 21 - July 5). As a result, just as was predicted by the forecast, the realized monsoon was very close to average.[35] The forecast was also extremely accurate in our study sample. All 25 of our rain gauges received rainfall by Mrigashira karte. As the forecast also predicted, we find that the amount of rain declined for approximately two weeks following onset, and began to increase again after June 29th. Appendix Figure A.4 shows rainfall across the weather gauges we installed in our sample.

[34] Appendix Figure A.3 and Appendix Table A.3 document the later the farmer thinks the monsoon is likely to be, the more likely farmers are to purchase each product when offered.

[35] In the endline survey, farmers report placing substantially more trust in the forecast after having received it once, presumably because the forecast was accurate (Appendix Figure A.5).

4.2 Data

Outcome data We collected detailed data on three main categories: beliefs, *ex ante* investment, and *ex post* outcomes.

Our first outcome of interest is farmer beliefs about the arrival of the coming monsoon. To measure this, we elicited the farmers' subjective probability distribution of when the monsoon would arrive this year. We did so by providing the farmers with 10 beans to distribute across kartes within a year, following Cole and Xiong (2017). We first asked them to place the beans according to the historical distribution for the past 10 years, where we told farmers to think of each bean as representing one year's monsoon. Once the historical distribution was laid out on the table in front of the farmer, we asked them to consider whether they believed the monsoon would arrive on time, early or late in the coming year. We then asked how they would like to move the beans around in light of their response. We gathered this information during baseline round I and baseline round II to establish whether (and by how much) the forecast changed farmers' priors.

Figure 3 takes the mean of the prior distribution for each farmer, and plots a histogram of these means. The forecast is represented as a purple dashed vertical line. The forecast in 2022 was for an average monsoon, close to the mean of the prior distribution. We divide this distribution into terciles: tercile 1 (indicated by the blue vertical line) are farmers who expected an early monsoon, and would receive bad news from the forecast. Tercile 2 (between the dashed vertical lines) are farmers who (correctly) expected an average monsoon, and would receive neutral news from the forecast. Finally, Tercile 3 (indicated by the green vertical line) are farmers who expected a late monsoon, and would receive good news from the forecast. Appendix Table A.7 and Appendix Figure A.6 suggest that these measures are informative: beliefs correlate with whether the farmer is in their home village and the farmer's land holdings, and control farmers' investments during our study year correlate strongly with their beliefs.[36]

The second main category of outcomes are *ex ante* agricultural investment decisions made by the farmers. We consider a number of choices that may be affected by our treatments, including the amount of land cultivated, which crops they cultivate, and the amount of inputs applied to each plot. For crop choice, we are particularly interested in whether farmers choose to plant cash crops and how these crop choices differ from what the farmer cultivated in the past. Our measure of inputs includes the amount each farmer spent on fertilizer, seed, fertilizer, labor, and total costs. Finally, to summarize these investment choices into one measure, we construct an investment index which is an inverse-covariance weighted index of total land cultivated, cash crop choice, and total input expenditure.

The third main group of outcomes includes downstream *ex post* outcomes for the household. Of primary interest are agricultural output (including profits), the household's financial position, non-agricultural business, and overall household welfare. We measure agricultural output by the weight of crops harvested, the value of crops sold, the value of crops they produced, average

[36] A simple linear regression of our standardized investment index on the mean of prior beliefs in the control group yields a coefficient of -0.14 and p-value of -0.039.

Figure 3: Farmers' priors and the monsoon forecast

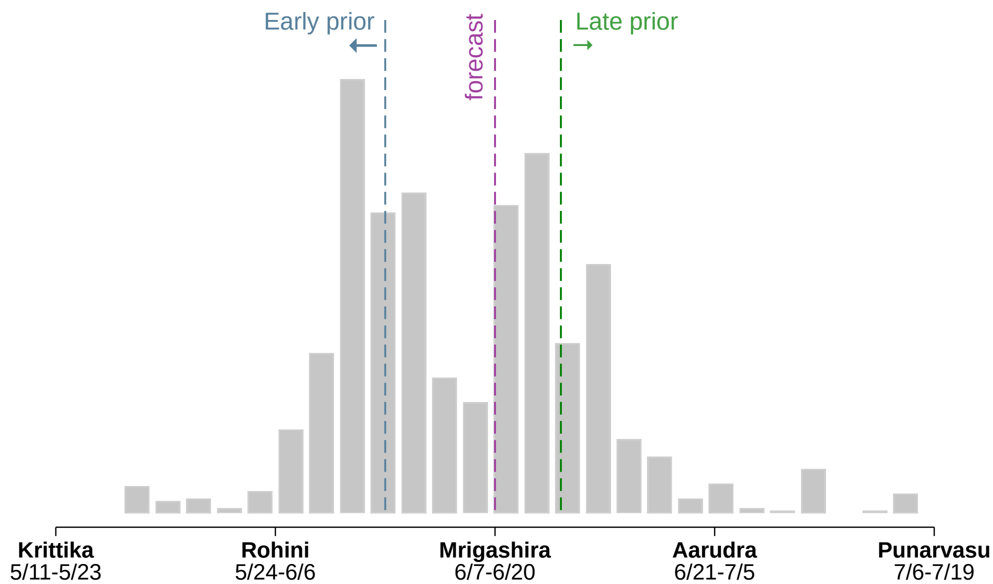

Notes: This figure presents the distribution of farmers' mean priors over the 2022 monsoon onset, measured in kartes (a local approximately two-week long unit of time). To elicit these priors, we use the beans task described in Section 4; we then take the mean of each farmer's prior distribution to form this histogram. The forecasted monsoon onset date is represented by the dashed purple vertical line. The 2022 forecast was for an average monsoon, and the forecast lies close to the mean of the prior distribution. We shade the terciles of beliefs. Tercile 1 (indicated by the blue dashed line) are farmers who expected an early monsoon, and receive bad news in the forecast group. Tercile 2 (white) are farmers who (correctly) expected an average monsoon, and receive neutral news in the forecast group. Finally, Tercile 3 (indicated by the green dashed line) are farmers who expected a late monsoon, and receive good news in the forecast group.

yield, and profits (defined as the value of crops produced minus total expenditure). Our main indicator of the household's financial position is savings net of debt. For non-agricultural business, we consider ownership, investment amount, and business profits. For overall well-being, we focus on two measures: household consumption per-capita across eight consumption categories over the past month and the PHQ-9 screening tool, a standard and locally-validated depression metric to measure mental health (Bhat et al. (2022)). We further consider effects of our treatments on other outcomes, including assets and migration.[37]

Attrition, descriptive statistics, and balance Before proceeding with main results, we test for differential attrition and balance between villages in the control group, forecast treatment group, and insurance treatment group. Appendix Table A.1 shows that overall attrition (defined as being present in baseline round I but absent from *either* baseline round II or endline) is extremely low: only 4% of households in the control group attrited from the study. Households in the insurance treatment arm are more likely to answer all surveys (if anything, this is likely to bias our insurance treatment effects downwards as we anticipate that those who do not respond are likely to have experienced worse outcomes).[38] Appendix Table A.2 explores the correlation between attrition and baseline characteristics. The mean of a farmer's beliefs about monsoon onset this year does not predict differential attrition, though we find that farmers with more diffuse priors (higher SD) are more likely to exit the sample. Taken together, these results imply that the offer of insurance likely retained some farmers with uncertain beliefs over this year's monsoon.

Appendix Table A.3 presents some descriptive statistics and our balance checks. As expected, we find that villages are similar between groups on a variety of characteristics. Villages contain approximately 400 households on average, and span 360 hectares of cultivated land. The share of irrigated land is low by design (approximately 30%). We also find balance across characteristics of our sample households. On average, households consist of five members. The head of the household is typically in their mid-40s and has received 6 years of education. Households have two plots of land on average and cultivate 2.5 hectares of land. The sample is broadly well-balanced, although we see statistically significant differences between the control and forecast treatment villages in terms of the standard deviation of the monsoon onset timing distribution and the standard deviation of expectations over this year's monsoon. However, these differences are quite minor, accounting for only 3% and 4% of the control mean, respectively. As such, we do not consider them to be a significant cause for concern.

Pre-registration This research was pre-registered at the AEA and the analysis plan was accepted via the pre-results review at the *Journal of Development Economics.* We include footnotes in the main text to discuss any changes in regression specification from our analysis plan. A full list of deviations from the PAP is described in Appendix C.

[37] We measure assets via a count of individual asset items, the self-reported value of these assets, and a count of livestock holdings. Finally, we measure migration by capturing how many individuals from the household migrated elsewhere over the cropping season and the value of remittances they sent home.

[38] Of the 495 control group households, 497 forecast group households, and 248 insurance group households, we were unable to conduct all three surveys with 21, 16, and 1 household(s), respectively.

Table 1: Effect of the forecast and insurance on beliefs

	(1) | posterior − forecast |	(2) | posterior − prior |	(3) K-S Stat
Forecast	-0.156*	-0.214**	-0.052*
	(0.087)	(0.098)	(0.028)
Insurance	0.037	-0.059	-0.017
	(0.119)	(0.131)	(0.035)
Control Mean	0.70	0.89	0.44
Observations	921	921	921

Notes: This table presents estimates of the treatment effects of forecasts and insurance on farmers' beliefs about the onset timing of the Indian Summer Monsoon, estimated using Equation (2). To compute priors and posteriors, we use the beans task described in Section 4. |posterior - forecast| is the absolute difference between a respondent's posterior and the forecast date for the monsoon onset. |posterior - prior| is the absolute difference between a respondent's prior and posterior belief for when the monsoon will arrive. K-S Stat is the Kolmogorov–Smirnov test statistic for the difference between a respondent's prior distribution and their posterior distribution. We exclude households where we were unable to speak to the same respondent when eliciting priors and posteriors. All regressions include strata fixed effects and baseline controls chosen by double-selection LASSO. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, ** $p < 0.10$. We present an IV analogue in Appendix Table D.19.

5 Results

5.1 Beliefs and willingness-to-pay

Impact on beliefs The "first stage" effect of a forecast should be to update a farmer's beliefs about monsoon onset. We test for this by comparing prior beliefs elicited at baseline with posterior beliefs measured during baseline round II in the treatment groups compared with the control group and the insurance group (which ought to act as a placebo group, having not received the forecast):

$$Y_{iv} = \beta_0 + \beta_1 \text{Forecast offer}_v + \beta_2 \text{Insurance offer}_v + \gamma \mathbf{X}_{iv} + \eta_{iv} \tag{2}$$

where Y_{iv} are various measures of beliefs for household i in village v; Forecast offer$_v$ is an indicator for being in a forecast offer village, Insurance offer$_v$ is an indicator for being in an insurance offer village, $\mathbf{X}_{iv}$ are strata fixed effects and a set of controls chosen by double-selection LASSO, and η_{iv} is an error term, clustered at the village level.[39]

Table 1 presents the results. We find that the absolute difference between the forecast and the prior is 22% lower in the forecast group than the control group (Column 1). As this year's forecast was for an average monsoon – and therefore close to the mean of the overall prior distribution – we also find that the distance between the posterior and prior distribution is smaller in the forecast treatment arm, measured both in absolute value (24% lower than control, Column 2) and in the Komolgorov-Smirnov test (12% lower than control, Column 3). Reassuringly, we find no evidence that the insurance treatment affected farmers' beliefs. As a result, we conclude that the forecast was successful in shifting farmers beliefs' about the monsoon's arrival. Figure 4 corroborates these

[39]Because takeup of the forecast and insurance products was not 100% (as documented in Appendix Figure A.3 and Appendix Table A.4, we present IV versions of all of the results in Section 5 in Appendix D, where we instrument for forecast (insurance) takeup with an indicator for being in a forecast (insurance) village. As expected, our estimated magnitudes increase somewhat, and significance is broadly unchanged.

Figure 4: Distribution of prior and posterior beliefs

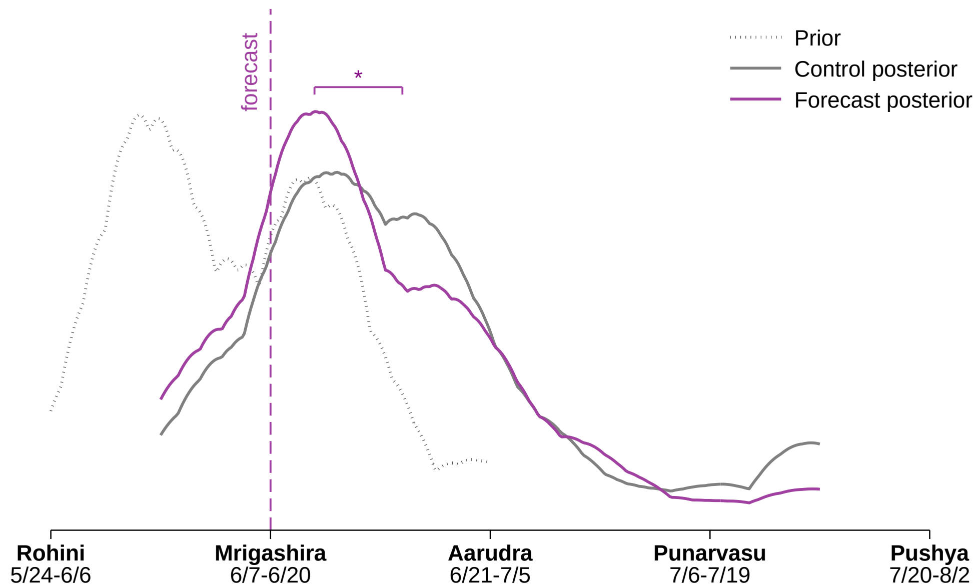

Notes: This figure plots prior and posterior beliefs over this year's monsoon onset, measured in kartes (a local approximately two-week long unit of time). To elicit these beliefs, we use the beans tasked described in Section 4; we then take the mean of each farmer's prior distribution to form distributions over priors and posteriors. The light gray dashed line plots the distribution of priors. The solid gray line plots the distribution of posteriors in the control group, and the solid purple line plots the distribution of posteriors in the forecast group. The vertical purple dashed line indicates the forecast. The overbrace represents the significance level on the test of the null hypothesis on the forecast coefficient in Equation (2), estimated using the posterior mean as the outcome variable (coefficient of -0.177 and p-value 0.061 without controlling for prior beliefs, coefficient -0.176 and p-value 0.063 when controlling for priors). We exclude households where we were unable to speak to the same respondent when eliciting priors and posteriors. We winsorize priors and posteriors at the 3rd and 97th percentile for display purposes, but this does not have a quantitative impact on the regression results nor statistical significance. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

results, showing that while prior beliefs (dashed gray) were approximately centered on the forecast (as shown in Figure 3), posterior beliefs in the control group (solid gray) shifted substantially later. In the forecast group (solid purple), the distribution of posterior beliefs is meaningfully earlier and therefore closer to the forecast.

Willingness-to-pay While we investigate farmers' willingness to pay for the forecast, we interpret the results with some caution. Forecasts can be readily disseminated within the village, leading farmers to potentially offer a lower price in the BDM game compared to their true valuation, under the assumption that they can obtain the same information from fellow farmers. Nevertheless, we find that WTP for the forecast and insurance are remarkably similar (Appendix Figure D.1).

Our theoretical model also suggests that a farmer's WTP will depend on the strength of their prior belief about the monsoon onset. Appendix Table D.5 presents correlations between farmers' WTP and different measures of the strength of farmers' priors.[40] We find suggestive evidence that

[40] In our pre-analysis plan, we erroneously included controls in these regressions. Because these regressions study a single experimental group at a time – rather than comparing treatment to control – this removes useful variation rather than adding precision. Therefore we present the unconditional correlations here. We will include the version with controls in the Appendix.

the strength of farmers' priors matters: farmers with more diffuse priors have a higher WTP for the forecast using some measures of prior strength (share of prior distribution before an on-time threshold, Column 2, or early threshold, Column 3) but not others (SD of prior, Column 1).[41] We do not find evidence of correlations between WTP and the difference between the average (reported) historical onset date and the farmer's prior – farmers' *ex ante* "sophistication" (Column 4) – nor between WTP and risk aversion (Column 5). In theory, the direction of this correlation is ambiguous (Blair and Romano, 1988). Finally, we check for similar correlations with insurance demand. Theory predicts a monotonic relationship between likelihood of a good year and insurance demand, but we find no strong relationship between any *ex ante* prior measures and WTP (Appendix Table D.7). These mixed results suggest that caution is warranted when interpreting our WTP measures, but the fact that WTP for forecasts and insurance is similar suggests farmers do find forecasts valuable.

Information spillovers Finally, we check whether our forecast treatment caused any spillover effects on beliefs. To do so, we compare monsoon beliefs from a sample of untreated farmers living in treated villages (where some farmers received our forecast) to a similar spillover sample in control villages (where no one did). Appendix Table A.8 shows no evidence of information spillovers. While this exercise is informative, it does not rule out the possibility of future information spillovers once farmers have more experience with the forecast, or spillovers in other dimensions (spillover farmers mimicking treated farmers crop decisions, price changes, etc.).

5.2 Effects on *ex ante* choices

Because our theory implies that the effect of the forecast will differ depending on a farmer's prior, our main specification is:

$$\begin{aligned} Y_{iv} = \beta_0 + \sum_{b=1,2,3} \beta_1^b \text{Forecast offer}_v \times [\text{Prior bin = b}]_i \\ + \beta_2 \text{Insurance offer}_v + \rho_b [\text{Prior bin = b}]_i + \gamma \mathbf{X}_{iv} + \eta_{iv} \end{aligned} \tag{3}$$

where $[\text{Prior bin = b}]_i$ are indicators which divide farmers into terciles on the basis of their priors. Those in the first tercile have priors that the monsoon will arrive relatively early (and therefore if they are in the forecast group they will receive bad news); those in the second tercile have priors that the monsoon will be average (and therefore receive neutral news from the forecast); and those in the third tercile have priors that the monsoon will arrive relatively late (and therefore receive

[41] We find no evidence of non-linearities in the relationship between WTP and prior SD (Appendix Table D.6).

good news from the forecast).[42] All other variables are as defined in Equation (2) above.[43]

Because the realized forecast was for an average monsoon, close to the mean of farmer beliefs (Figure 3), we generally expect that the treatment effects for the first tercile (bad-news) and third tercile (good-news) farmers will move in opposite directions, and effects for the second tercile (whose priors were close to the forecast) will be close to zero. Our results are broadly in line with these expectations. Across all four *ex ante* outcomes (land, an indicator for whether a farmer changed crop relative to the prior year, total expenditure on agricultural inputs, and an investment index), our point estimates imply that farmers who received bad news from the forecast substantially reduced their investments, farmers who received neutral news do not change their investments, and farmers who received good news substantially increased their investments.[44]

Land and crop choice We first investigate the impact of our treatments on land use and crop choice (Table 2). We find evidence in support of our theory. Farmers who received bad news *reduce* land under cultivation by 22% of the control mean.[45] We also see that farmers who received bad news were 32% less likely to add a crop type from last year to this year. While they were also less likely to plant a cash crop or change crops, these effects are not statistically significant. Farmers who received neutral news do not change their land under cultivation (point estimate of -3.7%), or their crop choices.

Farmers who received good news *increase* land under cultivation by 15% (though this is imprecisely measured). They are also 16 percentage points more likely to grow a cash crop (Column 2), 16 percentage points more likely to have changed a crop compared to last year, and 14 percentage points more likely to have added a new crop type compared to last year (Column 4), all compared to control group farmers with similar priors. We do not find evidence that these farmers replaced a previous-year crop with something else (Column 5), suggesting that the changes we see were driven by new crops being added to the mix, rather than substitution.

We again find statistically significant differences between farmers who received good news and bad news on land cultivation (p-value 0.01), cash cropping (p-value 0.032), changing crops from last year (p-value 0.013), and adding a crop between last year and this year (p-value 0.004). These

[42] In our pre-analysis plan, we specified that we would split the sample into bad-news and good-news farmers. Upon learning that the monsoon was average, with a large mass of farmers with priors right around the forecast, we chose to divide the sample into terciles to better reflect this heterogeneity. We present continuous treatment effects on our summary investment index in Figure 5. In Appendix D, we also present results from a pooled specification where we do not separate farmers by prior. The insurance effects remain quantitatively unchanged, but – as expected – the forecast results tend to aggregate to zero across a variety of outcome variables, as they average negative and positive treatment effects.

[43] Because we are testing multiple outcomes, in addition to reporting standard p-values, we present sharpened False Discovery Rate (FDR) q-values, which control for the expected proportion of rejections that are Type I errors, following Anderson (2008). We apply these q-values within outcome categories that we measure using multiple questions. This includes all *ex ante* agricultural investment choices, *ex post* agricultural productivity measures, *ex post* welfare measures, *ex post* asset measures, and *ex post* income-generating opportunity measures.

[44] To the extent that farmers did not fully trust the forecast this year, these results are lower bounds.

[45] Throughout the results section, for the sake of interpretation, we present results in percent of the control mean. To do so, we scale our treatment effects (which compare forecast group farmers in each prior tercile with control group farmers in each prior tercile) against the *overall* control mean, ensuring that the three tercile treatment effects remain comparable when converting into percent terms.

Table 2: Effect of the forecast and insurance on land use and cropping

	(1) Land Ha.	(2) Cash Crop	(3) Changed Crop	(4) Added Crop	(5) Sub Crop
Forecast × Ind Bin 1	-0.479**	-0.010	-0.037	-0.117*	-0.000
	(0.188)	(0.049)	(0.054)	(0.060)	(0.047)
Forecast × Ind Bin 2	-0.079	0.024	0.031	0.006	0.010
	(0.177)	(0.038)	(0.050)	(0.046)	(0.040)
Forecast × Ind Bin 3	0.330	0.159**	0.158***	0.141**	0.039
	(0.251)	(0.067)	(0.060)	(0.071)	(0.057)
Insurance	0.266	0.047	0.033	0.038	-0.019
	(0.166)	(0.039)	(0.046)	(0.049)	(0.037)
q-val Tercile 1	0.123	1.000	1.000	0.204	1.000
q-val Tercile 2	1.000	1.000	1.000	1.000	1.000
q-val Tercile 3	0.157	0.039	0.039	0.063	0.323
q-val Insurance	0.540	0.575	0.680	0.680	0.754
Test Tercile 1=3	0.009	0.032	0.013	0.004	0.574
Test Insur. = Ter. 3	0.819	0.120	0.059	0.183	0.348
Control Mean	2.18	0.51	0.57	0.36	0.39
Observations	1201	1201	1201	1201	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on farmers' land use and cropping decisions, estimated using Equation (3). Land Ha. is area cultivated, measured in hectares. Cash Crop is an indicator for the farmer planting at least one cash crop. Changed crop is an indicator for planting a different crop mix in the 2022 Kharif season than the 2021 Kharif season. Added Crop is an indicator for planting at least one additional crop in the 2022 Kharif season compared to 2021. Sub Crop is an indicator for planting at least one fewer crop in the 2022 Kharif season compared to 2021. Bins 1–3 indicate the prior tercile for a respondent. Prior bin 1 were the most optimistic, and received bad news. Prior bin 2 had their beliefs more or less confirmed, receiving neutral news. Prior bin 3 received good news. All regressions include strata fixed effects and baseline controls chosen by double-selection LASSO. "Test Tercile 1 = 3" is the p-value on the test of equality between the first and third coefficient; "Test Insur. = Ter. 3" is the p-value for the test of equality between the third and fourth coefficient. Sharpened q-values are adjusted across all outcomes in Tables 2 and 3 (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Appendix Table D.20 presents an IV analogue.

results are consistent with the forecast enabling tailored investments: farmers in this treatment group adjusted their crop mix to match their expectations about the upcoming growing season.

Farmers in the insurance offer group increased their land under cultivation by approximately 12 percent (p-value 0.11). This increase is similar in magnitude to previous evaluations of index insurance (Karlan et al. (2014)). We cannot reject that insurance farmers and good-news farmers saw the same increase (p-value 0.819). On the other hand, we find limited evidence that insurance changed what farmers chose to grow. We do not expect insurance to affect crop choice as much as the forecast, because insurance does not provide farmers with new information to help decide which crops will be most successful. Though the point estimates on cash crop (4.7 percentage points), changing crop mix from last year to this year (3.3 percentage points), and adding a crop (3.8 percentage points) are positive, they are much smaller than in the forecast group. We reject equality between the good-news and insurance coefficients on crop changes from last year (p-value 0.059), demonstrating that while insurance is useful at enabling farmers to increase overall investment, it does not allow farmers to optimize with respect to the specific upcoming growing season realization.

Inputs We also investigate the impact of our treatments on agricultural input expenditures

Table 3: Effect of the forecast and insurance on inputs

	(1) Fert	(2) Seed	(3) Irri	(4) Labor	(5) Total	(6) Invest Index
Forecast × Ind Bin 1	-143.74*	-79.96	0.60	1.09	-205.24	-0.12
	(75.26)	(199.18)	(16.60)	(92.59)	(281.22)	(0.08)
Forecast × Ind Bin 2	-58.33	-155.22	-1.04	-56.00	-216.97	-0.01
	(63.16)	(123.52)	(11.65)	(69.34)	(198.36)	(0.06)
Forecast × Ind Bin 3	63.74	180.22	22.52	279.31**	671.27***	0.29***
	(76.51)	(160.95)	(16.77)	(109.90)	(255.62)	(0.10)
Insurance	73.91	-126.92*	-7.33	155.39**	235.75	0.12*
	(69.23)	(76.73)	(11.40)	(73.64)	(194.35)	(0.06)
q-val Tercile 1	0.204	1.000	1.000	1.000	1.000	
q-val Tercile 2	1.000	1.000	1.000	1.000	1.000	
q-val Tercile 3	0.291	0.197	0.157	0.039	0.039	
q-val Insurance	0.575	0.540	0.680	0.540	0.575	
Test Tercile 1=3	0.043	0.298	0.355	0.055	0.018	0.001
Test Insur. = Ter. 3	0.911	0.058	0.102	0.340	0.133	0.122
Control Mean	492.51	434.41	54.05	761.96	1948.48	0.00
Observations	1201	1201	1201	1201	1201	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on inputs, estimated using Equation (3). Fert is the amount spend on fertilizer, Seeds the amount spent on seeds, Irri the amount spent on irrigation, and Labor the amount spent on labor throughout the cropping season. Total is the total amount spent on all inputs, including all previous outcomes and any other costs reported by farmers. All outcomes in Columns 1–5 are in USD. Invest Index is an inverse covariance weighted index of land cultivated, cash crop cultivation, and total input expenditure. It has been excluded from the MHT correction as it is a composite of three outcomes already included. Bins 1–3 indicate the prior tercile for a respondent. Prior bin 1 were the most optimistic, and received bad news. Prior bin 2 had their beliefs more or less confirmed, receiving neutral news. Prior bin 3 were the most pessimistic, and received good news. All regressions include strata fixed effects and baseline controls chosen by double-selection LASSO. "Test Tercile 1 = 3" is the p-value on the test of equality between the first and third coefficient; "Test Insur. = Ter. 3" is the p-value for the test of equality between the third and fourth coefficient. Sharpened q-values are adjusted across all outcomes in Tables 2 and 3 (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. We present an IV analogue in Appendix Table D.21.

(Table 3).[46] We again estimate treatment effects for farmers who received bad news, neutral news, and good news, as well as for farmers in the insurance group.

We see suggestive evidence that farmers who received bad news reduced total expenditures in agriculture by 10% (Column 5, a measure that includes expenditures on fertilizer, seeds, irrigation and labor). This reduction is driven by spending 30% less on fertilizer (Column 1), a noisily-measured 18% less on seed (Column 2), with null results on irrigation and labor (Columns 3 and 4).[47] We find no effect on farmers who received neutral news. However, good-news farmers increase their investments substantially, with total expenditures increasing by 34% of the control mean. These are driven by significant changes in labor expenditure (more than 37% of the control mean), and positive but noisy impacts for fertilizer, seed, and irrigation. We reject equality between good-news and bad-news farmers at the 5% level for total spending (p-value 0.018).[48]

[46] Appendix Table D.9 contains treatment effects on per-acre input use. Broadly, we do not find changes in per-acre input use. In theory, farmers would only change per-acre input use in response to treatment if another friction were preventing them for using the optimal amounts of inputs per acre. In this context, farmers appear to be using reasonable levels of input, which suggest that they are not far away from the optimal input per hectare ratio.

[47] We split labor into "early" (prior to planting and planting) and "late" (between planting and harvest and harvest) in Appendix Table D.8.

[48] We also reject equality between good-news and bad-news farmers at at least the 10% level for fertilizer expenditure

We also create an index from outcomes in Table 2 (land cultivation and cash crop) and Table 3 (total input expenditure), and find that bad-news farmers reduced investment by 0.12 standard deviations (p-value 0.13). We find no impacts on farmers who received neutral news, with a standardized treatment effect on the investment index of -0.01 SD. However, good-news farmers experience a 0.29 SD effect on the investment index. We reject equality between good-news and bad-news farmers at the 1% level.

Finally, we find suggestive evidence that farmers who received the insurance treatment increased their input expenditures. While the impact on total expenditures is insignificant, the point estimate suggests a 12% increase in expenditure, which is commensurate with the 12% increase in land cultivation that we reported for the insurance group in Table 2. This appears to be driven in part by a 20% increase in labor expenditures. While imprecisely estimated, we see fertilizer spending increases by 15%. We see no changes in irrigation expenditures, and spending on seeds falls. The overall investment index we create from Tables 2 and 3 also increases by 0.12 standard deviations for the insurance group. We again cannot reject similarity between insurance and good-news farmers for all outcomes except spending on seeds (p-value 0.058). If anything, the total expenditure and investment index treatment effects are *larger* for the good-news group.

These treatment effects suggest that the impact of forecasts differs significantly across farmers with different prior beliefs. Farmers who receive good news are more likely to increase their agricultural investment and experimentation (including land cultivation, crops, inputs), while farmers who receive bad news are more likely to reduce it. This is consistent with our theoretical model. Even in cases where we cannot reject zero for the good- or bad-news group on their own, we can often reject equality between them. Moreover, these results highlight the differences between the forecast and insurance treatments. While both treatments cause farmers to make substantial behavior changes, we only see evidence of changes in the crop mix for forecast farmers, in line with our theory.

As a final point, it is important to note that only some of the individual results discussed above remain statistically significant at conventional levels after applying the MHT correction. In particular, none of the reported increases for specific outcomes in the insurance group have q-values below 0.1, nor do the estimated reductions for forecast farmers in the bad news group. However, while these results should therefore be interpreted with caution, in the case of insurance, the fact that the investment index also shows a strong, highly significant standardized increase supports the conclusion that this treatment encouraged aggregate investment.

5.3 *Ex post* outcomes

Next we examine the extent to which changes in *ex ante* investments lead to improvements in *ex post* welfare (although these outcomes should be less directly impacted by farmers' decisions, because they are exposed to additional noise generated by the realized growing season weather). Our primary outcomes are agricultural output (including profits), non-agricultural enterprise, and

(p-value 0.043), and labor expenditure (p-value 0.055).

well-being.[49] We use the same regression specification we used for our *ex ante* outcomes (Equation (3)). Our *ex post* results generally follow from our *ex ante* findings. Farmers who received bad news did less farming, and as a result have lower agricultural profits, but we find suggestive evidence of substitution into non-agricultural business and meaningfully higher net savings. Farmers who received neutral news did not change their *ex ante* behavior, and we see no corresponding treatment effects on *ex post* outcomes. Finally, for the good-news group, we find a noisy zero on agricultural profits (which can be explained by low or negative returns to some investments this season, as we explore in Section 5.3.1 below), suggestive evidence of less non-agricultural business activity (consistent with doing more farming instead), and positive but noisy impacts on net savings. Finally, we estimate null results on per capita consumption for all groups, including insurance.

Agricultural output We examine five measures of agricultural productivity: total crop production in kilograms, the value of crop sales, the value of crop production (using average sale prices), yield per hectare, and profits (Table 4). The forecast treatment effects follow the broad pattern we documented in the *ex ante* results. We find negative effects for farmers receiving bad news, including a 27% decline in production and a noisy 20% decline in crop sales. We also estimate negative effects on agricultural profits – consistent with the *ex ante* result that these farmers are engaged in less agriculture as a result of the forecast. We find close to zero effects for farmers receiving neutral news. We also find positive effects for farmers who received good news. These effects are relatively noisy, but have point estimates of approximately 22% and 7% for production and yield, respectively. For these two outcomes, we reject equality between bad-news and good-news farmers (p-values of 0.009 and 0.078). Finally, for the insurance treatment, we estimate zero effects on *ex post* outcomes despite the increased investment in land, and inputs documented above.[50]

We document decreases in input expenditures for the bad-news group, which translate into reductions in agricultural profits. However, for the good-news group, we find a null result on agricultural profits, despite large increases in investment (Tables 2 and 3). Using a machine learning approach (Chernozhukov et al. (2023)), we show that the average treatment effects on agricultural profits mask significant heterogeneity, which can be predicted using observable characteristics. We discuss this in more detail below.[51]

Non-agricultural outcomes and welfare We conclude our main analysis with measures of non-agricultural activity and welfare: savings less debt, ownership of a non-agricultural business, non-agricultural investment, business profits, per-capita consumption, and standardized PHQ (Table 5). Column (1) shows that farmers who received bad news increased net savings by more than \$560, driven by a substantial reduction in debt (see Appendix Table A.11). We cannot reject zero effect on net savings for neutral news or good-news farmers. In Columns (2)–(4), we find

[49] We summarize additional impactsin the Appendix.

[50] We fail to reject equality between the good-news farmers and the insurance farmers on all output variables.

[51] It bears mentioning that, while these point estimates on *ex post* outcomes are in the neighborhood of our treatment effects on the *ex ante* outcomes in many cases, we may simply be underpowered to detect treatment effects on the *ex post* outcomes overall. This is perhaps unsurprising, as there is a longer lag between the forcast being provided and these outcomes being realized, during which the rainfall realization and other external forces such as temperature and pests impact agricultural production.

Table 4: Effect of the forecast and insurance on agricultural output

	(1) Prod (Kg)	(2) Crop Sold ($)	(3) Value Prod ($)	(4) Yield	(5) Profit ($)
Forecast × Ind Bin 1	-18.00**	-288.80	-1865.76**	-6.59	-1224.54**
	(8.34)	(257.23)	(923.64)	(4.35)	(602.10)
Forecast × Ind Bin 2	-14.03*	-464.24*	-170.46	-1.76	106.85
	(7.98)	(258.96)	(965.31)	(3.86)	(547.14)
Forecast × Ind Bin 3	14.98	-75.64	12.11	2.46	-588.69
	(10.69)	(283.21)	(891.95)	(3.76)	(529.63)
Insurance	3.80	35.43	-128.20	-2.17	-280.20
	(7.17)	(230.19)	(702.26)	(2.84)	(444.04)
q-val Tercile 1	0.074	0.096	0.074	0.074	0.074
q-val Tercile 2	0.310	0.310	1.000	1.000	1.000
q-val Tercile 3	1.000	1.000	1.000	1.000	1.000
q-val Insurance	1.000	1.000	1.000	1.000	1.000
Test Tercile 1=3	0.009	0.557	0.118	0.078	0.393
Test Insur. = Ter. 3	0.315	0.706	0.875	0.265	0.552
Control Mean	66.15	1428.46	5311.78	34.22	1365.46
Observations	1201	1201	1201	1178	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on agricultural output, estimated using Equation (3). Prod (Kg) is total agricultural production in kilograms. Crop sold ($) is the total value of crops that were sold in USD. Value Prod ($) is the value of all crops produced in USD, whether they were sold or not, using median village prices for each crop. Yield is kilograms of production per hectare. Profit is value of production less total expenditure in USD. Bins 1–3 indicate the prior tercile for a respondent. Prior bin 1 were the most optimistic, and received bad news. Prior bin 2 had their beliefs more or less confirmed, receiving neutral news. Prior bin 3 received good news. All regressions include strata fixed effects and baseline controls chosen by double-selection LASSO. "Test Tercile 1 = 3" is the p-value on the test of equality between the first and third coefficient; "Test Insur. = Ter. 3" is the p-value for the test of equality between the third and fourth coefficient. Sharpened q-values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. We present an IV analogue in Appendix Table D.22.

suggestive evidence that farmers who received bad news engaged in *more* non-agricultural activity, while farmers who received good news engaged in less. While not statistically significant, the point estimates imply that bad-news farmers were 37% more likely than control to own a non-agricultural business, increased non-agricultural investment by 12%, and increased business profits by 47%. In contrast, we see suggestive evidence that good-news farmers were less likely to own a non-agricultural business, reduced non-agricultural investment by more than 70% (significant at the 10% level), and saw a 22% decline in business profits. These results are consistent with good-news farmers investing more in agriculture and bad-news farmers investing less in agriculture, but instead turning to non-agricultural enterprise.

Our treatment effects on consumption mirror those for *ex ante* and *ex post* outcomes: a 13% decline for bad-news farmers, less than a 1% change for neutral-news farmers, and a 4% increase for good-news farmers. However, none of these estimates is statistically different from zero. We also see a noisy 8% reduction for the insurance group. We cannot reject equality between any groups of farmers. For our PHQ measure, we see a meaningful increase of 0.21 for bad-news farmers. This is consistent with stress they might have experienced upon receiving bad news, or their subsequent disappointment in not being able to grow as much as they had hoped for this year. We do not find statistically significant impacts for any other group.[52] While we do not see large impacts on consumption, the large shifts in *ex ante* behavior and *ex post* economic activity provide revealed-preference evidence that both the forecast and insurance were valuable to the farmers in our sample.

5.3.1 Heterogeneity in profits and the returns to agricultural investment

We next examine heterogeneity in the effect of the forecast on agricultural profits, and use our results to assess the returns to different types of agricultural investment in our setting. Because we have a large number of baseline characteristics, we use a generic machine learning approach to predict this heterogeneity (Chernozhukov et al. (2023)).[53]

[52]In Appendix A, we present results for additional *ex post* outcomes: household finances (Appendix Table A.11), other income sources (Appendix Table A.12) assets (Appendix Table A.13), and migration (Appendix Table A.14). We find that the insurance treatment caused households to save less and borrow more in order to fund farming investments. If anything, the forecast treatment appears to have reduced borrowing on net, with the strongest effects for bad-news farmers, who reduced their outstanding debt by nearly 50%. Insurance enabled farmers to conduct more non-agricultural business and increase their business profits; if anything, we find that good-news farmers *reduced* their business activities, consistent with a shift of finances out of business and into agriculture. Neither insurance nor forecasts had significant impacts on assets, though we see suggestive evidence that bad-news farmers increased their asset value, while good-news farmers reduced asset value – perhaps selling assets to fund their agricultural production. Finally, we see that both the forecast and insurance treatments reduced the number of migrants that left the household. These reductions are concentrated among bad news and good-news farmers, with the strongest effects on bad news households. To the extent that migration is a hedging strategy against agricultural risk, households ought to reduce their migration in response to lower risk exposure. In contrast, the reduction in migration rates among farmers with the earliest priors does not fit this explanation.

[53]In addition to this data-driven exercise, in our pre-analysis plan, we specified several unidimensional heterogeneity tests (Appendix Tables D.10 to D.17). We also examine (non-pre-specified) heterogeneity by insurance payout. Appendix Figure E.1, 115 of 247 insurance farmers were in payout-eligible villages, with all 94 who took up insurance receiving payouts of 9,100 INR, approximately $110. Because farmers were unaware of their payout status prior to making *ex ante* decisions, Appendix Tables E.1 and E.2 present results on *ex post* outcomes only. We find no evidence

Table 5: Effect of the forecast and insurance on savings, business activity, and welfare

	(1) Net savings	(2) Non-Ag Bus.	(3) Non-Ag Invest	(4) Bus Profit	(5) Cons Per Cap	(6) PhQ
Forecast × Ind Bin 1	565.68*	0.05	0.40	11.19	-4.87	0.21**
	(295.29)	(0.04)	(1.61)	(11.16)	(7.22)	(0.09)
Forecast × Ind Bin 2	-1.37	0.01	0.09	1.87	-0.58	0.07
	(247.01)	(0.03)	(1.22)	(6.73)	(5.68)	(0.08)
Forecast × Ind Bin 3	238.22	-0.06	-2.35*	-5.25	1.44	0.06
	(333.14)	(0.04)	(1.40)	(12.18)	(4.32)	(0.13)
Insurance	-384.55	0.09***	1.91	15.57*	-3.12	0.05
	(267.94)	(0.04)	(1.35)	(8.07)	(4.19)	(0.07)
q-val Tercile 1	0.186	0.398	0.745	0.463	0.668	0.186
q-val Tercile 2	1.000	1.000	1.000	1.000	1.000	1.000
q-val Tercile 3	1.000	1.000	1.000	1.000	1.000	1.000
q-val Insurance	0.193	0.052	0.193	0.156	0.311	0.311
Test Tercile 1=3	0.463	0.061	0.191	0.311	0.441	0.361
Test Insur. = Ter. 3	0.073	0.004	0.004	0.151	0.379	0.936
Control Mean	-1039.51	0.14	3.30	23.64	40.00	-0.02
Observations	1129	1197	1199	1197	1201	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on savings, business activity, and welfare, estimated using Equation (3). Net savings is savings less outstanding debt in USD. Non-Ag Bus. is a dummy for owning a non-agricultural business. Non-Ag Invest is investment outside of agriculture in USD. Bus Profit is business profit in USD. Cons per cap is consumption per household member in USD. PhQ is the standardized score of the PHQ-9 screening tool; higher values are worse. Bins 1–3 indicate the prior tercile for a respondent. Prior bin 1 were the most optimistic, and received bad news. Prior bin 2 had their beliefs more or less confirmed, receiving neutral news. Prior bin 3 received good news. All regressions include strata fixed effects and baseline controls chosen by double-selection LASSO. "Test Tercile 1 = 3" is the p-value on the test of equality between the first and third coefficient; "Test Insur. = Ter. 3" is the p-value for the test of equality between the third and fourth coefficient. Sharpened q-values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. We present an IV analogue in Appendix Table D.23.

We find meaningful heterogeneity in our sample (Appendix Figure A.7): some farmers increased their agricultural profits in response to the forecast, while others faced a decline. This heterogeneity in the effect of the forecast on agricultural profits is predictable on the basis of baseline characteristics. Broadly, farmers with high predicted (and therefore realized) profits were *less* well-off *ex ante*, having lower landholdings; lower investment in agriculture in 2022, including cash cropping, land under cultivation, and inputs; lower agricultural revenue in 2022; and higher membership in Scheduled Castes. This implies that the forecast may be a particularly pro-poor climate adaptation technology.

Next, we estimate the extent to which this predicted heterogeneity in the effect of the forecast *on profits* is correlated with treatment effect heterogeneity on land use and inputs. Table 6 shows that farmers with higher predicted profit effects have higher forecast treatment effect on land use. To put the heterogeneity in context, moving from the 25th to the 75th percentile of predicted treatment effects – -1414 to 675 USD – is associated with a 0.58 hectare larger treatment effect on land cultivation. Conversely, the farmers we predicted to have higher profit effects also have lower forecast treatment effects on cash cropping and total expenditure – in particular on fertilizer

of differential treatment effects on the basis of payouts.

and seeds – than forecast farmers with low predicted profit effects. Moving from the 25th to the 75th percentile of predicted profit treatment effects reduces the forecast treatment effect on cash cropping by 8.4 percentage points and total expenditure by 460 USD.[54] The fact that farmers with high *predicted* profit treatment effects had increased *realized* profit treatment effects, increased effects on land use, and reduced effects on cash cropping and expenditures suggests that investing in land delivered positive returns this season, while cash cropping and input expenditures had negative returns.

Table 6: Effect of the forecast on land use and inputs by predicted profit effects

A. Land use	(1) Land Ha.	(2) Cash Crop	(3) Changed Crop	(4) Added Crop	(5) Sub Crop
Forecast	−0.197***	0.061***	0.006	−0.009***	−0.027***
	(0.107)	(0.015)	(0.008)	(0.008)	(0.006)
Forecast × CATE	0.028***	−0.004***	−0.001	−0.001	−0.000
	(0.006)	(0.001)	(0.001)	(0.001)	(0.001)
CATE (00 USD)	−0.047***	−0.002	0.001	0.001*	0.000
	(0.007)	(0.001)	(0.001)	(0.001)	(0.001)
Control mean	2.51	0.51	0.57	0.36	0.39
Observations	955	955	955	955	955

B. Input use	(1) Fert	(2) Seed	(3) Irri	(4) Labor	(5) Total
Forecast	−38.46***	−160.70***	9.41***	43.89***	−161.81***
	(16.58)	(59.72)	(2.51)	(15.96)	(76.73)
Forecast × CATE	−4.20***	−18.81***	−0.45	1.66	−22.36***
	(1.87)	(4.39)	(0.33)	(2.20)	(7.50)
CATE (00 USD)	−1.94	8.74***	0.20	−9.51***	−5.76
	(1.61)	(3.40)	(0.21)	(1.82)	(6.44)
Control mean	492.51	434.41	54.05	761.96	1948.48
Observations	955	955	955	955	955
CATE runs	Mean	Median	SD	P25	P75
(00 USD)	−4.68	−2.79	19.31	−14.14	6.75

Notes: This table presents results from 100 splits of the generic machine learning algorithm for heterogeneity prediction described in Chernozhukov et al. (2023). In each split, we use baseline characteristics and a random forest with 100 trees to generate farmer-specific predictions of the effect of the forecast on agricultural profits (ML-predicted Conditional Average Treatment Effects, or CATEs) for all farmers in the control group and forecast treatment group. We then regress each outcome in the table on a forecast group indicator, CATEs (in hundreds of dollars), and the interaction between these two variables. The table presents the median and standard deviation (in parentheses) of the resulting three coefficients across the 100 splits. Land Ha. is area cultivated, measured in hectares. Cash Crop is an indicator for the farmer planting at least one cash crop. Changed crop is an indicator for planting a different crop mix in the 2022 Kharif season than the farmer planted during the 2021 Kharif season. Added Crop is an indicator for planting at least one additional crop in the 2022 Kharif season compared to 2021. Sub Crop is an indicator for planting at least one fewer crop in the 2022 Kharif season compared to 2021. Fert is the amount spend on fertilizer, Seeds the amount spent on seeds, Irri the amount spent on irrigation, and Labor the amount spent on labor throughout the cropping season. Total is the total amount spent on all inputs, including all previous outcomes and any other costs reported by farmers. All outcomes in Columns 1–5 of Panel B are in USD. Inference: ***: the 1st and 99th percentile of the coefficient distribution are either both negative or both positive; **: as above, but 5th and 95th percentile; *: as above but 10th and 90th percentile.

We corroborate this evidence by correlating realized agricultural profits against land under

[54] We also investigate why experimenting with cash crops could have reduced profits. We find that there are costs associated with experimentation: the planting of new crops is associated with lower yields, as farmers may need time to figure out optimal cultivation practices for these new crops (Appendix Table A.10). In addition, the prices these farmers receive may be lower, as they are newly marketing these crops. Appendix Table A.15 displays no effect of our treatments on crop prices.

cultivation and input use in the control group only in Appendix Table A.16. If expanding land and using fewer inputs was indeed the most profitable strategy this year, we would expect to see higher agricultural profits among control group farmers who planted more land but used fewer inputs. Appendix Table A.16 shows exactly this: returns to increasing land under cultivation were positive, while returns to expenditure (particularly seeds) were negative. Specifically, every dollar of extra expenditure is associated with a reduction in profits of 0.67 USD.[55] While input use is not random in the control, the results hold when adding a series of control variables.

Taken together, these findings help to explain the null results of the forecast on agricultural profits for good-news farmers. While we find suggestive evidence that farmers who received good news from the forecast responded by increasing land under cultivation, these farmers also engaged in more cash cropping and spent more on inputs, suggesting that these effects may have offset one another. In addition, these results are consistent with the negative impact on agricultural profits in the bad-news group: these farmers farmed significantly less land than their control-group counterparts, but did not change their crop mix nor their expenditures.

6 Forecasts vs. insurance

How do forecasts compare to insurance, the canonical risk-coping tool? These products have three main differences. First, we hypothesizes that forecasts and insurance work in fundamentally different ways. accurate forecasts provide farmers with information, allowing them to optimize their investments for a given state of the world, but providing no *ex post* assistance based on a specific realization. On the other hand, insurance allows farmers to shift consumption between states of the world, but does not enable farmers to tailor their investments to a particular state. Our treatment effects on investment and crop choice are consistent with these hypotheses. Farmers in the insurance group substantially increase overall investments, with a similar magnitude to farmers receiving good news from the forecast, while farmers who receive bad news from the forecast reduce their investments. In addition, while we find that the forecast impacts crop choice, we do not find that insurance group farmers change their cropping patterns, and reject equality between the two groups. Though receiving insurance could encourage farmers to choose higher-risk, higher-return crops (Karlan et al. (2014)), we expect this effect to be larger for the forecast group, as insurance lacks the forecast's season-specific tailoring property.

Second, in addition to inducing different types of investments, we find evidence that forecasts and insurance also induce investments from different farmers. While not typically modeled in the insurance literature, beliefs about the coming monsoon may also change how farmers respond to insurance. To illustrate this dynamic, we extend our base model presented above to incorporate access to an insurance product for farmers with different *ex ante* beliefs.[56] We use the model

[55]In addition to these tests, Appendix Table A.9 provides suggestive evidence that crop experimentation is associated with lower yields in the control group.

[56]We incorporate the insurance product by adding some fixed amount of income if the realized state falls below some pre-determined threshold (i.e., $i < T$). See Appendix B for more detail.

to predict how the size of the insurance treatment effect varies by farmers' priors. As discussed above, Panel B of Figure 1 shows simulated treatment effects across farmer priors for a farmer who receives an average forecast (purple) compared with a farmer who has access to an insurance product (black). The forecast leads to divergent treatment effects in farmers of different prior beliefs, leading these farmers to tailor their inputs to the coming growing season. In contrast, our theory predicts that insurance should weakly increase investment for all farmers.

However, there is also meaningful heterogeneity in the insurance treatment effects by prior beliefs. We predict that insurance will have strong positive impacts on *ex ante* investments for "optimistic" farmers with early priors (those who would receive bad news with a forecast). As farmers become more pessimistic, the positive treatment effects of insurance fall. We predict no treatment effect at all for the most "pessimistic" farmers with late priors (who would receive good news from a forecast). While insurance is meant to shield farmers from adverse events, pessimistic farmers are unlikely to find agricultural investments appealing, likely believing these investments will go to waste. In contrast, insurance enables optimistic farmers who were previously cautious and did not invest much for fear of the associated downside risk to respond by substantially increasing their investments in anticipation of a promising year. As a result, the difference between the two products is not just in the approach to dealing with weather risk – they also encourage different sets of farmers to invest depending on their priors.

Figure 5 presents an empirical test of the model's predictions, using the *ex ante* investment index as our key outcome of interest. As we describe above, the relationship between investment and prior is positive for the forecast: farmers with later priors (who receive good news) invest *more* with a forecast than farmers with earlier priors (who receive bad news). In contrast, the slope is negative for insurance farmers: farmers with later priors invest *less* with insurance than farmers with earlier priors. Appendix Table D.18 reports effects of the insurance treatment interacted with prior terciles. We focus on our core *ex ante* outcomes of the investment index and its three components – land cultivated, cash crops, and total input expenditure. We find that early-prior insurance farmers increase land under cultivation by 28 percent of the control mean, while late-prior insurance farmers, if anything, reduce land under cultivation. We do not find significant differences between these groups on cash cropping, but the insurance treatment effect on total input expenditure is more than twice as large for optimistic farmers than for pessimistic farmers (though both are noisy). The overall investment index reflects these differences: optimistic farmers see a 0.18 SD increase in investments (p-value 0.075), while we see no change for pessimistic farmers (point estimate 0.06 SD, p-value 0.72). Taken together, these results suggest that farmers' prior beliefs interact with the insurance treatment: insurance is more effective at encouraging investment among optimistic farmers than pessimistic farmers. This contrasts with the forecast, which ultimately corrects these beliefs, reducing investment among overly-optimistic farmers and encouraging investment among overly-pessimistic farmers. These results reinforce the different mechanisms underlying the impacts of forecasts and insurance, suggest that there may be useful complementarities between the forecast and insurance, and may help to explain the low demand for insurance in prior work (e.g., Cole and

Figure 5: Investment choice with a forecast or insurance (empirics)

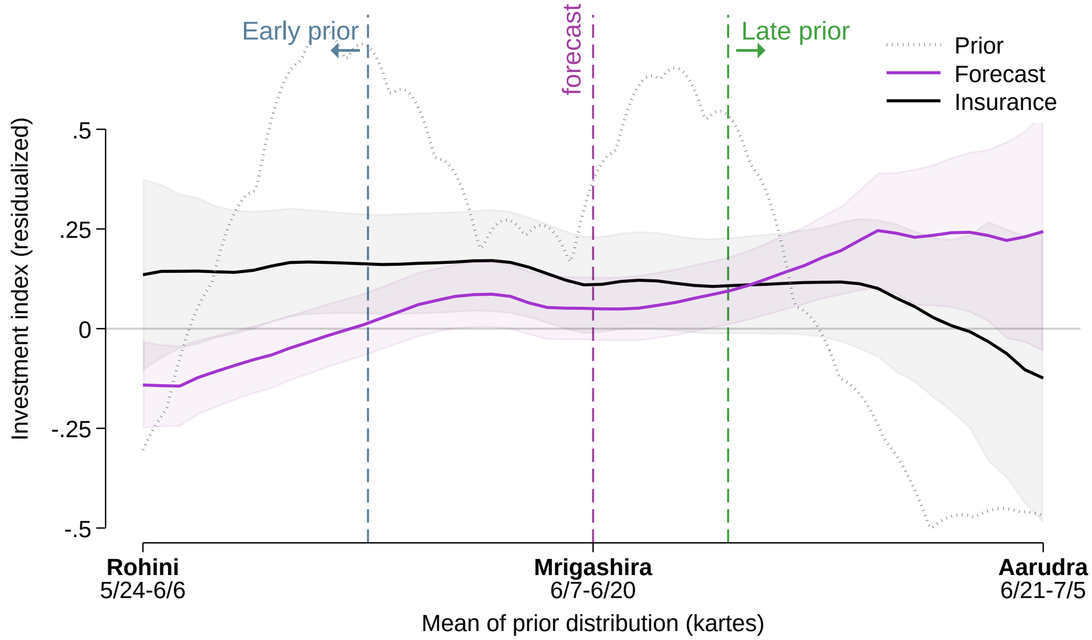

Notes: This figure plots the relationship between the treatment effect on investments and farmers' prior beliefs for the forecast and for insurance: the empirical analogue to Figure 1. We first residualize investments (measured as a standardized index over inputs and land use) using strata fixed effects, enumerator fixed effects, and crop choice from 2021. We then perform two local linear regressions of these residuals on the difference between the mean of the farmer's prior distribution and the forecast date: one for the forecast group vs. control (in purple) and one for insurance vs. control (in black). We winsorize priors at the 3rd and 97th percentile. The purple vertical line denotes the realized forecast (an average monsoon). The dotted gray line plots the prior distribution. The vertical blue and green dashed lines denote the terciles of this distribution.

Xiong (2017)).

Finally, the costs of providing forecasts and insurance are quite different in practice. Though an actuarily fair insurance product is, in theory, self-funding, in reality, insurance requires large subsidies to induce take-up.[57] In contrast, once an accurate weather forecast exists, the marginal cost of delivering it to farmers is very low, making this an attractive option for potential scale.

Despite the differences between the two approaches, however, forecasts and insurance need not be seen as substitutes. Both products enable farmers to make better investment decisions, but they operate on different margins. Instead, forecasts and insurance may be complements: access to insurance protects farmers from the possibility that the forecast is incorrect, enabling farmers to shift even more resources onto the farm under a forecast of a good state. It is therefore possible that access to a forecast could improve demand for insurance, and that both products together could substantially increase farmer welfare.[58] Understanding the size of this interaction is therefore an important topic for future research.

[57]Note that subsidized insurance makes households structural winners: even if they make no behavior changes, they will still receive a cash transfer in bad states of the world. However, this comes at the cost of the subsidy, as well as any costs required for program administration.

[58]Because of the potential for adverse selection into insurance on the basis of the realized forecast, the relevant question is whether knowing that a farmer *will receive* a forecast changes their demand for insurance.

7 Conclusion

In this paper, we use a cluster-randomized trial to study a novel approach for climate adaptation among farmers in low-income countries: long-range monsoon forecasts that provide information about the onset of the Indian Summer Monsoon at least one month in advance of its arrival. We find that receiving a forecast causes farmers to update their beliefs about the weather in the direction of the forecast. Consistent with theory, farmers who receive forecasts tailor their agricultural practices to the upcoming growing season. Farmers who receive good news (that the forecast is for an earlier monsoon than the farmer's prior) significantly increase their investment in agriculture. Good-news farmers also adjust their crop mix in anticipation of a better-than-expected season by doing more cash cropping and adding crops relative to 2021. On the other hand, farmers who receive bad news reduce their investments in agriculture, appearing to substitute into non-agricultural economic activity. These results demonstrate that forecasts are a useful tool for coping with a variable climate – and are likely to become increasingly important as the climate changes further.

While we study long-range forecasts in the context of one Indian state, their usefulness as a tool for climate adaptation likely extends much further. More than a third of the global population lives in the Asian monsoon region, and two thirds live in areas with monsoonal systems more broadly. There already exist similar forecasts elsewhere in India, and new advances in climate science are enabling their development around the world. Broadly representing the global meterological, humanitarian, and food sectors, the COP28 Presidency identified improved forecasts as one of seven priority areas, with "the potential to not only help address the impact of climate change on food security and agriculture, but also transform the lives and livelihoods of millions of farmers" (COP28 Presidency (2023)).

References

Ahmad, Husnain F., Matthew Gibson, Fatiq Nadeem, Sanval Nasim, and Arman Rezaee (2023): "Forecasts: Consumption, production, and behavioral responses," Working paper.

Amale, Hardeep Singh, Pratap Singh Birthal, and Digvijay Singh Negi (2023): "Delayed monsoon, irrigation and crop yields," *Agricultural Economics*, 54 (1), 77–94.

Anderson, Michael L. (2008): "Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training projects," *Journal of the American Statistical Association*, 103 (484), 1481–1495.

Bathiany, Sebastian, Vasilis Dakos, Marten Scheffer, and Timothy M Lenton (2018): "Climate models predict increasing temperature variability in poor countries," *Science Advances*, 4 (5), eaar5809.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964): "Measuring utility by a single-response sequential method," *Behavioral Science*, 9 (3), 262–232.

Berkouwer, Susanna B. and Joshua T. Dean (2022): "Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households," *American Economic Review*, 112 (10), 3291–3330.

Bhat, Bhargav, Jonathan de Quidt, Johannes Haushofer, Vikram H. Patel, Gautam Rao, Frank Schilbach, and Pierre-Luc P. Vautrey (2022): "The long-run effects of psychotherapy on depression, beliefs, and economic outcomes," NBER working paper No. 30011.

Blair, Roger D. and Richard E. Romano (1988): "The influence of attitudes toward risk on the value of forecasting," *Quarterly Journal of Economics*, 103 (2), 387–396.

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang (2022): "Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits," *The Quarterly Journal of Economics*, 137 (4), 2037–2105.

Carleton, Tamma A. and Solomon M. Hsiang (2016): "Social and economic impacts of climate," *Science*, 353 (6304), aad9837.

Carter, Michael, Alain de Janvry, Elisabeth Sadoulet, and Alexandros Sarris (2017): "Index insurance for developing country agriculture: A reassessment," *Annual Review of Resource Economics*, 9, 421–438.

Casaburi, Lorenzo and Jack Willis (2018): "Time versus state in insurance: Experimental evidence from contract farming in Kenya," *American Economic Review*, 108 (12), 3778–3813.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val (2023): "Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India," NBER working paper No. 24678.

Cole, Shawn A and Wentao Xiong (2017): "Agricultural insurance and economic development," *Annual Review of Economics*, 9, 235–262.

COP28 Presidency (2023): "COP28 partnership announces food innovation priorities for investment and scaling," .

Das, Sukanta Kumar, Sanjib Kumar Deb, CM Kishtawal, and Pradip Kumar Pal (2015): "Validation of seasonal forecast of Indian summer monsoon rainfall," *Pure and Applied Geophysics*, 172 (6), 1699–1716.

Deschênes, Olivier and Michael Greenstone (2007): "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather," *American Economic Review*, 97 (1), 354–385.

Donovan, Kevin (2021): "The equilibrium impact of agricultural risk on intermediate inputs and aggregate productivity," *Review of Economic Studies*, 88, 2275–2307.

Duflo, Esther, Michael Kremer, and Jonathan Robinson (2008): "How high are rates of return to fertilizer? Evidence from field experiments in Kenya," *American Economic Review*, 98 (2), 482–488.

Emerick, Kyle, Alain de Janvry, Elisabeth Sadoulet, and Manzoor H. Dar (2016): "Technological innovations, downside risk, and the modernization of agriculture," *American Economic Review*, 106 (6), 1537–1561.

Fabregas, Raissa, Michael Kremer, Matthew Lowes, Robert On, and Guilia Zane (2019): "SMS-extension and farmer behavior: Lessons from six RCTs in East Africa," Working paper.

FAO (2019): "Handbook on climate information for farming communities – What farmers need and what is available," Rome, Italy.

Fosu, Mathias, Dean Karlan, Shashidhara Kolavalli, and Christopher Udry (2018): "Disseminating innovative resources and technologies to smallholders in Ghana (DIRTS)," .

Gine, Xavier, Robert M. Townsend, and James Vickery (2015): "Forecasting when it matters: Evidence from semi-arid India," Working paper.

Government of Telangana (2020): "Socioeconomic Outlook," .

Hsiang, Solomon, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, D. J. Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, Kate Larsen, and Trevor Houser (2017): "Estimating economic damage from climate change in the United States," *Science*, 356 (6345), 1362–1369.

Hultgren, Andrew, Tamma Carleton, Michael Delgado, Diana R. Gergel, Michael Greenstone, Trevor Houser, Solomon Hsiang, Amir Jina, Robert E. Kopp, Steven B. Malevich, Kelly E. McCusker, Terin Mayer, Ishan Nath, James Rising, Ashwin Rode, and Jiacan Yuan (2022): "Estimating global impacts to agriculture from climate change accounting for adaptation," Working paper.

IPCC (2021): *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, vol. In Press, Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

Jensen, Nathaniel and Christopher Barrett (2017): "Agricultural index insurance for development," *Applied Economic Perspectives and Policy*, 39 (2), 199–219.

Jones, Maria, Florence Kondylis, John Loeser, and Jeremy Magruder (2022): "Factor market failures and the adoption of irrigation in Rwanda," *American Economic Review*, 112 (7), 2316–2352.

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry (2014): "Agricultural decisions after relaxing credit and risk constraints," *The Quarterly Journal of Economics*, 129 (2), 597–652.

Kirtman, Ben P, Dughong Min, Johnna M Infanti, James L Kinter, Daniel A Paolino, Qin Zhang, Huug Van Den Dool, Suranjana Saha, Malaquias Pena Mendez, Emily Becker, et al. (2014): "The North American multimodel ensemble: phase-1 seasonal-to-interannual prediction; phase-2 toward developing intraseasonal prediction," *Bulletin of the American Meteorological Society*, 95 (4), 585–601.

Kumar, K Niranjan, M Rajeevan, DS Pai, AK Srivastava, and B Preethi (2013): "On the observed variability of monsoon droughts over India," *Weather and Climate Extremes*, 1, 42–50.

Lane, Gregory (2024): "Adapting to floods with guaranteed credit: Evidence from Bangladesh," *Econometrica*.

Lybbert, Travis J. and Daniel A. Sumner (2012): "Agricultural technologies for climate change in developing countries: Policy options for innovation and technology diffusion," *Food Policy*, 37 (1), 114–123.

Mase, Amber Saylor and Linda Stalker Prokopy (2014): "Unrealized potential: A review of perceptions and use of weather and climate information in agricultural decision making," *Weather, Climate, and Society*, 6 (1), 47–61.

Meza, Francisco J., James W. Hansen, and Daniel Osgood (2008): "Economic value of seasonal climate forecasts for agriculture: Review of ex-ante assessments and recommendations for future research," *Journal of Applied Meteorology and Climatology*, 47 (5), 1269—-1286.

Ministry of Statistics and Programme Implementation, Government of India (2013): "Situation assessment survey of agricultural households," National Sample Survey, 70th round.

Mobarak, Ahmed Mushfiq and Mark Rosenzweig (2014): "Risk, insurance and wages in general equilibrium," NBER working paper No. 19811.

Molina, Renato and Ivan Rudik (2023): "The social value of hurricane forecasts," Working paper.

MORON, VINCENT AND ANDREW W ROBERTSON (2014): "Interannual variability of Indian summer monsoon rainfall onset date at local scale," *International Journal of Climatology*, 34 (4), 1050–1061.

MORON, VINCENT, ANDREW W ROBERTSON, AND DS PAI (2017): "On the spatial coherence of sub-seasonal to seasonal Indian rainfall anomalies," *Climate Dynamics*, 49 (9-10), 3403–3423.

MUÑOZ-SABATER, JOAQUÍN, EMANUEL DUTRA, ANNA AGUSTÍ-PANAREDA, CLÉMENT ALBERGEL, GABRIELE ARDUINI, GIANPAOLO BALSAMO, SOUHAIL BOUSSETTA, MARGARITA CHOULGA, SHAUN HARRIGAN, HANS HERSBACH, ET AL. (2021): "ERA5-Land: A state-of-the-art global reanalysis dataset for land applications," *Earth System Science Data Discussions*, 1–50.

NORDHAUS, WILLIAM D. (1993): "Reflections on the economics of climate change," *Journal of Economic Perspectives*, 7 (4), 11–25.

PINDYCK, ROBERT S. (2013): "Climate change policy: What do the models tell us?" *Journal of Economic Literature*, 51 (3), 860–872.

POTSDAM INSTITUTE FOR CLIMATE IMPACT RESEARCH (2021): "Delayed monsoon in Delhi: early warning for farmers," *PIK News*.

PREENU, PN, PV JOSEPH, AND PK DINESHKUMAR (2017): "Variability of the date of monsoon onset over Kerala (India) of the period 1870–2014 and its relation to sea surface temperature," *Journal of Earth System Science*, 126 (5), 76.

RAJEEVAN, M, DS PAI, R ANIL KUMAR, AND B LAL (2007): "New statistical models for long-range forecasting of southwest monsoon rainfall over India," *Climate Dynamics*, 28 (7-8), 813–828.

ROSENZWEIG, MARK AND CHRISTOPHER R UDRY (2019): "Assessing the benefits of long-run weather Forecasting for the rural poor: Farmer investments and worker migration in a dynamic equilibrium model," NBER working paper No. 25894.

ROSENZWEIG, MARK R. AND HANS P. BINSWANGER (1993): "Wealth, weather risk and the composition and profitability of agricultural investments," *Economic Journal*, 103 (416), 56–78.

SCHNEIDER, ROCHELLE, MASSIMO BONAVITA, ALAN GEER, ROSSELLA ARCUCCI, PETER DUEBEN, CLAUDIA VITOLO, BERTRAND LE SAUX, BEǴUM DEMIR, AND PIERRE-PHILIPPE MATHIEU (2022): "ESA-ECMWF Report on recent progress and research directions in machine learning for Earth System observation and prediction," *npj Climate and Atmospheric Science*, 5 (51).

SHRADER, JEFFREY G. (2023): "Improving climate damage estimates by accounting for adaptation," Working paper.

Song, Yuqi (2023): "The value of weather forecasts: Labor responses to accurate and inaccurate temperature forecasts in China," Working paper.

Stolbova, Veronika, Elena Surovyatkina, Bodo Bookhagen, and Jürgen Kurths (2016): "Tipping elements of the Indian monsoon: Prediction of onset and withdrawal," *Geophysical Research Letters*, 43 (8), 3982–3990.

The World Bank, UN (2022): "What you need to know about food security and climate change," .

Wang, Bin, Michela Biasutti, Michael P Byrne, Christopher Castro, Chih-Pei Chang, Kerry Cook, Rong Fu, Alice M Grimm, Kyung-Ja Ha, Harry Hendon, et al. (2021): "Monsoons climate change assessment," *Bulletin of the American Meteorological Society*, 102 (1), E1–E19.

Wang, Bin, Baoqiang Xiang, Juan Li, Peter J Webster, Madhavan N Rajeevan, Jian Liu, and Kyung-Ja Ha (2015): "Rethinking Indian monsoon rainfall prediction in the context of recent global warming," *Nature Communications*, 6 (1), 1–9.

Webster, Peter J (2006): "The coupled monsoon system," in *The Asian Monsoon*, ed. by Bin Wang, Springer, 3–66.

Yegbemey, Rosaine N, Gunther Bensch, and Colin Vance (2023): "Weather information and agricultural outcomes: Evidence from a pilot field experiment in Benin," *World Development*, 167, 106178.